Università della Svizzera italiana

Faculty of Informatics

Research Prospectus

May, 2012

Towards Adaptive and Autonomous Robots

Developing Visually Guided Object Manipulation on the iCub Humanoid

Jürgen Leitner

Abstract

Although robotics has seen advances over the last decades they are still not in wide-spread use outside industrial applications. Proposed scenarios for robots range from cleaning tasks, grocery shopping to elderly care, helping in a hospital, etc. These involve the robots working together, helping and coexisting with humans in daily life. From this the need to deal with a more unstructured environment arises. Object manipulation, which is of high importance in these scenarios, is still a hard problem in robotics. Humans, in contrast, are able to quickly, without much thought, perform a variety of object manipulation tasks on arbitrary objects.

The goal of this research is to overcome the limitations of current robotic object manipulation in unstructured environments. A developmental, step-wise approach is used to generate adaptive and autonomous grasping behaviours for novel objects on the *iCub* humanoid robot. To facilitate this a combined integration of computer vision, artificial intelligence and machine learning is employed.

Research Advisor	Research Co-advisor
Prof. Jürgen Schmidhuber	Dr. Alexander Förster
Academic Advisor	Review Committee

Research Advisor's approval (Prof. Jürgen Schmidhuber): Date: ....

PhD Director's approval (Prof. Antonio Carzaniga): Date: ....

1 Introduction

The overarching goal of the research proposed herein is to extend the capabilities of robotic systems to provide more autonomous and adaptive behaviours. In the last century robots have transitioned from science fiction to science fact. After centuries of imagining automated machines that would help humans, the last few decades brought an ever growing number of these programmable machines into existence. The field of robotics developed from interdisciplinary research in the fields of electronics, mechanics and computer science, with the first prominent robots appearing in academia during the 1960s, such as the famous Shakey at SRI. In the 1980s the digital revolution, yielding programmable personal computes and embedded systems, kick-started the creation of robotic systems to be used in large scale industry and automated production. Nowadays robotic arms are fulfilling automation tasks in factories all over the world. In recent years the field is moving towards other areas, such as household settings. Proposed applications range from cleaning tasks, grocery shopping to elderly care, helping in a hospital, etc. These involve the robots working together, helping and coexisting with humans in daily life.

To decide and act in these environments, the robot needs to perceive, as it is not possible to provide the robot with all the information it needs a priori. This might be because of hardware limitations, such as limited storage space, or simply because the information is plain not available or known (the most striking example for this is robotic space exploration). Therefore the need for more autonomous, more intelligent robots arises, to extend robotic applications. Sensory feedback is of importance to allow the robot to perceive the environment. This information though might be incomplete or inconsistent. This necessitates means of managing uncertainty, which has been extensively explored in Artificial Intelligence (AI) and Machine Learning (ML) research. A wide range of sensors have been used to build models of the environment the robot is in. Visual feedback, though it tends to be harder to interpret than other sensory information, is actively researched. One motivation is that our world is built around (human) visual perception, which also means that to allow 'natural' interaction of robots it needs to be able to understand its environment based on the camera images it receives.

Another aspect of these 'natural' interactions is the human capability of adaptation. This adaptation allows to reach, grasp and manipulate a wide range of objects in arbitrary positions. In order to work naturally in human environments such as offices and homes, robots will need to be much more flexible and robust in the face of uncertainty and novelty. During my research I will investigate AI and ML techniques to allow for better 'perception', 'motion' and 'coordination'. The aim is to generate reaching and grasping capabilities for general manipulation of objects perceived.

2 Background and Related Research

Manipulation of objects is still a very hard problem in robotics, yet it is one of the most important skills for robots [29]. It allows to interact with the world and become useful and helpful to humans. Although AI has developed algorithms to play chess on a level good enough to beat the average human player, the robotic manipulation of a chess piece, in contrast, is still not feasible at a human (children) level.

Object manipulation tasks are based on the precise and coordinated movements of the arm and hand. The understanding of how humans and animals control these movements is a fundamental research topic in cognitive and neurosciences. Despite the interest and importance of the topic, e.g. in rehabilitation medicine and robotics, the issues and theories behind how humans learn, adapt and perform reaching and grasping behaviours remain controversial. The development of reaching and grasping in humans is still not fully understood. Although there are many experimental studies on how reaching and grasping develops in humans no computational models of the process are yet available. Vision is seen as important to develop reaching and grasping skills [4, 39, 53], e.g. imitation of object manipulation has been reported for already 14-month-old infants [41]. In comparison most robots can currently at most perform simple grasping of certain objects.

Manipulating arbitrary objects is not a trivial thing. The development of hand control in children, for an apparently simple, prototypical grasping task is not matured until the age of 8-10 years [17]. Moreover, complexity, as can be seen by the number of neurons comprising the control of the arm and hand, is staggeringly high. Even after manipulation skills have been learnt they are constantly adapted by an perception-action loop to yield the desired result. To build a robot that is able to perform complex manipulation skills that can help users in their activities of daily living (ADL) is the aim of various on-going EU projects [11, 3, 19, 22].

2.1 Robot Learning: Approaches to Autonomy and Adaptation

In recent years the interest of the robotics community into AI and ML as tools to build *intelligent robots* has been increasing. At the same time, the learning community has seen a rise in using robots as motivating applications for new algorithms and techniques.

Already from the start of AI in the 1950s, researchers have been striving to create intelligent machines. Over the years AI research has developed tool and techniques that are able to solve complex problems, and notable successes in specialised domains, such as game playing. However, the vision of creating general-purpose, human-like intelligence has not yet been achieved. In the 1960s the first prominent robot in academia was developed at the Artificial Intelligence Center of the Stanford Research Institute (now called SRI International), named *Shakey* [48, 49]. It was built on the AI principles of building a 'General Problem Solver' and was the first mobile robot performing and reasoning about its own actions. A robot is the ideal test bench for many of AI and Machine Learning (ML) techniques. There is though the need to combine and integrate these specialised subfields. It uses these techniques, and tests them, in a physical entity and real environments to interact/move in. This 'embodiment'[1] is seen as an important condition for the development of cognitive abilities both in humans and robots [8, 64, 56].

How to design intelligent robots, that integrate and apply those techniques is though difficult and an emerging field of research as can be seen by the number of workshops and symposia on the topic. These problems arise from the combination of complexity, uncertainty, measurement errors and interaction with the physical world, while performing real-time processing. The recent advent of advanced, complex and humanoid robots, such as Honda's Asimo [26], NASA and GM's Robonaut2 [5, 15], Toyota's Partner robots [60], and the European *iCub* [61, 43], increases the need of AI and ML techniques to control the high physical degrees of freedom.

2.2 Cognitive, Developmental and Evolutionary Robotics

Cognitive Robotics, *Developmental Robotics* and *Evolutionary Robotics* are field that recently emerged to investigate how robots can 'learn' for themselves to generate more autonomous and adaptive capabilities. Especially in highly complex, high DOF robotic systems engineering solutions in all details increasingly difficult and cumbersome.

In *Cognitive Robotics* [2] the aim is to provide robots with cognitive processes, similar to humans and animals. An integrated view of the body, including the motor system, the perceptual system, the body's interactions with the environment (situatedness) is taken. The acquisition of knowledge, may it be through actions (e.g. motor babbling) or perception is a big part of cognitive robotics research. Various cognitive architectures have been proposed for robots [9, 58, 62, 12, 65], they have yet to yield the promised improvements in robotic applications. This can be attributed to the varying definitions of cognition and the highly complex human cognitive system whose workings are still not understood. Generally there are two approaches to building cognitive architectures, a top-down approach, which seems to have been the main focus of research and a bottom-up, which seems more fit for robotic purposes. An overview of cognitive architectures and the aim for an iCub cognitive architecture is presented in [62].

Developmental Robotics [2, 37, 63, 40, 31, 1] is aiming to put more emphasis on the development of skills. It is an interdisciplinary approach to developmental science. It differs from the previous approaches, as the engineer only creates the architecture and then allows the robot to explore and learn its own representation of its capabilities (sensory and motor) and the environment. Fundamental for the development of skills is the body and its interactions with the environment. Aims are to build adaptive robotic systems by exploration and autonomous learning, that is, as learning without a direct intervention from a human designer is a dedicated goal [37]. It tries to develop this by building on previous knowledge, while seeking out novel, interesting areas to explore.

Evolutionary Robotics [13, 25, 50, 16] is another approach to add adaptiveness and developmental processes to robots. It emerged as a new approach to overcome the difficulties of designing control systems for autonomous robots: (a) coordinating the (increasing) number of DOF both in mechanics and control is hard, especially since the complexity scales with the number of possible interactions between parts (or modules) [13] (see also the 'Curse of Dimensionality') (b) the environment and how the robot interacts with it are often not known (perfectly) beforehand. Its main focus is on using artificial neural networks (ANN) as the control system to be evolved. These neuro-controllers (NC) are inspired by the neuron activity in the human brain. NCs have been shown to work for diverse robotic applications, e.g. in mobile robots [51] and space applications [14, 33]. One of the issues with evolving such controllers is that it generally takes a larger number of iterations (or generations) to 'learn' certain behaviours. This works fine in simulation but is hard to achieve on a real robotic platform. Nolfi *et al.* showed that evolving a NC on hardware is, though time consuming, feasible [51]. Only a very simple mobile robot was used and it was stated that a hybrid approach, where the controllers are trained first in simulation and then then transferred to the real robotic system, seems preferential. An example of this was published by *Miglino et al.*: after evolving only in simulation the NC was transferred directly onto the robot. Although the trajectories differ significantly in the real setting their fitness correlations were high [44]. How to effectively train NCs (in simulation) and transfer them to real, high-DOF hardware is still an open research question.

Other Approaches to robot learning have been developed in the past. Recently the area of Reinforcement Learning (RL) [59] has appealed to many roboticists, especially for complex robotic systems. Here only a general RL algorithm

[1] i.e., the claim that having a body that mediates perception and affects behaviour plays an integral plays in the emergence of human cognition.

and the means to inform the robot whether its actions were successful (positive reward) or not (negative reward) are needed. RL and its applicability to humanoid robots has been investigated by Peters *et al.* [55]. Imitation Learning or Apprenticeship Learning is of importance in human skill development as it allows to transfer skills from one person to the next. In robotics Robot Learning from demonstration or Programming by Demonstration is a similar paradigm for enabling robots to autonomously perform new tasks. It takes the view that an appropriate robot controller can be derived from observations of a another agent's performance thereof [57].

3 Towards Vision-guided Object Manipulation on the *iCub*: A Research Plan

At IDSIA we are interested in learning object manipulation and how motivation effects the development of complex (inter)actions[2]. Combining the robot learning approaches described above with computer vision we aim to enable adaptive, autonomous grasping based on visual feedback. Grasping primitives, which can adapt based on the available sensory feedback about object locations, are therefore needed. The visual feedback enables to close the loop and perform grasping of objects while adapting to a dynamic environment. In my research I aim to combine low level robotic control with high level AI planning approaches to enable the robot to reason about how the overall task should influence manipulations of individual objects.

Our Experimental Platform: The *iCub* Humanoid

The *iCub* humanoid robot platform [61] is an open-system robotic platform developed within EU funded projects. It consists of two arms and a head attached to a torso roughly the size of a human child. The head and arms are designed anthropomorphic and provide a high degree-of-freedom (DOF) system that is used for researching human cognitive and sensorimotor development. The *iCub* is an excellent experimental platform for cognitive and sensorimotor development and embodied AI [42] and was designed for human-like object manipulation research. An interesting area of research on this robot is to test the scalability of learning methods towards these complex systems. These movements need to be coordinated with feedback from visual, tactile, and acoustic[3] perception.

3.1 Step 1: (Visual) Perception

The first step is to work on the robots's ability to perceive the environment: it needs to be able to to perceive, detect and locate objects in its environment. This spatial understanding is crucial for motion planning, obstacle avoidance and finally interacting with these environments and the objects therein. We are focussing on visual perception, which in humans and animals refers to the ability to interpret their environment based on the information available when visible lights reaches the eye. The reason for the focus on vision is twofold, firstly the limited sensing capabilities of the robotic platform and secondly, vision is the most important sense for humans, and using a humanoid robot we look at how humans do it.

Vision and the visual system are the focus of much research in psychology, cognitive science, neuroscience and biology. A major problem in visual perception is that what individuals 'see' is not just a simple translation of input stimuli (compare optical illusions). The research of Marr in the 1970s led to a theory of vision using different levels of abstraction. He described human vision as processing from a two-dimensional visual array (on the retina) to a three-dimensional description of the world as output. His stages are: a 2D (or primal) sketch of the scene (using feature extraction), a 2D sketch of the scene (using textures to provide more information), a 3D model [38]. In the last months I was working on a visual system for the *iCub* humanoid similar to the description of human perception by Marr. The framework developed at IDSIA [34], allows to filter out (known) objects in the 2D space of the images and can estimate the 3D location of the object from performing this filtering on both cameras.

Detection. Research on perception has been an active component of developing artificial vision (also known as computer vision or machine vision) systems, in industry and robotics. In the last few months I was working on a visual system for the *iCub* humanoid similar to the description of human perception by Marr. The first thing done by the human visual system, and investigated by us, is the segmentation (detection) in the visual space (the 2D images). There exists a vast body of work on all aspects of image processing [21], using both classical and machine learning approaches. Genetic programming [30] has often been used to solve problems in image processing, however previous attempts typically use a small set of mathematical functions to evolve kernels, or a small number of basic image processing functions (such as erode and dilate). Given the maturity of the field of image processing, it should be possible to construct programs that use much more complicated image operations and hence incorporate domain knowledge. We use a technique based on Cartesian Genetic Programming (CGP) [45, 46, 23], that allows

[2]This research is funded by the European Union grant 'IM-CLeVeR' [3].

[3]For grasping acoustic feedback (i.e. noise) might be used as feedback whether a grasp was performed successfully or not.

for the automatic generation of computer programs using a large subset of the functionality of the OpenCV image processing library [7]. It provides an effective method to learn new object detection algorithms (filters). An added feature is that if the training set is chosen correctly different lighting conditions are no problem for detecting.

Localisation. To enable the reaching behaviour visual object localisation in the world is of importance. Developing an approach to perform robust localisation to be deployed on a real humanoid robot is necessary to provide the necessary inputs for on-line motion planning & reaching, and object manipulation tasks. To localise objects in the environment the *iCub* has to rely solely, similarly to human perception, on a visual system based on stereo vision. The two cameras are mounted in the head of the robot. Their pan and tilt can jointly be controlled, with vergence providing a third DOF. To look around 3 DOF in the neck can be used to move the head. The *iCub*'s cameras are mounted in the head of the robot therefore the method for localisation must be able to cope with motion to work *while* the robot is controlling its gaze and upper body for reaching. The most logical choice, from a planning and control standpoint, of coordinate system is the world reference frame, originating from the robots hip. In order to transform coordinates from the coordinate frame of the eyes to world coordinates an accurate kinematic model of the robot is necessary. More formally, the fundamental matrix will vary as a function of pan and vergence of the eyes, and the position and orientation of the stereo camera unit (the head) will vary as a function of the state of the torso and neck.

Stereo Vision describes the extraction of 3D information out of digital images and is similar to the biological process of stereopsis in humans. Its basic principle is the comparison of images taken of the same scene from different viewpoints. To obtain a distance measure the relative displacement of a pixel between the two images is used [24]. While these approaches, based on projective geometry, have been proven effective under carefully controlled experimental circumstances, they are not easily transferred to robotics applications.

For the *iCub* platform several approaches have previously been developed. One of these methods is a biologically inspired approach that mimics the retina of the human eye. Another, the 'Cartesian controller module', provides basic 3D position estimation functionality [54] and gaze control. This module works well on the simulated iCub, however it is not fully supported and functional on the hardware platform, and therefore does not perform well. As part of my thesis work I am investigating how spatial representation can be learned from visual feedback [35, 36].

3.2 Step 2: Motion

To control the motion of the robot is the next step. As our platform is a highly complex system, learning and the approaches described above are deemed to be the most suitable. Our approach is starting with low complexity control, such as, controlling the robot's gaze and will slowly include more DOF until the full 7DOF arms are controlled for reaching. This has been extensively investigated previously, e.g. using a combination of open and closed loop control [47, 52] and motor babbling [10].

When attempting to synthesise behaviour on a complex robot like the *iCub*, the shortcomings of state-of-the-art learning and control theories can be discovered and addressed in subsequent research. We hope to integrate suitable learning methods for controlling the (arm of the) robot with existing software and research, such as, MoBeE [18].

3.3 Step 3: Coordination and Adaptation

The final step is to closely integrate the sensor and motor sides and have an action-perception loop. To approach to the goal of general object manipulation, where a robot can autonomously manipulate any given object within its workspace, it would ideally encode and reuse knowledge in terms of task features that are invariant to the object. Confronted with a novel instance of a specific task the robot needs to establish appropriate correspondences between objects and actions in its repertoire and their counterparts in the current task.

Sensorimotor Development is aiming to learn basic hand-eye coordination. Langdon & Nordin have shown this on a simple humanoid robot using GP techniques [32]. Various methods for learning this sensorimotor mapping have been investigated [27, 28]. These lead to biologically inspired mappings, yet applying these directly to control the robot is still an issue. *Adaptation* is needed for precise object manipulation, as no highly accurate model of the world and also of the robotic system can not be assumed in most interesting scenarios. An estimation of the robot kinematics might help in generating more precise motions. So far no module exists to estimate the kinematics of the *iCub*, this is partly due to the openly available CAD models and thorough calibration procedures that need to/should be applied regularly. For other robotic platforms machine learning has been used to estimate the kinematic model, e.g. Gloye et al. used visual feedback to learn the model of a holonomic wheeled robot [20] and Bongard et al. used low-dimensional sensory feedback to learn the model of a legged robot [6].

Our aim is to combine these previous approaches with our machine/robot learning frameworks. Possible cooperation (for research in this direction on the *iCub*) were discussed with the Italian Institute of Technology ('home of the *iCub*') (G. Metta) and Institute for Systems and Robotics (at IST in Lisbon) (A. Bernardino).

References

[1] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida. Cognitive developmental robotics: A survey. *IEEE Transactions on Autonomous Mental Development*, 1:12–34, 2009.

[2] M. Asada, K. MacDorman, H. Ishiguro, and Y. Kuniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37(2):185–193, 2001.

[3] G. Baldassare et al. The IM-CLeVeR project: Intrinsically motivated cumulative learning versatile robots. In *Proc. of the International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 2009.

[4] N. Berthier, R. Clifton, V. Gullapalli, D. McCall, and D. Robin. Visual information and object size in the control of reaching. *Journal of Motor Behavior*, 28(3):187–197, 1996.

[5] W. Bluethmann, R. Ambrose, M. Diftler, S. Askew, E. Huber, M. Goza, F. Rehnmark, C. Lovchik, and D. Magruder. Robonaut: A robot designed to work with humans in space. *Autonomous Robots*, 14(2):179–197, 2003.

[6] J. Bongard and V. Zykov. Resilient machines through continuous self-modeling. *Science*, 314(5802):1118–1121, 2006.

[7] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000.

[8] R. Brooks. *Cambrian intelligence: the early history of the new AI*. The MIT Press, 1999.

[9] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376–386, 2005.

[10] D. Caligiore, T. Ferrauto, D. Parisi, N. Accornero, M. Capozza, and G. Baldassarre. Using motor babbling and hebb rules for modeling the development of reaching with obstacles and grasping. In *International Conference on Cognitive Systems*, 2008.

[11] A. Cangelosi et al. The italk project: Integration and transfer of action and language knowledge in robots. In *Proc. of the International Conference on Human Robot Interaction (HRI)*, 2008.

[12] A. Chella, M. Frixione, and S. Gaglio. A cognitive architecture for robot self-consciousness. *Artificial Intelligence in Medicine*, 44(2):147–154, 2008.

[13] D. Cliff, P. Husbands, and I. Harvey. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):73–110, 1993.

[14] B. Dachwald. Optimization of interplanetary solar sailcraft trajectories using evolutionary neurocontrol. *Journal of Guidance, Control, and Dynamics*, 27(1):66–72, 2004.

[15] M. Diftler, J. Mehling, M. Abdallah, N. Radford, L. Bridgwater, A. Sanders, R. Askew, D. Linn, J. Yamokoski, F. Permenter, et al. Robonaut 2-the first humanoid robot in space. In *Proc. of the International Conference on Robotics and Automation (ICRA)*, pages 2178–2183, 2011.

[16] S. Doncieux, J. Mouret, N. Bredeche, and V. Padois. Evolutionary robotics: Exploring new horizons. *New Horizons in Evolutionary Robotics*, pages 3–25, 2011.

[17] H. Forssberg, A. Eliasson, H. Kinoshita, R. Johansson, and G. Westling. Development of human precision grip i: basic coordination of force. *Experimental Brain Research*, 85(2):451–457, 1991.

[18] M. Frank, J. Leitner, M. Stollenga, S. Harding, A. Förster, and J. Schmidhuber. The modular behavioral environment for humanoids and other robots (mobee). In *International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, 2012.

[19] GeRT Consortium. EU Project: Generalizing Robot Manipulation Tasks. `http://www.gert-project.eu/project/the-vision/`, Feb. 2012.

[20] A. Gloye, F. Wiesel, O. Tenchio, and M. Simon. Reinforcing the Driving Quality of Soccer Playing Robots by Anticipation. *IT - Information Technology*, 47(5), 2005.

[21] R. Gonzalez and E. W. Richard. Digital image processing, 2002.

[22] Handle Consortium. EU Project: HANDLE. http://www.handle-project.eu/, Feb. 2012.

[23] S. Harding, J. Leitner, and J. Schmidhuber. Cartesian genetic programming for image processing. In *Genetic Programming Theory and Practice X*. Springer, 2012.

[24] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2000.

[25] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, 20(2):205–224, 1997.

[26] R. Hirose and T. Takenaka. Development of the humanoid robot asimo. *Honda R&D Technical Review*, 13(1):1–6, 2001.

[27] H. Hoffmann, W. Schenck, and R. Möller. Learning visuomotor transformations for gaze-control and grasping. *Biological Cybernetics*, 93:119–130, 2005.

[28] M. Hülse, S. McBride, J. Law, and M. Lee. Integration of active vision and reaching from a developmental robotics perspective. *IEEE Transactions on Autonomous Mental Development*, 2(4):355–367, 2010.

[29] C. Kemp, A. Edsinger, and E. Torres-Jara. Challenges for robot manipulation in human environments [grand challenges of robotics]. *IEEE Robotics & Automation Magazine*, 14(1):20–29, 2007.

[30] J. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, 1992.

[31] B. Kuipers, P. Beeson, J. Modayil, and J. Provost. Bootstrap learning of foundational representations. *Connection Science*, 18(2):145–158, 2006.

[32] W. B. Langdon and P. Nordin. Evolving Hand-Eye Coordination for a Humanoid Robot with Machine Code Genetic Programming. In *Proc. of the European Conference on Genetic Programming (EuroGP)*, 2001.

[33] J. Leitner, C. Ampatzis, and D. Izzo. Evolving ANNs for spacecraft rendezvous and docking. In *Proc. of the International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2010.

[34] J. Leitner, S. Harding, M. Frank, A. Förster, and J. Schmidhuber. icVision: A Modular Vision System for Cognitive Robotics Research. In *Proc. of the International Conference on Cognitive Systems*, 2012.

[35] J. Leitner, S. Harding, M. Frank, A. Förster, and J. Schmidhuber. Learning spatial object localisation from vision on a humanoid robot. *International Journal of Advanced Robotics Systems*, 2012. submitted.

[36] J. Leitner, S. Harding, M. Frank, A. Förster, and J. Schmidhuber. Transferring spatial perception between robots operating in a shared workspace. IROS, 2012. submitted.

[37] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2003.

[38] D. Marr. *Vision: A Computational Approach*. Freeman & Co., San Francisco, 1982.

[39] M. McCarty, R. Clifton, D. Ashmead, P. Lee, and N. Goubet. How infants use vision for grasping objects. *Child development*, 72(4):973–987, 2001.

[40] L. Meeden and D. Blank. Introduction to developmental robotics. *Connection Science*, 18(2):93–96, 2006.

[41] A. Meltzoff. Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, 24(4):470, 1988.

[42] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, and L. Montesano. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks*, 23(8-9):1125–1134, Oct. 2010.

[43] G. Metta, G. Sandini, D. Vernon, L. Natale, and F. Nori. The icub humanoid robot: an open platform for research in embodied cognition. In *Proc. of the Workshop on Performance Metrics for Intelligent Systems*, pages 50–56, 2008.

[44] O. Miglino, K. Nafasi, and C. Taylor. Selection for wandering behavior in a small robot. *Artificial Life*, 2(1):101–116, 1994.

[45] J. Miller. An empirical study of the efficiency of learning boolean functions using a cartesian genetic programming approach. In *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1135–1142. Citeseer, 1999.

[46] J. Miller. Cartesian genetic programming. *Cartesian Genetic Programming*, pages 17–34, 2011.

[47] L. Natale, F. Nori, G. Sandini, and G. Metta. Learning precise 3d reaching in a humanoid robot. In *Proc. of the International Conference on Development and Learning (ICDL)*, pages 324–329, 2007.

[48] N. Nilsson. A mobile automaton: An application of artificial intelligence techniques. In *International Joint Conference on Artificial Intelligence*, 1969.

[49] N. Nilsson. Shakey the robot. Technical report, DTIC Document, 1984.

[50] S. Nolfi and D. Floreano. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. 2000.

[51] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In *Artificial Life IV*, pages 190–197. Cambridge, MA: MIT Press, 1994.

[52] F. Nori, L. Natale, G. Sandini, and G. Metta. Autonomous learning of 3d reaching in a humanoid robot. In *Proc. of the International Conference on Intelligent Robots and Systems (IROS)*, pages 1142–1147, 2007.

[53] E. Oztop, N. Bradley, and M. Arbib. Infant grasp learning: a computational model. *Experimental Brain Research*, 158(4):480–503, 2004.

[54] U. Pattacini. *Modular Cartesian Controllers for Humanoid Robots: Design and Implementation on the iCub*. PhD thesis, RBCS, Italian Institute of Technology, Genova, 2011.

[55] J. Peters, S. Vijayakumar, and S. Schaal. Reinforcement learning for humanoid robotics. In *Proc. of the International Conference on Humanoid Robots*, pages 1–20, 2003.

[56] R. Pfeifer, J. Bongard, and S. Grand. *How the body shapes the way we think: a new view of intelligence*. The MIT Press, 2007.

[57] S. Schaal. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.

[58] M. Shanahan. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition*, 15(2):433–449, 2006.

[59] R. Sutton and A. Barto. *Reinforcement learning: An introduction*, volume 1. Cambridge Univ Press, 1998.

[60] S. Takagi. Toyota partner robots. *Journal-Robotics Society of Japan*, 24(2):62, 2006.

[61] N. G. Tsagarakis et al. iCub: the design and realization of an open humanoid platform for cognitive and neuroscience research. *Advanced Robotics*, 21:1151–1175, Jan. 2007.

[62] D. Vernon, G. Metta, and G. Sandini. The icub cognitive architecture: Interactive development in a humanoid robot. In *Proc. of the International Conference on Development and Learning (ICDL)*, pages 122–127. Proc. of the, 2007.

[63] J. Weng. Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics*, 1(2):199–236, 2004.

[64] D. Wolpert, Z. Ghahramani, and J. Flanagan. Perspectives and problems in motor learning. *Trends in cognitive sciences*, 5(11):487–494, 2001.

[65] J. Wyatt and N. Hawes. Multiple workspaces as an architecture for cognition. In A. V. Samsonovich, editor, *AAAI Fall Symposium on Biologically Inspired Cognitive Architectures*. The AAAI Press, 2008.