

2014

Peripersonal Space in the Humanoid Robot iCub

Ramirez Contla, Salomon

<http://hdl.handle.net/10026.1/3050>

<http://dx.doi.org/10.24382/4045>

Plymouth University

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's prior consent.

**PERIPERSONAL SPACE IN THE
HUMANOID ROBOT ICUB**

by

Salomón Ramírez Contla

A thesis submitted to the University of Plymouth
in partial fulfilment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

School of Computing and Mathematics
Faculty of Science and Environment

May 2014

PERIPERSONAL SPACE IN THE HUMANOID ROBOT ICUB

by
Salomón Ramírez Contla

Abstract

Developing behaviours for interaction with objects close to the body is a primary goal for any organism to survive in the world. Being able to develop such behaviours will be an essential feature in autonomous humanoid robots in order to improve their integration into human environments. Adaptable spatial abilities will make robots safer and improve their social skills, human-robot and robot-robot collaboration abilities.

This work investigated how a humanoid robot can explore and create action-based representations of its peripersonal space, the region immediately surrounding the body where reaching is possible without location displacement. It presents three empirical studies based on peripersonal space findings from psychology, neuroscience and robotics. The experiments used a visual perception system based on active-vision and biologically inspired neural networks.

The first study investigated the contribution of binocular vision in a reaching task. Results indicated the signal from vergence is a useful embodied depth estimation cue in the peripersonal space in humanoid robots. The second study explored the influence of morphology and postural experience on confidence levels in reaching assessment. Results showed that a decrease of confidence when assessing targets located farther from the body, possibly in accordance to errors in depth estimation from vergence for longer distances. Additionally, it was found that a proprioceptive arm-length signal extends the robot's peripersonal space. The last experiment modelled development of the reaching skill by implementing motor synergies that progressively unlock degrees of freedom in the arm. The model was advantageous when compared to one that included no developmental stages.

The contribution to knowledge of this work is extending the research on biologically-inspired methods for building robots, presenting new ways to further investigate the robotic properties involved in the dynamical adaptation to body and sensing characteristics, vision-based action, morphology and confidence levels in reaching assessment.

Table of Contents

	Page
Abstract	iii
Table of Contents	v
List of Figures	ix
List of Tables	xiii
List of Code	xiii
Acknowledgements	xv
Author's declaration	xvii
1 Setting the study context	1
1.1 Introduction	1
1.2 Aims and objectives	9
1.2.1 Thesis aims	9
1.2.2 Objectives	9
1.3 Dissertation outline	10
2 Background	13
2.1 Robots	13
2.1.1 Humanoid robots	16
2.1.2 Human-Robot Interaction	18
2.1.3 Evolutionary robotics	20
2.1.4 Bottom-up approaches to adaptive behaviour	21
2.1.5 Developmental robotics	23
2.1.6 Relation to this work	25
2.2 Embodied intelligence	26
2.2.1 Embodied artificial intelligence	28
2.3 Reaching development and peripersonal space	30
2.3.1 The development of reaching in infants	30
2.3.2 Peripersonal space	33
2.3.3 Existing work on robotics reaching and peripersonal space	41
2.4 Vision	46

2.4.1	The classical approach to computer vision	46
2.4.2	Active vision	47
2.4.3	Visual vergence for depth estimation	54
2.4.4	Existing work on robot mono and stereo vision	55
2.5	Machine learning	59
2.5.1	Supervised learning	60
2.5.2	Unsupervised learning	63
2.5.3	Semi-supervised learning	65
2.5.4	Reinforcement learning	67
2.6	Neural networks	69
2.6.1	The multilayer perceptron	70
2.6.2	The backpropagation of error training algorithm	71
2.6.3	Recurrent neural networks	74
2.6.4	Self-organising maps	77
3	Methods: the iCub Robotic Platform and Custom Software	81
3.1	The iCub humanoid robotic platform	81
3.2	Communications middleware: YARP	85
3.3	Auxiliary software created for this project	88
3.3.1	iCub-S a GUI interface for iCub resources	88
3.3.2	Auxiliary Python libraries	90
4	Experiment 1: Monocular and Binocular Contributions in a Bimodal Reaching Task	93
4.1	Methods	94
4.1.1	A visual perception system and tracking behaviour for iCub	95
4.1.2	Experimental conditions	100
4.1.3	Arm neural controller	100
4.2	Results	105
4.2.1	Analysis	106
4.3	Discussion	109
4.3.1	Contribution of binocularity	109
4.3.2	Contributions of proprioception in vision	111
4.4	Conclusion	113
5	Experiment 2: Posture and Arm-Modification Contributions to Adaptive Reachability Assessment	115
5.1	Methods	119
5.1.1	Visuomotor system	119
5.1.2	Neural controller	121
5.1.3	Learning process	122

5.2	Results	126
5.2.1	Analysis	132
5.3	Discussion	135
5.3.1	Postural contributions during the task	135
5.3.2	Role of neural network used	136
5.3.3	Confidence in reaching assessment	139
5.3.4	Relation to tool-use and the plasticity of peripersonal space	139
5.4	Conclusion	141
6	Experiment 3: Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom	143
6.1	Methods	148
6.1.1	Robot perception and arm controller	148
6.1.2	Experimental conditions and description of task	150
6.1.3	Stages in the Dev condition	152
6.1.4	Single-stage NoDev condition	154
6.1.5	A note on time in respect to the iCub simulator	154
6.2	Results	156
6.2.1	Learning processes	156
6.2.2	Evaluation	156
6.2.3	Analysis	157
6.3	Discussion	161
6.3.1	Benefits of staged learning	162
6.3.2	Performance and generalisation	163
6.3.3	Observations with respect to the training processes	164
6.3.4	Spacial exploration dependency on time	165
6.3.5	Motor synergies exploration	166
6.4	Conclusion	167
7	Overall Discussion and Conclusions	171
	List of references	181
	Bound copies of published papers	209

List of Figures

	Page
2.1 Decompositions for a mobile robot control system. a) Functional decomposition b) Behaviour based decomposition (reproduced from Brooks 1986b).	53
2.2 The basic element of a neural network.	71
2.3 A simple recurrent neural network.	75
2.4 A simple recurrent neural network with an architecture typical of Elman networks.	75
2.5 A Hopfield network with 10 connections.	77
2.6 A SOM with a three-dimensional input layer and a two-dimensional, hexagonal lattice, output layer.	79
3.1 The iCub robot	82
3.2 The iCub Simulator	83
3.3 YARP allows to have several remote clients for different parts of the robot. Clients can be in different computers even with different operating systems.	87
3.4 iCub-S, a program developed for being an interface for the cameras and for carrying out image processing was a multi-threaded program that uses YARP libraries and Qt for the GUI.	89
3.5 The iCub simulator running along with typical programs for accessing the robot: a robot motor GUI, an application manager and command line terminals.	92
4.1 The iCub head showing 5 of it's 6 degrees of freedom. The sixth one is vergence.	96
4.2 Graph showing the distance to vergence angle relation the tracking system exhibits	99
4.3 The simulated iCub performing the reaching task. Colours of arm and target are used for image segmenting.	100
4.4 The four layer, partially connected feed-forward network used for the arm controller.	102
4.5 The active vision process with the neural controller.	104
4.6 iCub Simulator reference frame. The horizontal plane consists of XZ and the origin of the axes is on the floor between the location of the feet.	106

4.7	Comparison of the two visual systems' performance in terms of reaching distance error. Bars indicate standard deviation.	107
4.8	End effector orienting error for the two cases. Bars indicate standard deviation.	108
5.1	The iCub's arm with two end effector lengths used in the experience phase.	120
5.2	The neural controller architecture. Two output units use one-hot (winner-take-all) encoding which indicate whether the target is considered in reach or not.	122
5.3	Error during training and corresponding validation error for both conditions. Only one of the twelve networks for each condition is shown as the other eleven displayed very similar training profiles.	125
5.4	A view of the robot's peripersonal space with three different effector lengths (short, medium, long from top to bottom) in the two conditions. On the left the robot trained in condition A (arm and vision) and on the right condition B (vision only). Red points indicate the robot assessed the point as reachable. Axes XZ define the horizontal plane. Units are metres.	128
5.5	a) Vergence angles at which the robot considered the target became unreachable for the three different arm lengths. b) Confidence-decrease peaks were present just before the reachable limits. Confidence in reaching assessment was measured as the absolute difference of the one-hot output units used to indicate reachability. Confidence-decrement peaks became stronger as the length of the arm increased. . . .	129
5.6	View of the peripersonal space of one robot for each condition with the 3 arm lengths (short, medium, long from top to bottom in each subfigure). Condition A (arm posture experience) above, B (fix arm) at the bottom. Axes XZ define the horizontal plane. Units are simulator metres. Colours are used for better displaying 3D locations.	131
5.7	Descriptive statistics of the perceived maximum reachable distance for the two conditions and three different arm lengths. Means are indicated by larger dots.	133

5.8	Top-view of operational space showing levels of confidence in reaching assessment for both conditions. Blue squares are regions the robot was able to foveate, grey where it could not. Objects were located in different heights and the average confidence in reachability assessment of targets in each discretised column is shown. Confidence was measured as the absolute difference between reachable/unreachable output neurons. Darker regions indicate lower confidence. Confidence in reaching assessment decreased as a function of arm-length.	134
6.1	Arm joints used in the experiment and their rotation directions. In the Dev condition, only the two most proximal joints (0 and 1) were used in a first phase and later on the two most distal ones (farther from the torso, joints 2 and 3) were included in a second phase. All four joints were used in the single-phased NoDev condition.	149
6.2	The learning path the iCub followed in each of the two testing conditions.	150
6.3	Images from the robot’s two cameras once the controller has foveated the target. Above, the original images. Below, the low-resolution colour-segmented images.	151
6.4	Mean squared error during the training of the first stage of development. Training sets for this learning were generated using two degrees of freedom, the other two had fixed values of 0° and 50°.	157
6.5	Mean squared error during the training of the second stage of development. Training sets for this learning were generated using four degrees of freedom of the arm.	158
6.6	Mean squared error during the training of the single-stage NoDev condition. Training data for this condition was generated using four degrees of freedom of the arm.	159
6.7	The iCub performing the reaching task once it has foveated the target. The robot used two arm joints in the first stage of development and four in the second in the exploratory processes and the four for the evaluation.	159
6.8	Mean distance from the centre of the palm to the centre of the target for each robot in the three conditions after three hundred trials. White bars indicate the mean distance for the three robots in each condition. Error bars indicate standard deviation.	160

6.9 Percentage of touching success. Each robot made three hundred reaching attempts to targets in locations reachable with four degrees of freedom. Shaded columns report the success percentage of each robot in the three conditions. White columns correspond to mean success for each condition. 160

List of Tables

	Page
3.1 iCub Joints	84
3.2 The modules available in iCub-S auxiliary program.	89
4.1 Robot arm joints used in experiment 1.	95
4.2 Performance of the system in the two conditions compared using <i>t</i> -test analysis.	108
5.1 Data in the training sets for each of the two conditions.	124
5.2 Measurements of maximum perceived reachable distance.	132
5.3 Results of one-way analysis of variance for comparing the three arm lengths in both conditions found no statistical differences between them. Condition A included arm posture-for-reaching information and posture of head and eyes in the training. Condition B only used head and eyes postural information and a “neutral” arm posture.	133
6.1 Robot arm joints used in experiment experiment 3.	149
6.2 Description of the experimental conditions.	155

List of Code

	Page
2.1 Pseudo-code for the backpropagation of error training algorithm	74
3.1 Basic usage example of auxiliary library <i>pySalo.syparp</i> used in a Python script. A <i>DataSource</i> object is created.	90
3.2 Basic usage example of auxiliary library <i>pySalo.sutils</i> . The script creates a <i>RobotPart</i> object for accessing the robot’s head.	91
4.1 Pseudo-code for the tracking behaviour algorithms implemented in the motor-control module of iCub-S.	97

Acknowledgements

I want to thank Paola, for all her love and support, my family and all my friends. Also, I would like to express my gratitude to my supervisors and examiners.

This work has been carried out with the support of CONACyT (Mexico) under the supervision of Dr Davide Marocco and Prof Angelo Cangelosi.

Author's declaration

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award, without prior agreement of the Graduate Committee.

Work submitted for this research degree at the Plymouth University has not formed part of any other degree either at Plymouth University or at another establishment.

This work has been carried out by Salomón Ramírez Contla with the support of CONACyT (Mexico) and the ITALK project under the supervision of Dr Davide Marocco and Prof Angelo Cangelosi.

Conference Presentations and Posters

Adaptive Reachability Assessment in the Humanoid Robot iCub.

Poster presentation at the Twelveth International Conference on Epigenetic Robotics.

Osaka, Japan, 2013.

Peripersonal Space in the iCub Robot.

Poster presentation at the RobotDoc International Conference on Development of Cognition.

Osaka, Japan, 2013.

Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom on the iCub Humanoid Robot. (*Best PhD Student Author* RobotDoc award)

Oral presentation at the Postgraduate Conference on Robotics and Development of Cognition RobotDoC-PhD.

Lausanne, Switzerland, 2012.

Active Vision and Depth Estimation Toward a Peripersonal Space Encoding in a Humanoid Robot.

Poster presentation at the Tenth International Conference on Epigenetic Robotics.

Lund, Sweden, 2010.

Publications

Peniak, M., Marocco, D., Ramirez-Contla, S., and Cangelosi, A. (2009). An active vision system for navigating unknown environments: An evolutionary robotics approach for space research. In *Proceedings of IJCAI-09 Workshop on Artificial Intelligence in Space*, Pasadena

Ramirez-Contla, S., Marocco, D., and Cangelosi, A. (2010). Active vision and depth estimation: Towards a peripersonal space encoding in a humanoid robot. In *Proceedings of the Tenth International Conference on Epigenetic Robotics*, pp. 175–176, Lund. Lund University

Peniak, M., Morse, A., Larcombe, C., Ramirez-Contla, S., and Cangelosi, A. (2011). Aquila: An open-source GPU-accelerated toolkit for cognitive and neuro-robotics research. In *Proceedings of The International Joint Conference on Neural Networks (IJCNN) 2011*, pp. 1753–1760. IEEE

Ramirez-Contla, S., Cangelosi, A., and Marocco, D. (2012). Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom on the iCub Humanoid Robot. In *Proceeding of the Postgraduate Conference on Robotics and Development of Cognition (RobotDoC-PhD)*, Lausanne

Ramirez-Contla, S. and Marocco, D. (2013). Adaptive Reachability Assessment in the Humanoid Robot iCub. In *Proceedings of the IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob 2013)*, Osaka

Word count of the main body of this thesis: 41,290

Signed: _____

Date: _____

1 *Setting the study context*

1.1 *Introduction*

Robotics is concerned with the study of machines which can substitute humans in the execution of physical and decision making tasks (Sciavicco and Villani 2009). We build robots because they can be helpers in risky or dull tasks. However, seamless robot and human collaboration has not been achieved and is still a goal actively pursued by many roboticists (Bauer *et al.* 2008; Bicho *et al.* 2011; Sofge *et al.* 2005). Humans constantly create tools to facilitate the execution of difficult or tiring tasks (Jonassen 1992). Although robots have been seen as tools in the past, today the perception and expectations are changing. We want them to be collaborators in human social environments, and for this, the interaction with them “should not be a chore” (CMU Social Robots Project 2006). We have already seen important trends in robotics during the last decade, aiming to the possibility of robots present side by side with humans cooperating and interacting (Bekey 2005), and having robots that can understand and adapt to their surrounding in an intelligent manner. In order to achieve this goal the development of robots has to include research from cognitive sciences and psychology.

During the last decades, robots have slowly but constantly shifted towards new environments beyond the structured ones of manufacturing plants (Khatib *et al.* 1999). Robots have moved out from completely industrial

environments such as *black factories*¹ and made their way into unconfined environments. Also, robots are now employed in deep-sea exploration (Kunz *et al.* 2008) and on the surface of other planets (Bell 2012; Liang *et al.* 2012). However impressive the achievements of robotics might seem, it is worth noticing that the environments in which these types of robots operate are human-free, they are not human environments, and the tasks for which they were built have not required a human-like shape.

Today, one main ambition of robotics is bringing robots into human environments. The expectation is to have robots displaying human-like behaviours and skills, and which can directly interact with us in human society. Although having robots for entertainment is one of the motivations for this, there is also the need for human-assistance robots (Roy *et al.* 2000) that are able to naturally interact² with us, for instance, at home, in hospitals or offices. However, social environments display highly dynamical properties and operating within them requires more than automated movements or tele-operation. Adaptability and autonomy are therefore necessary features for robots to enter this “final frontier”.

According to the robot taxonomy proposed by Fong *et al.* (2003), robot morphology can be anthropomorphic, zoomorphic or functional. A functional morphology is neither human nor animal-like, but is related to the robot’s function Yanco and Drury (2004). Although this categories were meant to apply to socially interactive robots, they also apply to robots before the HRI emerged. Robot morphology, in contrast to robot controllers which can be

¹Black factories are industrial environments where work is done by robots and where due to the absence of humans, lights are not needed (Murphy 2000).

²In a natural interaction, humans should not behave differently when interacting with a robot than when they interact with other humans (Fong *et al.* 2003).

optimised more easily, has for a long time been an issue solved by heuristics (Pfeifer *et al.* 2007). From this, it is evident that industrial robots traditionally have been built following a functional morphological design: their shape is intended to facilitate the execution of the task they are expected to carry out. Their shape, therefore, determines their capabilities but also their limitations. Traditionally, robotic shapes have not been anthropomorphic. In the search of robots that can be present in society it is important to consider the shape they are going to be built with. A flat-shape vacuum cleaner robot is good for getting under tables and beds but if it was intended to be in society and interact with people its shape would be an obstacle for natural interaction. Even if we could make this robot completely autonomous, the paths it will take to tackle a problem are different to those a human would take. A question can be raised: when is it necessary for an autonomous robot to have a humanoid body?

In the history of robotic systems, our understanding of intelligence has played an important role. Cognitive science is the study of the mind and intelligence (Thagard 2005). It is a multidisciplinary field where computer science and psychology have been very influential since its beginnings (Bechtel *et al.* 2001). Cognitive science has contributed to advances in artificial intelligence (AI), and the creation of intelligent systems and robots as both fields are closely related. For example, theories of human behaviour and cognition have found good platforms in robots and computers for being put to the test. A milestone of early cognitive science was Shakey (Nilsson 1984), a mobile robot created in 1966 that served as a theoretical framework for studying computer vision, language processing and robotic engineering. Around that time traditional artificial intelligence was developed and it had

high-level cognitive skills, such as planning processes, as its centre of interest. Therefore, robot designers focused on these for their implementations. Criticism to the traditional AI began with Dreyfus (1972). Some years later Brooks (1987, 1990, 1991b) made many further criticisms and he is considered one of the main contributors to the called *new artificial intelligence*, proposing the concepts of situatedness—the here and now influencing behaviour of the agent, without representations, using the world as its own model—and embodiment—experience comes directly from the world, and actions provide immediate feedback. A recent view of cognition (Pfeifer and Bongard 2007) suggests that the way in which the human mind works and solves problems is largely determined by the shape of the body, and that artificial intelligence can only be achieved by a system with sensors and actuators connected by a body. In other words, if we are looking for robots that develop and exhibit human-like intelligence, we should provide them with a human-like body³. Currently this view is being investigated by many scientists (Asada 2011; Hoffmann and Pfeifer 2012; Kuniyoshi *et al.* 2004; Ziemke 2003).

One characteristic of humans' interaction with objects and with other humans is the close physical proximity at which it occurs (Edsinger and Kemp 2007). It is expected that in a near future robots present in society will also display this type of close proximity interactions with humans. Therefore, it is necessary to investigate how they can develop perceptive and understanding capabilities for their near—also called peripersonal—space.

³The scope of this work is the development of behaviours, adaptability and autonomy, a different debate and research regarding the responses humanoid and non-humanoid robots produce on people concerns the field of Human-Robot-Interaction.

Peripersonal space is defined as the space immediately surrounding the body (Rizzolatti *et al.* 1997). It is the space where we can move our limbs and reach objects without the need to walk or change location. This near space is distinguished from the far space—the one that is beyond the reach of our limbs—because in the former we perceive potential interaction, as when reaching or grasping an object, or simply when we need to avoid colliding something. This has been measured by monitoring motor preparation using EMG (Serino *et al.* 2009) or reaction times (Costantini *et al.* 2011). Our actions in the peripersonal space are the first contact we have with our environment, therefore, our brain needs to be aware and react to objects and events in it and so robots will need it.

Behavioural and neurophysiological studies (Holmes and Spence 2004 for a review) suggest the brain encodes peripersonal space differently from extrapersonal space, and the existence of a multisensory peripersonal space representation encoded in body part-centred coordinates in cortical and sub-cortical brain areas. The representation of peripersonal space adapts according to the possibilities offered by tools or the shape of the body (Berti and Frassinetti 2000; Farnè *et al.* 2007; Holmes *et al.* 2007). Also, as we grow, our body goes through changes of its physical characteristics (the length of our limbs extend) and at the same time the way in which we perceive potential interaction is shaped accordingly. In a similar manner, the use of tools also originates perceptual changes of what can be reached or not (Farnè *et al.* 2007).

Traditionally, robots have not used representations of their peripersonal space. One reason for this is that for automated tasks, in factories for instance,

the surroundings are tightly controlled⁴ and therefore it is not necessary to monitor the surroundings of a robot manipulator. Another reason is that until recently robotics have been more focused to the development of mobile systems, which need representations of the far space—*i.e.* spatial maps for navigation—and not of the close one. The evolutionary (Fernandez-Leon *et al.* 2009; Floreano and Mondada 1996; Mata *et al.* 2005; Miglino *et al.* 1995) and learning (Asada *et al.* 1995; Fox *et al.* 1998; Smart and Pack Kaelbling 2002) approaches are among the ones traditionally used in mobile robots. Because a new goal of robotics now focuses on human-like interaction, it is important to direct the attention to the recent findings on human cognition development and perception as they propose novel accounts on how peripersonal space is represented by natural systems, and how action and perception contribute to this effect.

In the classical view of action and perception, these two processes were considered to be separated from each other. Cognition was thought to happen in between these two processes, as if it was an interface for going from one to the other (Hurley 2001). However, recent evidence from neurophysiology (Di Pellegrino *et al.* 1992; Rizzolatti *et al.* 1997), the discovery of mirror-neurons, and new accounts of visual perception (Gibson 1966, 1979) and consciousness (O'Regan and Noë 2001) suggest that these two processes should be considered together. These are neurons located in the monkey brain's motor cortex that activate both when the animal performs a task and when it sees another monkey performing the same or a similar task. It is known that multiple sensory modalities are involved in this phenomenon. In

⁴Separation of human and robot workspaces is defined in guidelines for robots in industrial environments ISO10218

the study of peripersonal space, Rizzolatti *et al.* (1997) found there are neurons that respond to tactile stimulus on the limb but also to visual stimuli near that part of the body, regardless of the location of the limb. Also, Làdavas and Serino (2008) and Coello *et al.* (2008) have found psychophysical evidence of how the visual perception of peripersonal space is modulated by the motor representations acquired during action execution. There is clear evidence of the influence action has on perception and vice-versa.

In the fields of robotics and AI, the view that separated action from perception was also dominant for many years in a similar manner as it was in psychology. And also in these fields, later on a new approach considered these two elements as complementary to each other and simultaneous. The *hierarchical paradigm*⁵ (Murphy 2000 , p. 41-65) proposed that the three accepted primitives of robotics, sense, plan, act occurred in a hierarchical manner. Planning having sensing as a prerequisite and acting needing planning. Later on, around 1988, a new approach was proposed. The *reactive paradigm* (Murphy 2000 , p. 67-99) took into account the effects of new sensory data and directly fed these to the actuators. Examples of this latter paradigm are active-perception (Bajcsy 1988) and active-vision (Aloimonos *et al.* 1988). Active vision is an area of computer vision where simultaneous action and perception promote the exploration of both motor and perceptual capabilities simultaneously. Although the limitations of this new paradigm became evident after some years because, it took the no-planning idea too far, it led to new advances in robotics by bringing relevance to the action component of robotic-perception. The paradigm applied to the study or

⁵Robotic paradigms are the philosophies or ideas defining a way in which an artificial agent makes decisions and/or acts.

peripersonal space is still a subject that demands investigation because it embodies elements currently explored in other fields of science that are very relevant to robotics.

The investigation of how action and perception are closely related will be a step forward in the creation of autonomous robots and allow them to explore their surroundings in active ways and follow a development closer to the one that humans do.

We have mentioned how robots are expected to be present in human environments and interact in more human-like ways. We want them to be adaptable and develop skills as they interact with their surroundings. That is the motivation for this work. We want to investigate an embodied robotic system with mechanisms that allow it to develop bio-inspired representations that aid skills in the peripersonal space. With the present work I intend to shed light on the following research questions related to peripersonal space:

- How can the research findings on peripersonal space from psychology and cognitive science be implemented in the iCub?
- Are these implementations useful for endowing the iCub with adaptive capabilities for representing the space around it similar to those observed in humans?
- What are the contributions of morphology and sensing capabilities in the creation of action-based representations for peripersonal space assessment?

1.2 Aims and objectives

1.2.1 Thesis aims

1. To find mechanisms that can provide a humanoid robot with behaviours present in humans and described in psychological literature related to the development of cognition and the encoding of peripersonal space.
2. To explore bio-inspired methodologies and machine learning techniques on artificial systems that might give an insight to the origins of vision-based motor control strategies within the peripersonal space in humans.

1.2.2 Objectives

To achieve these aims, the following objectives were defined:

1. A review of the existing literature on robotics and peripersonal space. Study recent methodologies for the creation of intelligent robotic systems, classify and identify the original contribution of this thesis.
2. The integration of visual, proprioceptive and tactile sensory modalities in the iCub for creating reaching actions.
3. The implementation in iCub of methods for assessing the reachability of targets inside the peripersonal space using vision-based action.
4. The investigation, in the mentioned implementation, of contributions from different sensory modalities, learning conditions and body characteristics in the reaching performance and codification of the surrounding space.

1.3 *Dissertation outline*

- *Chapter 1* is the present introduction to the thesis.
- *Chapter 2* provides the necessary background on the topics related to this study and a literature review of relevant and recent research on cognitive robotics, peripersonal space in psychology, peripersonal space in robotics and active-vision. By doing this, the framework for the studies is presented and the terrain is prepared for going deeper into these topics in the following chapters.
- *Chapter 3* introduces the research platform and materials used for this study. It presents details of the robotic platform iCub, the communications middle-ware used, the software developed during the research work, and chosen artificial intelligence technology used in the experiments.
- *Chapter 4* presents an experiment that explores and compares monocular and binocular vision in a task involving the visual and proprioceptive modalities. In this chapter the visuomotor system used as basis for this and the two other experiments in this thesis is detailed and put to test.
- *Chapter 5* deals with the creation of implicit representations of peripersonal space. The experiment introduces a model for reachability assessment implemented in a simulated iCub humanoid robot. We explored the perceived reaching range in presence of three arm lengths and in two training conditions with different type of information provided. We implemented a model for reachability assessment on a

robotic system, and studied the implications of it having or not postural information related to reaching when categorising the reachable from the unreachable.

- *Chapter 6* explores the benefits of progressively unlocking degrees of freedom of the robot's arm in the development of motor control for interaction with object in the peripersonal space. Following a proximodistal development, a reaching skill learning process is modelled and compared to learning in a non-staged process.
- *Chapter 7* gives overall conclusions and presents a general discussion of the work.

2 Background

As we have discussed about in chapter 1, in the near future we expect to have humanoid robots collaborating with humans in our society. For the development of these robots it is important to take into account the empirical findings regarding how we, humans, move in the world. In the present chapter we detail the concepts and recent findings related to our study, including cognitive robotics and the embodiment approach, a strongly influential approach applied in robotics today that contributes to the present work. This review is essential for putting our research on peripersonal space into context. For our work we were based on psychological findings related to the integration of visual and motor capabilities (both faculties present at early cognitive development in natural systems) when implementing experiments in a robot, therefore, in this literature review some sections present a related psychological aspect followed by a corresponding robotics' aspect and a brief comparison of both.

2.1 Robots

In the search for understanding ourselves, humans have always been fascinated with the idea of an artificial person (Duffy 2003). This can be traced back to legends such as Pygmalion or the Golem, it is a recurring theme throughout classic and modern fiction literature (Brooks 1996). For a long period of human history the available means of production and the economic

and social conditions were not favourable for the construction of “intelligent” machines (Kurfess 2010). It was in the early eighteenth century that clock-makers and engineers started building automatons, mechanical devices that imitated animals or humans. Later on, it was the industrial revolution what provided new means for scientific and technological advances that would allow the inception and creation of artificial systems. Electronic and digital systems began being developed in the late nineteenth century. In the twentieth century, because of WWII these technologies were hugely developed (Kurfess 2010). Robots in the modern age came into scene thanks to these technological advances, and during the last five decades, robots have been increasingly present in modern life. In this period, they have slowly but constantly shifted towards new environments and are expected to be in human social environments in a near future. It is important to realise that throughout the history of robotics, the way in which robots have been built has been inspired by different ideas and so the target environments have resulted accordingly to them. Therefore, for the robots to reach human social environments it is necessary to build them with the concepts, knowledge and characteristics that these environments need. This important fact will be present throughout the rest of the chapter and it will be discussed.

In order to have a glance of how robots have changed during their history, lets first look at some definitions. According to the Robot Institute of America (RIA) an industrial robot is “a programmable, multifunctional manipulator designed to move material, parts or specialised devices through variable programmed motions for the performance of a variety of tasks” (Jablonowski and Posey 1985). In the face of recent developments on robotics

and the fact that they have moved into more environments, there are new definitions for what a robot is. Xie (2003) proposes that a robot “is the embodiment of manipulative, locomotive, perceptive, communicative and cognitive abilities in an artificial body, which may or may not have a human shape. It can advantageously be deployed as a tool, to make things in various environments”. In these newer definition we can notice important points. The inclusion of more human abilities and the mentioning of the shape of the robot. In scientific fields such as cognitive science, the idea of what cognition and intelligence are has changed. Those changes are reflected also in the field of robotics. Today, as more things are expected from robots, like communicative abilities, also the shape of the system has become important. Drawing a conclusion that incorporates previous ideas of this section, robots in human environments are expected to display more human capabilities. If we want them to interact with us and collaborate, the shape they will have will be an important factor. We will see later on, that not only the shape has relevance because it enables robots with physical abilities that are more human-like and improve face-to-face interaction with humans, but also because the way a system develops cognitive capabilities is influenced and shaped by the body it has.

In order to see robots as equal partners we have tried making them more human-like. As robots get more into social environments, they have acquired human-like bodies, these robots are called humanoids. Xie (2003), extends the definition of a robot and now defines a humanoid robot as “the embodiment of manipulative, locomotive, perceptive, communicative and cognitive abilities *in an artificial body similar to that of a human*, which

possesses skills in executing motions with a certain degree of autonomy, and can be advantageously deployed as agents to perform tasks in various environments”.

We have seen how the definitions of what a robot is display a tendency to make them more similar to us and also to be able to perform the type of tasks we carry out. For getting to that point, the approaches have also shown a tendency to get closer to biologically plausible methods for building them. Now, we will briefly review three different bio-inspired approaches taken in the search of intelligent and/or autonomous robots. This is for us to mention how each of these relates to our own approach and how they relate also to the theories of active-perception and embodiment.

2.1.1 Humanoid robots

Humanoid robots are robots with an anthropomorphic design. The motivation for creating human-like robots is having intelligent systems that can get around in a human world (Appin Knowledge Solutions 2007). Such systems are endowed with the functional mobility of the human body and therefore can deal with human-familiar objects and thus are considered good candidates for participating in collaborative tasks with humans in many human activities.

The quest for developing humanoid robots has been long. Leonardo da Vinci already drew and envisioned a humanoid mechanism in the late 15th century (Rosheim 2006). In the 18th century the Jaquet-Droz family built automata (self operating machines) able to carry out human-like activities such as writing or playing musical instruments (Rosheim 1994). Later on, in the 19th, humanoid robots Steam-man and Elektro, robots powered by steam

and electricity respectively, were developed (Akhratuzzaman and Shafie 2010) paving the road for the robots of the 20th century. Along with the creation of electronic devices and general-purpose digital computers, robotics saw important advances during the last century. During this period it can be observed that interest on creating human-like robots has been a constant. Starting in the late 60's, the robotics team at Waseda University in Japan developed a whole family of legged robots (Lim and Takanishi 2007; Yamaguchi *et al.* 1993) aimed to research bipedal walking and robot interaction with the environment. Integrated humanoid robots were later developed with the intention of building general purpose systems (Hirai *et al.* 1998; Morita *et al.* 1998; Takanishi *et al.* 1998).

In the last decade development of humanoids has continued. In Japan, Honda company developed ASIMO (acronym for Advanced Step in Innovative Mobility) (Hirose and Ogawa 2007) which was presented in 2000 and aims to provide a robot that can assist people in indoor environments (Sakagami *et al.* 2002). More recently other humanoids have been developed by the scientific research community taking an approach more bio-inspired, rather than an engineering one. MIT, for instance, in the 90's began building humanoid robots for exploring theories of human intelligence (Adams *et al.* 2000) starting with Cog (Brooks *et al.* 1999) and later with Kismet for studying social interaction with humans (Breazeal 2003). A recent example is the iCub robotic (Metta *et al.* 2008; Vernon *et al.* 2007). The iCub was developed in order to have a standard platform for studies of human intelligence. The idea is to have a ready-to-use system that for experimental replication and allow researchers in different locations to independently use and test the same

system.

Many scientists working with bio-inspired methodologies consider the development of humanoid robots as the ultimate goal of robotics. The idea that motivates this goal is that humanoids can be important in the research of human-like behaviour (object reaching, grasping, manipulation) (Floreano and Mattiussi 2008). Moreover, there are other motivations for building humanoid robots. Next we present a field of robotics that deals with the design of robots from a user-perception perspective.

2.1.2 Human-Robot Interaction

Human-Robot Interaction (HRI) is the area of robotics that studies the design of robots used by or that interact with humans. It also is dedicated to create methods for evaluating this interactions (Goodrich and Schultz 2007). Traditionally, robots have been used in environments that require very little or no interaction with humans, however, new applications have emerged recently that make the ability to interact with humans an important part of robots' functionality (Breazeal 2004). HRI draws inspiration from work on anthropomorphic interfaces, cognitive science, psychology, social sciences, artificial intelligence, human-computer interaction, computer science, robotics and engineering (Dautenhahn 2007; Kidd and Breazeal 2005) for investigating natural means by which a human can interact and communicate with a robot (Dautenhahn 2007).

Applications of HRI include: search and rescue, assistive robotics, personal service robotics, education, entertainment. These applications can benefit from robot designs that foster the interaction of robots with humans counterparts. In HRI, robots can play different roles. Scholtz (2003) lists these

roles as supervisor, operator, mechanic, peer and bystander, and Goodrich and Schultz (2007) add mentor to this list.

As we can see, a robot can play different roles when interacting with a human counterpart, therefore, the types of interaction also can change. According to Goodrich and Schultz (2007), depending on the type of communication taking place between humans and robots, there are two general categories of interactions: remote interaction, where the participants (human and robot) can be spatially or even temporally separated, and proximate interaction, where both participants are close to each other (in the same room, for instance). For many application areas of HRI, current research tends to focus on proximate interactions, as it is more interested on social interactions that by nature are proximate rather than remote (Goodrich and Schultz 2007).

Breazeal (2003) proposed four categories for social robots. Ordered by their ability to get around the human social environment, they are: socially evocative, social interface, socially receptive, and sociable robots. Socially evocative robots make use of the tendency to anthropomorphise objects in order to evoke an emotional response. Social interface robots can detect and produce natural interaction modalities such as gestures and speech. Socially receptive robots have at least some level of social cognition so that they can learn from social interaction, although they are not yet fully socially functional. Sociable robots, by contrast, will be able to seek interaction by themselves in order to satisfy internal states similar to those humans have.

For Goodrich and Schultz (2007), human-robot interaction problems depend on various aspects that integrate the interactions between the one or

more humans and one and more robots. This aspects are: the level of behaviour of autonomy of the robot, the nature of the information exchange, the structure of the team, the adaptation or leaning of people and robots and the shape of the task.

Although the present dissertation does not focus on human-robot interaction issues in the study of reaching and peripersonal space, HRI is mentioned due to its importance in future humanoid robots' design. Ultimately, this issues are essential for successfully introducing robots into the human environment and every robot designer should be fully aware of this area of robotics.

2.1.3 Evolutionary robotics

Evolutionary robotics is a field of research that investigates and uses artificial evolution in the creation or synthesis of robot sensors and controllers. In the 80's and 90's, ideas about automatic generation of control systems like the one Turin (genetic combination search) proposed in the 50's (Turing 1950) were brought back in an analogue to Darwinian natural evolution which seeks to design artificial agents' bodies and/or sensorimotor systems. This came as a response to the failure of GOFAI (pure symbol manipulation) and as a search for more biologically plausible computational processing. In this approach, and highly influenced by McClelland and Rumelhart (1986), connectionist models such as neural networks resurged for the exploration of cognition. The use of such models made some people consider brains to indeed perform computation but not computation in the sense proposed by Turing. Ideas like the thought experiments proposed by Braitenberg (1986) also contributed to this new view with his neurally

controlled *vehicles*, describing how intricate behaviours can emerge from small “nervous” circuitry. Shortly after, empirical work from Nolfi *et al.* (1994), Nolfi and Floreano (1998), Floreano and Mondada (1996) largely contributed to the field by testing the new hypothesis in real and simulated robots and using the robot platform *Khepera* (Mondada *et al.* 1999) as an important tool.

In the evolutionary robotics approach an agent can be considered a dynamical system which when coupled to the environment has to display an adaptive behaviour. To achieve this, a designer could make this by hand, however, this could be a task requiring an enormous amount of effort. In the first place because the initial design of the dynamical system is not intuitive, and secondly because it is even less intuitive once it is coupled to the environment. So why not leave this task to the environment? By means of Darwinian evolution an agent with a neural controller, for instance, can reach a stage where it adapts to a required task (Floreano and Mondada 1996). By having an appropriate genetic encoding and applying variation, heredity and selection to a population of agents, the system can achieve an artificial brain or nervous system that complies to a fitness function describing the task (Nelson *et al.* (2009) for a survey of fitness functions).

2.1.4 Bottom-up approaches to adaptive behaviour

The *animat approach* (e.g. Wilson 1985), *behaviour-based* (e.g. Maes 1994) and *behaviour-oriented* (e.g. Steels 1993) robotics are fields studying adaptive behaviour and which overlap to some extent (Ziemke 1998). Behaviour-based approaches emerged as an alternative to the previous approach of knowledge based. The animat approach uses the construction of artificial animals called animats (Wilson 1985) for understanding the

mechanisms that allow survival in changing environments. Design of animats was inspired by behaviours observed in real animals (Webb 2009). As a long term objective, the animat approach also seek to provide a contribution to the understanding of human intelligence by progressively taking the learnt lessons into the field of human behaviour (Brooks 1997).

In the pursuit of its aims, bottom-up approaches put together research from psychology, ethology, ecology, robotics, AI and engineering. Generally, it used a bottom-up approach which also became common at the time in cognitive science and had as key issues the autonomy and the adaptive control composition of behaviour (Ziemke 1998).

The animat approach used different techniques for its exploration of intelligent behaviour, including dynamical systems, learning and connectionist models. With their work researchers like Webb (1995, 1998) and Brooks (1986a, 1989) were able to synthesise interesting behaviours in insect-like creatures and propose models for these behaviours. They argue that by studying these artificial and non-existing animals we can learn from how real animals adapt to uncertain environments by modulating simultaneous behaviours and building on top of persistent ones. Eventually, Brooks *et al.* (1999) built the humanoid robot *COG* using principles for further investigating the possible applications of the bottom-up approach and following principles such as the distributed control humans exhibit, the dispensability of complete representations or models, embodiment and the role of social interaction in development.

Guillot (2001) considers the contribution of the animat approach to cognitive science to be showing that these artificial animals were able to

become active information processors like any real animal, and they were not only passive reflex devices. Moreover, animats seek for useful information in their environment. In summary, animats provides a bottom-up viewpoint that complemented traditional AI but it also gave rise to the question of how far can researchers get by synthesising behaviours.

One of behaviour-based robotics was the development of the *subsumption architecture* (Brooks 1986a). It describes layered asynchronous and parallel behaviours that contribute to achieve a goal or a set of goal behaviours. The design of the layers can be one difficult aspect of the architecture, however, genetic algorithms are often used successfully to this end (*i.e.* Togelius 2003). This architecture has the advantages of being situated and modular, allowing to build real-time behaviours which had not been obtained before and was an important contribution to bottom-up approaches to adaptive behaviour because it introduced robot mechanisms to solve a problem in distributed and cooperative way and provided basis for further independent intelligent agents (Yu 2005).

2.1.5 Developmental robotics

Developmental robotics is a broad discipline that encompasses psychology, biology, artificial intelligence and robotics interested on the extended periods of development that biological systems must undergo in order to reach their adult form and abilities (Meeden and Blank 2006). Developmental robotics origins can be found in (Brooks 1991a) and Brooks and Stein (1994). It recognises the importance of taking into account development for constructing intelligent robotic agents (Sandini *et al.* 1997). Developmental robotics, looks to the properties that are acquired during the

interaction with the physical and social environment and the increasing complexity that the cognitive structures display (Berthouze and Ziemke 2003). A very important aspect of this approach is the attention paid to temporal elements of development (Schlesinger *et al.* 2008).

Developmental robotics and epigenetic robotics are considered to be mostly identical disciplines. However, it is considered that developmental robotics encompasses a broader spectrum of issues (Lungarella and Berthouze 2002). The diversity of work in developmental robotics is broad because the discipline is concerned with many aspects of development, ranging from the sensorimotor (*e.g.* Hulse *et al.* 2010; Sandini *et al.* 1997) to the cognitive (Asada *et al.* 2009; Lee *et al.* 2012b) and morphology (Jin and Meng 2011) facets. In order to better identify work within the discipline, Lungarella *et al.* (2003) proposed that for a study to contribute to the field of developmental robotics, it must provide clear evidence for experiments on physical robots and has to clearly put forward an hypotheses addressing developmental psychology or developmental neuroscience. Although in Lungarella *et al.* (2003) do not discard the contributions that computer models and simulations can provide to the field and the advantages they have over real-world experiments, they stress the importance of the influence from the interaction with the environment. Nonetheless, research on developmental robotics has also been carried out on simulated environments in many cases for studying sensorimotor grounding (Cangelosi and Riga 2006), mental imagery (Di Nuovo *et al.* 2013), integration of speech and action (Tikhanoff *et al.* 2011), perceptual development (Schlesinger *et al.* 2012), for instance.

To some extent, the developmental robotics approach contrasts with

evolutionary robotics. In evolutionary robotics population of individuals are studied to see their development over generations. By contrast, developmental robotics is more concerned with the changes that occur to a single individual during its lifespan. Moreover, although in evolutionary robotics sometimes the genetic development is investigated (Cangelosi 1999; Cangelosi *et al.* 1994), it is the result from genetic encoding development the focus of interest. On the other hand, in epigenetic robotics, the focus of interest is on ongoing emergence, autonomous development and/or bootstrapping of initial skills (Prince *et al.* 2005).

Work in epigenetic robotics has been fostered by engineers, psychologists and neuroscientists (Lungarella *et al.* 2003), and it has provided contributions to their fields by offering novel methodologies for more advanced robotics either in entertainment, industry or any other application, and by creating new tools for scientific investigation.

2.1.6 Relation to this work

In *evolutionary robotics* a population is evolved to find for agents that can solve a task. This is analogous as reaching a stage in natural evolution where we can observe the result of a process on that population, which has the ability to solve a task in the genetic information it received as result of the shaping evolutionary process. Although during the evolutionary process the individuals of the population recreate the series of steps that led to its solution, in this work the focus was on development at the level of single individual. Therefore, we did not employ generic algorithms.

Regarding the *animat* approach, in our studies, we took an approach similar to the *animat* one. We attempted to explore how the behaviours

resulting from simple visual and motor systems led to the emergence of the representation and use of an inexplicit representation of space that allowed a humanoid robot to explore, probe and adapt to bodily modifications.

For this study, we were interested in an approach that can lead to accounts of adaptation to body constraints through exploration and learning. We investigated the developmental process of a synthesised agent that carries out reaching tasks with different arm lengths in a *epigenetic robotics* approach.

2.2 Embodied intelligence

Throughout the history of intelligence research it has been difficult to come to a single or standard definition of intelligence (Legg and Hutter 2007; Sternberg 2000). Even today, many exist that regard either one or other aspects of behaviour or the mind for assessing its presence. Logic, self-awareness, planning, adaptation, learning, communication among many aspects are taken into account when investigating intelligence.

Representational-Computational (or traditional) cognitive science viewed the process of thinking as an abstraction of the physical mechanisms that allow us to interact in the world and then a manipulation of the resulted symbols. However many researchers in the late 70's started noticing and investigating phenomena that requires the inclusion of the body for an explanation, suggesting that cognition processes were deeply rooted in bodily interactions with the world.

The ideas of the body involvement for explaining intelligence became known as "the embodied theory". In its origins there are Lakoff and Johnson

(1980) and Varela *et al.* (1991), who proposed how the physical interaction which forcibly requires a body is a condition for systems to show intelligence. The embodied theory has been useful for further understanding cognitive phenomena as well as providing roboticists with a framework for implementing artificial systems which in turn also helps philosophy and psychology for testing their hypothesis.

The embodied cognition theory also tries to explain how changes during development have an effect in the way we perceive and react to stimuli. Thelen and Smith (1994) introduced dynamical systems theory into developmental psychology and proposed non-representational theories (and also anti-nativism) when they argued that the body in action can be a processing distributor and a guiding constraint for the development of behaviours like reaching or walking. They proposed bodily interaction in changing contexts as a source of development in contrast to pure physical maturation processes. This changing context was already in Gibson (1979) when he talks about vision perception taking place in a feature-rich environment. In this context, we can observe how when humans start our lives, our motor abilities are very limited. Our movements seem to have no aim or seem like uncontrolled reactions to the stimuli that the world out of the womb, our new environment, presents. Before achieving the degree of control a child shows, a baby has to follow a path of development that will help him to give sense to all the sensory input he receives and relate it to his own movements and reactions. For relating all this information humans possess a body, and of course its brain. Both immersed in a vast stimuli-rich environment but with specific characteristics that set constraints to the

relations to be made by the child in order adapt his behaviours and move in the world, adjust his actions to the environment.

Next we will present how these concepts have been applied to the field of robotics research that inspired the present work.

2.2.1 Embodied artificial intelligence

In the search of artificial intelligence a lot of work considered the capability of symbols manipulation a sign of intelligence. This view, although it allowed to create very sophisticated systems capable of mastering playing chess (Campbell *et al.* 2002) or *expert systems* able to give accurate diagnostics of problems and recommend a course of action, faced the problem of making those same systems perform tasks in the real world. This, due to symbols lacked a grounding in real physical objects, which is a necessary condition to really understand them. The machine did not understand the symbols it used, it was simply manipulating them, like someone in Searle's (1980) Chinese room, where a person without any knowledge of the Chinese language, provided with instructions on how to arrange the symbols he is provided with, can give written answers indistinguishable from ones from a people who knows the language. Finally, because the machine did not possess a body with which it could make any direct effect on the world, it is also impossible for it to solve practical, more everyday human problems, such as playing football, for instance.

The classical view of cognitive science and artificial intelligence (now known as GOFAI, Good Old-Fashioned AI) supported the idea that intelligence was in the controlling part of a system. The brain was seen as a main computer that processed inputs and produced outputs such as motor commands simply

by manipulating the symbols, in that vision brain and environment could be decoupled and the complexities of perception were overlooked. This way of thinking in AI was influenced by the current cognitive science at the time and therefore had a strong legacy from Cartesian dualism.

More recent research on artificial intelligence has emphasised the role that physical interactions with the environment play for the understanding of intelligence in natural systems. Pfeifer has been a big proponent of this new consideration in robotics. Under this new approach, he points the boundaries from low-level motor control and cognition become very fuzzy but the important issue is to look at the process of development from which this abilities emerge (Pfeifer and Bongard 2007). That is, we should look at the relation between the agent and its surroundings. For this to be studied, the physical shape of the system has to be taken into account, its embodiment. Different morphologies result in particular behaviours but also the type of sensors and the materials itself provide different perception capabilities.

Embodiment has been basis for studies on self-stabilisation during locomotion in bipedal (Blickhan *et al.* 2007) and insect-like walking (Dürr *et al.* 2003) but also the theory is being used to explore higher level behaviours such as categorisation (Morse *et al.* 2010a). On the same grounds, Hoffmann and Pfeifer (2012) propose the use of the bottom-up approach available in embodied agents as a pathway to follow for studying not only locomotion but also concepts such as the *body-schema*. The body-schema is, in classical neurology, a constantly updated sensorimotor picture of the body shape and posture (Head and Holmes 1911) that could allow also for creating a “representation” of the agent’s action possibilities (Bermúdez *et al.* 1995).

Hoffmann and Pfeifer argue that the representation of the body is the basis of cognition and it is in it where traditionally considered high-level processes like categorisation, memory and perception couple with the environment.

In summary, the research mentioned suggests elements of cognition such as memory, perception and categorisation can be grounded in behaviour and sensorimotor loops. The study of the representations of the body itself is a first stepping stone for the building of intelligent behaviours. The body schema is closely related to another concept which is the object of the present study: peripersonal space¹. Next we will present and discuss this topic as a way to further introduce our research.

2.3 Reaching development and peripersonal space

2.3.1 The development of reaching in infants

As we have seen, developmental robotics tries to provide new understanding of human development following a synthetic approach. Ongoing emergence is the continuous development and integration of new skills by an agent (Prince *et al.* 2005). Providing artificial agents the properties that give way to the emergence of skills and cognition is one of the aims of the field. For this, many developmental studies focus their attention on children.

In order to achieve successful use of vision for reaching, the development of human infants shows constraints at different ages. These constraints limit the possible actions the child can do as well as his perceptual capabilities. However, they also serve as a frame or platform that aids them in the process of understanding their environment (Rutkowska 1994). Hulse

¹Farnè *et al.* (2009) discuss if the two concepts are the same.

et al. (2010) and Law *et al.* (2011) review in detail several of those constraints and other elements that become present during the development in the first 12 months of newborns. Following, a list focusing on the aspects related to vision and reaching in the newborn is presented.

At birth, newborn children sensorimotor constraints include having a restricted visual system (Hainline 1998). Horizontal visual field width angle in neonates was found to increase very quickly after birth (Harris and MacFarlane 1974) and it has been shown it increases from 30° to 60° from week two to week 10 (Tronick 1972). Eyes' movement is the most controlled motor ability of newborn infants. There is debate regarding motor control of the arms in the newborn. Hand-to-mouth movements at birth were considered reflexive (Piaget 1952) but today there is evidence suggesting it might be intentional (Butterworth and Hopkins 1988; Rochat *et al.* 1988). Moreover, the random motor experimentation with the arms known as “random babbling” (Meltzoff and Moore 1997) is one step towards the sophisticated forms of imitation of 18-months-old infants (Rao *et al.* 2004).

Neonates also possess a grasping reflex. they can reach to targets in a ballistic manner. However, this visually initiated goal directed reaching known as “visually elicited reaching” (Rader and Stern 1982) is done through uncoordinated and ballistic movements. This kind of reaching is present during the first seven weeks after birth. This contrasts with the visually directed movements learned later on.

Newborn infants can only focus maximum distances of 21 cm. At that stage it might not be necessary to focus farther and interestingly, that distance is approximately the same at which their mother face is when they are carried

(Butterworth 2014).

Colour perception is poor in newborn infants but reaches the adult categorisation levels at about four weeks but with lack of clarity in the centre of the visual field. It is possible that this lack of clarity is why newborns are mostly attracted by diffuse lights, colours and moving objects within their focal range (Sheridan *et al.* 2007). In the first month, infants fixate on objects' edges (Maurer and Maurer 1988).

At two months of age focusing abilities increase and acuity is improved greatly (Oates *et al.* 2005) and the width of the field of view is increased to 40° around week 10 (Tronick 1972). They can perform visual search similar to adults in terms of saccades and fixation of the centre of objects Maurer1988. At this age infants can follow moving targets but very little head movement is done to this end.

At three months the infants can initiate reaching movements at will but the movements are still elicited by visual stimuli and not yet guided. At that moment they can employ head movements for visual search all the time for gaze shifts greater than 30° (Goodkin 1980).

At around four to five months infants' reaching and grasping are guided through visual feedback (White *et al.* 1964).

At month six the visual search abilities are mature and infants are very visually active, gazing for novel stimuli (Sheridan *et al.* 2007).

When infants reach nine months of age, hand-eye coordination is well developed (White *et al.* 1964) and they .

A characteristic of the development of motor control for reaching is its flow direction. For reducing the complexity of reaching movements, infants

use a limited number of degrees of freedom when they start learning the skill. The direction followed by the degrees of freedom used is from the centre of the torso towards the extremities. This is called a “proximodistal development”. This is due to the control of distal muscles of arm and hand is gradually acquired as the cortico-spinal tract, which is not functional at birth, matures (Kuypers 1981; White *et al.* 1964).

As we can observe from this summary, during the first twelve months of the infant’s life, the abilities, present from birth but not developed, go through a series of stages. From random experimentation to coordinated and visually guided movements. Due to the fact that action is a component in the creation and modification of peripersonal space representation (Brozzoli *et al.* 2009), it has been important to review the elements involved for the achievement of reaching.

2.3.2 Peripersonal space

If someone was asked to recall all the objects that he came across and had to avoid while walking he would very likely fail (Droll and Eckstein 2009). Our body constantly performs motor responses in order to avoid collisions and we are usually unaware of the processing taking place for this (Làdavas and Serino 2008). When walking, for instance, we don’t always have to pay attention to the irregularities of the floor or spend much time dealing with a step, or when someone come close to us rapidly in the public transport we naturally move slightly to avoid contact.

A related effect is when we feel uncomfortable if there is something near our face or our hand, even after no longer seeing the object causing the discomfort (Hall 1969; Kennedy *et al.* 2009). Response to these type of stimuli

are stronger in the space close to our body than further away, as evidence from extinction patients demonstrate (Farnè *et al.* 2005). Related to these phenomena and following research in neuroscience, (Rizzolatti *et al.* 1997) defined *peripersonal space* as the immediately space surrounding the body that can be reached with our limbs. The space beyond it is defined as extrapersonal space. A differentiation of these two spaces can be made on the basis of tasks performed in it: we move (or navigate) in extrapersonal space while within peripersonal space we examine things, use tools or perform movements to avoid harmful objects and protect our body (Làdavias and Serino 2008). When we move through space our brain needs information from the peripersonal space in order to avoid objects. In a similar way, our brain uses the information from the peripersonal space when we manipulate objects for guiding our limbs to desired locations and reach targets (Holmes and Spence 2004). Behavioural and neurophysiological studies suggest that the brain encodes peripersonal space differently from far space (extrapersonal) (Halligan and Marshall 1991; Rizzolatti *et al.* 1997). This discovery has turned attention because it also suggests that these two encodings are made in at least partially separate neural systems.

A finding that suggests an encoding of peripersonal space in the brain are the visual receptive fields for neurons in the ventral premotor cortex responsive to certain tactile receptive fields, regardless of the position of the eyes or the position of the body, there seems to be a coordinate system determined not by the absolute operational space but by the body position itself (Rizzolatti *et al.* 1997). The visual fields around the hands follow parts of the body no matter the location of the hands indicating a possible transformation of coordinates

from a body centred to a hand centred system of reference associating what the eyes see and what the body senses as the location of it's parts.

Moreover, in monkeys, some areas in the brain have been found to be strongly responsive to visual stimuli in the peripersonal space and are very likely involved in the coding of it. Several studies have found neurons in area F4 in the premotor area 6 and the putamen, in the intraparietal area and the medial intraparietal area and area 7b of parietal lobe to be strongly responsive to visual stimuli in the space near the body (Colby and Duhamel 1996; Colby *et al.* 1993; Graziano and Gross 1993, 1995; Rizzolatti *et al.* 1988). Some brain areas thought to be involved in the representation of near space are the premotor area 6, the putamen, parts of the parietal cortices (Graziano and Gross 1993).

For a proper interaction with the objects around us the brain constructs a representation using the auditory, visual and tactile modalities (Macaluso and Maravita 2010). A right localisation of a tactile stimulus in the hand for instance, needs information of the posture of the hand. Studies have found neurons that respond to multi-sensory stimuli in cat and monkey brains (Aspell *et al.* 2010; Holmes and Spence 2004; Maravita *et al.* 2003). Studies in monkeys have shown most of this multi-modal neurons respond to visual and tactile stimuli around or close to the hand, no matter the location of the hand in space and the response of these neurons can also be maintained after the stimuli is no longer visible (Graziano *et al.* 1997). One more time, research on the topic seems to indicate that it is a representation of the objects with respect to our own body, especially our hands (Schicke 2007).

Other studies have shown that auditory stimuli also contributes to the

representation of space around the body in monkeys (Farnè and Làdavas 2002) as well as in humans (Serino *et al.* 2009) although the present project will not address that sensory modality.

Most evidence of a peripersonal space representation in humans comes from neglect and extinction patients. Làdavas *et al.* (1997) argue that these two perception disorders occur when there exists a imbalanced competition between two or more spatial representations. However, there is also evidence from functional magnetic resonance imaging on healthy subjects showing that in the human intraparietal sulcus (IPS) and lateral occipital complex (LOC) there are neurons that respond to visual stimuli near the hand (Zohary *et al.* 2007). That study suggests neurons in the LOC and posterior IPS represent space around the hand in a more visual way while anterior IPS does it using multisensory information. For investigating sensory contribution from different modalities the stimuli presented individually by using a rubber hand or changing the position of the unseen hand. This kind of studies might also reveal the reliance on certain usual conditions and the miscalculations product of unusual ones. Bremner *et al.* (2008) propose a framework for studying the development of spatial representations in infants and argues that in infants around 6 months old have a high reliance on the visual information of their arms position when attempting to perform reaching, later, around 9 months the reliance on this information decreases and proprioceptive information gets a more weight. According to this, space representation goes through these two stages in it's early development and this explains why infants around 6 months perform well on spatial recognition tasks but not on visual-spatial

orienting tasks² as the use of proprioceptive information is not yet present. A characteristic of peripersonal space is its plasticity (Làdavas and Serino 2008) experience will extend or contract the space the system knows is able to reach. There are also studies that show that looking at someone else performing certain tasks can lead to modifications in one's peripersonal space (Heed *et al.* 2010). Modifications to the arm, extensions provided by tools, for instance, will produce these modulations to the peripersonal space representation.

Peripersonal space is usually limited to 30-60 cm from the body, approximately the length of the arms (Holmes and Spence 2004). Neurons in parts of the brain thought to be related to peripersonal space representation get activated significantly stronger by visual stimuli inside that volume. However, perceived reachable space, and therefore peripersonal space representation, has also been found to have plastic properties. Modulated by experience, tool use can extend it. This was first found in monkeys by Iriki *et al.* (1996) and it has also been investigated in human brain-damaged (Farnè *et al.* 2007; Làdavas and Serino 2008) and healthy subjects (Maravita 2002). In humans, Lourenco and Longo (2009) also found it can get contracted following an increase in the effort required to perform a task with a tool. Their study suggests the assertion of our near space is conscious. However, another study on peripersonal space in humans by Ambrosini *et al.* (2011) suggest that our reaching ability relies on actual motor potentialities and not in our cognitive estimates. When observing objects, accessibility estimates are constructed at the moment of observation and are not necessarily accessible to conscious representations/estimations, meaning that

²Visual-spatial orienting involves the turning the eyes towards a stimulus, it is specific to vision, and involves no verbal component (Posner and Cohen 1984).

it is the actual reaching ability what modulates behaviour. An interesting finding that promotes discussion on the notion of objects and embodiment. In a higher level cognitive behaviour such as language, peripersonal space studies have been done to investigate words as peripersonal space extension tools (Borghi and Scorolli 2012) based on recent research on the similarities between tool-use and language (Clark 2008). Another study using language (Costantini *et al.* 2011) explored how observation of objects in the peripersonal space resulted in fastest response times (RT) after showing observation or manipulation verbs in comparison with when objects were out of reach. This also suggest that objects are encoded in the brain in terms of affordances related to the physical possibilities of action on them.

Moving a little bit beyond the fields of neurophysiology or psychology and getting more into *computational approaches* to the study of peripersonal space, Magosso *et al.* (2010) has developed a model of visuotactile representation of the peripersonal space around the hands. The model tries to account for phenomena observed in neurophysiological studies with a series of interconnected neural maps that represented uni and bimodal sensory maps from simulated visual and haptic sensors for each hand. Inhibitory and excitatory connections following a Mexican hat function connected the bimodal maps to the unimodal ones with the argument that multimodal regions in the brain can elicit or inhibit response of single modality neurons. The authors also suggest using Hebbian rules in this model for exploring plasticity of peripersonal space. This model was successful in reproducing some of the phenomena that natural systems exhibit, including tactile extinction.

As we mentioned above, Magosso *et al.* (2010) succeed in reproducing phenomena present in peripersonal space studies. The model accounts for the activation of certain neurons in the brain in a biological plausible way. However, these kind of models overlooks the complexities of perception. Due to the unembodied nature of the models themselves, the inclusion of physical interaction is simplified. On the other hand, for roboticists it is still very difficult to cope with all the uncertainties that real sensors show in a real physical environment as well as with the control of the robot itself. For the robot Asimo, Goerick *et al.* (2005) developed a perception framework based on findings on peripersonal space, specifically that work is aimed at providing a robot with an attention shifting and object recognition capabilities. One main component in the system is the computer vision sub-system, composed by various saliency maps including stereo disparity for depth information acquisition. The system is able to track and recognise objects that enter it's peripersonal space, which is defined beforehand. Looking at these two cases, one modelling the neural basis of peripersonal space and another implementing peripersonal space in a technical system we can see a big contrast. That is because phenomena related to peripersonal space seem to be present at many levels of the working pieces of natural systems behaviour.

Chinellato *et al.* (2011) also explored the possibility of creating a representation of the peripersonal space based on visuomotor associations using a radial basis function framework. The representation is implicit and emerges through the interaction with the agent environment. This study therefore advocates for non-explicit representations and ones that the system itself constructs and thus has meaning and are useful for it and not necessarily

to the designer. Contrasting with Goerick *et al.* (2005), here the visual part is not as elaborated due to it is the interplay with the rest of the exploration elements what contributes to the resulting representation.

Peripersonal space has come into attention in recent years because its relation to the start of development and the close connection it has to embodiment (Lewis *et al.* 2012). Additionally, we have seen in this review that scientific findings demonstrate the necessity of action for developing the perception of reachable space (Coello *et al.* 2008), and in accordance, studies on robotics suggest gazing and reaching can contribute to create a representation of it on artificial systems , making these topics a centre of interest of roboticists. Under the new, embodied, approach on intelligence, the body of the system is very relevant to it's development and if we want to generate systems with a similar intelligence to what we humans have, the series of physical interactions an artificial system experience should be similar to those a human body experience. For many decades mobile robots have been a design platform and test bed for AI systems. Today, by contrast, the tendency is to explore AI in more human-like systems. In the beginning of our lives most of our interaction occur in the peripersonal space and therefore the study of it is relevant to current research on robotics. The review we have done on peripersonal space is important because we start our study from these psychological an neurophysiological findings. Many of them involve the visual modality, therefore we will now review some aspects of vision important for our research and also the applications and approaches taken on it in robotics.

Because vision is a perceptual modality has been widely employed in peripersonal space studies, it is important to make a quick review of some

aspects related to it. Computer vision is the most flexible sensing method available currently (Mata *et al.* 2005) and, as we will see, vision has been approached in different ways for integrating it to robots.

2.3.3 Existing work on robotics reaching and peripersonal space

Following, a review of the work related to the research topic of this dissertation is presented. Work on reaching learning, as well as on both peripersonal space and body-schema is exposed. These two latter topics are closely related (Farnè *et al.* 2009) and therefore considered both to build up in the development of reaching and space representation.

Robotics offers the possibility to study reaching, peripersonal space and body schema because it allows putting to the test hypothesis about how action is involved in the creation and modification of space representation. In order to study peripersonal space representations, the a first step has been to study action creation and development. Endowing a robot with reaching behaviours has been researched by many authors in the last twenty years (Fagg 1994; Guenter *et al.* 2007; Metta *et al.* 1999; Trullier *et al.* 1997; Vahrenkamp *et al.* 2008) and continues to be a challenging topic despite of the advances in technology. This review focuses on research related to developmental approaches in the study of the topic.

Developmental approaches include Schlesinger *et al.* (2000), who explored progressive expansion of the search space in order to allow a simulated two-dimensional robot arm learn reaching movements. The approach used evolutionary robotics with a neural network controller that encoded what they call styles or strategies for reaching. It also took a developmental approach based on initially limiting or constraining the

conditions in which a robot learns a skill. In another evolutionary study (Massera *et al.* 2007) simulated an anthropomorphic robotic arm that used and evolved neural network controller to reach and grasp target objects with different shapes. The simulated manipulator had seven degrees of freedom in the arm plus twenty in the hand for a total of 27 degrees of freedom. The authors argue that effective reaching can be developed through a trial an error process if there exists a fine grained interaction between the robot and its environment, emphasising on the physical interaction. A recent study on development of reaching using a simulated humanoid robotic iCub (Savastano and Nolfi 2013) also used an evolutionary approach that included accurate modelling of infant reflexes and maturational processes. The work reported achieving behaviours that closely match those by children.

With a very clear emphasis on development, Lee *et al.* (2012a) used a robot manipulator with visual capabilities and followed stages from a developmental timeline (Law *et al.* 2011) as guidelines for a robot shaping methodology that endowed their robot with reaching and grasping skills. The robot was provided with a top-mounted camera that extracted a two-dimensional view of the workspace. A saccade learning stage was put in place before motor babbling and the construction of a proprioceptive mapping was achieved. The approach took into account the proximodistal flow of motor development by means of enabling distal joints only after control over proximal ones has been attained. Their work provides insights of how a robot can follow a developmental pathway similar to that in humans and shows that it can be applied in robots as well. In their work motor babbling was a key element in learning and they argue that it has close links to behaviour because

of the vital sensorimotor data and rehearsing it entails that is useful in later stages of action and experience. They also consider the proprioceptive data acquired during initial learning to traditionally been under-rated and that it is the main feedback on limb positioning. Their work stresses the importance of refinement of certain abilities before other in order to achieve a complex skill.

As we have seen in section 2.1.1, today robotics is at a stage where building humanoid robots for scientific research is possible. Although peripersonal space in neuroscience and psychology has been studied in primates and humans, peripersonal space and body schema in robotics are not always explored using humanoids. Roschin *et al.* (2011), for instance, investigated a method for body-schema development in a seven degrees-of-freedom simulated 3D agent. The method employed a self-organising neural network connected to several haptic receptive fields on the body of the robot. The system successfully created representations that integrated both the position of the end-effector in operational space and in joint space.

Sturm *et al.* (2009) presented an approach for allowing a robotic manipulator to learn its kinematic model based on observation. The approach employs Bayesian networks that learn the structure of the robot by analysing geometrical relations. The experiments were carried out both in simulation and in the physical manipulator. It is worth noticing the use of vision in their approach. By providing the robot a monocular camera, it was able to learn based on exploration of its own characteristics. Direct and inverse kinematic functions were learnt by the robot. The Bayesian network learnt in an online mode and they propose that their methodology is useful for life-long

adaptation.

Work on peripersonal space with humanoid robotics includes the one from (Hersch *et al.* 2008). Their work also is interested in providing the robot the means to learn kinematic functions that can be useful as a body schema. Their work is interesting because they allow the sensorimotor contingencies shape a representation in a simulated robot with 24 degrees of freedom with similar morphology to that of the Fujitsu Hoap3 robot. The simulated robot was endowed with tactile and visual sensors so that it could correlate the sensorimotor contingencies based on different modalities and create a coherent image of itself. The learning algorithm used in this approach was online and self-supervised, although in the simulations the algorithm not always converged to a good representation of the robot's structure. In this work, the approach was also tested on a real Hoap3 robot. They presented results of a robot that was initialised with its "real" body schema but with an extension attached to one of its hands. The robot was able to adapt the initial body schema to include the elongated limb. On another peripersonal space study on humanoids, Antonelli *et al.* (2012) explored an adaptive representation of reachable space in a study using the NAO humanoid robot. They proposed a model in which the robot has an implicit representation of the target position encoded in the positions of the arms. The robot had neural network controllers for calculate direct and inverse transformations of head and arms positions by means of radial basis function networks. The authors argue that a representation of peripersonal space is encoded by the radial basis neural network by integrating the redundant cues regarding the target position, to which the robot has access through spherical coordinates coming

from the head, and also by the position of the arm. In a similar approach, Chinellato *et al.* (2011) explored the emergence of implicit sensorimotor mappings on a simulated two-dimensional robot manipulator provided with a visual vergence mechanism that endowed the robot with visual-like capabilities for acquiring proprioceptive information once a target was foveated. They argue that the representation of peripersonal space the robot generated was never made explicit but instead emerged from gazing and reaching actions.

As we have seen in this review of work on robotic peripersonal and related topics, most work in this line of research follows or is inspired by developmental robotics section 2.1.5. This approach is also explored in the work presented in this dissertation. Missing in the study of peripersonal space in robotics is more work in more realistic environments such as three-dimensional environments and humanoid platforms. Although humanoid robots have been used previously, this is only a recent approach. The use of integrated humanoid robotics will foster the shifting to working in three-dimensional environments. Another aspect of peripersonal space still to be researched extensively is its plastic nature during body modification and the effect it has in the confidence levels the agent has when attempting a reach. This issues are explored in the present dissertation in a simulated humanoid robot.

2.4 Vision

2.4.1 The classical approach to computer vision

Computer vision is the transformation of data taken from a still image or series of images or video into a decision or a representation to be used in solving some task by a machine. For many years the approach taken on computer vision was that an artificial system would process an image in order to perceive the world and from this processing alone be able to generate useful and/or meaningful information about objects and the environment. This was an intuition inspired on studies of the mammal retina that followed a reductionist approach. These studies investigated the structures and connections of the mammal ocular system and produced accurate models that could be implemented on machines therefore computational results were treated on par with neurobiological findings (Marr 1974). This was the computational perspective on vision on which Marr was a big proponent. Marr and Poggio (1979) proposed three levels for a visual processing system: the top is an abstract computational part involving the conversion from one type of information to another, a middle representational and algorithmic part defining what the input and output are and for generating the output, and a hardware implementation.

Many techniques were developed in the 80's and 90's for scene analysis and shape extraction from images. Shape from shading (Horn 1989) and other shape-from-X techniques as well as stereo matching which have been widely used for depth perception. Most of them adhered to Marr's approach of visual perception of capturing aspects of reality and assigning symbols to them for building a representation of the object from these symbols (Marr and

Poggio 1979). It's important to note that this perspective on visual perception does not take into account an embodiment of the system, that is the possible dynamical interactions with the environment the agent could have. A more recent approach to computer vision is active vision.

2.4.2 Active vision

In the field of mobile robots, AI has been extensively studied and applied alongside computer vision. Part of the problem mobile robots have to solve is the extraction of information coming from sensors, usually vision being sensing modality used in many of them. Computer vision following the traditional approach proposed by Marr and Poggio (1979) has been used with successful results in map-based, navigation algorithms and representations, like x, y, z coordinates or used by the robot have been generally imposed by the designers and artificial intelligence has been used for navigation and mapping. One example of these methodologies are map-based models (Filliat 2003 for a review). Very often geometric calculations were used to get the representation in space of landmarks and generate the internal representation of the space. An approach for the vision problem more in line with the topics reviewed for this work is *active vision*.

Active vision is a bottom-up approach to visual perception consisting in allowing the visual input an agent receives directly influence it's movement actions which in turn can provide complementary information of the observation target. This approach recognises that vision involves more than acquiring and processing images. This methodology contrasts with the classical sense-plan-act of the hierarchical paradigm that dominated robot implementations for some decades around the 80's, starting with the robot,

Shakey (Nilsson 1984). Active vision, and in general active perception was a precursor to more recent embodied intelligence theories, like the sensorimotor account of vision and visual consciousness (O'Regan and Noë 2001).

Initially proposed by Bajcsy (1988), Aloimonos (1991) and (Ballard 1991), active vision has studied the control strategies for perception. Perception is a process that requires interaction with the environment. Therefore, active vision and in general active perception, suggests that perception has to be done by exploring, probing and searching. Active sensing—not necessarily related to the use of active sensors—deals with controlling strategies applied to data acquisition differently depending on the state of the data interpretation and the goal or the task of the process. Active vision mechanisms tries to extend the visual capabilities of an agent allowing a continuous feed of information linked to the movement generated that is only possible if the agent has a body embedded in an information-rich environment. Bajcsy (1988) and Aloimonos (1991) first suggested the idea, while (Ballard 1991) and others extended the concept and showed that it can be used for controlling multiple behaviours at the same time in a similar way a natural system like humans has to control neck and eye movements for centring an object of interest in the field of view. Further work on active vision combined with evolutionary robotics has provided systems that can carry out complex shape discrimination (Floreano *et al.* 2004) from very simple visual capabilities and locomotion control.

Active vision differs from control theory in that the feedback the system gets is performed not only the data received but also on processed sensory data. Also, the feedback the system produces depends on the models built in

the system that combine symbolic and numeric information of the task to solve. It also contrasts with Marr and Poggio (1979)'s proposed approach to computer perception where a static system should be able to recognise objects, in contrast with active vision a physical interaction when performing a visual task is essential for carrying out the perceptual task. Around 1970 block worlds were very popular in the study of AI, they were used even in vision research because they provide constraints for the perceptual processing and greatly simplified perception problems. With help in the abstraction creation, search was an easy problem. This kind of environments were good for studying planning problems but implementing the same ideas in systems that interacted with the real world was still a great challenge. The active paradigm has been a way to cope with the problems of robot perception that complexities real physical systems pose. In a physical system, recognition, spatial understanding, sensor noise, etc. are no longer ignored or delegated to black boxes.

Aloimonos (1991) argues that in the classical approach vision was regarded as a recovery problem, the task was to reconstruct a three-dimensional scene from image cues (shading, contours, motion, stereo colour, etc.). This approach has provided many mathematical techniques for image analysis such as the ones mentioned above. However, in natural systems images might not be completely analysed, vision serves for visual tasks, that is vision involves action (Milner and Goodale 1998), an idea which relates to Gibson's (1979) ecological approach to visual perception. In the ecological approach, cognitivism and information processing are criticised in favour of direct perception generated by light reaching the *ambient optic array*

in the eye, which provides unambiguous information about the arrangement of objects in the environment. Gibson stressed the relation of perception and action and his ideas gave rise to ecological psychology which holds significant appeal to robot designers.

Animals use vision for detecting predators or mating partners, for instance. If we aim to having robots interacting with the environment and other agents, it might be that complete reconstruction is not a necessary condition. Aloimonos also makes a remark on this by posing the question: “what do you need vision for?”. If a natural system needs to know only if something is getting closer to or away from it, scene reconstruction could take more computational power, that is energy, and time where reaction time could be critical for survival. Aloimonos’ purposive paradigm on vision argues in favour of simple, robust algorithms based on qualitative techniques that could be simple comparisons of quantities and discrete classifications.

In the real world, perception (or abstraction) and planning (or reasoning) are not clearly divided (Brooks 1991a). There is psychological evidence that many of the processes involved in the representation of the world used by an intelligent system are intrinsically connected: spatial understanding, recognition, coping with sensor noise (Klatzky and Lederman 2010; Lacey *et al.* 2007; Shams *et al.* 2000). Bajcsy (1988) argues that it should be axiomatic that perception is an active process and that modelling perception problems is more relevant than modelling biological systems when investigating vision. Models of control strategies that include sensors, objects, the environment, and the interaction between all these elements should be given a purpose such as manipulation, mobility, and recognition. She finishes

by defining active perception as “a problem of an intelligent data acquisition process... [where] one needs to define and measure parameters and errors from the scene which in turn can be fed back to control the data acquisition process”. Again we see the methods that active vision proposed regarding the use of simple algorithms. Something that also Ballard (1991) mentions is that visual computation is less expensive in active systems than in passive ones. This is possible because of the use of the same simple algorithms which also can be based on simple sensors such as low resolution cameras because the movements of a the camera provides the system with “virtual high resolution cameras”.

A part of the methodology of active vision is also present in the animat approach and this is the recognition of the existence of several systems needed for achieving visual tasks. Moreover, that these systems are closely intertwined and collaborate for the task. According to Ramachandran (1993), psychophysical evidence shows that the visual system can be considered as many different algorithms exploiting several cues but not all algorithms always work and may not be simultaneously satisfiable. Figure 2.1 shows a way of decomposing an intelligent behaviour into a collection of simpler behaviours. On previous approaches the decomposition was made from the functional steps to take for solving a task (a) where all were considered to be sequential, while in the animat approach the decomposition is made on the basis of the behaviours needed at the same time in order to achieve a task (b). For example, Ballard (1991) mentions one experiment with a robot with 5 independent visually guided behaviours that kept a balloon in the air. The processes had no communication between each other apart from the effects

on the robot and the environment. He argues how these kind of strategy can originate a coherent behaviour.

Computer vision has been used extensively in mobile robots. However, the kind of intelligence this kind of systems can exhibit is far different from the intelligence that, for instance, an object manipulation task would require. Brooks (1997) suggests that if a robot shall have human-like intelligence it must have a human-like body. That way the experience that body has from the interactions with the world will build similar sort of representations a human does. The active approach to perception is good candidate for providing artificial systems with means of interacting and perceiving the world. This project differs from much of the active vision work previously done because we will use active vision for sensing the environment and adapting behaviours in a system that has a humanoid embodiment. Also, we are interested in applying active vision for interaction in the near space around a the body of the system, whereas active vision in mobile robots was used for navigation and mostly involved perceiving the extrapersonal space of the robot.

Although we can only build crude approximations to the human body that may miss the essential characteristics of it the study of this systems might reveal some of the aspects that are important for human-like intelligence. This ideas share a lot in common with the *animats* approach (Wilson 1991) that has been used for studying animal behaviour for some years. An animat is a real or simulated robot embedded that continuously interacts with it's environment through it's sensors and actuators. Some researchers think at this point AI is incapable of understanding and reproducing human intelligence. However, we think, along with Brooks (1991a), Guillot (2001), Hoffmann and Pfeifer (2012)

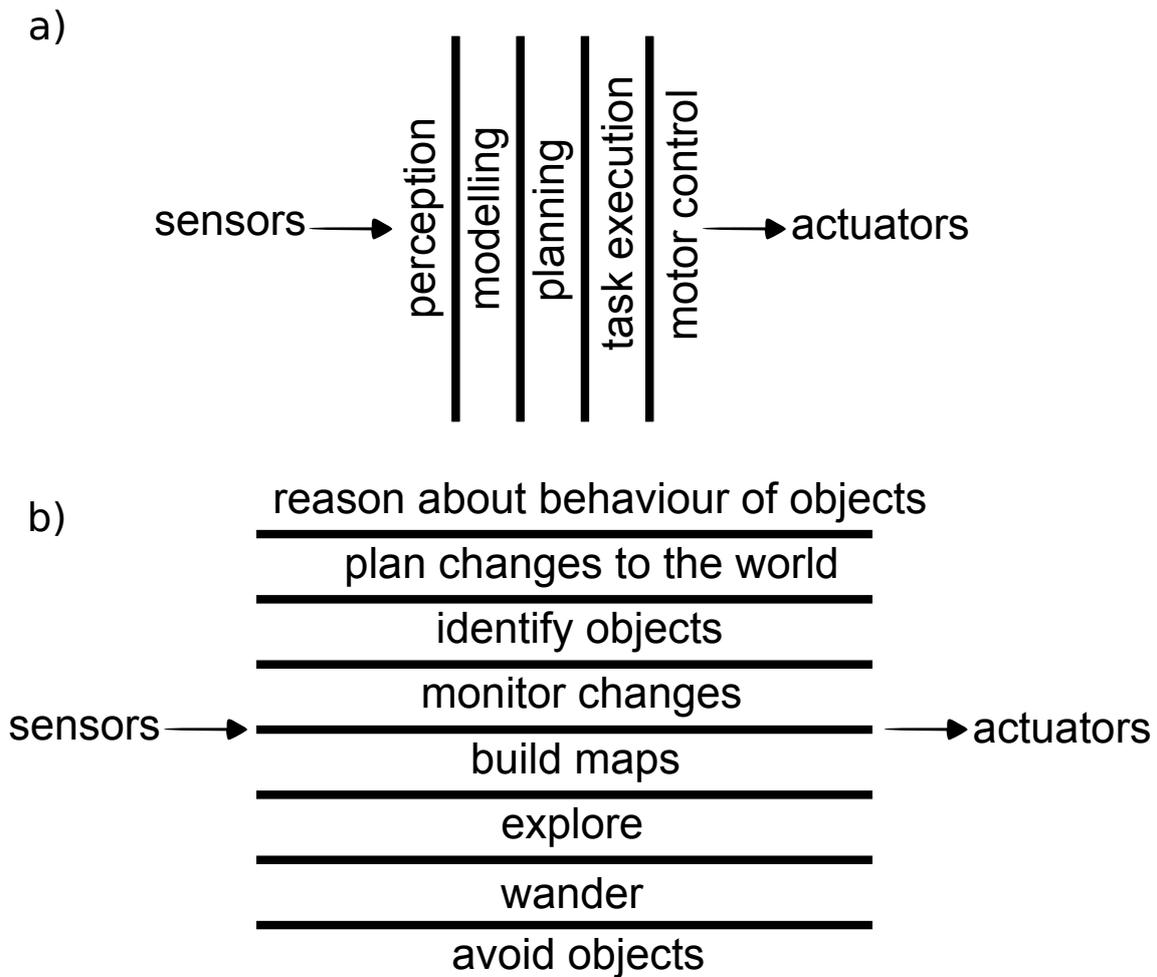


Figure 2.1: Decompositions for a mobile robot control system. a) Functional decomposition b) Behaviour based decomposition (reproduced from Brooks 1986b).

and other researchers, that the animat approach can help in understanding human cognition with a bottom-up strategy, and that it is necessary to produce robust layers of control for low-level intelligence tasks such as vision, mobility and manipulation in order to achieve high-level intelligence as we observe in ourselves. This, independently of if there can be such a distinction between sensorimotor and cognitive, or high level, intelligence in human behaviour.

In traditional computer vision great attention was put into the reconstruction of three-dimensional scenes. The approach was to extract as much information as possible from an image or stereo image and then create

a depth map. However, it has been argued that humans do not do this (Ballard 1989). Instead exploration of the scene is actively done and thus computation is done locally in the regions of interest and not on the whole image. For depth perception, the human visual system has the vergence mechanism, used for decreasing disparities from stereo images. This is an eye movement that can be used for active exploration of a scene and offers a way to perceive and get depth information, therefore, next we are going to present a short introduction to vergence for depth estimation.

2.4.3 Visual vergence for depth estimation

In humans, the development of frontal vision allowed to foveate objects of interest with both eyes. The disjunctive movement of the eyes to do so is called vergence (Leigh and Zee 1999). Hering (1977) noted that the two eye's move either the same amplitude in the same direction (version) or the same amplitude in the opposite direction (vergence). Hering's "Law of Equal Innervation" says that all eye movements should be generated by a linear combination of largely independent versional (conjugate) commands and vergence (disconjugate) commands. According to this law, conjugate commands operate on an imaginary eye located between the two actual eyes.

Many vergence studies adhere to Hering's theory where versional movements occur simultaneously and symmetrically with same amplitudes on both eyes as if there was a line between them and perpendicular to a line described by the location of them. Other theories suggest that the brain has independent control over each of the eyes (Enright 1998; Zhou and King 1998). Enright (1998) suggests that when binocular fixation is shifted an unbalanced saccade on one eye occurs and is later followed by a vergence

movement by the other eye to get full binocular foveation and that this mechanism allows quick high-resolution monocular view.

The debate on this issue has been on for more than a century. However, for our experiments with iCub, Hering's theory is followed as it has been set as a constraint for the vision system to have symmetrical versional movements by design. Vergence was used for depth estimation as it is considered a bottom-up mechanism—image based—that is a natural representation of depth in biological systems.

It is also discussed whereas vergence is affected by higher-level cues or vice-versa. It has been shown that altering vergence affects depth perception coming from higher-level , horizontal disparities (Cumming *et al.* 1991). We adhere to this view and make the use of vergence as the main cue for depth. Also, compatible with this view is the study of Tresilian *et al.* (1999) where it was found that the vergence signal for depth estimation is given more importance as the target is closer to the observer and also when information from pictorial cues decreases.

2.4.4 Existing work on robot mono and stereo vision

Robotic vision systems can be implemented in two ways. The camera can work in a stand-alone configuration. In that case the camera serves as a global sensor and is not mounted on the robot. The other configuration is when the camera is mounted on the robot as in a in-hand sensor. For both configurations, a typical application has been position-based servoing, which means estimating and tracking three-dimensional orientation and position of a target object based on the camera images. Extracting features of the images is a common approach in robot vision in both of the mentioned implementations.

Feature extractors like the one developed by Harris and Stephens (1988) and Canny (1986) are usually implemented and employed in many robot vision systems due to the importance in many tasks of consistent image edge filtering, line and contours detection.

Implementations of monocular vision using the stand-alone configuration were prominent in early studies of robot vision (Feddema *et al.* 1992; Shirai and Inoue 1973; Tonko *et al.* 1997; Yoshimi and Allen 1994). The typical task was to estimate the target's pose based on the static images obtained by the camera sometimes in order to prepare a grasp action on the object. Shirai and Inoue (1973) used the sense-plan-act paradigm was used for precise positioning of a block into a box. Buttazzo *et al.* (1994) created a real time system capable of catching a fast moving object in a two-dimensional plane using colour segmentation.

Most common these days is the in-hand configuration, where the camera is rigidly mounted on an end-effector. In the in-hand configuration, the transformation from camera to end-effector coordinate frames is typically known (in section 2.3.3 recent body-schema and peripersonal space studies were presented that do not require *a priori* knowledge of the transformation). Optical flow (analysis of changes the visual scene caused by relative motion between the observer and the scene) has been used by Papanikolopoulos and Smith (1995) and (Brandt *et al.* 1994) among others. For achieving the same task Colombo *et al.* (1995) used an active contours approach to estimate the parameters of the required end-effector transformation then used for calculating a motion parallel matrix used for setting the position of the four degrees-of-freedom of the robot. A neural network approach is presented by

(Wunsch *et al.* 1997) where a Sobel filter was used for enhancing the target's state information from the image features and a Kohonen self-organising neural network was trained with computer generated object views. The network was able to estimate the position of the end-effector for both simulated and real images with the required accuracy in 81% of all cases. The network topology was chosen according to the representation of 3D orientation.

Binocular vision consists in mounting a two camera array used for visual tasks. This approach uses more computational resources but it has the advantage of providing the visual system with depth information. This approach has been implemented in many robots in the last decades, being disparity the most common way for estimating depth (Barnard and Fischler 1982; Marr and Poggio 1976). An in-hand configuration approach that used disparity and vergence is Olson and Potter (1989) where a control mechanism was designed. The mechanism controlled the cameras' vergence angle (see section 2.4.3) in the Rochester robot, an industrial robot arm with a custom built head, (Brown 1988). They presented an algorithm based on a discrete control loop for the motors and algorithms for calculating image disparity based on the cepstral filter. Arlotti and Granieri (1991) created a robot with a mounted camera that create a three-dimensional wire-frame reconstruction of an object by moving to different positions and triangulating the location of the target's features. Li *et al.* (1994) studied hand-eye calibration

The camera mounted, in-hand, configuration usually does not achieve the high accuracy that stand-alone configurations does. However, it is widely used in humanoid robots studies due to its biological plausibility. Also,

different theories from those of control are explored in this type of systems. The robot “Richard the First” (Mowforth *et al.* 1990) looked to recreate a system with the same degrees of freedom and reflex times of a human head. The interest of the researchers was achieving fast execution times for the movements of the head. Natale *et al.* (2002) built a robotic model of visuo-acoustic integration in a binocular head. The visual part acquired and processed the images in a space-invariant format known as log-polar (Sandini and Tagliasco 1980) with biologically plausible higher resolution foveas and lower resolution periphery of the visual receptive fields. The binocular Yorick stereo head (Eklundh and Björkman 2005) was able to use several visual cues including disparity, motion, local texture and colour for identifying objects using low resolution, wide field cameras for scene analysis and high resolution foveal cameras for individual object recognition. Their system combining three-dimensional and monocular cues to execute the task. The Medusa stereo head (Santos-Victor *et al.* 1994) was built for studying active-vision. Bernardino and Santos-Victor (1999) used log-polar images and implemented fast low-level vision algorithms in the Medusa and achieved real-time performance and high accuracy in tracking tasks.

The stand-alone configuration, with the binocular system is not mounted on the robot, provides more accurate depth estimations. For example, in one of the earliest stereo servoing systems a ping-pong playing system was built using colour segmenting and a dynamic model for calculating the ball trajectory (Andersson 1989). Other studies also used dynamical models of a target for catching or placing it in a desired location (Burridge *et al.* 1995) and a juggling task (Rizzi and Koditschek 1994) analysing the motion of colour segmented

blobs. Stereo triangulation of optic flow in real-time was used by Allen *et al.* (1993) for making a robot with two top-mounted cameras reach and grasp a moving target on a planar surface. Grosso *et al.* (1996) used optic flow for a three-dimensional reaching task in a five degrees of freedom robot.

In this review we have seen the varied implementations of robot vision. Vision holds an important place in robotics research because it is one of our most developed senses. From the HRI perspective, understanding it will allow us to build robots suitable for operating in unstructured human environments. From a scientific perspective, understanding how visual activities are carried out by biological systems can provide insights to other aspects of intelligence.

2.5 Machine learning

Humans and other biological systems display behaviours that we label as intelligent. Learning allows modifying these behaviours and also to learn new ones. Being capable of learning makes an agent flexible to changes in the environment. Since computers were invented, designers have always tried to find ways for make them learn from experience, so that they can improve themselves. Detailed understanding of how to make computers learn autonomously could open up new uses for computers and improve their competences in many fields of human activity.

Adapting and generalising are capabilities that a learning agent would display and therefore ones that machine learning researchers are interested in providing computers with. Logical deduction and reasoning are aspects of intelligence that were the focus of attention of early artificial intelligence, which studied symbolic methods for creating intelligent systems. By contrast,

the methods used in machine learning are sometimes called sub-symbolic.

Machine learning is a set of computational methods that use experience to improve performance in a task. The types of tasks solved by machine learning are varied, although traditionally most of them have been predictions, decision making and categorisation. Examples of learning tasks are optical character recognition (OCR), text and document classification (used for detecting email spam, for instance), speech recognition, speech synthesis, speaker identification, image and facial recognition, games, fraud detection, medical diagnosis, internet search engines. For carrying out all these tasks machine learning, as an inter-disciplinary field, draws concepts from biology, neuroscience, computer science, mathematics, physics and mathematics.

Depending on the task in machine learning problems are addressed using different algorithms. A broad classification of machine learning algorithms considers the form in which experience information (data) is provided. Following we will review different machine learning types, explain their characteristics and present algorithms they include.

2.5.1 Supervised learning

In this category a training set consisting of examples with the correct responses (targets) is used. This is called a labelled data set. Based on this training set, the algorithm generalises to respond correctly to all possible inputs. If we had examples of every possible input-output data a lookup table could be created and machine learning would not be required at all but it is generalisation what makes machine learning powerful. By generalising, the machine learning algorithm can produce sensible outputs for inputs it had

never exposed to. That way, the algorithm is also robust to noise so that small inaccuracies in the input do not affect its performance.

Let the domain of instances be X , the domain of labels be Y and let $P(x, y)$ be a unknown joint probability distribution on instances and labels $X \times Y$. With a given sample $\{(x_i, y_i)\}_{i=1}^n$ drawn in a independently and identically distributed (i.i.d.) way from $P(x, y)$, supervised learning trains a function $f : X \rightarrow Y$ in some function family F so that $f(x)$ predicts the true label or output value y on unseen yet data x , where the points (x, y) are drawn in an independently and identically distributed way from $P(x, y)$.

Common supervised learning tasks:

- *Classification* assigns a category to an object. In a typical classifications problem we have a dataset with input vectors and we have to decide which of N classes they belong to based on training from exemplars of each class. This is a discrete problem as each vector (object) belongs to just one class and the set of classes covers the whole output space. In this case the function f is called a classifier and the labels are Y .
- *Regression* predicts a real value for an object. In this case Y is continuous and the f is called a regression function.

Supervised learning algorithms include:

Logic based algorithms like *decision trees*. Classifying trees that based on feature values sort instances. Each node in the tree represent a feature of an instance to be classified and each branch a value that the node can assume. The classification of instances start from the root node and its based on their feature values.

Perceptron-based techniques. *Perceptrons* are linear classifiers that

combine a set of weights with an input feature vector. They can be described as follows: if $\{x_1, \dots, x_n\}$ are input feature values and $\{w_1, \dots, w_n\}$ are real valued connection weights (typically in the range $[-1, 1]$). The perceptron computes the sum of weighted inputs $\sum_i x_i w_i$ and the output goes into an adjustable threshold function. If the calculated sum is above a threshold, output is 1, otherwise it is 0. The perceptron weights are usually adjusted (trained) using a batch of training instances that are fed repeatedly until the connection weights are correct for all the training set. A test set is used for verifying the perceptron predicts correct labels for instances that were not in the training set.

Perceptrons are only capable of classifying linearly separable sets of instances. *Artificial neural networks* (ANN) solve this problem by combining perceptrons in a multi-layered array (a multilayer perceptron). Multi-layer neural networks is composed by a large number of units (neurons) joined together by connection weights and typically these units are divided into three classes: input units that receive the values vector to be processed, output units that give the result of the processing and hidden units that are in the layers between those of the input and output units.

Radial basis functions (RBF) networks are three-layered feedback network in which hidden units implement a radial activation function and output units a weighted sum of the hidden units' outputs. The training of these networks consists of two parts: first a clustering algorithm is used for deciding the centres and widths of the hidden units, and then the weights connecting the hidden units to the output ones are adjusted using least mean squared or singular value decomposition algorithms.

Other supervised learning algorithms include **statistical learning algorithms** (*naive Bayes classifiers* and *Bayesian networks*) and **support vector machines**.

2.5.2 Unsupervised learning

These type of learning do not use training sets that include the correct answers to a problem. The algorithm instead tries to identify similarities between the inputs in order to categorise together those inputs that share something in common. Density estimation is a statistical approach to unsupervised learning.

Unsupervised learning algorithms use a training sample with n elements $\{x_i\}_{i=1}^n$ and no teacher providing supervision regarding how the elements should be handled.

Common unsupervised learning tasks:

- *Clustering* partitions large data sets into homogeneous groups. Clustering partitions $\{x_i\}_{i=1}^n$ into k clusters so that elements in the same cluster are similar and elements in different ones are dissimilar. The number of clusters k can be inferred from the training data or can be assigned arbitrarily.
- *Dimensionality reduction* finds a lower-dimensional manifold that preserves some properties of the original data.
- *Density estimation* learns a probability distribution according to the sampled data.

Some unsupervised learning algorithms are:

The **k-Means** clustering algorithm generates an arbitrary number k of disjoint, non-hierarchical clusters. The clustering is done by minimising the

sum of squares of the distances between the training data and the corresponding cluster centroid. In that way the data is classified. The initial positions of the cluster centres μ_j is chosen randomly. Then for each instance x_i the distance to each centre is computed and the data point is assigned to the nearest cluster centre. Following, each cluster centre is moved to the position of the centre of the mean of the points assigned to that cluster. This two last steps are repeated until the cluster centres stop moving. Once the training is finished, for new instances of x the distance to each cluster is calculated and the instance is assigned to its nearest cluster centre.

Vector quantisation is based on the competitive learning paradigm (where the nodes compete to respond to a subset of the input data) and maps k -dimensional vectors in the vector space R^k into a finite set of vectors $\{y_1, \dots, y_n\}$. Each y_i is a code vector or codeword. The set of all the codewords is the codebook. Each codeword y_i is associated to a nearest neighbour region called Voronoi region defined by $V_i = \{x \in R^k : \|x - y_i\| \leq \|x - y_j\|\}$ for all $j \neq i$. This algorithm allows modelling probability density functions by the distribution of prototype vectors and was originally used for data compression.

Self-organising map (SOM) is the most used competitive learning algorithm. It was proposed by Teuvo Kohonen and therefore it is also called Kohonen map (Kohonen 1982). Self-organising maps keep a topological organisation inspired in the cerebral cortex where some neurons are close to some and far from others. The idea behind self-organising maps was to find how sensory signals get mapped into the cerebral cortex with certain order, so that similar sensory input excites neurons that are close to each other, while

two neurons that are excited by very different sensory input are located very far apart. Organisation of data in the Kohonen map training uses a cooperative process where because of topological neighbourhood neurons close to excited neurons are also excited, a competitive process where the neuron whose weight vector comes closest to the input vector is declared a winner and gets activated.

2.5.3 Semi-supervised learning

Semi-supervised learning is halfway between supervised and unsupervised learning. In this type of learning, the training set is unlabelled, however additional supervised information is given to the algorithm but not necessarily to all the examples. This type of learning is used when unlabelled data is easy to get and some labelled data is available. Sometimes labelled data is difficult or expensive to get.

Tasks of this type of learning:

- *Semi-supervised classification* is a classification that uses labelled and unlabelled data. This classification is an extension of the supervised learning problem. The idea is to get a better classifier f by training from both labelled and unlabelled data. Training data includes both l labelled elements $\{(x_i, y_i)\}_{i=1}^l$ and u unlabelled ones $\{x_j\}_{j=l+1}^{l+u}$ and typically there is much more unlabelled data than labelled therefore $u \gg l$.
- *Constrained clustering* is an extension of unsupervised clustering. In this case the training data consists of unlabelled data $\{x_i\}_{i=1}^n$ and some kind of supervised information regarding the clusters. “Must-link” constraints is one of such kinds of information that, for instance,

indicates that elements x_i and x_j must be in the same cluster.

“Cannot-link” constraints indicates x_i and x_j cannot be in the same cluster.

For semi-supervised learning to work it is necessary to make certain assumptions. Supervised learning also relies on assumptions. One of the most common assumptions in supervised learning is the *smoothness assumption for supervised learning*: if two points x_1 and x_2 are close, then so should be the corresponding outputs y_1 and y_2 . Without this assumption it would not be possible to generalise from a finite training data set to a possibly infinite set of unseen points. Semi-supervised learning uses a generalisation of the smoothness assumption for supervised learning that takes into account the density of the inputs. It assumes that the label function is smoother in high-density regions than in low-density ones. The *semi-supervised smoothness assumption* is that if two points x_1 and x_2 in a high-density region are close, then so should be the corresponding outputs y_1 and y_2 . By transitivity this assumption implies that if two points are connected by a high-density path then their corresponding outputs will very likely be together. This assumption applies for both regression and classification tasks.

Another assumption used for semi-supervised algorithms is the *cluster assumption*. If we knew that the points of each class tended to form a cluster then the unlabelled data could help in finding the boundaries of each cluster with more accuracy. By using a clustering algorithm we could use the labelled points to assign a class to each cluster. The cluster assumption is that if points are in the same cluster they very possibly belong to the same class. Put in other terms, this assumption defines the so called *low density separation*, which

indicates the decision boundary should lie in a low-density region.

Another assumption that forms the basis of semi-supervised learning algorithms is the *manifold assumption*. This assumption indicates that the high-dimensional data lie approximately on a low-dimensional manifold. This assumption tries to tackle the curse of dimensionality by approximating the input space to a lower-dimensional space and trying to learn the manifold using both the labelled and unlabelled data. After the manifold has been learnt, learning can continue using distances and densities defined in that manifold.

Semi-supervised learning methods include *generative models*, *low-density separation*, *graph-based methods* and *heuristic approaches* which are not detailed here because they are beyond the scope of this review.

2.5.4 Reinforcement learning

Reinforcement learning fills the space between supervised learning, where the correct answers are given to the algorithm, and unsupervised learning, where the algorithm has to discover common features in the training data. In the middle, reinforcement learning algorithms receive information about whether the output is correct or not, but no information is given regarding how to improve it. The learning algorithm has to explore different strategies in order to find the path through a *state space* that leads to the correct answer to a problem. The reinforcement learner has to get some feedback from the environment in order to find the correct strategy for achieving the goal. This feedback is provided by a *reward function*.

Applications of reinforcement learning are varied and it has been successfully applied to many problems and has been of interest to

psychologists and computer scientists because the similarities it has with biological learning. Reinforcement learning has had an important place in robotics because it has allowed robots to learn tasks like moving objects to clear a room or navigation. The Sarsa algorithm for learning Markov decision making process policies and Q-learning are two of the traditional and most common implementations of reinforcement learning. For this review only the Q-learning algorithm is described.

Q-learning is a specific kind of reinforcement learning. At each step s the learning agent chooses an action a which maximises the function $Q(s, a)$. Q is the *estimated utility function* and it indicates how good an action is for a given state. Therefore $Q(s, a)$ is equal to the immediate reward for an action plus the best utility (Q) for the resulting state.

The formal definition of Q-learning: $Q(s, a) = r(s, a) + \gamma \max_{a'}(Q(s', a'))$, where r is the immediate reward, γ is the relative value of delayed vs. immediate rewards, s' is the new state after action a , and a and a' are the action states s and s' , respectively. The action is selected according to $\pi(s) = \arg \max_a Q(s, a)$

The algorithm is as follows. For each state-action pair (s, a) , the table entry $\hat{Q}(s, a)$ to zero. Observe current state s and repeat: select and execute an action a , receive the immediate reward r , observe the new state s' , update the entry for $\hat{Q}(s, a)$ according to $\hat{Q}(s, a) = r + \gamma \max_{a'} \hat{Q}(s', a')$ and make $s = s'$. Iterating this process will make the estimate value of $\hat{Q}(s, a)$ converge to the real $Q(s, a)$.

2.6 Neural networks

Neural networks are information processing systems, an artificial intelligence technology present now for over fifty years. They were developed as mathematical models of biological nervous systems, and brought into widespread interest by McCulloch and Pitts (1943). Neural networks are considered connectionist models. The connectionist approach considers to be no separation between knowledge and the inference mechanism, contrasting with the symbolic approach where knowledge acquisition is separate from the inference mechanism. They also connectionist models because they are comprised of a set of interconnected nodes. Each connection is associated to a value called weight that modulates the potential going through it before connection to the destination neuron. The nodes mentioned are the basic unit of a neural network which are models of biological neurons. Each of these neurons carries out a calculation every time step. Neural networks are parallel distributed computing systems because the whole calculation process that they perform is carried out by all the neuron conforming it.

Neural networks offer the neuron-like processing that natural systems exhibit and because are useful for adaptation and learning. Neural networks can learn to recognise patterns or to approximate any function. Learning involves connection's weights adjustments and there exist several algorithms to this end. Learning algorithms are divided in supervised and unsupervised. In supervised learning the input and the desired output are provided to the network so that at the end of the process the network can reproduce a training set. Unsupervised, or adaptive learning the network is provided with inputs and then decide how to group or categorise them. This is referred as

self-organisation or adaptation.

Figure 2.2 shows an example of a single neuron in a network. The neuron U_i has connections coming from inputs X_1 and X_2 , that can be originated from other neurons' output or from external stimuli. It also has a connection from a bias unit, which always has a value of one. Each connections has a weight that modulates the value going into the neuron. The computational process of the output is in two steps: calculation of the overall input to the neuron (weighted sum of inputs), and the activation of the neuron using a function for getting the output value. This activation function can be a discrete or a continuous one, however, this depends on the application the network will be used on and also on the training algorithm that will be used for adjusting the weights. A common continuous function used is the logistic function (equation (2.1)). The mathematical expression of a neuron's activation is shown in equation (2.2).

$$f(x) = \frac{1}{1 + \exp(-x)} \quad (2.1)$$

$$output = f\left(\sum_{i=0}^n X_i \times w_i\right) \quad (2.2)$$

2.6.1 The multilayer perceptron

There are several types of networks which are defined by their patterns of connections between neurons (architecture), the way the connection weights are adjusted (learning algorithm) and their activation function. A multi layer perceptron consists of an arrangement of one input layer of one or more neurons, one or more intermediate or hidden layers each with one or more neurons and an output layer. The nodes in the input layer encode the

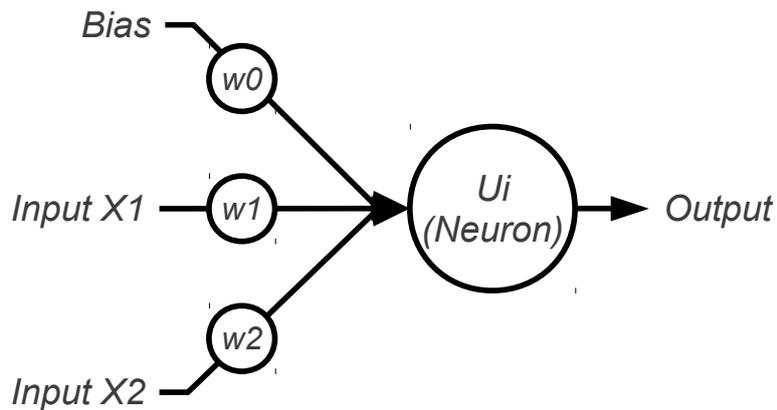


Figure 2.2: The basic element of a neural network.

data presented to the network for processing but do not carry out any calculation, they simply distribute the information to the neurons in the next layer. The hidden layers provide nonlinearity for the data and compute an internal distributed representation of it. The output neurons also carry out processing and encode the output of the system when activated or encode the desired output while being trained. This type of network is usually trained using the *backpropagation of error* algorithm (Rumelhart *et al.* 1986).

2.6.2 The backpropagation of error training algorithm

The backpropagation of error algorithm is a supervised learning method for multilayer feed-forward networks. It is a form of gradient descent that uses the errors generated by the network and the derivative of the activation function for calculating the weights variation. For the case of the function of the sigmoid activation function, the derivative is shown in equation (2.3). Calculated error of each layer is sent backwards (as opposed as when forward-activated) for calculating and applying the weight adjustments.

$$f'(x) = f(x)[1 - f(x)] \quad (2.3)$$

The backpropagation algorithm is useful for training a multilayer feed-forward networks to approximate arbitrary non-linear functions, for regression and for classification problems. It aims to model a specified function by modifying the weights of input signals in order to produce a specified output. The error between the system's output and the known one is used by the algorithm to change the weights that connect the layers. Once trained, the weights in each layer represent abstractions of the mapping between input and output vectors for patterns the system was presented with during training. Each layer of the network abstracts the information processed by the previous layer and therefore the network can combine an arbitrary number of functions for high order modelling.

For teaching the network a single supervised data pair from a training set, the algorithm first makes a forward pass using the input. The input is put in the first layer and is used for activating the next layer, the activation of this layer is then used as input for the next one until the output layer is reached. Hidden and output neurons use the activation function on their input data sums. The sum of input is equal to the sum of the products of the incoming signals from the previous layer and the weights of the input connections from those neurons. This sum of inputs is expressed in equation (2.4), where N_i are the input neurons to neuron x_j .

$$x_j = \sum (w_{ij} * x_i) \quad (2.4)$$

The activation of each neuron x_j is defined in equation (2.2) has been defined earlier (page 70) and uses the non-linear activation function that in this case is the sigmoid. This is calculated for every neuron in the hidden using the input values, then using the hidden layer activation the output layer activation is calculated.

The next step is the backpropagation. Once the last layer has the output values, this activation is compared to one known output value of the training set that corresponds to the input fed to the system. The difference between this vectors is the error which is back-propagated for adjusting the weights. For calculating the error the algorithm uses the sum-of-squares error function, so that errors all have the same sign and the minimisation of the function can be done. The sum-of-squares error function calculates the difference between the system's output and desired output for each output node, squares them and then adds them together. This error function is presented in equation (2.5), where X_k is the system's output and Y_k is the expected known output. This error is used for obtaining the first Δ that will be used for calculating the weight adjustment of the hidden layer.

$$E = \frac{1}{2} \sum_{k=1}^n (X_k - Y_k)^2 \quad (2.5)$$

The calculation of values used for adjusting the weights (ΔW) of each layer is done by applying equation (2.6). The first calculated delta is for the hidden layer, this is then used for calculating ΔW of subsequent hidden layers successively until the first hidden layer is adjusted. The amount of change in each backpropagation process can be modulated by changing the learning rate

η and is applied at the end of the backpropagation.

$$\Delta W_{ij}(t + 1) = -\eta \frac{\delta E}{\delta W_{ij}} \quad (2.6)$$

For teaching a network a whole training set, an iterative process is used. Code 2.1 shows the pseudo-code for this training using the backpropagation training algorithm.

Code 2.1: Pseudo-code for the backpropagation of error training algorithm

```
1 Initialisation of weights with random values.
2 For a specified number of training epochs (iterations) or while the error is
  above a specific value do:
3   For each input/output known pattern from the training set
4     Calculate the system's output feeding it with training pair's input
5     Calculate output neurons error
6     Calculate hidden neurons error
7     Calculate weights variations ( $\delta W_{ij}$ )
8     Apply the change of weights variations
```

2.6.3 Recurrent neural networks

When a neural network has connections that form a loop, unlike feed-forward networks, it is called a recurrent neural network (RNN). A simple architecture of recurrent neural networks is that where the output units are connected to the input ones (figure 2.3). Feedback to the network is given with one time unit delay, although this can be changed. The characteristic of possessing loops in its structure enables the networks to do temporal processing and learn sequences, so that they can perform sequence recognition or reproduction, and temporal association or prediction. Recurrent neural networks are used mainly in supervised learning. However, algorithms for unsupervised learning have been recently explored and implemented.

Another common architecture of recurrent neural networks is where the

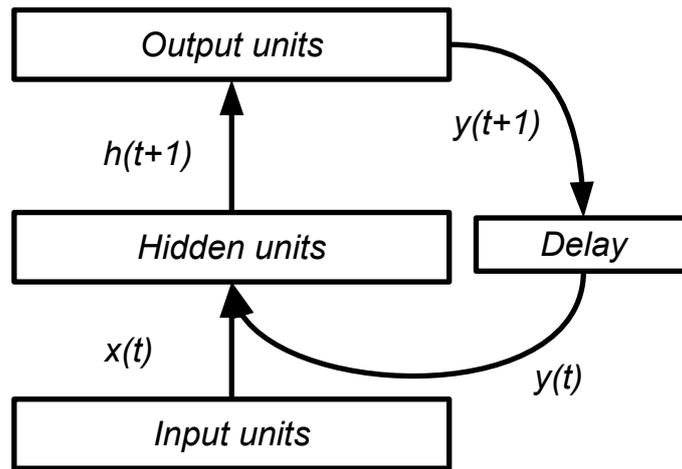


Figure 2.3: A simple recurrent neural network.

activation of the hidden units for one activation pass feeds back into the network along with the inputs in the next step (figure 2.4). If the connection weights in this architecture are fix, then the network is an example of an *Elman network* and the units that store the previous hidden activation are called *context units*. For the architecture shown in figure 2.3, the units that store the activation of the output are called state units and that type or recurrent network is called a *Jordan network*. These two types of architectures are called *simple recurrent networks (SRN)*.

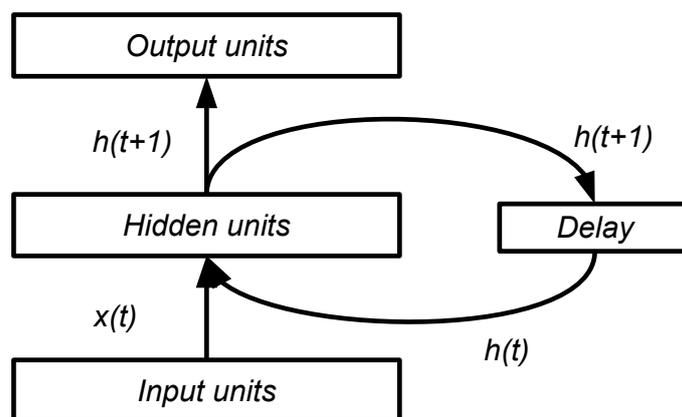


Figure 2.4: A simple recurrent neural network with an architecture typical of Elman networks.

Two of the most simple architectures of recurrent networks have been

presented. There are many more architectures, however, they all share the incorporation of some form of multilayer perceptron as a sub-system, and that they exploit the non-linear capabilities of it in addition to another form of memory.

For training recurrent neural networks gradient descent techniques can be used such as the back-propagation algorithm used in feed-forward networks. Another possible training algorithm to be used on them is a natural variation of back-propagation called back-propagation through time. This algorithm performs gradient descent on a complete unfolded version of the network which contains no loops.

Another application for recurrent networks is creating *associative memories*. The goal of associative memories is to recognise input vectors that have been previously learnt even in the presence of noise or when the input is incomplete. Although associative memories can be created with networks without recurrence, in general recurrent networks produce better results. The *Hopfield network* is an implementation of such memories. These networks are fully-connected (every network is connected to all the others) and have a single layer and units do not have self-loops (figure 2.5). These networks are trained using Hebbian learning and converge to “remember” a given value when part of that value is presented to the units as input.

Seen as dynamical systems, three aspects of recurrent neural networks can be studied: their stability, which concerns the boundaries of the output over time and the response to small changes of input or connection weights, their controllability, concerning the possibility to control their dynamic behaviour so that it can be taken to a desired state, and observability, that tries to find

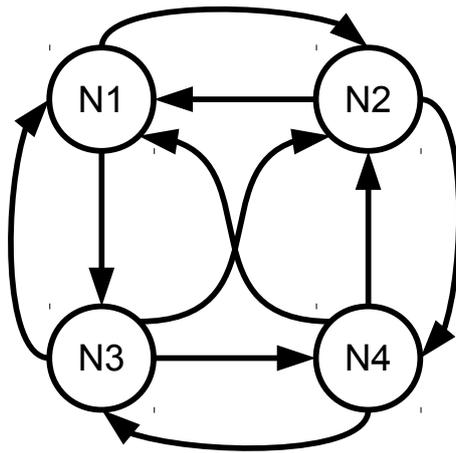


Figure 2.5: A Hopfield network with 10 connections.

whether it is possible to observe the results of the control applied with a finite set of input and output measurements.

2.6.4 Self-organising maps

The self-organising map (SOM) proposed by Kohonen (1982), is a network that implements an unsupervised learning algorithm. SOM networks are inspired in the cerebral cortex, where neurons that get excited by certain type of sensory input are clustered together whereas neurons that respond to very dissimilar input are far away from each other. Self-organising networks learn to cluster groups of similar input vectors from a high dimensional input space in a non-linear way into a low-dimensional, usually two-dimensional, discrete array of neurons on an output layer. The SOM does this in a way such that neurons topologically located close to each other are activated by similar input patterns. The reason why the output layer is usually two-dimensional is for visualisation purposes.

The most common architecture of a SOM consists of two layers. An input layer and a Kohonen or output layer. The SOM output layer is usually a rectangular or hexagonal lattice. The input layer consists of one neuron for

each dimension of the input space (feature) and each of these neurons is connected through adjustable weights to each neuron in the output layer. The weight vectors in the output layer store a representation of the distribution of the input vectors in a topologically preserved order.

For the self-organising map training it is necessary to normalise the samples from the input space. At the start of the training process the weights are randomly initialised. After initialisation, a competitive unsupervised learning algorithm is applied repeatedly to the SOM. In this process, the network compares each presented input with each of its weight vectors (the set of all these weight vectors is called the reference or codebook). The neuron with the better match to the presented input is called the best matching unit (BMU). The most common match measure is the euclidean distance between input vector and SOM weight vector. Once a the BMU has been found, its weight vector and the weight vector of neighbouring neurons are modified in such way that they become more similar to the input vector. The change on each neuron's weight vector is proportional to the topological distance from the BMU. The most common neighbourhood function used is the Gaussian $N_{j'j} = e^{-\frac{\|r_{j'} - r_j\|^2}{2\sigma^2(t)}}$, where $N_{j'j}$ is the neighbourhood function for the BMU j' at iteration t and $\|r_{j'} - r_j\|$ is the distance between neurons j' and j on the output layer. The training process is repeated until the weights change is negligible.

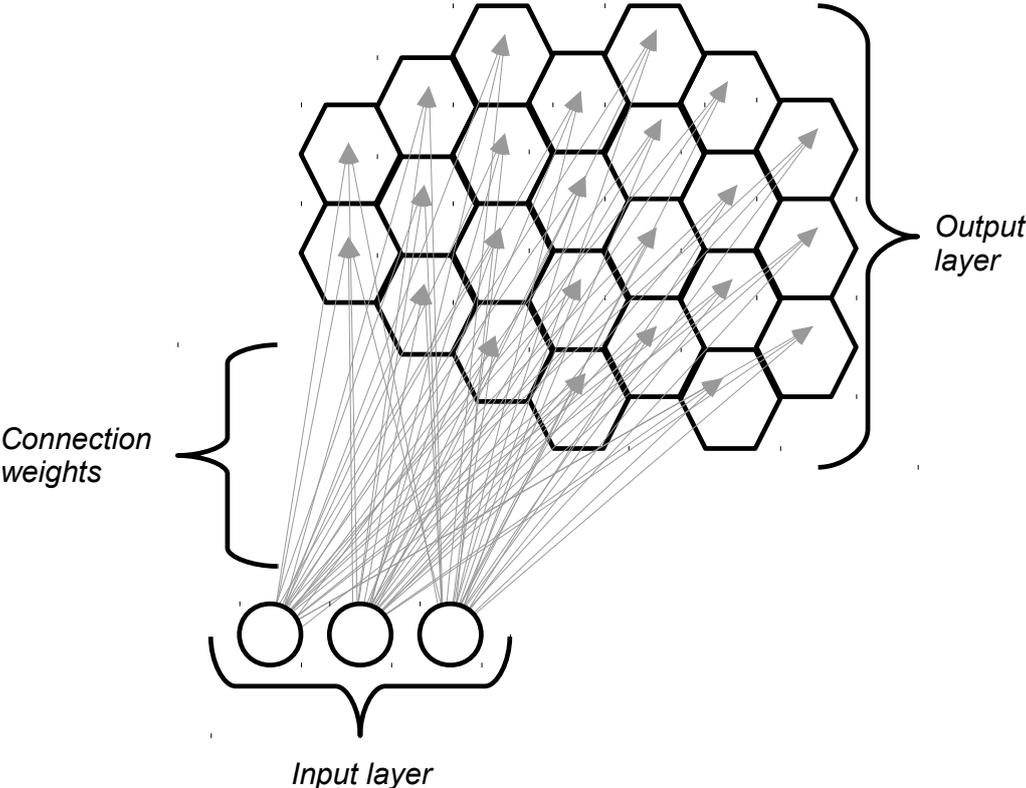


Figure 2.6: A SOM with a three-dimensional input layer and a two-dimensional, hexagonal lattice, output layer.

3 Methods: the iCub Robotic Platform and Custom Software

3.1 The iCub humanoid robotic platform

Experiments for this research were carried out on a simulated iCub robotic platform (Metta *et al.* 2008). The iCub robot is one of the most advanced research platform for humanoid robots (Greggio *et al.* 2008; Narioka *et al.* 2009). Its full software and hardware architectures are distributed as open-source (GPL/FDL licenses) and available from the iCub website¹.

Epigenetic robotics (see section 2.1.5) has been a proposed as a principle for creating a new generation of robots, especially humanoid ones because this approach aim to explain the processes an intelligent robot would need to go through (Ishiguro *et al.* 2011). Recently, infant humanoid robots have been used for the object of this study, reaching behaviour (Lee *et al.* 2012a) as well as for other type or studies, like gait development (Degallier *et al.* 2008; Narioka *et al.* 2009), social interaction with humans (Breazeal and Scassellati 2000; Kozima and Yano 2001) and are proposed to be an investigation platform for language acquisition (Asada 2012).

The iCub robot is an infant-size humanoid robot with 53 degrees of freedom driven by electric motors. It has two high frame-rate, high-resolution

¹<http://www.icub.org/>

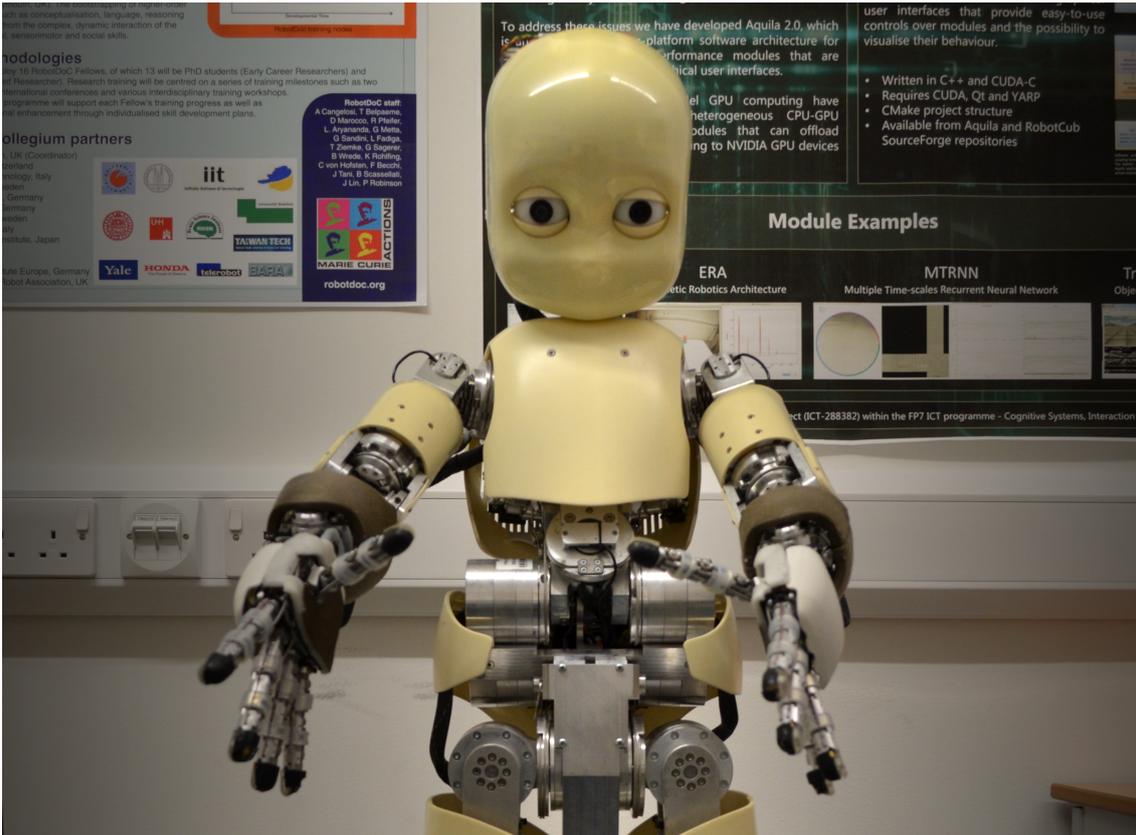


Figure 3.1: The iCub robot

cameras (model: PointGrey Dragonfly2), condenser electret microphones, gyroscopes and linear accelerometers, torque and touch sensors. One great advantage above other similar robotic platforms is the inclusion of an also open source simulator (figure 3.2) displaying realistic physical interactions (using the Open Dynamics Engine) and full 3D rendering (with OpenGL) for testing most aspects of the real robot and research. It allows for studying an embodied agent even without the whole robot platform. The idea is that controllers, applications or frameworks for the robot can be developed and tested in the robot and then ported to the physical robot.

The iCub is a platform suitable for studying object interaction/manipulation behaviours because of its 53 degrees, of these, 41 are for the upper body (see table 3.1 for a detailed description of the DOF's).

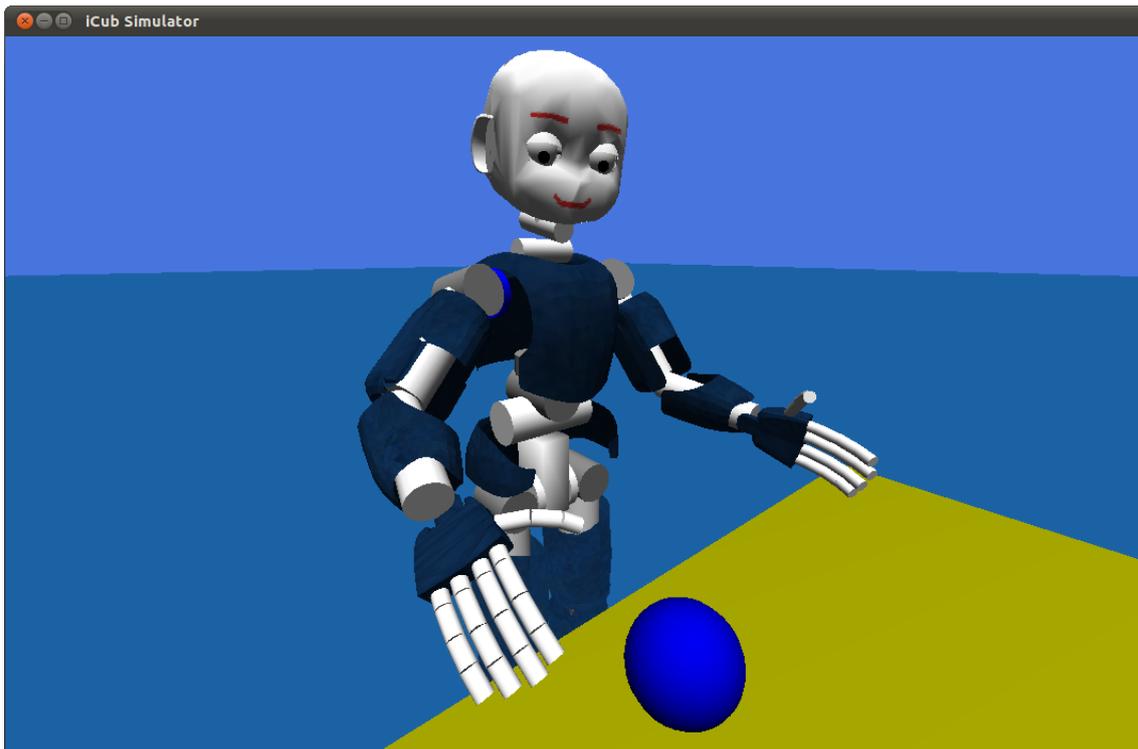


Figure 3.2: The iCub Simulator

Special attention was paid to the arms and head. The head has six degrees of freedom including one for vergence movements of the eyes. The iCub robot was specifically designed for serving as a research tool or platform for embodied cognition (Beira *et al.* 2006) and since it was built, it has been one of the most complete and advanced humanoid robots available.

The iCub platform robotic platform resulted convenient for the studies presented in this dissertation because it allowed testing and measuring the generated behaviours in order to compare them with a natural counterpart. The iCub platform and machine learning methods were used for teaching a humanoid to perform movements related to the basic human skill of reaching. Also, the resulting organisation of the learning structures involved and the learned end-effector positions generated by the movements were analysed to get an insight of the artificial spacial notion of the system. The research

Table 3.1: Details of joints in the iCub humanoid robot.

Body Part	Joint Number	Description
Head	0	Head tilt
	1	Head roll
	2	Head pan
	3	Eyes tilt
	4	Eyes pan (version)
	5	Eyes vergence
Arms/Hands	0	Shoulder pitch
	1	Shoulder roll
	2	Shoulder yaw
	3	Elbow
	4	Wrist pronosupination
	5	Wrist pitch
	6	Wrist yaw
	7	Hand finger
	8	Thumb opposition
	9	Thumb flexion/extension of the most proximal joint
	10	Thumb flexion/extension of the must distal joint
	11	Index finger flexion/extension of the must proximal joint
	12	Index distal flexion
	13	Middle proximal flexion/extension
	14	Middle distal flexion
15	Ring and little finger flexion	
Torso	0	Torso yaw
	1	Torso roll
	2	Torso pitch
Legs	0	Hip pitch
	1	Hip roll
	2	Hip yaw
	3	Knee
	4	Ankle pitch
	5	Ankle roll

experiments were carried out on the simulated version of an iCub humanoid platform (Tikhanoff *et al.* 2008). A decision was made for using the simulated iCub due to the nature of the experiments to be undertaken and the continuous availability of a computer simulation. A serious difficulty of using the physical robot is the mechanical stress the motors have to sustain while

performing random exploratory movements. Moreover the available physical robot suffered from a manufacturing fault that easily provoked breakage of the cables that moved the arm and repair was expensive in terms of time. Additionally, although the computer vision algorithms developed in the simulator worked on the real robot, that was not the case of the active vision system, which used velocity control to move head and eyes and the real robot motors do not perform as in the simulator in this mode of operation. An advantage of using the simulated robot is that it was possible to independently control all the joints of the arm with minor modifications to the source code, which was not the case in the real robot that had the shoulder joints coupled for certain position values.

3.2 Communications middleware: YARP

YARP (Yet Another Robotic Platform) is not a robot operating system, instead, it is the “glue” for many computers an operating systems working together in order to make use of the iCub, either the simulated or the physical one. It is a communications library and server available as open-software. YARP is very portable because it relies in very few dependent libraries. In Linux and MacOS the ACE library—“The ADAPTIVE Communication Environment”, a framework that implements many core patterns for concurrent communication across a range of OS platforms—can even be omitted.

Robot control requires communications among the many parts that constitute the system. For this project, those communications requirements are fulfilled by YARP (Metta *et al.* 2006). YARP has been selected as

communication middle-ware mainly because the whole iCub robotic platform has been developed in parallel with YARP, and also because it allows distributing processing in several computer systems through a network. YARP is written in C++ but can be easily interfaced with scripting languages such as Python and Matlab. Figure 3.3 shows a YARP scenario where two robot parts (head and arm) are controlled from different clients using YARP.

YARP provides easy inter-process and inter-system communications for connecting different modules that need to interact in a task. For example, in parts of my research the visual processing needed was done on a separate computer due to limitations of computing resources, then on another computer the simulator was being run.

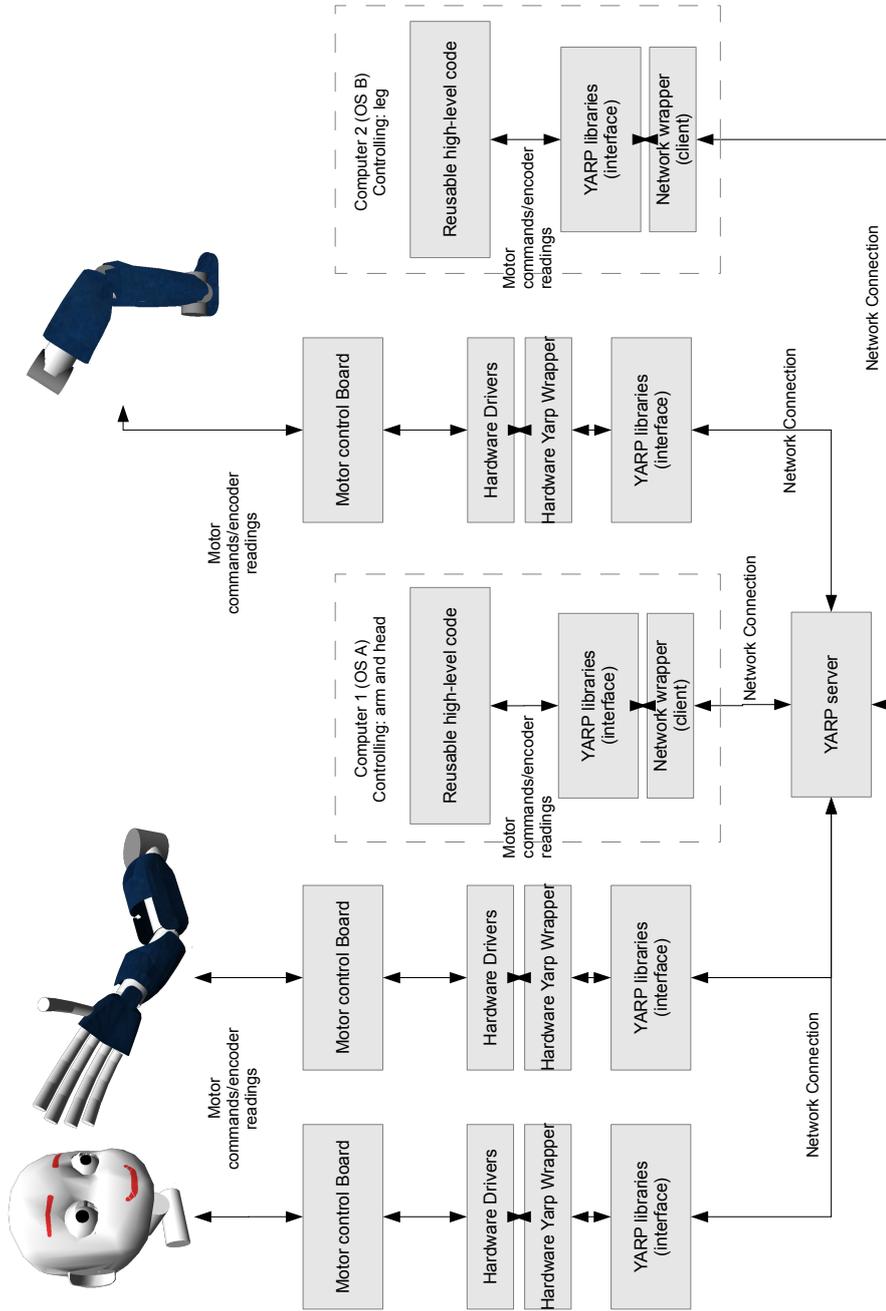


Figure 3.3: YARP allows to have several remote clients for different parts of the robot. Clients can be in different computers even with different operating systems.

By using YARP, several processes can communicate with the resources available in the robot or the simulator. YARP and the iCub platform include a set of programs readily available for this purpose. Figure 3.5 displays the simulator along with some of these programs, like the motor-control GUI² used for easily read the encoders of each motor in the robot and also send commands to it or change velocities. Additionally, purpose-specific programs can be coded and built using the YARP libraries and its several *language bindings*³ as the ones created for this work and presented following.

3.3 Auxiliary software created for this project

3.3.1 iCub-S a GUI interface for iCub resources

For this research several pieces of software were created using the libraries provided by YARP and the iCub. The main ones were an interface for the head and visual system of the robot used for the project and named *iCub-S* is shown in figure 3.4. *iCub-S* is a multi-threaded program (it usually ran on a 8 multi-threading cores computer) which provides modules for accessing different resources of the robot/simulator. The iCub-S program has a module that offers a GUI for the cameras that interacts with another module that performs image processing and easy access to the parameters via visual controls. The program has a tracking behaviour module used for this work's experiments (whose implementation and algorithm will be detailed in section 4.1.1). Available modules developed for the program are shown in table 3.2 The whole multi-threaded program was coded in C++ language and

²Graphical User Interface.

³A *language binding* is an interface between two different programming languages. Essentially they are wrapper libraries bridging the two languages so that a library written in one language can also be implicitly used from the other language.

it's GUI used the Qt application framework⁴.

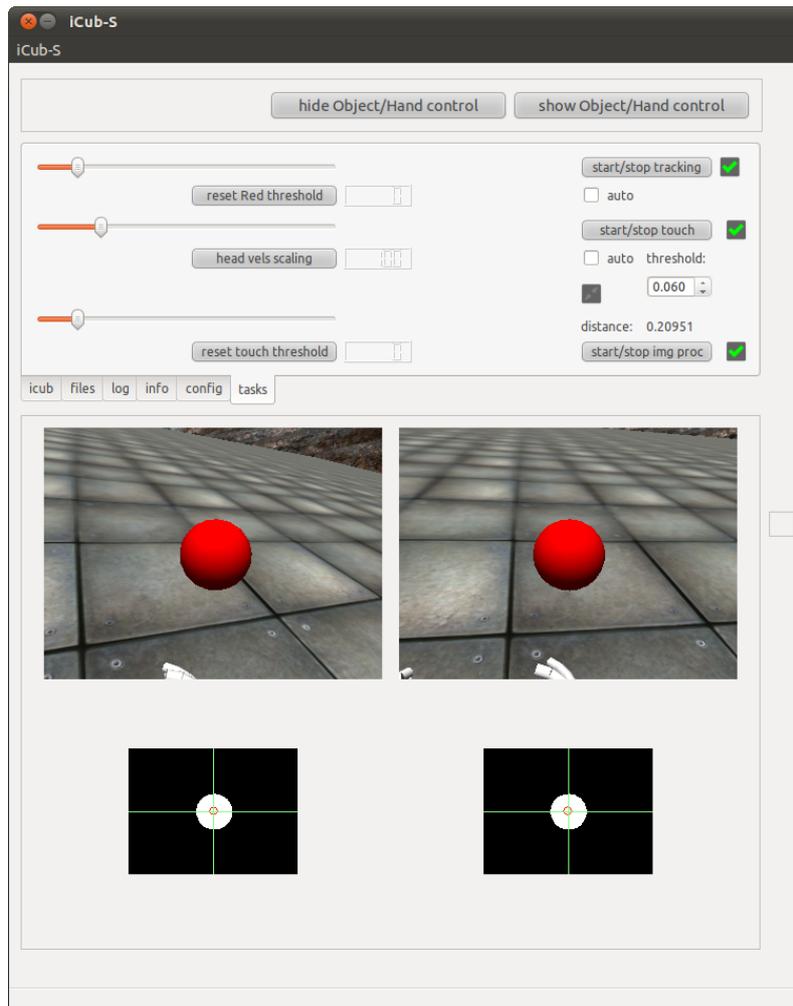


Figure 3.4: iCub-S, a program developed for being an interface for the cameras and for carrying out image processing was a multi-threaded program that uses YARP libraries and Qt for the GUI.

Table 3.2: The modules available in iCub-S auxiliary program.

Module Name	Purpose
ImgMain	Image acquisition
ImgProc	Image processing
WorldControl	Object creation and manipulation in the simulator
EyesNeckTrack	Visual object foveation and tracking

⁴Qt is a cross-platform application framework that includes a widget library for creating graphical user interfaces among many other features.

3.3.2 Auxiliary Python libraries

Other software created for this project were two auxiliary Python libraries: *pySalo.syarp* and *pySalo.sutils*. Both provide easy access from Python scripts to YARP and iCub resources and are also suitable for tests and interactive sessions using iPython. Usage examples of these libraries are shown in code 3.1 and code 3.2. Full source code and documentation of these libraries can be found at <http://sourceforge.net/users/salo9000>.

Code 3.1: Basic usage example of auxiliary library *pySalo.syarp* used in a Python script. A *DataSource* object is created.

```
1 # our program uses YARP
2 import yarp
3
4 # network initialisation is required in every YARP program
5 yarp.Network.init()
6
7 # for using the auxiliary library
8 import pySalo.syarp
9
10 # reader will be an object for easy access to the head of the robot
11 # it will create a port called /head_reader connected to /icubSim/head/state:o
12 reader = pySalo.syarp.DataSource('/head_reader', '/icubSim/head/state:o')
13
14 # reading the port:
15 current_value = reader.getNArrayRead()
16
17 #deleting the object and closing the YARP port:
18 del reader
```

Code 3.2: Basic usage example of auxiliary library *pySalo.sutils*. The script creates a *RobotPart* object for accessing the robot's head.

```
1 # our program uses YARP
2 import yarp
3
4 # network initialisation is required in every YARP program
5 yarp.Network.init()
6
7 # calling the auxiliary library sutils
8 import pySalo.sutils
9
10 # creating a RobotPart object used for accessing the head of the robot
11 head = pySalo.sutils.RobotPart('/icubSim/head')
12
13 # getting a dictionary with the limits of the encoders
14 head_limits_dictionary = head.limitsDict
15
16 # getting the number of degrees of freedom of the robot part
17 head_number_of_axes = head.nAxes
18
19 # reading the current positions of the encoders
20 current_encoders_positions = head.getEncoderPositions()
21
22 # setting velocities for all the degrees of freedom of the robot part
23 head.setVels(30)
24
25 # sending a position control command
26 head.setPositionDict({0:10,1:15,2:5,3:10,4:25,5:0})
27
28 # stopping the movement of the robot part by sending a velocity command with null
    velocities
29 head.velocityMove([0,0,0,0,0,0])
30
31 # sending and rpc command to the simulator world port and printing the result
32 b = pySalo.sutils.rpcSend('/icubSim/world','world_get_rhand')
33 print b.toString()
```

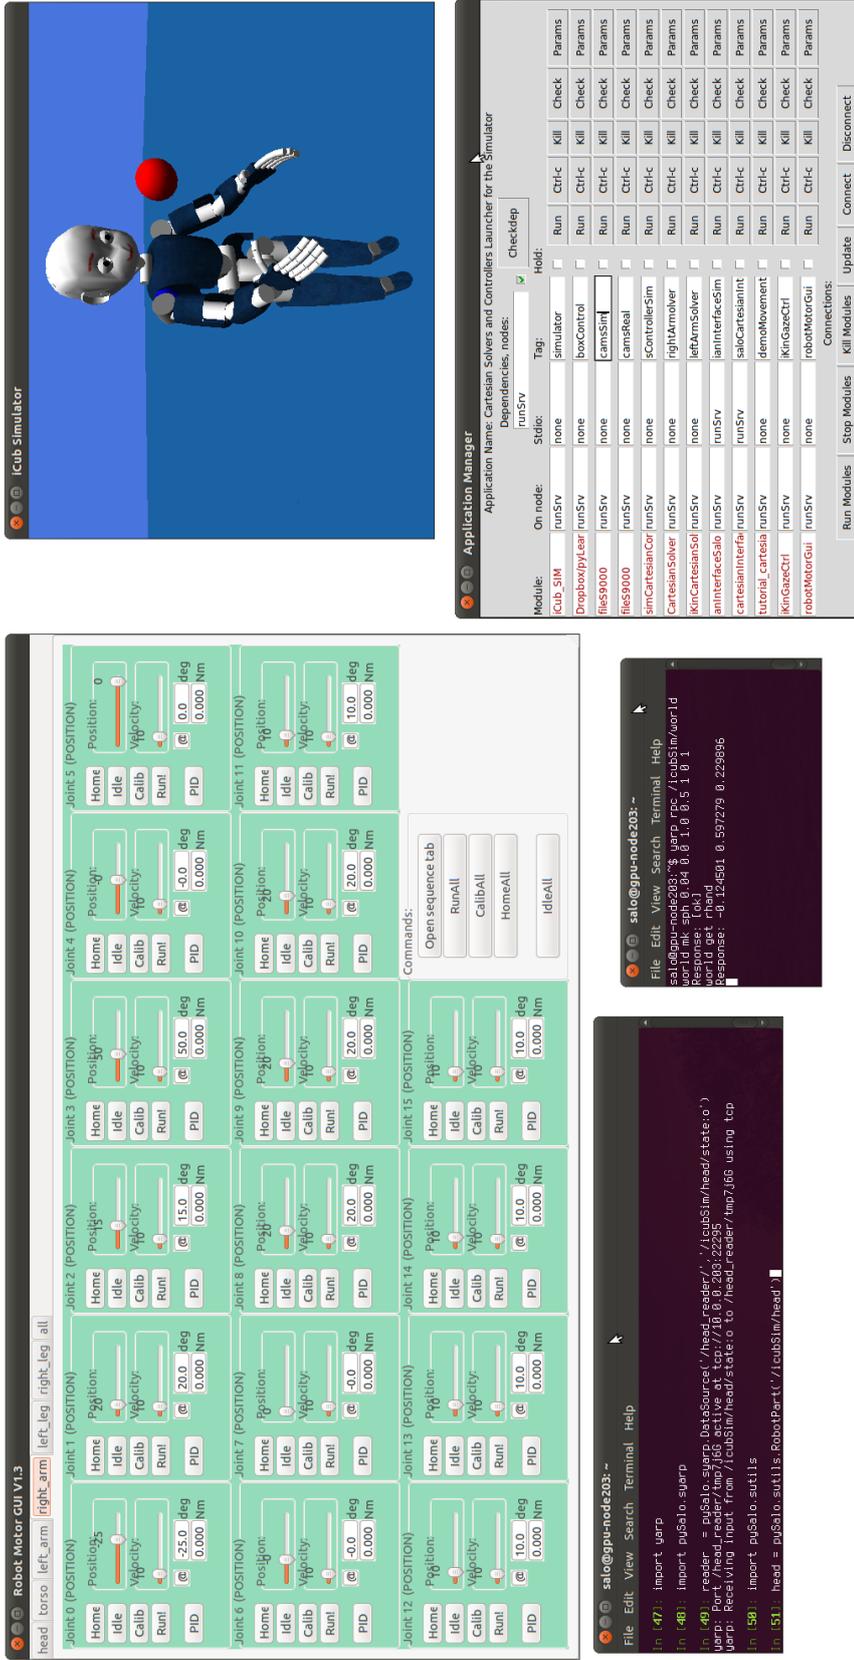


Figure 3.5: The iCub simulator running along with typical programs for accessing the robot: a robot motor GUI, an application manager and command line terminals.

4 Experiment 1: Monocular and Binocular Contributions in a Bimodal Reaching Task

Depth estimation in biological systems can be of great importance for a good performance in crucial tasks such as reaching, grasping or avoiding obstacles (Mon-Williams and Dijkerman 1999). From the literature it is known that monocular vision provides cues for depth perception, including motion parallax, accommodation effort, casted shadows by near objects and contrast (Howard and Rogers 1995). Nevertheless, those cues can only be used in certain circumstances and in most of the cases the use of monocular depth indicators needs complex processing on the acquired image (Saxena *et al.* 2005). For this reason, processes and algorithms that extract depth information from vision are widely focused on stereo images, images of the same scene taken from two slightly different positions (Reichelt *et al.* 2010).

Vergence is the oculomotor adjustment needed to foveate the same point in space with both eyes (Leigh and Zee 1999). Besides the possible algorithms that can be applied to stereo images, vergence is an additional proprioceptive information that is available to organisms endowed with two movable eyes. Recent studies show that in humans, vergence occurs well before the actual depth perception (Wismeijer *et al.* 2008) and therefore it is an important cue even in the absence of more complex monocular cues or processing of stereo

images.

Being one of the early motor skills developed by infants, reaching origins could shed light on other motor developments that arise later. Moreover, it is a skill that emerges from reflexive behaviours, as it has been demonstrated that the visuomotor mechanisms of reaching and prehension are, to an extent, independent from the perceptual (the ones we are more aware of) ones and visuomotor system priority of binocular cues over pictorial ones (Marotta *et al.* 1997). This findings are possibly related to phenomena found to be present in peripersonal space such as involuntary movements for defence, object avoidance and/or reaching (Graziano and Cooke 2006). The shaping of these behaviours and later control of them that leads to grasping can be explored in a humanoid robot. In the present experiment we studied the possible relevance of vergence in the development of a peripersonal space representation. We also explored the contributions that two different perceptual modalities, visual and proprioceptive have in a system for reaching implemented in the iCub simulator.

4.1 Methods

Experiments were carried out in the simulated version of an iCub humanoid (Tikhanoff *et al.* 2008). The robot's task was to reach a red cube placed in front of the robot with the right hand. The five most proximal degrees of freedom of the iCub arm were used. These are specified in table 4.1. For this experiment and subsequent ones in the present work, a visuomotor system was developed.

Table 4.1: Robot arm joints used in experiment 1.

Arm joint number	Joint	Limits[degrees]
0	Shoulder pitch	[-95.0, 10.0]
1	Shoulder roll	[0.0, 160.8]
2	Shoulder yaw	[-22.0, 95.0]
3	Elbow	[15.5, 106.0]
4	Wrist prosupination	[-90.0, 90.0]

4.1.1 A visual perception system and tracking behaviour for iCub

The models used in this work made use of information from the visual modality of the robot. The iCub's head, both in the real robot (Beira *et al.* 2006) and in the simulated version, is provided with six degrees of freedom for the head: one for tilting and one for panning the eyes, one for tilting and one for panning the head, a degree of freedom head rolling, and finally, eyes' vergence (figure 4.1). For the experiments, the robot was required to use vision and perform a tracking behaviour. This behaviour was implemented as a module in the iCub-S program (an interface developed as part of this thesis, see section 3.3.1 on page 88), it consisted of a simple but effective heuristic for foveating targets with both eyes in real-time and low computational resources consumption.

The tracking behaviour is achieved by three processes: one process for each eye image processing and one for motor control that uses information from the two. Targets used in the experiments were red objects. For each incoming video stream, images were colour-segmented using the OpenCV library¹. These three processes were implemented as modules in iCub-S program (section 3.3.1 on page 88). The motor control process for tracking

¹<http://opencv.org/>

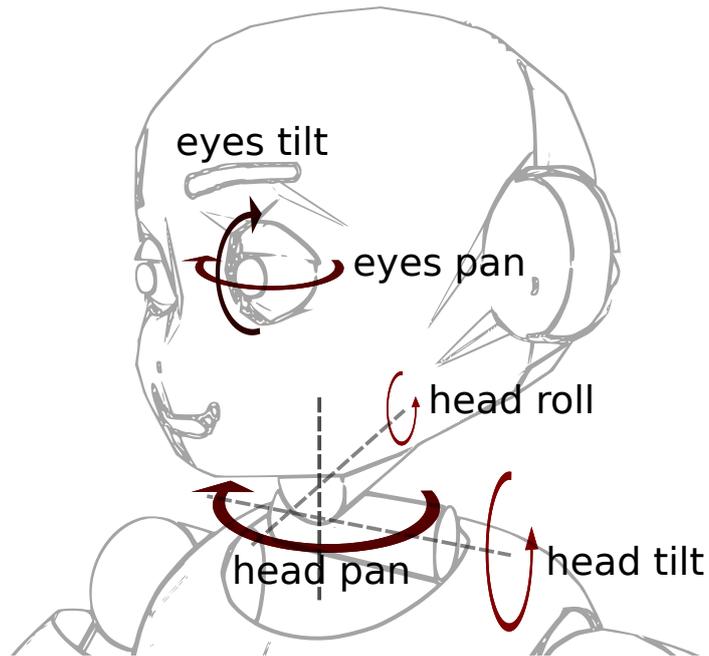


Figure 4.1: The iCub head showing 5 of it's 6 degrees of freedom. The sixth one is vergence.

behaviour is presented in code 4.1 and it consists of detecting the distance from the centroids of the segmented target to the centre of the image and modulating (tuning) the speed of eye and neck movement for placing these centroids in the centre of the images from both cameras. The modulation uses the Gaussian function equation (4.1) ($\mu = 0$ and $\sigma = 2$) and the centroid-to-centre distances for scaling a reference speed that was chosen based on manual tests.

$$f(x) = \frac{1}{\sigma \sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)} \quad (4.1)$$

Red blobs detected in the images are identified, then the biggest one is taken as the target to foveate. Each image processing locates the centroid of the blob in corner-centred coordinates which are then transformed to the middle-

Code 4.1: Pseudo-code for the tracking behaviour algorithms implemented in the motor-control module of iCub-S.

```

1 while 1:
2   read coordinates of object centroid from image processing (in each eye)
3   transform coordinates to have origin in the centre
4   if(object in both eyes)
5     calculate distance from the centroids to the origin for each eye
6     calculate mean distance for both centroids
7     normalize means between [-1,1]
8     # for vertical eyes movement:
9     if abs(verticalMeanDist) > VERTICAL_OFFSET_THRESHOLD
10      velCommands[eyes_vertical] = modulate(BASE_VERTICAL_EYES_SPEED)
11    else
12      velCommands[eyes_vertical] = 0
13    # neck vertical movement follows eyes
14    if(fabs(encVals[eyes_vertical]) > MAX_EYES_POSITION_OFFSET)
15      velCommands[neck_vertical] =
16        scaled_current_position_of_eyes_horizontal
17    else
18      velCommands[neck_vertical] = 0 # do not move neck
19    # for horizontal neck movement
20    if(( HORIZ_LEFT_OFFSET >= 0 AND HORIZ_RIGHT_OFFSET >= 0 ) OR (
21      HORIZ_LEFT_OFFSET <= 0 AND HORIZ_RIGHT_OFFSET <= 0 )) AND abs(
22      horizontlaMeanDist) > HORIZONTAL_EYES_OFFSET_THRESHOLD ) {
23      # move the neck
24      velCommands[neck_horizontal] = (-modulate(horizontlaMeanDist)) *
25        velNeckHor
26    else
27      velCommands[neck_horizontal] = 0
28    # for version and vergence:
29    if (abs(HORIZ_LEFT_OFFSET) > 1 OR abs(HORIZ_RIGHT_OFFSET) > 1)
30      if (HORIZ_LEFT_OFFSET>=0 AND HORIZ_RIGHT_OFFSET<=0 )
31        velCommands[eyes_vergence] = VERGENCE_SPEED
32      else if (HORIZ_LEFT_OFFSET<=0 && HORIZ_RIGHT_OFFSET>=0)
33        velCommands[eyes_vergence]= - VERGENCE_SPEED
34      else
35        velCommands[eyes_vergence] = 0
36    else
37      velCommands[eyes_vergence] = 0 # null speed
38    send velocity commands contained in velCommands

```

centred ones. This vertical and horizontal locations are normalised to be in the range $[0, 1]$ and sent trough a YARP port. Another process is responsible for generating motor control from the data reads from the two streams of data generated by the image-processing processes.

The motor control process sends commands to the robot head and eyes for foveating the target. This process issues velocity commands to the head of

the robot and controls the motors of three joints of it: neck tilt, neck pan and eyes' vergence. Eyes' pan and tilt are not used. Tests were made before deciding which degrees of freedom to use and although using the eyes' degrees of freedom allows more human like movements, they don't contribute to precision in the foveating task and once the robot has foveated the target, position data from these joints is redundant. Because this process receives normalised values of the target's position in the images from the eyes, this process is robust for information coming from different size images processing, that is, image processing and motor control are loosely coupled. Although we are using neck movements for tracking, this mechanism was a first step in this project because, as Gunnar and Nelson (1992) mention, "eye movements may be the most sophisticated behavioural capability of the neonate, in that they permit the infant to actively control the acquisition of information about the visual environment, long before other purposeful behaviours such as active reaching and grasping emerge".

The controller endows the system with the possibility of determining target depth because for foveating the target with both eyes, the line of sight of each of them is rotated in the horizontal plane. The angle these lines form with the resting position line of sight is a measure of the target distance from the face. This simple way of getting depth information is going to be tested and exploited by the system in experiments that involve reaching tasks.

Figure 4.2 shows data of 24,700 measures collected during twenty four hours of continuous operation. They show the distance of a target and the correspondent vergence angle the system generated. A regression using a logarithmic fit was used to obtain the fitting model shown in equation (4.2)

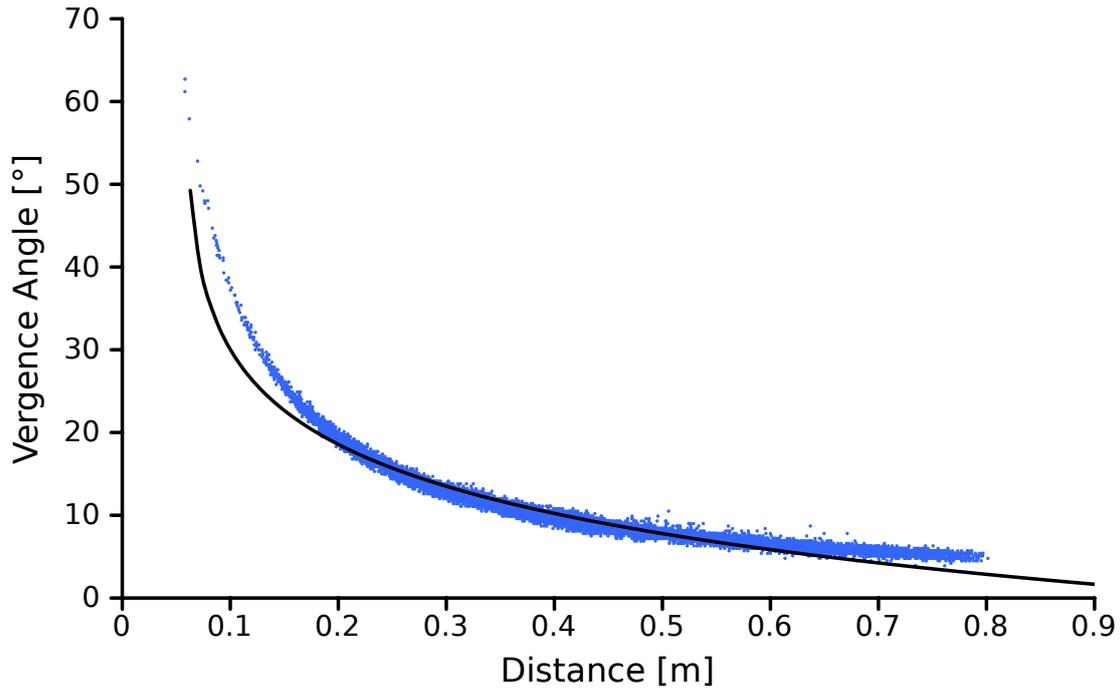


Figure 4.2: Graph showing the distance to vergence angle relation the tracking system exhibits

which has a squared correlation coefficient of $R^2 = 0.958474$.

$$y = 0.0344746 - 9.51428 \ln(x - 0.0573347) \quad (4.2)$$

The graph shows the growing uncertainty as distance is increased. The system displays an inconsistency in its depth estimation this can be due distorted images coming from the cameras or the precision of the simulated encoders and PID controllers, leading to incorrectly estimating centroids of the blobs considered to be the target, however, this kind of error in depth perception is shared by infants in the firsts months of life (Aslin 1977). In infants this is due to a larger Panum's fusion area, which consists of a range of disparities within two images received eyes, experienced as a single object. Objects outside of Panum's fusional range are perceived as double or diplopic.

4.1.2 *Experimental conditions*

Two different conditions of the task were considered: monocular vision and binocular vision. Tracking and foveation the object was achieved by the active vision controller described in section 4.1.1 which moved the head and the eyes of the robot so to locate the target's centroid in the centre of the right eye image, or in both eyes' images, in the monocular and the binocular conditions, respectively. For the binocular vision condition, modulation of vergence was needed in order to foveate the target in both images coming from the eyes. Figure 4.3 shows the simulated iCub performing the task.

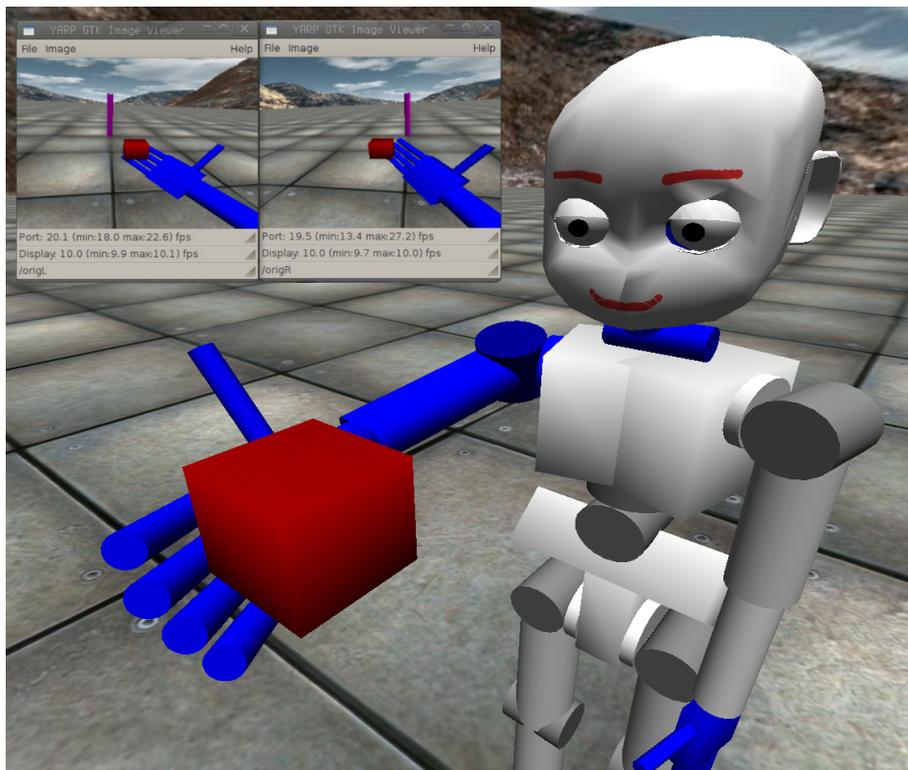


Figure 4.3: The simulated iCub performing the reaching task. Colours of arm and target are used for image segmenting.

4.1.3 *Arm neural controller*

A neural controller was used for moving the right arm of the simulated robot. Input for this controller was the proprioceptive information (pan and tilt

joint positions from the head, and tilt, pan and vergence joints from the eyes) and preprocessed visual information. Preprocessing was colour based image segmentation: red was used in the case of the target object and blue for the arm of the robot. Image segmentation provided then two sets of data that were fed into the neural controller.

The controller was a feed-forward, partially connected neural network with the following architecture: one input layer (three units corresponding to tilt, pan and vergence joints, plus 160×320 units for each eye); one hidden layer hA which received connections from visual input; an additional hidden layer hB (10 units) which received connections from proprioceptive input and hidden layer hA (10 units); an output layer (5 units, one for each arm joint in the robot) which received connections from the proprioceptive input and the two hidden layers hA and hB . This architecture was devised in order to analyse *unimodal* and *bimodal* contributions to depth perception. A diagram of this partially connected network is presented in figure 4.4. The hidden output layers used a sigmoid activation function. Also, the architecture takes inspiration to some extent from the hypothesis on how the brain processes vision proposed Goodale and Milner (1992), where the dorsal stream, corresponding here to the head and eyes, including vergence, proprioceptive components (left part in the diagram), is not used for abstract planning but for control of elementary movements and the ventral stream corresponds to the use of retinal information (left part in the diagram), that in the Goodale and Milner's model is used for determining identification of the object that in this case is only the perceived size. However, in our model additionally there is a central pathway that was included in order to partially integrate both

mentioned pathways. Note that the output layer has connections from the direct proprioceptive inputs, visual feature selection done by layer hA , and the integrating layer hB .

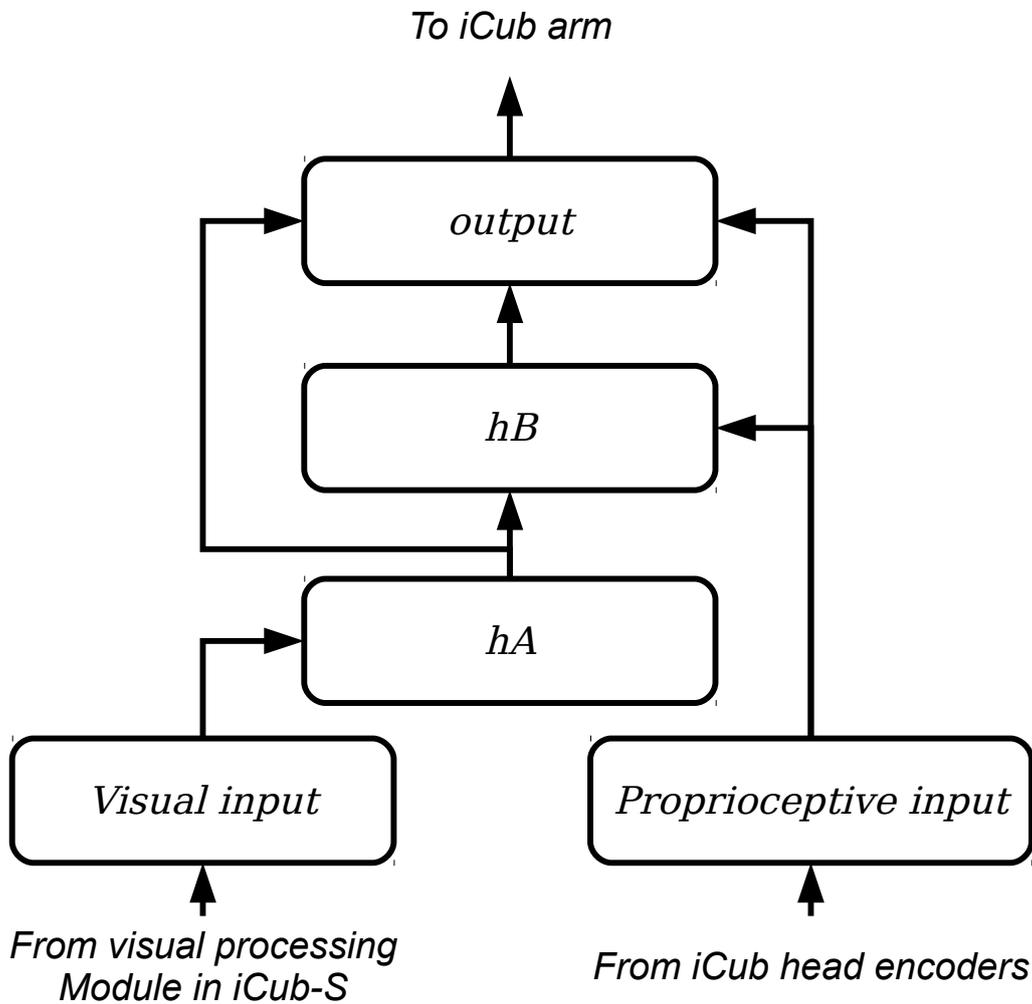


Figure 4.4: The four layer, partially connected feed-forward network used for the arm controller.

Training data consisted of 120 input/output pairs. Inputs were proprioceptive data coming from head and eyes plus a vector containing image data coming from the visual processing module of the iCub-S program (which acquires the images from the cameras and process them see section 3.3.1) which performed colour-segmenting on the red target. The images from the cameras were 160×320 pixels. Desired outputs

corresponded to a vector of arm joints positions needed for reaching the target. Collection of training data was done by allowing the robot to move its arm and placing the target in the hand once the movement was finished, real time tracking of a red marker in the hand produced the correspondent foveating head and eye postures. In order to avoid self-collisions affect the experiment, the simulator was provided with a collision detection mechanism that stopped the robot movement when the arm or hand touched other parts of the body (head, torso). In self-collision cases data was discarded and the trial restarted. Sixty data pairs (half of the dataset) were obtained using monocular vision, with the tracking algorithms using only one eye, and the other sixty using binocular vision. In the monocular data pairs, the values of the right eye input were null. The neural controller was trained using error backpropagation (Rumelhart and McClelland 1986) for 10,000 epochs using $M = 0.1$ as momentum and a learning rate of $\alpha = 0.01$. The general system is shown in figure 4.5. For investigating the main focus of the present dissertation, which is peripersonal space, the behaviour associated with documented differences in the representations (either implicit or explicit) of near and far space had to be put in place as the ground level for the research. Namely, the reaching behaviour already present in infants in their first year of age. To attain this it was decided to use the supervised learning algorithm of back-propagation of error for training the controller. It was not required that the robot discovered latent variables in the output space (the reaching positions), which is what unsupervised learning do, but to map sensory data from the visual and proprioceptive data to motor commands to reach to known locations in space.

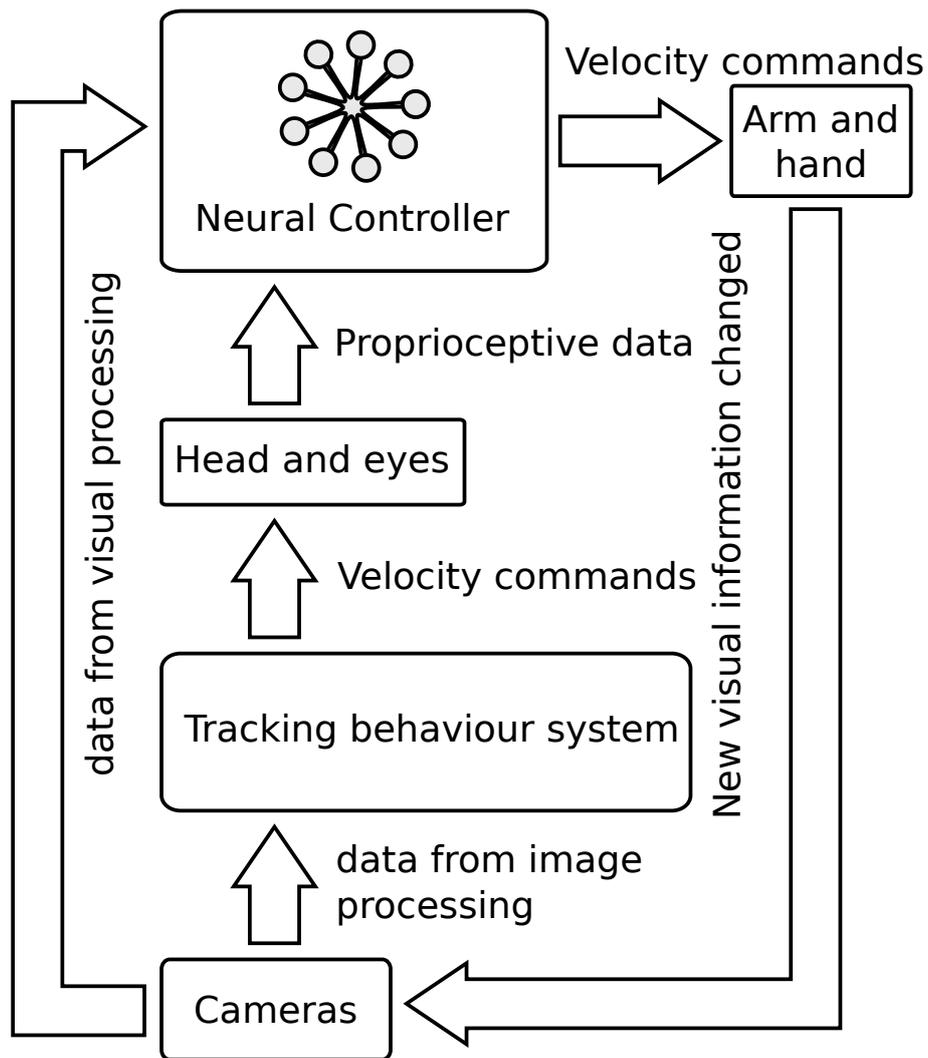


Figure 4.5: The active vision process with the neural controller.

4.2 Results

Test data was collected by placing the target object in 18 positions located inside a volume that the robot was able to reach with its right hand. The positions were distributed evenly in two planes, at $Y = 0.86m$ and $Y = 0.96m$ correspondingly, having therefore nine points in each plane. The points on the X axis ran from $-0.12m$ to $0.12m$ and along Z from $0.16m$ to $0.32m$. For each position the vision controller was activated in order to track and foveate the target object and then the arm motor controller activated to perform a reaching attempt. The reaching movement was always started with the arm in a home position parallel to the body of the robot. Two cases were tested using the same trained controller: *binocular* vision, that used vergence and *monocular* vision, where only the left eye was used. In the latter case the tracking behaviour algorithm used only the image coming from one eye and the visual information fed to the network also was only coming from that eye, that is, the inputs to the network belonging to it were set to zero. For both cases the same target points were used. For each target location, the arm motor-control network was continuously activated until movement of the arm stopped. Reaching accuracy, depth perception and end effector orienting were measured for each trial.

Reaching accuracy was measured as the distance between the target and the centre of the palm of the hand. Depth perception error was taken as the absolute difference between the distance from the head and the target and the distance between the head and the hand, both along the horizontal plane (XZ). Orienting error was measured as the angle between a vector from the head (a point between the eyes) to the target and a vector between the head and the hand in the horizontal plane (figure 4.6 shows the reference frame used by the

simulator).

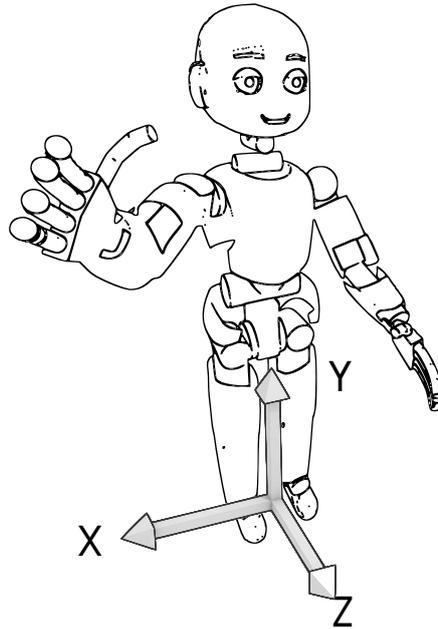


Figure 4.6: iCub Simulator reference frame. The horizontal plane consists of XZ and the origin of the axes is on the floor between the location of the feet.

4.2.1 Analysis

Three variables were measured: reaching distance error, depth estimation error and orienting angle error. For all three of them, error decreased when vergence was used (binocular vision). Reaching distance error in the monocular case showed a mean of $M = 0.1243$, $SD = 0.0371$ which was larger than in binocular, where the mean was $M = 0.1040$, $SD = 0.0296$. Depth estimation error mean was $M = 0.0580$, $SD = 0.0368$, larger than $M = 0.0503$, $SD = 0.0344$ in the binocular. Finally, orienting angle error measured on the Y plane had a mean of $M = 12.8644$, $SD = 8.8475$ in the monocular, larger than $M = 8.1419$, $SD = 6.7262$ in the binocular.

Results showed that when the system had access to binocular data there was a decrease of 16.3% in distance to target error and a 13.3% decrease in

depth estimation error measured in the horizontal plane. These two measurements can be seen in figure 4.7. The other measurement taken, orienting error is shown in figure 4.8. This measure reported the highest decrease which was of 36.70%.

Independent *t*-tests were carried out using the free-software R language for statistical analysis² to compare the monocular and binocular conditions. For the depth estimation error clearly there was no statistical difference. However, for reaching distance and also for the orienting angle the result of the test was close to the margin to be different. Table 4.2 details the *t*-test results.

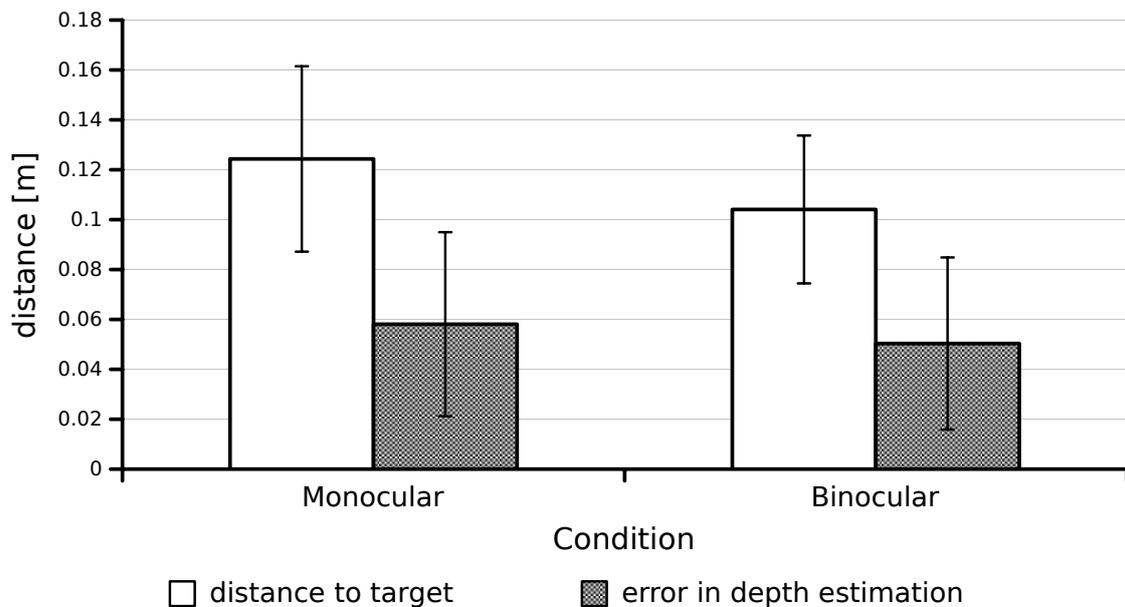


Figure 4.7: Comparison of the two visual systems' performance in terms of reaching distance error. Bars indicate standard deviation.

²The GNU R Project for Statistical Computing: <http://www.r-project.org/>

Table 4.2: Performance of the system in the two conditions compared using t -test analysis.

Reaching distance error	Depth estimation error	Orienting error
$t(34) = 1.8076$	$t(34) = 0.6493$	$t(34) = 1.8027$
$p = 0.0795$	$p = 0.5204$	$p = 0.0802$

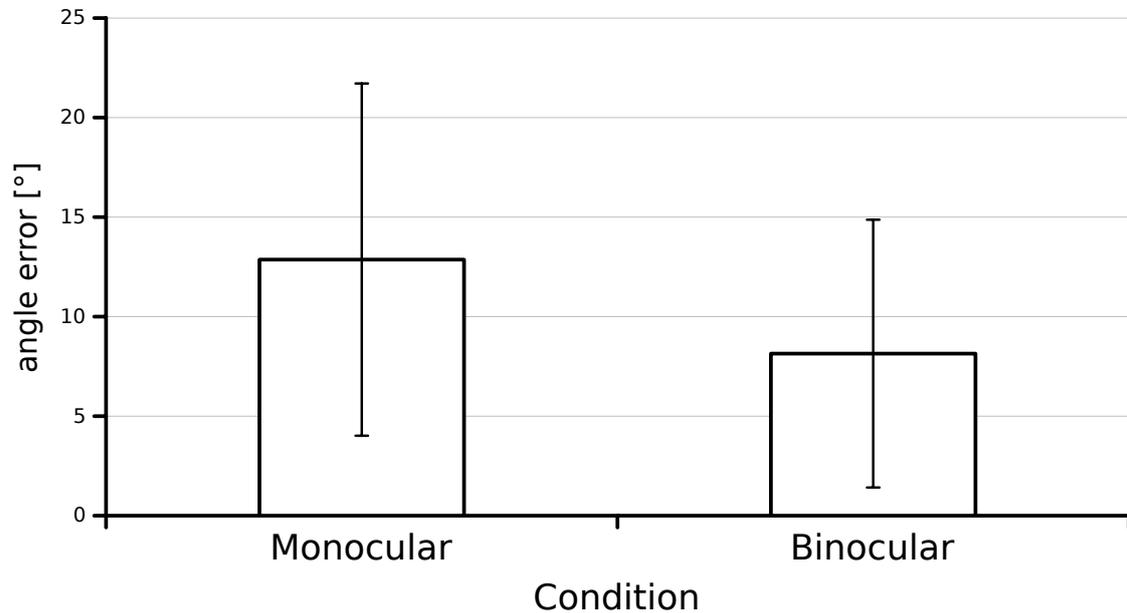


Figure 4.8: End effector orienting error for the two cases. Bars indicate standard deviation.

4.3 Discussion

In this study we examined the contribution of binocular vision to the control of the robot grasping. The relevance of binocular vision is a well known subject of study in humans that continues to be extensively researched (Banks and Salapatek 1983; Barela *et al.* 2011; Ekberg *et al.* 2013; Hu and Knill 2011; McKee and Taylor 2010). One aspect associated to binocular vision currently debated is whether we need it for depth perception or if monocular vision is sufficient, and also what are the contributions from each of them in prehension tasks. To investigate this Marotta *et al.* (1998) asked participants to reach a self-illuminated target in darkness. The target was located on the same horizontal plane as the eyes and two conditions were tested: monocular and binocular vision. They found less on-line corrections in the monocular condition when participants were allowed to move their head than where their head were fix, while in the binocular condition no difference was found. This suggests that under normal circumstances (when we can move the head) the visuomotor system “prefers” or relies more on binocular vision and uses monocular ones (motion parallax) as a last resort in prehension tasks. Similarly, it is likely that a system based on our model would select binocular capabilities due to the improved accuracy.

4.3.1 Contribution of binocularity

Granrud (1986) studied the relation of binocular vision and spatial perception in four to five month-old infants and found more consistent reaching for the binocular viewing conditions than for the monocular one, infants in the binocular condition tried reaching more often. Moreover, in another study binocular vergence was found to be contributor in performance

and even be advantageous parts of the task by Bradshaw and Elliott (2003). In that study participants were asked to reach and grab an object. When they started the movement, binocular vision was artificially inhibited by means of goggles with liquid crystal lenses. At some point during the translation movement, binocularity was enabled (the darkened lens of the goggles became clear). Kinematic analysis presented no differences for the transport component of the movement but final phase elements like grip aperture were affected in good measure. They concluded that, in a prehension task, binocular vision contributes mostly to the non-ballistic phase of the movement, where feedback is needed for fine adjustment. In addition, in another study Servos *et al.* (1992) found that binocular vision contributes to visuomotor control in the premovement phase of the task and that performance is better with binocular vision. Servos *et al.* (1992) argued that the difference in performance between the two viewing conditions is due to differences in estimates of the target's size and distance calculated prior to the movement. Melmoth and Grant (2006); Melmoth *et al.* (2007) also found advantages, and in line with Bradshaw and Elliott's 2003, reported an extension during the end phase of the reaching movement. He also reported that premovements grip aperture inaccurately matched the target size and therefore argued that binocular performance was better. He concluded that binocular disparity provides depth information used for grip aperture calculation but that vergence is used for calculating distance for grasp. Even though we did not measure kinematic parameters in our experiment, we can see in figure 4.7 a quantitative similar pattern in the lower accuracy of reaching in the robot for the monocular condition. Moreover, we used a

simplified robotic model of these two visual conditions and our robot also had access to target size information from the retinal projections. This information was incorporated to the general output of the system by the layer *hB* and still the results were similar although the network was capable of using segmented-image information for calculating size of the target (a feature in the input data that hidden layers can easily recognise and exploit).

4.3.2 Contributions of proprioception in vision

We also investigated the use of vergence angle as a proprioceptive signal. Vergence has been found to provide information of perceived relative depth of objects and is especially useful in short distances (Lie 1965; von Hofsten 1976) and thus in the peripersonal space. The robot was tested in monocular and binocular conditions with the output from the neural controller fed with contributions from two different sensory modalities: proprioception and vision. For the monocular case, it was expected that the visual modality, that is the connections providing input to the neural controller with retinal data, would compensate for the lack of a vergence signal. It did not happen like that or if it did, not to a degree in which the compensation was enough to equal or surpass the good performance achieved by the binocular case, where the vergence signal was present. A possibility is that the characteristics of the robot potentiated the use of a vergence signal, as we can assess by the 13% decrease of depth estimation error seen in figure 4.7. Vergence, therefore, resulted in a natural representation of depth and seems closely related to the physical properties of the iCub's body. In this case the proprioceptive modality signal from vergence resulted much more useful than the visual one.

Our experiment indicated the effectiveness of the use of vergence for depth estimation in a reaching task in a simple active vision system implemented on the iCub simulator. These results indicated to us the suitability of these depth estimation system for our further development of a peripersonal space representation. In this case, results suggested proprioceptive information (from, vergence thus binocular vision) was a much stronger cue than the visual one, similarly to Marotta *et al.* (1997), on the conditions for the experiment and it was present in both test scenarios.

The experimental results show that proprioception made an important contribution for the reaching performance. This pointed out the importance of vergence for obtaining more information of the environment on an anthropomorphic robotic system like the one we used. Relating to the tracking controller used for our experiment, Gibaldi *et al.* carried out experiments on vergence eye movements in the iCub robot head using a neuromorphic control module for visual stimulus within peripersonal space. Their controller is based on full image, intensive disparity process, we on the other hand, decided to use lighter image processing, therefore designed our own controller (section 4.1.1). Our algorithm also worked in real-time and was stable by damping oscillations (*modulate* function in code 4.1). Our controller did not use the binocular energy model and we aimed at visual segmenting instead of pixel disparity, but nevertheless resulted suitable for our purposes in the simulator. In addition, in the present study we started exploring the role of vergence during environmental interaction.

4.4 Conclusion

It has been shown in several studies (Han and Lennerstrand 1995, 1998; Lennerstrand *et al.* 1996; Maxwell and Schor 1996) that neck posture and stimulation of neck muscles (proprioception) have an influence in eye movement visual tasks. We have shown how the signal coming from the proprioceptive modality, although not independently analysed, seems to be an important component of the output in our robotic system, meaning that the positions used for learning helped creating a motor memory. Here we presented an exploration of multimodality in a monocular/binocular visual task. We consider this finding a step to account for the role of embodiment in the building up of cognitive processes, in particular that of understanding or implicitly representing the space around the body in artificial systems enabled with visual vergence.

5 Experiment 2: Posture and Arm-Modification Contributions to Adaptive Reachability Assessment

The near-space or peripersonal space is the region surrounding the body where reaching is possible without translation. Being able to tell what is immediately reachable—in the peripersonal space—and what is out of reach is an action-based skill that humans develop in the sensorimotor period in early infancy (Piaget 1952). It has been reported that four month-old infants systematically do not attempt to reach objects out of reaching range (Granrud 1986). The ability to recognise the reachable space has been demonstrated to be associated to motor representations (Coello *et al.* 2008) and needs neural representations of the body and the space around it (Holmes and Spence 2004). Investigating how these representations are created and modified, has been the object of recent research in psychology (Costantini *et al.* 2011), neuroscience (Berti and Frassinetti 2000; Ursino *et al.* 2007) and robotics (Chinellato *et al.* 2011; Goerick *et al.* 2005).

For robots to operate in unstructured environments, they will need to create and dynamically adapt their space constraints in order to effectively interact with objects, other robots and/or humans. For important safety reasons, this kind of adaptation will be an essential requirement as robots

become more present in human environments, offices or hospitals, for instance. Some research has been done on the issue of body representation, peripersonal space and the relation of these with action. For instance, Sturm *et al.* (2009), implemented a robotic manipulator system that uses Bayesian networks for learning and adapting an internal kinematic model according to body changes. Chinellato *et al.* (2011) endow a simulated agent with a radial basis function system that allows it to create a sensorimotor map after interaction with the environment. In another related study, (Hersch *et al.* 2008) presented a model for body schema of a simulated humanoid robot that uses a hierarchy of reference frames transformations that adapt to visual shifts of the end effector.

We have mentioned studies where natural and robotic systems develop representations of their own bodies and/or the space around them and have also pointed out the importance of the physical properties of the body in order to do so. In addition to the geometrical characteristics of the system the posture it displays during an action also conveys information that is associated to the task in question. As discussed in section 2.4.2, it has been argued that vision does not only consist in acquiring images from the environment, that analysing static images is far from being the totality of what a system needs for operating the environments. This is recognised by recent approaches to perception which suggest action and perception work together (O'Regan and Noë 2001; Pfeifer and Bongard 2007). In the process of seeing, action and perception are closely related. Acting involves changing body posture, having kinematic (and also dynamic) experiences thus, in the process of perceiving and eventually being able to predict, the body plays an important role. During

all our experiences the body holds a certain posture or set of postures related to the activity we are performing. Consider how Strack *et al.* (1988) found that participants who are induced a certain facial expression similar to a smile evaluated cartoons as funnier in comparison to participant with no induced smile. The role of postural information in cognitive tasks in robots has been recently investigated by (Morse *et al.* 2010b). They propose an architecture for cognitive robotics which uses a self-organising map with that encodes postural information as a main “hub” for associating more information coming from other sensory modalities. Their architecture has been successful in replicating psychological studies on categorisation (Morse *et al.* 2010a) and suggests the important role the body can convey for developing humanoid cognitive robotics. In the present study we explored the use of body postural information when assessing reachability of a target. We were interested in the contribution of arm posture in the categorisation of reachable and unreachable targets.

In the present experiment we investigated how the iCub can develop a representation of its reachable space by using postural and body-shape information—originated from vision and action—as components in the learning process. Proprioception of the body-shape characteristics, in this study, considered the length of the forearm (for convenience, we will refer to the forearm as the arm from now on). The arm extension is also like artificially extending the reaching range with a tool. After training, the robot was expected to be able to tell when something was reachable or not after foveating at it, without requiring to perform the reaching movement, that is, assess its own perceived reaching range, which was modulated by the

arm-length proprioceptive information. Furthermore, in this experiment we were interested in verifying if the robot perceived reachable space differently when arm postural information needed for reaching was involved in the learning process, or if neck and eyes posture provide with enough information for assessing the reachability of a target. We expected the system to display an effect and perceive reachable space either as shorter or extended as a result of access to supplementary bodily information during training. In other words, we were interested in investigating if more body-posture information has an effect in the perceived reachable space as more motor-action becomes involved for the creation of this information. This consideration was made under the light of findings that show the necessity of action in the development of perception of the reachable space (Coello *et al.* 2008), and of studies in robotics suggesting gazing and reaching can contribute to create a representation of it in artificial agents (Chinellato *et al.* 2011). We intended to verify if our robotic system would present such behaviour. The interest also originates from the first experiment presented in this work (chapter 4). There we found that proprioceptive signal from the eyes and neck was effective for a reaching task, and when more body information was available (vergence) the robot displayed smaller reaching distance error. Here we investigated if a learnt reaching signal is affected by including or not the arm-postures used for reaching into the learning process.

Understanding how robots can create representations of the space around them and their relation with it in terms of what they can reach or not, will allow the creation of safe robots for human environments. This knowledge could also be a step in the process of providing them with

capabilities for understanding the reach of others, a possibility suggested by studies on shared representation of near space on humans looking at mirror neuron activation (Brozzoli *et al.* 2013) and the fact that allocentric perception starts developing in children as young as three years old. This allocentric perception is a skill needed for cooperation with others and is still a topic of which little is known about (Fischer 2003). By studying multimodal representations of space and, this study is a step in the development of such systems which also contributes to existing literature on the use of binocular vision for peripersonal space in robotics. A novel contribution in this research is the inclusion of confidence levels in reaching assessment.

5.1 Methods

For our experiment we used the iCub robot simulator (Tikhanoﬀ *et al.* 2008), also described in section 3.1. We provided the iCub with a neural network for performing a reaching task with the arm. The robotic controller was trained with self-generated data. In the process of training data generation, the shape of the body, namely the arm, was modified in its length. Three lengths in total were used in the experiment: two during training which are shown in figure 5.1, corresponding to the short (0.137 m) and long arm (0.277 m). Then, a third one, with a length between the short and long ones (0.207 m), was also used during the evaluation phase of the experiment.

5.1.1 Visuomotor system

The robot was provided with a gazing behaviour by means of the head/eyes active vision controller described in section 4.1.1, with visual input directly influencing movement of the head and originating the proprioceptive

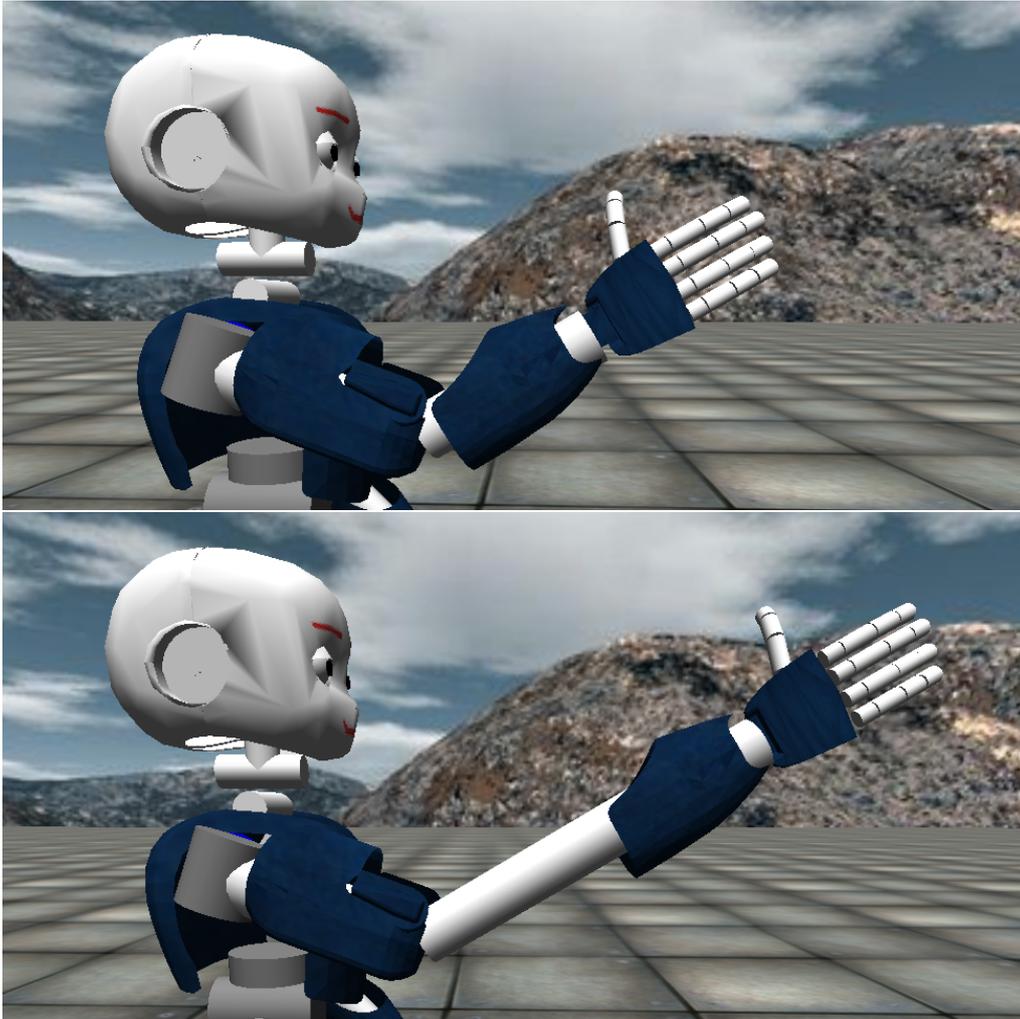


Figure 5.1: The iCub's arm with two end effector lengths used in the experience phase.

data. Images were colour-segmented and used to modulate the velocity of two joints of the neck and the eyes' vergence angle (the controller operation is described in code 4.1). Movement leading to foveation with both eyes (to the centre of the red target in the retinal image) is produced by this controller when a target is presented inside the iCub's field of view. The visuomotor system used does not need prior camera calibration nor disparity calculation for producing correct vergence movements. Still, it was a reliable way for depth estimation easily implemented in the simulated robot and easy to transfer to the physical robot. As we have mentioned earlier in this thesis (section 2.4.3

and also in chapter 4), we considered vergence as an early-stage estimator for depth based on existing literature demonstrating that correct vergence response to static targets is present in babies as early as they are one or two-months old (Hainline and Riddell 1995). The visuomotor system generated neck and eyes postures used in the training and evaluation phases of the experiment.

5.1.2 Neural controller

For implementing our model, a neural network was designed for accomplishing two tasks. In the first place, it was used by the robot as motor-controller to generate correct arm postures for touching a foveated target. The second task was assessing perceived reachable space by predicting if a target was reachable or not. The controller was a feed-forward neural-network with three layers. The first layer was an input layer that received proprioceptive input from the head and eyes' angle positions with one unit for each of these values and an extra input unit that received proprioceptive information about the end effector length. This last unit was activated to one when the arm was the longest, zero when it was the shortest (top and bottom in figure 5.1, respectively). The hidden layer h consisted of ten units. Finally, the output layer had four units for control of arm joints' angles (two in the shoulder and two in the elbow) plus two more for indicating the predicted reachability. These last two units use an encoding very similar to a one-hot¹ encoding. It differs from one-hot encoding in that the network uses continuous activation functions instead of discrete functions,

¹*One-hot* is a type of encoding used in neural networks. It uses a vector with zeros in all its elements except for one element. Each element of the vector corresponds to a different category. The non-zero element indicates the category of the data input in a winner-take-all manner

therefore the values provided can still be real values. However, for simplicity in this experiment the output encodings of these neurons will be referred as being one-hot. Hidden and output units used a sigmoid activation function. A diagram of the controller is shown in figure 5.2.

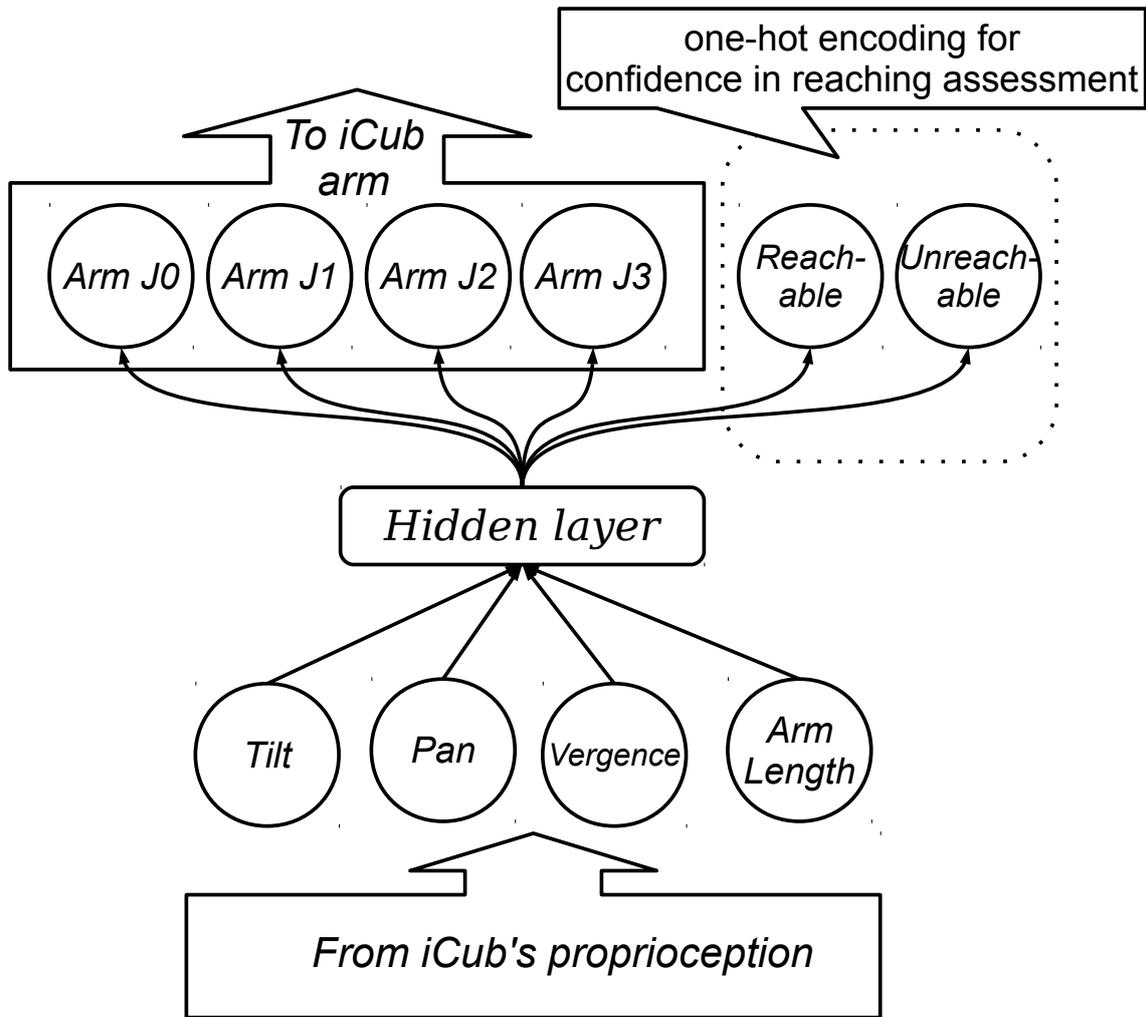


Figure 5.2: The neural controller architecture. Two output units use one-hot (winner-take-all) encoding which indicate whether the target is considered in reach or not.

5.1.3 Learning process

The collection of data was designed to allow the robot acquire motor experience. The robot performed random movements after foveating a target set in many different locations. Motor babbling (Meltzoff and Moore 1997)

allowed the robot to explore the space around implicitly considering the physical constraints imposed by its own body. When self collision was detected by the simulator, the trial was discarded and a new one executed. During this stage, if the robot touched the target, the current posture in that moment was added to a set of training data as input, along with the corresponding head/eyes angles, arm-length signal. The associated output in that case would be a “reachable” signal equal to one (analogous to a touch signal when reaching something) and arm posture used for reaching the target. Otherwise, if the robot could not touch the target within a certain lapse of time (thirty minutes, which was equivalent to twenty trials approximately, but this number varied due to the realistic mass and weight properties simulated), the robot went to a “neutral” arm posture (arm extended parallel to the torso) and the touch signal was set to zero before being appended to the training set. In other words, if the robot could not touch the target, the arm adopted a neutral position and the touch signal was set to zero, however, neck and eyes’ angle values for foveating still would be correct as well as arm-length signal. This cases when the robot could not touch a target also went into the training set in order to provide experience of unreachable space. However, the position stored in the set was the “neutral” one explained above.

For comparing the effect that the postural information from the experience phase had in the perceived reachable space, the network was trained in two conditions. Twelve robots were trained and tested in each condition. In the training phase each of them had access to different proprioceptive information. The two conditions were:

- **Condition A:** During training, the network was taught arm postural information needed for reaching.
- **Condition B:** The network received a “neutral” arm posture as part of the output every time. As if the arm had been disabled during the experience phase.

Additionally, table 5.1 summarises the data contained in the training sets used for the arm controller.

Table 5.1: Data in the training sets for each of the two conditions.

Condition	A	B
Arm postural	✓ arm posture for reaching	✗ (arm was in “neutral” position)
Arm Length	✓	✓
Head/Eyes posture	✓	✓
Touch signal	✓	✓

All training data was normalised in the range zero to one. The training algorithm used was backpropagation of error with learning rate $lr = 0.01$ and momentum $M = 0.0$, final error was 0.0064 and 0.0031 for condition A and B respectively. Figure 5.3 displays the training error for one of the networks

in each condition. For both conditions, training error of the twelve networks had a very similar profile, therefore in the graph only one of them is shown as a typical training profile for the condition. Training tests were performed with five, ten and fifteen units with no difference in the training process and it was decided to use ten units in this controller.

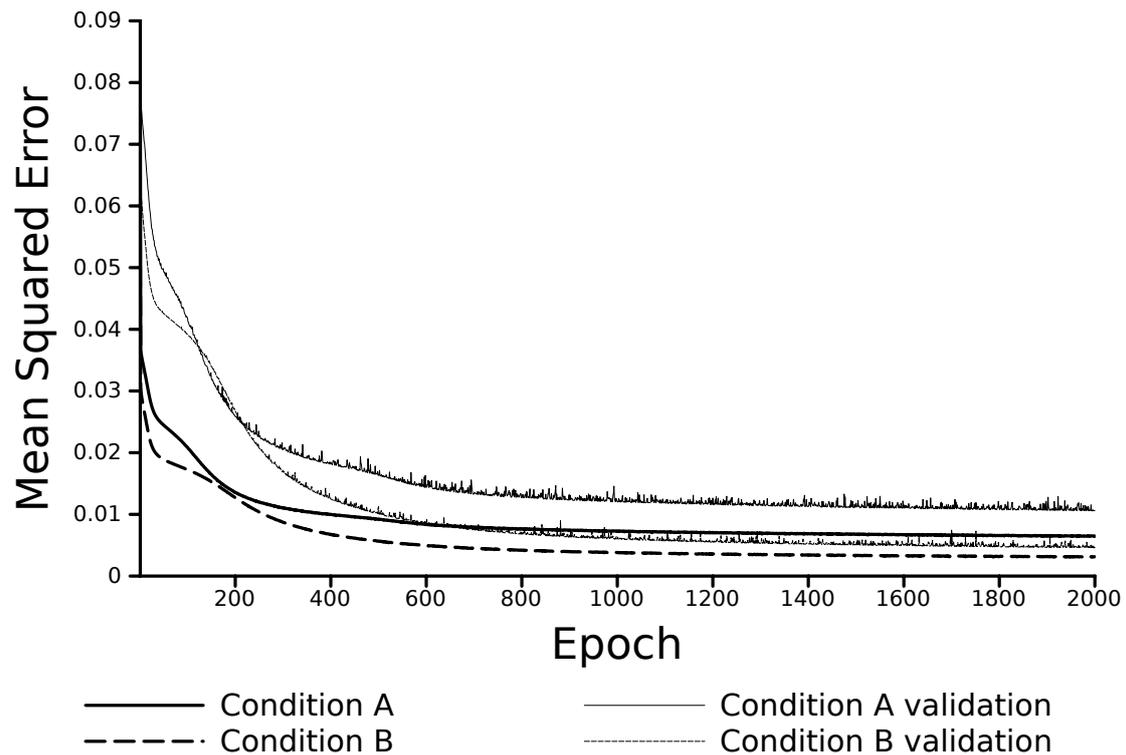


Figure 5.3: Error during training and corresponding validation error for both conditions. Only one of the twelve networks for each condition is shown as the other eleven displayed very similar training profiles.

5.2 Results

In general, evaluation was done by presenting the robot with targets in different locations. After foveating the target, the robot used the network to determine if the target was reachable or not. Reachability prediction, present in a one-hot encoding, was given by the network using two output units, indicating reachable or unreachable. With this encoding, confidence level in reaching assessment was measured as the absolute difference of these two outputs. In the following we will detail each of the measurements and explain the considerations taken for collecting results.

For measuring the perceived reachable distance in operational space coordinates, the robot was presented with uniformly placed single targets in front of it, in volume delimited by $[-0.75m, 0.75m]$ on the X axis (left, right), $[0.4m, 1m]$ on Y (height) and $[0.1m, 0.8m]$ on Z (depth), with a resolution of ten points on each dimension, that is, ten points were evenly located along each axis. Each time, the robot had to foveate to the target and afterwards assess if it could be reached or not. This was done by activating the network with the head and eyes postural information along with the arm-length signal and recording the outputs of the network that indicated perceived reachability. These were the outputs that used a one-hot encoding shown in figure 5.2.

For visualising confidence levels in reaching assessment in the space in front of the robot, the robot was presented with 600 points individually evenly distributed in the space in front of it in six planes parallel to the XZ that it was able to foveate with both eyes. It assessed each point individually one after the other. For figure 5.4 operational space was discretised in squares of 16×16 cm along the horizontal plane and the mean confidence level obtained for the

targets across that column was calculated. Figure 5.8 shows the confidence in the resulting perceived reachable space for one network in each condition for the three arm lengths. The measurement of confidence was the difference of the values of the two outputs that categorised a point as reachable or not. Absolute values of this differences were recorded as the sign of the difference was of no interest due to that it was always positive when the robot decided the target was reachable and always negative when it was not (the difference was the value of output unit that indicated reachable, minus the value of the output unit that indicated unreachable). This absolute value only indicated how sure the robot was of the categorisation but did not indicated if the robot was overestimating its reach or underestimating it. For areas of space which were completely outside the work space but that the robot was able to see (too far from the robot or so close to the body that the end effector could not reach due to mechanical constraints) the robot, never having reached anything in those areas, categorised them as not reachable with full confidence values. Those areas are shown from a top view in figure 5.8 as the lightest blue squares.

In addition, standalone network tests were carried out. In these tests, the network was activated independently from the robot visuomotor signal. In this manner it was not needed to wait until the robot performed the search and gaze head/eyes movements. The network was fed with a range of input values corresponding to the head and eyes' joints limits. The reaching assessment output units' activation was measured for a large number of inputs and plotted for the two conditions and the three arm lengths in figure 5.6. In the figure, the maximum reaching range shown with coloured spheres is congruent with the actual reaching range the end effector displays product of

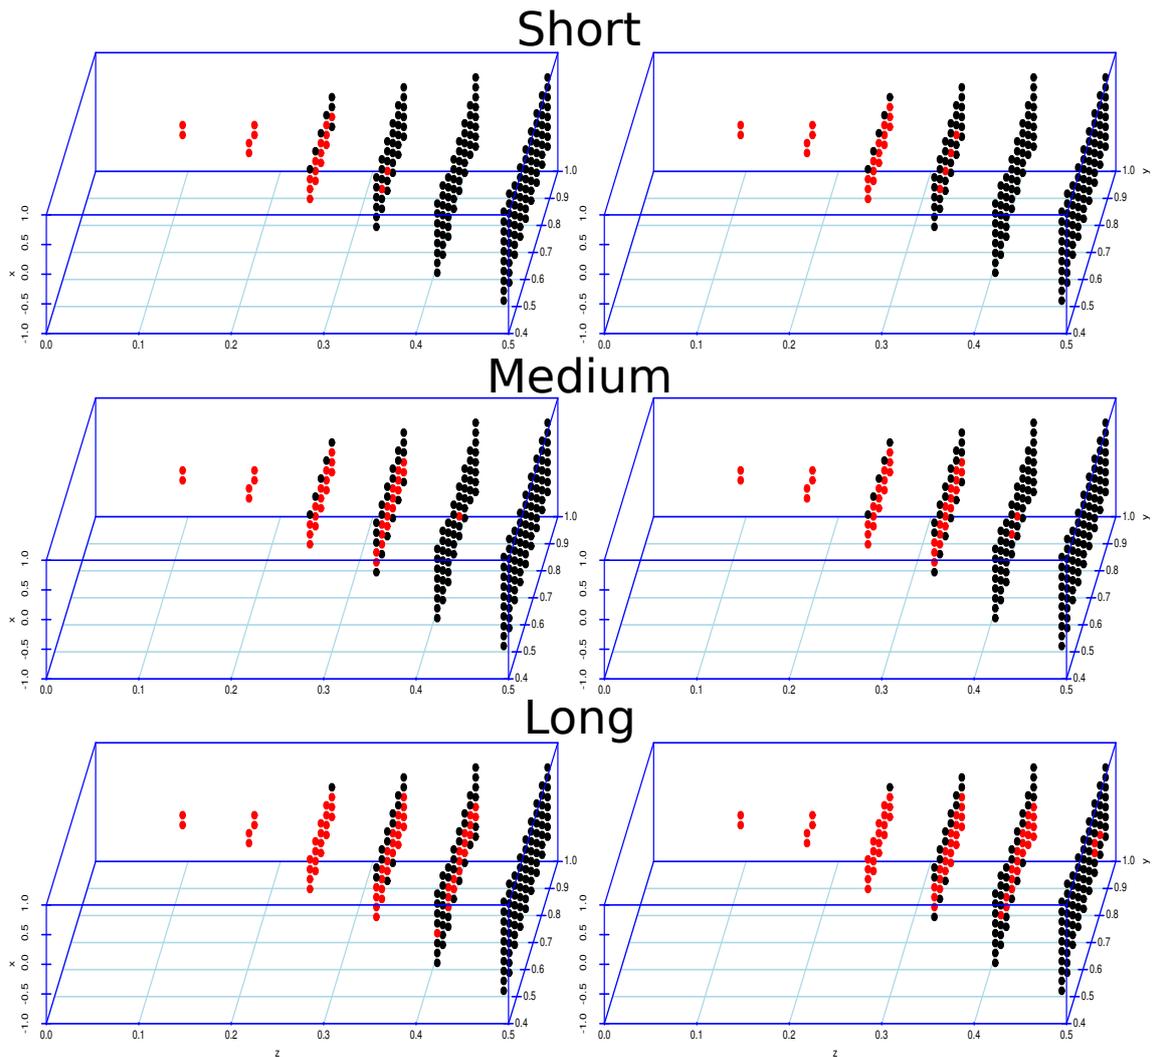


Figure 5.4: A view of the robot's peripersonal space with three different effector lengths (short, medium, long from top to bottom) in the two conditions. On the left the robot trained in condition A (arm and vision) and on the right condition B (vision only). Red points indicate the robot assessed the point as reachable. Axes XZ define the horizontal plane. Units are metres.

its kinematic properties. Because for this measurements the network was activated in a standalone mode, for obtaining the distance at which the hypothetical target would be located, the inverse of the fitting model detailed in equation (4.2) was provided with the vergence angle fed to the network. The mentioned inverse of the fitting model is shown in equation (5.1). This provided the distances on the Z axis (depth) shown in the visualisations of the reachable space for each of the lengths and conditions of the experiment.

Using the same model to determine the limits of the reaching space, the network was activated with fixed tilt and pan values so that only vergence was modulated. In this manner a profile of reachable distance was obtained. The profile for one of the robots of condition A is shown in figure 5.5, subfigure a). Using the same method, maximum perceived reaching distance was measured for all twelve networks in each condition.

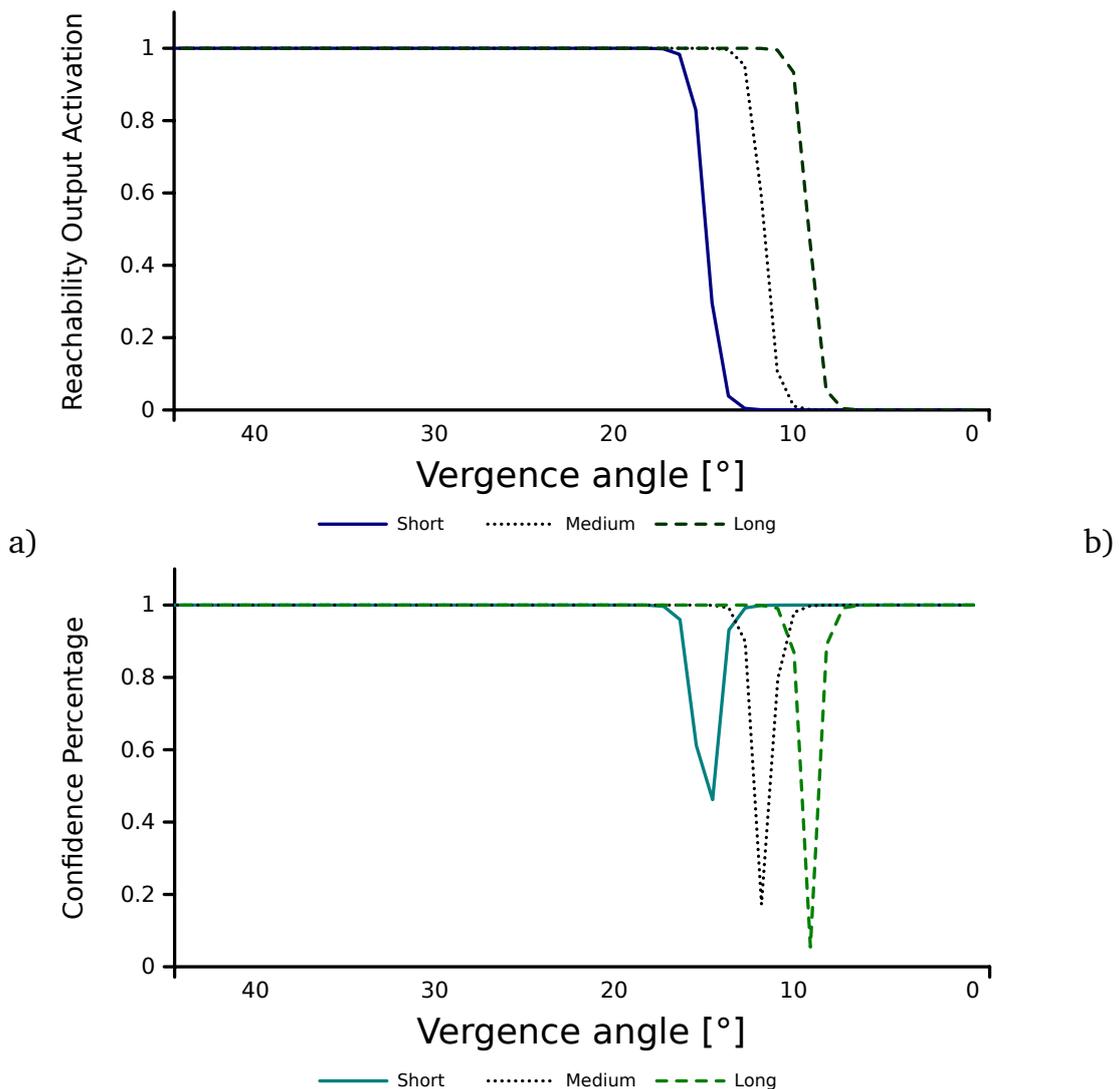


Figure 5.5: a) Vergence angles at which the robot considered the target became unreachable for the three different arm lengths. b) Confidence-decrease peaks were present just before the reachable limits. Confidence in reaching assessment was measured as the absolute difference of the one-hot output units used to indicate reachability. Confidence-decrement peaks became stronger as the length of the arm increased.

$$y = (1.00363 + 0.0573347 \times e^{(0.105105x)}) \times e^{(-0.105105x)} \quad (5.1)$$

As it has been mentioned, confidence displayed by the robot when assessing reachability was also recorded. Confidence level in reaching assessment was measured as the absolute difference of the two output neurons used to indicate reachability. Because one-hot encoding was used, ideally, when a target was reachable the output of the one used for indicating reachability would be one and the other one zero. However, confidence decreased for certain target positions/vergence angles. This is shown in figure 5.5, subfigure b) for one of the networks of condition A. This profile was obtained fixing tilt and pan values and modulating vergence. The profile presented a decrease of confidence as it approached to the limits of the reachable distance. Figure 5.8 also displays this confidence modulation by presenting a top-view of the operational space and the mean value of confidence for foveated points along columns perpendicular to the Y axis.

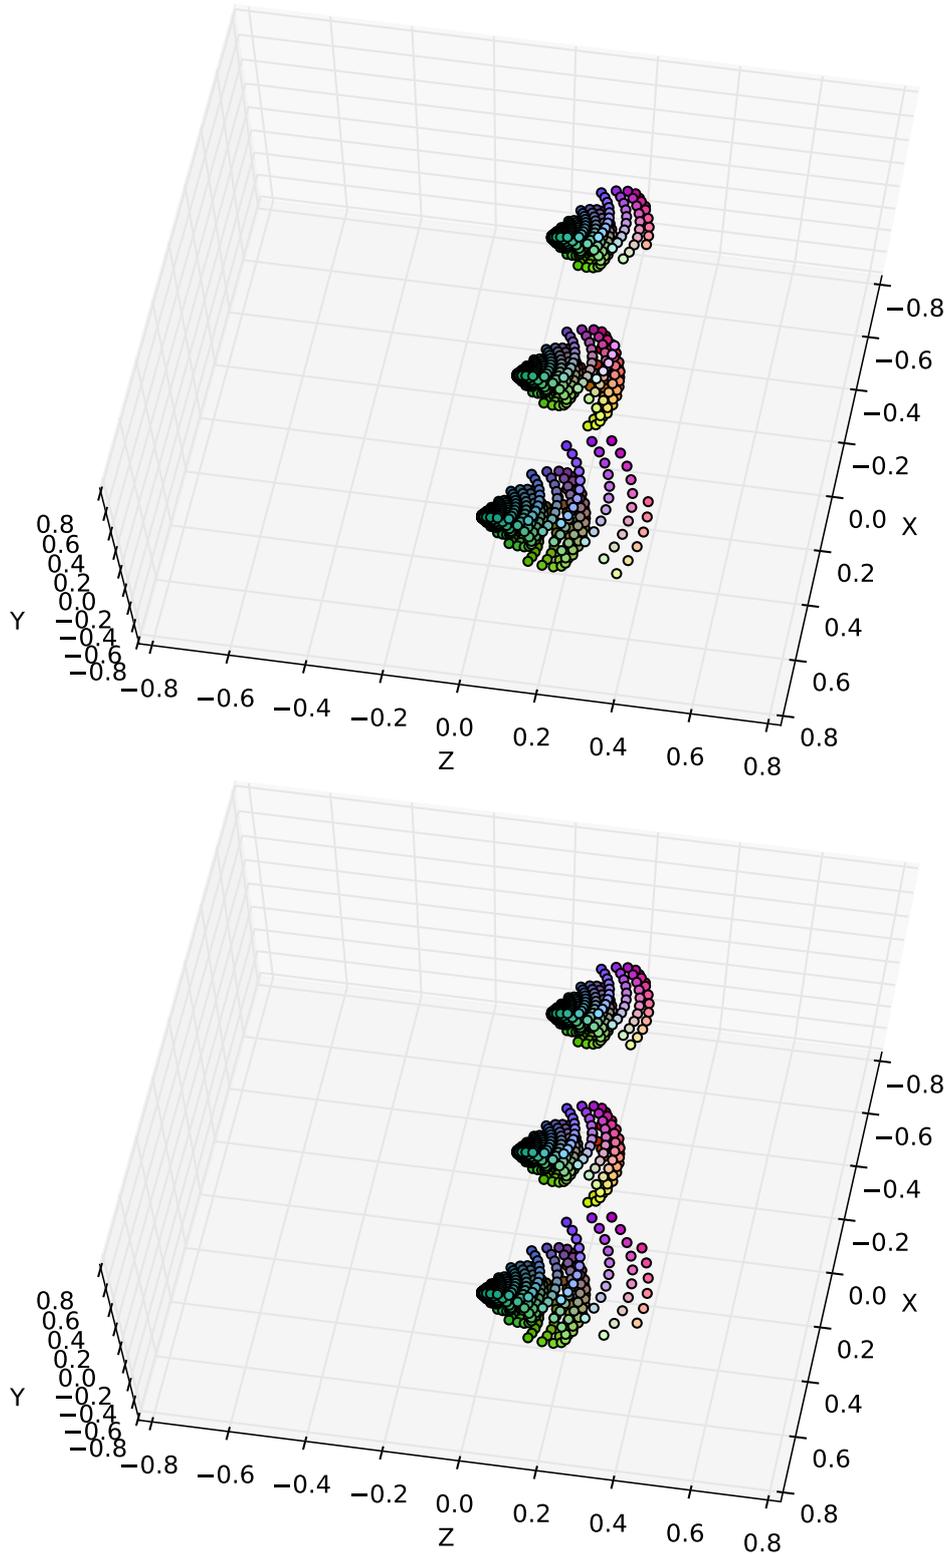


Figure 5.6: View of the peripersonal space of one robot for each condition with the 3 arm lengths (short, medium, long from top to bottom in each subfigure). Condition A (arm posture experience) above, B (fix arm) at the bottom. Axes XZ define the horizontal plane. Units are simulator metres. Colours are used for better displaying 3D locations.

5.2.1 Analysis

The twelve controllers of each condition were tested and measurements of maximum perceived reachable space for the three different arm lengths was recorded. This measure was obtained by isolating the neural network controller and feeding it with simulated proprioceptive values that the head would have by placing an object on a line in front of the robot parallel to the Z axis and shifting it along that line. When fed with this values the network would still give an estimation of reachability for an object in the corresponding position. Descriptive statistics for these measurements are shown in figure 5.7 and summarised in table 5.2. As it was mentioned, confidence was measured as the absolute difference of the one-hot reachable reachability indicator output units. This confidence values were measured for points in three-dimensional space and then plotted in two-dimensional space from a top-view in figure 5.8, which presents the mean confidence in reaching assessment recorded for the points in each column of foveated points.

Table 5.2: Measurements of maximum perceived reachable distance.

Condition	Arm length	Max. reaching dist. mean	Std. Dev.
A	Short	0.2785	0.0093
B	Short	0.2728	0.0136
A	Medium	0.3587	0.0090
B	Medium	0.3554	0.0149
A	Long	0.4443	0.0053
B	Long	0.4451	0.0058

One-way analysis of variance were carried out (using the R language) for comparing the two conditions (robots having used arm posture in the training phase and robots that held the arm in the “neutral” position) finding

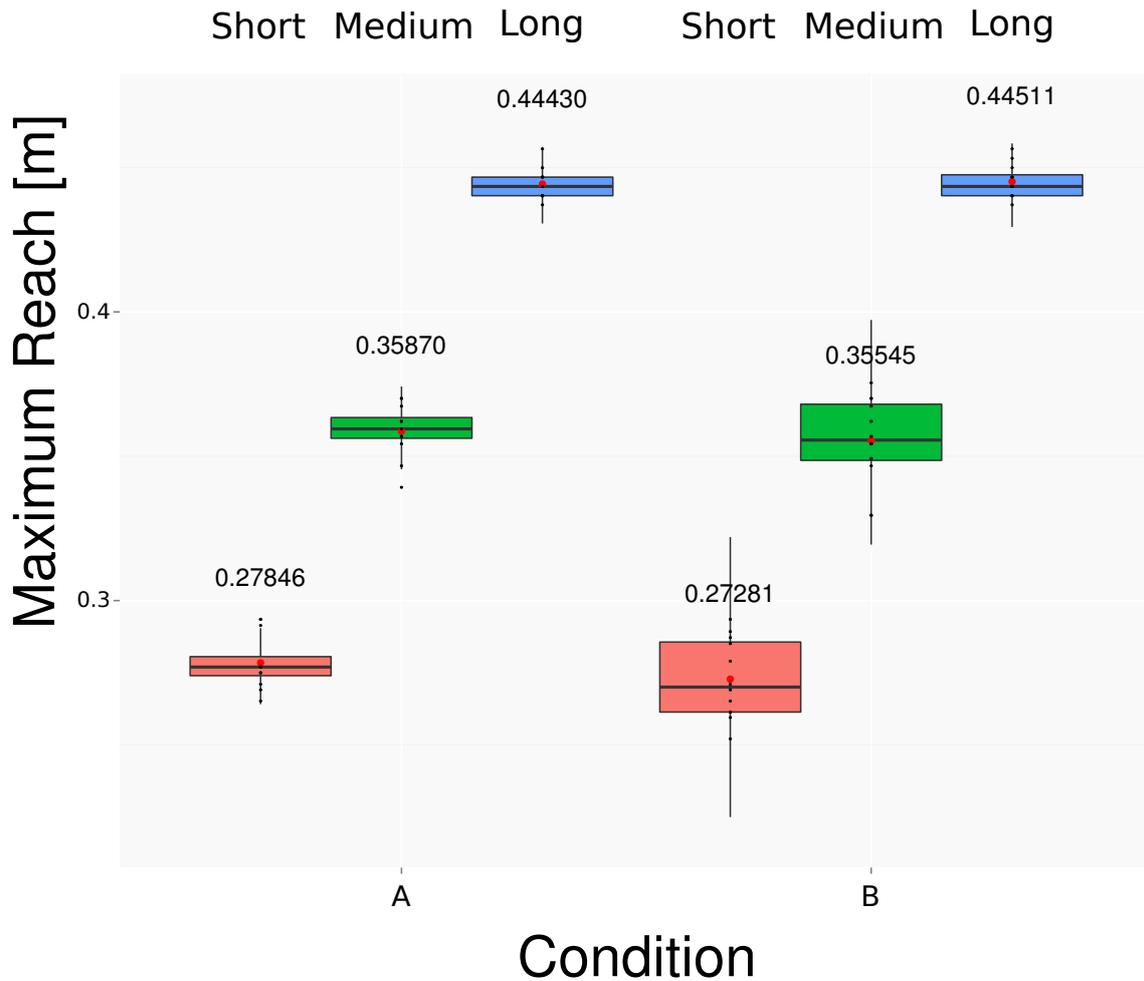


Figure 5.7: Descriptive statistics of the perceived maximum reachable distance for the two conditions and three different arm lengths. Means are indicated by larger dots.

not statistical difference between them. Table 5.3 shows the result of the test.

Table 5.3: Results of one-way analysis of variance for comparing the three arm lengths in both conditions found no statistical differences between them. Condition A included arm posture-for-reaching information and posture of head and eyes in the training. Condition B only used head and eyes postural information and a “neutral” arm posture.

Arm length:	Short	Medium	Long
A vs B:	$F(1, 22) = 1.400$ $p = 0.249$	$F(1, 22) = 0.414$ $p = 0.526$	$F(1, 22) = 0.126$ $p = 0.725$

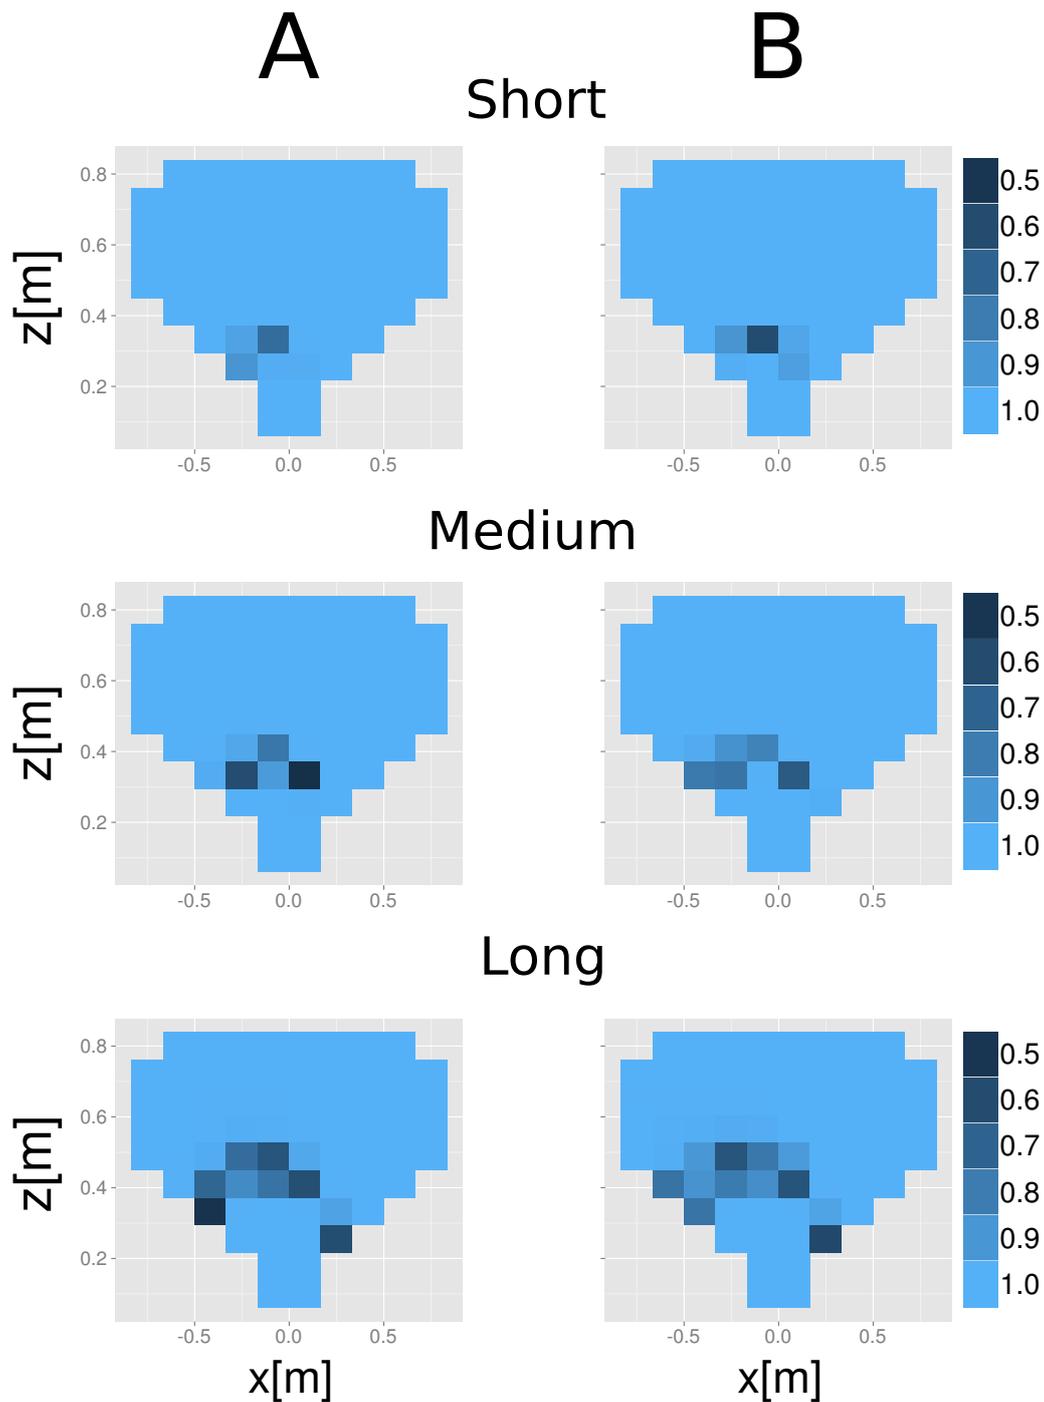


Figure 5.8: Top-view of operational space showing levels of confidence in reaching assessment for both conditions. Blue squares are regions the robot was able to foveate, grey where it could not. Objects were located in different heights and the average confidence in reachability assessment of targets in each discretised column is shown. Confidence was measured as the absolute difference between reachable/unreachable output neurons. Darker regions indicate lower confidence. Confidence in reaching assessment decreased as a function of arm-length.

5.3 Discussion

The experiment presented suggests that explicit representations of the space around the body are not necessary. Instead, experiencing action within reachable space allows or constraints the robot for acting on certain elements of the world that come into attention by means of vision and touch. Therefore, the model suggests it is the sensorimotor interactions, or action-oriented representations, what are used for monitoring or assessing peripersonal space. In this experiment, having had the experience of interacting with unreachable objects (by foveating at them, attempting to reach and then going to the neutral arm posture in the experience phase) in addition to the reachable ones, contributed in learning not to perform the reaching movement when assessing reachability, a behaviour present in three to four-month old children (Granrud 1986). Hoffmann *et al.* (2010) define two categories for synthetic studies for understanding multimodal body representations. According to their categories, the model presented in this work corresponds to the *action-oriented* category, as the generated representation is used for controlling the robot's behaviour. However, it could be argued that condition B of the experiment would partly correspond to the *nonaction-oriented* category, where individual modalities are associated independently of any specific robotic action or task (with the use of Hebbian learning, for example, although we used another learning technique).

5.3.1 Postural contributions during the task

It was expected that having access to the postural information necessary for reaching would have an effect in the perceived peripersonal space, arm posture working as a kind of extra memory for spatial categorisation for

reachable and the unreachable. It was expected the use of the arm for reaching would provide a more “embodied” experience, and therefore play a bigger role in the creation of peripersonal space in this type of robotic system. However, analysis of the assessed reachable distance for the robots in the two conditions presented no statistical difference for none of the three arm-lengths. P values were $p = 0.249$, $p = 0.526$, $p = 0.725$ for short, medium and long arm, respectively. The network was potentially able to identify relevant patterns of head and eyes’ angles values that indicate whether a point in space can be reached. It was also noted that the training errors were fairly similar, although for condition A output data presented a higher variance due to condition B supervised output indicated the “neutral” arm posture every time. Again, the network might have extracted necessary features only and discarded part of the redundant information, potentially arm postures in this case. Possibly, another factor limiting the contribution of postural information is the type of network used in our models. A feed-forward network was used for a one-shot activation, in this manner, the network is working as a memory that associates the position of a point in space—encoded in the head/eyes’ joints—to a yes/no response, and the participation of the body therefore limited to the arm length input. More investigation on other type of network controllers has to be done in this respect.

5.3.2 Role of neural network used

Research on related topics currently explores other network types for the representation of space and/or the body schema. Roschin *et al.* (2011) studied a method for developing a body-schema for a seven

degrees-of-freedom simulated 3D agent using a self-organising neural network with a number of receptive fields, their system produced the same representation for its end effector in operational space by using joints' proprioception and also by using a tactile stimulus in the same spot of the body it was touching. Although the work of Roschin *et al.* (2011) contrasts with ours in that we investigate humanoid robots, their model is an interesting one that should be studied in humanoid robots with a high number of degrees of freedom. Another neural network approach to peripersonal space is the work by Magosso *et al.* (2010). Strongly based on neurophysiological findings, they make an extensive study of suggested underlying neural circuitry involved in peripersonal space phenomena. Their model consists of uni and bimodal maps that correspond to receptive fields of visual and tactile modalities and replicates neuronal activity in cross-modal extinction/facilitation in brain-damaged patients. The approach Magosso *et al.* used sheds light on aspects related to the contribution of different modalities for peripersonal space representation. The plastic properties of peripersonal space (Iriki *et al.* 1996), however, were not explored by Magosso *et al.*, contrasting with our present study. The work of Magosso *et al.* (2010) also contrasts with ours in that we presented an embodied model, which uses perceptual short-cuts. In the present work we were interested in not skipping any of the perceptual steps when implementing our model. Although with a simplified system, we addressed vision and we also took into account the geometrical limitations the robot can face by using a simulated environment with realistic computer-simulated physics. Still, it would be very interesting to put that model into an embodied robot and work adding learning properties

to further investigate the potential of their system.

Regarding the way in which the neural network was used, making the arm adopt a neutral pose was used as a way to force the mapping of points in the input space that were unreachable to a general output (the neutral pose for all of them) and to a category labelled by the one-hot units. The labelling is not only encoded in the one-hot units but also in the pose that is sent as a motor command by the units connected to the arm controller. An expected effect of this was that the arm would be farther away from the neutral position in cases when the robot was less confident about its assessment of reachability. However, although the one-hot neurons showed variations indicating less confidence, the expected arm effect was not observed in neither of the two training conditions. This could mean that the arm neutral position was a very strong attractor when classifying the unreachable locations. The inspiration for following approach came from the hypothesis of *encoding specificity* (Thomson and Tulving 1970): memory retrieval is better when the conditions under which a memory is retrieved are similar to the conditions under which the item was originally encoded. Distributed encoding of memory and embodiment provide a framework for trying to understand this effect. For example, Morse *et al.* (2010a) modelled spatial biases in categorisation in the iCub implementing an unsupervised learning algorithm which encoded a classifier mechanism in a distributed and multi-modal manner along several self-organising maps, one of them corresponding to body posture. In this case the result did not allow to collect any behavioural output apart from when the robot decided something was reachable and made the movement but it provided the robot with a very good classifier that

correctly labelled the inputs as reachable and unreachable.

5.3.3 Confidence in reaching assessment

Confidence levels in reaching assessment was a novel element in the present work. Confidence (or certainty) in reaching assessment in robotics is a topic that has not been previously studied. In figure 5.8 it can be observed that a gradient of confidence is present indicating that the robot is less confident when vergence values get close to those corresponding to the reachable space limits. There is a peak in the decrease of confidence close to those limits, as it is shown also in figure 5.5. This is in accordance to studies indicating that vergence signal provides good depth estimation especially for closer distances (Lie 1965; von Hofsten 1976). In the experiment, the confidence in the reaching assessment decreased following a depth-gradient. As peripersonal space was extended, the confidence in reaching decreased. Reaching confidence levels are shown in figure 5.8, and figure 5.4 shows the extension of peripersonal space for both conditions as dependant on the arm length. Points in space which the robot classified as reachable are shown in red for each of the three arm lengths. The confidence decrease most likely was originated in depth estimation errors from the vergence signal.

5.3.4 Relation to tool-use and the plasticity of peripersonal space

The extension of the arm present in the experiment is to some extent similar to putting a tool at the end of its hand. There are multiple studies that explore how tool use also modulates peripersonal space. Iriki *et al.* (1996) made single-cell activity studies in macaque monkeys and found that bimodal neurons encoding hand schema presented a modulation after the monkey having used a rake to extend its reaching distance. Berti and Frassinetti

(2000) investigated this phenomenon patient who presents a dissociation of far a near space and displays neglect only in the near space, after using a tool presented neglect with objects previously in his far space. The system proposed in the present experiment was able to generalise for varying lengths of the robot arm, suggesting it to be useful also for tool-aided reaching. Moreover, Holmes *et al.* (2004) found that when holding the tool, the effect it has in peripersonal space extension decreased with the length of the tool. This is in line with what we found regarding confidence levels for reaching, as it can be seen in figure 5.8, where the confidence decreases proportionally to the length of the arm. There, darker squares indicate less confidence and it is evident that the reachable space becomes lighter as it get farther from the coordinates origin.

Regarding the modulation of peripersonal space, this experiment suggests that the model implemented on our simulated robot is capable of using proprioceptive information about body characteristics in an efficient manner in conjunction with a touch signal for assessing its peripersonal space. Similarly to the first experiment, where the visuomotor system provided good signals for generating reaching positions, in this case the visuomotor system produced good reachability predictions as can be seen in figures 5.4 and 5.6. Moreover, the embodied proprioceptive signal of arm length might have contributed to this effect. We have mentioned how confidence was modulated by vergence, moreover, it is plausible length or the arm had an effect in the confidence. The robot, having one of its inputs a proprioceptive signal about one characteristic of its body (in this case arm length), was enabled for assessing its reaching potential depending on its morphology. This worked for

the three arm lengths which included the medium one which the robots was not trained for. This happened for both conditions, possibly indicating that the proprioception of the arm's length was used during the reachability assessment.

5.4 Conclusion

This work contributes to the existing literature on peripersonal space in robotics and extends it by including three aspects missing or not commonly present in previous studies. First, the work presented here uses binocular vision which provides more biological plausibility to the model, as vergence is a natural, embodied representation of depth (Lehar 2003). Antonelli *et al.* (2012); Hersch *et al.* (2008); Sturm *et al.* (2009) have used monocular vision. Moreover, the study by Sturm *et al.* (2009) also contrasts with ours in the use of an externally mounted camera on a robot manipulator and the use of visual markers, opposed to the study presented in this chapter, where a humanoid robot was used and provided with binocular cameras in the head and no visual markers were used. Second, the model was embodied and situated. Actually the visuomotor system, being an active-vision one, requires situatedness to operate. The iCub simulator provides a three-dimensional environment to work with. Therefore, the present study contrasts with the disembodied model in Magosso *et al.* (2010) (2010) but also with the work by Chinellato *et al.* (2011), who studied the emergence of implicit sensorimotor mapping on a two-dimensional simulation. Although working in three dimensions presents difficulties such as controlling manipulators with many degrees of freedom, the study of working models in three-dimensional

environments can be more easily transferred to real world robots. Third and last, our work introduces the concept of confidence in reaching assessment, absent in previous works.

Overall, the model implemented in this work was able to generalise for varying arm lengths even without explicitly encoding in any particular frame of reference this morphological characteristic. Arm length and depth information from vergence allowed the robot to, in an action-based approach—“neutral” postures used for the unreachable vs. reaching positions for the reachable—represent the space around it. We believe it is important to further investigate the role of body related signals and how an artificial system can exploit them for developing a way for discerning reachable from unreachable. The approach here presented introduced a model which can be used independently of the morphology and visual capabilities of the system that can be implemented on other robots for more investigation.

6 Experiment 3: Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom

Developing behaviours for interaction with objects close to the body must be a primary goal for any organism for it to survive in the world (Holmes and Spence 2004). Partly because of the infant's limited sensory capacities, it is within peripersonal space—the region of space in which an agent can immediately interact with objects—where the first contact with the world occurs and therefore where the development of motor control starts. According to Pfeifer *et al.* (2006), the origins of intelligence might lie in the interplay between brain, morphology and the environment. Thelen and Smith (1994) also has highlighted the importance of the environment in the development of cognition, action and the development of motor control. However, motor control has not always been considered a relevant element in the study of cognition (Rosenbaum 2005) although it is essential in the creation of behaviours and adaptation. Behaviour and cognition are processes that develop constantly while the agent is immersed in and interacting with the environment (Nolfi 2011), therefore it is necessary to investigate how the motor control for exploring and knowing about the environment emerges and develops, how the biomechanical parts of a system are coordinated to

accomplish a goal. Moreover, behaviours might lead to the creation of sensorimotor, low-level distributed representations that can be at the basis of intelligence, as it has been proposed by Kuniyoshi *et al.* (2004); Pfeifer and Bongard (2007). On this respect, Bernstein (1967) was the first to highlight the issue of motor control and the possibility of behaviour being grounded in the development of it.

Bernstein considered information-processing systems' operation would be very difficult to explain with closed-loop models or any other model of motor control that held the idea of a central unit involved in the production of all decisions necessary for moving every single muscle in every task (Schmidt and Lee 1988). A system like that would have too many independent states that would need to be controlled, too many *degrees of freedom* (DOF).

In the field of robotics, motor control for manipulators has typically been carried out using *inverse kinematics*. In a robot, each of its joints is a degree of freedom. There exist a relationship that allows to know the point in space of the end-effector starting from the configuration of those joints. This process is known as *forward kinematics*. On the other hand, for moving the end-effector of a manipulator to a specific location in space, a configuration in joint space is needed to be found, usually by calculating the *inverse kinematics*. Finding a joints' configuration for a specific end-effector location is a difficult problem and mathematical solutions are costly in terms of computing resources, especially as the number of DOF's becomes higher. This is where the *degrees of freedom problem* (Bernstein 1967) becomes relevant, as it poses difficulties for artificial systems control. Especially when looking for alternatives to the mathematical solutions for motor control, such as the synthetic approaches to

robotics. Addressing it might contribute to the creation of systems that can learn to exploit complex morphological properties of their bodies. However, Bernstein proposed a solution to the degrees of freedom problem which is applicable to robots. He introduced the concept of synergies, groups of muscles working together by the presence of constraints on them, as opposed as each of them acting and being controlled independently. One of such synergies from morphological changes in a system proposed by Bernstein as a way to reducing the problem's complexity is the locking of joints of a redundant biomechanical system.

The locking and unlocking of degrees of freedom has been found to be present in human motor development. Arutyunyan *et al.* (1968) found that people learning to shoot a pistol held the elbow and wrist fixed as opposed to experts who aim using all the joints of the arm. Pianists and drummers can achieve high independence in the movements of the hands (Pressing *et al.* 1996; Shaffer 1976). This evidence suggests that unlocking degrees of freedom can be a path to develop skills and that learning allows the decoupling of joints previously locked together. The degrees of freedom problem and the solution proposed by (Bernstein 1967) has been extensively investigated in psychology (Konczak *et al.* 2009; Newell and Vaillancourt 2001; Vereijken *et al.* 1992) and has been recently studied in robotics (Berthouze and Lungarella 2004; Gomez *et al.* 2004; Rohde and Paolo 2005) as it offers a bio-inspired approach to the development of controllers for humanoid robot's arms.

An equally significant aspect in the development of the human body is the direction it follows. Motor development flows from the top and centre of the body to the tips of the limbs. The spinal cord is the starting point, arms

and legs are controlled second, wrists and ankles next and finally, hands and toes are the last parts of the body to develop. This development of the motor system is said to follow a proximodistal and cephalocaudal direction (Sharma 2005). This phenomenon is easily observed in infants. Berthier *et al.* (1999) demonstrated infants on the onset of reaching mainly use their shoulder and torso rotation for reaching. Their findings support the idea that infants reduce the complexity of movements by reducing the number of degrees of freedom used in the initial stages of learning to reach, possibly simplifying and accelerating the skill learning. The corticospinal tract goes through a large development in the first year after birth, gradually improving the infant's control of the trunk and proximal joints like the shoulder, and later on distal ones in the arm and hand (Armand *et al.* 1997; Kuypers 1981; White *et al.* 1964). This is why younger infants move their limbs in broad, apparently uncontrolled movements in the beginning: they can only control their shoulders. Later on, it is observed that the elbow and wrist also come into play. The corticospinal tract development is also the reason that control of the lower part of the body comes after that of the upper part. The proximodistal and cephalocaudal development has also been found in studies on the usefulness of locking (or freezing) degrees of freedom during skill acquisition (Arutyunyan *et al.* 1968; McDonald *et al.* 1989; Vereijken *et al.* 1992).

Research on robotics and motor control related to unlocking degrees of freedom has been mainly done on manipulators lacking more human-like perceptual abilities. Rohde and Paolo (2005) used a three-dimensional simulated robotic arm, Gomez *et al.* (2004) used a real robotic arm. More recent research aims to investigate the phenomenon by embedding it in

humanoid robots, therefore including morphological and perceptual constraints that allow more direct comparisons with observations of human subjects. Berthouze and Lungarella (2004) used a humanoid robot that developed a swinging behaviour by increasingly using more degrees of freedom. Another example with a humanoid robot is Savastano and Nolfi (2012). The authors used an evolutionary robotics approach and modelled characteristics of the perceptual and motor capabilities of infants from four months to the first year of age in a computer simulated iCub, obtaining babbling and exploratory movements similar to the observed in real infants. The changes included incrementing visual acuity and unlocking degrees of freedom. A study related to arm control and degrees of freedom exploration using evolutionary robotics for development of reaching and manipulation behaviours is Massera *et al.* (2007), where networks capable of fine-grained interaction with objects were successfully evolved by exploiting the morphological constraints of a robotic arm.

In this study we explore the development of motor skills for reaching in the iCub robot. A methodology following the synergy approach as suggested by Bernstein was implemented for the development of the arm controller. We tested the capabilities of robot's neural network controller to learn progressively by locking some degrees of freedom on the iCub's arm before allowing it to explore the space with more degrees of freedom. We believe exploration using bio-inspired mechanisms can aid in the development of precise reaching, necessary for interaction with objects in the peripersonal space. We expected a progressive development to be advantageous over an initial full training that made no use of synergy constraints.

6.1 Methods

For testing our hypothesis, experiments were planned and carried out on the iCub robot simulator (section 3.1). The simulator provided a good working model for the tests, especially because, as it will be detailed later on, parts of the experiment included long exploratory periods. However, due to some limitations in the simulator program, motor exploration stages could not be speeded-up more than to a certain extent.

6.1.1 Robot perception and arm controller

The robot was provided with the visual perception and tracking behaviour system described in section 4.1.1. This system allows the robot to look for and foveate targets around it. The arm of the robot was controlled by a neural network. The network was a three layered feed-forward multilayer perceptron. The input layer had three input neurons connected to the proprioceptive sensors of the robot's head corresponding to tilt, pan and vergence angles' sensors. The hidden layer consisted of forty units whose activity was described by a sigmoid function. The output layer consisted of four units connected to four of the seven joints of the iCub's arm. Output units were also activated according to a sigmoid function. The arm joints controlled were the four most proximal (closer to the torso) ones: two from the shoulder, two from the elbow. See table 6.1 and figure 6.1 for details on the joints used. In the training phases of the experiment, the back-propagation of error algorithm (Rumelhart *et al.* 1986) was used with a learning rate $lr = 0.01$ and a momentum $M = 0.1$ for one thousand epochs. All inputs were previously scaled to be in the range $[0, 1]$.

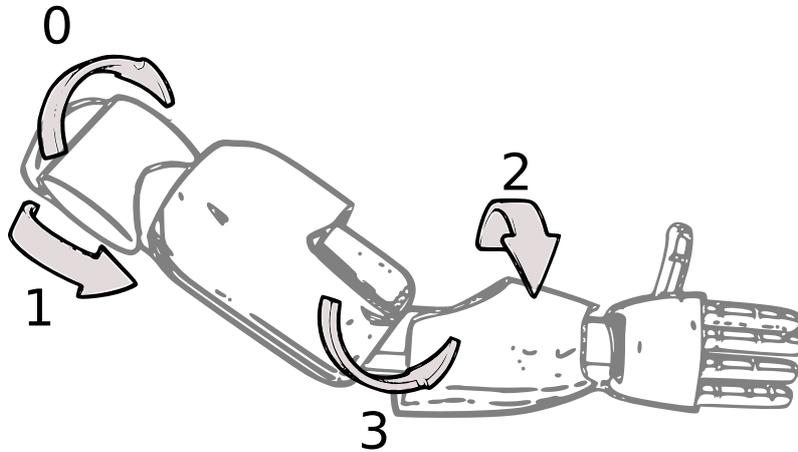


Figure 6.1: Arm joints used in the experiment and their rotation directions. In the Dev condition, only the two most proximal joints (0 and 1) were used in a first phase and later on the two most distal ones (farther from the torso, joints 2 and 3) were included in a second phase. All four joints were used in the single-phased NoDev condition.

Table 6.1: Robot arm joints used in experiment experiment 3.

Arm joint number	Joint	Limits[degrees]
0	Shoulder pitch	[-95.0, 10.0]
1	Shoulder roll	[0.0, 160.8]
2	Shoulder yaw	[-22.0, 95.0]
3	Elbow	[15.5, 106.0]

6.1.2 Experimental conditions and description of task

For the experiment, the robot was tested using nine neural controllers for the arm. The controllers were divided in two experimental groups plus a control group. Each condition group had $n = 3$. The individuals in the first experimental condition group followed a staged or developmental learning process. This was named the **Dev** condition. The second condition group, named **NoDev**, did not follow a staged learning but one consisting of a single learning phase. Finally, a control group condition named **NoTrain** consisted of randomly initialised controllers that did not go through any learning process. The **Dev** condition required the robot to perform the task in an intermediate stage of its learning process, therefore, before continuing with detailed descriptions of the two conditions, we will describe the task the robot was required to perform. Figure 6.2 shows a diagram of learning paths followed by the two experimental condition groups.

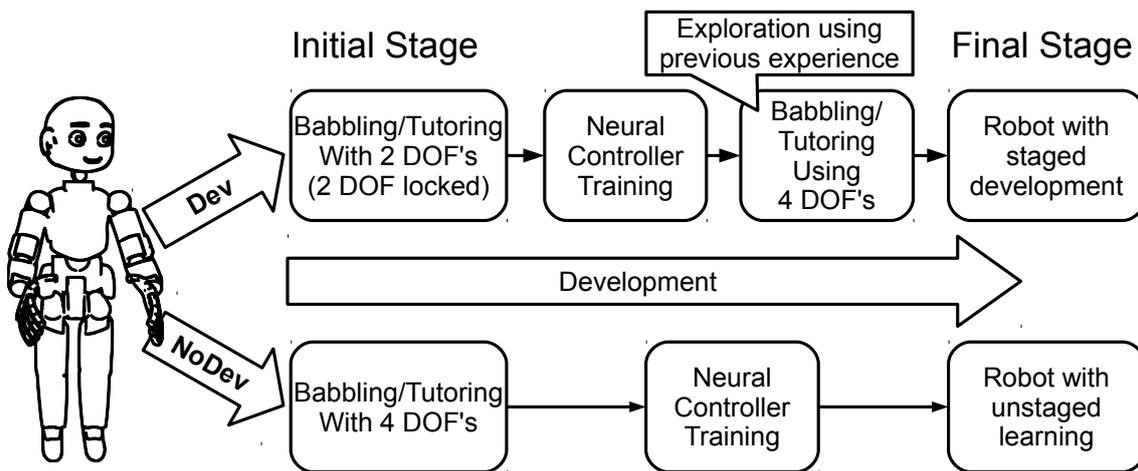


Figure 6.2: The learning path the iCub followed in each of the two testing conditions.

The task for the robot was to gaze to the target and reach for it. The target was a red ball. The visual perception system performed colour segmenting for

locating and foveating it. Figure 6.3 shows how both eyes perceived the target before and after the colour segmenting process and once the target is foveated. In the figure the target is in the middle of the cross-hairs in both images at the bottom. In every trial during either training data generation or evaluation, the red ball had to be positioned at the centre of the retina before any exploration with the arm or reaching attempt. Through the visual system the robot could use vergence to acquire depth information, or distance at which the target was. Earlier in this thesis we have mentioned the usefulness of vergence for depth estimation in humans (in section 2.4.3 and also in chapter 4) as well as its contribution for a reaching task in the previous experiment (section 4.3.1 on page 109) therefore we used it again in this experiment. The visuomotor system provided head/eyes' joint values for the training sets and for activation of the arm neural controller.

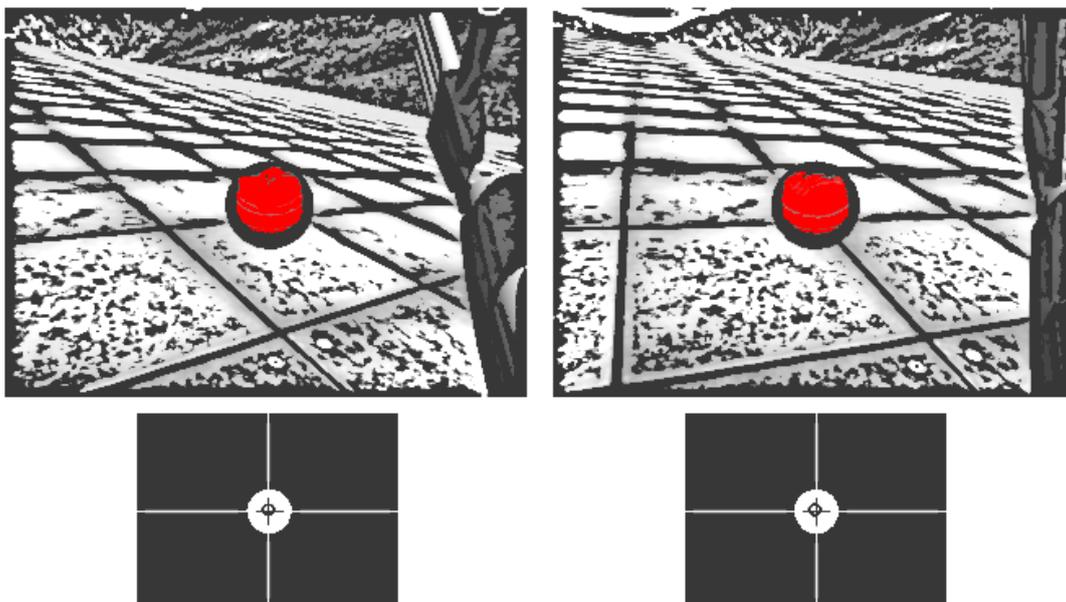


Figure 6.3: Images from the robot's two cameras once the controller has foveated the target. Above, the original images. Below, the low-resolution colour-segmented images.

6.1.3 Stages in the Dev condition

The Dev condition modelled a proximodistal development for motor control in the arm. Its learning process consisted in two phases. Many researchers have suggested infants use exploration and discovery processes for finding solutions to the problem of learning to reach (Berthier 1996; Thelen *et al.* 1993). Following this idea, in the initial phase for this condition, the robot was left to explore the space using motor-babbling (von Hofsten 1982) (for this experiment self-collision was handled as in the two previous experiments, starting a new trial when detected). However, following a proximodistal development, the robot used only joints zero and one from the arm for this initial exploration, corresponding to degrees of freedom located in the shoulder. During this operation the two most distal joints of the arm (located in the elbow) were kept in constant values considered similar to those of a semi-extended arm ($j_2 = 0.0^\circ$ and $j_3 = 50.0^\circ$). The elbow was not fully extended following the findings reported by White *et al.* (1964) indicating rare elbow extension during reaching in around three months old infants.

First stage. The robot was presented targets in the area in front of it where if it had perfect reaching skills it would be able to reach. Once the target was foveated, the robot would attempt to reach it by performing exploration movements with the arm for a certain amount of time. If hand and target collided, a data pair consisting of the joints' values of the head/eyes and those of the arm needed for touching the target were appended to a training-data set. The locations where the targets were placed during this stage were verified to be reachable using only the unlocked

degrees of freedom. Each of the three robots went through this exploratory phase and each generated their own training set. The number of data pairs collected in this manner was five hundred by each robot. Each of these explorations took several hours to be completed due to the simulator program cannot be speeded-up more than to a certain extent. Those sets were used for training the neural network for the arm. The process of exploration and subsequent neural network training was the first stage of development in the **Dev** condition.

Second stage. After the network controller was trained for reaching using the two most proximal joints in the arm, a test phase was carried out in ecological conditions. The robot was presented with the target in different locations that could be reached using four degrees of freedom. Again, in each trial, the robot gazing mechanism was used for foveating the target. Then the trained arm neural controller was activated with proprioceptive data coming from head and eyes. When the robot successfully reached the target, the arm went to its resting position and a new trial was started by presenting the target at a new position. Otherwise, the robot, still in the posture adopted for the reaching attempt, moved the two degrees of freedom that were locked in phase one of the development (2 and 3). Motor babbling for this stage was done in the following manner: the two new joints were added an angle offset corresponding to a random value in the range $[-10,10]$ that followed a normal distribution while the two most proximal degrees of freedom (0 and 1) were kept in the position values the neural controller produced. The sampling distribution and value's range for this stage were decided considering a biological agent in a more advanced stage of development would have better

motor control and, in a process similar to simulated annealing, the exploration of the solution space would be done with smoother movements. This was repeated for twenty times and if it still did not touch the target, a new trial was started. With the additional movement of joints 2 and 3, the robot was sometimes able to reach the target. When that was the case, the data pair of head/eyes' and arm posture was appended to a new training-data set that was used for training in this second phase of development.

6.1.4 *Single-stage NoDev condition*

The **NoDev** condition group went through a learning for reaching postures that used the four degrees of freedom of the arm from the beginning. Similarly to the first stage of development in the **Dev** condition, the robot had to explore the movements capabilities of its arm for generating training data. However, in this condition there were no locked degrees of freedom. The points presented to the robot for the trials were verified to be reachable using four degrees of freedom. Again, each of the three robots generated their own training set, in this condition consisting of one thousand data pairs of values of head/eyes and arm postures.

A summary of the description of the experimental conditions is shown in table 6.2 and a diagram of the different learning paths the robot followed in each condition and the processes involved is shown in figure 6.2.

6.1.5 *A note on time in respect to the iCub simulator*

A time-related aspect in our experiments was the exploration as main component in the learning process and the role of the robotic platform used. The iCub simulator made possible the motor exploration that, if implemented in the real robot would have put it under intense mechanical stress. The

Table 6.2: Description of the experimental conditions.

Condition	Dev	NoDev	NoTrain
Learning process:	Staged (two phases)	Single-stage	No learning (control group)
Degrees of freedom used during learning:	Joints 0,1 in the first phase, joints 2 and 3 added in second phase.	Joints 1,2,3 and 4.	Not applicable.

exploration would have been infeasible for the number of data pairs required if we had used the real robot. However the iCub simulator program was limited regarding the maximum time speed-up. Therefore, the exploratory phases took very large amounts of time for completion. However, this also stresses the importance of the necessity of time in motor learning processes.

6.2 Results

6.2.1 Learning processes

In the first phase of the development of the **Dev** condition, each of the three networks was trained using the data each robot generated during the first exploratory process. In these data-sets, the two outputs corresponding to joints 2 and 3 of the arm had constant values corresponding to the normalised values for 0° and 50° . In this training the mean square error (MSE) became stable after around three hundred epochs. Figure 6.4 shows the training error series for this stage of learning for the three networks. In the second phase of the **Dev** condition, the training-sets included the data collected in the first and second exploratory processes. Figure 6.5 shows the training error series for this second phase. Training error series for the **NoDev** condition are shown in figure 6.6. The mean squared error in this case became stable after around five hundred epochs.

All graphs show error series for the three networks of the corresponding condition and phase.

6.2.2 Evaluation

Once all the controllers reached the final stage (see figure 6.2), tests for comparing the two conditions were performed. The robot was tested with the three controllers of each condition by being presented with a target in front of it. Three hundred target locations were used for the trials, these locations were reachable using four degrees of freedom and were not included in any of the training-sets used during learning. In each trial, the task for the iCub was to foveate the target and attempting to reach as it can be seen in figure 6.7. After arm movement finished, distance from the centre of the palm to the centre

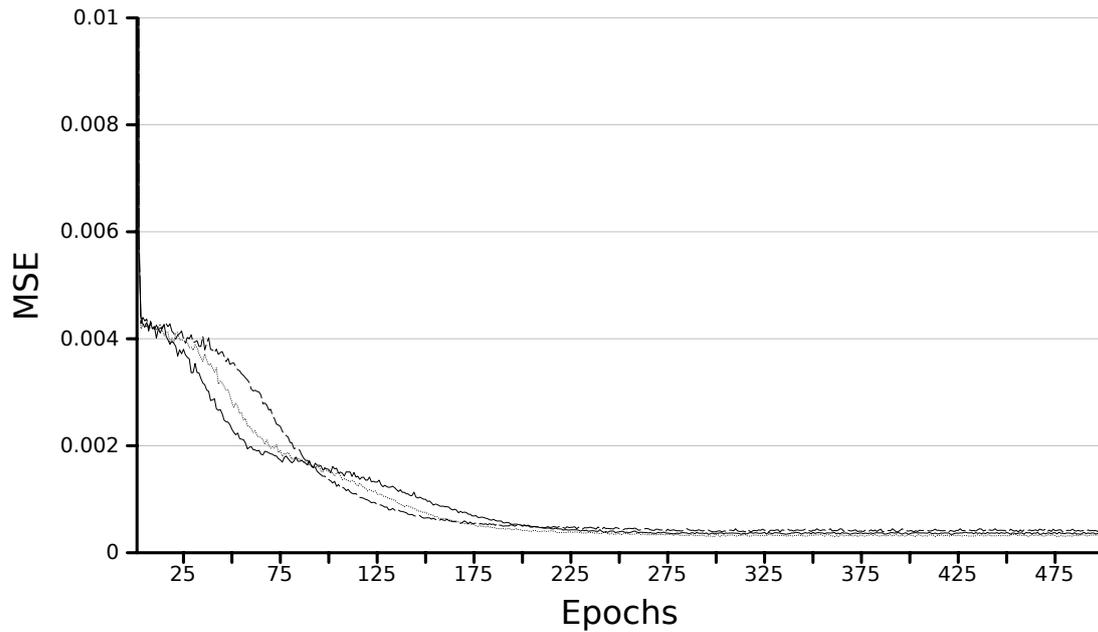


Figure 6.4: Mean squared error during the training of the first stage of development. Training sets for this learning were generated using two degrees of freedom, the other two had fixed values of 0° and 50° .

of the target was recorded. These measurements are shown in figure 6.8. Mean distances in each condition were **NoTrain** ($M = 0.197\text{ m}$, $SD = 0.073$), **NoDev** ($M = 0.074\text{ m}$, $SD = 0.040$) and **Dev** ($M = 0.038\text{ m}$, $SD = 0.019$). For each robot controller, a record of the number of attempts that resulted in touching the target was also taken. The mean success percentage for each condition was **NoTrain** = 0.01, **NoDev** = 0.38 and **Dev** = 0.88 (see figure 6.9).

6.2.3 Analysis

Analysis of the output data indicates the controllers belonging to the staged or developmental learning performed better in terms of final distance to the target as well as in the percentage of success (figure 6.8 and figure 6.9). An analysis of variance was performed on the distance to target to check for statistical difference between the three conditions. This test reported statistical differences $F(2, 897) = 850.45$, $p = 0.000$. Post hoc comparisons

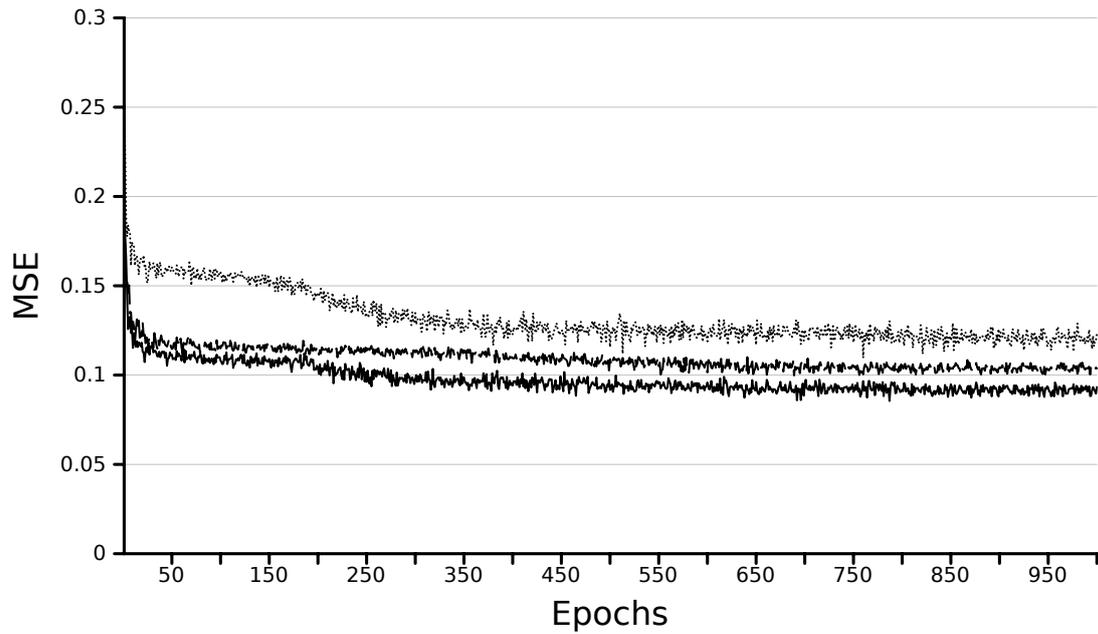


Figure 6.5: Mean squared error during the training of the second stage of development. Training sets for this learning were generated using four degrees of freedom of the arm.

using Tukey HSD test indicated that the mean distance for the three conditions were significantly different, with $p = 0.000022$ between each other.

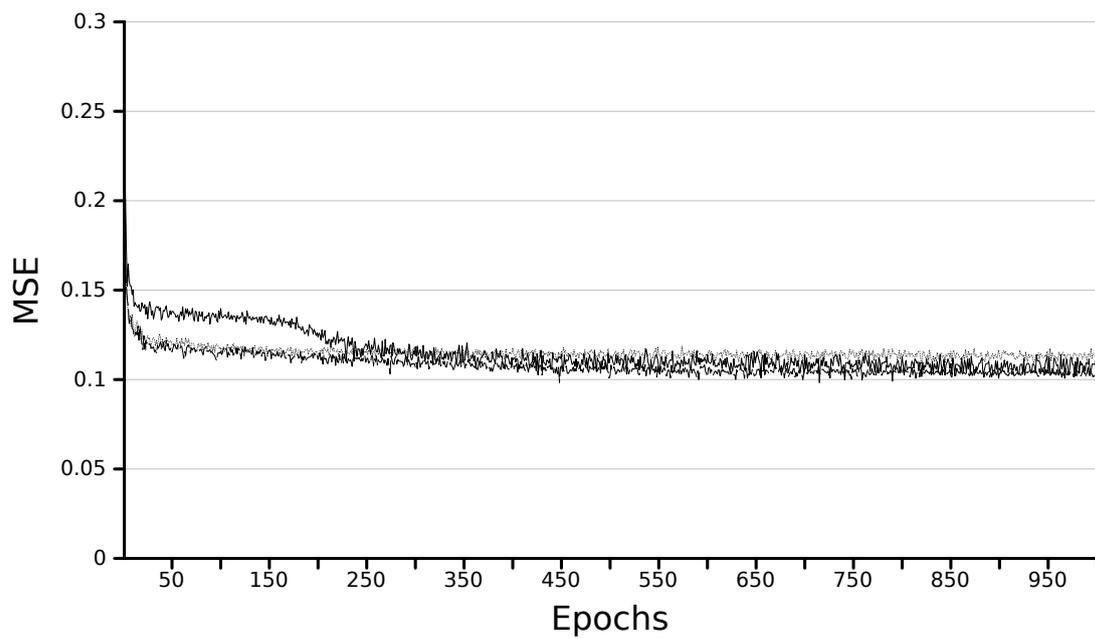


Figure 6.6: Mean squared error during the training of the single-stage NoDev condition. Training data for this condition was generated using four degrees of freedom of the arm.

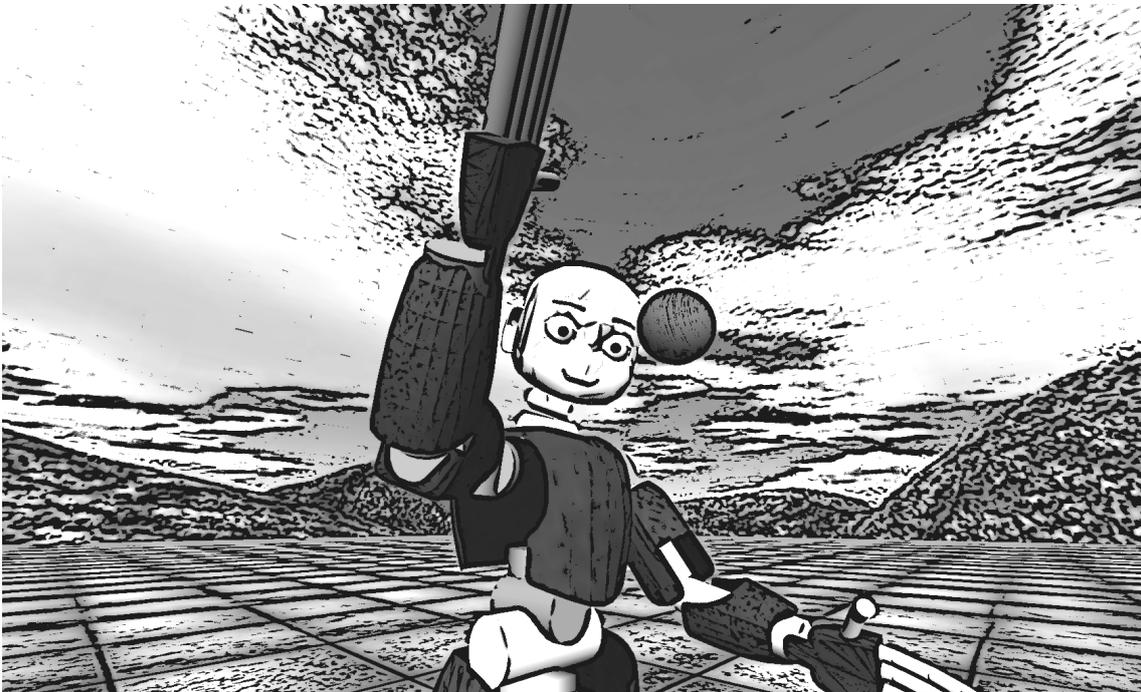


Figure 6.7: The iCub performing the reaching task once it has foveated the target. The robot used two arm joints in the first stage of development and four in the second in the exploratory processes and the four for the evaluation.

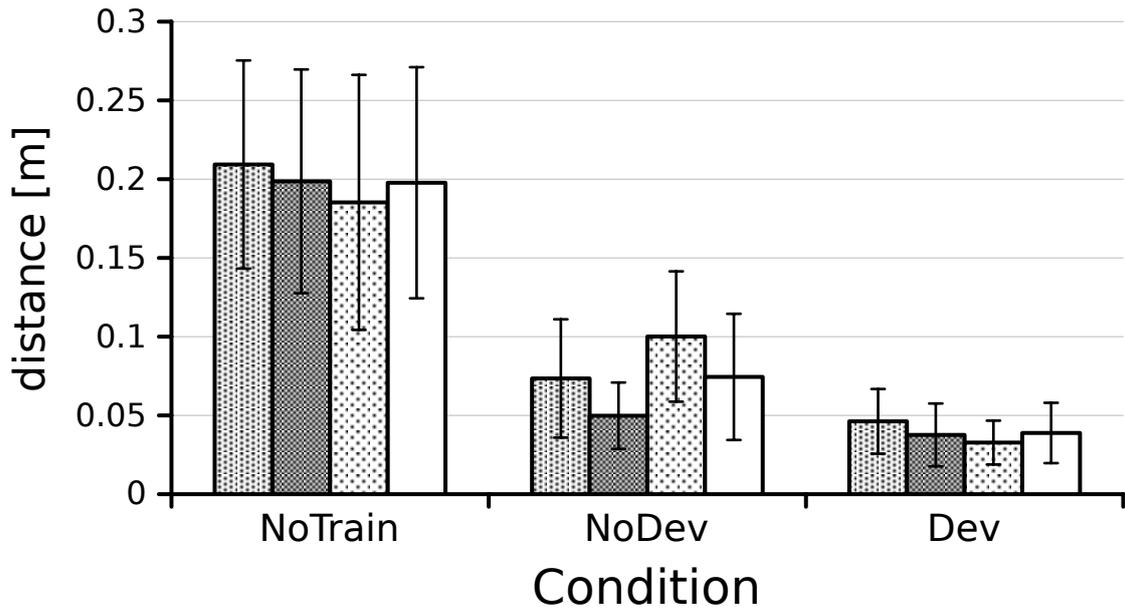


Figure 6.8: Mean distance from the centre of the palm to the centre of the target for each robot in the three conditions after three hundred trials. White bars indicate the mean distance for the three robots in each condition. Error bars indicate standard deviation.

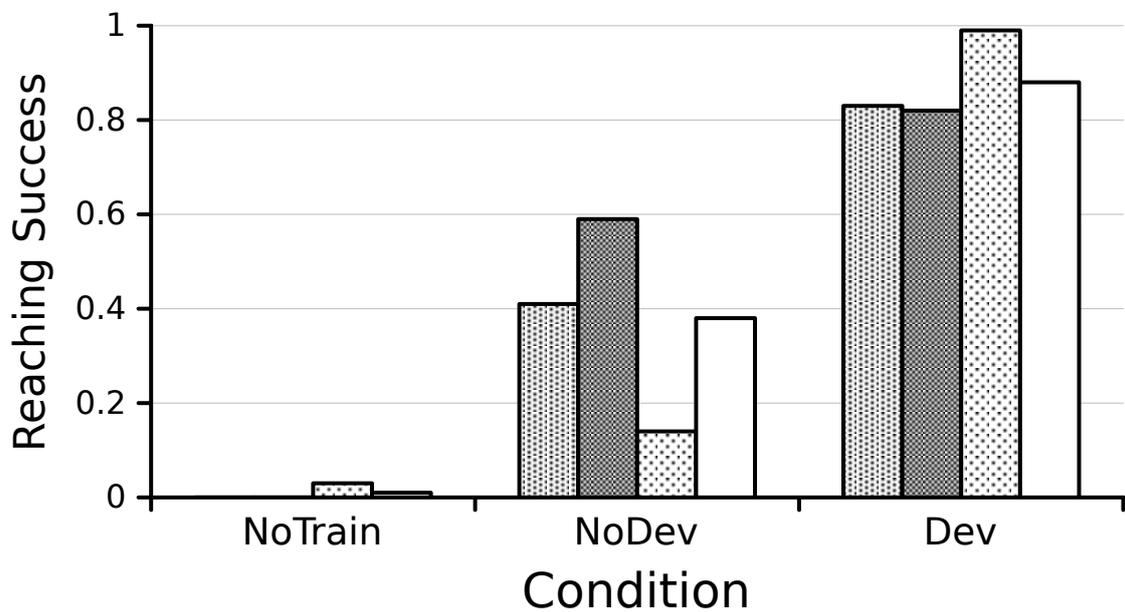


Figure 6.9: Percentage of touching success. Each robot made three hundred reaching attempts to targets in locations reachable with four degrees of freedom. Shaded columns report the success percentage of each robot in the three conditions. White columns correspond to mean success for each condition.

6.3 Discussion

In the present study, we implemented Bernstein (1967) suggestions for simplifying the degrees-of-freedom problem: in our experiment the robot arm controller goes through a developmental progression in order to find a first but simpler solution to the problem and later on, increasing the complexity of the problem. Other roboticists have implemented similar ideas in recent years. Ivanchenko and Jacobs (2003) simulated a three degrees of freedom robot arm that tries to learn the dynamics of the arm while moving on trajectories on a two-dimensional space. The difference with our approach is that in our case the architecture of the networks is the same for every condition, it is the presence or absence of experience that shapes the performance at their final stages. Ivanchenko and Jacobs had a special architecture, devised from the idea that this decouples dynamic interactions among the joints and therefore allows to separately train the joints. Unlike them, for our experiments we decided to keep the same architecture. We want to explore uncoupling of dynamics without changing the internal (not directly exposed to the environment) characteristics of the system. In the work of Ivanchenko and Jacobs, results indicated that a developmentally trained controllers only outperformed the non-developmentally ones when the developmental path matched the nature of the task executed. Similarly, in our experiments, a possible explanation for the better performance of the **Dev** condition over the **NoDev** is that training, as well as the exploratory phases, matched the final task.

As Massera *et al.* (2007), we have started this exploration on a robotic arm with just four degrees of freedom. Our approach contrasts with that one

in that we are interested in the epigenetic development of the skill instead of an evolutionary one. Moreover, in our case, experiments look to include vision into the development of the task instead of direct pass of coordinates or distances to the system without visual processing, as we consider that working towards the implementation of this type of skill development will need real-life sensory capabilities. The head controller for our experiments employs vision as a simple processing but action-related task. Schlesinger *et al.* (2000) have also explored with the locking of DOF's but again, using a non-realistic vision mechanism and a 2D environment and using evolutionary algorithms. Our work has pushed this type of exploration to a more realistic environment and explores the interaction on fixed architecture systems. We showed that even in this circumstances, a developmental approach can lead to better performance. Using the iCub simulator has proven to be a good test-bed for this type of research, as it allowed to implement and test controllers and visual sensors and explore performance in a controlled environment and free of mechanical strain issues.

6.3.1 Benefits of staged learning

We presented a robotic model for learning a reaching skill using motor synergies development in the iCub. The results indicated the advantage of this learning path over one that does not use a staged process. The improved performance of the robots from the **Dev** condition can be observed in figure 6.9 which shows a higher success rate in the reaching attempts for the three robots in that experimental condition, and figure 6.8, which shows the shorter distances achieved when attempting to reach.

According to Berthouze and Lungarella (2004), a lower number of

degrees of freedom reduces the sensorimotor search space and allows a more efficient exploration of it. The **Dev** condition was able to refine its skills during the collection of the second half of samples because the first exploration must have been more efficient than that of the **NoDev** condition. In accordance to the literature on motor control and the degrees of freedom problem, empirical data from this study suggests that for a humanoid robot such as iCub, having redundant degrees of freedom can facilitate the control development of certain task, in spite of the much larger search space, if the exploration of that search space is made gradually by locking some joints in the earlier stages.

6.3.2 Performance and generalisation

The robots in both learning conditions were tested with target locations that were not included in any of their training sets, therefore, during the tests the networks had to generalise for these new input data. The use of four degrees of freedom from in the exploration stage of condition **NoDev** was faced with a large search space, making it more difficult to make the best from the available resources. In other words, the task-relevant knowledge extracted from the samples during exploration was less because it was not guided by any mechanism. On the same line of ideas, it is known that redundancy of degrees of freedom is important because it allows smooth movements and improved dexterity (Hammond III and Shimada 2011; Shadmehr and Wise 2004) however, it also makes the search space too large and it is the origin of the degrees-of-freedom problem. On the other hand, the **Dev** condition, throughout the whole learning process, collected the same number of samples for training (one thousand) but only collected and used for training half this number in the first phase, and later on, when collecting the other half, it was

able to exploit the previous learning making the exploration and learning more relevant to the task. On this aspect, the way in which the second exploration was done must have played a role in the total time required by the data collection. The second random babbling the robot performed was based on normally distributed random position values. The decision of using a normally distribution instead of a uniform one was made from the beginning, considering that the first exploration would already approximate a good solution and the agent, being in a more advance stage, now would have more refined movements in proximal joints. The consequences of using a uniform distribution in the second stage of exploration was not investigated. However, a possibility that arises in that situation is that during the exploration the robot would have missed more times and therefore require more time, although this would have not affected final performance. This due to that in the first attempt the hand was already close to the solution and a normally distributed distribution would likely make small shifts hopefully towards the right positions and by contrast, a uniform distribution is more likely to make a large movements prone to miss the target. The method implemented for making smaller movements in the second stage is analogous to a simulated annealing algorithm that gradually focuses in on a area of the search space where hopefully a solution is to be found.

6.3.3 Observations with respect to the training processes

Regarding the learning processes there some aspects worth looking at. Despite different learning paths were taken in each condition, due to both aimed to solve the same task, between the second training stage of **Dev** and the only training in **NoDev** error series profiles were similar. Figures 6.5

and 6.6 show error during training in the second stage in condition **Dev** and during training of condition **NoDev**, respectively. Training error for the **Dev** condition shows more variability between its three neural network controllers. However, the way in which the training sets were generated played a role in the final behaviour of the system, which was made clear by the data analysis showing statistical difference for reaching distance error. Additionally, training in the second phase of the development condition started from particular conditions. This second learning process began with a weight configuration that already approximated the sought solution. It is possible that both, destructive interference—the “forgetting” neural networks go through when re-trained—was mitigated to some extent for two reasons. First, because the second training set included training data from the first stage, and second, the new data introduced data which shared elements (*i.e.* the first two output values) with the data learnt previously.

6.3.4 Spatial exploration dependency on time

The study presented has been an opportunity to investigate a developmental process in a biologically-inspired methodology focused on the motor control theories of Bernstein and the results can also be considered under more recent theories. The dynamical systems approach to the study of cognition and motor control (Smith and Thelen 2003; Thelen and Smith 1994) investigates the cooperating components that produce stability or elicit changes during development. Under this approach, and following Newell *et al.* (2001), in our experiment the developmental process should be regarded as a function of time instead of a single function for behavioural change across contexts and tasks. In a similar study Corbetta *et al.* (2000),

studied how infants reached for an object and observed the relation of perception and action to be closely linked to the changes in the motor system. The authors found that before the eight months of age, systemic motor tendencies conflict with the perceptual-motor mapping needed for the task. Later on this tendencies disappears, allowing the infants to more successfully reach for the target. Likewise, in the experiment presented in this chapter the characteristics of the exploration the robot performed was time dependant. In $t = 0$, the initial stage, exploration covered a volume of space described by the intersection of two spherical sectors with slightly different radii, on time $t = 1$, the second stage, the volume covered by exploration was widened due to the inclusion of the two new joints. However, contrasting with what happens in natural systems, these stages were discrete. There is still much to be investigated on how to implement continuous-time changes in our approach. Ideally, we should find a way in which robots go through a completely autonomous maturation process. While we achieve that, we have to try and investigate how more development stages can be put in place. In the synthetic approach we took, it was up to the designer deciding when to the constraints on motor properties were lifted, while in natural systems this arises from the interaction with the environment, where temporal processes can affect the agent and its relations with the surroundings.

6.3.5 Motor synergies exploration

The study presented in this chapter investigated the usefulness of unlocking degrees of freedom for developing a reaching skill in the iCub. Intrinsic dynamics of a physical system is another motor synergy to be explored. This motor synergy has been studied in bipedal walking (Collins

et al. 2005; McGeer 1990) and, in closer context to the present study, arm movement by Gottlieb *et al.* (1997). In their study, the authors who found a quasi-linear relation between shoulder and elbow torques in human arm movements on a plane. In this respect, it was argued by Tuller *et al.* (1982) that motor learning is about integrating muscular forces with forces coming from other sources, such as the dynamics originated in the physical properties of the body. Embodied cognition theories suggest that the body offloads work to the intrinsic dynamics of the system, thus performing physical processes rather than computation (Pfeifer and Bongard 2007; Smith and Gasser 2005; Ziemke 2003). In robotics these study approaches are still largely missing due to the materials used for building artificial systems. Traditional materials are rigid and actuators are usually electrical motors, while natural systems use elastic parts and actuators that exhibit damping. One reason for not using these more biologically-plausible materials is the difficulty to mathematically model them due to their non-linear nature. Nevertheless, this seem to be the way forward for developing Tondu *et al.* (2005) have developed a platform with elastic, compliant and damping properties useful for this kind of studies.

6.4 Conclusion

Reaching is an important step in the development of motor and cognitive skills. Exploring this essential skill in many contexts and approaches will give insights of the series of processes emerging in infants (Corbetta *et al.* 2000). Our work on development of reaching tries also to consider the fact that for acquiring a skill it is necessary to has trial-and-error processes where time constraints cannot be avoided. In our experiments, the generation of the

second training set for the staged learning condition, an experience phase that used the motor knowledge acquired in the first stage, took considerably longer than any other part of the experiment. However, we believe this was a very important step due to each network will generate different outputs for the same inputs so the set is particular to each of them. Our system has used two and then four degrees of freedom to explore and then improve a motor skill but the human arm has seven degrees of freedom. Thus it would be valuable investigating the approach taken for more kinematically complex manipulators and more stages in development. The presented study contributes to the study of motor synergies in humanoid robots. However, having a continuous-time development is still an open question and a possible line for future research. Nevertheless, the resulting behaviours were satisfactory. Behaviours In both stages, and the performance obtained in the tests are in accordance to behaviours exhibited by infants in the initial explorations of their peripersonal space.

The modification to the body extended the peripersonal space of the robot depending on arm-length, similarly as what happens following tool use in research done by (Iriki *et al.* 1996) on macaque monkeys and studied in healthy humans by Berti and Frassinetti (2000) in a line-bisecting task and in neglect patients with cross-modal, visual-tactile extinction by Maravita (2002) and in healthy human participants by Longo and Lourenco (2006). Our study seems to be congruent to an interesting finding related to tool-use made by (Holmes *et al.* 2004). The authors found that it is the tips of the tools what is incorporated to the representation of the body and peripersonal space. Moreover, they also reported that when holding the tool, the extension effect

decreased following the length of the tool, similarly to what we saw in the confidence levels in our study. We believe that the decrease of confidence observed in our findings from chapter 5 point towards a possible direction for further research.

7 Overall Discussion and Conclusions

The present chapter begins by summarising the research carried out for completing this thesis, then presents the contributions to knowledge it produced. Finally, it discusses various aspects of the work done, the findings and general observations. It also points to possible lines of research this work can lead to.

The study set out to explore the concept of peripersonal space in a humanoid robot. It identified methodologies for providing a robot with behaviours that can allow it to assess its reachable space, create and adapt implicit representations of it. Using a synthetic methodology, we carried out three studies on the iCub simulator that explored different aspects of peripersonal space, namely, the use of visual vergence, its plastic nature elicited by end-effector length, and the benefits of synergic motor development for action within it.

We investigated the use of simple vision and action mechanisms for learning motor skills and in this way implemented peripersonal space research findings in the iCub. These mechanisms endowed the robot with the ability of creating implicit representations of the space around it. The resulting behaviours were modulated in terms of performance, confidence in reaching assessment and improved motor knowledge acquisition by the morphology of

the visuomotor system and the arm length.

This thesis contributes to the knowledge of the mechanisms that can allow humanoid robots to adapt to their peripersonal space. It took inspiration from the literature of developmental science, focusing on the episodes of the reaching behaviour. Then, it implemented aspects of those on the iCub robotic platform. It produced models that can be applied to any humanoid robot, contributing in this way to the development and technological implementation of autonomous robots. The three experiments involved vision-related reaching tasks. In the case of the first experiment and as a first step, we provided the robot with a visuomotor tracking behaviour that employed vergence, and we focused more on the role of vision (binocularity) in the reaching task.

As observations regarding the work carried out, we would like to start by talking about some topics close to the line of investigation we took. Throughout the thesis work we have presented and discussed several aspects of peripersonal space, and in doing so, we have found a close relation of these to some aspects of embodiment. For our first experiment we developed a visuomotor system (section 4.1.1) that was used from then on in the rest of the thesis. In that experiment we compared the reaching performance of the robot in two visual modalities that used the referred developed visuomotor system. Namely, a monocular one and a binocular one. The results indicated a better performance in the binocular condition. It is worth noticing that although the visuomotor system only uses one feature for foveating the target (*i.e.* colour) which produced very similar retinal images for both eyes, it is the geometrical relation between each eye, considered as individual sensors, what

endows the system with the possibility of extracting more information through its own body characteristics.

On the same line of research, in the second experiment (chapter 5) we explored the plastic properties of peripersonal space in the iCub by providing the controllers with information of its own body characteristics that were not encoded in any particular frame of reference. Length of the arm was a single input into the network. This fact displayed the potential of neural controllers with direct access to embedded sensors in a robot's body, much as the nervous system in the body, of adapting and making use of implicit information from the body. This draws our attention to work on sensory substitution and neuroplasticity. In our experiment the sensor uses only a proprioceptive modality (which provided arm-length). However, it is worth revisiting work like that of Bach-y Rita *et al.* (1969), where the authors studied how the neural circuitry of the body can adapt to sensory information and exploit it when motor activity is involved in the process. In our experiment, action used to learn the body postures needed for reaching, and the head and eyes' angles values that indicated if something was reachable or not, converged the sensory information of having certain length arm into motor knowledge for assessing the reachable space. The neural controller adaptation to body characteristics played a role in producing motor knowledge. Similarly, we observed that in the third experiment the embodiment of the agent was an important factor in the resulting behaviours. The degrees of freedom of the arm played a role in the exploration needed for learning a motor skill. Moreover, changes in the available degrees of freedom through the stages of the learning process, allowed the exploration to be simplified and therefore

more easily exploited by the neural controller.

Overall, we can see the importance of the geometrical properties of the robot closely related to the peripersonal space aspects we investigated. It is clear that in the research on robotic peripersonal space and in humanoid robotics in general, a review of the literature on embodied theories of cognition can contribute in the interpretation of results. Our work using the simulated iCub allowed us mainly to explore geometrical considerations on embodiment. Other aspects of embodiment, that were not possible to explore with the simulator, such as materials and intrinsic dynamics could provide other lines for further investigations. More on the simulator and why these lines of research were not taken is presented further on in this discussion.

An interesting finding in our second experiment was the reaching assessment confidence levels displayed by the iCub. These levels were modulated when targets were farther from the body and also when the robot had a longer arm. In the implementation of this experiment, the robot's arm neural controller was activated as soon as the target was foveated, producing simultaneously both the position control signal for reaching the target (in one of the conditions) and the signal that encoded reaching assessment. This is a reactive behaviour, working similarly to the active-vision system that produces the foveation behaviour of the robot. However, the confidence signal could be used for inhibiting the reaching attempt for saving resources, for instance. In natural systems, body growth provides plenty of contexts, these being the mix of internal and environmental conditions. In these contexts, the elements composing a behaviour interact to, for example, inhibit a reaching behaviour. Finding mechanisms for providing a robot with rich environments that also

provide proper contexts for development is an interesting challenge, and one that makes us, robot designers and scientists, realise the broader extent in which artificial agents should be studied.

Another aspect in the second experiment (chapter 5) dealt with body modifications during the learning of reaching. There, we explored the plastic properties of peripersonal space. The way in which we carried out the experiment, by extending the forearm only, is equivalent to the robot handling a tool in its hand. The modification to the body extended the peripersonal space of the robot depending on arm-length, similarly as what happens following tool use (Berti and Frassinetti 2000; Iriki *et al.* 1996; Longo and Lourenco 2006; Maravita 2002). Moreover, our study was also congruent with the findings reporting that the extension effect in peripersonal space (the sensation of being able to reach farther) decreases as function of the length of the tool. In our study, this can be observed in the confidence levels. The decrease of confidence observed in our findings point towards a possible direction for further research.

In the present work, the three experiments dealt with the creation of space representations, however, these do not explicitly use any reference frame. They encode *sensorimotor mappings* that the robot used for approximating inverse kinematics (especially in chapters 4 and 6) and evaluating the perceived space in order to delimit the reachable region (chapter 5). The representations were encoded in feed-forward neural networks. These type of networks were suitable for our experiments, however, there are other types of networks that allow for the inclusion of temporal dynamics in their activation. Using recurrent neural networks allows the

robotic controller system to have a series of activations that could recreate the movements involved in the reaching. A possible line for investigation could use recurrent neural networks for investigating temporal dynamics in behaviours within the peripersonal space, possibly resulting in interesting outcomes from kinematic analysis.

The bottom-up approach we followed throughout the thesis provided satisfactory. We decided to take this approach partly for avoiding perceptual short-cuts, for example, in the vision capabilities. Contrasting with this approach, in some studies this kind of short-cuts are used. For instance, by feeding the robot controllers with Cartesian coordinates of targets. Instead, we provided the robot with the a visuomotor behaviour based on active-vision and vergence, that doesn't need any visual markers. Which was a contributing element in the results obtained: the confidence levels gradient in chapter 2 were likely to be originated in the inherent depth estimation error of the vergence signal, a result which would have not been observed with a visual perception short-cut like the one mentioned and, although we used a simulated robot, we wanted to make perception more realistic. The simulator allowed using a realistic three-dimensional environment, something also valuable because eventually this research can inspire future studies implemented on the real iCub. The robotic platform iCub has improved a lot since we started this work. Especially, the physical robot is more mechanically robust and the controlling software has been refined.

We have also mentioned numerous times that the simulator provided a good environment for the experiments. It was also very advantageous being able to implement ideas and experiment with the iCub platform resources and

therefore we used it in the three experiments presented. However, some aspects have to be considered when using a simulator like the iCub's, and we would like to mention them for other researchers in robotics to take into account. When selecting a robotic platform or a simulator it is important to fully know its capabilities, features and limitations. If dynamics are to be explored, robustness of the physical simulations has to be verified. For this studies we did not exploited the dynamics of the system but it is worth mentioning that implicit dynamics of mechanical system of a humanoid robot, as we mentioned in the discussion of our last experiment (section 6.3.5 on page 166), are a very promising topic to study for future studies inspired by this thesis. This applies for both motor control and peripersonal space studies, due to the close relation between these two aspects of human space representation. If future researchers are to explore this topic, maybe the best would be to do so directly in real physical environments, because current simulators are still be far from being able to handle the required number of components necessary for complex explorations (material dynamics, elastic robot parts, damping elements, and many more). Nevertheless, we do not to imply that simulators can be discarded. We found in them a great testbed for our research, and the empirical data presented here can hopefully provide with ideas for new simulators or improvements in the current ones. In this work we did not focus on building or improving a robotic platform, which allowed us to focus on the research on peripersonal space. However, we hope the work carried out here will be found useful and inspire roboticists, consecrating on robot design and building, who want to make their robots more suitable for human development research.

Now I want to make a defence of the animat approach and clarify why I have chosen to take an approach in line with it in my research even when having an advanced integrated humanoid. Although the animat approach is seldom used in humanoid robotics in an explicit way nowadays, simple behaviours are still being actively studied on many advanced robotic platforms. The decision of focusing my work on this particular approach was in order to apply the concepts of animats on an humanoid robot so that the internal structures and behaviours exhibited are a layer over which more structures and behaviours can be built upon. I consider important to keep studying this approach. It is my impression that today's cognitive robotics research is tending far towards high level cognition or physiologically based currents. This makes aside once again the simple, basic, closer to sensorimotor behaviours that if are not addressed it seems very hard we can build humanoid robots with real human-like capabilities.

As we have reviewed in chapter 2, there is a large amount of research on peripersonal space done and being carried out in the fields of psychology and neuroscience. This research presents roboticists with several models of human behaviour which are an important source of inspiration for robot design. After all, creating artificial humans has been and still is the dream many of us have, and approaching other sciences always contributes to synthesise novel solutions for old problems. In turn, humanoid robotics have impacted cognitive sciences by offering new tools for research. New tools provide alternative methods for investigation but more importantly, they foster the development of theories (Gigerenzer 1991). The epigenetic robotics approach taken in this work for studying peripersonal space in humanoids successfully

integrated different sensing modalities for synthesising reaching behaviours which can complement empirical studies of peripersonal space in other fields of study. Here we have explored early stages in the development of motor control in peripersonal space in a humanoid robot, however, this is a topic that still needs to be extensively investigated. Robotics research is still far from being able to produce systems with the expected human-like abilities of spatial adaptation and from fully integrating robots in modern human environments. This work contributed to the development towards that aim by presenting a platform and methods for further investigations, and yielded a set of behaviours to be used in further studies and in technological implementations.

List of references

- Adams, B., Breazeal, C. L., Brooks, R. A., and Scassellati, B. (2000). Humanoid robots: a new kind of tool. *IEEE Intelligent Systems and Their Applications: Special Issue on Humanoid Robotics*, 15(4):25–31.
- Akhratuzzaman, M. and Shafie, A. A. (2010). Advancement of Android and Contribution of Various Countries in the Research and Development of the Humanoid Platform. *International Journal of Robotics and Automation (IJRA)*, 1(2):43–57.
- Allen, P. K., Timcenko, A., Yoshimi, B., and Michelman, P. (1993). Automated tracking and grasping of a moving object with a robotic hand-eye system. *Robotics and Automation, IEEE Transactions on*, 9(2):152–165.
- Aloimonos, J., Weiss, I., and Bandyopadhyay, A. (1988). Active vision. *International Journal of Computer Vision*, 1(4):333–356.
- Aloimonos, Y. (1991). Purposive and qualitative active vision. In *Proceedings of the 6th Mediterranean Electrotechnical Conference*, volume 1, pp. 24–27, Ljubljana. IEEE.
- Ambrosini, E., Scorolli, C., Borghi, A. M., and Costantini, M. (2011). Experiencing Objects: the Role of the Body. In Kokinov, B., Karmiloff-Smith, A., Nersessian, N. J., editor, *EuroCogSci 2011, the European Cognitive Science Conference*, Sophia. New Bulgarian University Press.
- Andersson, R. L. (1989). Dynamic sensing in a ping-pong playing robot. *Robotics and Automation, IEEE Transactions on*, 5(6):728–739.
- Antonelli, M., Grzyb, B. J., Castelló, V., and Del Pobil, A. P. (2012). Plastic representation of the reachable space for a humanoid robot. In *From Animals to Animats 12*, pp. 167–176. Springer.
- Appin Knowledge Solutions (2007). *Robotics*. Infinity Science Press.
- Arlotti, M. A. and Granieri, M. N. (1991). A perception technique for a 3-D robotic stereo eye-in-hand vision system. In *Fifth International Conference on Advanced Robotics/Robots in Unstructured Environments' 91 ICAR*, volume 2, pp. 1626–1629, Pisa. IEEE.

-
- Armand, J., Olivier, E., Edgley, S. A., and Lemon, R. N. (1997). Postnatal development of corticospinal projections from motor cortex to the cervical enlargement in the macaque monkey. *The Journal of neuroscience*, 17(1):251–266.
- Arutyunyan, G. H., Gurfinkel, V. S., and Mirskii, M. L. (1968). Investigation of aiming at a target. *Biophysics*, 13:536–538.
- Asada, M. (2011). Can cognitive developmental robotics cause a paradigm shift? In *Neuromorphic and Brain-Based Robots: Trends and Perspectives*, chapter 11, pp. 251–273. Cambridge University Press, Cambridge.
- Asada, M. (2012). Towards language acquisition by cognitive developmental robotics. *Scott-Phillips, Thomas C.; Tamariz, Mónica; Cartmill, Erica A*, pp. 19–26.
- Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M., and Yoshida, C. (2009). Cognitive developmental robotics: A survey. *Autonomous Mental Development, IEEE Transactions on*, 1(1):12–34.
- Asada, M., Noda, S., Tawaratsumida, S., and Hosoda, K. (1995). Vision-based reinforcement learning for purposive behavior acquisition. In *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, volume 1, pp. 146–153. IEEE.
- Aslin, R. N. (1977). Development of binocular fixation in human infants. *Journal of experimental child psychology*, 23(1):133–150.
- Aspell, J. E., Lavanchy, T., Lenggenhager, B., and Blanke, O. (2010). Seeing the body modulates audiotactile integration. *The European journal of neuroscience*, 31(10):1868–1873.
- Bach-y Rita, P., Collins, C. C., Saunders, F. A., White, B., and Scadden, L. (1969). Vision substitution by tactile image projection. *Nature*, 221(5184):963–964.
- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, 76(8):966–1005.
- Ballard, D. (1991). Animate vision. *Artificial intelligence*, 48(1):57–86.
- Ballard, D. H. (1989). Reference Frames for Animate Vision. In *IJCAI*, volume 89, pp. 1635–1641.
- Banks, M. S. and Salapatek, P. (1983). Infant visual perception. *Handbook of child psychology*, 2:435–571.

- Barela, J. A., Sanches, M., Lopes, A. G., Razuk, M., and Moraes, R. (2011). Use of monocular and binocular visual cues for postural control in children. *Journal of Vision*, 11(12):1–8.
- Barnard, S. T. and Fischler, M. A. (1982). Computational stereo. *ACM Computing Surveys (CSUR)*, 14(4):553–572.
- Bauer, A., Wollherr, D., and Buss, M. (2008). Human–robot collaboration: a survey. *International Journal of Humanoid Robotics*, 5(01):47–66.
- Bechtel, W., Abrahamsen, A., and Graham, G. (2001). Cognitive science: History. In Baltes, N. J. S. and B., P., editors, *International Encyclopedia of the Social and Behavioral Sciences*, chapter Cognitive, pp. 2154–2158. Elsevier.
- Beira, R., Lopes, M., Praca, M., Santos-Victor, J., Bernardino, A., Metta, G., Becchi, F., and Saltaren, R. (2006). Design of the robot-cub (iCub) head. In *Proceedings of the International Conference on Robotics and Automation ICRA 2006*, pp. 94–100. IEEE.
- Bekey, G. A. (2005). *Autonomous robots: from biological inspiration to implementation and control*. The MIT Press.
- Bell, J. (2012). Mars exploration: Roving the red planet. *Nature*, 490(7418):34–35.
- Bermúdez, J. L. E., Marcel, A. J., and Eilan, N. E. (1995). *The body and the self*. The MIT Press.
- Bernardino, A. and Santos-Victor, J. (1999). Binocular tracking: integrating perception and control. *IEEE Transactions on Robotics and Automation*, 15(6):1080–1094.
- Bernstein, N. A. (1967). *The coordination and regulation of movements*. Pergamon Press, London.
- Berthier, N. E. (1996). Learning to reach: A mathematical model. *Developmental psychology*, 32(5):811–823.
- Berthier, N. E., Clifton, R. K., McCall, D. D., and Robin, D. J. (1999). Proximodistal structure of early reaching in human infants. *Experimental Brain Research*, 127(3):259–269.
- Berthouze, L. and Lungarella, M. (2004). Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behavior*, 12(1):47–64.

-
- Berthouze, L. and Ziemke, T. (2003). Epigenetic robotics—modelling cognitive development in robotic systems. *Connection Science*, 15(4):147–150.
- Berti, A. and Frassinetti, F. (2000). When far becomes near: remapping of space by tool use. *Journal of cognitive neuroscience*, 12(3):415–20.
- Bicho, E., Erlhagen, W., Louro, L., Costa e Silva, E., Silva, R., and Hipólito, N. (2011). A dynamic field approach to goal inference, error detection and anticipatory action selection in human-robot collaboration. In Dautenhahn, K. and Saunders, J., editors, *New Frontiers in Human-Robot Interaction*, pp. 135–164. John Benjamins Publishing Company.
- Blickhan, R., Seyfarth, A., Geyer, H., Grimmer, S., Wagner, H., and Günther, M. (2007). Intelligence by mechanics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1850):199–220.
- Borghi, A. M. and Scorolli, C. (2012). Parole come strumenti che estendono il corpo. *Sistemi intelligenti*, 24(1):117–126.
- Bradshaw, M. F. and Elliott, K. M. (2003). The role of binocular information in the 'on-line' control of prehension. *Spatial vision*, 16(3-4):295–310.
- Braitenberg, V. (1986). *Vehicles: experiments in synthetic psychology*. MIT Press.
- Brandt, S. A., Smith, C. E., and Papanikolopoulos, N. P. (1994). The Minnesota Robotic Visual Tracker: a flexible testbed for vision-guided robotic research. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 1994 'Humans, Information and Technology'*, volume 2, pp. 1363–1368. IEEE.
- Breazeal, C. (2003). Toward sociable robots. *Robotics and Autonomous Systems*, 42(3):167–175.
- Breazeal, C. (2004). Social interactions in HRI: the robot view. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 34(2):181–186.
- Breazeal, C. and Scassellati, B. (2000). Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*, 8(1):49–74.
- Bremner, A. J., Holmes, N. P., and Spence, C. (2008). Infants lost in (peripersonal) space? *Trends in cognitive sciences*, 12(8):298–305.
- Brooks, R. A. (1986a). A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of*, 2(1):14–23.

- Brooks, R. A. (1986b). Achieving Artificial Intelligence through Building Robots. AI Memo number 899. MIT AI Laboratory, Cambridge, MA.
- Brooks, R. A. (1987). Planning is just a way of avoiding figuring out what to do next. Technical report, MIT AI Laboratory, Cambridge, MA.
- Brooks, R. A. (1989). A robot that walks; emergent behaviors from a carefully evolved networkSocial interactions in HRI: the robot view. *Neural computation*, 1(2):253–262.
- Brooks, R. A. (1990). Elephants don't play chess. *Robotics and autonomous systems*, 6(1):3–15.
- Brooks, R. A. (1991a). Intelligence without representation. *Artificial Intelligence*, 47(1-3):139–159.
- Brooks, R. A. (1991b). New approaches to robotics. *Science*, 253(5025):1227–1232.
- Brooks, R. A. (1996). Prospects for human level intelligence for humanoid robots. In *Proceedings of the First International Symposium on Humanoid Robots (HURO-96)*, pp. 17–24.
- Brooks, R. A. (1997). From earwigs to humans. *Robotics and Autonomous Systems*, 20(2-4):291–304.
- Brooks, R. A., Breazeal, C., Marjanović, M., Scassellati, B., and Williamson, M. M. (1999). The Cog project: Building a humanoid robot. In *Computation for metaphors, analogy, and agents*, pp. 52–87. Springer.
- Brooks, R. A. and Stein, L. A. (1994). Building brains for bodies. *Autonomous Robots*, 1(1):7–25.
- Brown, C. M. (1988). The Rochester robot. Technical report, Rochester University.
- Brozzoli, C., Gentile, G., Bergouignan, L., and Ehrsson, H. H. (2013). A shared representation of the space near oneself and others in the human premotor cortex. *Current Biology*, 23(18):1764–1768.
- Brozzoli, C., Pavani, F., Urquizar, C., Cardinali, L., and Farne, A. (2009). Grasping actions remap peripersonal space. *Neuroreport*, 20(10):913–917.
- Burrige, R. R., Rizzi, A. A., and Koditschek, D. E. (1995). Toward a dynamical pick and place. In *Proceedings of the IEEE International Conference on*

-
- Intelligent Robots and Systems 95 'Human Robot Interaction and Cooperative Robots'*, volume 2, pp. 292–297. IEEE.
- Buttazzo, G. C., Allotta, B., and Fanizza, F. P. (1994). Mousebuster: A robot for real-time catching. *Control Systems, IEEE*, 14(1):49–56.
- Butterworth, G. (2014). *Principles of Developmental Psychology: An Introduction*. Psychology Press.
- Butterworth, G. and Hopkins, B. (1988). Hand-mouth coordination in the newborn baby. *British Journal of Developmental Psychology*, 6:303–314.
- Campbell, M., Hoane Jr, A. J., and Hsu, F.-h. (2002). Deep blue. *Artificial intelligence*, 134(1):57–83.
- Cangelosi, A. (1999). Heterochrony and adaptation in developing neural networks. In Banzhaf, Wolfgang, Daida, Jason, Eiben, A. E., Garzon, Max H., Honavar, Vasant, Jakiela, Mark, and Smith, R. E., editor, *Proceedings of GECCO99 Genetic and Evolutionary Computation Conference*, pp. 1241–1248, San Francisco, CA. Morgan Kaufmann.
- Cangelosi, A., Parisi, D., and Nolfi, S. (1994). Cell division and migration in a "genotype" for neural networks. *Network: computation in neural systems*, 5(4):497–515.
- Cangelosi, A. and Riga, T. (2006). An embodied model for sensorimotor grounding and grounding transfer: Experiments with epigenetic robots. *Cognitive science*, 30(4):673–689.
- Canny, J. (1986). A computational approach to edge detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 8(6):679–698.
- Chinellato, E., Antonelli, M., Grzyb, B. J., and del Pobil, A. P. (2011). Implicit Sensorimotor Mapping of the Peripersonal Space by Gazing and Reaching. *IEEE Transactions on Autonomous Mental Development*, 3(1):43–53.
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension: Embodiment, Action, and Cognitive Extension*. Oxford University Press.
- CMU Social Robots Project (2006). Social Robots Project. [on-line] <http://www.cs.cmu.edu/~social/> Date accessed: 5th Jan 2013.
- Coello, Y., Bartolo, A., Amiri, B., Devanne, H., Houdayer, E., and Derambure, P. (2008). Perceiving what is reachable depends on motor representations:

- evidence from a transcranial magnetic stimulation study. *PloS One*, 3(8):e2862.
- Colby, C. L. and Duhamel, J. R. (1996). Spatial representations for action in parietal cortex. *Brain research. Cognitive brain research*, 5(1-2):105–115.
- Colby, C. L., Duhamel, J. R., and Goldberg, M. E. (1993). Ventral intraparietal area of the macaque: anatomic location and visual response properties. *Journal of neurophysiology*, 69(3):902–914.
- Collins, S., Ruina, A., Tedrake, R., and Wisse, M. (2005). Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085.
- Colombo, C., Allotta, B., and Dario, P. (1995). Affine visual servoing: A framework for relative positioning with a robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pp. 464–471. IEEE.
- Corbetta, D., Thelen, E., and Johnson, K. (2000). Motor constraints on the development of perception-action matching in infant reaching. *Infant Behavior and Development*, 23(3):351–374.
- Costantini, M., Ambrosini, E., Scrolli, C., and Borghi, A. M. (2011). When objects are close to me: affordances in the peripersonal space. *Psychonomic bulletin & review*, 18(2):302–308.
- Cumming, B. G., Johnston, E. B., and Parker, A. J. (1991). Vertical disparities and perception of three-dimensional shape. *Nature*, 349(6308):411–413.
- Dautenhahn, K. (2007). Methodology and themes of human-robot interaction: a growing research field. *International Journal of Advanced Robotic Systems*, 4(1):103–108.
- Degallier, S., Righetti, L., Natale, L., Nori, F., Metta, G., and Ijspeert, A. (2008). A modular bio-inspired architecture for movement generation for the infant-like robot icub. In *Proceedings of the International Conference on Biomedical Robotics and Biomechatronics, BioRob*, pp. 795–800. IEEE.
- Di Nuovo, A. G., Marocco, D., Di Nuovo, S., and Cangelosi, A. (2013). Autonomous learning in humanoid robotics through mental imagery. *Neural Networks*, 41:147–155.
- Di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., and Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. *Experimental brain research*, 91(1):176–180.

-
- Dreyfus, H. L. (1972). *What computers can't do: A critique of artificial reason*. Harper & Row New York.
- Droll, J. and Eckstein, M. (2009). Gaze control and memory for objects while walking in a real world environment. *Visual Cognition*, 17(6):1159–1184.
- Duffy, B. R. (2003). Anthropomorphism and the social robot. *Robotics and autonomous systems*, 42(3):177–190.
- Dürr, V., Krause, A. F., Schmitz, J., and Cruse, H. (2003). Neuroethological concepts and their transfer to walking machines. *The International Journal of Robotics Research*, 22(3-4):151–167.
- Edsinger, A. and Kemp, C. C. (2007). Human-robot interaction for cooperative manipulation: Handing objects to one another. In *Proceedings of the 16th IEEE International Symposium on Robot and Human interactive Communication, RO-MAN 2007*, pp. 1167–1172. IEEE.
- Ekberg, T. L., Rosander, K., von Hofsten, C., Olsson, U., Soska, K. C., and Adolph, K. E. (2013). Dynamic reaching in infants during binocular and monocular viewing. *Experimental Brain Research*, 229(1):1–12.
- Eklundh, J.-O. and Björkman, M. r. (2005). Recognition of objects in the real world from a systems perspective. *Kuenstliche Intelligenz*, 19(2):12–17.
- Enright, J. T. (1998). Monocularly programmed human saccades during vergence changes? *The Journal of physiology*, 512(1):235–250.
- Fagg, A. H. (1994). Reinforcement learning for robotic reaching and grasping. *Advances in psychology*, 105:281–308.
- Farnè, A., Brozzoli, C., and Cardinali, L. (2009). Peripersonal Space and Body Schema: Two Labels for the Same Concept? *Brain Topography*, 21(3-4):252–260.
- Farnè, A., Demattè, M. L., and Làdavas, E. (2005). Neuropsychological evidence of modular organization of the near peripersonal space. *Neurology*, 65(11):1754–1758.
- Farnè, A. and Làdavas, E. (2002). Auditory peripersonal space in humans. *Journal of cognitive neuroscience*, 14(7):1030–1043.
- Farnè, A., Serino, A., and Làdavas, E. (2007). Dynamic Size-Change of Peri-Hand Space Following Tool-Use: Determinants and Spatial Characteristics Revealed Through Cross-Modal Extinction. *Cortex*, 43(3):436–443.

- Feddema, J. T., Lee, C. S. G., and Mitchell, O. R. (1992). Model-based visual feedback control for a hand-eye coordinated robotic system. *Computer*, 25(8):21–31.
- Fernandez-Leon, J. A., Acosta, G. G., and Mayosky, M. A. (2009). Behavioral control through evolutionary neurocontrollers for autonomous mobile robot navigation. *Robotics and Autonomous Systems*, 57(4):411–419.
- Filliat, D. (2003). Map-based navigation in mobile robots: I. A review of localization strategies. *Cognitive Systems Research*, 4(4):243–282.
- Fischer, M. H. (2003). Can we correctly perceive the reaching range of others? *British Journal of Psychology*, 94(4):487–500.
- Floreano, D., Marocco, D., Sauser, E., and Kato, T. (2004). Coevolution of active vision and feature selection. *Biological cybernetics*, 90(3):218–228.
- Floreano, D. and Mattiussi, C. (2008). *Bio-inspired artificial intelligence: theories, methods, and technologies*. MIT press.
- Floreano, D. and Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 26(3):396–407.
- Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and autonomous systems*, 42(3):143–166.
- Fox, D., Burgard, W., and Thrun, S. (1998). Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25(3):195–207.
- Gibaldi, A., Canessa, A., Chessa, M., Sabatini, S. P., and Solari, F. (2011). A neuromorphic control module for real-time vergence eye movements on the iCub robot head. In *Proceedings of the 11th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2011)*, pp. 543–550. IEEE.
- Gibson, J. J. (1966). *The senses considered as perceptual systems*. Houghton Mifflin, Oxford.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological review*, 98(2):254–267.

-
- Goerick, C., Wersing, H., Mikhailova, I., and Dunn, M. (2005). Peripersonal space and object recognition for humanoids. In *Proceedings of the 5th IEEE-RAS International Conference on Humanoid Robots*, pp. 387–392. IEEE.
- Gomez, G., Lungarella, M., and Eggenberger-Hotz (2004). Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity. In *Proceedings of the 4th International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pp. 119 – 122.
- Goodale, M. A. and Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, 15(1):20–25.
- Goodkin, F. (1980). The development of mature patterns of head-eye coordination in the human infant. *Early human development*, 4(4):373–386.
- Goodrich, M. A. and Schultz, A. C. (2007). Human-robot interaction: a survey. *Foundations and trends in human-computer interaction*, 1(3):203–275.
- Gottlieb, G. L., Song, Q., Almeida, G. L., Hong, D.-A., and Corcos, D. (1997). Directional control of planar human arm movement. *Journal of Neurophysiology*, 78(6):2985–2998.
- Granrud, C. E. (1986). Binocular vision and spatial perception on 4-and 5-month-old infants. *Journal of Experimental Psychology: Human Perception and Performance*, 12(1):36–49.
- Graziano, M. and Gross, C. (1993). A bimodal map of space: somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields. *Experimental Brain Research*, 97(1):96–109.
- Graziano, M. and Gross, C. (1995). The representation of extrapersonal space: A possible role for bimodal, visual-tactile neurons. *The cognitive neurosciences*, pp. 1021–1034.
- Graziano, M. S., Hu, X. T., and Gross, C. G. (1997). Coding the locations of objects in the dark. *Science*, 277(5323):239–241.
- Graziano, M. S. A. and Cooke, D. F. (2006). Parieto-frontal interactions, personal space, and defensive behavior. *Neuropsychologia*, 44(6):845–859.
- Greggio, N., Manfredi, L., Laschi, C., Dario, P., and Carrozza, M. C. (2008). Robotcub implementation of real-time least-square fitting of ellipses. In *Proceedings of the 8th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2008)*, pp. 174–181. IEEE.

- Grosso, E., Metta, G., Oddera, A., and Sandini, G. (1996). Robust visual servoing in 3-D reaching tasks. *IEEE Transactions on Robotics and Automation*, 12(5):732–742.
- Guenter, F., Hersch, M., Calinon, S., and Billard, A. (2007). Reinforcement learning for imitating constrained reaching movements. *Advanced Robotics*, 21(13):1521–1544.
- Guillot, A. (2001). The animat contribution to cognitive systems research. *Cognitive Systems Research*, 2(2):157–165.
- Gunnar, M. R. and Nelson, C. A. (1992). *Developmental behavioral neuroscience*. Routledge.
- Hainline, L. (1998). The development of basic visual abilities. In *Perceptual development: Visual, auditory, and speech perception in infancy*, pp. 5–50. Psychology Press.
- Hainline, L. and Riddell, P. M. (1995). Binocular alignment and vergence in early infancy. *Vision research*, 35(23-24):3229–3236.
- Hall, E. T. (1969). *The Hidden Dimension*. Doubleday, Garden City, NY.
- Halligan, P. W. and Marshall, J. C. (1991). Left neglect for near but not far space in man. *Nature*, 350(6318):498–500.
- Hammond III, F. L. and Shimada, K. (2011). Multi-objective weighted isotropy measures for the morphological design optimisation of kinematically redundant manipulators. *International Journal of Mechatronics and Manufacturing Systems*, 4(2):150–170.
- Han, Y. and Lennerstrand, G. (1995). Eye movements in normal subjects induced by vibratory activation of neck muscle proprioceptors. *Acta Ophthalmologica Scandinavica*, 73(5):414–416.
- Han, Y. and Lennerstrand, G. (1998). Effects of neck muscle proprioceptive activation on the dynamics of monocularly driven horizontal vergence movements. *Acta Ophthalmologica Scandinavica*, 76(3):283–288.
- Harris, C. and Stephens, M. (1988). A combined corner and edge detector. In *Proceedings of the 4th Alvey Vision Conference*, pp. 147–151, Manchester.
- Harris, P. and MacFarlane, A. (1974). The growth of the effective visual field from birth to seven weeks. *Journal of Experimental Child Psychology*, 18(2):340–348.

-
- Head, H. and Holmes, G. (1911). Sensory disturbances from cerebral lesions. *Brain*, 34(2-3):102–254.
- Heed, T., Habets, B., Sebanz, N., and Knoblich, G. (2010). Others' Actions Reduce Crossmodal Integration in Peripersonal Space. *Current biology*, 20(15):1345–1349.
- Hering, E. (1977). *The Theory of Binocular Vision*. Plenum Publishing Corporation.
- Hersch, M., Sauser, E., and Billard, A. (2008). Online learning of the body schema. *International Journal of Humanoid Robotics*, 5(02):161–181.
- Hirai, K., Hirose, M., Haikawa, Y., and Takenaka, T. (1998). The development of Honda humanoid robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 2, pp. 1321–1326. IEEE.
- Hirose, M. and Ogawa, K. (2007). Honda humanoid robots development. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1850):11–19.
- Hoffmann, M., Marques, H., Arieta, A., Sumioka, H., Lungarella, M., and Pfeifer, R. (2010). Body schema in robotics: a review. *IEEE Transactions on Autonomous Mental Development*, 2(4):304–324.
- Hoffmann, M. and Pfeifer, R. (2012). The implications of embodiment for behavior and cognition: animal and robotic case studies. *arXiv preprint arXiv:1202.0440*.
- Holmes, N. P., Calvert, G. A., and Spence, C. (2004). Extending or projecting peripersonal space with tools? Multisensory interactions highlight only the distal and proximal ends of tools. *Neuroscience letters*, 372(1):62–67.
- Holmes, N. P., Calvert, G. A., and Spence, C. (2007). Tool use changes multisensory interactions in seconds: evidence from the crossmodal congruency task. *Experimental Brain Research*, 183(4):465–476.
- Holmes, N. P. and Spence, C. (2004). The body schema and the multisensory representation(s) of peripersonal space. *Cognitive Processing*, 5(2):94–105.
- Horn, B. K. P. (1989). *Obtaining shape from shading information*. MIT press.
- Howard, I. P. and Rogers, B. J. (1995). *Binocular vision and stereopsis*. Oxford University Press.

- Hu, B. and Knill, D. C. (2011). Binocular and monocular depth cues in online feedback control of 3D pointing movement. *Journal of Vision*, 11(7):23.
- Hulse, M., McBride, S., Law, J., and Lee, M. (2010). Integration of active vision and reaching from a developmental robotics perspective. *IEEE Transactions on Autonomous Mental Development*, 2(4):355–367.
- Hurley, S. (2001). Perception and action: Alternative views. *Synthese*, 129(1):3–40.
- Iriki, A., Tanaka, M., and Iwamura, Y. (1996). Coding of modified body schema during tool use by macaque postcentral neurones. *Neuroreport*, 7(14):2325–2330.
- Ishiguro, H., Minato, T., Yoshikawa, Y., and Asada, M. (2011). Humanoid platforms for cognitive developmental robotics. *International Journal of Humanoid Robotics*, 8(03):391–418.
- Ivanchenko, V. and Jacobs, R. A. (2003). A developmental approach aids motor learning. *Neural computation*, 15(9):2051–2065.
- Jablonowski, J. and Posey, J. W. (1985). Robotics terminology. In Nof, S. Y., editor, *Handbook of Industrial Robotics*, pp. 1271–1303. John Wiley & Sons.
- Jin, Y. and Meng, Y. (2011). Morphogenetic robotics: An emerging new field in developmental robotics. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 41(2):145–160.
- Jonassen, D. H. (1992). What are cognitive tools? In *Cognitive tools for learning*, pp. 1–6. Springer.
- Kennedy, D. P., Gläscher, J., Tyszka, J. M., and Adolphs, R. (2009). Personal space regulation by the human amygdala. *Nature neuroscience*, 12(10):1226–1227.
- Khatib, O., Yokoi, K., Brock, O., Chang, K., and Casal, A. (1999). Robots in human environments: Basic autonomous capabilities. *The International Journal of Robotics Research*, 18(7):684–696.
- Kidd, C. D. and Breazeal, C. (2005). Human-robot interaction experiments: Lessons learned. In *Proceeding of the Conference on Artificial Intelligence and Simulation of Behavior (AISB 2005)*, volume 5, pp. 141–142.
- Klatzky, R. L. and Lederman, S. J. (2010). Multisensory Texture Perception. In Kaiser, J. and Naumer, M. J., editors, *Multisensory Object Perception in the Primate Brain*, pp. 211–230. Springer, New York, NY.

-
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69.
- Konczak, J., vander Velden, H., and Jaeger, L. (2009). Learning to play the violin: motor control by freezing, not freeing degrees of freedom. *Journal of motor behavior*, 41(3):243–252.
- Kozima, H. and Yano, H. (2001). A robot that learns to communicate with human caregivers. In *Proceedings of the First International Workshop on Epigenetic Robotics*, pp. 47–52.
- Kuniyoshi, Y., Yorozu, Y., Ohmura, Y., Terada, K., Otani, T., Nagakubo, A., and Yamamoto, T. (2004). From humanoid embodiment to theory of mind. In *Embodied artificial intelligence*, pp. 202–218. Springer.
- Kunz, C., Murphy, C., Camilli, R., Singh, H., Bailey, J., Eustice, R., Jakuba, M., Nakamura, K.-i., Roman, C., and Sato, T. (2008). Deep sea underwater robotic exploration in the ice-covered arctic ocean with AUVs. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, pp. 3654–3660. IEEE.
- Kurfess, T. R. (2010). *Robotics and automation handbook*. CRC press.
- Kuypers, H. (1981). Anatomy of the descending pathways. In Brookhart J M, M. V. B., editor, *Handbook of physiology. The nervous system, part II*, pp. 597–666. American Physiological Society.
- Lacey, S., Peters, A., and Sathian, K. (2007). Cross-modal object recognition is viewpoint-independent. *PloS one*, 2(9):269–274.
- Làdavas, E., Berti, A., Ruozzi, E., and Barboni, F. (1997). Neglect as a deficit determined by an imbalance between multiple spatial representations. *Experimental Brain Research*, 116(3):493–500.
- Làdavas, E. and Serino, A. (2008). Action-dependent plasticity in peripersonal space representations. *Cognitive neuropsychology*, 25(7):1099–1113.
- Lakoff, G. and Johnson, M. (1980). *Metaphors we live by*. University of Chicago Press, Chicago.
- Law, J., Lee, M., Hulse, M., and Tomassetti, a. (2011). The infant development timeline and its application to robot shaping. *Adaptive Behavior*, 19(5):335–358.

- Lee, M., Law, J., Shaw, P., and Sheldon, M. (2012a). An infant inspired model of reaching for a humanoid robot. In *Proceedings of the IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, pp. 1–6. IEEE.
- Lee, M., Nehmzow, U., and Rodrigues, M. (2012b). Towards cognitive robotics: robotics, biology and developmental psychology. In *The Complex Mind: An Interdisciplinary Approach*, pp. 103–125. Palgrave Macmillan.
- Legg, S. and Hutter, M. (2007). A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and Applications*, 157:17–24.
- Lehar, S. M. (2003). *The world in your head: A gestalt view of the mechanism of conscious experience*. Psychology Press.
- Leigh, R. J. and Zee, D. S. (1999). *The neurology of eye movements*. Oxford University Press, New York.
- Lennerstrand, G., Han, Y., and Velay, J.-L. (1996). Properties of eye movements induced by activation of neck muscle proprioceptors. *Graefe's archive for clinical and experimental ophthalmology*, 234(11):703–709.
- Lewis, E., Lloyd, D. M., and Farrell, M. J. (2012). The role of the environment in eliciting phantom-like sensations in non-amputees. *Frontiers in psychology*, 3(600):1–10.
- Li, M., Betsis, D., and Lavest, J.-M. (1994). Kinematic calibration of the KTH head-eye system. Technical report, Computational Vision and Active Perception Laboratory (CVAP), Royal Institute of Technology (KTH), Stockholm.
- Liang, Y.-s., Zhang, X.-l., Huang, H., Yang, Y.-f., Jin, W.-t., and Sang, Z.-x. (2012). A deformable spherical planet exploration robot. In *Proceedings of the International Conference on Graphic and Image Processing (ICGIP 2012)*. International Society for Optics and Photonics.
- Lie, I. (1965). Convergence as a cue to perceived size and distance. *Scandinavian Journal of Psychology*, 6(4):109–116.
- Lim, H.-O. and Takanishi, A. (2007). Biped walking robots created at Waseda University: WL and WABIAN family. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1850):49–64.

-
- Longo, M. R. and Lourenco, S. F. (2006). On the nature of near space: Effects of tool use and the transition to far space. *Neuropsychologia*, 44(6):977–981.
- Lourenco, S. F. and Longo, M. R. (2009). The plasticity of near space: evidence for contraction. *Cognition*, 112(3):451–456.
- Lungarella, M. and Berthouze, L. (2002). On the interplay between morphological, neural, and environmental dynamics: a robotic case study. *Adaptive Behavior*, 10(3-4):223–241.
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- Macaluso, E. and Maravita, A. (2010). The representation of space near the body through touch and vision. *Neuropsychologia*, 48(3):782–795.
- Maes, P. (1994). Modeling adaptive autonomous agents. In Langton, C. G., editor, *Artificial Life*, pp. 135–162. The MIT Press, Cambridge, MA.
- Magosso, E., Zavaglia, M., Serino, A., di Pellegrino, G., and Ursino, M. (2010). Visuotactile representation of peripersonal space: a neural network study. *Neural computation*, 22(1):190–243.
- Maravita, A. (2002). Tool-use changes multimodal spatial interactions between vision and touch in normal humans. *Cognition*, 83(2):25–34.
- Maravita, A., Spence, C., and Driver, J. (2003). Multisensory integration and the body schema: Close to hand and within reach. *Current biology*, 13(13):531–539.
- Marotta, J. J., Behrmann, M., and Goodale, M. A. (1997). The removal of binocular cues disrupts the calibration of grasping in patients with visual form agnosia. *Experimental Brain Research*, 116(1):113–121.
- Marotta, J. J., Kruyer, A., and Goodale, M. A. (1998). The role of head movements in the control of manual prehension. *Experimental Brain Research*, 120(1):134–138.
- Marr, D. (1974). The computation of lightness by the primate retina. *Vision Research*, 14(12):1377–1388.
- Marr, D. and Poggio, T. (1976). Cooperative computation of stereo disparity. *Science*, 194(4262):283–287.

- Marr, D. and Poggio, T. (1979). A Computational Theory of Human Stereo Vision. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 204(1156):301–328.
- Massera, G., Cangelosi, A., and Nolfi, S. (2007). Evolution of prehension ability in an anthropomorphic neurorobotic arm. *Frontiers in neurorobotics*, 1:4.
- Mata, M., Armingol, J. M., and de la Escalera, A. (2005). Learning Visual Landmarks for Mobile Robot Topological Navigation. In *Machine Learning and Robot Perception*, pp. 1–55. Springer Verlag.
- Maurer, D. and Maurer, C. (1988). *The world of the newborn*. Basic Books, New York, NY.
- Maxwell, J. S. and Schor, C. M. (1996). Adaptation of vertical eye alignment in relation to head tilt. *Vision research*, 36(8):1195–1205.
- McClelland, J. L. and Rumelhart, D. E. (1986). *Parallel Distributed Processing, Volume 2 Explorations in the Microstructure of Cognition: Psychological and Biological Models*. MIT Press, Cambridge, MA.
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4):115–133.
- McDonald, P. V., van Emmerik, R. E., and Newell, K. M. (1989). The effects of practice on limb kinematics in a throwing task. *Journal of motor behavior*, 21(3):245–264.
- McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82.
- McKee, S. P. and Taylor, D. G. (2010). The precision of binocular and monocular depth judgments in natural settings. *Journal of vision*, 10(10):1–13.
- Meeden, L. A. and Blank, D. S. (2006). Introduction to developmental robotics. *Connection Science*, 18(2):93–96.
- Melmoth, D. R. and Grant, S. (2006). Advantages of binocular vision for the control of reaching and grasping. *Experimental Brain Research*, 171(3):371–388.
- Melmoth, D. R., Storoni, M., Todd, G., Finlay, A. L., and Grant, S. (2007). Dissociation between vergence and binocular disparity cues in the control of prehension. *Experimental Brain Research*, 183(3):283–298.

-
- Meltzoff, A. N. and Moore, M. K. (1997). Explaining facial imitation: a theoretical model. *Early development and parenting*, 6(34):179–192.
- Metta, G., Fitzpatrick, P., and Natale, L. (2006). YARP: Yet Another Robot Platform. *International Journal of Advanced Robotic Systems*, 3(1):43–48.
- Metta, G., Sandini, G., and Konczak, J. (1999). A developmental approach to visually-guided reaching in artificial systems. *Neural networks*, 12(10):1413–1427.
- Metta, G., Sandini, G., Vernon, D., Natale, L., and Nori, F. (2008). The iCub humanoid robot: an open platform for research in embodied cognition. In *Proceedings of the 8th workshop on performance metrics for intelligent systems (PerMIS'08)*, pp. 50–56, Washington DC. ACM.
- Miglino, O., Lund, H. H., and Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial life*, 2(4):417–434.
- Milner, A. D. and Goodale, M. A. (1998). *The Visual Brain in Action*. Oxford University Press, Oxford.
- Mon-Williams, M. and Dijkerman, H. C. (1999). The use of vergence information in the programming of prehension. *Experimental brain research*, 128(4):578–582.
- Mondada, F., Franzi, E., and Guignard, A. (1999). The development of khepera. In *Proceedings of the 1st international Khepera workshop*, volume 64, pp. 7–13, Paderborn.
- Morita, T., Shibuya, K., and Sugano, S. (1998). Design and control of mobile manipulation system for human symbiotic humanoid: Hadaly-2. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 2, pp. 1315–1320. IEEE.
- Morse, A. F., Belpaeme, T., Cangelosi, A., and Smith, L. B. (2010a). Thinking with your body: Modelling spatial biases in categorization using a real humanoid robot. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pp. 1362–1368, Portland.
- Morse, A. F., de Greeff, J., Belpaeme, T., and Cangelosi, A. (2010b). Epigenetic Robotics Architecture (ERA). *IEEE Transactions on Autonomous Mental Development*, 2(4):325–339.
- Mowforth, P., Siebert, J. P., Jin, Z., and Urquhart, C. W. (1990). A head called Richard. In *Proceedings of the British Machine Vision Conference*, pp. 361–366, Oxford.

- Murphy, R. (2000). *An Introduction to AI Robotics*. MIT Press.
- Narioka, K., Niiyama, R., Ishii, Y., and Hosoda, K. (2009). Pneumatic musculoskeletal infant robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems WorkShop (IROS2009)*, Louis. IEEE.
- Natale, L., Metta, G., and Sandini, G. (2002). Development of auditory-evoked reflexes: Visuo-acoustic cues integration in a binocular head. *Robotics and Autonomous Systems*, 39(2):87–106.
- Nelson, A. L., Barlow, G. J., and Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57(4):345–370.
- Newell, K. M., Liu, Y.-T., and Mayer-Kress, G. (2001). Time scales in motor learning and development. *Psychological review*, 108(1):57–82.
- Newell, K. M. and Vaillancourt, D. E. (2001). Dimensional change in motor learning. *Human movement science*, 20(4-5):695–715.
- Nilsson, N. (1984). Shakey the Robot. Technical Note 323. SRI AI Center.
- Nolfi, S. (2011). Behavior and cognition as a complex adaptive system: Insights from robotic experiments. In Gabbay, D. M., Thagard, P., Woods, J., and Hooker, C. A., editors, *Philosophy of Complex Systems*, chapter 4, pp. 443–466. Elsevier, Amsterdam.
- Nolfi, S. and Floreano, D. (1998). Coevolving predator and prey robots: Do “arms races” arise in artificial evolution? *Artificial Life*, 4(4):311–335.
- Nolfi, S., Floreano, D., Miglino, O., and Mondada, F. (1994). How to evolve autonomous robots: Different approaches in evolutionary robotics. In Brooks, R. and Maes, P., editors, *Proceedings of the IV International Workshop on Artificial Life*, pp. 190–197, Cambridge, MA. MIT Press.
- Oates, J., Grayson, A., and Wood, C. (2005). *Psychological Development and Early Childhood*. Wiley-Blackwell.
- Olson, T. J. and Potter, R. D. (1989). Real time vergence control. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’89)*, pp. 404–409. IEEE.
- O’Regan, J. K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5):939–972.

-
- Papanikolopoulos, N. P. and Smith, C. E. (1995). Computer vision issues during eye-in-hand robotic tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pp. 2989–2994. IEEE.
- Peniak, M., Marocco, D., Ramirez-Contla, S., and Cangelosi, A. (2009). An active vision system for navigating unknown environments: An evolutionary robotics approach for space research. In *Proceedings of IJCAI-09 Workshop on Artificial Intelligence in Space*, Pasadena.
- Peniak, M., Morse, A., Larcombe, C., Ramirez-Contla, S., and Cangelosi, A. (2011). Aquila: An open-source GPU-accelerated toolkit for cognitive and neuro-robotics research. In *Proceedings of The International Joint Conference on Neural Networks (IJCNN) 2011*, pp. 1753–1760. IEEE.
- Pfeifer, R. and Bongard, J. (2007). *How the body shapes the way we think: a new view of intelligence*. MIT Press.
- Pfeifer, R., Iida, F., and Gómez, G. (2006). Morphological computation for adaptive behavior and cognition. In *International Congress Series*, volume 1291, pp. 22–29. Elsevier.
- Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853):1088–1093.
- Piaget, J. (1952). *The origin of intelligence in the child*. Routledge.
- Posner, M. I. and Cohen, Y. (1984). Components of visual orienting. *Attention and performance X: Control of language processes*, 32:531–556.
- Pressing, J., Summers, J., and Magill, J. (1996). Cognitive multiplicity in polyrhythmic pattern performance. *Journal of Experimental Psychology: Human Perception and Performance*, 22(5):1127–1148.
- Prince, C., Helder, N., and Hollich, G. (2005). Ongoing emergence: A core concept in epigenetic robotics. In *Proceedings of the Fifth International Workshop on Epigenetic Robotics*, pp. 63–70, Lund. Lund University Cognitive Studies.
- Rader, N. and Stern, J. D. (1982). Visually elicited reaching in neonates. *Child development*, 54(4):1004–1007.
- Ramachandran, V. S. (1993). Interactions between motion, depth, color and form: The utilitarian theory of perception. In *Vision: coding and efficiency*. Cambridge University Press.

- Ramirez-Contla, S., Cangelosi, A., and Marocco, D. (2012). Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom on the iCub Humanoid Robot. In *Proceeding of the Postgraduate Conference on Robotics and Development of Cognition (RobotDoC-PhD)*, Lausanne.
- Ramirez-Contla, S. and Marocco, D. (2013). Adaptive Reachability Assessment in the Humanoid Robot iCub. In *Proceedings of the IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob 2013)*, Osaka.
- Ramirez-Contla, S., Marocco, D., and Cangelosi, A. (2010). Active vision and depth estimation: Towards a peripersonal space encoding in a humanoid robot. In *Proceedings of the Tenth International Conference on Epigenetic Robotics*, pp. 175–176, Lund. Lund University.
- Rao, R. P. N., Shon, A. P., and Meltzoff, A. N. (2004). A Bayesian Model of Imitation in Infants and Robots. In Nehaniv, C. and Dautenhahn, K., editors, *Imitation and Social Learning in Robots, Humans, and Animals: Behavioural, Social and Communicative Dimensions*, pp. 217–247. Cambridge University Press.
- Reichelt, S., Haussler, R., Fütterer, G., and Leister, N. (2010). Depth cues in human visual perception and their realization in 3D displays. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*.
- Rizzi, A. A. and Koditschek, D. E. (1994). Further progress in robot juggling: Solvable mirror laws. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2935–2940. IEEE.
- Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G., and Matelli, M. (1988). Functional organization of inferior area 6 in the macaque monkey. *Experimental Brain Research*, 71(3):491–507.
- Rizzolatti, G., Fadiga, L., Fogassi, L., and Gallese, V. (1997). The space around us. *Science*, 277(5323):190–191.
- Rochat, P., Blass, E. M., and Hoffmeyer, L. B. (1988). Oropharyngeal control of hand-mouth coordination in newborn infants. *Developmental Psychology*, 24(4):459–463.
- Rohde, M. and Paolo, E. (2005). t for Two Linear Synergy Advances the Evolution of Directional Pointing Behaviour. In Capcarrère, M., Freitas, A., Bentley, P., Johnson, C., and Timmis, J., editors, *Advances in Artificial Life*, Lecture Notes in Computer Science, pp. 262–271. Springer, Berlin.

-
- Roschin, V. Y., Frolov, A. A., Burnod, Y., and Maier, M. A. (2011). A neural network model for the acquisition of a spatial body scheme through sensorimotor interaction. *Neural computation*, 23(7):1821–1834.
- Rosenbaum, D. A. (2005). The Cinderella of psychology: the neglect of motor control in the science of mental life and behavior. *American Psychologist*, 60(4):308–317.
- Rosheim, M. E. (1994). *Robot evolution: the development of anthrobotics*. John Wiley & Sons.
- Rosheim, M. E. (2006). Leonardo's Knight. In *Leonardo's Lost Robots*, pp. 69–113. Springer, Berlin.
- Roy, N., Baltus, G., Fox, D., Gemperle, F., Goetz, J., Hirsch, T., Margaritis, D., Montemerlo, M., Pineau, J., and Schulte, J. (2000). Towards personal service robots for the elderly. In *Proceedings of the Workshop on Interactive Robots and Entertainment (WIRE 2000)*, volume 25, p. 184, Pittsburgh.
- Rumelhart, D. and McClelland, J. (1986). *Parallel Distributed Processing, Volume 1 Explorations in the Microstructure of Cognition: Foundations*. MIT Press, Cambridge, MA.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088):533–536.
- Rutkowska, J. C. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior*, 2(4):349–373.
- Sakagami, Y., Watanabe, R., Aoyama, C., Matsunaga, S., Higaki, N., and Fujimura, K. (2002). The intelligent ASIMO: System overview and integration. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pp. 2478–2483. IEEE.
- Sandini, G., Metta, G., and Konczak, J. (1997). Human sensori-motor development and artificial systems. In *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Intellectual Human Activity Support for Applications (AIR&IHAS '97)*, pp. 303–314, Riken.
- Sandini, G. and Tagliasco, V. (1980). An anthropomorphic retina-like structure for scene analysis. *Computer Graphics and Image Processing*, 14(4):365–372.
- Santos-Victor, J., van Trigt, F., and Sentieiro, J. (1994). Medusa - A stereo head for active vision. In *Proceedings of the International Symposium on Intelligent Robotic Systems*, Grenoble. Citeseer.

- Savastano, P. and Nolfi, S. (2012). Incremental learning in a 14 dof simulated icub robot: Modeling infant reach/grasp development. In *Biomimetic and Biohybrid Systems*, pp. 250–261. Springer.
- Savastano, P. and Nolfi, S. (2013). A robotic model of reaching and grasping development. *IEEE Transactions on Autonomous Mental Development*, 5(4):326–336.
- Saxena, A., Chung, S. H., and Ng, A. (2005). Learning depth from single monocular images. In *Proceedings of the 19th Annual Conference on Neural Information Processing Systems*, volume 18, pp. 1161–1168.
- Schicke, T. (2007). Human peripersonal space: Evidence from functional magnetic resonance imaging. *The Journal of neuroscience*, 27(14):3616–3617.
- Schlesinger, M., Amso, D., Johnson, S. P., Hantehzadeh, N., and Gupta, L. (2012). Using the iCub simulator to study perceptual development: A case study. In *Proceedings of the IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL 2012)*. IEEE.
- Schlesinger, M., Berthouze, L., and Balkenius, C. (2008). Heterochrony: It's (all) about time! In *Proceedings of the Eighth Interantional Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pp. 111–117, Lund. Lund University Cognitive Studies.
- Schlesinger, M., Parisi, D., and Langer, J. (2000). Learning to Reach by Constraining the Movement Search Space. *Developmental Science*, 3(1):67–80.
- Schmidt, R. A. and Lee, T. (1988). *Motor Control and Learning: A Behavioral Emphasis*. Human kinetics.
- Scholtz, J. (2003). Theory and evaluation of human robot interactions. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences*. IEEE.
- Sciavicco, L. and Villani, L. (2009). *Robotics: modelling, planning and control*. Springer.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and brain sciences*, 3(3):417–457.
- Serino, A., Annella, L., and Avenanti, A. (2009). Motor properties of peripersonal space in humans. *PloS one*, 4(8):e6582.

-
- Servos, P., Goodale, M. A., and Jakobson, L. S. (1992). The role of binocular vision in prehension: a kinematic analysis. *Vision Research*, 32(8):1513–1521.
- Shadmehr, R. and Wise, S. P. (2004). *The Computational Neurobiology of Reaching and Pointing: A Foundation for Motor Learning (Computational Neuroscience Series)*. MIT Press.
- Shaffer, L. H. (1976). Intention and performance. *Psychological Review*, 83(5):375–393.
- Shams, L., Kamitani, Y., and Shimojo, S. (2000). Illusions. What you see is what you hear. *Nature*, 408(6814):788.
- Sharma, S. (2005). *Principles Of Growth And Development*. Gyan Publishing House.
- Sheridan, M., Sharma, A., and Cockerill, H. (2007). *From birth to five years: children's developmental progress*. Routledge.
- Shirai, Y. and Inoue, H. (1973). Guiding a robot by visual feedback in assembling tasks. *Pattern recognition*, 5(2):99–108.
- Smart, W. D. and Pack Kaelbling, L. (2002). Effective reinforcement learning for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'02)*, volume 4, pp. 3404–3410. IEEE.
- Smith, L. and Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1-2):13–29.
- Smith, L. B. and Thelen, E. (2003). Development as a dynamic system. *Trends in cognitive sciences*, 7(8):343–348.
- Sofge, D., Bugajska, M., Trafton, J. G., Perzanowski, D., Thomas, S., Skubic, M., Blisard, S., Cassimatis, N., Brock, D., and Adams, W. (2005). Collaborating with humanoid robots in space. *International Journal of Humanoid Robotics*, 2(02):181–201.
- Steels, L. (1993). The artificial life roots of artificial intelligence. *Artificial Life*, 1(1):75–110.
- Sternberg, R. J. (2000). The concept of intelligence. In *Handbook of intelligence*, pp. 3–15. Cambridge University Press.

- Strack, F., Martin, L. L., and Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: a nonobtrusive test of the facial feedback hypothesis. *Journal of personality and social psychology*, 54(5):768.
- Sturm, J., Plagemann, C., and Burgard, W. (2009). Body schema learning for robotic manipulators from visual self-perception. *Journal of Physiology*, 103(3):220–231.
- Takanishi, A., Hirano, S., and Sato, K. (1998). Development of an anthropomorphic head-eye system for a humanoid robot-realization of human-like head-eye motion using eyelids adjusting to brightness. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 2, pp. 1308–1314. IEEE.
- Thagard, P. (2005). *Mind: Introduction to cognitive science*. Bradford Books.
- Thelen, E., Corbetta, D., Kamm, K., Spencer, J. P., Schneider, K., and Zernicke, R. F. (1993). The transition to reaching: mapping intention and intrinsic dynamics. *Child development*, 64(4):1058–98.
- Thelen, E. and Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. MIT Press.
- Thomson, D. M. and Tulving, E. (1970). Associative encoding and retrieval processes in episodic memory. *Psychological Review*, 80:352–373.
- Tikhanoff, V., Cangelosi, A., Fitzpatrick, P., Metta, G., Natale, L., and Nori, F. (2008). An open-source simulator for cognitive robotics research: the prototype of the iCub humanoid robot simulator. In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, pp. 57–61. ACM.
- Tikhanoff, V., Cangelosi, A., and Metta, G. (2011). Integration of Speech and Action in Humanoid Robots: iCub Simulation Experiments. *IEEE Transactions on Autonomous Mental Development*, 3(1):17–29.
- Togelius, J. (2003). *Evolution of the layers in a subsumption architecture robot controller*. University of Sussex.
- Tondu, B., Ippolito, S., Guiochet, J., and Daidie, A. (2005). A seven-degrees-of-freedom robot-arm driven by pneumatic artificial muscles for humanoid robots. *The International Journal of Robotics Research*, 24(4):257–274.
- Tonko, M., Schurmann, J., Schafer, K., and Nagel, H. H. (1997). Visually servoed gripping of a used car battery. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'97)*, volume 1, pp. 49–54. IEEE.

-
- Tresilian, J. R., Mon-Williams, M., and Kelly, B. M. (1999). Increasing Confidence in Vergence as a Cue to Distance. *Proceedings: Biological Sciences*, 266(1414):39–44.
- Tronick, E. (1972). Stimulus control and the growth of the infant's effective visual field. *Perception & Psychophysics*, 11(5):373–376.
- Trullier, O., Wiener, S. I., Berthoz, A., and Meyer, J. A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51(5):483–544.
- Tuller, B., Turvey, M. T., and Fitch, H. L. (1982). The Bernstein perspective: II. The concept of muscle linkage or coordinative structure. In Kelso, J. S., editor, *Human motor behavior: An introduction*, pp. 253–270. Psychology Press.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236):433–460.
- Ursino, M., Zavaglia, M., Magosso, E., Serino, A., and di Pellegrino, G. (2007). A Neural Network Model of Multisensory Representation of Peripersonal Space: Effect of tool use. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pp. 2735–2739. IEEE.
- Vahrenkamp, N., Wieland, S., Azad, P., Gonzalez, D., Asfour, T., and Dillmann, R. (2008). Visual servoing for humanoid grasping and manipulation tasks. In *Proceedings of the 8th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pp. 406–412. IEEE.
- Varela, F. J., Thompson, E. T., and Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. The MIT Press.
- Vereijken, B., van Emmerik, R. E. A., Whiting, H. T. A., and Newell, K. M. (1992). Free(z)ing Degrees of Freedom in Skill Acquisition. *Journal of Motor Behavior*, 24(1):133–142.
- Vernon, D., Metta, G., and Sandini, G. (2007). The iCub Cognitive Humanoid Robot: An Open-System Research Platform for Enactive Cognition. In *50 years of artificial intelligence*, pp. 358–369. Springer, Berlin.
- von Hofsten, C. (1976). The role of convergence in visual space perception. *Vision Research*, 16(2):193–198.

- von Hofsten, C. (1982). Eye-hand coordination in the newborn. *Developmental Psychology*, 18(3):450–461.
- Webb, B. (1995). Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*, 16(2):117–134.
- Webb, B. (1998). Robots, crickets and ants: models of neural control of chemotaxis and phonotaxis. *Neural Networks*, 11(7):1479–1496.
- Webb, B. (2009). Animals Versus Animats: Or Why Not Model the Real Iguana? *Adaptive Behavior*, 17(4):269–286.
- White, B. L., Castle, P., and Held, R. (1964). Observations on the development of visually-directed reaching. *Child Development*, 35(2):349–364.
- Wilson, S. W. (1985). Knowledge growth in an artificial animal. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pp. 16–23, Hillsdale. Lawrence Erlbaum Associates.
- Wilson, S. W. (1991). The animat path to AI. In Meyer, J.A., Wilson, S., editor, *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pp. 15–21. The MIT Press.
- Wismeijer, D. A., van Ee, R., and Erkelens, C. J. (2008). Depth cues, rather than perceived depth, govern vergence. *Experimental Brain Research*, 184(1):61–70.
- Wunsch, P., Winkler, S., and Hirzinger, G. (1997). Real-Time pose estimation of 3D objects from camera images using neural networks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pp. 3232–3237. IEEE.
- Xie, M. (2003). *Fundamentals of robotics: linking perception to action*. World Scientific.
- Yamaguchi, J.-I., Takanishi, A., and Kato, I. (1993). Development of a biped walking robot compensating for three-axis moment by trunk motion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93)*, volume 1, pp. 561–566. IEEE.
- Yanco, H. A. and Drury, J. L. (2004). Classifying human-robot interaction: an updated taxonomy. In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, pp. 2841–2846. IEEE.

-
- Yoshimi, B. H. and Allen, P. K. (1994). Visual control of grasping and manipulation tasks. In *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI'94)*, pp. 575–582. IEEE.
- Yu, C. (2005). *GeoAgent-based knowledge systems*. ProQuest.
- Zhou, W. and King, W. M. (1998). Premotor commands encode monocular eye movements. *Nature*, 393(6686):692–695.
- Ziemke, T. (1998). Adaptive behavior in autonomous agents. *Presence*, 7(6):564–587.
- Ziemke, T. (2003). What's that thing called embodiment. In *Proceedings of the 25th Annual meeting of the Cognitive Science Society*, pp. 1305–1310, Mahwah, NJ. Lawrence Erlbaum.
- Zohary, E., Holmes, N. P., and Makin, T. R. (2007). Is That Near My Hand? Multisensory Representation of Peripersonal Space in Human Intraparietal Sulcus. *The Journal of neuroscience*, 27(4):731–740.

Bound copies of published papers

Active Vision and Depth Estimation Toward a Peripersonal Space Encoding in a Humanoid Robot

Salomon Ramirez-Contla Davide Marocco Angelo Cangelosi
School of Computing and Mathematics
University of Plymouth

1. Introduction

Peripersonal space is defined as the space around a person's body, which is the space that defines the region of interactions between an agent and its environment. Estimating the distances at which objects are, in other words delimiting reachable (peripersonal) and not reachable (extrapersonal) space, is very important in order to properly interact with the environment and learn from it. Interestingly, behavioural and neurophysiological studies suggest that the brain encodes the peripersonal space differently from the extrapersonal space and that the coding of the former is achieved through the integration of different modalities (Farné et al., 1998).

This abstract presents a preliminary result of an ongoing work for developing an agent capable of learning a peripersonal space representation, the integration of different modalities and the use of active vision for interacting with objects in the peripersonal space.

Depth estimation in biological agents can be of great importance for a good performance in crucial tasks such as reaching, grasping or avoiding obstacles (Mon-Williams and Dijkerman, 1999). From literature it's known that monocular vision provides indirect cues for depth perception: motion parallax, accommodation effort, casted shadows by near objects and contrast to name some. Nevertheless, those cues can only be used in certain circumstances and in most of the cases the use of monocular depth indicators requires complex processing on the acquired image. For this reason, processes and algorithms that extract depth information from vision are widely focused on stereo images. That is, images of the same scene taken from two slightly different position (Reichelt et al., 2010). Besides the possible algorithms that can be applied to stereo images, vergence, an additional proprioceptive information is available to organisms endowed with two movable eyes. Vergence is the oculomotor adjustment needed to foveate the same point in space with both eyes. Recent studies show that in humans, vergence occurs well before the actual depth estimation (Wismeijer et al., 2008) and therefore it can be an important cue even in the absence of complex monoc-

ular cues or processing of stereo images. In this part of the work we study the possible relevance of vergence in the development of a peripersonal space representation.

2. Material and Methods

The experiment was carried out with a simulated version of an iCub humanoid platform (Tikhanoff et al., 2008). The task was to reach a red cube placed in front of the robot with the right hand. Only 5 DoF of the iCub arm were used. Two different conditions of the task were considered: using monocular or binocular vision. Tracking and foveating the object was achieved by an closed-loop pre-programmed controller that moves the head and the eyes of the robot so to locate the target's centroid in the centre of the right eye image, or in both eyes' images, respectively. In the binocular vision case, modulation of vergence was required.

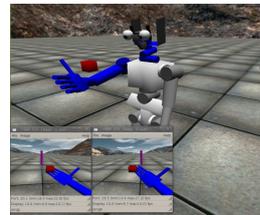


Figure 1: The simulated iCub performing the reaching task. Note the red target object and the blue painted arm.

A neural controller was used for moving the right arm of the simulated robot. Input for this controller was the proprioceptive information (pan and tilt joint positions from the head, and tilt, pan and vergence joints from the eyes) and pre-processed visual information. Pre-processing was colour based image segmentation: red was used in the case of the target object and blue for the arm of the robot (see figure 1). Image segmentation provided then two sets of data that were fed into the neural controller.

The controller was a feed forward, partially connected neural network with the following architecture: one input layer; one hidden layer *hA* which re-

ceives connections from visual input; an additional hidden layer hB which received connections from proprioceptive input and hidden layer hA ; an output layer which receives connections from the proprioceptive input and the two hidden layers hA and hB . This architecture was devised in order to analyse *unimodal* and *bimodal* contributions to depth perception, as described in (Farné et al., 1998), for the reaching task in a later stage of the study. The neural controller was trained using backpropagation (Rumelhart and McClelland, 1986). Training data consisted of 120 input/output pairs. Inputs corresponded to visual and proprioceptive data and desired outputs corresponded to a set of arm joints positions. For collecting the training data set the robot was pre-programmed to perform a reaching and grasping action followed by motor babbling. Half of the data was obtained using monocular vision and the rest using binocular vision because the controller was expected to generalise and perform well in both conditions.

3. Results

Data for this preliminary test was collected by placing the target object in 18 different positions. For each position the controller was activated in order to track and reach the object, starting from the arm in a home position along the body of the robot. Two test cases were used: *binocular* vision using vergence, and *monocular* vision both with the network continuously activated. Accuracy, direction of the arm respective to the target’s position and depth perception were measured for the 18 target positions. Accuracy was measured as the distance between the target and the palm of the hand. Direction was measured as the angle between a line from the head to the object and a line between the head and the hand in the horizontal plane. This measure gives an indication about the orientation of the arm with respect to the object. Depth perception was measured as the difference between the distance from the head and the target and the distance between the head and the hand. Results show that the robot is able to generalise well

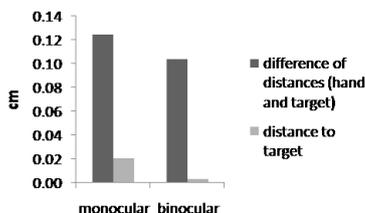


Figure 2: Comparison of the two visual systems’ performance. Distances are in cm.

from the training set and shows smooth transitions from different positions when continuously activated.

Analysis of the data shows that the effectiveness of the system to move the hand to the required direction in order to reach the target is very similar for both cases. However, we were more interested on depth accuracy rather than direction. On this matter of the data shows tendencies of binocular data being more accurate when reaching the target and generally better for depth estimation as it can be seen in figure 2.

4. Conclusion

Experiments indicated the effectiveness of the use of vergence for depth estimation in a reaching task in a simple active vision system implemented on the iCub simulator. These results make the use of these kind of depth estimation system suitable for our further development of peripersonal space representation. Additional studies will be carried out for investigations of the contribution of the different modalities used (proprioception and vision) as well as the implementation of a peripersonal space encoding that utilises future versions of this multimodal system.

References

- Farné, A., Ládavas, E., Zeloni, G., and Pellegrino, G. (1998). Neuropsychological evidence of an integrated visuotactile representation of peripersonal space in humans. *Journal of Cognitive Neuroscience*, 10(5):581–589.
- Mon-Williams, M. and Dijkerman, H. (1999). The use of vergence information in the programming of prehension. *Experimental Brain Research*, 128(4):578–582.
- Reichelt, S., Häussler, R., Fütterer, G., and Leister, N. (2010). Depth cues in human visual perception and their realization in 3d displays. volume 7690, page 76900B. SPIE.
- Rumelhart, D. and McClelland, J. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. M.I.T. Press, Cambridge.
- Tikhanoff, V., Cangelosi, A., Fitzpatrick, P., Metta, G., Natale, L., and Nori, F. (2008). *Proceedings of IEEE Workshop on Performance Metrics for Intelligent Systems Workshop (PerMIS’08)*, chapter An open-source simulator for cognitive robotics research: The prototype of the iCub humanoid robot simulator. Washington, D.C.
- Wismeijer, D., van Ee, R., and Erkelens, C. (2008). Depth cues, rather than perceived depth, govern vergence. *Experimental Brain Research*, 184(1):61–70.

Developing Motor Skills for Reaching by Progressively Unlocking Degrees of Freedom on the iCub Humanoid Robot

Salomón Ramírez-Contla¹, Angelo Cangelosi¹, Davide Marocco¹

¹University of Plymouth, Adaptive Behaviour and Cognition Lab
Drake Circus, Plymouth, England

salomon.ramirez-contla@plymouth.ac.uk, angelo.cangelosi@plymouth.ac.uk,
davide-marocco@plymouth.ac.uk

Abstract

To explore development of motor skills for reaching in the iCub robot, we test the capabilities for a neural network controller to learn progressively by locking some degrees of freedom (DOF) of the robot's arm before allowing it to explore the space with more DOF's. We consider exploration and bio-inspired mechanisms can aid in the development of control of the iCub robot arm. Results suggest the advantage of progressive development over an initial full training, also, these pointed out the importance of interaction with the world and the necessity of trial and error occurring in a time lapse for developing of reaching skills. **Index Terms:** degrees of freedom, motor skills, development, epigenetic robotics

1. Introduction

Proposed by Bernstein, the degrees of freedom problem [1] poses difficulties for autonomous skills learning and has drawn attention recently in the psychology field [2, 3]. Recent research on robotics [5, 4] has addressed this problem as well and tried to implement some of the ideas proposed by Bernstein due to the nature of recent advances in robotics and the need of developing controllers for redundant robot arms, specially of those of humanoids. Current cognitive robotics research has focused on the importance of the embodiment of an agent in order to richly interact within a world plenty of stimuli and cues that can aid in processes and reduce workload for a central controller such as the brain. The body plays an important role for this interaction and roboticists constantly look for new and better ways to control it.

Studies with evolutionary robotics approaches have been carried out with success for reaching and manipulation tasks. Massera et al. [6, 7] successfully evolved networks capable of fine-grained interaction with objects by exploiting the morphological constraints of a robotic arm. In this work, however, we are interested on the epigenetic development of such tasks.

Development of the human body flows from the top and centre of the body to the limbs. The spinal cord is the starting point, arms, legs, hands and toes take longer to develop. It is said that it follows a proximo-distal and cephalo-caudal direction [8] and this can be appreciated in infants: younger infants move their limbs in broad uncontrolled movements because only the most proximal joints of the limbs have been developed, like the shoulders. Later on, it can be seen that the elbow and wrist also come into play. Also, control of the lower part of the body comes after that of the upper part. In experimental psychology and motor development of humans there is evidence indicating that for learning new skills, adults freeze

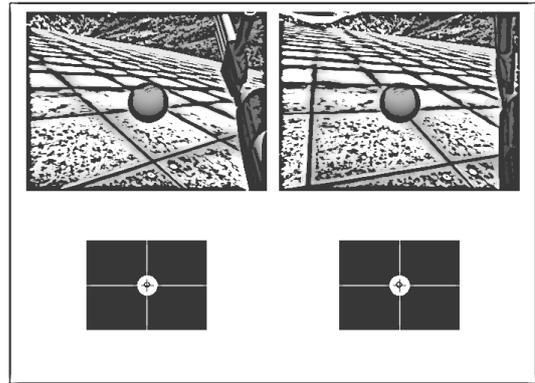


Figure 1: *Images from the robot once the controller has foveated the target. Above, the original images. Below, the low-res colour-segmented images.*

some of the distal joints involved in the new task until some degree of performance has been achieved, then some more degrees of freedom are used for achieving better performance [9, 10]. In the present work we test if that interaction with the world along with experience limited by constraints imposed by the physical characteristics of the arm, can help the learning process if this is segmented. We use a simulated iCub robot with neural controllers for the arm.

2. Methods

For testing the hypothesis, experiments were planned and carried out on the iCub robot simulator [11]. The iCub robot [12] is a humanoid robot about the size of a four years old child with 53 degrees of freedom designed for cognitive development research. The iCub's head subsystem consists of six degrees of freedom and is capable of vergence (the oculomotor adjustment needed to foveate the same point in space with both eyes). Three degrees of freedom in the head (tilt, pan and eyes' vergence) and four degrees on the arm (two from the shoulder, two from the elbow) were used. The robot head was provided with a visual tracking controller that locates and gazes at a specific target. For the experiments the target was a red ball. The gazing controller performed colour segmenting for the target's colour on the images coming from both eyes. This processing allowed it to track the centre of the target and adjust the position of three joints in the head in order to have the target in the centre of each eyes' field of view (Fig 1). By this mechanism, the robot

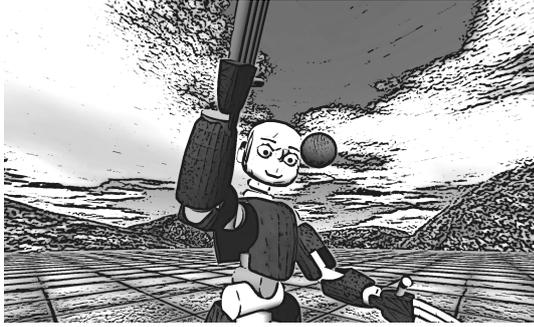


Figure 2: *The iCub performing the reaching task once it has foveated.*

gets information about the depth or distance at which the target is and together with the tilt and pan joints positions, it encodes the positions of the target in space. We use vergence as a depth measure following recent findings [13, 14] that indicate that vergence is in fact one strong signal for depth estimation and programming of prehension movements of humans.

Three different learning conditions (with three networks each) were tested on the robot to test our hypothesis. The first two were: staged learning or development, involving learning control of two DOF's and then the two other (DEV condition), and learning the head-hand associations involving four DOF's from the beginning (NO-DEV condition). For the last condition (NO-TRAIN), a group of three randomly initialized networks were created. These did not go through any learning process and are the control group.

With the help of the gazing mechanism a dataset was captured consisting of joint values of the head and eyes and the joint values of an arm position suitable for locating the end effector (the hand) in the point where the target was. This process can be considered a tutoring stage where the ball was put in the hand every time the robot executed random babbling [15] with the arm, then the gaze controller moved the head for foveating the target. For the cases the head was not able to move to a position where the target could be gazed no data was captured. This train set is equivalent of one acquired by performing random babbling while foveating the target. Reduction of the time required by this process is of course reduced when this kind of tutoring is present, as it happens with infants helped by parents when they start trying to reach objects that are usually out of reach or the baby simply fails to reach.

The controller for the robot arm was a feed-forward network with three inputs (one for each joint of the head controlled by the gazing controller), forty hidden units and four outputs, each of these output units controlled one joint of the arm. Two of these joints are in the shoulder and two in the elbow of the robot. During the initial training set creation, random babbling only occurred for the two most proximal joints of the arm, that is, for the ones in the shoulder. The other two joints were kept in constant values, in positions that we considered natural for an extended, similar to those when performing reaching for objects not very close to the body. Therefore, the positions that can be reached after the initial training are determined by the physical characteristics of the arm and by the generalization capabilities of the network. All learning for the networks was with the back-propagation of error algorithm using a learning rate of 0.01 and a momentum of 0.1.

For the development condition (DEV), each of the three net-

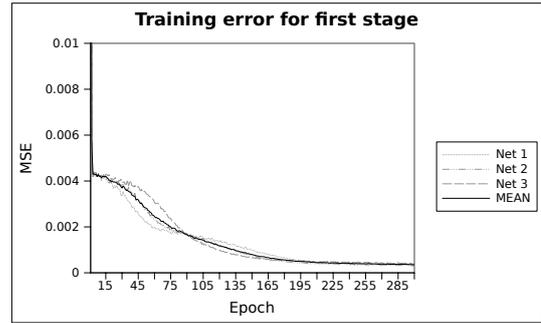


Figure 3: *Error during training of the first stage of development.*

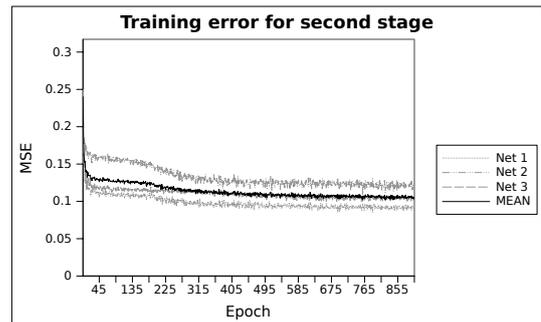


Figure 4: *Error during training of the second stage of development.*

works was trained using the set acquired via tutoring until the mean square error (MSE) became stable. For the three networks this was around the three hundred epochs. Figure 3 shows the training error for this stage of learning. After the initial training for reaching using two degrees of freedom, a test phase was carried out in ecological conditions. For an extended period of time, the robot was presented with the target in different locations, each time, the robot gazing mechanism was used for gazing the target, then the arm neural controller was activated with the inputs coming from the position of the head and eyes. When the robot successfully reached the target, that is, it touched it, the arm went to its resting position and the next test target position was presented. Otherwise, the two degrees of freedom that were initially locked (remember their values were constant for the first phase of learning) were randomly moved while the two most proximal degrees of freedom were kept constant with the values the neural controller produced. With this movement the robot was sometimes able to reach the target. When that was the case, the position that enabled it to achieve reaching was stored in a new set that was used for later training. This phase will be called from now on "experience phase". Figure 4 shows the training error during the second phase of learning for the three networks and the mean of the three of them.

The training using the new set generated in the experience consisted of 900 epochs. Figure 4 shows the error during training of this second stage of the learning.

The second condition (NO-DEV) consisted of using the training set generated during the experience phase on randomly initialised networks without going through an initial, partial, learning phase nor an experience phase. That is, these controllers were trained with the set that uses four degrees of freedom from the beginning. The training was for nine hundred

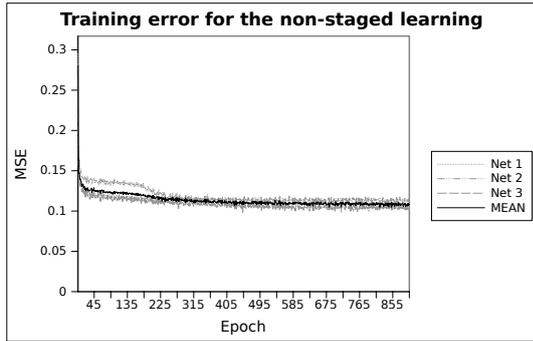


Figure 5: Error during training of the non-developed controllers.

epochs, at that point the MSE was stable. Figure 5 shows the error for this no-development learning.

Measurements for comparing the two conditions were performed during the execution of a reaching task similar to the task executed in the experience phase. Figure 2 shows the iCub executing the task once it has foveated the target. Final distance from hand to target was saved for each of the trails of the three controllers. Also, the number of times the controller successfully reached was recorded for having a percentage of success for each of the networks.

3. Results and discussion

Analysis on the output data indicates the controllers belonging to the staged or developmental training performed better in terms of final distance to the target as well as in the percentage of success (Figs 6 and 7). An analysis of variance test was performed to check for statistical difference between conditions (including the non-trained condition). This test reported statistical difference: current effect $F_{2,897}=850.45, p=0.0000$.

This can be due to various factors: following a developmental training, consisting of tutoring, experience during operation in its environment and learning based on that experience could have shaped the weights of the controller's networks to a stage that was able to find a solution for the second training set. Even when the training error of the final training in both conditions is very similar, in test conditions an advantage of the developed can be appreciated.

Because reaching is an important step in the development of motor and cognitive skills, it is also a skill explored to get an insight of the series of processes emerging in infants [16]. Our work on development of reaching tries also to consider the fact that for acquiring a skill it is necessary to have trial-and-error processes where time constraints cannot be avoided. In our experiments, the generation of the second training set for the staged learning condition, the "experience phase", took considerably longer than any other part of the experiment. But we believe this was a very important step due to the fact that each network will generate different outputs for the same inputs so the set is particular to each of them.

We have tried to implement what Bernstein [1] suggested for simplifying the degrees-of-freedom problem: in our experiments the robot arm controller goes through a developmental progression in order to find a first but simpler solution to the problem and later on, increasing the complexity of the problem. Other roboticists have implemented similar ideas in recent years. Ivanchenko and Jacobs [17] simulated a three degrees of

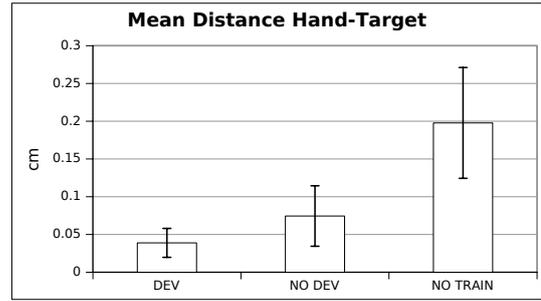


Figure 6: Mean distance from the center of the palm to the hand to the center of the target. Each controller was tested with 60 different target locations, none of them belonging to any of the sets used for training. Bars indicate standard deviation for all tests on each condition.

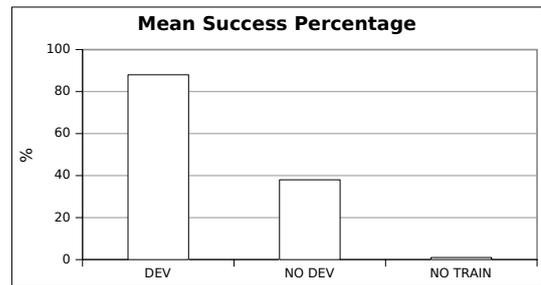


Figure 7: Results show the percentage of times the end effector touched the target. Each controller was tested with 60 different target locations, none of them belonging to any of the sets used for training.

freedom robot arm that tries to learn the dynamics of the arm while moving on trajectories on two dimensional space. The difference with our approach is that in our case the architecture of the networks is the same for every condition, it is the presence or absence of experience what shapes the performance at their final stages. Ivanchenko has a special architecture, devised from the idea that this decouples dynamic interactions among the joints and therefore allows to separately train the joints. Unlike Ivanchenko, for our experiments we decided to keep the same architecture. We want to explore uncoupling of dynamics without changing the internal (not directly exposed to the environment) characteristics of the system. In Ivanchenko's, results indicate that a developmentally trained controllers only outperformed the non-developmentally ones when the developmental path matched the nature of the task executed. In the case of our experiments, training as well as the "experience" phase matched the final task. This could explain the obtained results.

As Massera et al. in [6], we have started this exploration on a robotic arm with just four degrees of freedom. Our approach contrasts with that one in that we are interested on the epigenetic development of the skill instead of an evolutionary one. Moreover, in our case, experiments look to include vision into the development of the task instead of direct pass of coordinates or distances to the system without visual processing, as we consider that working towards the implementation of this type of skill development will require real-life sensory capabilities. The head controller for our experiments employs vision as a simple processing but action-involving task. Schlesinger et al. [18]

have also explored with the freezing of DOF's but again, using a non-realistic vision mechanism and a 2D environment and using evolutionary algorithms. Our work has pushed this type of exploration to a more realistic environment and explores the interaction on fixed architecture systems. We showed that even in this circumstances, a developmental approach can lead to better performance. Using the iCub simulator has proven to be a good test-bed for this type of research, as it allowed to implement and test controllers and visual sensors and explore performance in a controlled environment and free of mechanical strain issues.

3.1. Future work

In this study we have investigated the advantages of a progressive unlocking of joints to achieve better reaching performance. Our system has used two and then four degrees of freedom to explore and then improve a motor skill. However, limbs of natural systems, such a humans, display the property of overcompleteness. Overcompleteness implies that even though only 4 degrees of freedom are required for navigating a limb through three dimensional space [19], limbs on many vertebrates usually exhibit more than 4 degrees of freedom. This property turns the problem of controlling a limb more complex in computational terms (at least for traditional control) but also can represent and advantage in terms of the possibility of finding solutions that allow to reach a target at the same time that an obstacle is avoided. This could keep a relation with the representation of the reachable space. Also, constraints in other sensory or mechanical parts will be explored in further work.

4. Acknowledgements

This research has been funded by the National Council on Science and Technology of Mexico (CONACyT) and EU FP7 Project ITALK.

5. References

- [1] Bernstein, N., "The coordination and regulation of movements", Pergamon, 1967.
- [2] Karl M. Newell, David E. Vaillancourt, "Dimensional change in motor learning", *Human Movement Science*, 20(4-5), pp. 695-715, 2001.
- [3] Rohde, M., and Di Paolo, E., "t for Two: Linear Synergy Advances the Evolution of Directional Pointing Behaviour", *Proc. Advances in Artificial Life: 8th European Conference, ECAL 2005*, pp.262-271, Springer, 2005.
- [4] Berthouze, L., and Lungarella, M., "Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom" *Adaptive Behavior*, 12(1), pp. 47-64, 2004.
- [5] Gómez, G., Lungarella, M., Eggenberger Hotz, P., Matsushita, K., and Pfeifer, R., "Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity", *Proc. of 4th Int. Workshop on Epigenetic Robotics*, pp. 119-122, 2004.
- [6] G. Massera, A. Cangelosi, S. Nolfi, "Developing a Reaching Behaviour in a Simulated Anthromorphic Robotic Arm Through an Evolutionary Technique", in *Artificial Life X: Proceeding of the Tenth International Conference on the simulation and synthesis of living systems* MIT Press, 2006.
- [7] Massera G., Cangelosi A., Nolfi S., "Evolution of prehension ability in an anthropomorphic neurobotic arm", *Frontiers in Neuro-robotics*, 1(4), 2007.
- [8] Sharma, S.K., "Principles Of Growth And Development", Isha Books, 2005.
- [9] Vereijken, B., Van Emmerik, R., Whiting, H.T.A. and Newell, K., "Free(z)ing degrees of freedom in skill acquisition", *Journal of Motor Behavior* 24, pp. 133-142, 1992.
- [10] McDonald, P. V., Van Emmerik, R. E., Newell, K. M., "The effects of practice on limb kinematics in a throwing task", *Journal of Motor Behavior*, 21(3), pp. 245-264, 1989.
- [11] Tikhonoff, V., Cangelosi A., Fitzpatrick, P., Metta, G., Natale, L. and Nori, F., "An open-source simulator for cognitive robotics research: The prototype of the iCub humanoid robot simulator", in *Proc. IEEE Workshop Perform. Metrics Intell. Syst.*, Washington, DC, 2008.
- [12] Metta G., Natale L., Nori F., Sandini G., Vernon D., Fadiga L., von Hofsten C., Rosander K., Santos-Victor J., Bernardino A., Montesano L., "The iCub Humanoid Robot: An Open-Systems Platform for Research in Cognitive Development", *Neural Networks, special issue on Social Cognition: From Babies to Robots*. 23(8-9), 2010.
- [13] Mon-Williams, M. and Dijkerman, H., "The use of vergence information in the programming of prehension", *Experimental Brain Research*, 128(4), pp. 578-582, 1999.
- [14] Wismeijer, D., van Ee, R., and Erkelens, C., "Depth cues, rather than perceived depth, govern vergence" *Experimental Brain Research*, 184(1), pp. 61-70, 2008.
- [15] Von Hofsten, C., "Eye-hand coordination in the newborn", *Developmental Psychology*, 18(3), pp. 450-461, 1982.
- [16] Corbetta, D., Thelen, E. and Johnson, K., "Motor constraints on the development of perception-action matching in infant reaching", *Infant Behavior and Development*, 23, pp. 351-374, 2000.
- [17] Ivanchenko, V. and Jacobs, R. A., "A developmental approach aids motor learning", *Neural Computation*, 15(9), Pages 2051-2065, 2003.
- [18] Schlesinger, M., Parisi, D. and Langer, J., "Learning to Reach by Constraining the Movement Search Space", *Developmental Science*, 3, pp. 67-80, 2000.
- [19] Shadmehr, R. and Wise, S. P., "The Computational Neurobiology Of Reaching And Pointing: A Foundation for Motor Learning", *Computational Neuroscience*, MIT Press, 2005.

Adaptive Reachability Assessment in the Humanoid Robot iCub

Salomón Ramírez-Contla

School of Computing and
Mathematics

University of Plymouth

Drake Circus, Plymouth, UK

Email: salomon.ramirez-contla@plymouth.ac.uk

Davide Marocco

School of Computing and
Mathematics

University of Plymouth

Drake Circus, Plymouth, UK

Email: davide.marocco@plymouth.ac.uk

Abstract—We present a model for reachability assessment implemented in a simulated iCub humanoid robot. The robot uses a neural network both for estimating reachability and as a controller for the arm. During training, multi-modality information including vision and proprioception of the effector’s length was provided, along with tactile and postural information. The task was to assess if a target in view was at reach range. After training with data from two different effector’s lengths, the system generalised also for a third one, both for producing reaching postures and for assessing reachability. We present preliminary results that show good reachability predictions with a decrease in confidence that display a depth gradient.

I. INTRODUCTION

Being able to tell what is immediately reachable and what is out of reach is a skill that humans develop early in infancy [1]. The near-space or peripersonal space is the region surrounding the body where reaching is possible without translation. Body and brain must work together to create a flexible representation of this space in order to properly interact with objects. Investigating how this representation is created and modified is the object of recent research in psychology [2], neuroscience [3] and robotics [4].

Robots in unstructured environments will need to create and adapt their space constraints dynamically in order to safely and meaningfully interact with objects, other robots and/or humans. Acquiring images from the environment is far from being the totality of what a system needs for operating in these environments. A recent approach to perception tells us that action and perception work together [5]. In the present work we investigate how a robot could develop a representation of its reachable space by means of incorporating visual, postural and non-explicit body-shape information coming from action as components in the learning process.

II. MATERIALS AND METHODS

For our experiment we used the iCub robot simulator [6]. For investigating how the shape of the body can modify the peripersonal space of the robot, modifications in the robot’s arm were used. Figure 1 shows two of the three end effector lengths: short/no-tool and long/tool. These were used during the learning process. A third one, with a length between these two was used also during tests. A priori studies show that

reaching behaviour emerges at around 3 to 4 months of age [7] and we intend to investigate how from this early behaviour a peripersonal space is created and modified.

The robot was provided with a gazing behaviour by means of a head/eyes active vision controller. Images are colour-segmented and used to modulate the velocity of two joints of the neck and the eyes’ vergence angle. Movement leading to foveation in both eyes is produced by this controller when a target is presented. Correct vergence response to static targets is present in babies as early as the 1-2 months old [8]. In our model we consider vergence as an early-stage estimator for depth. Our visuo-motor system does not need prior camera calibration nor disparity calculation for producing correct vergence movements. Still, it is a reliable way for depth estimation easily implemented in the simulated and real robots. We believe it is important to investigate this kind of strategies that an artificial system can exploit for developing a way for discerning reachable from unreachable in a dynamic way to some extent independent of the morphology and visual capabilities.

A neural network was used by the robot to generate correct arm postures for touching a foveated target. The same network also assessed if a target was reachable or not. A feed-forward network with 3 layers was used. Inputs to it were: one proprioceptive unit for arm length/tool-use and three for tilt, pan and vergence angles. The hidden layer consisted of 10 units and the output was four units for arm joint angles and two for a 1-of-n encoding for predicted reachability. Hidden and output units used a sigmoid activation function.

During an ”experience” stage, the robot performed random movements after foveating a target set in many different locations. Random babbling allowed the robot to explore the space around it. If the robot touched the target, the current

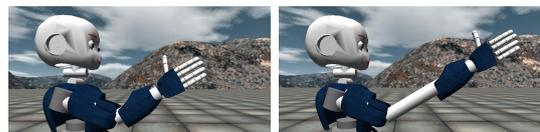


Fig. 1. End effector shapes used for training: short/no-tool and long/tool.

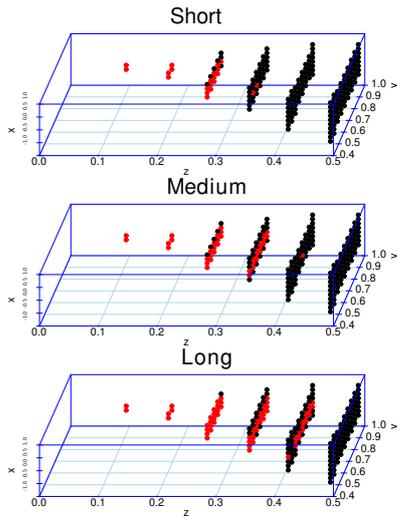


Fig. 2. A “view” of the robot’s peripersonal space with three different effector lengths. Red points indicate the robot assessed the point as reachable.

posture was added to a set of training data associated to a touch signal equal to 1 and the head/eyes angles. Otherwise, after a certain time, the robot would go to a “neutral” posture and the touch signal was set to 0 before appending to the training set. The learning process used backpropagation of error algorithm. The network learned a mapping of head and eyes postures to the arm posture for reaching it.

III. RESULTS AND CONCLUSIONS

In the experiment, the robot was presented with targets in different locations. After foveating the target, it used the network to determine if the target was reachable or not. Reachability prediction, using 1-of-n encoding allowed to measure confidence as the absolute difference of these two outputs. The confidence displayed a depth-gradient: as peripersonal space is extended, the confidence decreased. This can be due to depth estimation errors from the vergence signal. Confidence levels in fig. 4 and fig. 2 show, in red, the points in space the robot determined it could reach for three effector lengths.

Results suggest that this type of neural controller, trained with data coming from visuo-motor interactions, allowed the robot to create a non-explicit representation of the space close to its body. This representation was modulated by proprioceptive body shape information. Even without being explicitly encoded in a particular frame of reference, arm length and depth information from vergence allowed the robot to, in an action-based approach (“neutral” postures for not reachable target vs. reaching positions when in reach) perceive the space around it. As seen in fig. 3 vergence elicited response was modulated for each of the three different effector lengths.

We believe that explicit representations of the space around the body are not necessary, instead, experiencing moving in space allows or constraints the robot for acting on certain elements of the world that come into attention by other means of perception, namely, vision and touch.

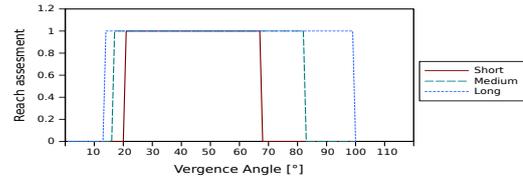


Fig. 3. Vergence angles at which the robot considered the target as reachable for three different arm lengths. The training was done only for the short and long ones.

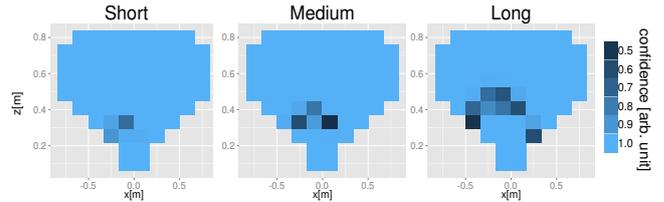


Fig. 4. Top view of working space that the robot was able to foveate. Average confidence in reachability assessment of target locations is shown. Confidence decreases in a depth/vergence gradient.

As further work, more experiments, currently in progress, will extend the understanding of how this simple embodiment like body shape properties along with signals like vergence elicit the response of the system. Comparisons against a similar system but without using posture in the training signal, and with different body related input will also shed light on this matter.

ACKNOWLEDGEMENTS

This research has been funded by the National Council of Science and Technology of México (CONACyT).

REFERENCES

- [1] J. Piaget, “The origins of intelligence in children,” *Journal of Consulting Psychology*, vol. 17, no. 6, p. 467, 1953.
- [2] M. Costantini, E. Ambrosini, C. Scorolli, and A. M. Borghi, “When objects are close to me: Affordances in the peripersonal space.” *Psychonomic bulletin & review*, Feb. 2011.
- [3] M. Ursino, M. Zavaglia, E. Magosso, A. Serino, and G. di Pellegrino, “A Neural Network Model of Multisensory Representation of Peripersonal Space: Effect of tool use,” in *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*. IEEE, 2007, pp. 2735–2739.
- [4] E. Chinellato, M. Antonelli, B. J. Grzyb, and A. P. del Pobil, “Implicit Sensorimotor Mapping of the Peripersonal Space by Gazing and Reaching,” *IEEE Transactions on Autonomous Mental Development*, vol. 3, no. 1, pp. 43–53, Mar. 2011.
- [5] J. K. O’Regan and A. Noë, “A sensorimotor account of vision and visual consciousness,” *Behavioral and brain sciences*, vol. 24, no. 5, pp. 939–972, 2001.
- [6] V. Tikhonoff, A. Cangelosi, P. Fitzpatrick, G. Metta, L. Natale, and F. Nori, “An open-source simulator for cognitive robotics research: the prototype of the iCub humanoid robot simulator,” in *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*. ACM, 2008, pp. 57–61.
- [7] E. Thelen, D. Corbetta, K. Kamm, J. P. Spencer, K. Schneider, and R. F. Zernicke, “The transition to reaching: mapping intention and intrinsic dynamics,” *Child development*, vol. 64, no. 4, pp. 1058–98, Aug. 1993.
- [8] L. Hainline and P. M. Riddell, “Binocular alignment and vergence in early infancy,” *Vision research*, vol. 35, no. 23-24, pp. 3229–36, Dec. 1995.