VII

Artificial Intelligence Aspect of Cognitive Robotics

CHAPTER 10

A Bottom-Up Integration of Vision and Actions To Create Cognitive Humanoids

Jürgen Leitner

Dalle Molle Institute for AI (IDSIA) & SMRTRobots,
(juxi@smrtr.cc), Switzerland

CONTENTS

10.1	Introduction	192
	Artificial Intelligence, Machine Learning and Robotics	192
	Robot Learning	193
10.2	A Cognitive Robotics Approach	196
10.3	Perceiving the Environment	199
	10.3.1 Object Detection: icVision & CGP-IP	200
10.4	Interacting With the Environment	201
	10.4.1 Collision Avoidance and World Model: MoBeE	203
	10.4.2 Action Repertoire: TRM & LEOGrasper	204
10.5	Integration	206
	10.5.1 Closing the Action-Perception Loop	206
10.6	Results	207
	10.6.1 Avoiding a Moving Obstacle	207
	10.6.2 Reaching and Grasping Objects	207
10.7	Conclusions	209

In recent years more and more complex humanoid robots have been developed. On the other hand programming these systems has become more difficult. There is a clear need for such robots to be able to adapt and perform certain tasks autonomously, or even learn by themselves how to act. An important issue to tackle is the closing of the sensorimotor loop. Especially when talking about humanoids the tight integration of perception with actions will allow for improved behaviours, embedding adaptation on the lower-level of the system.

10.1 INTRODUCTION

Object manipulation in real-world settings is a very hard problem in robotics, yet it is one of the most important skills for robots to possess [29]. Through manipulation they are able to interact with the world and therefore become useful and helpful to humans. Yet to produce even the simplest human-like behaviours, a humanoid robot must be able to see, act, and react continuously. Even more so for object manipulation tasks, which require 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- [49] and neuro-sciences [27]. Despite the interest and importance of the topic, e.g. in rehabilitation and medicine, the issues and theories behind how humans learn, adapt and perform reaching and grasping behaviours remain controversial. Although there are many experimental studies on how humans perform these actions, the development of reaching and grasping is still not fully understood and only very basic computational models exist [45]. Vision is seen as an important factor in the development of reaching and grasping skills in humans [5, 39]. For example, imitation of simple manipulation skills has been observed already in 14-month-old infants [40]. Current robots in contrast are only able to perform (simple) grasps in very limited, specific settings. To enable more autonomous object manipulation, more specifically how to enable some level of eye-hand coordination to perform actions more successfully, is of high interest to the robotics community (see e.g. NASA's Space Technology Roadmap calls for "Real-time self-calibrating hand-eye System" [1]).

Artificial Intelligence, Machine Learning and Robotics

The research in the fields of Artificial Intelligence (AI) and robotics were strongly connected in the early days, but have diverged over the last decades. Although AI techniques were developed to play chess on a level good enough to win against (and/or tutor) the average human player [51], the robotic manipulation of a chess piece, in contrast, the creation of intelligent machines has lacked quite a bit behind the algorithmic side. It is still not feasible to control a robot on a similar level of precision, adaptation and success as a human — not even comparative to children level. To produce even the simplest autonomous, adaptive, human-like behaviours, a humanoid robot must be able to, at least:

- Identify and localize objects in the environment, e.g. the chess pieces and board
- Execute purposeful motions for interaction, e.g. move a piece to a desired position

At the beginning of AI research a clear goal was to build complete, intelligent, autonomous robotic system [50]. As with the example of the above example of chess, it has proven to be quite challenging. Not helping the cause was the fractioning of the fields into many distinct facets of research. While there was progress in each of the sub-fields and the both disciplines (AI and robotics) separately, it has now become clear that a closer integration is again needed. There has been a renewed interest, from both research communities, to work together again towards the goal of intelligent robotic systems.

The field of robotics has clearly matured over the last few years. Current humanoid robots are stunning feats of engineering as mention above. To embed this systems with some sense of 'intelligence' and use the full versatility of advanced robotic systems, a bigger collaboration with the research community in Artificial Intelligence and Machine Learning is required.

The idea of the 'embodied mind' stems from philosophy. It claims that the nature of the human mind is determined by the form of the human body. Philosophers, psychologists, cognitive scientists, and artificial intelligence researchers who study embodied cognition and the embodied mind argue that all aspects of cognition are shaped by aspects of the body. The embodied mind thesis is opposed to other theories of cognition. Embodied cognition reflects the argument that the motor system influences our cognition, just as the mind influences bodily actions. Roboticists have argued that to understand intelligence and build artificial system that comprise intelligence can only be achieved by machines that have both sensory and motor skills. Furthermore they need to be interacting with the world through a body. This 'embodiment' is seen as an important condition for the development of cognitive abilities both in humans and robots [9, 62, 48]. The insights of these robotics researchers have in return also influenced philosophers.

Machine Learning algorithms, have been applied in experimental robotics to acquire new skills, however the need for carefully gathered training data, clever initialization conditions, and/or demonstrated example behaviours limits the autonomy with which behaviours can be learned. To build robots that can perform complex manipulation skills that help users in their activities of daily living is the aim of various research projects in Europe (e.g. [11, 60]).

Robot Learning

As mentioned above the programming of these highly complex robot systems is a cumbersome, difficult and time-consuming process. Current approaches tend to describe each precise each precise movement in detail, allowing little to no flexibility or adaptation during execution. This obviously has issues with scaling to highly complex robots in complicates settings. Therefore the robotics community has focused on methods to provide robots with the ability to act autonomously, adapt or 'learn' how to behave without the need of hard-coding every possible outcome.

Autonomous robots research is aimed at building systems that do not require the pre-programming of every possible situation encountered. Many kinds of robots have some degree of autonomy and different robots can be autonomous in different ways. In fields, such as space exploration, a high degree of autonomy is desirable. For an autonomous robot one generally assumes the following capabilities [14]:

- Gain information about the environment (Rule #1)
- Work for an extended period without human intervention (Rule #2)
- Move either all or part of itself throughout its operating environment without human assistance (Rule #3)
- Avoid situations that are harmful to people, property, or itself unless those are part of its design specifications (Rule #4)
- Maintain its own survival at the expense of the previous rules (Sentient Robot Mandate) (Rule #5)
- Learn or gain new capabilities like adjusting strategies for accomplishing its task(s) or adapting to changing surroundings (Rule #6)

In the early 90s of the last century *Behavioural Robotics* (or behaviour-based robotics) was introduced as a way to deal with more and more complex robots and application areas [8]. This research area focuses on flexible switching mechanisms to change the robots main behaviours based only on a very simple internal model. The basic idea is that close (and probably simple) sensor-motor connections can result in behaviours that appear complex and sophisticated. Due to the fact that these models used a simple approach, rather than a computational complex model and the relatively low cost of development, popularised this approach in the mid-1990s. This paradigm has had a wide range of application in multi-robot teams [4] yet the scaling to complex robots, such as humanoid, has not been successful so far.

Robot Learning generally refers to research into ways for a robot to learn certain aspects by itself. Instead of providing all information to the robot a priori, for example, possible motions to reach a certain target position, the agent will through some process 'learn' which motor commands lead to what action. The research field is placed at the intersection of machine learning and robotics and studies how robots can acquire new skills through experimentation. The earlier mentioned 'embodiment' plays an important role here. Example include the learning of sensorimotor skills (for example locomotion, grasping, object manipulation), as well as interactive skills such as manipulation of an object in collaboration with a human. In addition the learning of linguistic skills, especially the grounding of words or phrases in the real world, is of interest to the research community. The field of 'robot learning' is closely related to other disciplines, for example, adaptive control. Learning in realistic environments requires algorithms that can deal with high-dimensional states, e.g. to detect events in the stream of sensory inputs, change and uncertainty. Note that while machine learning is nowadays often used for computer and robot vision tasks (like in this dissertation), this area of research are usually not referred to as 'robot learning'. The fields of *Cognitive Robotics*, *Developmental Robotics* and *Evolutionary Robotics* emerged with the specific aim to investigate how robots can 'learn' for themselves and thereby generate more autonomous and adaptive capabilities.

In *Cognitive Robotics* [3] the aim is to provide robots with cognitive processes, similar to humans and animals. An integrated view of the body is taken, including the motor system, the perceptual system and the body's interactions with the environment. The acquisition of knowledge, may it be through actions (e.g. motor babbling) or perception is a big part of cognitive robotics research. Another is the development of architectures for these tasks. A variety has been proposed [54, 58, 12, 63], but the promised improvements in robotic applications still need to be shown. This can be attributed to the varying definitions of cognition and the complex human cognitive system, whose workings are still not fully understood. To build cognitive architectures two distinct approaches have been tried. The research seems to mainly focus on top-down architectures. A bottom-up approach has been described as suitable for the use with robots (e.g. the proposed *iCub* cognitive architecture [59]).

Developmental Robotics [3, 61, 2] 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. As above, the body and its interactions with the environment are seen as being fundamental for the development of skills. Aims are to build adaptive robotic systems by exploration and autonomous learning, i.e. learning without a direct intervention from a designer [37]. Here interesting areas to explore are selected by building on previous knowledge, while seeking out novel stimuli.

Evolutionary Robotics [23, 43] 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 (see 'Curse of Dimensionality' [15]) (b) the environment and how the robot interacts with it are often not known before. Its main focus is on evolve a control system based on artificial neural networks. These neuro-controllers (NC), inspired by the neuron activity in the human brain, have been shown to work in a wide range of applications [44, 17, 31]. An important issue is that to 'learn' behaviours, a large number of iterations (or generations) is required. This works fine in simulation but is hard to achieve on a real robotic platform. [44] showed that evolving a NC on hardware is, while time consuming, feasible, at least for simple mobile robots. Hybrid approaches, where NCs are trained first in simulation and then transferred to the real hardware, seem preferential. The performance of the controllers in the real world can then be used to improve the simulation [6]. How to effectively train and apply NCs to real, high-DOF hardware is still an open research question.

Other Approaches to robot learning have been developed in the past. The area of Reinforcement Learning (RL) [56] has appealed to many roboticists, especially for learning to control complex robotic systems. A general RL algorithm and the means to inform the robot whether its actions were successful (positive reward) or not (negative reward) is all that is required. RL and its applicability to humanoid robots has been investigated by [47]. Imitation Learning or Apprenticeship Learning is of importance in human skill development as it allows to transfer skills from one person to another. In robotics Robot Learning from demonstration or Programming by Demonstration is a similar paradigm for enabling robots to learn to perform novel tasks. It takes the view that an appropriate robot controller can be derived from observations of a another agent's performance thereof [53].

10.2 A COGNITIVE ROBOTICS APPROACH

One of the most important problems in robotics currently is arguably to improve the robots' abilities to understand and interact with the environment around them: a robot needs to be able to perceive, detect and locate objects in its surrounding and then then have the ability to plan and execute actions to manipulate these objects detected.

The described approach herein was developed to extend the capabilities of the *iCub* humanoid robot, especially to allow for more autonomous and more adaptive – some would say, more 'intelligent' – behaviours. The *iCub* is a state-of-the-art, high degree-of-freedom (DOF) humanoid (see Figure 10.1) [57]. It consists of two arms and a head attached to a torso roughly the size of a human child. The head and arms follow an anthropomorphic design and provide a high DOF system that was designed to investigate human-like object manipulation. It provides also a tool to investigate human cognitive and sensorimotor development. To allow for safe and 'intelligent' behaviours the robot's movements need to be coordinated closely with feedback from its sensors. The *iCub* is an excellent experimental platform for cognitive, developmental robotics and embodied artificial intelligence [42].

The aim is to generate a not before seen level of eye-hand coordination on the *iCub*. Pick-and-place operations were chosen as they require intelligent behaviour in a complex environment, i.e. perceiving which objects are in its vicinity, reaching for a specific object, while avoiding obstacles. The cognitive

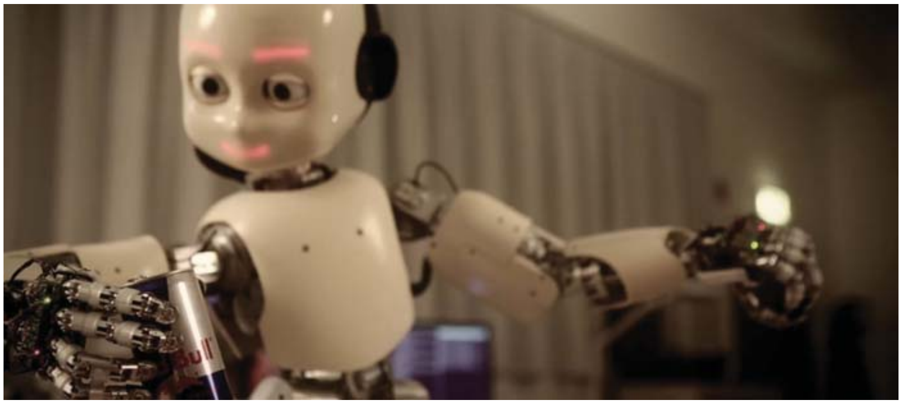

FIGURE 10.1 The experimental platform used: the *iCub* humanoid robot.

skills, from learning what an object is and how to detect it in the sensory stream, to adapting the reach if the environment changes, are embedded using a variety of frameworks. First functional motion and vision subsystems are developed, which are then integrated to create a closed action-perception loop. The vision side detects and localises the object continuously, while the motor-side tries to reach for target objects avoiding obstacles at the same time.

A combination of robot learning approaches with computer vision and actions is used to improve adaptivity and autonomy in robot grasping based on visual feedback. The next section contains the description of the robot vision frameworks and techniques developed for and implemented on the *iCub*, shown in the top row of Figure 10.2 (in green). It includes the modules for the detection and identification of objects (in the images), as well as, the localization (in 3D Cartesian space). The bottom half, in yellow, shows the action and motion side. To generate motion using machine learning techniques a crucial feature is avoiding collisions, both between the robot and the environment *and* the robot and itself.

The various modules developed and interacting in this cognitive approach are the following:

- *Object Models, Detection and Identification:* as mentioned above, the detection and identification of objects is a hard problem. To perform these tasks CGP-IP (Cartesian Genetic Programming for Image Processing) [22] is used. It provides a machine learning approach to building visual object models, which can be converted into executable code, in both supervised and unsupervised fashion [32]. The resulting program performs the segmentation of the camera images for a specific object.

- *Object Localisation:* by using the image coordinates of the detected object from the two cameras together with the current robot's pose, the posi-

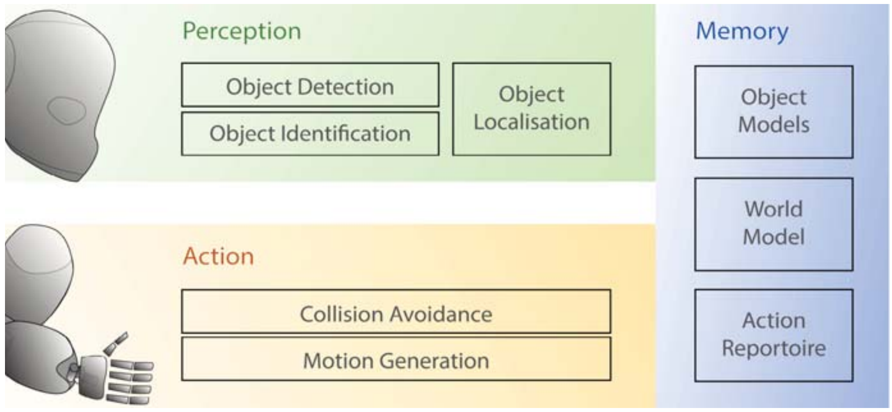

FIGURE 10.2 Overview of the proposed architecture for a functional eye-hand coordination on the *iCub* humanoid. The *object detection* and *identification* is currently solely based on the camera images (2D) received. The *object localization* uses the information from the two cameras to calculate an operational space (3D) position. This is the same space in which *collision avoidance* is applied and the *world* is modelled. The *object model* contains the information of how to detect the object in the 2D images (see Section 10.3). The *motion generation* and *action repertoire* can use the full configuration space of the humanoid (41 DOF).

tion of the object can be estimated in Cartesian space wrt. the robot's reference frame. Instead of a calibration for each single camera, the stereo system and the kinematic chain, a module that learns to predict these from a training set is incorporated. The system has been shown to estimate these positions with a technique based on genetic programming [35] and an artificial neural network estimators [34]. After the object is detected in the camera images the location of an object is estimated and the world model is updated.

- *Action Repertoire:* a light-weight, easy-to-use, one-shot grasping system (LEOGrasper[1]), which has been used extensively at IDSIA (Figure 10.6), provides the main grasping subsystem. It can be configured to perform a variety of grasps, all requiring to close the fingers in a coordinated fashion. A variety of more complex actions/roadmaps can be generated offline and later executed on the *iCub* [55], to e.g. lead to improved perception [36].

[1]Source code available at: `https://github.com/Juxi/iCub/`

- *World Model, Collision Avoidance and Motion Generation:* the world model keeps track of the robot's pose in space and the objects it has visually detected. Figure 10.8 shows this model including the robot's pose the static table, and the two objects localised from vision. MoBeE is used to safeguard the robot from (self-)collisions. It furthermore allows to generate motion by forcing the hand in operational space.

Section 10.4 describes the developed techniques to control the *iCub*. All these subsystems are supported by memory (in blue) enabling the persistent modelling of the world and providing a repertoire of actions to be triggered. In Section 10.5 the tight integration of these two sides, the perception and the motion, is described. Section 10.6 presents a proof-of-concept, highlighting a level of eye-hand coordination not previously seen on the *iCub*.

10.3 PERCEIVING THE ENVIRONMENT

To be useful in the above proposed scenarios a robot must be able to see, act, and react continuously. Perception is a key requirement in order to purposefully adapt robot motion to the environment, allowing for more successful, more autonomous interactions. The first important step towards this is to understand the environment the robot is embedded in. Coming back to the example of playing chess, this would compare to finding the chess board and each of the chess pieces (e.g. in a camera image) or even just to to realise that there is a chess board and pieces in the scene.

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*). One important area of research to build robots that can understand their surroundings is the development of artificial vision. *Computer Vision* – sometimes referred to as *Robot Vision* when applied in a robotic system – generally describes the field of research dealing with acquiring, processing, analysing, and understanding images in order to produce decisions based on the observation. The fields of computer vision and AI have close ties, e.g. autonomous planning or decision making for robots requires information about the environment, which could be provided by a computer vision system. AI and computer vision share other topics such as pattern recognition and learning techniques.

Even though no clear definition of the areas of computer vision and image processing exists, the latter is commonly used to refer to a subsection of computer vision. Image processing techniques generally provide ways of extracting information from the image data, for example, noise reduction, feature extraction, segmentation, etc.[21] Another important topic in computer vision is 'image understanding'. With the aid of geometry, statistics, and learning the goal is to mimic the abilities of the human (visual) perception system.

Research into vision for the special requirements of robotic systems is referred to as *robot vision* or *machine vision* [24, 25]. For example, visual feedback has extensively been used in mobile robot applications, for obstacle avoidance, mapping and localization. With the advancement of humanoids and the increased interest in working around humans, object detection and manipulation are more and more driving the development of robot vision systems. An important problem is that of determining whether or not the image data contains some specific object, feature, or activity. While this has been researched for quite some time already, the task seems harder than expected and no solution for the general case of detecting arbitrary objects in arbitrary situations exists. From a robot vision point of view, this means that the robot is required to detect previously unknown objects in its surroundings and be able to build models to memorise and identify them in the future. Most of the work is heavily relying on artificial landmarks and fiducial markers to simplify the detection problem. Furthermore existing methods can at best solve it for specific objects (simple geometries, faces, printed or hand-written characters, or vehicles) and in specific situations (in terms of well-defined illumination, background, and pose of the object wrt. the camera). For a detailed introduction and overview of the foundations and the current trends the reader is referred to the excellent survey by [30].

10.3.1 Object Detection: icVision & CGP-IP

Aiming at eye-hand coordination and object manipulation the following frameworks were implemented to enable the learning and real-time operation of object detection and identification.

icVision [33] is an open-source[2], biologically-inspired framework consisting of distributed YARP [41] modules performing computer vision related tasks in support of cognitive robotics research (Figure 10.3). It includes the modules for the detection and identification of objects (in the camera images, referred to as *Filters*), as well as the localisation of the objects in the robot's operational space (3D Cartesian space). At the centre is the *icVision* core module, which handles the connection with the hardware and provides housekeeping functionality (e.g., extra information about the modules started and methods to stop them). Currently available modules include object detection, 3D localisation, gazing control (attention mechanism) and saliency maps. Standardised interfaces allow for easy swapping and reuse of modules.

The main part in object detection, the binary segmentation of the object from the background (see Figure 10.3 on the right), in the visual space, is performed in separate *icVision* filter modules. Each one is trained using the Cartesian Genetic Programming for Image Processing (CGP-IP) framework [22], in combination with the OpenCV [7] library, to detect and identify specific objects in a variety of real-life situations (e.g. a tea box as shown in Figure 10.4). The framework can run multiple filter modules in parallel. A

[2]Code available at: `https://github.com/Juxi/icVision/`

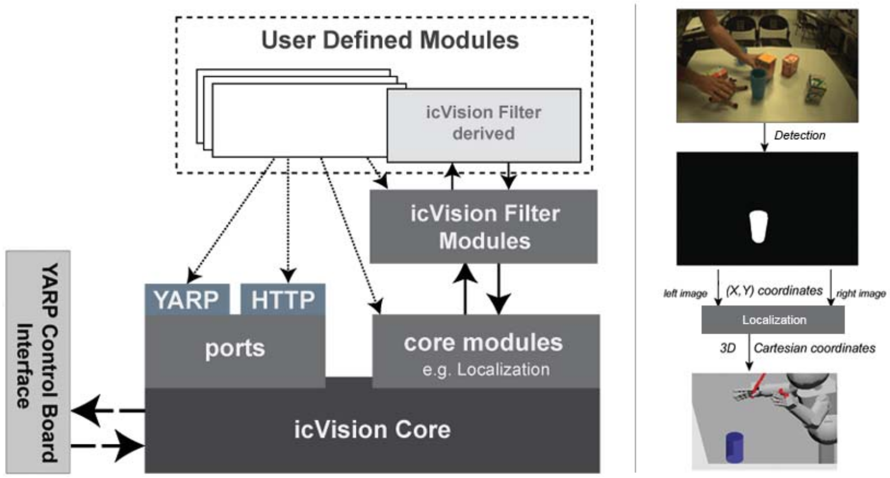

FIGURE 10.3 The icVision Architecture: The core module (bottom) is mainly for housekeeping and accessing & distributing the robot's camera images and motor positions. Object detection is performed in the *filter* modules (top), by segmenting the object of interest from the background. A typical work flow is shown (right): input images are retrieved from the cameras, the specific object is detected by a trained *Filter Module*, before the outputs (together with the robot's current pose) are used to estimate a 3D position using a localisation module. The object is then placed in the world model.

variety of filters have been learnt and most of them are able to perform the object detection in near real-time.

To interact with the objects the robot also needs to know where the object is located. 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. *icVision* provides modules to estimate the 3D position based on the robot's pose and the location of object in the camera images.

10.4 INTERACTING WITH THE ENVIRONMENT

Computer vision has become a more and more prominent topic of research over the past decades, also in the field of robotics. Like humans and animals, robots are able to interact with the world around them. While most robot vision research tends focus on understanding the world from just passive observations, these interactions with the environment provide and create valuable information to build better visual systems. Connecting manipulation commands with visual inputs allows for a robot to create methods to actively explore its surroundings. These connections between motor actions and observations exist in the human brain and are an important aspect of human development [5].

Only after the scene is observed and the robot has an idea about which objects are in the environment, can it start interacting with these in a safe fashion. In the chess example, even if the state of the board and where it is located are known, to move a certain chess piece from one field to another without toppling other pieces is still a hard problem by itself. In fact, children even at a very young age, have significantly better (smoother, more 'natural', 'fluent' and controlled) hand movements than all currently available humanoid robots. But manipulating arbitrary objects is not a trivial thing, even for humans. The development of hand control in children, for an apparently simple, prototypical precision grasp task is not matured until the age of 8-10 years [18]. 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 desired results. In infants various specializations in the visual pathways may develop for extracting and encoding information relevant for visual cognition, as well as, information about the lo-

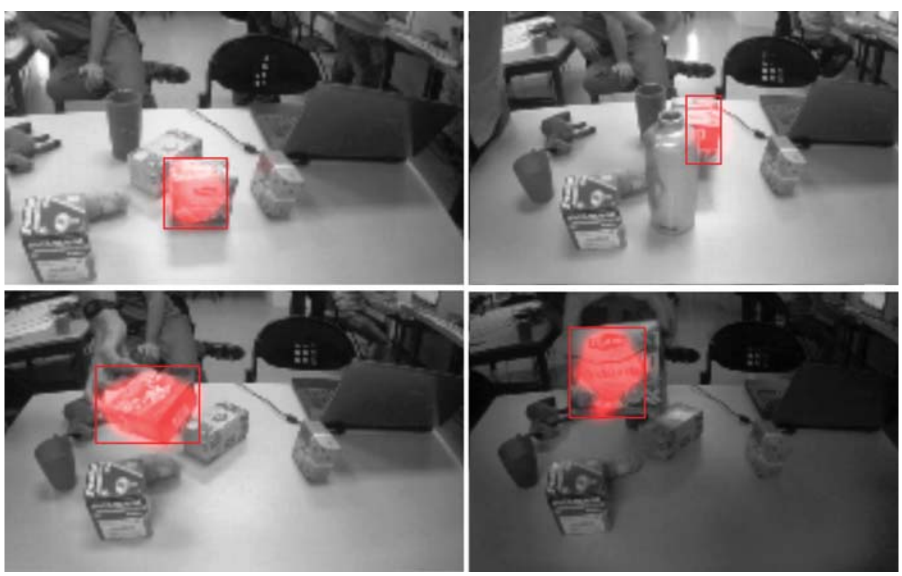

FIGURE 10.4 Detection of complex objects, e.g. a tea box, in changing poses, different light and when partially occluded is a hard problem in robot vision.

cation and graspability of objects [28]. This hints at the very close integration of vision and action in the human brain.

In recent years good progress was made with robotic grasping of objects. The various manipulators, mainly hands and grippers, and techniques clearly improved. Also novel concepts of 'grippers' have been designed and some are quite ingenious solutions to a number of issues. One such example is the granular gripper made by [10], which is made out of grounded coffee beans which are able to 'flow' around the object and then fixed in position by creating a vacuum. This concept has recently been extended to a full sized elephant-trunk-style arm [13]. Also in terms of how to grasp objects with regular grippers and 'hands' recent results highlight the advanced state of research in grasping. For example, [38], with their research showed that robots are able to pick up non-rigid objects, such as, towels. Their robot is able to reliably and robustly pick up a randomly dropped towel from a table by going through a sequence of vision-based re-grasps and manipulations-partially in the air, partially on the table. On the other hand, ways for a robot to learn from only a small number of real world examples, where good grasping points are on a wide variety of previously unknown objects have been presented [52].

The progress on performing grasping operations in the last years shows that one can now use these grasping subroutines and further integrate them in autonomous systems. The direct interface between various components, which makes robotics such a hard but interesting field, clearly needs to improve to allow for robust object manipulation. Only by combining sensing and control of the whole robotic platform a fully functional 'pick-and-place' capable system will appear. To allow for a variety of objects to be picked up from various positions the robot needs to see, act and react within a control system in which these elements are tightly integrated.

10.4.1 Collision Avoidance and World Model: MoBeE

An important issue is to ensure the safe operation of our humanoid. Modular Behavioral Environment (MoBeE) [19] is a software infrastructure to realise complex, autonomous, adaptive and foremost safe robot behaviours. It acts as an intermediary between three loosely coupled types of modules: the Sensor, the Agent and the Controller. These correspond to abstract solutions to problems in Computer Vision, Motion Planning, and Feedback Control, respectively. An overview of the system is depicted in Figure 10.5. The framework is robot independent, and can exploit any device that controlled via YARP [41]. It also supports multiple interacting robots, and behavioural components are portable and reusable thanks to their weak coupling. MoBeE controls the robot constantly, according to the following second order dynamical system:

$$M\ddot{q}(t) + C\dot{q}(t) + K(q(t) - q^*) = \sum f_i(t) \quad (1)$$

where $q(t) \in R^n$ is the vector function representing the robot's configuration, M, C, K are matrices containing mass, damping and spring constants

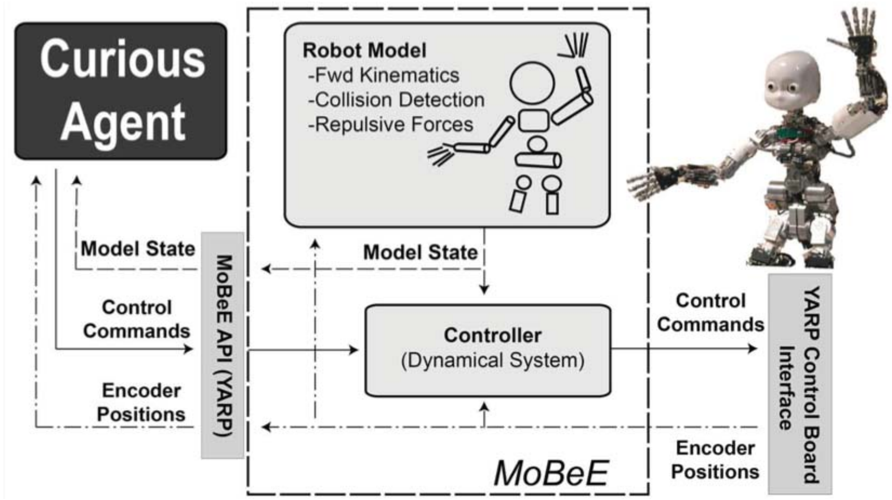

FIGURE 10.5 The Modular Behavioral Environment Architecture: MoBeE implements low-level control and enforces necessary constraints to keep the robot safe and operational in real-time. Agents (left) are able to send high-level commands, while a kinematic model (top) is driven by the stream of encoder positions (right). The model computes fictitious constraint forces, which repel the robot from collisions, joint limits, and other infeasibilities. These forces, $f_i(t)$, are passed to the controller (middle), which computes the attractor dynamics that governs the actual movement of the robot.

respectively. q^* denotes an attractor (resting pose) in configuration space. Constraints on the system are implemented by forcing the system via $f_i(t)$, providing automatic avoidance of kinematic infeasibilites arising from joint limits, cable lengths, and collisions.

An agent can interact with MoBeE, instead of directly with the robot, by sending arbitrary high-level control commands. For example, when a new attractor q^* is set to a desired pose by an agent, e.g. by calculating the inverse kinematics of an operational space point, $q(t)$ begins to move toward q^*. The action then terminates either when the dynamical system settles or when a timeout occurs, depending on the constraint forces $f_i(t)$ encountered during the transient response.

10.4.2 Action Repertoire: TRM & LEOGrasper

The action repertoire for the scenario herein consists mainly of a grasping subsystem and a framework to generate full-body motions. LEOGrasper is a light-weight, easy-to-use, one-shot grasping system (Figure 10.6).[3]), It can be configured to perform a variety of grasps, all requiring to close the fingers in a coordinated fashion. The *iCub* incorporates touch sensors on the fingertips, due to the high noise, we use the error reported by the PID controllers of the finger motors to know when they are in contact with the object. A variety of more complex actions/roadmaps can be generated offline and later executed on the *iCub* [55], to e.g. lead to improved perception [36].

To generate a set of more-complex actions to execute MoBeE's kinematic model was extracted and connected with a machine-learning based, black-box optimizer. The system aims to find find a robot pose $q_{goal} \in C$, where C describes the robot's configuration space, that satisfies some operational space constraints, with planning, i.e. find a feasible configuration-space trajectory, $Q \subset C$, which is the curve from the current pose, $q_{initial}$ and the target pose q_{goal}.

Natural Evolution Strategies (NES) are applies to find a set of task-related poses yielding Task-relevant Road Map (TRM) [55]. It finds a family of postures that are optimized under constraints defined by arbitrary cost-functions, and at the same time maximally covers a user-defined task-space. Connecting these postures creates a rather dense, traversable graph, which we call roadmaps. In other words, the task-relevant constraints are built directly into the TRM, and motion planning is reduced to graph search. This allows to build TRMs that can perform useful tasks in the 41-dimensional configuration space of the upper body of the *iCub* humanoid. Additionally these maps can be stored to create an action repertoire that can be recalled when a certain tasks needs to be executed. Figure 10.7 shows time-lapse snapshots of motions, planned within TRMs. It provides an idea of what kind of motions can be generated.

[3]Source code available at: `https://github.com/Juxi/iCub/`

FIGURE 10.6 Grasping a variety of objects successfully, such as, tin cans, plastic cups and tea boxes. The module works for both the right and left hand.

10.5 INTEGRATION

To allow for a variety of objects to be picked up from various positions the robot needs to see, act and react within an integrated control system.

For example, methods enabling a 5 DOF robotic arm to pick up objects using a point-cloud generated model of the world and objects are available to calculate reach and grasp behaviours [52]. In 2010 a technique for robots to pick up non-rigid objects, such as, towels was presented [38]. It allows to reliably and robustly pick up a towel from a table by going through a sequence of vision-based re-grasps and manipulations-partially in the air, partially on the table. Even when sufficient manipulation skills are available these need to be constantly adapted by an perception-action loop to yield desired results. 'Robotics, Vision and Control' [16] puts this close integration of the mentioned components into the spotlight and describes common pitfalls and issues when trying to build such systems with high levels of sensorimotor integration. In the DARPA ARM project, which aims to create highly autonomous manipulators capable of serving multiple purposes across a wide variety of applications, the winning team showed an end-to-end system that allows the robot to grasp and pick-up diverse objects (e.g. a power drill, keys, screwdrivers, etc.) from a table by combining touch and LASER sensing [26].

10.5.1 Closing the Action-Perception Loop

The aim is to generate a pick-and-place operation for the *iCub*. For this, functional motion and vision subsystems are integrated to create a closed action-perception loop. The vision side detects and localises the object continuously, while the motor-side tries to reach for target objects avoiding obstacles at the same time. A grasping of the object is triggered when the hand is near the target. The sensory and motor sides establish quite a few capabilities by themselves, yet to grasp objects successfully while avoiding obstacles they

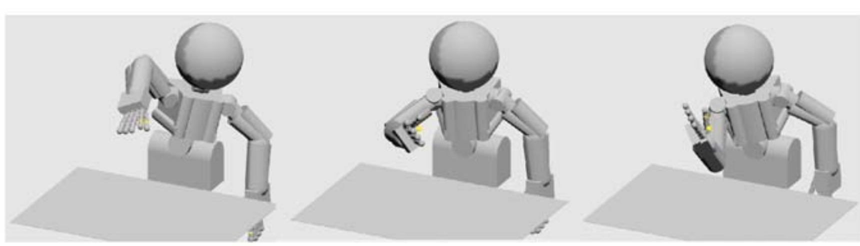

FIGURE 10.7 'Curiously inspect something small': the 3D position of the hand is constrained, and the task space is its angle with respect to the gaze direction. The resulting map rotates the hand (and any grasped object) in front of the eyes.

need to work closely together. The continuous tracking of obstacles and the target object is required to create a reactive reaching behaviour which adapts in real-time to the changes of the environment.

By creating interfaces between *MoBeE* and *icVision* the robot is able to continuously perform a visual based localisation of the detected objects and propagated this information into the world model. This basic eye-hand coordination allows for an adaptation while executing the reaching behaviour to changing circumstances, improving our robot's autonomy.

10.6 RESULTS

The first experiment shows that the herein presented system is able to reactively move the arm out of harms way when the environment changes. Then it is shown how this system can be used to reactively reach and grasp objects.

10.6.1 Avoiding a Moving Obstacle

Static objects in the environment can be added directly into *MoBeE*'s world model. Once, e.g. the table, is in the model, actions and behaviours are adapted due to computed constraint forces. These forces, $f_i(t)$ in (1), which repel the robot from collisions with the table, governs the actual movement of the robot. This way we are able to send arbitrary motions to our system, while ensuring the safety of our robot (this has recently been shown to provide a good reinforcement signal for learning robot reaching behaviours [46, 20]). The presented system has the same functionality also for arbitrary, non-static objects. After detection in both cameras the object's location is estimated (*icVision*) and propagated to *MoBeE*. The fictional forces are calculated to avoid impeding collisions. Figure 10.8 shows how the localised object is in the way of the arm and the hand.[4] To ensure the safety of the rather fragile fingers, a sphere around the end-effector can be seen. It is red, indicating a possible collision, because the sphere intersects with the object. The same is valid for the lower arm. The forces, calculated at each body part using Jacobians, push the intersecting geometries away from each other, leading to a forcing of the hand (and arm) away from the obstacle. Figure 10.9 shows how the the robot's arm is avoiding a non-stationary obstacle.[5] The arm is 'pushed' aside at the beginning, when the cup is moved close to the arm. It does so until the arm reaches its limit, then the forces cumulate and the end-effector is 'forced' upwards to continue avoiding the obstacle. Without an obstacle the arm starts to settle back into its resting pose q^*.

[4]See video: `https://www.youtube.com/watch?v=w_qDH5tSe7g`
[5]See video: `https://www.youtube.com/watch?v=w_qDH5tSe7g`

10.6.2 Reaching and Grasping Objects

This next experiment is on a simple reactive pick-and-place routine for the *iCub*. Similarly to the above experiment we are using *MoBeE* to adapt the reaching behaviour while the object is moved. To do this we change the type of the object within the world model from 'obstacle' into 'target'. Due to this change there is no repelling force calculated between the object and the robot parts. In fact we can now use the vector from the end-effector to the target object as a force that drives the hand towards a good grasping position.

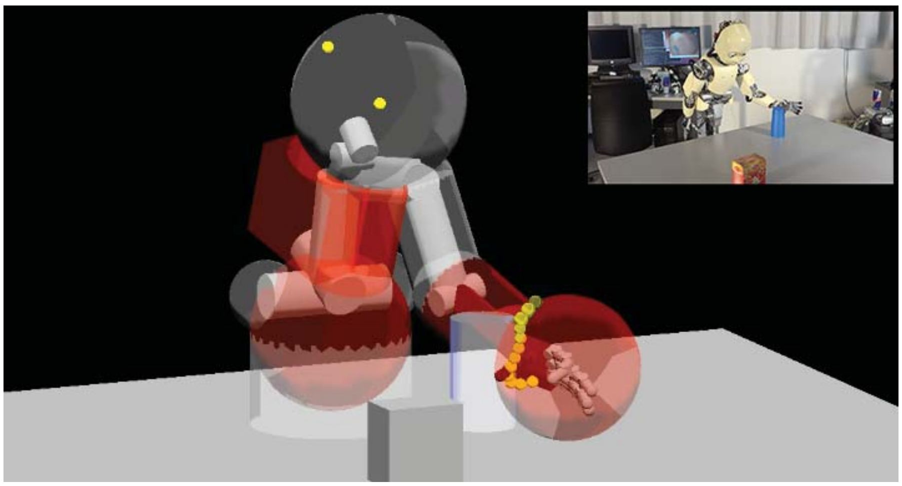

FIGURE 10.8 Showing the visual output of the *MoBeE* world model during one of our experiments. Parts in red indicate (an impeding) collision with the environment (or itself). The inset shows the actual scene.

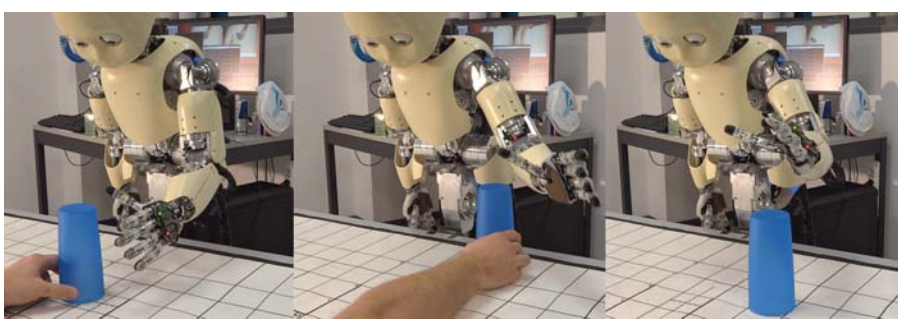

FIGURE 10.9 The reactive control of the left arm, permitting the iCub to stay clear of the non-static 'obstacle', as well as the table.

MoBeE also allows to trigger certain responses when collisions occur. In the case, when we want the robot to pick-up the object, we can active a grasp subsystem whenever the hand is in the close vicinity of the object. We are using a prototypical power grasp style hand-closing action, which has been used successfully in various demos and videos.[6] Figure 10.6 shows the *iCub* successfully picking up (by adding an extra upwards force) various objects using our grasping subsystem, executing the same action.

Our robot frameworks are able to track multiple objects at the same time, which is also visible in Figure 10.8, where it tracks both the cup and the tea box. By simply changing the type of the object within *MoBeE* the robot reaches for a certain object while avoiding the other.

10.7 CONCLUSIONS

Herein a cognitive robotics approach towards visual guided object manipulation with a humanoid was presented. A tightly integrated sensorimotor system, based on two frameworks developed over the past years, enables the robot to perform a simple pick-and-place task. The robot reaches to detected objects, placed at random positions on a table.

The implementation enables the robot to adapt to changes in the environment. Through this it safeguards the *iCub* from unwanted interactions – i.e. collisions. This is facilitated by a tight integration of the visual system with the motor side. Specifically an attractor dynamic based on the robot's pose and a model of the world. This way a level of eye-hand coordination not previously seen on the *iCub* was achieved.

In the future more integration of machine learning to further improve the object manipulation skills of our robotic system is planned. Improving the predication and selection of actions will lead to a more adaptive, versatile robot. Furthermore it might be of interest to investigate an even tighter sensorimotor coupling, e.g. avoiding translation into operational space by working in vision/configuration space.

10.8 Bibliography

[1] R. Ambrose, B. Wilcox, B. Reed, L. Matthies, D. Lavery, and D. Korsmeyer. NASA's Space Technology Roadmaps (STRs): Robotics, telerobotics, and autonomous systems roadmap. Technical report, National Aeronautics and Space Administration (NASA), 2012.

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

[6] See videos at: `http://robotics.idsia.ch/media/`

[3] M. Asada, K.F. 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.

[4] Tucker Balch and Ronald C Arkin. Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, 1998.

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

[6] J. Bongard, V. Zykov, and H. Lipson. 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.A. Brooks. Intelligence without representation. *Artificial intelligence*, 47(1):139–159, 1991.

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

[10] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M.R. Zakin, H. Lipson, and H.M. Jaeger. Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences (PNAS)*, 107(44):18809–18814, 2010.

[11] A. Cangelosi, T. Belpaeme, G. Sandini, G. Metta, L. Fadiga, G. Sagerer, K. Rohlfing, B. Wrede, S. Nolfi, D. Parisi, C. Nehaniv, K. Dautenhahn, J. Saunders, K. Fischer, J. Tani, and D. Roy. The ITALK project: Integration and transfer of action and language knowledge in robots. In *Proceedings 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] N.G. Cheng, M.B. Lobovsky, S.J. Keating, A.M. Setapen, K.I. Gero, A.E. Hosoi, and K.D. Iagnemma. Design and analysis of a robust, low-cost, highly articulated manipulator enabled by jamming of granular media. In *Proceedings of the International Conference of Robotics and Automation (ICRA)*, pages 4328–4333, 2012.

[14] Peter C. Chu. SMART Underwater Robot (SUR) Application & Mining, 2011. Technical Presentation.

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

[16] P.I. Corke. *Robotics, Vision and Control*, volume 73 of *Springer Tracts in Advanced Robotics*. Springer, 2011.

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

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

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

[20] Mikhail Frank, Jürgen Leitner, Marijn Stollenga, Alexander Förster, and Jürgen Schmidhuber. Curiosity driven reinforcement learning for motion planning on humanoids. *Frontiers in Neurorobotics*, 7(25), 2014.

[21] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Prentice-Hall, 3rd edition, 2006.

[22] Simon Harding, Jürgen Leitner, and Jürgen Schmidhuber. Cartesian genetic programming for image processing. In Rick Riolo, Ekaterina Vladislavleva, Marylyn D Ritchie, and Jason H. Moore, editors, *Genetic Programming Theory and Practice X*, Genetic and Evolutionary Computation, pages 31–44. Springer New York, 2013.

[23] 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.

[24] B.K. Horn. *Robot Vision*. MIT Press, 1986.

[25] A. Hornberg. *Handbook of machine vision*. Wiley, 2007.

[26] N. Hudson, T. Howard, J. Ma, A. Jain, M. Bajracharya, S. Myint, C. Kuo, L. Matthies, P. Backes, P. Hebert, T. Fuchs, and J. Burdick. End-to-end dexterous manipulation with deliberate interactive estimation. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2012.

[27] M. Jeannerod. *The cognitive neuroscience of action.* Blackwell Publishing, 1997.

[28] Mark H Johnson and Yuko Munakata. Processes of change in brain and cognitive development. *Trends in cognitive sciences*, 9(3):152–158, 2005.

[29] C.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] D. Kragic and M. Vincze. Vision for robotics. *Foundations and Trends in Robotics*, 1(1):1–78, 2009.

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

[32] J. Leitner, P. Chandrashekhariah, S. Harding, M. Frank, G. Spina, A. Forster, J. Triesch, and J. Schmidhuber. Autonomous learning of robust visual object detection and identification on a humanoid. In *Proceedings of the International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, November 2012.

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

[34] J. Leitner, S. Harding, M. Frank, A. Förster, and J. Schmidhuber. Learning spatial object localization from vision on a humanoid robot. *International Journal of Advanced Robotic Systems (ARS)*, 9, 2012.

[35] J. Leitner, S. Harding, M. Frank, A. Förster, and J. Schmidhuber. Transferring spatial perception between robots operating in a shared workspace. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, October 2012.

[36] Jürgen Leitner, Alexander Förster, and Jürgen Schmidhuber. Improving robot vision models for object detection through interaction. In *2014 International Joint Conference on Neural Networks, IJCNN 2014, Beijing, China, July 6-11, 2014*, pages 3355–3362. IEEE, 2014.

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

[38] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel. Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 2308–2315, 2010.

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

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

[41] G. Metta, P. Fitzpatrick, and L. Natale. YARP: Yet Another Robot Platform. *International Journal of Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics*, 3(1), 2006.

[42] Giorgio Metta, Lorenzo Natale, Francesco Nori, Giulio Sandini, David Vernon, Luciano Fadiga, Claes von Hofsten, Kerstin Rosander, Manuel Lopes, JosÃl' Santos-Victor, Alexandre Bernardino, and Luis Montesano. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks*, 23(8-9):1125–1134, October 2010.

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

[44] 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. MIT Press, 1994.

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

[46] Shashank Pathak, Luca Pulina, Giorgio Metta, and Armando Tacchella. Ensuring safety of policies learned by reinforcement: Reaching objects in the presence of obstacles with the icub. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[47] J. Peters, S. Vijayakumar, and S. Schaal. Reinforcement learning for humanoid robotics. In *Proceedings of the International Conference on Humanoid Robots*, 2003.

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

[49] M.I. Posner. *Foundations of cognitive science*. The MIT Press, 1989.

[50] Stuart Jonathan Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, 2010.

[51] A. Sadikov, M. Možina, M. Guid, J. Krivec, and I. Bratko. Automated chess tutor. In H.J. Herik, P. Ciancarini, and H.H.L.M. Donkers, editors, *Computers and Games*, volume 4630 of *Lecture Notes in Computer Science*, pages 13–25. Springer Berlin Heidelberg, 2007.

[52] A. Saxena, J. Driemeyer, and A.Y. Ng. Robotic grasping of novel objects using vision. *The International Journal of Robotics Research*, 27(2):157, 2008.

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

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

[55] M. Stollenga, L. Pape, M. Frank, J. Leitner, A. Fõrster, and J. Schmidhuber. Task-relevant roadmaps: A framework for humanoid motion planning. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[56] R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*. Cambridge Univ Press, 1998.

[57] N.G. Tsagarakis, G. Metta, G. Sandini, D. Vernon, R. Beira, F. Becchi, L. Righetti, J. Santos-Victor, A.J. Ijspeert, M.C. Carrozza, and D.G. Caldwell. iCub: the design and realization of an open humanoid platform for cognitive and neuroscience research. *Advanced Robotics*, 21:1151–1175, 2007.

[58] D. Vernon, G. Metta, and G. Sandini. A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation*, 11(2):151–180, 2007.

[59] D. Vernon, G. Metta, and G. Sandini. The iCub Cognitive Architecture: Interactive Development in a Humanoid Robot. In *Proceedings of the International Conference on Development and Learning (ICDL)*, 2007.

[60] WAY. EU Project Consortium: Wearable Interfaces for Hand Function Recovery. `http://www.wayproject.eu/`, January 2012.

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

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

[63] 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.