Proceedings of Humanoids 2012
Workshop on Developmental Robotics

Can developmental robotics yield human-like cognitive abilities?

Editors

Emre Ugur, Yukie Nagai
Erhan Oztop, Minoru Asada

Business Innovation Center, Osaka, Japan

November 29th, 2012

Preface

The state-of-the-art robots designed for specific tasks exhibit abilities that surpass those of human. However, no robot can pass the motor and cognitive capabilities of a 3 years old child. To create machines that parallel or even pass the motor and cognitive capabilities of humans, the Developmental Robotics field was born as an alternative to previous robot-learning or AI programming approaches. This new approach argued that the continual development of an embodied robotic agent following the development steps of a human is key to achieve motor and cognitive skills that of a human. Consequently research efforts based on this view produced impressive results by focusing on different developmental stages of the embodied "infant robots" that develop through interaction with the environment.

The aim of this one-day workshop is two-fold. First, we provide an overview of the current state-of-the-art in this field, and remind ourselves about the promises of the developmental robotics and the achievements obtained until now. Our speakers from developmental psychology also discuss cognitive capabilities of human infants in different stages of their development and the possible mechanisms of acquiring these capabilities.

Second, we motivate our speakers to comment on (1) how different developmental stages and computational models developed so far can be combined to achieve a coherent model that explains all different developmental stages, and (2) how higher-level cognitive competence can emerge in this developmental progression. The majority of computational models and learning methods developed until now correspond to skills of infants 2 years old or younger. We try to answer the question what is needed for developmental robotics to make the leap to enable infant robots to acquire the higher-level cognitive abilities, such as complex reasoning, symbolic planning and mental state inference.

Five regular papers and four extended abstracts were accepted as contributions to the workshop after peer-reviewing. We expect that the six invited talks from distinguished scientists on developmental robotics and psychology, together with the contributed talks will elucidate initial answers for the questions posed above, and emphasize the challenges ahead.

We thank all submitting authors for choosing this workshop to disseminate their work. We thank keynote speakers who considerably contributed to the quality and the impact of the workshop. Needless to say, the program committee members have a big role in making the workshop a success; we thank them for their fine reviewing efforts. Finally, we would like to thank the Humanoids 2012 Organization Committee for facilitating the workshop execution, and Prof. Shin Ishii, Kyoto University and head of Dynamic Brain Imaging Department, ATR for giving support to the preparation of the workshop.

November 2012

Emre Ugur
Yukie Nagai
Erhan Oztop
Minoru Asada

Organizing committee

- Emre Ugur, ATR, Kyoto, Japan.
- Yukie Nagai, Osaka University, Osaka, Japan.
- Erhan Oztop, Ozyegin University, Istanbul, Turkey.
- Minoru Asada, Osaka University, Osaka, Japan.

Program committee

- Eren Erdal Aksoy, University of Gottingen, Germany
- Maya Cakmak, Georgia Institute of Technology, USA
- Hande Celikkanat, University of Plymouth, UK
- Renaud Detry, KTH Royal Institute of Technology, Sweden
- Mehmet Remzi Dogar, Carnegie Mellon University, USA
- Frank Guerin, University of Aberdeen, Scotland
- Sinan Kalkan, Middle East Technical Univesity, Turkey
- Jun Morimoto, ATR, Japan
- Ludovic Righetti, University of Southern California, USA
- Erol Sahin, Middle East Technical Univesity, Turkey
- Jivko Sinapov, Iowa State University, USA

Contact

Emre Ugur, PhD
ATR, NIA Labs.
2-2-2 Hikaridai Seika-cho Soraku-gun
Kyoto 619-0288 Japan
email: emre@atr.jp
phone: +81 774 95 2403

Table of Contents

Invited Talks

Regular Papers

Extended Abstracts

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Invited Talk I

Action as a founding principle of cognitive development in humans and robots

Prof. Claes von Hofsten

Uppsala University, Uppsala, Sweden

Abstract: It is argued that action constitutes the foundation for cognitive development. It reflects the motives of the child, the problems to be solved, the goals to be attained, and the constraints and possibilities of the childs body and sensory-motor system. Actions are directed into the future and their control is based on knowledge of what is going to happen next. This is possible because the stream of events in the world is governed by rules and regularities. Infants are endowed with innate predispositions that make them able to use those rules to their advantage. However, the most important developmental principle is a set of motives that makes infants do certain things rather than others. These motives provide the goals of actions and the urge to fulfill them. In early development, infants rapidly acquire knowledge about external events, their own body, and other peoples actions that enable them to interact intelligently with the outside world. By acting on the world, infants develop their cognition. If robots can be endowed with similar motives, they could, in principle, develop human-like cognitive abilities.

Speaker Bio: Dr. Claes von Hofsten is a Full Professor of Psychology at the University of Oslo and the University of Uppsala. He received his PhD in psychology at Uppsala University in Sweden in 1973. Between 1998 and 2011 he was a professor in perception at Uppsala. He has spent several extended periods at American universities; as visiting professor at University of Minnesota and University of Virginia and as visiting scientist at MIT and the Center for Advanced Study in the Behavioral Sciences at Stanford. He is also Honoris Causa at University of Normandy in France and honorary member of the American Academy of Arts and Sciences. Dr von Hofsten research interests are focused on the development of action in young children.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Invited Talk II

Developmental Approach to Robotic Intelligence

Prof. Alexander Stoytchev

Iowa State University, Iowa, USA

Abstract: This talk will focus on recent research results that show how a robot can solve multiple tasks based on what it learns during a developmental period similar to a childs play. During this period the robot actively tries to grasp, lift, shake, touch, scratch, tap, push, drop, and crush objects. At the end of this period the robot knows what different objects sound like when they are dropped, feel like when they are squeezed, etc. Because these properties are grounded in the robots sensorimotor repertoire the robot can autonomously learn, test, and verify its own representations without human intervention. The talk will demonstrate how the robot can use this information to recognize objects, separate objects into functional categories, and even find the odd-one-out in a set of objects. The talk will also demonstrate how the robot can use sensorimotor interactions to bootstrap the development of its visual system in the context of a button-pressing task. Results and videos will be presented for two different humanoid platforms.

Speaker Bio: Dr. Alexander Stoytchev is an Assistant Professor of Electrical and Computer Engineering and the Director of the Developmental Robotics Laboratory at Iowa State University, USA. He received his MS and PhD degrees in computer science from the Georgia Institute of Technology in 2001 and 2007, respectively. His research interests are in the areas of developmental robotics, autonomous robotics, computational perception, and machine learning.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of
Humanoids 2012 Workshop on Developmental Robotics: Can
developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Invited Talk III

Developmental Robotics: Where from and Where to (a personal view).

Prof. Giulio Sandini

Italian Institute of Technology (IIT), Genova, Italy

Speaker Bio: Dr. Giulio Sandini Director of Research at the Italian Institute of Technology and full professor of bioengineering at the University of Genoa. Main research activities are in the fields of computational and cognitive neuroscience and robotics with the objective of understanding the neural mechanisms of human sensory-motor coordination and cognitive development. After graduating in Electronic Engineering (Bioengineering) he was research fellow and assistant professor at the Scuola Normale Superiore in Pisa and at the Laboratorio di Neurofisiologia of the CNR where he investigated aspects of visual processing at the level of single neurons as well as aspects of visual perception in human adults and children. He has been Visiting Research Associate at the Department of Neurology of the Harvard Medical School and Visiting Scientist at the Artificial Intelligence Lab of MIT. Since July 2006 he has been appointed Director of Research at the Italian Institute of Technology where he has established and is currently directing the department of Robotics, Brain and Cognitive Sciences.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Invited Talk IV

Integrating visual perception and manipulation for autonomous learning of object representations

Dr. Ales Ude

Jozef Stefan Institute, Ljubljana, Slovenia

Abstract: The human ability to discern objects in the scene is not innate but rather acquired during the early development. From birth on, children are constantly exposed to events caused by the effects of their own actions. The information thus gained can be used to evolve the agents perceptual judgements, including the way how objects are perceived. Reliable object perception is still difficult to achieve in artificial systems because it is not clear how to define the concept of objectness in its full generality. In our research we follow the paradigm that integrates the development of perceptual representations with the robots manipulation capabilities and tactile sensing. In this way the robot can introduce additional information that can be utilized to reliably separate previously unknown objects from the background and learn their representations and affordances.

Speaker Bio: Ale Ude studied applied mathematics at the University of Ljubljana, Slovenia. He received the Ph.D. degree for work on robot programming by demonstration from the University of Karlsruhe, Germany. He was an STA fellow in the Kawato Dynamic Brain Project, which was conducted at ATR in Kyoto, Japan. Currently he is the head of Humanoid and Cognitive Robotics Lab at Joef Stefan Institute, Ljubljana, Slovenia and is also associated with the ATR Computational Neuroscience Laboratories in Kyoto, Japan. His current research interests include learning in humanoid systems, especially imitation learning and learning by exploration, humanoid robot vision, and humanoid cognition.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of
Humanoids 2012 Workshop on Developmental Robotics: Can
developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Invited Talk V

Development of functional hierarchy for actions and motor imageries: a synthetic neurorobotics experiment

Prof. Jun Tani

Korean Advanced Science and Technology (KAIST), Korea

Abstract: In this talk I introduce a neuro-robotics experiment on developmental learning of goal-directed actions. The robot was trained to predict visuo-proprioceptive flow of achieving a set of goal-directed behaviors through iterative tutor training processes. The learning was conducted by employing a dynamic neural network model which is characterized by their multiple time-scales dynamics. The experimental results showed that functional hierarchical structures emerge through stages of developments where behavior primitives are generated in earlier stages and their sequences of achieving goals appear later stages. It was also observed that motor imagery is generated in earlier stages compared to actual behaviors. Our claim that manipulable inner representation should emerge through the sensory-motor interactions is corresponded to Piaget's constructivist view.

Speaker Bio: Jun Tani received a B.S. in Mechanical Engineering from Waseda University, a dual M.S. in Electrical Engineering and Mechanical Engineering from the University of Michigan, and a Dr. Eng. from Sophia University. He started his research career in Sony Computer Science Laboratory in 1990. He worked as a PI of the Lab. for Behavior and Dynamic Cognition, Brain Science Institute, RIKEN in Tokyo from 2001 to 2012. He also held the position of Visiting Associate Professor at the Univ. of Tokyo between 1997 and 2002. He became a full professor in Electrical Engineering Dept. in KAIST 2012 where he started cognitive neurorobotics. He is interested in neuroscience, psychology, phenomenology, complex adaptive systems, and robotics.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Invited Talk VI

Constructive Developmental Science Based on Understanding the Process from Neuro-Dynamics to Social Interaction

Prof. Minoru Asada

Adaptive Machine Systems, Graduate School of Engineering, Osaka University, Japan

Speaker Bio: Minoru Asada is Professor of the department of Adaptive Machine Systems at the Graduate School of Engineering, Osaka University (Suita, Japan). He received his Ph.D. in control engineering from Osaka University in 1982. Professor Asada was elected to and served as General Chair of the IEEE/RSJ 1996 International Conference on Intelligent Robots and Systems (IROS96). Since early 1990, Professor Asada has been involved in RoboCup activities and his team was the inaugural champion (shared with USC), in the mid- sized league of the first RoboCup competition held in conjunction with IJCAI-97 (Nagoya, Japan). Since 2002, Professor Asada has served as President of the International RoboCup Federation. In 2005, Professor Asada was elected Fellow of the IEEE for Contributions to Robot Learning and Applications. Also in 2005, Professor Asada was elected to serve as Research Director of the ASADA Synergistic Intelligence Project of ERATO (Exploratory Research for Advanced Technology) by the Japan Science and Technology Agency and he continued to serve as Research Director until the Project was completed in 2012. In 2007, Professor Asada was awarded The Okawa Publications Prize (The Okawa Foundation) and, in 2008, he received The Good Designs Award for VoCal - Vivid Oral Conversation through Acquiring Language (Japan Industrial Design Promotion Organization) . In 2009, Professor Asada again received the Best Paper Award of the Robotics Society of Japan. And, in 2012, The Japan Society for Promotion of Science (JSPS) named Professor Asada to serve as Research Leader for the Specially Promoted Research Project (Tokusui) on Constructive Developmental Science Based on Understanding the Process From Neuro-Dynamics to Social Interaction.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

The Object Pairing and Matching Task: Toward Montessori Tests for Robots

Connor Schenck and Alexander Stoytchev
Developmental Robotics Laboratory
Iowa State University
{cschenck, alexs}@iastate.edu

***Abstract*—The Montessori method is a popular approach to education that emphasizes student-directed learning in a controlled environment. Object matching is one common task that children perform in Montessori classrooms. Matching tasks also occur quite frequently on intelligence tests for humans, which suggests that intelligence correlates with the skills required to solve these tasks. This paper describes robotic experiments with four Montessori matching tasks: sound cylinders, sound boxes, weight cylinders, and pressure cylinders. The robot grounded its representation for the twelve objects in each task in terms of the auditory and proprioceptive outcomes that they produced in response to a set of ten exploratory behaviors. The results show that based on this representation, it is possible to identify task-relevant sensorimotor contexts (i.e., exploratory behavior and sensory modality combinations) that are useful for performing matching on a given set of objects. Furthermore, the results show that as the number of sensorimotor contexts used to perform matching increases, the robot's ability to match the objects also increases.**

I. Introduction

The Montessori method is a 100-year-old method of schooling that was developed by Maria Montessori (1870-1952), an influential Italian educator. It is characterized by a special set of educational materials and student-directed learning activities [1] [2] [3]. One of its core principles is that of *embodied cognition*, tying movement of the body and learning together. It focuses on stimulating the development of different skill sets, including sensory development, language development, and numeracy skills. Most Montessori tasks require that the children actively touch, move, relate, and compare objects [2].

One task typical for a Montessori classroom is object matching. Children are given two sets of objects and asked to find the matches from one set to another. Sample tasks include matching colored tiles, matching 3-dimensional shapes, and matching pieces of textured cloth [4]. All these tasks are designed to stimulate a child's ability to perceive object properties and to allow the child to learn about the nature of objects and their similarities.

The skills required to perform matching are also useful for other tasks such as object grouping, category recognition, and object ordering. At a fundamental level, these skills require the ability to find differences between similar objects and similarities between different objects. Recent work in robotics has found that robots are able to recognize objects and their categories [5], [6], group objects in an unsupervised manner [7], and find the odd one out in a set of objects [8]. These studies all strongly suggest that a robot should be able to solve object pairing tasks.

Fig. 1. The robot and the four Montessori matching tasks that were used in the experiments. In clockwise order, the four tasks were: sound cylinders, weight cylinders, pressure cylinders, and sound boxes.

This paper describes a method that allows a robot to identify and match object pairs within a set of objects based on their sensorimotor properties. To do this, the robot first interacted with the objects using a set of exploratory behaviors (grasp, lift, hold, shake, rattle, drop, tap, poke, push, and press) in order to ground the properties of the objects in the robot's behavioral repertoire. After interacting with the objects, the robot performed feature extraction on the raw sensory data to create sensory feedback sequences for each interaction. For each object, the robot recorded both proprioceptive feedback in the form of joint torques and auditory feedback in the form of an audio spectrogram. Next, the robot generated similarity scores for all possible object pairs and used these scores to match the objects. To combine information from different sensorimotor contexts (e.g., *audio-drop* and *proprioception-shake*), the robot used three different methods: uniform-weight combination, recognition accuracy based weight combination, and pairing accuracy based combination. These methods were evaluated for their ability to match standard Montessori objects.

This study used four typical Montessori matching tasks. In each task there were two groups of six objects and the goal was to find the matching pairs of objects between the two groups. The results indicate that the estimated object similarities were sufficient to adequately pair objects. The robot was able to solve the object matching task with a high degree of accuracy. Furthermore, the robot was able to identify the functionally meaningful sensorimotor contexts in which it can distinguish between objects. To the best of our knowledge, this is the first study that has applied Montessori learning techniques in a robotic setting.

II. Related Work

A. Psychology

Recent studies have found that students educated using the Montessori method often outperform students educated by traditional methods. For example, one study found that middle school students from Montessori schools had higher intrinsic motivation when it came to academic activities as compared to students from traditional schools [9]. This suggests that the Montessori method is more effective at fostering learning in young children than the traditional methods. This conclusion was supported by another study [3], which found that, by the end of kindergarten, Montessori students outperformed traditional students on standardized tests of reading and math and also showed more advanced social skills and executive control.

One task commonly used in the Montessori style of teaching for younger children is the matching task [4]. In this task, a child is given a set of objects (sometimes split into two subsets and sometimes not) and asked to pair the objects. A variant of that task was used by Daehler *et al.* [10] who used both objects and pictures of objects in their experiments. They found that children around the age of two are able to correctly match objects from both pictures and objects to sets of pictures or objects. One interesting result of this experiment was that the children performed significantly better on tasks where they were asked to match an object to a set of objects, versus picture to object, object to picture, or picture to picture matching. They suggested that this was due to the ability of the children to perceive the objects from multiple angles, thus giving them more reliable information about the objects than they could extract from the pictures.

Other studies have shown infants' ability to identify object pairs and group objects into categories. A study by Leslie *et al.* [11] demonstrated that eleven-month-old infants can individuate pairs of objects only when there is a large amount of physical similarity between objects in the same pair (in this study they used identical objects) and a large physical difference between objects of different pairs. Younger [12] showed that ten-month-old infants can form object categories and determine the variants and invariants of the objects within a category and based on that information they can determine the inclusion of a novel object in the given category. These studies show that even at an early age, humans are able to identify object properties and use them to compare objects, which suggests that this is a fundamental part of intelligence.

Another experiment by McPherson and Holcomb [13] examined event-related brain potentials. Participants were shown a picture of an object, then a picture of an object from one of three categories: related, moderately related, or unrelated. The electroencephalogram (EEG) results showed that across all participants, there was a large negative spike in the N400 family of potentials in the participants' brain shortly after being shown the second picture. The study found that the magnitude of the spike was related to the similarity between the two objects in the pictures. This suggests that, at least at some level, the brain makes a quantitative measure of how similar the two objects are.

B. Robotics

Several studies have demonstrated that robots can measure perceptual as well as functional object similarities for a variety of tasks [14], [15], [16], [17], [18], [8]. The ability to measure the similarity between two objects is extremely useful for tasks such as category recognition and object grouping. Several studies [16], [5] have used unsupervised approaches for object categorization, in which objects were categorized by the similarity of their perceptual features. Their results showed that when the robot was allowed to use all of its sensory modalities, its object categorizations closely resembled the human-provided ones. This suggests that allowing robots to perceive more features about objects can improve their ability to detect similarities between the objects.

Sinapov and Stoytchev [8] showed how a robot can solve the odd-one-out task. The robot picked the object in the group that was least similar to the rest and resulted in the rest of the objects being maximally similar. In this paper we use a similar method to generate similarity scores between objects. We then use this similarity measure to perform object matching rather than solving the odd-one-out task, though they are fundamentally related problems.

III. Experimental Platform

A. Robot and Sensors

The experiments in this study were performed with the upper-torso humanoid robot shown in Fig. 1. The robot has as its actuators two 7-DOF Barrett Whole Arm Manipulators (WAMs), each with an attached Barrett Hand. Each WAM has built-in sensors that measure joint angles and torques at 500 Hz. An Audio-Technica U853AW cardioid microphone mounted in the robot's head was used to capture auditory feedback at the standard 16-bit/44.1 kHz over a single channel.

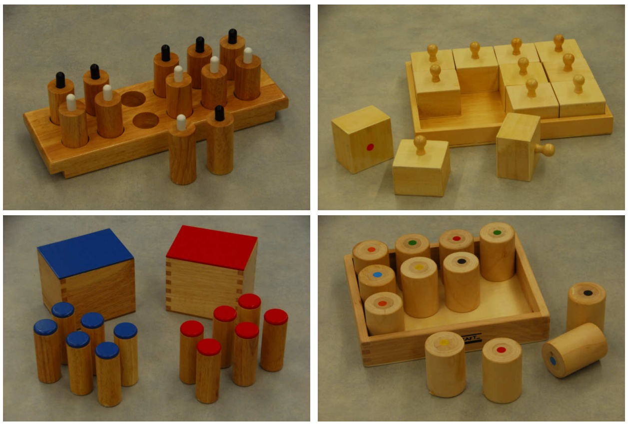

Fig. 2. The four sets of Montessori objects used in the experiments. From left to right and top to bottom the object sets are: *pressure cylinders*, *sound boxes*, *sound cylinders*, and *weight cylinders*. All the objects are marked with colored dots on the bottom to indicate the correct matches; other than that, the objects in each set are all visually identical (except for the *pressure cylinders* and the *sound cylinders*, which also have different colors for the tops to indicate the two sets of six objects).

B. Objects

The robot explored four standard Montessori sets of objects: *pressure cylinders*, *sound boxes*, *sound cylinders*, and *weight cylinders* (Fig. 2). Each set is composed of six pairs of objects. The objects in each pair are functionally identical to each other. The objects in each set are designed to vary in one specific dimension and be identical in all other dimensions. The *pressure cylinders* vary in the amount of force required to depress the rod, with pairs requiring the same amount of force. The *sound boxes* vary in the sounds they make when the contents move around inside the box, with pairs making the same sounds. The *sound cylinders* vary in the same way as the *sound boxes*, but are cylindrical in shape and have different contents than the boxes. The *weight cylinders* vary by weight, going from light to heavy, with pairs having the same weight.

C. Exploratory Behaviors

The robot used ten behaviors to explore the objects: *grasp*, *lift*, *hold*, *shake*, *rattle*, *drop*, *tap*, *poke*, *push*, and *press*. All of these exploratory behaviors, except rattle, have been used in our previous work [19], i.e., they were not specifically designed for the Montessori objects used in this paper. The behaviors were performed with the robot's left arm and encoded with the Barrett WAM API as trajectories in joint-space. The default PID controller of the WAM was used to execute the trajectories. Figure 3 shows images of the robot performing each behavior on one of the sound boxes. All the behaviors were performed identically on each object, with only minor variations due to the initial placement of the objects by the experimenter.

D. Data Collection

The robot interacted with the objects by performing a series of exploration trials. During each trial, an object was placed at a marked location on the table by the experimenter and the robot performed all ten of its exploratory behaviors on the object. The experimenter then picked another object and the robot repeated this process. This was done until each object had been explored ten times. During each interaction, the robot recorded proprioceptive information in the form of joint torques applied to the arm and auditory data captured by the microphone. The robot also recorded visual data, but it was not used in this experiment. In the end, the robot performed all ten behaviors ten times on each of the twelve objects in the four sets, resulting in $10 \times 10 \times 12 \times 4 = 4800$ behavior executions. This resulted in 18 GB of data, which was stored for off-line analysis. It took approximately 20 hours to collect this dataset.

IV. Feature Extraction

We used the method and the publicly available source code for proprioceptive and auditory feature extraction that is described in [5]. It is briefly summarized below. Proprioceptive data was recorded as joint torques over time resulting in a $7 \times m$ matrix, in which each column represents one set of torque readings for all joints and m is the number of readings. To reduce noise, a moving-average filter was applied over each row in the matrix, which corresponds to the torques from one joint. Audio data was recorded as wave files, one for each interaction. A log-normalized Discrete Fourier Transform was performed on each audio file using $2^5+1 = 33$ frequency bins resulting in a $33\times n$ matrix, where each column represents the activation values for different frequencies at a given point in time and n is the number of samples in the interaction. The Growing Hierarchical Self-Organizing Map (SOM) toolbox [20] was used to map each column to a single state. Two 6×6 SOMs were trained (one for audio and one for proprioception) using 5% of the columns that were randomly selected from all the joint torque and auditory data recorded by the robot. Each joint torque and auditory record was then mapped to a discrete sequence of states, where each column in the record was represented by the most highly activated SOM state for that column. For more details see [5].

V. Experimental Methodology

A. Estimating Similarity

Given a set of objects $\mathcal{O}$ the robot must be able to estimate the pairwise similarity for any two objects $i, j \in \mathcal{O}$ in a given sensorimotor context (i.e., exploratory behavior and sensory modality combination). Let $\mathcal{X}_c^i = [X_1, ..., X_D]$ be the set of sensory feedback sequences detected while interacting with object $i \in \mathcal{O}$ in sensorimotor context $c \in \mathcal{C}$ (where $\mathcal{C}$ is the set of all contexts) and let $sim(X_a, X_b)$ be the similarity between two sequences X_a and X_b. The similarity between objects i and j can be approximated with the expected pairwise similarity of the sequences in $\mathcal{X}_c^i$ and $\mathcal{X}_c^j$:

$$s_{ij}^c = \mathbf{E}[sim(X_a, X_b)|X_a \in \mathcal{X}_c^i, X_b \in \mathcal{X}_c^j]$$

In this paper we used the Needleman-Wunsch global alignment algorithm [21] to calculate $sim(X_a, X_b)$. The algorithm calculates the cost of aligning two discrete sequences

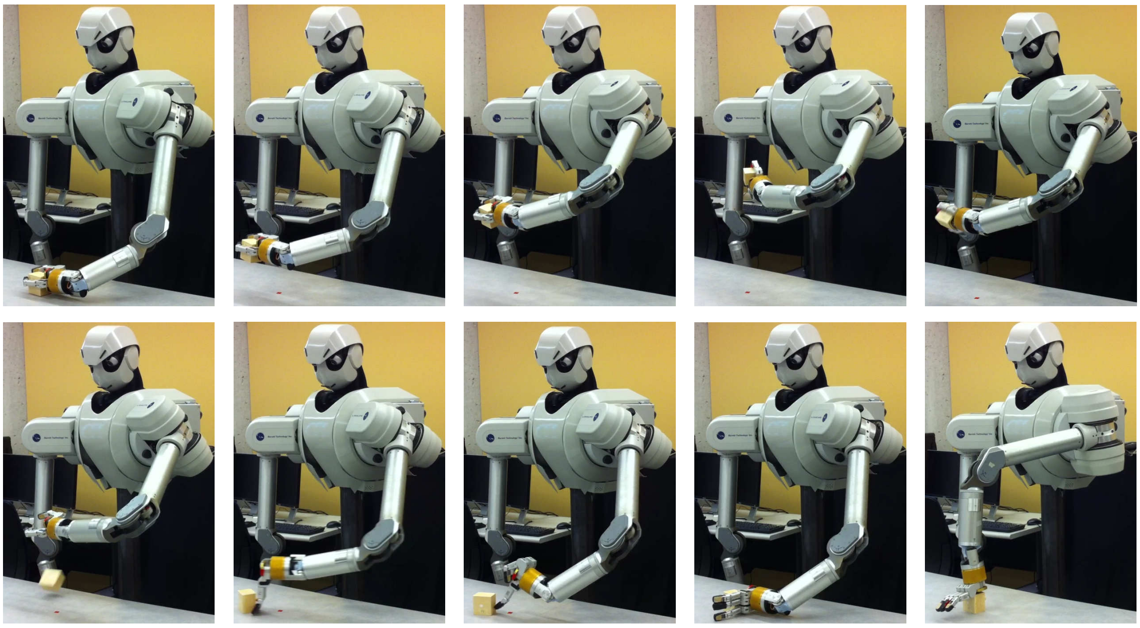

Fig. 3. The ten exploratory behaviors that the robot performed on all objects. From left to right and top to bottom: *grasp*, *lift*, *hold*, *shake*, *rattle*, *drop*, *tap*, *poke*, *push*, and *press*. The object in this figure is one of the sound boxes. The red marker on the table indicates the initial position of the objects at the beginning of each trial. The object was placed back in that position by the experimenter after some of the behaviors (e.g., drop).

(strings), which in our case correspond to sequences of most highly-activated SOM states (see the previous section). The expected similarity s^c_{ij} is estimated as

$$\frac{1}{|\mathcal{X}^i_c| \times |\mathcal{X}^j_c|} \sum_{X_a \in \mathcal{X}^i_c} \sum_{X_b \in \mathcal{X}^j_c} sim(X_a, X_b)$$

Next, the robot estimates the $|\mathcal{O}| \times |\mathcal{O}|$ pairwise object similarity matrix $\mathbf{W}^c$ for a specific sensorimotor context $c \in \mathcal{C}$. Each entry W^c_{ij} in $\mathbf{W}^c$ is defined as the similarity s^c_{ij} between two objects i and j in the specific context c. Figure 4 shows the similarity matrices for the *sound cylinders* for each of the 20 contexts.

B. Combining Sensorimotor Contexts

It has been shown that combining information from different sensorimotor contexts has a boosting effect for tasks such as object recognition [22]. Since object matching is a similar task, it is likely that combining contexts will be useful in this case as well. Thus, in this paper, we propose three methods to combine sensorimotor contexts: uniform combination, recognition accuracy based combination, and pairing accuracy based combination. The result of combining different contexts is a consensus matrix $\mathbf{W}$ that represents the similarity between object pairs for the specific set of contexts that was used to create it.

1) Uniform Combination: Given some set of contexts $\mathcal{C}'$, where $\mathcal{C}' \subseteq \mathcal{C}$, the similarity matrices $\mathbf{W}^c$ for each of these contexts can be used to construct the consensus matrix $\mathbf{W}$ by simply averaging their individual values, i.e.,

$$W_{ij} = \frac{1}{|\mathcal{C}'|} \sum_{c \in \mathcal{C}'} W^c_{ij}$$

for all pairs of objects i and j.

2) Recognition Accuracy Based Combination: This method assumes that contexts that are useful for object recognition will also be useful for object pairing. The object recognition accuracy r_c for context c is estimated by performing 10-fold cross validation on all the data from context c using a classifier that attempts to recognize object identities from sensory feedback sequences. To create the consensus matrix for a given set of contexts $\mathcal{C}'$ ($\mathcal{C}' \subseteq \mathcal{C}$), a weighted combination was used:

$$W_{ij} = \sum_{c \in \mathcal{C}'} \alpha_c \times W^c_{ij}$$

where α_c is the normalized recognition accuracy r_c for context c such that $\sum_{c \in \mathcal{C}'} \alpha_c = 1.0$. The classifier used in this paper was the k-nearest neighbor classifier with k set to 3 and using the global alignment similarity function as a similarity metric.

3) Pairing Accuracy Based Combination: The third combination method allowed the robot to get feedback on its attempts to pair some of the objects to refine its ability to pair the remaining objects. In order to determine the usefulness of each context, the robot split the set of objects such that either 2, 3, or 4 of the six pairs were in the training set and the rest remained in the testing set. Then, for each context c, using the objects in the training set, the robot would attempt to pair them (using the pairing method described below) and evaluate the pairing accuracy p_c for that context. To construct the consensus matrix $\mathbf{W}$, a weighted combination was used similar to the previous method:

$$W_{ij} = \sum_{c \in \mathcal{C}'} \alpha_c \times W^c_{ij}$$

where α_c is the normalized pairing accuracy p_c for context c such that $\sum_{c \in \mathcal{C}'} \alpha_c = 1.0$. After generating the consensus matrix $\mathbf{W}$, the robot would then attempt to pair only the

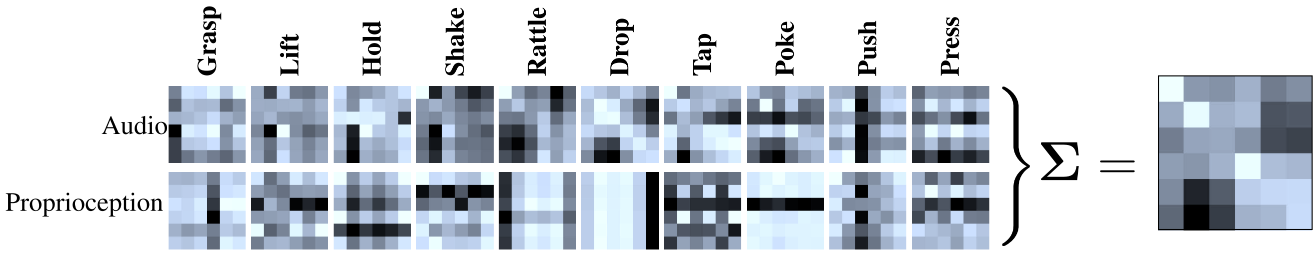

Fig. 4. The similarity matrices used to perform matching given two sets of six objects each for the *sound cylinders*. The matrices for each individual context are shown as well as the consensus matrix for all 20 contexts. The pairing accuracy combination method using four pairs for training was used to combine the individual matrices.

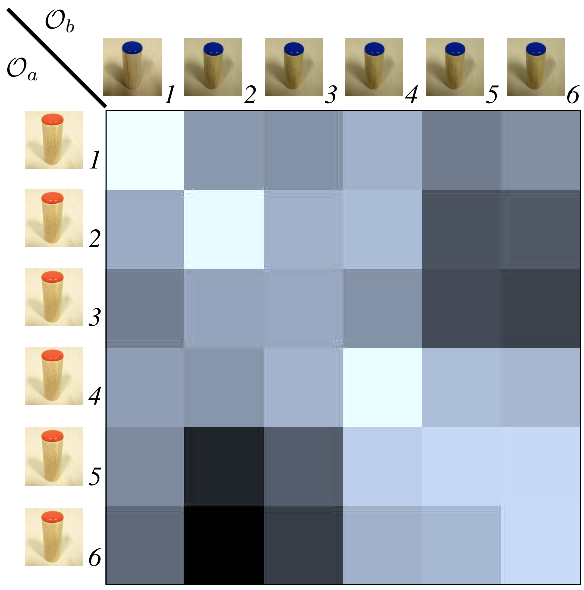

Fig. 5. The consensus weight matrix for the *sound cylinders* using all 20 sensorimotor contexts for matching two groups of six objects. The pairing accuracy combination method using four pairs to train was used to combine the individual similarity matrices for each context. The subscripts indicate correct matches.

objects from the testing set. Figures 4 and 5 show a consensus matrix generated by combining the similarity matrices from all 20 contexts when training using 4 pairs of objects.

C. Generating Matchings

The robot was tasked with generating matchings among the objects in the four Montessori toys. The objects were split into two groups of six and the robot was tasked with selecting one object from each group to generate a match. This split into two groups of six is naturally suggested by the Montessori toys. For example, the sound cylinders have either red or blue caps; the pressure cylinders have either black or white buttons (see Fig. 2).

More formally, given a 6x6 non-symmetric similarity matrix $\mathbf{W}^c$ or a consensus matrix $\mathbf{W}$ and objects $\mathcal{O}$ partitioned into two sets of equal size $\mathcal{O}_a$ and $\mathcal{O}_b$, matches were generated by picking pairs that maximized similarity between the objects in the pair and minimized similarity between those objects and the remaining objects. One such matrix is shown in Fig. 4. Formally, the objects $i \in \mathcal{O}_a$ and $j \in \mathcal{O}_b$ that maximize

$$q(i, j, \mathbf{W}) = W_{ij} - \gamma \left[\sum_{k \in \mathcal{O}_b / j} W_{ik} + \sum_{k \in \mathcal{O}_a / i} W_{kj} \right]$$

were selected and then removed from $\mathcal{O}_a$ and $\mathcal{O}_b$. The first term captures the pairwise similarity between objects i and j; the last term captures the pairwise similarity between objects i and j and the rest of the objects. The constant γ is a normalizing weight, which ensures that this function is not biased toward any of the terms. In our case, it was set to

$$\gamma = \frac{1}{2(|\mathcal{O}| - 1)}.$$

This process was repeated until no more objects remained to be paired.

D. Evaluation

Given a set of objects (e.g., the weight cylinders), the robot's model was queried in order to group the objects into pairs. Five interactions were randomly picked for each object from the set of ten interactions that were performed on each object and used to create the weight matrix $\mathbf{W}^c$ for each sensorimotor context $c \in \mathcal{C}$. Consensus matrices $\mathbf{W}$ were generated using the three methods described above for a given set of contexts. Matchings were then generated using the method described above. This process was repeated 100 times for every group of contexts. For each size from 1 to $|\mathcal{C}|$, 100 sets of contexts were randomly generated and tested $(1,721$ in total)[1]. Results are reported as the average accuracy or as Cohen's kappa statistic [23] over all 100 iterations. Accuracy is computed as

$$\%Accuracy = \frac{\#correct\ matchings}{\#total\ matchings} \times 100.$$

The kappa statistic is computed as

$$kappa = \frac{P(a) - P(e)}{1 - P(e)}.$$

In our experiments, $P(a)$ is the pairing accuracy of the robot and $P(e)$ is the accuracy a random matching would be

[1]For sets of size 1, $|\mathcal{C}| - 1$, and $|\mathcal{C}|$ all sets of that size were tested since there were fewer than 100 sets of those sizes.

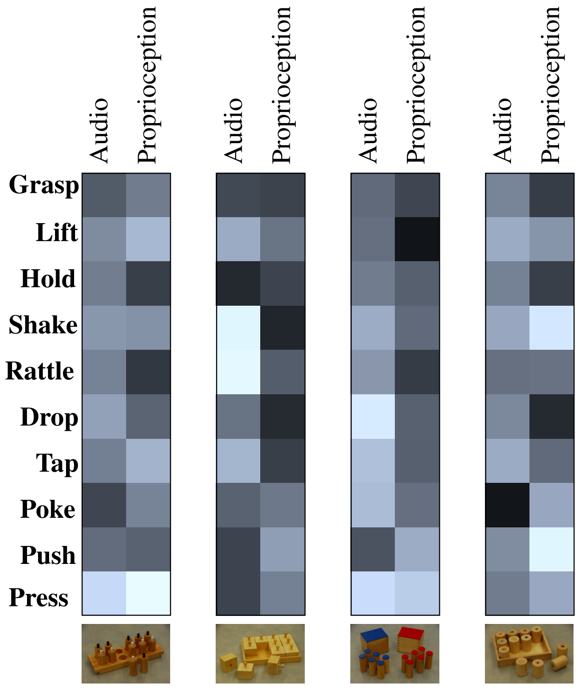

Fig. 6. The accuracy of each context when matching between two sets of six objects. Lighter values indicate higher accuracy with completely white being 100%. Darker values indicate lower accuracy with completely black being 0%. The images from left to right are: *pressure cylinders*, *sound boxes*, *sound cylinders*, and *weight cylinders*.

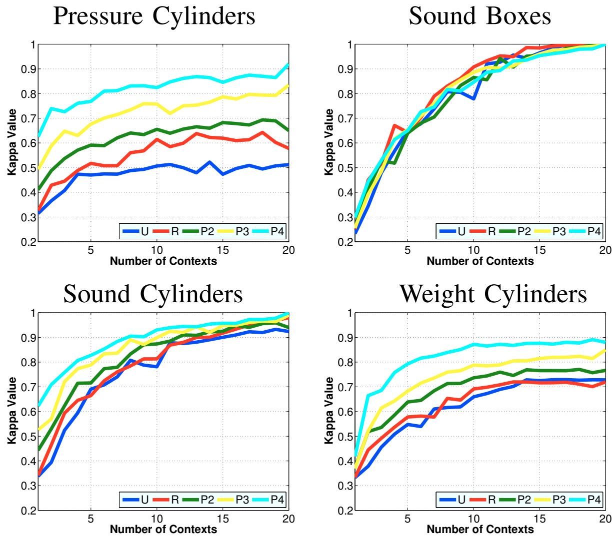

Fig. 7. The kappa statistic for each set of objects. Each line represents a different method for combing the sensorimotor contexts. The line labels are as follows: U-uniform combination; R-recognition accuracy based combination; P2-pairing accuracy using two pairs for training; P3-pairing accuracy using three pairs for training; P4-pairing accuracy using four pairs for training.

expected to get. Kappa is used to allow for direct comparisons between the different sensorimotor context combination methods, since for the pairing accuracy based method, chance accuracy is different than it is for the other methods. The kappa statistic controls for chance accuracy.

The evaluation was performed off-line after the robot interacted with all 48 objects (4 Montessori tasks $\times$ 12 objects in each).

VI. Results

A. *Object Matching with a Single Context*

Figure 6 shows the matching accuracy for each context for all four Montessori tasks. For the *pressure cylinders*, the best sensorimotor context was *proprioception-press* (97.5% pairing accuracy), which was expected. Surprisingly, *audio-press* also did well (80.7%), which was not expected since (at least to the authors' ears) all the cylinders sound the same when pressed. Also interesting is the *audio-drop* context for the *sound cylinders* (89.3% accuracy), which outperformed both *shake* (60.3%) and *rattle* (51.3%) behaviors for *audio*. *Audio-press* (82.3%) for the *sound cylinders* also did well, which is likely due to the fact that they would fall over while being pressed. It is also worth noting that for the *weight cylinders*, the best contexts were *proprioception-shake* (87.7%) and *proprioception-push* (94.3%) rather than contexts that more directly measure the weight such as *proprioception-lift* (50.7%) and *proprioception-hold* (18.8%).

In summary, the robot was able to identify the relevant behaviors and sensory modalities and use them to pair the objects in each of the four Montessori tasks with a high degree of accuracy.

B. *Object Matching with Multiple Contexts*

Figure 7 shows the kappa statistic for each set of objects as the number of contexts is varied from 1 to 20. The graphs show that as the number of sensorimotor contexts used to perform matching increases, so does the kappa statistic. In all cases, the pairing accuracy based combination using four pairs for training (the cyan line) outperforms all the other combination methods. The only exception to this is for the *sound boxes*, since accuracy reaches 100%, all methods reach a kappa value of 1.0. In most cases, the pairing accuracy based combination using three pairs for training (the yellow line) also outperforms the other methods (except for the method that uses four pairs for training). The pairing accuracy based combination using two pairs for training performs about the same as the recognition accuracy combination method, which usually performs slightly better than the uniform combination method. All the combination methods perform better than chance for all object sets, which is indicated by a 0.0 kappa value.

C. *Repeating the Same Behavior*

In all results reported up to this point, five interactions were randomly chosen from the ten for each object during each iteration. Figure 8 shows the average kappa statistic as the number of trials vary, averaged over all the sets of objects and number of contexts. The accuracies quickly converge after only a few trials, implying that repeating the same behavioral repertoire multiple times on an object has quickly diminishing returns. In most cases and for all combination methods, after four repetitions there is very little gain. Diminishing returns is most quickly realized for the pairing accuracy combination method using four pairs for training. The largest gain when increasing interactions was realized by the uniform combination method. This suggest that the uniform combination method benefited the most from a decrease in noise due to its lack of weighted preferences between the contexts, whereas the pairing accuracy combination methods didn't benefit as much because the weights assigned to each context already decreased the noise.

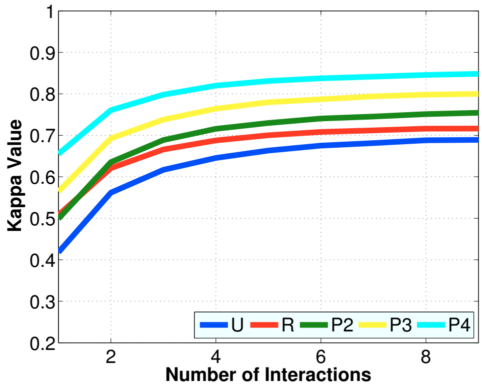

Fig. 8. The kappa statistic averaged across all four sets of objects while varying the number of interactions used to generate the similarity matrices $\mathbf{W}^c$ for each context $c \in \mathcal{C}$. The number of randomly sampled interactions was varied from 1 to 9. The line labels are the same as in Fig. 7.

VII. Conclusion and Future Work

This paper demonstrated a framework that allows a robot to solve object matching tasks by estimating the pairwise similarity of objects in specific sensorimotor contexts. The performance of this framework was evaluated with four standard Montessori tasks that require pairing a set of objects based on their perceived similarities across multiple sensory modalities. The results showed that for a given set of objects, certain contexts are best suited to extract the information necessary to perform object pairing (e.g., *audio-shake* for the *sound boxes*), while others are not useful for that set of objects (e.g., *proprioception-lift* for the *sound cylinders*).

The robot was also able to combine similarity measures from different contexts using three different methods: uniform combination, recognition accuracy based combination, and pairing accuracy based combination. The robot was able to achieve the best performance in almost every case when it was allowed to train on four of the six object pairs before being tested on the remaining two. These results show that embodied sensorimotor similarity measures between objects can be extremely useful for performing matching tasks.

This paper introduced the domain of Montessori tasks to the field of robotics and showed how embodied learning could be used to solve object pairing tasks. For each set of objects the robot learned which set of contexts are most useful for pairing the objects and which are not. The objects in each Montessori task implicitly capture an important concept that the robot can discover on its own through sensorimotor exploration. In the future similar tasks could be used to teach robots not only matching skills, but also important concepts such as ordering, sorting, and relating.

Future work can also expand upon this research by improving the feature extraction methods, the similarity measure, the combination methods, or by using a better matching algorithm. It would also be useful to develop methods that can discover novel exploratory behaviors. This framework can also be applied to other tasks such as object categorization and object recognition. For example, a robot could match previous experiences with objects with new experiences in order to label the objects.

References

[1] M. Montessori, *The Montessori method*. Frederick A. Stokes Co., New York City, USA, 1912.
[2] A. Lillard, *Montessori: The science behind the genius*. Oxford University Press, New York City, USA, 2008.
[3] A. Lillard and N. Else-Quest, "The early years: Evaluating Montessori," *Science*, vol. 313, no. 5795, pp. 1893–1894, 2006.
[4] M. Pitamic, *Teach Me to Do It Myself*. Elwin Street Productions, London, UK, 2004.
[5] J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and A. Stoytchev, "Interactive object recognition using proprioceptive and auditory feedback," *The International Journal of Robotics Research*, vol. 30, no. 10, pp. 1250–1262, 2011.
[6] K. Saenko and T. Darrell, "Object category recognition using probabilistic fusion of speech and image classifiers," in *Machine Learning for Multimodal Interaction: 4th International Workshop, MLMI 2007, Brno, Czech Republic, June 28-30, 2007, Revised Selected Papers*, vol. 4892. Springer-Verlag, New York, 2008, pp. 36–47.
[7] F. Endres, C. Plagemann, C. Stachniss, and W. Burgard, "Unsupervised discovery of object classes from range data using latent Dirichlet allocation," in *Proc. of RSS*, 2009, pp. 113–120.
[8] J. Sinapov and A. Stoytchev, "The odd one out task: Toward an intelligence test for robots," in *Proceedings of the 9th IEEE International Conference on Development and Learning (ICDL)*, 2010, pp. 126–131.
[9] K. Rathunde and M. Csikszentmihalyi, "Middle school students motivation and quality of experience: A comparison of Montessori and traditional school environments," *American Journal of Education*, vol. 111, no. 3, pp. 341–371, 2005.
[10] M. Daehler, R. Lonardo, and D. Bukatko, "Matching and equivalence judgments in very young children," *Child Dev.*, pp. 170–179, 1979.
[11] A. Leslie and M. Chen, "Individuation of pairs of objects in infancy," *Developmental Science*, vol. 10, no. 4, p. 423, 2007.
[12] B. Younger, "The segregation of items into categories by ten-month-old infants," *Child Development*, pp. 1574–1583, 1985.
[13] W. McPherson and P. Holcomb, "An electrophysiological investigation of semantic priming with pictures of real objects," *Psychophysiology*, vol. 36, no. 01, pp. 53–65, 1999.
[14] S. Nolfi and D. Marocco, "Active perception: A sensorimotor account of object categorization," in *From Animals to Animats: The Seventh International Conference on the Simulation of Adaptive Behavior*, 2002, pp. 266–271.
[15] L. Natale, G. Metta, and G. Sandini, "Learning haptic representation of objects," in *Proceedings of the International Conference on Intelligent Manipulation and Grasping*, 2004.
[16] T. Nakamura, T. Nagai, and N. Iwahashi, "Multimodal object categorization by a robot," in *Proceedings of the IEEE/RSJ IROS*, 2007, pp. 2415–2420.
[17] S. Takamuku, K. Hosoda, and M. Asada, "Object category acquisition by dynamic touch," *Advanced Robotics*, vol. 22, no. 10, pp. 1143–1154, 2008.
[18] J. Sun, J. Moore, A. Bobick, and J. Rehg, "Learning visual object categories for robot affordance prediction," *The International Journal of Robotics Research*, vol. 29, no. 2-3, p. 174, 2010.
[19] J. Sinapov, C. Schenck, K. Staley, V. Sukhoy, and A. Stoytchev, "Grounding semantic categories in behavioral interactions: Experiments with 100 objects," *Robotics and Autonomous Systems (to appear)*.
[20] A. Chan and E. Pampalk, "Growing hierarchical self organizing map (GHSOM) toolbox: visualizations and enhancements," in *Proc. of the 9th NIPS Conference*, vol. 5, 2002, pp. 2537–2541.
[21] G. Navarro, "A guided tour to approximate string matching," *ACM computing surveys (CSUR)*, vol. 33, no. 1, pp. 31–88, 2001.
[22] J. Sinapov and A. Stoytchev, "The boosting effect of exploratory behaviors," in *Proc. of the 24th National Conference on Artificial Intelligence (AAAI)*, 2010, pp. 1613–1618.
[23] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and Psychological Measurement*, vol. 20, no. 1, pp. 37–46, 1960.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Using Depth to Increase Robot Visual Attention Accuracy during Tutoring

Christian I. Penaloza, Yasushi Mae, Kenichi Ohara and Tatsuo Arai
Graduate School of Engineering Science
Osaka University
1-3 Machikaneyamacho, Toyonaka, Japan
http://www-arailab.sys.es.osaka-u.ac.jp

***Abstract*—This paper explores the problem of attention models for robot tutoring as related to the cognitive development of infants. We discuss the factors that have an important influence in infants' attention and the way these factors can be taken into consideration to develop robot attention models that simulate infants' cognitive stimuli. In particular, we focus on the attention given to objects that appear closer to the infant when they are shown by an adult. Using the distance of an object as an important factor to increase visual attention, our model uses depth information along with the well-known Bottom-Up Visual Attention Model Based on Saliency (Itti & Koch, 2001) in order to increase attention accuracy even if non-salient feature objects are shown to the robot or if tutoring activity takes place under cluttered environments. Our model also considers the presence or absence of a human tutor to decide whether a tutoring activity might take place. Experimental results suggest that depth information is a key factor to emulate effective infants' attention.**

I. Introduction

One of the main objectives of researchers in the area of cognitive robotics is to design mechanisms to provide robots with human-like abilities in perception, decision making, reasoning, and action execution. Among the many challenges of developmental cognition for robots, attention is perhaps one of the most important challenges that needs to be addressed since it plays a very important role in the process of learning.

The study of attention of infants can provide important clues to develop systems that emulate this important ability. This is because even before babies with normal vision can talk or walk, they are able to perceive and parse their visual environment and are able to move their eyes and head to select visual targets (objects or people) [7]. Moreover, by observing the cognitive development of infants when they interact with their parents, it has been shown that infants' attention and learning are favorably influenced by factors such as motionese (e.g., exaggeration of parent's actions) [8] or contingent reactions (helping infants find proper association) [10]. In this direction, researchers have developed robotic systems that emulate attention and learning processes of infants by using socially guided exploration [2], dialogs [4], or motionese [6]. Human guided instruction plays a very important role in infants learning and can be simulated in robot systems by presenting a visual task or object within robot's visual field as shown in Fig. 1 (a).

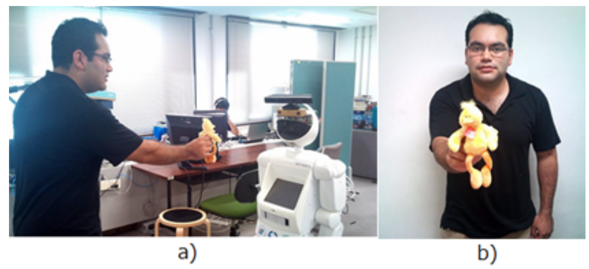

Fig. 1. 1) Human Guided Robot Learning. 2) Controlled environment (salient object and plain background) commonly used for experiments.

One of the main difficulties of robot learning is the fact that robots do not know where to look at when observing a demonstration [6]. Researchers have proposed computer vision models of attention that enable the robot to selectively choose a relevant visual segment while ignoring others (e.g., [19]-[22]). The *Bottom-Up Visual Attention Model Based on Saliency* originally proposed by Itti & Koch [20] is perhaps one of the most used and widely accepted models of attention. This model proposes the idea that visual attention is attracted by salient stimuli that 'pop out' from their surroundings due to primitive features such as color, intensity and orientation. This model is commonly used in robotic systems to achieve visual attention during a tutoring activity (e.g., [6], [25]). However, robotic systems that use this model are usually evaluated with objects that have strong salient features or in controlled environments with plain backgrounds (e.g., Fig. 1 (b)). This certainly facilitates the tutoring activity but limits its applicability to experimental setups and cannot be used in real environments. Moreover, since the Bottom-Up Attention Model uses 2D images to find 'salient' features to define focus of attention, distance of the object is not considered due to the lack of depth (3D) information.

The distance of an object to the infant is also a very important factor that effects visual attention. Smith *et al.* [1] provide experimental evidence that demonstrates that visual attention is largely increased by bringing objects close to the child. Therefore, this factor should certainly be used to improve attention mechanisms for robots. In this paper, we put this concept into practice and developed an attention model that uses depth information along with the Bottom-Up Visual Attention Model Based on Saliency in order to

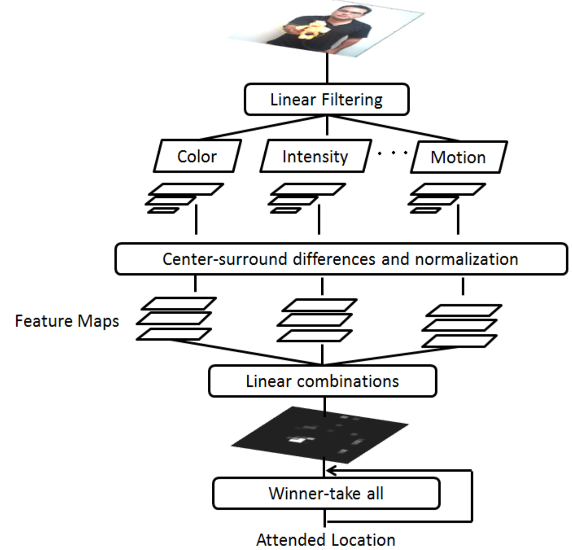

Fig. 2. Schematic of salience model proposed by Itti & Koch [20].

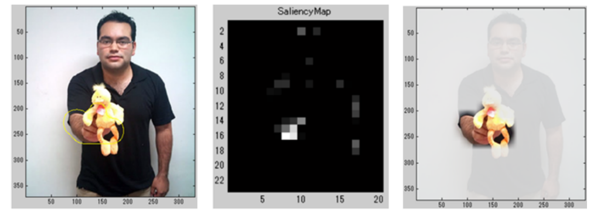

a) b) c)

Fig. 3. Tutoring environment: a) Plain background - salient-feature object. b) Saliency Map c) Focus of attention. Attention model correctly locates object of interest.

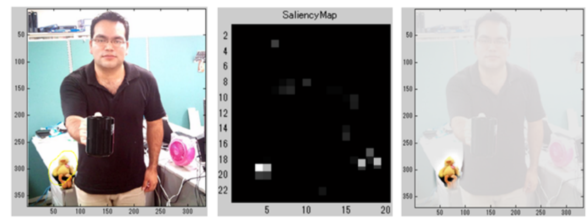

a) b) c)

Fig. 4. Tutoring environment: a) Cluttered background, non-salient object, presence of distracters (salient-feature objects in background) b) Saliency Map c) Focus of attention. Attention model fails to locate the object of interest that is shown by the human tutor.

increase attention accuracy when the robot observes closer objects with or without salient features. We also take into consideration the presence or absence of a human teacher to activate the attention model when a tutoring activity takes place. Most importantly, our attention model can be used in real environments with cluttered backgrounds.

The remaining part of this work is organized as follows. Section II introduces a discussion of the Bottom-Up Visual Attention Model Based on Saliency and the need to use depth information for attention models. A description of our attention model is described in the subsequent section. Experiment design and results are given in section IV. Finally, in the last section we present some conclusive remarks.

II. A Bottom-Up Visual Attention Model Based On Saliency

Inspired by the behavioral and neuronal mechanism of primates, the *Bottom-Up Visual Attention Model Based on Saliency* uses the ”outstandingness” of primitive features of an image to be able to detect salient locations in a scene [20]. For example, a yellow object in a black background is detected as salient because of its distinctive color. A person moving to the right direction among other persons moving to the left direction is detected as salient with respect to motion direction.

This model is probably the most influential attention model, since it has been extensively used in many research fields including computer vision and robotics [11]. This model (Fig. 2) uses several concepts (e.g., feature map, saliency map) and proposes a well-structured process for calculation of the saliency map which defines attention focus. As a brief summary, multi-scale analysis of an input image is performed to evaluate five primitive features: color, intensity, orientation, flicker, and motion. Individual feature maps are combined to create a centralized saliency map that is used to identify the focus of attention. Refer to the original paper [20] for a more detailed description.

Contrary to the top-down attention model (an active scan of the visual field in search of a pre-specified object or stimuli), the bottom-up approach guides visual exploration focusing on the most salient stimuli - in a similar way babies do in early stages of development - and therefore it is more appropriate for emulating infant behavior. However, due to the native process of saliency computation from 2D images, the Bottom-Up Visual Attention Model Based on Saliency is unable to cope with attention focus based on depth information, which is also a key factor to effectively emulate infant attention.

A. *Importance of Depth Information in Tutoring Activities*

When an adult is tutoring an infant about an object or a particular task, the principle of *overt attention* (to place an object of interest at the center of visual field), along with the distance of the object are generally used to increase visual attention [13] - as demonstrated by experimental evidence of Smith *et al.* [1].

For tutoring activities, several researchers have emulated infants' attention by successfully applying the Bottom-Up Visual Attention Model Based on Saliency in robotic systems. Since this model is meant to find 'salient' features in the scene, most of the times the tutoring activity takes place in experimental environments (plain backgrounds scenes or use of objects with 'salient' features such as bright colors), which certainly facilitate the learning task (i.e. Fig. 3). However, a real environment such as the one presented in Fig. 4 (a) (cluttered background, presence of multiple objects with salient features, or teaching an object that lacks salient features) presents a bigger challenge to the Bottom-Up visual model, which is unable to locate the object presented by the human tutor since other salient-featured objects are present in the background. In this case, depth information plays a very important role in defining where the focus of attention should be located.

III. Attention Model using Depth

In order to deal with the difficulties presented in the previous section, our model uses depth information along with the Bottom-Up Visual Attention Model to be able to cope not only with feature saliency but also with object proximity. The main purpose is to emulate infants' attention when objects are presented - by a human tutor - at a close distance during a tutoring activity even if the objects lack strong feature saliency.

Attention models that use depth information to define focus of attention have been previously introduced by researchers (e.g., [14]-[17]) for scene analysis applications. In order to define saliency, mentioned techniques commonly use the 3D structural information of objects or object's relative position to other objects. Since none of these models takes into consideration the presence or absence of a person, it is difficult to implement them to emulate infant tutoring in robots because these models would encounter important difficulties such as focusing on the person vs object or detecting saliency effectively in extreme cluttered environments.

Since the main objective of this research is to emulate infants' stimuli, the particularity of our model is that we also take into consideration the presence or absence of a human teacher to activate the attention model, and we use *proxemics* theory according to Hall [23] to pre-define a distance range to which robots should pay attention when a tutoring activity takes place.

A. Development Process

During the tutoring activity, a human actively teaches an object to a robot by using the principle of *overt attention* (the object is presented within the visual field of the robot at a close distance). On the robot's behalf, our attention system that uses depth information along with the Bottom-Up Visual Attention Model is activated when a human teacher is found within its field of view. This is when the robot knows that a tutoring activity might take place. On the other hand, when a human teacher is not present or when there is no object close to the robot, only the Bottom-Up Visual Attention Model is activated.

Another way to look at our approach is by considering depth as an extra channel of the saliency map defined in the Bottom-Up model, but using a *binary* weight applied to the depth channel as described in Fig 5. This binary weight would be 0 when no human is present within the field of view of the robot and 1 otherwise - this can be considered as top-down influence of a human-detection. In other words, when a human is present, depth pixels within the personal space of the robot (if any) will be taken into consideration along with the pixels of salient features (color, intensity, orientation, flicker, and motion) of the Bottom-Up model. However, if no human is present, only the salient features of the Bottom-Up model are used to define attention location.

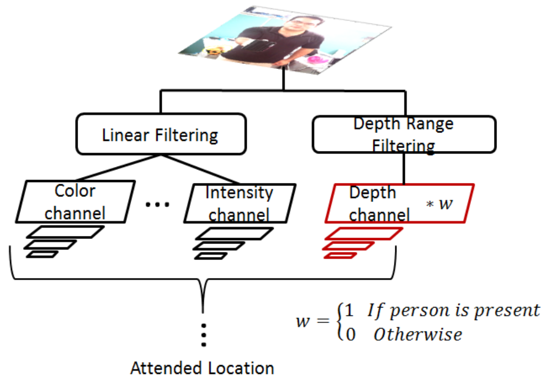

Fig. 5. Considering depth as an extra channel to saliency map defined in the Bottom-Up model, but using a *binary* weight applied to the depth channel.

Designation	*Specification*	*Usage*
Intimate distance	0 - 0.45m	Embracing or touching
Personal distance	**0.45 - 1.20m**	**Friends**
Social distance	1.20 - 3.60m	Acquaintances and strangers
Public distance	>3.60m	Public speaking

TABLE I
THE FOUR SPHERES OF PHYSICAL DISTANCE CORRESPONDING TO SOCIAL DISTANCE ACCORDING TO HALL [23].

In order to be able to recognize a human teacher, we used the built-in capabilities of our Kinect sensor through the Software Development Kit freely provided by Microsoft [24]. The main process of human body (pose) detection is explained in detail in [18]. As a brief overview of their method, the authors use a single depth image to accurately predict 3D positions of body joints by designing an intermediate body parts representation that maps the difficult pose estimation problem into a simpler per-pixel classification problem. Subsequently, they use a 'dictionary' of 3D pose proposals and find the closest match. Finally, they generate confidence-scored 3D proposals of several body joints by re-projecting the classification result and finding local modes.

In our model, once the human teacher is recognized, the depth-based attention is activated and the principle of depth-based saliency is performed. In this principle, a particular distance range is pre-defined and objects that appear within that range are given attention priority over objects that appear farther away from the robot even if they have stronger salient features than the ones of the object within the pre-defined distance range.

In order to define the most appropriate distance range for attention focus, we looked into social robotics literature and refer to the principle of *proxemics* — physical and psychological distancing from others. According to Hall [23], the four spheres of physical distance corresponding to social distance can be defined as described in Table I.

We chose personal distance as the most appropriate range for the tutoring activity, since objects within the intimate space appear too close to the camera and too far in the social space. The Kinect sensor has a depth distance limitation in which the minimum detection distance is 0.45m. Therefore, for our

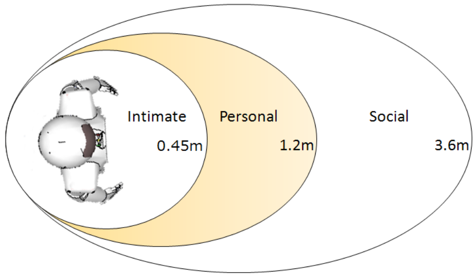

Fig. 6. Proxemics - From the four spheres of physical distance, personal distance was chosen as the most appropriate range for the tutoring activity.

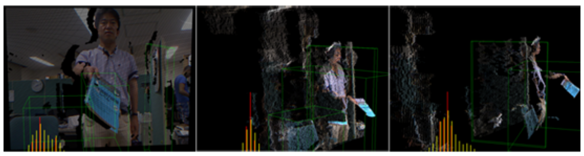

Fig. 7. Depth-RGB graphic representation as perceived by robot when human teacher demonstrates an object.

tutoring activity we defined the robot's personal space from 0.45m to 1.20m, as seen in Fig 6.

Kinect sensor provides valuable depth data that can be easily analyzed. Figure 7 shows a graphic representation of the visual Depth-RGB combination in which a human teacher is demonstrating an object to the robot. Our approach consists in extracting the RGB information corresponding to the object that appears within the personal distance range.

One way to detect objects within personal space is to perform depth-based thresholding. This involves estimating the depth value of each of the pixels that appear in the depth image and labeling those pixels whose z-value (depth) appears within the predefined distance range, as shown in Fig. 8 (d). Finally, target depth value pixels conform a pixel region that serves as a visual mask to the RGB input image to extract original color pixels of the attended object as observed in Fig. 8 (c).

The architecture of our attention model is described in Fig. 9. Our system integrates Bottom-Up Visual Attention Model with depth information by receiving input RGB and Depth images and use them to define whether a person is present or not, and decide the attention location based on the salient features of objects and object's distance to the robot.

IV. Experiment and Evaluation

A. Experimental Setting and Task

This section presents the experiment carried out to validate our attention model. The main objective of the experiment is to compare the attention accuracy during a tutoring activity in three cases: 1) using only the Bottom-Up Visual Attention Model, 2) using only depth-based attention model, and 3) using our model that combines both approaches.

The tutoring task consisted on a human volunteer presenting two types of objects to the robot: 1) objects with salient features (e.g. bright colors) and 2) objects with non-salient

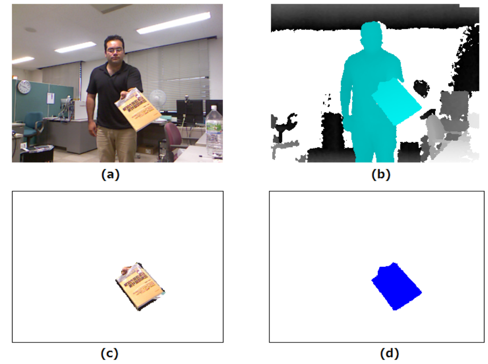

Fig. 8. Visual representation of our Depth-Based Attention Model: (a) Input RGB image, (b) Depth view - Human Detection, (c) Focus of Attention, (d) Personal space view.

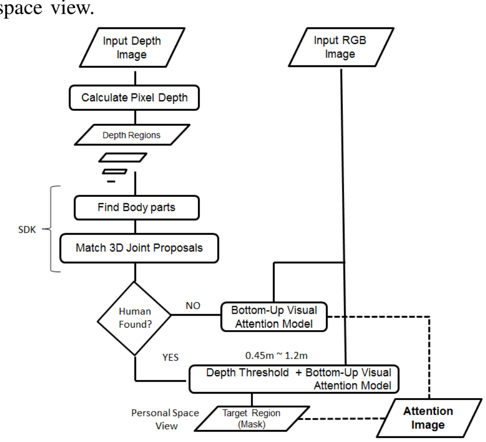

Fig. 9. Attention model that integrates Bottom-Up Visual Attention Model with depth information in order to define focus of attention.

features. In total, six objects (3 salient and 3 non-salient) shown in Fig. 10 were used in the experiment.

The experiment was performed in an ordinary room with no special pre-arranged settings such as plain backgrounds. In fact, our experimental setting contains cluttered background and objects with salient features (i.e. lamp, monitor) that may serve as distractors during the tutoring task.

The experiment was divided into two phases: 1) demonstrating objects with salient features and 2) demonstrating objects with non-salient features. Fig. 11 shows actual experiment images in which the volunteer holds the objects in front of the robot. It can be noticed that non-salient feature objects are difficult to distinguish from 2D image.

In each experimental phase, 3 tutoring tasks with corresponding objects were performed. Each tutoring task lasted 10 seconds and consisted on the following actions:

1) Volunteer stood 1.2m∼1.5m distance from the robot (2 sec).
2) Volunteer performed object demonstration by presenting the object within robot's personal space (6 sec).
3) Volunteer finished demonstration and stepped out of robot's field of view (2 sec).

It is worth mentioning that volunteer was not previously instructed how to perform object demonstration. Volunteer was

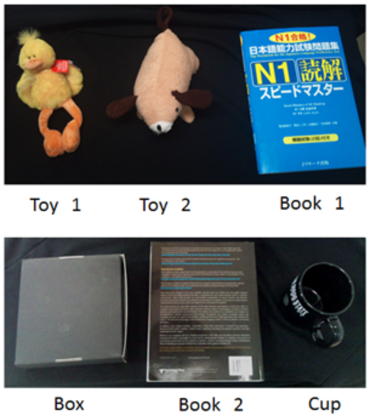

Fig. 10. Experiment objects - upper: objects with salient features, lower: objects with non-salient features.

Fig. 11. Phase 1: experiment using objects with salient features. Phase 2: experiment using objects with non-salient features.

free to present the object by lifting the object to desired height, holding the object with one or two hands, move the object in front of the robot, or keep the object still.

B. Evaluation

In order to evaluate which aspects of the demonstration were detected by the attention model, attention locations were classified into four regions: object, tutor's hand, tutor's face, and others (i.e., background objects). Figure 12 (c) shows examples of the classification regions.

Region classification was performed for every frame by examining the center region (20x20 pixel) of the attention image obtained by each attention model. Figure 12 shows the attention region result obtained by the Bottom-Up Attention Model (a) and Depth-base attention model (b). Each center region was classified as object or not depending on whether it was the same color as the object (for salient objects) and by visual inspection (for non-salient objects). Center regions with skin color were categorized as face or hands. Face and hands were then distinguished by the relative position in which hand position is usually lower than the face.

Attention analysis was performed by comparing how often the focus of attention was brought to object, tutor's hand, face or other, using salient and non-salient feature objects.

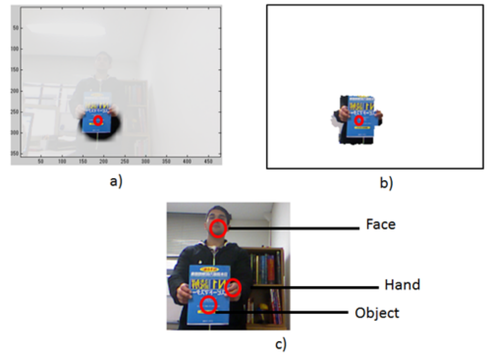

Fig. 12. Example attention regions detected by (a) Bottom-Up Attention Model and (b) Depth-based Attention Model. (c) Classification of attended locations.

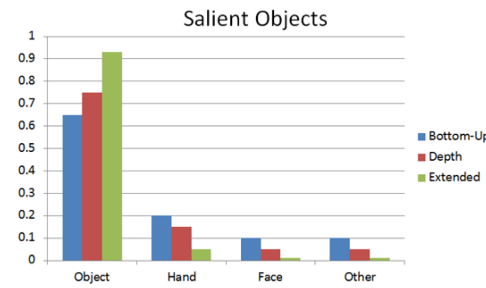

Fig. 13. Results of experiment using objects with salient features. Note that Bottom-Up attention model had a better accuracy in focusing on the demonstrated object compared to the performance of the same model in Fig. 14.

C. Results

Figures 13 and 14 present the proportion of attention of both phases: 1) using salient feature objects and 2) using non-salient feature objects. Each color bar represents the mean proportion of the attention during the three tutoring tasks using a particular type of object: Blue- using only Bottom-Up Visual Attention Model based on Saliency, Red- Depth-based attention model and Green- attention model that combines both approaches.

In Fig. 13 it can be noticed that Bottom-Up attention model had a better accuracy in focusing on the demonstrated object compared to the performance of the same model in Fig. 14. This was mainly due to the fact that salient-feature objects were easier to detect as compared to non-salient feature objects. An interesting point is that Bottom-Up model was able to focus on the object for some short period of time even with non-salient objects. This may be the result of the object movement done by the volunteer during the demonstration. Therefore, we can confirm that motionese is also a very important factor that defines the visual focus of attention.

In Fig. 14, it can also be noticed that the Bottom-Up attention model was highly distracted by the hand, face and background. On the other hand, depth-based attention model alone performed fairly well during the demonstration of salient and non-salient feature objects. This result seems reasonable since most of the times the object was demonstrated within the robot's pre-defined depth threshold distance. However, we

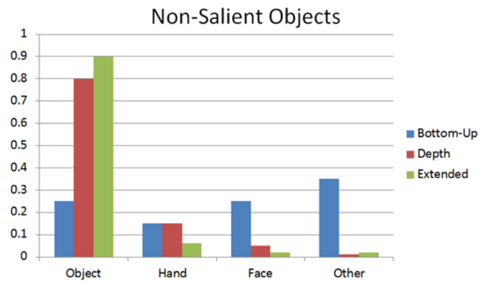

Fig. 14. Results of experiment using objects with non-salient features. Note that Bottom-Up attention model was highly distracted by the hand, face and background. Depth-based attention model alone performed fairly well during the demonstration of salient and non-salient feature objects. Our approach that combines depth information along with the Bottom-Up Attention Model represented by the green bar demonstrates higher attention accuracy located in the demonstrated object.

can notice a small proportion of attention directed to the hands of the volunteer that may have held the object with two hands or with one hand covering part of the object.

Finally, our proposed attention model that uses depth information along with the Bottom-Up Attention Model represented by the green bar demonstrates higher attention accuracy located in the demonstrated object. While the performance does not differ too much from the depth based attention model, the improvement may have been caused by using the Bottom-up attention model to find the salient feature object even when the volunteer did not present the object within the pre-defined depth threshold distance.

V. Conclusion

In this paper we discussed the factors that have an important effect in infants' attention and the way these factors can be taken into consideration to develop robot attention models that simulate infants' cognitive stimuli. We proposed an attention model that uses depth information along with the Bottom-Up Attention Model based on Saliency to increase attention accuracy of objects during a tutoring task when a human tutor is present. Our model can be used for robots to locate the focus of attention in objects that are presented at a close distance, even if objects do not have salient features. Experimental results show that depth information plays an important role for defining the focus of attention of systems that emulate the cognitive development of infants.

VI. Future Work

In future work we will perform experiments with a richer variety of objects, and we will compare saliency performance across multiple volunteer tutors.

Acknowledgment

This work was supported in part by Grant-in-Aid for Scientific Research (C) 23500242.

References

[1] Smith, L.B., Yu, C. and Pereira, A.F. (2007). From the Outside-In: Embodied Attention in Toddlers. In Proceedings of 9th European Conference of Artificial Life (ECAL2007) (pp. 445-454).
[2] C. Breazeal and A. L. Thomaz. "Learning from human teachers with socially guided exploration." In Proceedings of the International Conference on Robots and Automation (ICRA), 2008
[3] A. L. Thomaz and M. Cakmak. "Learning about Objects with Human Teachers." In Proceedings of the International Conference on Human Robot Interaction (HRI), 2009
[4] A. Vogel, K. Raghunathan, D. Jurafsky "Dialog with Robots". In *AAAI* 2005.
[5] Mansur, A. Sakata, K. Rukhsana, T. Kobayashi, Y. Kuno. "Human robot interaction through simple expressions for object recognition". The 17th IEEE International symposium on Robot and Human Interactive Communication. *RO-MAM* 2008.
[6] Y. Nagai and K. J. Rohlfing, "Computational Analysis of Motionese Toward Scaffolding Robot Action Learning," IEEE Transactions on Autonomous Mental Development, vol. 1, no. 1, pp. 44-54, May 2009
[7] Y. Nagai, A. Nakatani, and M. Asada, "How a robot's attention shapes the way people teach," in Proceedings of the 10th International Conference on Epigenetic Robotics, pp. 81-88, November 2010
[8] R.J. brand, D.A. Baldwin, and L.A. Ashburn, "'Evidence for 'motionese':Modifications in mothers' infant-directed action," Developmental Science, vol. 5, no.1, pp.72-83, 2002.
[9] K.J. Rohlfing, J.Fritsch,B. Wrede, and T. Jungmann, "How can multimodal cues from child-directed interaction reduce learning complexity in robots?" Adv, Robot., vol.20, no. 10. pp. 1183-1199, 2006.
[10] Y. Nagai and K. J. Rohlfing, "Parental action modification highlighting the goal versus the means," in Proc. IEEE 7th Int. Conf. Develop. Learning, 2008.
[11] Begum, M.; Karray, F.; , "Visual Attention for Robotic Cognition: A Survey," Autonomous Mental Development, IEEE Transactions on , vol.3, no.1, pp.92-105, March 2011
[12] C.Koch and S. Ullman, "Shifts in selective visual attention: Toward the underlying neural circuitry," Human Neurobiol., vol. 4, pp. 219-227, 1985.
[13] R. Bajscy, "Active perception," Proc. IEEE, vol. 76, pp. 996-1005, 1988.
[14] T.Jost, N.Ouerhani, R.Wartburg, R. Muri, H. Hugli, "Contribution of Depth to Visual Attention", Proc. Early Cognitive Vision Workshop, 2004.
[15] T.Jost, N.Ouerhani, R.Wartburg, R. Muri, H. Hugli, "Computing Visual Attention from Scene Depth", Proc. International Conference on Pattern Recognition, pp. 375-378, 2000
[16] M.Z. Aziz and B. Mertsching, "Fast Depth Saliency from Stereo for Region-Based Artificial Visual Attention", in Proc. ACIVS, pp.367-378, 2010.
[17] A. Maki, P. Nordlund, J. O. Eklundh, "A computational model of depth-based attention", In Pattern Recognition, 1996., Proceedings of the 13th International Conference on, Vol. 4 (Aug 1996).
[18] Jamie Shotton, Andrew W. Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, Andrew Blake: Real-time human pose recognition in parts from single depth images. CVPR 2011: 1297-1304
[19] L. Itti and C. Koch, "Computational modeling of visual attention," Nature Rev.: Neurosci., vol. 2, pp. 194-203, 2001.
[20] L. Itti C. Koch and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," IEEE Trans. Pattern Analysis and Machine Intelligence. vol 20, no.11, pp.1254-1259, Nov. 1998.
[21] Y.Sun and R. Fisher, "Object-based visual attention for computer vision," Artificial Intelligence vol. 146, pp.77-123, 2003.
[22] S. Frintrop, "VOCUS: A Visual Attention System for Object Detection and goal-directed Search". Heidelberg, Germany: Springer-Verlag, 2006, vol. 3899. LNAI 3-540-32759-2.
[23] Hall, E. T. (1966). *The Hidden Dimension*. New York: Anchor Books.
[24] Microsoft Kinect for Windows SDK BETA from Microsoft Research, http://research.microsoft.com/en-us/um/redmond/projects/kinectsdk
[25] Jonas Ruesch, Manuel Lopes, Alexandre Bernardino, Jonas Hornstein, Jose Santos-Victor, Rolf Pfeifer, "Multimodal saliency-based bottom-up attention a framework for the humanoid robot iCub" In proceeding of: 2008 IEEE International Conference on Robotics and Automation, ICRA 2008, May 19-23, 2008, Pasadena, California, USA.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Keyword Detection in Human-Robot Tutoring Scenarios

Christian Dondrup[§], Katrin Solveig Lohan[*], Joe Saunders[‡], Hagen Lehmann[‡], Chrystopher Nehaniv[‡] and Britta Wrede[§]
[§]Bielefeld University, Research Institute for Cognition and Robotics, Applied Informatics Group
[*]Instituto Italiano di Tecnologia , Robotics Brain & Cognitive Sciences
[‡]University of Hertfordshire, Adaptive Systems Research Group

***Abstract*—We describe a way of narrowing the search space for descriptive keywords during a human-robot tutoring scenario, where the tutor is explaining names and characteristics of objects to the robot, by employing interaction detection techniques. This system detects attention getting behaviour which is derived from mother-infant interactions and extracts the verbal information during these specific time periods, segmenting it and building up histograms to estimate word frequencies and thus word importance. This method should allow us to create a system that does not rely on a dictionary or normal speech recognition to acquire novel word-object relations but only relies on the pure interaction between the robot and a human tutor.**

I. Introduction

In human-robot interaction speech is an important way of communication. To achieve a natural interaction between a human and a robot we have followed the developmental robotics approach [1] with the intend to create a model for keyword acquisition gained from previous research on adult-child interactions. We have previously studied preverbal infants (6 to 8 months) in an interaction with their parents for clues on how infants learn words (see [2], [3]).

Most speech acquisition approaches typically use a predefined dictionary and a common speech recognition algorithm or manual annotation, see for example [4], [5]. These methods are well suited to their application and more or less accurate in their results, however, we want to build a system that is able to learn important keywords on its own. Such a system should be capable of learning online, so we cannot rely on manual annotation. In addition we want to build a system that profits from the learning behaviour shown by preverbal infants to prevent the need for predefined words which are then recognized by the robot.

In this paper we will show a way of reducing the search space for important words in a human-robot tutoring scenario by emulating the behaviour of preverbal infants, thus trying to achieve a tutoring behaviour in the human tutor which is, as similar as possible, to the behaviour of a mother playing with and teaching her child [6]. We will present a specific scenario where the human tutor is teaching the robot some objects, specifically their names, colours and shapes. We will try to describe a way to encourage the tutor to use more descriptive words (e.g. red, small etc.) than filler words (this, and, here etc.). By this we hope to achieve a search space that allows us to identify important words without knowing their meaning and thereby create a way of learning them with some kind of unsupervised learning algorithm.

Fig. 1. One of the participants explaining shapes and colours to the iCub [9] robot. The shape explained is the blue sun inside of an ARToolKit [10] marker for the object detection.

To recognize the situations where the tutor is more likely to use keywords for the object description we will employ a detector that relies purely on the interaction between the tutor and the robot. Afterwards we will have to segment [7] the recorded speech data and identify similar words [8] to determine a word scale to find the most important ones.

II. Adult-Infant Interaction

The following section describes visual clues that allow us to narrow down the keyword search space. All these clues are derived from the interaction between a mother and her preverbal infant (6 to 8 month) [6]. This is due to the fact that we want to teach our robot novel-words in relation to an object, therefore, we will have a look at how mothers teach their infants such object-word relations. The preverbal infant condition was chosen because we want to learn words that are new for our robot (no internal dictionary), so we need to have a look at infants that need very much assistance with their language acquisition.

A. Maternal Object Naming

As Matatyaho and Gogate describe in their paper [6] there is a connection between the movement of the object and naming it as a part of our natural behaviour in a mother-infant teaching scenario.

Matatyaho and Gogate [6] state that preverbal infants learn an object-word relation better if their mothers use attention getting gestures like forward motions and shaking (or waggling) while synchronously naming the object, compared to infants where the mothers did not use such techniques. So Matatyaho and Gogate [6] have shown that uni-modal (visual) properties occur in combination with inter-modal (synchrony) properties or maternal naming. As described above the mothers used showing gestures like forward and shaking motions more often in synchrony with naming the object than in asynchrony (these findings are consistent with the field studies of Zukow-Goldring [11], [12]). As a result of this Matatyaho and Gogate [6] state that these gestures in synchrony with words are naturally effective tools for conveying novel word-referent relations because they likely elicit greater infant joint attention and thereby facilitate the word mapping.

We hope to exploit this teaching technique for our human-robot tutoring scenario by relying on the natural attention getting behaviour of the human tutor. For this purpose we will try to implement a behaviour for the robot that encourages such attention getting gestures and hopefully synchronous object naming, thus narrowing down the search space for meaningful words without knowing what these words actually mean.

B. Looming

As stated by Matatyaho and Gogate [6] mothers use forward and upward/downward[1] and shaking or waggling gestures as attention getting movements. Since we do not care about the position of the tutor in respect to the robot we ignore the upward/downward movement which is combined with the forward movement (Matatyaho and Gogate [6] collapsed forward/downward movements into one because they often co-occurred at the same time) but will just concentrate on the forward movement itself. These forward movements which intend to bring an object into the line of sight of the infant (or robot in our case) are also called *looming*.

If we use this looming behaviour to narrow down our search space, we are more likely to get meaningful information as a result. As a logical consequence we disregard all the other verbal information that is given during the non-looming phases and just process the verbal information given during the looming phases.

C. Robot Behaviour

To induce looming gestures we will have to design a behaviour for our robot that shows some kind of reaction to the looming itself. One way of giving such a feedback would be gazing at the loomed object. This gaze switching to the object and thereby creating a state of joint attention is one of the main features that helps infants learn the relation between the spoken word and the described object [6]. As a result the obvious choice to reward looming would be the joint attention to the object by looking at it. However, to encourage the tutor to use as much looming as possible the robot has to reach a habituation[2] [13] state at some point during a looming gesture and thereby loose interest in the object and show its lack of attention by looking away (at random points for example). This is supposed to trigger as much looming in the tutor as possible, and thus help us gather more meaningful information about the object, by being sensitive to the ostensive stimuli and giving feedback about the capabilities of the robot and thereby creating an environment where the robot is treated infant like [3].

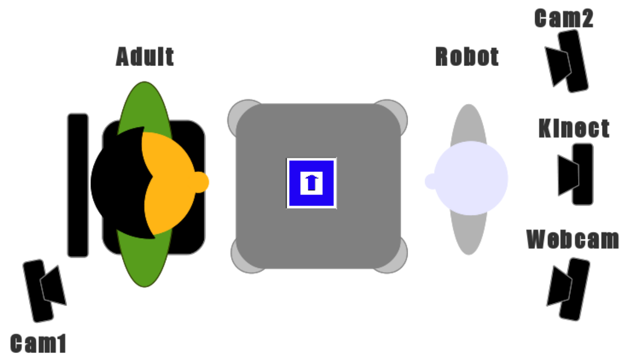

Fig. 2. This figure illustrates the study setting where the human tutor and the robot are placed on opposite sides of a table. The scene is captured by two cameras one facing the human and one facing the robot. In addition we need a Microsoft Kinect for the looming detection (see III-B), a webcam for the object detector and a headset for the voice recording (not shown in this figure). On top of the table is one of the cubes showing the blue arrow shape inside an ARToolKit [10] marker for the object detection.

III. Study

After we have shown how a robot should react and behave to facilitate the acquisition of meaningful data (see II-C) we will now describe a related study which was carried out in the italk project and conducted at the University of Hertfordshire in the beginning of 2012.

A. Parameters

1) Set-Up: We observed 19 participants, which are native English speakers, teaching the iCub robot [9]. The participants were divided into 2 groups which differed in the behaviour the robot showed. The first group was confronted with a random gaze switching, non-responsive[3] robot and the other half taught a robot showing a behaviour according to the Tutor Spotter [14]. The Tutor Spotter tries to create a contingent tutoring environment by showing joint attention according to the gazing

[1] Depending on the position of the tutor in respect to the infant.

[2] Our definition of habituation differs in the way that not repeated but persistent stimulus triggers the habituation and after the stimulus vanishes the system will immediately recover from said habituation.

[3] Less contingent, does not respond to gazing and looming behaviour of the tutor.

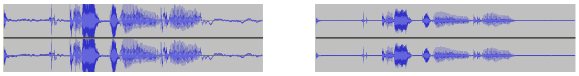

(a) Original signal of three words in one utterance recorded by a headset. Still pretty noisy and distorted. The boarders between the last two words are not very clear.

(b) Signal from Figure (a) after passing through the pre-emphasis and band-pass filter. We can now even make out the borders between all three words.

Fig. 3. Original and filtered Input signal.

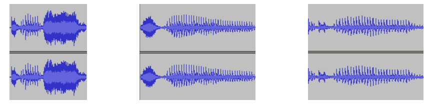

(a) (b) (c)

Fig. 4. The resulting signal segments after the automatic segmentation into single words.

behaviour of the tutor [14]. Looming behaviour is rewarded by pointing at the loomed object, thus trying to heighten the joint attention.

The participants had to partake in 3 sessions which had at least one day in between them. In the second and third session the robot spoke back to the participant [15], but this is only mentioned for the sake of completeness and will not be important for our analysis since we will only regard the first session.

2) Task: The task for the participants was to teach the robot about different shapes, sizes and colours. As objects they were given 3 different sized cubes (small, medium and large) with different shapes (sun, heart, cross, circle, arrow and crescent moon) in different colours (red, green and blue) on them.[4] The participants were then advised to explain these 3 different characteristics (sizes, colours, shapes) to the robot in any way they like for about two minutes in each session. In Figure 1 we can see one of the participants explaining the medium sized cube with the blue sun shape facing towards the iCub [9] and Figure 2 shows the general setting during the experiment.

B. Looming Detection

In Section II-B we defined what part of the interaction is meant to help us distinguishing important from less important words. To utilise this we have to detect the looming behaviour of the tutor. We will not talk about the object detection since the object tracker always has to fit to the specific problem[5], but will just focus on the hands to not go beyond the scope of this paper.

We used a Microsoft Kinect camera (as seen in Figure 2) to get a 3D image of the scenery and used the ability to track the position of the hands in 3D space provided by the OpenNI [16] framework. The only value we will observe is the distance (z-coordinate) of the hands to the Kinect camera which is placed behind the robot. These distances are the only important informations for our looming detection since every movement of bringing the object into the line of sight of the robot includes a forward movement [6]. Our approach in the mentioned study (see III) was to use just a fixed distance δ which had to be undershot to trigger the looming detector:

$$L = \begin{cases} \text{true} & \text{if } d < \delta \\ \text{false} & \text{else} \end{cases} \tag{1}$$

[4]On each side of the cubes was only one shape in one colour to make it unambiguous which object is explained.

[5]We used a standard ARToolKit [10] marker tracker.

Where d is the current distance of the hand to the Kinect and δ is the threshold for the looming detection which has to be obtained by manual calibration and testing.

Now that we can detect looming behaviour we have to record the voice during these parts of the experiment and segment it into words [7].

IV. Segmentation into Words

When recording sound we will always get some kind of noise that is distorting the signal we want to process. We can of course try to minimize that by using a headset or unidirectional microphone arrays (in our case we used a headset to enhance the audio track of one of the cameras) but, nevertheless, we will always get some kind of background noise or distortion. To get rid of almost all of the unwanted information in the signal we used two algorithms of noise reduction as suggested in the paper by Waheed et al. [7].

1) Pre-emphasis Filter: At first the incoming speech signal is preprocessed using a so called pre-emphasis filter: $y(n) = x(n) - \alpha \cdot x(n-1)$ where n is a discrete time step and $x(n)$ is the corresponding value. The α represents the pre-emphasis factor which usually is 0.95. The pre-emphasis filter in general is used to reduce differences in power of different components of the signal. In speech recognition the pre-emphasis filter is used to *"[...]reduce the effects of the glottal pulses and radiation impedance."* [7] and *"It takes the focus to the spectral properties of the vocal tract."* [7].

2) Band-Pass Filter: The second algorithm which is designed to reduce low frequency background noise and remove high frequency noise spikes [7] is a band-pass filter. This filter basically consists of a high-pass filter and a low-pass filter. So the band-pass filter passes through frequencies in between an upper and a lower border.

A. Segmentation

To segment the signal we used an algorithm based on an entropic contrast suggested by Waheed et al. [7]. After the signal has been filtered (see IV-1, IV-2) it is divided into windows of 1024 frames (at a signal frequency of $44100Hz$) with a 25% overlap which is then passed into a histogram with 100 bins to determine the probability distribution for that individual frame. The entropy of each of these individual windows is then computed by the standard entropy formula [7]: $H = -\sum_{k=1}^{N} p_k \log_2 p_k$. This gives us a list of entropies which are used to construct the entropy profile $\xi = [H_1\, H_2 \cdots H_m]$ with m total windows of 1024 frames in the signal. From this entropy profile we can now choose a

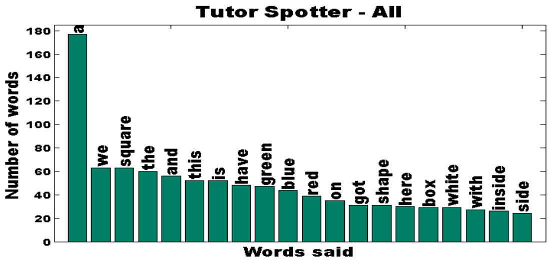

(a) Words said during whole session.

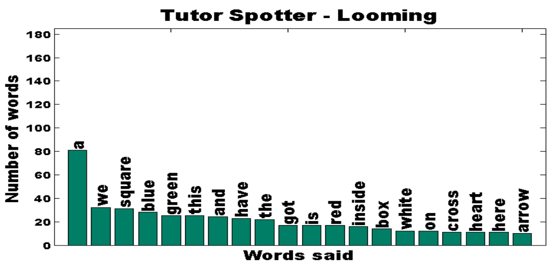

(b) Words said during looming phases.

Fig. 5. Two histograms of the 20 most said words during the first session of the study. Both histograms are taken from the Tutor Spotter [14] condition.

biased threshold to *"[...]minimize excessive influence of the background noise."* [7]:

$$\gamma = \frac{\max(\xi) - \min(\xi)}{2} + \mu \cdot \min(\xi) \qquad (2)$$

The bias is defined by $\mu \cdot \min(\xi)$ where $\mu > 0$ and $\min(\xi)$ represents the residual noise floor. After defining the threshold we can consider every window with an entropy above the threshold as speech and every window with an entropy below the threshold as noise [7]. The problem with that assumption is that in many cases non-speech data can be reported as speech data due to artefacts. Also some valid speech data may be ignored because of its physio-vocal characteristics. So Waheed et al. [7] suggest two further criterions in addition to the threshold to determine whether a segment contains speech or not.

The first criterion is the size of the found speech segment $\lambda_i > \kappa$ where κ symbolizes the duration of the shortest phoneme in the target language. Because, *"Humans generally do not produce very short duration sounds."* [7]. The second criterion is the inter-segment distance d_{ij} between the segments i and j. This criterion is required because there can be parts of speech that have been separated into two segments due to its pronunciation [7]. So the criterion is $d_{ij} < \delta$ where δ is the maximum inter-segment distance.

As our final distinguishing criterions to determine speech segments we now have our threshold and if λ_i or $\lambda_j > \kappa$ and $d_{ij} < \delta$ the two segments i and j are merged and the space in between will be considered part of the speech, too. On the other hand if $\lambda_i < \kappa$ and $d_{ij} > \delta$, then the segment i will be discarded and thereby considered noise.

The problem with this automatic segmentation algorithm is that the algorithm will just find sentence boundaries or the shortest utterances if the speech is very continuous [7]. Since we are expecting one or two word sentences during the looming phases we hope to achieve good results, nonetheless.

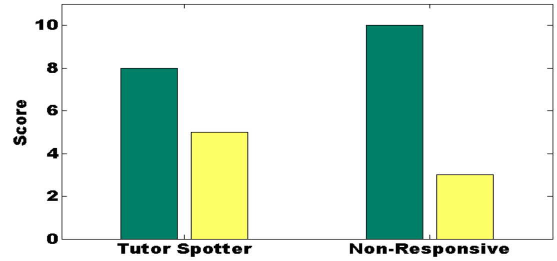

Fig. 6. The scores for the histograms of the two different conditions for the first session of the experiment. The histograms gained a point for every keyword (14 keywords in total) that was listed first. So e.g. if *blue* turned up first in the looming phase histogram then looming gained a point and vice versa. Green: looming, yellow: whole session.

V. Identifying Similar Words

To determine which words are most important for the object description, and by that which words we have to learn, we will need a way to make an assertion about which words are similar to previously heard words. By that we can construct a histogram of words and hope that the most used words are the most descriptive ones. As one possibility to do so we suggest an approach that is similar to the audio fingerprinting algorithm introduced by Yan Ke et al. [8] which is used by the music industry.

This approach was chosen because of its high reliability, the insensibility to noise and the possibility to find single words in longer utterances which compensates for the segmentation where not all of the words can be segmented due to continuous speech. This algorithm Fourier transforms the sound signal and treats it as a 2D image. By that they try to *"[...]employ geometric verification in conjunction with an EM-based occlusion model to identify the song that is most consistent with the observed signal."* [8]. These 2D images represent spectrograms of the given signal and could be compared directly by using correlation. This however would be too slow and inaccurate so Yan Ke et al. [8] suggest to use a small set of filters that are robust to small distortions and still give us enough information to distinguish between two different signals. After viewing the spectrogram images Yan Ke et al. [8] suggested that the filters introduced by Viola and Jones [17] are most suitable for their needs. To select a descriptive subset of these filters Yan Ke et al. [8] use a pairwise boosting algorithm that differs from the standard Adaboost [18], [19] in the fact that they only re-weight pairs of filters instead of single filters since their suggested weak classifiers cannot do better than chance on their own. After creating this subset of filters they are used to create a set of descriptors for overlapping windows of the signal.

These descriptors are written to a file and stored in a

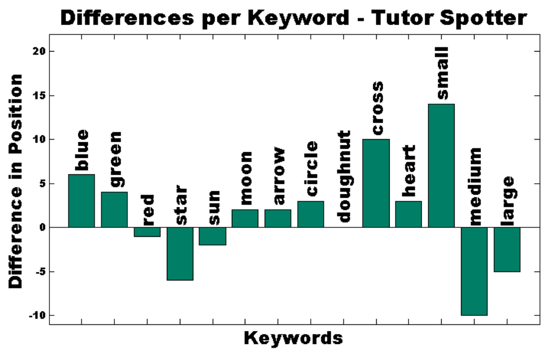

Fig. 7. Tutor Spotter [14] condition: The difference in position for each single keyword. Positive means the keyword moved up this number of ranks in the histogram when comparing looming with the whole session. Negative means the keyword moved down this number of ranks.

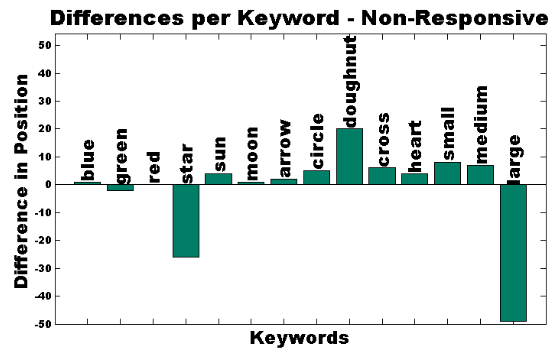

Fig. 8. Non-Responsive condition: The difference in position for each single keyword. Positive and negative values are defined as described in Figure 7.

database and this database can then be queried using other descriptor files. So we are now able to construct a histogram of words (by running cross references) without knowing the actual word itself but just knowing its number of occurrences in the recorded signals.

VI. Results

After we have seen how to construct a system that detects the looming behaviour and builds histograms of the said words we will go back and have a look at our study again (see III).

To show that it will be more likely to learn an important keyword[6] during a looming phase than during the whole session we will employ a histogram based analysis as suggested in Section V. For our analysis we just considered the first session to prevent any kind of learning effect from falsifying our results. In Figure 5 we can see two histograms which resulted from the first session of all participants facing the Tutor Spotter [14] condition. Figure 5b still shows that even during the looming phases the most said word is *a* but the amount is considerably smaller than during the whole session as we can see in Figure 5a. We can also see from Figure 5 that the first 3 words are identical in regard to their position in both histograms. The first real change is the 4th word which is *the* for the whole session but *blue* for the looming phases. To highlight this effect of keywords moving up the ranks in the histogram we can have a look at Figure 6 which shows scores depending on the position of the keywords in the histogram. The higher the score the more keywords are mentioned first in the related histogram. So the looming phases generate histograms that contain more keywords on higher ranks than the histogram over the whole session. The difference in ranks which the keywords moved up or down to is pictured in Figure 7 and 8. In Figure 7 the rank difference is 20 which means that we have a gain of 1.43 ranks per keywords on average. In Figure 8 we have a rank gain of -19 which is dominated by one outlier. If we disregard the outlier we achieve a rank gain of 30 which means an average rank gain of 2.31 per keyword.

[6]Keywords used for this analysis are: red, green, blue, star, sun, moon, arrow, circle, doughnut, cross, heart, small, medium and large.

VII. Conclusion

We have seen in the Results Section VI that we have achieved to move the keywords up in the histogram if only regarding the looming phases. This is the essential result if we want to rely on these histograms to create a database of words associated with one particular object. On the other hand we have also seen that we still get a lot of filler words that are meaningless for the object description. These filler words will still have to be filtered out by running a cross reference between the different objects like the Inverse Document Frequency [20] algorithm. The presence of the unwanted information may be explained by the experiment design which was not built to induce looming behaviour in particular but to evaluate the Tutor Spotter [14] (see VIII-2). However, it still narrows down the search space, nonetheless. Figure 6 shows that the Tutor Spotter [14] condition yielded less keywords that moved up in rank than the Non-Responsive condition. This could be due to the reason that the Tutor Spotter [14] itself already creates a state of joint attention triggered by the gazing behaviour of the tutor. Thus, it implies higher cognitive function which results in less use of the synchronous naming behaviour as implied by [6]. But there still is a gain in ranks which we can see in the Results Section (VI). The Non-Responsive condition tends to yield better results because of the general inattention which induces attention getting behaviour (see Figure 6 and 8).

This leads us to the conclusion that considering the looming phases as a clue for meaningful keywords will narrow down the search space and improve the possibility of finding keywords at the top of the histogram if we use an experimental set-up which is either inattentive and/or rewards looming.

In addition to the narrowed search space we gain an unambiguous clue which object these keywords are describing because, as Matatyaho and Gogate [6] state, the attention getting gestures also help to highlight the object-word relation by highlighting the object through movement. So we found a way of combining meaningful words with objects without knowing anything about the object or the word.

VIII. Future Work

1) Looming Detection: The looming detection method we used in the study (see III-B) is very robust concerning the actual detection of looming but also very fragile in regard to calibration and changing environments, e.g. a different position of the tutor that brings him closer to the Kinect could trigger the looming more easily. Due to this problem of exact calibration and adaptability we will follow a new approach of looming detection for future studies which relies on the mean distance of the hands and the variance of that distance. The idea behind this is that the hands (or at least one of the hands) of the tutor will be in front of him while explaining the object of interest. During the explanation phase the tutor will create his own *explanation space* where he moves the object about freely in his normal way of describing it. This so called *explanation space* will be defined by the mean distance over the time and a certain variance to compensate for normal purposeless movement.

Looming will now be triggered if the tutor moves the object in his hand out of the *explanation space* towards the robot (Kinect). To create a more robust detection the hand has to move towards the robot for a minimal distance of twice the radius of the variance sphere. So we end up with the following conditions:

$$L = \begin{cases} \text{true} & \text{if } d < \mu \wedge |d - \mu| > \sigma \cdot 2 \\ \text{false} & \text{else} \end{cases} \quad (3)$$

Where d is the current distance of the hand to the Kinect, μ is the mean distance over time and σ is the standard deviation which equals the square root of the variance $\sqrt{\sigma^2}$.

With this method of detection we hope to achieve a more natural looming detection and an easier experimental set-up and are hopefully able to construct a system that induces looming in a more robust fashion.

2) Future Studies: We believe that the presented study (III) was not optimal to test the real abilities of our system since it was designed to show the benefits of the Tutor Spotter [14], and therefore hope to conduct a new study where we can test an experimental set-up that is tailored to our needs with a system that obeys the rules of creating joint attention when looming is detected like stated in [6] (see Section II-C). We hope to achieve better results and show a more significant difference in rank gain for the desired keywords by doing so.

Also, our system was not implemented and running at the actual study. So we hope to show that with a running system during an experiment we can actually learn at least some of these found keywords and associate them with the given object.

Acknowledgement

We would like to thank for the support of Katharina Rohlfing. This research has been supported by the EU project RobotDoC (235065) from the FP7 Marie Curie Actions ITN and by the EU FP7 Project ITALK (ICT-214668) within the Cognitive Systems and Robotics unit, as well as from the Citec and the IIT.

References

[1] A. Cangelosi, G. Metta, G. Sagerer, S. Nolfi, C. Nehaniv, K. Fischer, J. Tani, T. Belpaeme, G. Sandini, F. Nori *et al.*, "Integration of action and language knowledge: A roadmap for developmental robotics," *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 3, pp. 167–195, 2010.

[2] K. Rohlfing, J. Fritsch, B. Wrede, and T. Jungmann, "How can multi-modal cues from child-directed interaction reduce learning complexity in robots?" *Advanced Robotics*, vol. 20, no. 10, pp. 1183–1199, 2006.

[3] K. Lohan, A. Vollmer, J. Fritsch, K. Rohlfing, and B. Wrede, "Which ostensive stimuli can be used for a robot to detect and maintain tutoring situations?" in *Affective Computing and Intelligent Interaction and Workshops, 2009. ACII 2009. 3rd International Conference on*. IEEE, 2009, pp. 1–6.

[4] T. Plötz and G. Fink, "Robust time-synchronous environmental adaptation for continuous speech recognition systems," in *Proc. ICSLP*, vol. 2. Citeseer, 2002, pp. 1409–1412.

[5] G. Fink, "Developing hmm-based recognizers with esmeralda," *Lecture notes in computer science*, pp. 229–234, 1999.

[6] D. J. Matatyaho and L. J. Gogate, "Type of maternal object motion during synchronous naming predicts psreverbal infants' learning of word-object rselations," *Infancy*, no. 13:2, pp. 172–184, 2008.

[7] K. Waheed, K. Weaver, and F. Salam, "A robust algorithm for detecting speech segments using an entropic contrast," in *Circuits and Systems, 2002. MWSCAS-2002. The 2002 45th Midwest Symposium on*, vol. 3. IEEE, 2002, pp. III–328.

[8] Y. Ke, D. Hoiem, and R. Sukthankar, "Computer vision for music identification," in *CVPR (1)*, 2005, pp. 597–604.

[9] G. Metta, G. Sandini, D. Vernon, L. Natale, and F. Nori, "The icub humanoid robot: an open platform for research in embodied cognition," in *Proceedings of the 8th workshop on performance metrics for intelligent systems*. ACM, 2008, pp. 50–56.

[10] H. Kato and M. Billinghurst, "Marker tracking and hmd calibration for a video-based augmented reality conferencing system," in *Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*, San Francisco, USA, Oct. 1999.

[11] P. Zukow-Goldring, "A social ecological realist approach to the emergence of the lexicon: Educating attention to the amodal invariants in gesture and speech." *Evolving explanations of development: Ecological approaches to organism-environment systems*, pp. 199–252, 1997.

[12] P. Zukow-Goldring and K. R. Ferko, "An ecological approach to the emergence of the lexicon: Socializing attention," *Sociocultural approaches to language and literacy: An interactionist perspective*, pp. 170–190, 1994.

[13] C. Rankin, T. Abrams, R. Barry, S. Bhatnagar, D. Clayton, J. Colombo, G. Coppola, M. Geyer, D. Glanzman, S. Marsland *et al.*, "Habituation revisited: an updated and revised description of the behavioral characteristics of habituation," *Neurobiology of learning and memory*, vol. 92, no. 2, p. 135, 2009.

[14] K. Lohan, K. Rohlfing, K. Pitsch, J. Saunders, H. Lehmann, C. Nehaniv, K. Fischer, and B. Wrede, "Tutor spotter: Proposing a feature set and evaluating system," *International Journal of Social Robotics*, 2012.

[15] J. Saunders, H. Lehmann, F. Foerster, and C. Nehaniv, "Robot acquisition of lexical meaning - moving towards the two-word stage," 2012.

[16] "Openni - the openni organization is an industry-led, not-for-profit organization formed to certify and promote the compatibility and interoperability of natural interaction (ni) devices, applications and middleware." [Online]. Available: http://openni.org

[17] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," in *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, vol. 1. IEEE, 2001, pp. I–511.

[18] Y. Freund and R. Schapire, "Experiments with a new boosting algorithm," in *Machine Learning-International Workshop then Conference*. Morgan Kaufmann Publishers, Inc., 1996, pp. 148–156.

[19] R. Schapire and Y. Singer, "Improved boosting algorithms using confidence-rated predictions," *Machine learning*, vol. 37, no. 3, pp. 297–336, 1999.

[20] K. Church and W. Gale, "Inverse document frequency (idf): A measure of deviations from poisson," in *Proceedings of the third workshop on very large corpora*, 1995, pp. 121–130.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Towards robots with teleological action and language understanding

Britta Wrede and Katharina Rohlfing
Citec
Bielefeld University
bwrede@techfak.uni-bielefeld.de
kjr@uni-bielefeld.de

Jochen Steil and Sebastian Wrede
CoR-Lab
Bielefeld University
{swrede,jsteil@cor-lab.uni-bielefeld.de}

Pierre-Yves Oudeyer
Flowers Team
INRIA, Bordeaux
pierre-yves.oudeyer@inria.fr

Jun Tani
Cognitive Neuro-Robotics Lab
KAIST
tani1216jp@gmail.com

***Abstract*—It is generally agreed upon that in order to achieve generalizable learning capabilities of robots they need to be able to acquire compositional structures - whether in language or in action. However, in human development the capability to perceive compositional structure only evolves at a later stage. Before the capability to understand action and language in a structured, compositional way arises, infants learn in a holistic way which enables them to interact in a socially adequate way with their social and physical environment even with very limited understanding of the world, e.g. trying to take part in games without knowing the exact rules. This capability endows them with an action production advantage which elicits corrective feedback from a tutor, thus reducing the search space of possible action interpretations tremendously. In accordance with findings from developmental psychology we argue that this holistic way is in fact a teleological representation encoding a goal-directed perception of actions facilitated through communicational frames. This observation leads to a range of consequences which need to be verfied and analysed in further research. Here, we discuss two hypotheses how this can be made accessible for action learning in robots: (1) We explore the idea that the teleological approach allows some kind of highly reduced one shot learning enabling the learner to perform a meaningful, although only partially "correct" action which can then be further refined through compositional approaches. (2) We discuss the possibility to transfer the concept of "conversational frames" as recurring interaction patterns to the action domain, thus facilitating to understand the meaning of a new action. We conclude that these capabilities need to be combined with more analytical compositional learning methods in order to achieve human-like learning performance.**

I. Introduction

One prominent goal in robotics is to endow robots with the capability of making sense of a situation. To achieve this, a popular approach is to build systems that perceive pieces of situation (objects, persons, their movements) first and then to combine them for a bigger picture on this situation. There is, however, increasing evidence in the developmental literature that children have the bigger picture first before they actually understand the individual pieces: Imagine a toddler who is eager to participate in a game that the older siblings are playing: S/he will not be able to understand every move in this game, let alone the proper rules but (a) s/he is definitely motivated to participate i.e. to act as the others do and (b) s/he knows that this participation is about holding some cards, having a turn with dicing and the ultimate goal of receiving other cards. This toddler seems to have amazing skills according to which s/he analyzes the situation instantly and this enables her or him to produce a behavioral turn which will provide the child with data and solicit positive or negative input from the physical and/or social environment. This way, learning within an interaction can be stimulated [1]. We know from our previous research that disabled persons also draw from similar capabilities to engage in an interaction even if not every component is understandable [2].

Todays artificial cognitive systems lack the capabilities to actively participate in an interaction on a very low level of understanding having a very rough and preliminary idea of the goal. But we know that without such capabilities an interaction with a user cannot start or will break down soon when facing unforeseen situations.

By teleological understanding we mean that a system can apply pragmatic or contextual knowledge about the task or goals. In contrast, by compositional understanding we mean an approach that is already able to parse an action into its components thus allowing for generalisation to new actions and situations. In the action modality, Gergely and Csibra (2003) [3] suggested that young children develop a teleological understanding of actions that is based on a situational analysis. Seeing a ball moving from one side to the other, one-year-olds are sensitive to whether this ball is moving in a rational way or not. A rational way is achieved when (1) the moving of the ball brings about a future goal state, and (2) the goal state is realized by the most rational action available to the ball within the constraints of the situation ([3], p. 289). Thus, a straight path that the ball takes from one side to a particular place is better acceptable than a curved path. Such a rationality assumption (ibid: 290) refers to the infants ability to understand actions without attributing intentional mental states to others. In the language modality, some scholars (e.g. [4], [5], [6], [7]) acknowledged that children need to establish some interaction protocols that provide a general pragmatic frame for them to act appropriately without knowing every linguistic detail. Social signals such as contingency and ostension ([8], [3], [7], [9]) guide children towards understanding of such frames and result in childrens ability to participate in an interaction, even though they might not understand every detail of it. The term “frame”, thus, refers to a multi-modal interaction structure that is already established. Fogel et al. ([5] p. 3) characterize frames as “regularly recurring patterns of communication”. For example a “labeling-frame” consists of looking at and pointing to an object and labeling it possibly with an excited facial expression and an according intonation. Thus, once the structure is established in recurrent interactions, the only new element within it is the new label for a particular objects - simplifying the understanding process tremendously.

II. Teleological action

We understand teleology in the sense that a child (or a robot) will pick up an element of the action (such as manner or goal) to be important in his/her observation which will then be applied for the action production. Although in the following we make no explicit difference between action understanding and production we are aware that a child (or robot) may be able to understand the ”goal” of an observed action (either a ”manner” goal or a ”target position” goal), but may not be able to reproduce this goal at this point. Also, the concrete mechanisms underlying perception and production may be different but at a higher level we do not exclude that a common representation exists.

Developmental psychology provides evidence that young infants are able to understand actions without decomposing them. Rather, they seem to apply some sort of goal-oriented interpretation enabling them to carry out an action in order to achieve that goal - even if in the first attempts they do not achieve the goal.

In our attempt to characterize a teleological action, we think that two questions need to be separated: The first question concerns the development of teleological reasoning, i.e. what kind of representation of the goal is formed when experiencing an action in a situation. Another question concerns how the experience of several actions can contribute top-down to an analysis of a new situation. We would like to exemplify these questions on recent discussion about findings from developmental studies. Gergely and colleagues [10] compared childrens imitation abilities in two different conditions. More specifically, children saw the experimenter switching on light by pressing a button with his forehead. In one condition, the experimenter had his hands put visibly free on a table. In another condition, the experimenter seemed to be cold and wrapped himself and his hands in a blanket. Children imitated the action of switching the light on with the forehead in the hands-free condition predominantly. The authors suggested that it is due to a situational analysis that the children conduct: In the hands-occupied condition, they might conclude that while the experimenter chooses the most rational action of pressing the button with his forehand (because his hands are occupied), for the children, the most rational action would be to press the button with their hands. Only in the condition, in which the action of pressing the button with a forehand seems to be a rational one, seem the children to imitate it. Recently, however, this interpretation has been challenged by a more low-level explanation. Conducting some controlling conditions, in the study by Beisert et al. [11], the authors found that children in the hands-occupied condition might have been distracted by the unfamiliarity of the experimenters appearance. Thus, it is not due to the fact that they judge an action as rational or not. Instead, children seem to imitate in cases, in which they have the cognitive resources to do so. The unfamiliarity of the experimenter distracted the childrens attention. Accordingly, only the goal (switching the light on) was emulated, regardless of the means (the forehead). We think that Beiserts findings actually perfectly fit with the idea of a teleological action: Children perceive an action as goal-directed, and depending on their cognitive resources they can perceive (and/or produce) further details of how to achieve the goal. If their resources are limited, they will choose to reach the goal with an action they have in their repertoire. [1] This repertoire is built upon the childs experience (the second question). Thus, teleological action understanding is about formulating a quick behavioral plan and to further learn from an interaction. However, we still know little about what action experience leads to this repertoire (the first question).

A. Representation of Goals

How are goals represented and how are concepts of goals formed? In accordance with Mandler [12] we believe that the “goal” of an action is a concept inflicted upon the action by the adult teacher who has already learned that it is helpful to understand the world with respect to “goals”. Consider for example a person moving an object from point A to point B. What is the “goal” of this action? The goal could be for the object to be in the final position B. However, the goal could also be to simply remove the object from its current position A. Yet another goal could be that the object needs to be moved in order to change its state e.g. shaking to make snowflakes move in the water as in a snowdome. There is an infinite number of potential goals which an action may have. It is therefore almost impossible to derive its goal (or meaning) from one single observation. So how do infants learn the meaning or goals of actions?

Mandler [12] argues that although young infants are already able to produce predictions about ongoing actions and to behave in a goal-oriented manner, the abstract concept of a goal only forms later. Instead, she proposes basic conceptual elements with which children might perceive a motion as meaningful. More specifically, in the example outlined above, infants will most likely perceive a START and an END of this motion as meaningful units. It is still of question whether and

[1] However, note that the results of this experiment can not explain if the infants did indeed not perceive the manner or if they were simply not able to incorprate it in the computation of their behavior generation. One could also argue that infants perform actions according to an optimality criterion, i.e. to produce the action with the least effort possible which depends on the manner. However, if the cognitive load is too high the infant may not be able to take a different manner (i.e. using the head) into account but rather chooses the manner that is generally accompanied with the least effort (i.e. using the hands).

how these units can be learned or are innate concepts that the children are equipped with.

Following this line of argumentation, we believe that in a very early representation of action the goal will not be explicitly represented. Rather, physical effects that occur in the scene during the action can be noticed and represented. Thus, in order to derive the invariant ”goal” of an action the learner needs to look out for effects that always occur with this action. We elaborate later on how such potentially goal-bearing, meaningful effects may be perceived. [2]

B. Teleological processing

We assume that infants separate between goal and manner and - if under cognitive load - are only capable of reproducing one aspect (e.g. only the goal, not the manner). While at first sight this may appear as a disadvantage or compensation strategy it might well be that this resource-boundedness actually helps the infant to learn more efficiently, as it produces an action that is partially “correct” and also partially “wrong”. This will provoke the tutor (or the physical environment) to give the learner a corrective feedback e.g. by pointing out the missing parts (or by reinforcing the achieved part). This is what we understand as the “Action production advantage”.

III. Action Frames

When compared to similar findings in language, for action learning, the scholars have just begun to acknowledge the value of recurrent structures and how their recognition might facilitate action understanding. While in language learning, “frames” refer to an established interaction structure that is based on regularly recurring patterns of communication, in action learning the value of recognizing the recurrent structure is implicit to the assumption of a rational action. In other words, one can assume that infants will consider a movement as rational if they recognize it as a familiar movement within a particular situational configuration. We think that this aspect needs to be more recognized: The memory of events is the basis against which we perceive and predict actions. However, to date, we know little about how young infants remember structural properties of the motion stream [13] - yet this capability of making sense of actions seems to be basic [12]. Loucks et al. [13] (p. 233) suggest that “as people take action to carry out their object-directed intentions, certain physical and temporal properties predictably coalesce within the motion stream [] The time course of this sequences gives the concatenation of elements a ballistic quality, with characteristics acceleration and deceleration parameters”. However, we need more evidence in developmental studies of how children’s memory allow to form such elements.

[2]Another issue related to this is the question how the learner knows that this is the same action as before, especially if some aspects have changed. We argue that it is the communicative frame - be it verbally by specifying a label e.g. “Look, I *shake* it.” or be it non-verbally by providing a repetition frame e.g. through prosody - that indicates whether the new action is the same or a different one. In the latter case it is likely that we find some kind of contrastive marker, indicating what it is that is different from the previous action (e.g. “Look, now I *don’t shake* it - see, there are no snowflakes if you move it carefully”.

These insights strongly suggest that we need to endow artificial systems with the capability to memorise and recognise recurrent structures. How the recognition of recurrent structures is integrated with the learning system we do not know. But we speculate that there is an intimate relationship especially in the (frequent) cases where a structure is only partially known and the unknown part needs to be learned and associated with meaning.

How can actions or the situation provide a recurrent structure that help the learner to better understand a new action? This may have different dimensions. One recurring cue may be the structure of a situation: consider for example a situation where only one agent is present (be it in the form of an agent toy such as a teddy bear or in the form of a real agent e.g. the tutor) but no manipulatable object in his/her reach. This is a frame indicating that an intransitive action has to happen, where the movement of the agent him/herself is of relevance. However, if there is an object present, a transitive action can be expected, where the important thing will happen to the object. A different dimension of action frames might be related to the structure of actions itself: consider for example the process of manipulating an object. This will always involve a sequence of reaching, grasping, manipulating. From this one might condense the restriction, imposed by physics, that in order to manipulate an object it needs first to be grasped. A more sophisticated notion of such an action structure has been explored and formalised by Pastra & Aloimonos [14]. Although we can formulate hypotheses on how actions and situations may frame actions and thus render them more understandable, these hypotheses need to be carefully tested in experiments with infants.

IV. Fast learning

How can we achieve the capability in an artificial system to learn some meaningful part of an action based on only one first observation? We argue that it needs to be capable of perceiving some sort of physically observable effect of the action which it will try to imitate (or emulate) based on its existing experience.

A. Goal identification

One bottom-up way for the learner to infer or detect the (potential) goal of an action is to identify recurring features in repeated action demonstrations. From a modeling perspective it is reasonable to assume that bottom-up biases exist that help the learner to identify goal-relevant features such as:

- **Goal-**position of manipulated object (the standard “goal-directed” action where the spatial goal matters e.g. when putting a cup on a saucer, or opening of a bottle where the goal is that the cap is away from the bottle)
- **Source-**position of manipulated object (e.g. move the object away)
- **Orientation** of manipulated object (e.g. turning an object upside-down)
- **Movement** itself (e.g. shaking leading to movement of snowflakes)

- **Visibility** of manipulated object (e.g. a cup that "vanishes" when nested in another cup)
- **Form** of manipulated object (e.g. open vs closed; scrunched; folded vs unfolded;)
- **Sound** (on vs off; intensity; frequency)
- **Social reaction** of a person (e.g. smiling, waving, nodding, speaking,)

A further consideration then concerns the process how such features may be observed. Again, we propose very basic strategies, such as:

- Detection of physical **effects** (for goal identification)
- Detection of **deviations** from "normal", that is highly frequent, patterns (e.g. a normal trajectory might be a straight one between two points in space, if the trajectory is deviating from this by a cyclic shaking pattern this will be noticed and interpreted as important) (for manner identification)

In order to produce the identified goal (which might be wrong) the system will choose the most frequent means it knows in order to achieve this goal (which might be wrong as well) (note that in contrast, Csibra & Gergeley [3] argue that infants follow a rationality principle: it is most rational to move from A to B in a straight line instead of taking a deviation; it is most rational to push a button with the hands instead of the head). Instead, here we argue for a frequency effect. For example:

Goal	Means
Position of object	Movement with hand in a direct line to B
Orientation of object	Grasp object with one hand and turn it with the other hand
Visability of object	Movement with hand towards larger object in environment

B. Carry out action

In order to achieve the action production advantage we need a system that is capable of carrying out an action regardless of its current learning state. This invites to concentrate on movements and actions that are relevant to engage in interaction and facilitate communication with a tutor on the one hand and which can be applied to a wide range of situations on the other hand. This early repertoire of action shall rely on minimal requirements for internal modeling and rather have the form of direct sensory-motor patterns that are easy to learn and to explore. An example are pointing gestures with the arm, which reference directions and can initiate and maintain much interaction, but need not be sophisticated in exactly referencing a particular object e.g. with an extended finger. Such elementary actions can easily be learned without representing exact positions in space and without an explicit kinematic model of the body in a holistic fashion [15]. This means also that there is a mechanism allowing the system to take part in an interaction e.g. by taking into account minimal turn-taking rules. The challenge is to devise respective simple but behavioral meaningful and flexible sensori-motor represenions of such elementary actions that can further be specified and differentiated as development progresses.

C. Develop compositional structure

A further desideratum accrueing from these considerations is the capability of the learning architecture to switch between teleological and compositional learning. While the system needs to be able to build an initial structure of an action it then needs to apply the emerged structure towards new incoming action demonstrations and parse them accordingly. This does not necessarily need to happen in an explicit way. It is possible that depending on the current representation different processing takes place implicitly.

V. Comparison of compositional and teleological learning

Why is teleological learning not only beneficial but even necessary? Why shouldn't a learning system start with compositional learning? By teleological learning we mean a process that exploits the production advante through continuous engagement in interaction frames and which progresses from teleological behavior to more differentiated (compositional) perception, situated action and situation understanding. We argue that it is the use of pragmatic and action frames that help the learner to bootstrap an initial action representation and to tie observed new cues to a meaning-frame, i.e. some kind of category that helps the learner to better understand. Such categories may be syntactic categories such as objects, verbs or adjectives or they may be semantic categories such as agent or patient.

We further argue that the teleological approach may achieve one-shot learning resulting in a partially correct action representation by focussing its resources on one single aspect of the action. This way the complexity of decomposing an observed action in all its parts and potential meanings is highly reduced.

The predominant approach in movement learning towards motor skills and complex actions currently is compositional. It considers movement primitives in low dimensional task spaces, i.e. elementary movement patterns that can be either goal directed [16], [17], [18] or oscillatory [19]. The respective learning methods have developed sophisticated schemes to represent and generalize the manner by encoding the movement in dynamical systems that follow learned transients towards goals and can even model the relevance of the manner according to demonstrations through probability based methods [20], [18]. These primitives are assumed to form the basic building blocks of a later to be composed complex action, however, this very process of composition is hardly considered and taken for granted. In practice the composition is far from easy to achieve [21] and typically employs very heuristic schemes [22] or needs in itself sophisticated modeling approaches [23]. The transformation from task spaces into the motor domain is solved model-based through standard inverse kinematics. Movement primitive based methods need to be complemented by approaches to segment complex actions into basic primitives, which is also far from easy e.g. using NMF [24] or sequential primitives ([25], or taking into account the already learned primitives [26]. On the other hand, it has been shown that neural networks (e.g. MTRNNs [27]) are capable of evolving a structure that allows to learn action primitives and their sequencing in the same neural structure.

This compositional approach can neither explain how in the course of development the coordination of motor function comes about, which needs to orchestrate a highly complex body in extremly high dimensional motor spaces. Nor does is offer any route to interaction and communication and thereby mostly ignores the behaviroal function of the action, as was recently observed also in [28]. More recently, first learning approaches to directly tackle the coordination problem in high dimensions towards complex sensory-motor behavior have been very successful [29], [30], [31], [32] motivated by the observation that infants also perform reaching movements from the very beginning [33]. While impressive, these approaches share the feature that they need a set of "goals" to drive learning, i.e. in motor coordination tasks a number of spatial positions to reach. These appraoches provide some means to follow a more behavior and goal oriented, i.e. a more teleological approach to action learning.

Fig. 1 illustrates this restriction: in the example the target position is specified ('B') whereas the other potential goals (source, manner, form) of the action are not specified in the sense that they can take any value ('*'). Although this is generally not represented in an explicit way, it is often encoded implicitly in the learning algorithm which, for example, does not allow to switch between different types of goals.

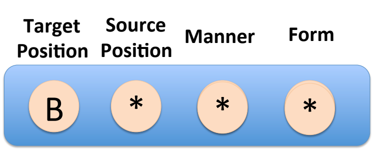

Fig. 1. Implicit structure underlying compositional learning approaches: each entry needs to be specified, generally only one entry is specified explicitly while the others are specified implicitly through the learning algorithm.

In contrast, in the teleological approach the learner would simply select (based on which strategy ever) one "goal" and try to carry it out with the easiest or most frequent means it has at its disposal (cf. Fig. 2. If this turns out to be unsuccessful – which can be tested through interaction with the tutor – the hypothesis can be refined or corrected according to the feedback. Also, new constraints can be added if necessary.

In short, what is needed is an approach that facilitates the acquisition of structure, i.e. the discovery that in an action not only the end position matters but also the way in which it is demonstrated or or executed and that constraints may exist that modulate the action such as obstacles.

Based on the structure that teleological learning has thus achieved, compositional learning strategies can be applied.

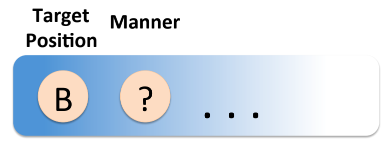

Fig. 2. Structure underlying teleological learning approaches: only one entry needs to be specified in order to enable the system to carry out the action. This yields in a corrective feedback of the tutor which can immediately be added to the teleological representation.

Note, that this is a circular relationship: at every point in time it must be possible to extend the "filled" compositional structure by a teleological approach.

VI. Conclusion

We have argued that infants follow a teleological action learning approach which is characterised by the ability to assign some sort of meaning to an observed action (which may be wrong) facilitated through conversational and action frames and the ability to imitate part of the action (possibly leading to the non-achievement of the goal).

This perspective on action and language learning has a range of consequences when considered as a basis for robotic learning approaches which need to be explored in future in order to better understand their functions. In this paper we have explored two potential consequences: (1) Communicational and action frames allow the learner to make use of recurrent structures, be they syntactic, semantic or pragmatic, in order to assign meaning to a newly encountered entity. Especially the concept of action frames is largely unexplored, although first attempts exist [14]. Further research needs to tackle the question how such action frames look like in more detail and how they can be exploited for learning. (2) Through the capability to extract only one part (e.g. only manner or only goal) from a single action demonstrationg and to reproduce it by means of extrapolation from previous experience the learner builds an action concept in an incremental way, by eliciting feedback of the tutor and thus gaining specifically tailored information to increase his/her action concept.

Both capabilities have not yet been shown by existing compositional approaches but might be achieved through different applications. We argue that in order to achieve a system that is able to learn new actions from the very first steps towards a level of expertise based on existing capabilities a combination of teleological and compositional approaches is needed.

References

[1] B. Wrede, K. J. Rohlfing, M. Hanheide, and G. Sagerer, *Towards Learning by Interacting*, ser. LNAI. Springer, 2009, pp. 139 – 150.

[2] B. Wrede, S. Kopp, K. J. Rohlfing, M. Lohse, and C. Muhl, "Appropriate feedback in asymmetric interactions," *Journal of Pragmatics*, vol. 42, pp. 2369–2384, 2010.

[3] G. Gergely and G. Csibra, "Teleological reasoning in infancy: the naive theory of rational action," *Trends in Cognitive Science*, vol. 7, no. 7, 2003.

[4] J. S. Bruner, *Childs talk: Learning to use language*. New York: Norton, 1983.

[5] A. Fogel, A. Garvey, H. C. Hsu, and D. West-Stroming, *Change Processes in Relationships: A Relational-Historical Research Approach*. Cambridge: Cambridge University Press, 2006.

[6] M. Tomasello, *Constructing language*. Cambridge: A usage-based theory of language acquisition, 2003.

[7] A. Senju and G. Csibra, "Gaze following in human infants depends on communicative signals," *Current Biology*, vol. 18, pp. 668–671, 2008.

[8] G. Gergely and J. Watson, "Early social-emotional development: contingency perception and the social biofeedback model," *Early social cognition: understanding others in the first months of life*, pp. 101–137, 1999.

[9] A. Cangelosi, G. Metta, G. Sagerer, S. Nolfi, C. Nehaniv, K. T. Fischer, B. J., S. T., F. G., W. L., R. B., T. K., D. E., S. K., J., and A. Zeschel, "Integration of action and language knowledge: A roadmap for developmental robotics," *IEEE Transactions on Autonomous Mental Development*, vol. 2, pp. 167–195, 2010.

[10] G. Gergely, H. Bekkering, and I. Kiraly, "Rational imitation in preverbal infants." *Nature*, vol. 415, p. 755, 2002.

[11] M. Beisert, N. Zmyj, R. Liepelt, F. Jung, W. Prinz, and M. M. Daum, "Rethinking rational imitation in 14-month-old infants: A perceptual distraction approach," *PLoS ONE*, vol. 7, 2012.

[12] J. M. Mandler, "On the spatial foundations of the conceptual system and its enrichment," *In: Cognitive Science*, vol. 36, pp. 421–451, 2012.

[13] J. T. Loucks and D. Baldwin, "When is a grasp a grasp? Characterizing some basic components of human action processing," in *Action meets word. How children learn verbs. Oxford: Oxford University Press*, K. Hirsh-Pasek and R. M. Golinkoff, Eds. 228–261, 2006.

[14] K. Pastra and Y. Aloimonos, "The minimalist grammar of action," *Philos Trans R Soc Lond B Biol Sci*, vol. 367, no. 1585, pp. 103–17, 2012.

[15] A. Lemme, A. Freire, G. Barreto, and J. Steil, "Kinesthetic teaching of visuomotor coordination for pointing by the humanoid robot icub," *Neurocomputing*, to appear.

[16] S. Calinon, F. Guenter, and A. G. Billard, "On Learning, Representing, and Generalizing a Task in a Humanoid robot," *IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics: a publication of the IEEE Systems, Man, and Cybernetics Society*, vol. 37, no. 2, pp. 286–98, Apr. 2007.

[17] A. Ijspeert, J. Nakanishi, and S. Schaal, "Learning Attractor Landscapes for Learning Motor Primitives," *Advances in Neural Information Processing Systems*, pp. 1547–1554, 2003.

[18] M. Muehlig, M. Gienger, S. Hellbach, J. J. Steil, and C. Goerick, "Task-level imitation learning using variance-based movement optimization." Kobe, Japan: IEEE, 2009, pp. 1177 – 1184.

[19] S. Degallier and A. Ijspeert, "Modeling discrete and rhythmic movements through motor primitives: a review," *Biological cybernetics*, vol. 103, no. 4, pp. 319–338, 2010.

[20] S. Calinon, *Robot Programming by Demonstration: A Probabilistic Approach*. EPFL Press, 2009.

[21] B. Nemec and A. Ude, "Action sequencing using dynamic movement primitives," *Robotica*, vol. 30, no. 5, p. 837, 2012.

[22] R. Reinhart, A. Lemme, and J. Steil, "Representation and generalization of bi-manual skills from kinesthetic teaching," in *IEEE Humanoids*. IEEE, 2012.

[23] D. Kulić, C. Ott, D. Lee, J. Ishikawa, and Y. Nakamura, "Incremental learning of full body motion primitives and their sequencing through human motion observation," *The International Journal of Robotics Research*, vol. 31, no. 3, pp. 330–345, 2012.

[24] O. Mangin and P.-Y. Oudeyer, "Learning to Recognize Parallel Combinations of Human Motion Primitives with Linguistic Descriptions using Non-Negative Matrix Factorization," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)*, 2012.

[25] M. Pardowitz, R. Haschke, J. J. Steil, and H. Ritter, "Gestalt-based action segmentation for robot task learning," in *IEEE-RAS 7th International Conference on Humanoid Robots (HUMANOIDS)*, 2008, pp. 347 – 352.

[26] A. L. A. Soltoggio and J. Steil, "Primitive-based iterative decomposition of complex trajectories by maximum error in kinesthetic learning." IEEE, to appear.

[27] Y. Yamashita and J. Tani, "Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment," *PLoS computational biology*, vol. 4, no. 11, p. e1000220, 2008.

[28] F. Guerin, N. Kruger, and D. Kraft, "A survey of the ontogeny of tool use: fromsensorimotor experience to planning," *IEEE Trans. Autonomous Mental Development*, to appear.

[29] A. Baranes and P. Oudeyer, "Active learning of inverse models with intrinsically motivated goal exploration in robots," *Robotics and Autonomous Systems*, 2012.

[30] ——, "Intrinsically motivated goal exploration for active motor learning in robots: A case study," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 2010, pp. 1766–1773.

[31] M. Rolf, J. J. Steil, and M. Gienger, "Goal babbling permits direct learning of inverse kinematics," *IEEE Trans. Autonomous Mental Development*, vol. 2, no. 3, pp. 216 – 229, 09/2010 2010.

[32] ——, "Online goal babbling for rapid bootstrapping of inverse models in high dimensions," in *IEEE Int. Conf. Development and Learning and on Epigenetic Robotics*, IEEE. Frankfurt: IEEE, 08/2011 In Press.

[33] V. Hofsten and C., "Action in development," *Developmental Science*, vol. 10, pp. 54–60, 2007.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Cognitive Development Through a Neurologically-Based Learning Framework

John Nassour and Gordon Cheng
Institute for Cognitive Systems
Technische Universität München
Munich, Germany
Email: {nassour, gordon}@tum.de

***Abstract*—We develop through continuous interactions with the world, learning from our experiences of successes as well as our failures. Neurological understanding of the involved mechanisms is beginning to emerge to a level where they can be validated on robots functioning in the real world. For instances: 1) the Anterior Cingulate Cortex "ACC" has shown to contribute to Cognitive control by modulating the error-related signals for both positive as well as negative past experiences, thus, acting as an early warning system "EWS". The notion of Vigilance helps formulate such a mechanism in the manner we learn and develop, thus the way we make decisions. 2) the Orbitofrontal Cortex "OFC" is said to play a role in the way we learn by representing the effective value of reinforcements, thus, regulating decision-making and expectation. These neural mechanisms play an underpinning role in Cognitive Development. In this paper, we show that computational models of these mechanisms have been realized on a robot that can acquire and develop new skills (e.g. walking, throwing).**

I. Introduction

The notion of "cognition" aims to capture the capability of mental activities of human beings for abstracted information from real world. It refers furthermore to their representation and storage in memory. It includes various mental processes like perception, attention, reasoning, learning, recognition, decision, as well as task coordination. All cognitive models have in general the same objective, analyze how human think, reason, remember, perceive, and learn. By studying the behavioral consequences of the brain, cognitive neuroscience promises to delineate the connections between the brain anatomy and the functionality of the human mind that is studied in cognitive psychology.

The research on cognitive neuroscience adds a biologically-inspired intelligent dimension into robotics research. In contrast to traditional robotics control, which focuses on programming robots to solve one specific task in one environment (sense-act, sense-act, sense-act, ...), cognitive robotics control aims to generate intelligence and adaptive behaviors based on animal or human thinking and learning processes (sense-learn-predict-act, sense-learn-predict-act,...).

To be considered as a human partner, human-like robots and human interactive robots should be provided with sophisticated cognitive systems based on open-ended learning toward cognitive developmental robotics. Learning can be considered to be open-ended if it handles tasks that are unknown or even not well-defined previously [1]. Nevertheless, the

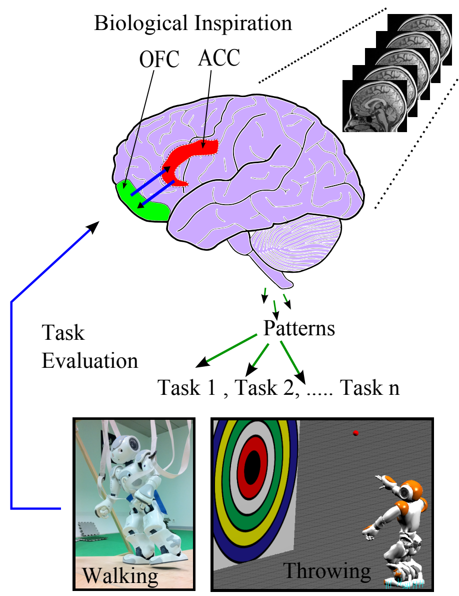

Fig. 1: The conceptual overview of our work. ("ACC" is the Interior Cingulate Cortex; "OFC" is the OrbitoFrontal Cortex.)

learning mechanisms based on fitness function minimization are limited to solve such tasks. Human develops and learns in open-ended way across its lifespan. A human child can learn tasks that he never did before, this can be referred to the mental and physical development. Traditional robots can percept and act only with the external environment, while robot based cognition can be intrinsically motivated by the internal environment, therefore, it can percept and act with the external and the internal environment to reach an intrinsic motivation (e.g. search for missing knowledge in the word-model and trigger a learning process when needed.) [2].

Human learns tasks from their own experiences by self-exploration and observation of others' actions. The evaluation of the achieved task is driven by rewards. Human can improve their skills in order to gain more rewards. Neurobiological studies suggest that the orbitofrontal cortex (OFC) is related to reward dealing in the human brain [3]. Neurons of OFC are the key reward structure of the brain, where reward is coded in an adaptive and flexible way [4].

By observing its cortical activities, studies of the Anterior Cingulate Cortex (ACC) suggest that it is responsible to avoid repeating mistakes [5]. This cortical area acts as an early warning system (EWS) that adjusts the behavior to avoid dangerous situations. It responds not only to the sources of errors (external error feedback), but also to earliest sources of error information available (internal error detection) [6]. EWS has shown to be affected by the tolerant to risks, psychological studies provide further evidences of people's strategies into two classes as in taking or aversion risks [7]. A Developmental study of risky decisions suggests that risk taking behavior is related to human age and development [8].

"NeuroRobotics" research draw on human learning methods in order to improve the autonomy and the robustness of robots for their dealing with environment changes. In connection with these neurological studies, we proposed a learning method based on human learning from experiences (ACC) and inspired by the way the human brain code rewards (OFC), in order to allow a humanoid robot to learn a walking task. With the vigilance threshold concept that represents the tolerance to risk, the method guaranteed the balance between exploration and exploitation. Most task learning methods based on reward use predefined parameters in their reward function [9], [10], which cannot be obtained without previous experiences to achieve the desired task. Learning based adaptive reward don't require any previous information about the reward, it is able to build the experience only based on the reward available information after starting from scratch.

Our approach has been implemented on the NAO humanoid robot, controlled by a bio-inspired neural controller based on a central pattern generator "CPG" [11], [12], see Figure. 1. The learning system adapts three intrinsic parameters in the CPG (the frequency of oscillation, and the motor neuron gain in pitch and roll) in order to walk on flat and sloped terrains.

II. Neuropsychological Inspiration

A. ACC contribution in Cognitive Control

The ACC and neighbouring areas are involved in controlling and monitoring goal-direct-behavior to avoid repeating mistakes. Brown and Braver develop a computational model that shows how ACC not only detects errors, it may predict error likelihood before the error occurs [5]. The ACC is activated proportionally to the observed likelihood of the error. The error-likelihood hypothesis assumes the training signal that affects the ACC is acquired and dopaminergic. The phasic suppression of dopamine, which drives the error-related negativity (ERN) [13], [14], may play the role of a training signal that make ACC activation stronger for contexts with more frequent error. As a result of FMRI observation of subjects' ACC, the ACC cells learn to respond with more activation for cues with high error likelihood. The results suggest that the ACC is involved in cognitive control through its risk-related cortical activity.

B. OFC contribution in Motivational Control

The primate orbitofrontal cortex (OFC) is involved in the motivational control of goal-directed behaviour [15]. It has an essential role in controlling and correcting reward-related and punishment-related behaviour [16]. The neurons of OFC are involved in the processing of motivational values of voluntary action rewards. OFC neurons increase their activities during the expectation of reward and after receiving it [15]. Subjects select more frequently rewards when they have to chose between different rewards at the same time, however, such frequently rewards can be ignored when more delectable rewards become available. It seems that motivational values are not fixed to defined rewards, unlike physical properties. The reward discrimination in some OFC neurons is based on the relative preference rather than the physical properties [15].

C. OFC-ACC Connectivity During Decision-Making

Many neuroscience research study the primate brain regions involved in decision-making and other neurobiological mechanisms [17], [18]. The challenge was not only to detecting the brain regions that exhibit significantly during such mechanisms, but also to understand how different brain regions interact between each other. Cohen et al. designed a FMRI study that separates experimentally the neural activity related high-risk and low-risk choosing from other processes such as reward anticipation and evaluation during the general framework of decision making [17]. They showed that choosing high-risk over low-risk decision was related with increased activity in both ACC and OFC. It seems that OFC carries on reward associations for stimulus [19], and ACC contains mechanisms that control the selection of appropriate behaviors [20]. According to [17], no ACC activities were observed during low risk decision, while both ACC and OFC show a high activation when subjects made high-risk. However, this study was not able to distinguish whether ACC activation are related with small chance of a large reward or large chance of a failure. ACC and OFC exhibited similar patterns for activation and time courses and distinct patterns of functional connectivity. This suggests that they may play different and complementary roles in decision making [17].

D. Risk Taking Behavior

Psychological studies show the probability of sampling with experience based learning in human is reduced with poor past outcomes [21], [22]. They show how adaptive sampling could lead to risk-averse as well as risk-seeking behaviors. Risk tendency may change according to the distribution of the uncertain alternatives and the information about foregone payoffs. Denrell et al. assume that the decision maker mostly samples the ambiguous substitution if its estimated value is positive. However, it explores eventually substitution with negative estimated value [21]. Erev and Barron have shown the role of adaptive sampling in modeling risk taking behavior for systems where decision is based on experience [23].

Based on previous studies in learning from mistakes, coding reward, and adaptive sampling for risk taking modulation; we introduce a learning mechanism that is able to learn and to evaluate humanoid tasks and to optimize its performance.

III. SUCCESS-FAILURE LEARNING

The objective of this learning mechanism is to adapt parameters of a low-level controller and to detect their domain of viability. We designated by Ω the state space of those influential parameters. The mechanism must be able to learn from negative feedback (failure) and positive feedback (success). Therefore it must have experience of success and other experience of failure in the state space Ω. As each action vector $\overrightarrow{v}$ from Ω leads to either success or failure, the mechanism will evaluate whether this vector belongs to a success case or to a failure case. The decision mechanism "go" or "no-go" described in [24] works as an early warning system similar to that in the Anterior Cingulate Cortex [5], [25]. The learning architecture is then based on these two mechanisms and works as shown in Fig.2.

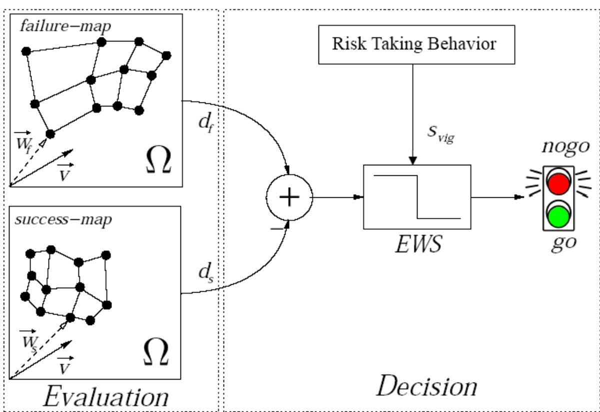

Fig. 2: Success-Failure Learning mechanism with evaluation and decision phases.

We have proposed previously a preliminary model for success-failure mechanism [26].

A. Evaluation Phase

To represent the knowledge in success and in failure, we define two independent neural networks that are well-known Self Organizing Maps (SOM), proposed by Kohonen [27]. Success map S_m learns in case of success trials, and failure map F_m learns in case of failure trials. During the learning, the two maps will be self-organized in the state space that will be therefore divided into three zones: 1) a zone of success represented by success map; 2) a zone of failure represented by failure map; and 3) a zone of conflict that corresponds to the overlapping between the two maps. The evaluation of any vector $\overrightarrow{v}$ from space Ω belonging to success or failure is defined by the distance between $\overrightarrow{v}$ and each map. The distance of a vector with a map is the minimal Euclidean norm between this vector and the closest neuron's weights vector in the state space (the winning neuron). For each $\overrightarrow{v}$ we have therefore two distances: one to S_m called d_s, and another to F_m called d_f. d_s and d_f are then used for the decision process.

B. Decision mechanism

For a vector $\overrightarrow{v}$, the comparison between the distance with success map d_s and the distance with failure map d_f leads to an expected result in the case where the vector was passed to the low level controller (trial). According to expected results, if it may lead to failure, then an Early Warning Signal (EWS) becomes active to avoid the passing into the lower level, and the decision will be "no-go". When EWS is inactive the decision is "go". The decision mechanism is affected by the threshold of vigilance s_{vig}.

C. Vigilance-Related Development

Psychological research studies suggest that some people are more tolerant to risk than others who are more cautious [28][29][30]. The vigilance is related to human learning in connection with decision making [31]. In the standard psychological assessment of risk taking, people are classed as risk seeking or risk averse [7].

Through the observation of particular areas located in cerebral cortex in the brain responsible for cognitive control, neuropsychological studies demonstrated a switching in human learning strategies around the age of twelve years. This switch from learning with positive feedback to learning with negative feedback probably comes from the combination of brain maturing and experience[8].

In our study for robot tasks learning by success and failure maps, we introduced the concept of vigilance in order to control the learning process in the two maps (success and failure) and manage the learning cycle while avoiding or taking risks according to the system's needs. The vigilance is represented by a threshold s_{vig} that is used to adjust the early warning signal in the decision mechanism. This threshold describes the tolerance of risk.

In our previous work, the vigilance threshold was modulated according to the number of trials [32], see Figure 3. The

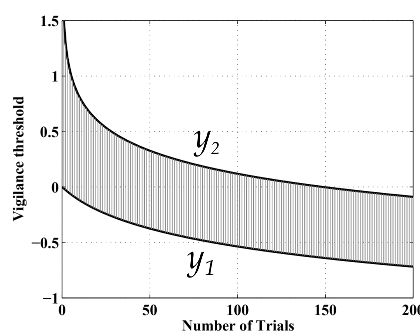

Fig. 3: Vigilance Model related to learning iterations, $y_1 \leq s_{vig} \leq y_2$. The risk behavior will change from prudence at the beginning of learning to adventure at the end [32].

vigilance was modulated to change the risk behavior from prudence at the beginning of learning to adventure at the end. An example of vigilance threshold modulation is given as following (see Fig.3):

$$y_1 \leq s_{vig} \leq y_2 \qquad \begin{cases} y_1 = a_1 - b_1 * log((x + c_1)^2) \\ y_2 = a_2 - b_2 * log((x + c_2)^2) \end{cases} \tag{1}$$

The coefficients values are ($a_1 = 0.9, a_2 = 1.47, b_1 = b_2 = 0.15, c_1 = c_2 = 20$) and were chosen after several attempts.

y_1 and y_2 chosen curves ensure smooth change between the prudence and adventure above mentioned behaviors.

Vigilance modulation is an important approach that can drive the open end learning, it can be modulated in opposite way, start with taking risk then switch to risk avoiding behavior.

However, vigilance can be also modulated according to the distribution of the uncertain alternatives. In the next section, the vigilance threshold was adapted to drive the sampling process in the way that ensures success and failure maps learn and converge together.

D. Reward Coding

Most reinforcement learning based robotic walking studies use predefined constants to determine the maximum and the minimum reward or to determine the multiplier factors [9], [10]. Reward coding is a way that qualifies succeeded trials differently according to an optimized criterion. Each trial will have its own weighted reward representing the objective criterion to be optimized. During each learning step, neurons will get closer to trials with high rewards rather than to trials with low rewards. If the optimized criterion was the efficiency of learned task, the minimum value of the reward related multiplication factor r_{min} matches the trial with lowest efficiency. While the maximum value of the reward related multiplication factor r_{min} matches the trial with highest efficiency. This matching is done in adaptive way during the learning process, see Figure 4.

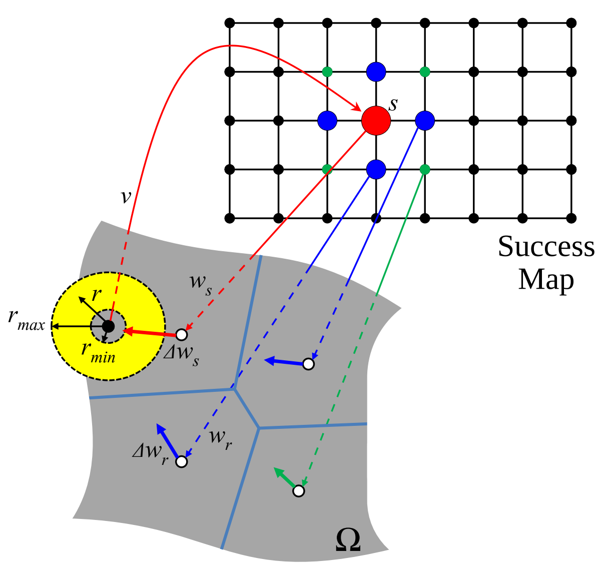

Fig. 4: Success map adaptation. r is a reward-related multiplication factor. r_{min} and r_{max} are the minimum and the maximum values for the multiplication factor.

After enough numbers of trials, this will result in a shift of the map into a spatial area associated with the highest rewards.

By introducing the concept of adaptive coding of reward it will be possible to scale the quality of a trial according to the quality in previous experiences even when starting from scratch. After learning, the optimal parameter is presented by the success map neuron that is close to the trial with maximum reward in training set.

The general diagram of the success-failure learning is presented in Figure 5.

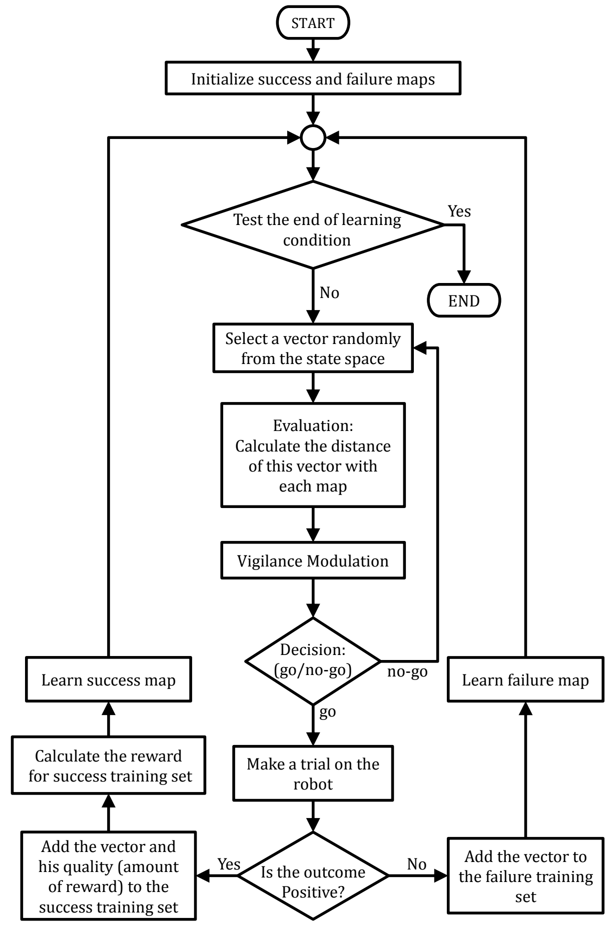

Fig. 5: Flow diagram for success-failure learning.

IV. Learning with Adaptation

We apply the learning mechanism proposed in the previous sections in order to learn efficient walking for a bipedal humanoid robot, NAO. We used success-failure learning to learn in a space of intrinsic parameters of the CPG controller (the frequency of oscillation, and the motor neuron gain in pitch and roll), the basic CPG model is presented in [32]. The optimization of walking efficiency was studied in term of energy as in [33]. The efficiency with which a muscle operates is defined in [33] by

$$efficiency = \frac{mechanical\ work\ done}{metabolic\ energy\ consumed} \tag{2}$$

This study is also generalized from a muscle to whole body movements like walking, and running [34], [35]. With inspiration from biomechanical studies, the efficiency of walking for a humanoid robot can be described in a similar fashion.

Our objective is to simultaneously learn and optimize walking with success-failure on-line learning. The robot learns to walk a $1.5[m]$ trajectory with start and end lines. In case of succeeded trials, the trainer sends a reward signal to the robot by caressing the head equipped with electrostatic sensors. The walking efficiency is calculated for each trial as:

$$\eta = \frac{E_k}{E_e} \tag{3}$$

Where E_k is the kinetic energy of a trial, E_e is the required electric energy for the entire trajectory.

The introduction of the efficiency for success map learning will shift the neurons of this map into the area in which the walking efficiency is high. Figure 6 shows the reward coding for success map in the beginning of learning (after four successful trials), and at the end of learning. Each sphere corresponds to a succeeded trial whose diameter represents the reward of this trial in the success map. The interest of using this technique is to make success-failure learning search for new trials in the space area where walking efficiency in term of energy is high. In other words, this leads to learn and optimize in a defined space. Figure 7 shows success maps after learning to walk on flat terrain with and without the technique of reward coding adaptively. In Figure 7(a), the success map learns all successful trials with the same opportunity, i.e. with the same reward. In Figure 7(b), the success map learns successful trials in accordance with their adaptive rewards. Trials with high reward influence success map neurons more than trials with low reward. Therefore, the success map will be attracted to the area where reward is high. This is influenced by the differences between highest and lowest rewards (scaling range limits: $[r_{min}, r_{max}]$). In this study, r_{min} and r_{max} are set to 0.1 and 2.5.

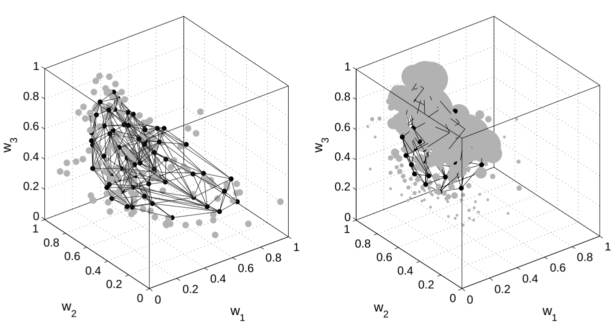

(a) Learning with same reward. (b) Learning with adaptive reward coding.

Fig. 7: The effect of reward coding on success map. Success map after learning with the same reward for all successful trials (a). Success map after learning with adaptive reward coding (b). Gray spots represent successful trials reward. Note that the map on the right moves into the area where rewards are high (representing high efficiency).

Regarding the learning frameworks with and without the application of reward coding adaptively shown in Figure 7, performance was increased by 60%. This was calculated by the ratio of the highest efficiencies matched by neurons of both success maps. The ratio of the lowest efficiency of the neurons of success maps has increased by 40%. In order to provide sufficient precision in the network for our task, we have empirically selected a $5 \times 5 \times 5$ dimensional network space to represent the success and failure maps. Learning occurred with 500 trials for each case.Computationally, all the processing of this learning framework in simulations as well as on the real robot can be performed in real-time, thus making our approach feasible for training on the real robot. Within the same cycle, joint angle commands are calculated in real-time and sent to joint motor circuit boards of NAO every $10[ms]$. This is done inside a high priority thread on the robot. Physically each trial requires about 3 minutes, which includes learning and the experimental set up. A complete learning session in the robot usually takes about one week. Both Learning frameworks shown in Figure 7 start from scratch.Adaptive sampling driven by vigilance threshold ensures to have the same size of training sets to learn success map and failure map.(A video shows NAO humanoid robot achieving the walking task is available on: http://web.ics.ei.tum.de/~nassour/naowalking.wmv.)

V. Conclusion

This paper has brought several ideas from different bodies of research. Research in machine learning, neuroscience, psychology, and robotics are involved in cognitive development. Understandings within human brain research help provide human-like learning mechanisms that can be implemented on robots. Our neurologically grounded learning framework imitates part of the functionality of the anterior cingulated cortex involved in learning from mistakes, and the orbitofrontal cortex involved in reward coding adaptively. This Success-Failure learning cycle forms an important part of our cognitive development architecture in order for robots to learn and acquire different physical and mental skills.

In this paper, we showed how vigilance is modulated differently as our robot develops over time under different conditions during decision making. These vigilance modulations can be said to play an important role throughout a lifetime of human as well as robot developments.

References

[1] R. Gomes, "Towards open ended learning: budgets, model selection, and representation," Ph.D. dissertation, California Institute of Technology, 2011.

[2] M. Oubbati and G. Palm, "Self-motivated learning robot," *9th Epigenetic Robotics International Conference (Epirob09): Modeling Cognitive Development in Robotic Systems, Venice, Italy, November 12-14*, 2009.

[3] S. D. Iversen and M. Mishkin, "Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity," *Experimental Brain Research*, vol. 11, pp. 376–386, 1970.

[4] S. Kobayashi, O. P. de Carvalho, and W. Schultz, "Adaptation of reward sensitivity in orbitofrontal neurons," *The Journal of Neuroscience*, vol. 30, no. 2, pp. 534–544, January 2010.

[5] J. W. Brown and T. S. Braver, "Learned predictions of error likelihood in the anterior cingulate cortex," *Science*, vol. 307, pp. 1118–1121, 2005.

[6] R. B. Mars, M. G. Coles, M. J. Grol, C. B. Holroyd, S. Nieuwenhuis, W. Hulstijn, and I. Toni, "Neural dynamics of error processing in medial frontal cortex." *Neuroimage*, vol. 28, no. 4, pp. 1007–1013, December 2005.

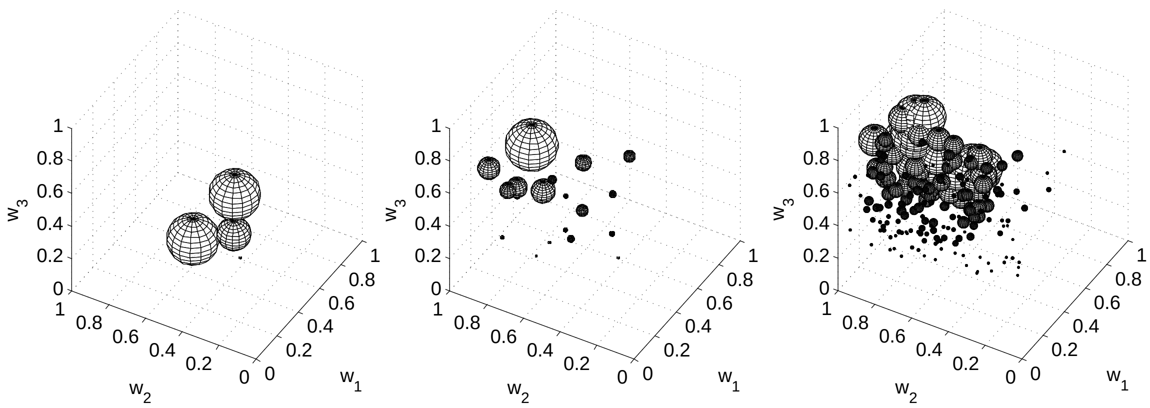

(a) Reward after 4th success. (b) Reward after 50 trials. (c) Reward after 500 trials.

Fig. 6: Succeeded trials' reward related to walking efficiency for learning success map. Where w_1 is motor neuron gain in pitch, w_2 is motor neuron gain in roll, and w_3 is σ_s that is related to the oscillation frequency.

[7] X. Wang, D. Kruger, and A. Wilke, "Towards the development of an evolutionarily valid domain-specific risk-taking scale," *Evolutioniary Psychology*, vol. 5, no. 3, pp. 555–568, 2007.

[8] L. Van Leijenhorst, P. M. Westenberg, and E. A. Crone, "A developmental study of risky decisions on the cake gambling task: Age and gender analyses of probability estimation and reward evaluation," *Developmental Neuropsychology*, vol. 33, no. 2, pp. 179–196, 2008.

[9] T.-H. S. Li, Y.-T. Su, S.-W. Lai, and J.-J. Hu, "Walking motion generation, synthesis, and control for biped robot by using pgrl, lpi, and fuzzy logic," *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 41, no. 3, pp. 736–748, 2011.

[10] G. Endo, J. Morimoto, T. Matsubara, J. Nakanishi, and G. Cheng, "Learning cpg-based biped locomotion with a policy gradient method: Application to a humanoid robot," *International Journal of Robotics Research*, vol. 27, pp. 213–228, February 2008.

[11] J. Nassour and G. Cheng, "Biologically-inspired neural controller based on adaptive reward learning," *Front. Comput. Neurosci. Conference Abstract: Bernstein Conference*, 2012.

[12] F. B. John Nassour, Vincent Hugel and G. Cheng, "Qualitative adaptive reward learning with success failure maps: Applied to humanoid robot walking," *IEEE Transactions on Neural Networks and Learning Systems [In press]*, 2012.

[13] W. J. Gehring, M. G. H. Coles, D. E. Meyer, and E. Donchin, "The error-related negativity: An event-related potential accompanying errors," *Journal of Psychophysiology*, vol. 27, p. S34, 1990.

[14] W. J. Gehring, Y. Liu, J. M. Orr, and J. Carp, *Oxford handbook of event-related potential components*. New York: Oxford University Press., 2012, ch. The error-related negativity (ERN/Ne), pp. 231–291.

[15] L. Tremblay and W. Schultz, "Relative reward preference in primate orbitofrontal cortex," *Nature*, vol. 398, no. 6729, pp. 704–8, April 1999.

[16] E. T. Rolls, "The orbitofrontal cortex," *Philosophical Transactions of the Royal Society B*, vol. 351, pp. 1433–1444, 1996.

[17] M. Cohen, A. Heller, and C. Ranganath, "Functional connectivity with anterior cingulate and orbitofrontal cortices during decision-making," *Cognitive Brain Research*, vol. 23, no. 1, pp. 61 – 70, 2005.

[18] E. Bicho, W. Erlhagen, L. Louro, and E. C. e Silva, "Neuro-cognitive mechanisms of decision making in joint action: A human-robot interaction study," *Human Movement Science*, vol. 30, no. 5, pp. 846 – 868, 2011.

[19] E. T. Rolls, "The orbitofrontal cortex and reward," *Cerebral Cortex*, vol. 10, no. 3, pp. 284–294, 2000.

[20] V. Van Veen, J. D. Cohen, M. M. Botvinick, V. A. Stenger, and C. S. Carter, "Anterior cingulate cortex, conflict monitoring, and levels of processing," *NeuroImage*, vol. 14, no. 6, pp. 1302–1308, 2001. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/17604651

[21] J. Denrell, "Adaptive learning and risk taking," *Psychol Rev.*, vol. 114, no. 1, pp. 177–187, 2007.

[22] J. G. March, "Learning to be risk averse," *Psychological Review*, vol. 103, no. 2, pp. 309–319, Apr 1996.

[23] G. Erev, Ido; Barron, "On adaptation, maximization, and reinforcement learning among cognitive strategies," *Psychological Review*, vol. 112, no. 4, pp. 912–931, Oct 2005.

[24] K. Matsumoto, W. Suzuki, and K. Tanaka, "Neuronal correlates of goal-based motor selection in the prefrontal cortex," *Science*, vol. 301, no. 5630, pp. 229–32, 2003.

[25] J. W. Brown and T. S. Braver, "A computational model of risk, conflict, and individual difference effects in the anterior cingulate cortex," *Brain research*, vol. 1202, pp. 99 –108, 2008.

[26] J. Nassour, P. Henaff, F. B. Ouezdou, and G. Cheng, "Experience-based learning mechanism for neural controller adaptation: Application to walking biped robots," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, October 11-15, St. Louis, MO, USA*, 2009, pp. 2616–2621.

[27] T. Kohonen, *Self-Organizing Maps*, ser. Springer Series in Information Sciences. Berlin, Heidelberg: Springer, 1995, vol. 30.

[28] J.-L. van Gelder, R. E. de Vries, and J. van der Pligt, "Evaluating a dual-process model of risk: affect and cognition as determinants of risky choice," *Journal of Behavioral Decision Making*, vol. 22, no. 1, pp. 45–61, 2009.

[29] B. Pawlowski, R. Atwal, and R. I. M. Dunbar, "Sex differences in everyday risk-taking behavior in humans," *Evolutionary Psychology*, vol. 6, no. 1, pp. 29–42, 2008.

[30] P. Horvath and M. Zuckerman, "Sensation seeking, risk appraisal, and risky behavior," *Personality and Individual Differences*, vol. 14, no. 1, pp. 41–52, 1993.

[31] H. Ahn and R. W. Picard, "Affective-cognitive learning and decision making: A motivational reward framework for affective agent," in *The 1st International Conference on Affective Computing and Intelligent Interaction. October 22-24*, 2005.

[32] J. Nassour, P. Hénaff, F. B. Ouezdou, and G. Cheng, "A study of adaptive locomotive behaviors of a biped robot: patterns generation and classification," in *Proceedings of the 11th international conference on Simulation of adaptive behavior: from animals to animats*, ser. SAB'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 313–324.

[33] B. Abernethy, *The biophysical foundations of human movement*, 2nd ed. HUMAN KINETICS, 2005.

[34] N. Berryman, M. Gayda, A. Nigam, M. Juneau, L. Bherer, and L. Bosquet, "Comparison of the metabolic energy cost of overground and treadmill walking in older adults," *European Journal of Applied Physiology*, pp. 1–8, 24 August 2011.

[35] R. Margaria, *Biomechanics and energetics of muscular exercise*, first edition ed. Oxford University Press, USA, November 18 1976.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Forming Goal-directed Memory for Cognitive Development

Erhard Wieser and Gordon Cheng

Keywords: embodied representation, meaningful association, coaching

I. INTRODUCTION

A. The Challenge of Skill Acquisition

Cognitive development [1] gives rise to the major challenge of skill acquisition, *i.e.*, the learning of a new physical skill for object manipulation. This learning is very difficult because the system designer cannot specify *a priori* all the necessary robot actions depending on the latest states of objects and environmental conditions. Even slightly different environments or facing new objects lead to an undesirable re–programming of the action programs of the robot.

B. Related Work

At the perceptual level, approaches to skill acquisition are imitation learning and coaching. Researchers created an imitation learning system [2], [3] controlling a humanoid robotic hand. Their imitation system learns hand postures by observing the hand of a human with a camera. Their imitation system uses a higher order Hopfield network (HHOP) as the main mechanism. In [3], Chaminade *et al.* showed that the HHOP was able to generalize between the learned patterns to a limited extent, *i.e.*, it could generate a few new gestures correctly even though they were *not* trained *a priori*. On the way from the perceptual level to motor control, system designers have to deal with object manipulation. In embodied cognition, objects are represented by sensorimotor patterns to reduce the symbol grounding problem [4]. Wörgötter *et al.* introduced their concept of object–action complexes [5] to describe possible actions, which a robot can perform on a given object.

C. Our approach

Our long term goal is the creation of a new cognitive architecture for skill acquisition. A cognitive architecture [6] is fundamental to any intelligent robot. In this paper, we present a first part of our future architecture. That part is based on our idea of *meaningful associations*. So far, a meaningful association is the link between a given percept, a learned goal state, and a corresponding action leading to that goal state, similar to the concept of object–action complexes. But at a later developmental stage, a meaningful association also includes cross links to abstract values (*good* or *bad* percepts / actions) and memories, which bias the current actions of the robot. For these associations, we do *not* provide *a priori* symbolic knowledge at all, instead we put the emphasis on the close interaction between the robot, its local sorrounding, and its human coach. A human coach shows the robot these meaningful associations by giving tactile feedback. Through continuous interaction with objects and a coach, the robot increases the amount of such associations, representing an increase of knowledge. Knowledge is internally stored by forming goal–directed memory contents. So far, these contents ground themselves in associations between sensorimotor and neural patterns, representing percepts, goals, and goal–directed actions (later, also values). Our approach has the potential to exploit many cross–modal associations, *e.g.*, visual, tactile, which can in turn bias the behaviour of the robot in a useful manner. Therefore, we developed the foundations [2], [3] in the following ways:
As a part of our cognitive architecture, we created a perception system with goal–directed memory to trigger goal–directed physical actions of the robot (a Humanoid Robot NAO). Our perception system processes latest visual data, enables visual servoing, and influences the behaviour of a robot by using previous experiences stored in an episodic memory module. The episodic memory is implemented by Hopfield networks. First, in contrast to [2] and [3], we extend the feature space of the Hopfield networks in order to capture not only simple shapes, but also basic colours. Second, we combine the memory output with a pattern associator, in order to link a recalled percept to a learned goal state. This goal state can in turn trigger a corresponding goal–directed action of the robot resulting into a new percept.

II. SYSTEM DESCRIPTION

A functional diagram of our perception system with goal–directed memory is depicted in Fig. 1. The sensor modalities of our system are vision and tactile feedback. However, the main modality is vision using any colour camera built into a robot. Tactile feedback only initiates the storage of the latest percept into the episodic memory. We implemented several fields of simple receptive cells, each responsive to certain visual features, such as shape, contour, and colour. Currently, we use four types of simple receptive cells, so called *simple retina cells*, as well as higher level receptive cells. Simple receptive cells are sensitive to shape and contour, mainly of objects in the foreground, and to each of the basic colours red, green, and blue. Each of these cells corresponds to a bipolar neuron, *i.e.*, it fires (activation value $+1$) when a certain feature is present, or it does not fire (activation value -1) when the feature is absent. A higher level receptive cell is only active, when both a shape cell and a corresponding colour cell are active at the same time. We implemented a simple, but robust and flexible visual servoing module, which directs the head of the robot towards an object of interest. Our visual servoing

Institute for Cognitive Systems, Technische Universität München, Karlstrasse 45/II, 80333, München, Germany, Email: see http://www.ics.ei.tum.de/

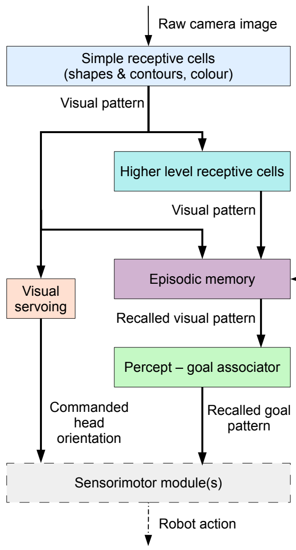

(a) Perception system with goal–directed memory.

(b) NAO robot memorizes meaningful percepts by a tap on its head.

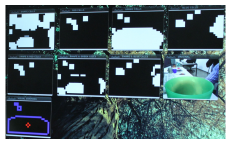

(c) Display of the current state of simple cells (first row) and higher level cells (second row), as well as visual servoing (bottom left) at the time when photo 1(b) was taken.

Fig. 1. Our perception system with goal–directed memory, depicted in functional diagram 1(a). Simple receptive cells emulate bio–inspired vision. Visual servoing directs the head of the robot to an object of interest, *e.g.*, a green cup. A human coach guides the learning by giving tactile feedback, see fig. 1(b). Tactile feedback activates the storage of visual patterns through the episodic memory. The percept–goal associator links a recalled visual pattern to a goal, which in turn triggers goal–directed actions of the robot. Goal–directed actions are realized by sensorimotor modules. During an executed action, the system is in a closed loop with its environment, and open to new (recalled) percepts and tactile feedback. These can influence the executed action at any time.

module moves the robot head, so that the object is in the middle of the field of view of its camera. Here, an important aspect is to note that at this stage, our perception system does *not* regard an object as an object. Our overall system will bootstrap this skill of higher level categorization at a later developmental stage, after enough interactions between the robot and environment have occurred. The activation signals from the simple as well as higher level receptive cells run into the episodic memory module. The storage of percepts is triggered by tactile reinforcement on the robot through a human coach. The episodic memory uses either the classic Hopfield network [7], or the higher order Hopfield network (HHOP) presented in [2], [3]. The drawback of the HHOP is its huge need of computer memory (given N neurons, then N^3 weights need to be saved, compared to N^2 weights in a classic Hopfield net). However, the advantage of HHOP is its limited ability to generalize to new patterns based on the already learned ones [3]. The episodic memory recalls a known visual pattern, *e.g.*, of an object the robot has experienced before. The percept–goal associator links an abstract goal (*e.g.*, *lift the object up*) to that recalled visual pattern by using a feedforward neural network. The goal state is represented by a neural pattern, which self–emerges through the interaction with a human coach. In sum, once the robot sees an already known object, our perception system recalls a visual pattern representing that object. The recalled visual pattern is in turn associated with a suitable goal (affordance, *e.g.*, *lift the object up*) enabling the robot to initiate action programs leading to that goal.

III. RESULTS

Our perception system is a part of our new cognitive architecture and is still an ongoing project. The simple and higher level receptive cells as well as the visual servoing module are fully implemented. So far we focused only on that part of the episodic memory, which responds to shape and contour. Within that part, we compared the classic Hopfield and higher order Hopfield network (HHOP) performance. We validated that the classic Hopfield network as well as the HHOP are not sufficient for usage as an episodic memory for a cognitive architecture when only the feature of shape and contour is considered. This is due to the consideration of only one feature (shape respectively contour) on the one hand, and to memory interferences on the other hand. We showed that the addition of colour features enhanced their performance by increasing the dimensionality of the overall stored pattern. Like Chaminade *et al.* [3], we observed that new patterns emerged within the higher order Hopfield network, which were *not* stored previously. All the implemented modules of our system run successfully on a NAO robot according to the descriptions in part II.

IV. CONCLUSION

We presented a perception system with goal–directed memory forming an important part of our cognitive architecture. Our perception system uses basic visual features and relies on tactile feedback given by a human coach to create meaningful associations between percepts and goals triggering goal–directed actions.

REFERENCES

[1] 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*, vol. 37, no. 2-3, pp. 185–193, 2001.

[2] E. Oztop, T. Chaminade, G. Cheng, and M. Kawato, "Imitation bootstrapping: Experiments on a robotic hand," *Proceedings of the 5th IEEE-RAS International Conference on Humanoid Robots*, pp. 189–195, 2005.

[3] T. Chaminade, E. Oztop, G. Cheng, and M. Kawato, "From self-observation to imitation: Visuomotor association on a robotic hand," *Brain research bulletin*, vol. 75, no. 6, pp. 775–784, 2008.

[4] R. Pfeifer, "Building fungus eaters: Design principles of autonomous agents," *From animals to animats: Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, pp. 3–12, 1996.

[5] F. Wörgötter, A. Agostini, N. Krüger, N. Shylo, and B. Porr, "Cognitive agents – a procedural perspective relying on the predictability of object-action-complexes (oacs)," *Robotics and Autonomous Systems*, vol. 57, no. 4, pp. 420–432, 2009.

[6] P. Langley, J. Laird, and S. Rogers, "Cognitive architectures: Research issues and challenges," *Cognitive Systems Research*, vol. 10, no. 2, pp. 141–160, 2009.

[7] J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?
November 29, 2012. Osaka, Japan

Goal Babbling: a New Concept for Early Sensorimotor Exploration

Matthias Rolf, Jochen J. Steil

I. Coordination Problems

The human body possesses more than 600 skeletal muscles [1]. Performing purposeful actions to achieve some behavioral goal requires a high degree of coordination of these many degrees of freedom. Yet, human infants are born without the most basic coordination skills like reaching for an object [2], which poses the *learning* of sensorimotor coordination as a fundamental problem in human development. Understanding this ability to learn, and utilizing it for modern robotics systems is one of the major goals of the research fields of *cognitive* [3] and *developmental robotics* [4], [5].

We investigate the learning of reaching skills as an exemplary coordination skill. The problem of reaching is to find motor commands (e.g. joint angles of a robot arm) that move the hand, or the robot's end-effector towards some desired position in space. Thereby motor commands q and outcomes x are connected by a causal relation which is denoted as the forward function $f(q) = x$. Learning needs to invert this relation in order achieve some desired outcome x^*. This problem setup is not only illustrative, but very *prototypical* for other coordination problems: it asks the very general question of *how* to achieve some behavioral goals by means of actions. The skill of reaching itself is also *fundamental* for both robots and humans, since the positioning in space is necessary for any use of the robot's gripper or the human's hand.

Successful reaching skills can be well understood with the notion of *internal models* [6], [7], whereas forward models predict the outcome of an action and inverse models suggest actions in order to achieve a desired outcome. The bootstrapping of internal models without explicit prior-knowledge requires experience that has to be generated by *exploration*. Machine learning approaches thereby traditionally rely on an exhaustive exploration of all possible motor commands, frequently generated by means of an entire random procedure, which is referred to as "motor babbling" [8], [9]. After the data generation phase, learning and coordination can be phrased in a variety of ways [10], [11], [12]. Yet, exhaustive exploration can not be achieved on high-dimensional motor systems such as the human body, modern humanoid robots, or biomimetic robots like elephant trunks. The sheer number of combinations of commands for different actuators is too large to be explored in the lifetime of any learning agent. Understanding human motor development, as well as the successful application of future robotic systems like the *Bionic Handling Assistant* (see Fig. 3), demands for concepts and methods that succeed in sensorimotor learning even without fully exploring the space of possible motor commands.

Matthias Rolf and Jochen J. Steil are with the Research Institute for Cognition and Robotics (CoR-Lab) at Bielefeld University, Germany. Mail: {mrolf,jsteil}@CoR-Lab.Uni-Bielefeld.de

II. Infant Development

The standard models for the learning of coordination skills demand either an exhaustive exploration of all actions [8], [10], [9], [11], [12], or prior knowledge about the action space and forward function [13], [14], [15]. Therefore the acquisition of the coordination skill is divided into separate stages of (random) exploration, learning, and exploitation of the learned mechanisms. However, exhaustive exploration does neither provide an explanation of infant's efficiency in sensorimotor development, nor does it provide a feasible approach for artificial agents to learn in high-dimensional domains.

Nevertheless, the generation of random actions by means of motor babbling has been repeatedly motivated [16], [8], [17] by Piaget's view on infant development [18]. Piaget suggested that development is organized in distinct stages and that, at first, infants do not perform purposeful actions. "The implication of [Piaget's] proposal is that the early behavior of the neonate is essentially random and insensitive to contextual information. Recent research suggests that some re-thinking of this extreme position is necessary" [19]. Contrary to Piaget's suggestions, and the random motor babbling approach, infant developmental studies over the last three decades have found conclusive evidence for coordinated behavior even in newborns. Examples include orienting towards sounds [20], tracking of visual targets [21], and apparent reflexes that have been re-discovered as goal-directed actions [22], [23]. "These behaviors are fragile and inconsistent, which explains why they were overlooked for quite some time" [19].

In the case of reaching, it has been shown that newborns attempt goal-directed movements already few days after birth [24], [25]. Von Hofsten showed that, when salient objects are in the visual field, infants produce more arm movements towards that object, than movements away from it. This indicates a strong role of "learning by doing" instead of random exploration and that infants learn to reach by trying to reach: "Before infants master reaching, they spend hours and hours trying to get the hand to an object in spite of the fact that they will fail, at least to begin with" [26]. From a machine learning point of view, these findings motivate to devise methods that closely intertwine exploration, learning, and exploitation, instead of organizing these aspects in distinct and subsequent stages.

Findings of early goal-directed actions are complemented by studies investigating the structure of infants' reaching attempts over the course of development. When infants perform the first *successful* reaching movements around the age of four months, these movements are controlled in an entire *feedforward* manner [27], [28]. This strongly indicates the use of an inverse model for motor control, which selects one solution and applies it without corrections. The importance of feedforward control does not diminish over the course of development, which is well known from prism-glass experiments [29], but the skill is later on augmented by mechanisms that allow for more adaptive movements and error corrections by means of visual feedback [30]. Moreover, the earliest reaching movements are rather jerky and suboptimal in the sense that the distribution and timing of muscular forces is more complicated than actually necessary [31], [2], [32].

In short, infants appear to follow a very efficient pathway, on which one initial solution is learned, and directly used for goal-directed behavior. Only later on these movements are gradually optimized and become more adaptive. While this pathway is very intuitive, it is orthogonal to the motor-babbling approach which first attempts to gather full knowledge about the sensorimotor space, from which particular solutions can be derived afterwards.

III. A New Concept: Goal Babbling

The general idea that connects early goal-directed movements and initial feedforward control is to take redundancy as an *opportunity* to reduce the demand for exploration, instead of a burden that has to be dealt with. If there are multiple ways to achieve some behavioral goal, there is no inherent need to know all of them. Of course, this requires an exploration mechanism that can generate relevant training data without exhaustive exploration. Our main hypothesis is that infants' early goal-directed movements do not only reflect an early exploitation of knowledge, but that they constitute the very mechanism to *generate* that knowledge by exploration, and therefore enable an efficient learning of valid solutions for the coordination problem. Consequently, our main research goal concerns the general mechanism of goal-directed exploration:

> **Research goal 1**: Conceptualize and understand early goal-directed movements as mechanism for the bootstrapping of coordination skills.

As a basis for this investigation, we have introduced the notion of "goal babbling" [33]:

> **Definition: Goal babbling** is the bootstrapping of a coordination skill by repetitively trying to accomplish multiple goals related to that skill.

A central aspect is, of course, trying to accomplish goals, which corresponds to infants' attempts to perform goal-directed movements. Several other aspects of this definition need to be highlighted in order to distinguish this concept from other approaches: Goal babbling aims at the *bootstrapping* of coordination skills such that a skill can be learned without prior knowledge, or non-goal-directed prior exploration. Goal babbling defines this as a *repeated* process, which implies that the skill acquisition is incremental and ongoing, as opposed to stage-like organizations of exploration and learning [8], [9]. Goal babbling applies to domains with multiple related goals. For reaching problems this is naturally given by a

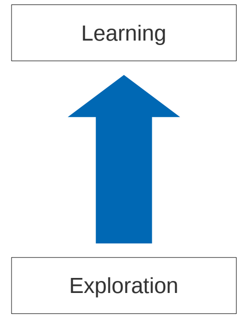

(a) Uninformed motor babbling

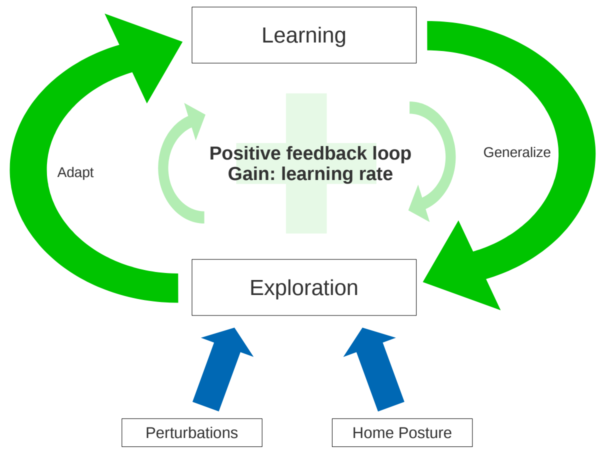

(b) Goal babbling

Fig. 1. In contrast to uninformed exploration processes like motor babbling, exploration and learning mutually inform each other in goal babbling. This organization constitutes a positive feedback-loop during bootstrapping which substantially accelerates learning.

continuous space of possible hand positions. This exploration across multiple goals stands in contrast to typical scenarios in reinforcement learning, in which only a single desired behavior is considered [34]. Given this research goal and the definition of goal babbling, we address several conceptual questions:

- Is goal babbling *possible* at all, and what are the mechanisms necessary to enable it?
- Does it actually permit a bootstrapping that is *scalable to high dimensions*?
- What are observable characteristics of such a bootstrapping process?

IV. Approach and Results

Goal babbling does not refer to a particular algorithm, but to a concept that can be methodically investigated by various means. A recent approach that is compatible with the concept of goal babbling has been introduced in [35]. Baranes' model attempts to learn a partial *forward model*. In this scenario, goal-directed movements are performed by analytically inverting the iteratively learned forward model and performing conventional feedback control. Goal babbling then generates a distribution of actions that lies in the typical regime

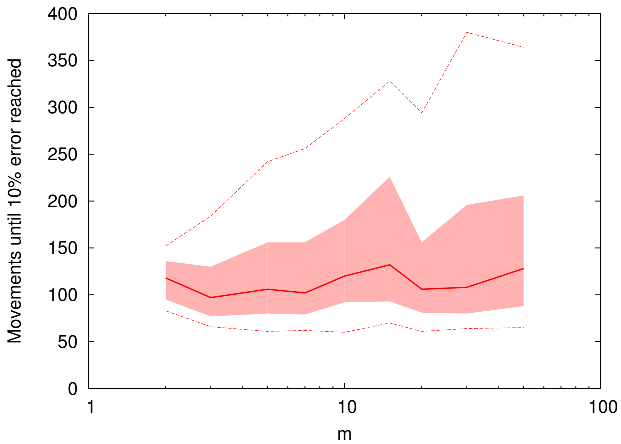

Fig. 2. Goal babbling scales to very high-dimensional problems, as shown by the only marginal increase of exploratory cost for reaching with between $m = 2$ and $m = 50$ degrees of freedom.

Fig. 3. Goal babbling allows to efficiently learn reaching with the *Bionic Handling Assistant*. The feedforward control with an inverse model allows to cope with intense sensory noise and delays.

of the feedback controller when trying to reach for goals, such that not the entire action space needs to be explored.

In contrast to Baranes' model, we investigate the learning of *inverse models* by means of goal babbling, and therefore focus on learning the coordination skill directly, without relying on analytical inversion mechanisms. This approach resembles infants' developmental pathway, which serves as an example of efficiency, by acquiring at first one valid solution that can be used for feedforward control. Learning inverse models in high-dimensional, redundant domains has, so far, only been possible with error-based mechanisms [13], [14] that use prior knowledge in order to generate a learning signal. The demand for a bootstrapping mechanism clearly disqualifies error-based methods due to their inherent need for prior knowledge. Instead, we focus on learning from autonomously self-generated examples. Learning inverse models from examples was believed to be impossible due to the non-convex solution sets in non-linear redundant domains [14]. Consequently, the second research goal concerns this methodological aspect:

> **Research goal 2**: Enable the learning of inverse models from examples in non-linear and redundant domains.

Finding an exploration scheme that can realize this goal clearly needs to cope with non-convex solution sets. Previous studies have only shown how to deal with non-convexity locally, either by reformulating the problem into a differential one [10], or by using prior knowledge to start learning from a well-initialized state [36]. We show that goal babbling provides an elegant account [33] for this long-standing problem, and demonstrate an algorithm that can utilize *goals as reference structure* in order to resolve possibly inconsistent solutions. However, non-convexity is not the only problem to deal with. While non-convexity makes it difficult to handle multiple solutions for the same outcome, the initial problem is to *find at least one correct solution* to realize the desired outcomes and, hence, to invert the causal relation of the forward function in a reliable manner. This *inversion of causality* is a general problem for exploration schemes, since the direction $x \rightarrow q$ can not be directly queried within coordination problems. Random motor babbling can theoretically solve the problem because it simply explores all actions, such that the necessary ones are also explored. This, however, is practically not feasible in high-dimensional domains. The inversion of causality has a distinct characteristic in goal-directed exploration schemes which tend to get stuck in only partial solutions of the coordination problem [37], [38], [39], in which only a subset of goals can be successfully reached. The general pattern to solve that problem is to introduce exploratory noise into the process [40], [36]. We show novel results for a mathematical theory of example-based learning of inverse models, and provide a proof [41] that goal babbling succeeds in linear domains.

Given a goal babbling method that can learn inverse models from examples, the consequential goal is to practically prove the success and usefulness of goal babbling in high-dimensional domains:

> **Research goal 3**: Devise a practical algorithm for goal babbling that is scalable, fast, and applicable in real-world scenarios.

We demonstrate an online goal babbling algorithm [42] and show that the method does indeed permit enormous scalability (see Fig. 2). This can be achieved because goal-directed exploration allows to leave out redundant choices of actions as soon as there are known ways that solve the problem. The experiments point out that goal babbling constitutes a positive feedback loop during bootstrapping, in which exploration and learning reinforce each other. This positive feedback loop is identified as an important conceptual property of goal babbling that is in line with the *dynamic systems perspective* [43] on development. Experiments demonstrate that it allows to achieve human-level [44] learning speed. We finally demonstrate the practical use of the approach to learn the inverse kinematics of the *Bionic Handling Assistant* (see Fig. 3), which is a relevant and very challenging use case [45].

The results demonstrate the theoretical as well as practical validity of our algorithmic approach. It thereby provides a coherent and constructive explanation of how infants' early goal-directed movements might lead to their rapid initial mastery of feedforward-controlled reaching movements. Other implementations of goal babbling have recently been proposed, and confirm the success of goal babbling, as well as its superiority over motor babbling in terms of bootstrapping efficiency [46], [47], [48]. These results also demonstrate the general validity of the goal babbling concept, which provides a new framework to foster research on infant development as well as robotic systems.

References

[1] J. P. Welsh and R. Llins, "Some organizing principles for the control of movement based on olivocerebellar physiology," *Progress in Brain Research*, vol. 114, pp. 449–461, 1997.

[2] J. Konczak, M. Borutta, and J. Dichgans, "The development of goal-directed reaching in infants ii. learning to produce task-adequate patterns of joint torque," *Experimental Brain Research*, vol. 113, no. 3, pp. 465–474, 1997.

[3] S. Kopp and J. J. Steil, "Special corner on cognitive robotics," *Cognitive Processing*, vol. 12, no. 4, pp. 317–318, 2011.

[4] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini, "Developmental robotics: A survey," *Connection Science*, vol. 15, no. 4, pp. 151–190, 2003.

[5] 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*, vol. 1, no. 1, pp. 12–34, 2009.

[6] D. M. Wolpert, R. C. Miall, and M. Kawato, "Internal models in the cerebellum," *Trends in Cognitive Science*, vol. 2, no. 9, pp. 338–347, 1998.

[7] C. Castellini, F. Orabona, G. Metta, and G. Sandini, "Internal models of reaching and grasping," *Advanced Robotics*, vol. 21, no. 13, pp. 1545–1564, 2007.

[8] D. Bullock, S. Grossberg, and F. H. Guenther, "A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm," *Cognitive Neuroscience*, vol. 5, no. 4, pp. 408–435, 1993.

[9] Y. Demiris and A. Dearden, "From motor babbling to hierarchical learning by imitation: A robot developmental pathway," in *Int. Conf. Epigenetic Robotics (EpiRob)*, 2005.

[10] A. D'Souza, S. Vijayakumar, and S. Schaal, "Learning inverse kinematics," in *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2001.

[11] G. Sun and B. Scassellati, "A fast and efficient model for learning to reach," *International Journal of Humanoid Robotics (IJHR)*, vol. 2, no. 4, pp. 391–413, 2005.

[12] F. R. Reinhart and J. J. Steil, "Neural learning and dynamical selection of redundant solutions for inverse kinematic control," in *IEEE-RAS Int. Conf. Humanoid Robots*, 2011.

[13] M. Kawato, "Feedback-error-learning neural network for supervised motor learning," in *Advanced Neural Computers*, R. Eckmiller, Ed. Elsevier, 1990, pp. 365–372.

[14] M. I. Jordan and D. E. Rumelhart, "Forward models: supervised learning with a distal teacher," *Cognitive Science*, vol. 16, no. 3, 1992.

[15] J. Porrill, P. Dean, and J. V. Stone, "Recurrent cerebellar architecture solves the motor-error problem," *Proc Biol Sci*, vol. 271, no. 1541, pp. 789–796, 2004.

[16] M. Kuperstein, "Neural model of adaptive hand-eye coordination for single postures," *Science*, vol. 239, no. 4845, pp. 1308–1311, 1988.

[17] 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 *Int. Conf. Cognitive Systems (CogSys)*, 2008.

[18] J. Piaget, *The Origin of Intelligence in the Child*. Routledge and Kegan Paul, 1953.

[19] B. I. Bertenthal, "Origins and early development of perception, action, and representation," *Annual Reviews Psychology*, vol. 47, pp. 431–459, 1996.

[20] R. K. Clifton, B. A. Morrongiello, W. Kulig, and J. Dowd, "Developmental changes in auditory localization in infancy," in *Development of Perception, Vol. 1, Psychobiological Perspectives*, R. Aslin, J. Alberts, and M. Petersen, Eds. New York: Academic Press, 1981, pp. 141–160.

[21] H. Bloch and I. Carchon, "On the onset of eye-head coordination in infants," *Behavioral Brain Research*, vol. 49, no. 1, pp. 85–90, 1992.

[22] A. van der Meer, F. van der Weel, and D. Lee, "The functional significance of arm movements in neonates," *Science*, vol. 267, no. 5198, pp. 693–695, 1995.

[23] A. van der Meer, "Keeping the arm in the limelight: Advanced visual control of arm movements in neonates," *European Journal of Paediatric Neurology*, vol. 1, no. 4, pp. 103–108, 1997.

[24] C. von Hofsten, "Eye-hand coordination in the newborn," *Developmental Psychology*, vol. 18, no. 3, pp. 450–461, 1982.

[25] L. Rönnqvist and C. von Hofsten, "Neonatal finger and arm movements as determined by a social and an object context," *Early Development and Parenting*, vol. 3, no. 2, pp. 81–94, 1994.

[26] C. von Hofsten, "An action perspective on motor development," *Trends in Cognitive Science*, vol. 8, no. 6, pp. 266–272, 2004.

[27] R. K. Clifton, D. W. Muir, D. H. Ashmead, and M. G. Clarkson, "Is visually guided reaching in early infancy a myth?" *Child Development*, vol. 64, no. 4, pp. 1099–1110, 1993.

[28] L. Out, A. J. van Soest, G. P. Savelsbergh, and B. Hopkins, "The effect of posture on early reaching movements," *Journal of Motor Behavior*, vol. 30, no. 3, pp. 260–272, 1998.

[29] J. Baily, "Adaptation to prisms: Do proprioceptive changes mediate adapted behaviour with ballistic arm movements?" *The Quarterly Journal of Experimental Psychology*, vol. 24, no. 1, pp. 8–20, 1972.

[30] E. W. Bushnell, "The decline of visually guided reaching during infancy," *Infant Behavior and Development*, vol. 8, no. 2, pp. 139–155, 1985.

[31] J. Konczak, M. Borutta, H. Topka, and J. Dichgans, "The development of goal-directed reaching in infant: hand trajectory formation and joint torque control," *Experimental Brain Research*, vol. 106, no. 1, pp. 156–168, 1995.

[32] N. E. Berthier and R. Keen, "Development of reaching in infancy," *Experimental Brain Research*, vol. 169, no. 4, pp. 507–518, 2005.

[33] M. Rolf, J. J. Steil, and M. Gienger, "Goal babbling permits direct learning of inverse kinematics," *IEEE Trans. Autonomous Mental Development*, vol. 2, no. 3, 2010.

[34] E. Theodorou, J. Buchli, and S. Schaal, "Reinforcement learning of motor skills in high dimensions: A path integral approach," in *IEEE Int. Conf. Robotics and Automation (ICRA)*, 2010.

[35] A. Baranes and P.-Y. Oudeyer, "Maturationally-constrained competence-based intrinsically motivated learning," in *IEEE Int. Conf. Development and Learning (ICDL)*, 2010.

[36] W. Schenck, "Adaptive internal models for motor control and visual prediction," in *MPI Series in Biological Cybernetics*. Logos Verlag: Berlin, 2008, doctoral Thesis.

[37] C. G. Atkeson, "Learning arm kinematics and dynamics," *Annual Reviews Neuroscience*, vol. 12, pp. 157–183, 1989.

[38] E. Oyama and T. Tachi, "Goal-directed property of online direct inverse modeling," in *Int. Joint Conf. Neural Networks (IJCNN)*, 2000.

[39] T. D. Sanger, "Failure of motor learning for large initial errors," *Neural Computation*, vol. 16, no. 9, pp. 1873–1886, 2004.

[40] J. Peters and S. Schaal, "Reinforcement learning by reward-weighted regression for operational space control," in *Int. Conf. Machine Learning (ICML)*, 2007.

[41] M. Rolf and J. J. Steil, "Explorative learning of right inverse functions: theoretical implications of redundancy," in *New Challenges in Neural Computation*, 2012.

[42] M. Rolf, J. J. Steil, and M. Gienger, "Online goal babbling for rapid bootstrapping of inverse models in high dimensions," in *IEEE Int. Joint Conf. Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, 2011.

[43] L. B. Smith and E. Thelen, "Development as a dynamic system," *Trends in Cognitive Sciences*, vol. 7, no. 8, pp. 343–348, 2003.

[44] U. Sailer, J. R. Flanagan, and R. S. Johansson, "Eye–hand coordination during learning of a novel visuomotor task," *Journal of Neuroscience*, vol. 25, no. 39, pp. 8833–8842, 2005.

[45] M. Rolf and J. J. Steil, "Efficient exploratory learning of inverse kinematics on a bionic elephant trunk," 2012, submitted.

[46] L. Jamone, L. Natale, K. Hashimoto, G. Sandini, and A. Takanishi, "Learning task space control through goal directed exploration," in *IEEE Int. Conf. Robotics and Biomimetics (ROBIO)*, 2011.

[47] P. O. Stalph and M. V. Butz, "Learning local linear jacobians for flexible and adaptive robot arm control," *Genetic Programming and Evolvable Machines*, vol. 13, no. 2, pp. 137–157, 2012.

[48] C. Hartmann, J. Boedecker, O. Obst, S. Ikemoto, and M. Asada, "Real-time inverse dynamics learning for musculoskeletal robots based on echo state gaussian process regression," in *Robotics: Science and Systems (RSS)*, 2012.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Uterine environment guides organization of somatosensory area: a computational approach

Akimasa Nakashima* Yasunori Yamada*† and Yasuo Kuniyoshi*
*Grad. School of Info. Sci. and Tech., The Univ. of Tokyo, Japan
†Research fellow of the Japan Society for the Promotion of Science
Email: {nakashima, y-yamada, kuniyosh}@isi.imi.i.u-tokyo.ac.jp

I. Introduction

Converging developmental studies have emphasized the significance experience within the uterine environment from as early as the fetal period for motor and cognitive development [1]. Notably, these studies have emphasized the importance of sensory feedback due to spontaneous movements for early development. It is therefore important to reveal how the fetus' interaction with the uterine environment guides its development in order to deepen our understanding of the underlying mechanisms for development.

Among all sensory experiences within the uterine environment, the somatosensory modality plays a central role in early development. In fact, this modality starts functioning over the whole body from as early as the 17th gestational age, before other sensory modalities [2].

Several researchers have suggested the importance of sensory stimulation generated by spontaneous fetal movements for the formation of the body map in the primary somatosensory area (S1) [3]. However, there are few studies on the mechanisms of how the S1 map is generated and what components shape its organization.

In this paper, we argue that uterine environment contributes to the guidance of the formation of somatosensory representations. We investigated the relationship between the uterine environment and the organization of S1 shaped by sensory information gathered via interaction with the environment.

II. Materials and Methods

We ran computer simulations of human fetus models within and outside uterine environment. This fetus model have biologically plausible musculoskeletal bodies, a spinal neural network and a primary somatosensory area.

A. Body Model and Environment Model

We used human fetus models, which undergo 30 gestational weeks [4] [5] (Fig.1). The model had parameters based on actual fetus data such as size, mass, moment of inertia of each body part, joint angle limits, muscle configuration and force. The human fetus models had 198 muscles in the whole body excluding the finger and face muscles, and 1500 tactile sensor cells, whose distribution was based on human two point discrimination (Fig.1B, Table.I). To simulate tactile sensation, we used the Merkel cell model. Merkel cells are mechanoreceptors which mainly detect continuous pressure. The Merkel cell model used in this simulation detected continuous pressure by low-pass filtering the pressure input (< 50 Hz) [6].

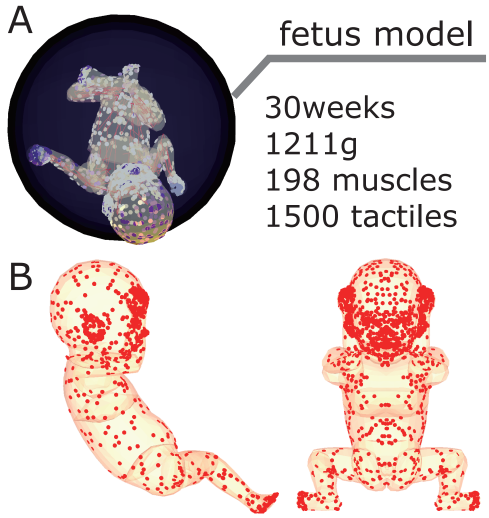

Fig. 1. Fetus model overview. (A) Fetus model appearance and fetus data. Blue circle represents uterus and white and red circles represent tactile sensors. Red one is responsing tactile. (B) Tactile distribution on the fetus model.

TABLE I
THE DISTRIBUTUION OF TACTILE SENSORS ON THE FETUS MODEL'S LEFT SIDE.

head	neck	shoulder	upper arm	lower arm
377	7	14	16	14

hand	chest	stomach	hip	thigh	calf	foot
132	34	48	22	24	15	47

Inside the uterus, pressure inputs to the fetus come from its embryonic and fetal environments. We used the amniotic fluid and uterine wall models produced by Mori and Kuniyoshi [4]. In our simulations, pressure inputs could be due to (1) physical contact, (2) the uterine wall, and (3) amniotic fluid resistance. Pressure due to physical contact between body parts was distributed according to the tactile sensors distance from the colliding body part. Pressure due to the uterine membrane depends on the sensor's distance from the center of the uterus and as well as its orientation. Pressure due to amniotic fluid resistance is calculated by taking the inner product of the velocity of the body part and directional unit vector of the tactile sensor. Outside the uterus, the fetus model was only subject to pressure due to physical contact between body parts and the ground.

B. Motion Generation Model

The neural basis for fetal spontaneous whole-body movements is believed to be Central Pattern Generators (CPGs), which are circuits mediating rhythmic behaviors such as walking and swimming in the spinal cord or brain stem [7]. We employed the spinobulbar model developed by Kuniyoshi and Sangawa [8], which includes a CPG model for generating various whole-body movements. This model receives muscle length and tension as sensory input, and outputs the degree of muscle activation as motor command.

C. Somatosensory Area Model

S1 has a somatotopic representation of the body, which largely presents the spatial organization of body parts [9]. Similar cortical representations are observed in other primary sensory areas such as the primary auditory cortex (A1) and the primary visual cortex (V1). Recently, Terashima and Okada suggested that A1 and V1 cortical representations can be explained by the common neural network model [10]. We applied the neural network model, Topographic Independent Component Analysis (TICA), to simulate the organization of the somatosensory map [11].

TICA takes the sensory inputs from tactile sensors and not only extracts the independent components using Independent Components Analysis (ICA), but also constructs a two-dimensional map in such a way that adjacent elements in the map have similar sensory representations. In other words, TICA is a variant of ICA in which the output is a sparse and topographically organized representation of the sensory inputs. To construct a two-dimensional map of m elements, an independent components vector $\boldsymbol{s}_t = [s_{1t}, \cdots, s_{jt}, \cdots, s_{mt}]^{\mathrm{T}}$ is calculated as

$$\boldsymbol{s}_t = \boldsymbol{W}\boldsymbol{x}_t, \tag{1}$$

where $\boldsymbol{x}_t$ is the vector of sensory inputs from tactile sensors, and $\boldsymbol{W}$ is the weight matrix. The weight matrix $\boldsymbol{W} = [\boldsymbol{w}_1, \cdots, \boldsymbol{w}_m]^{\mathrm{T}}$ is estimated using the gradient method, which maximizes the likelihood function L for the observed time series of tactile information $\boldsymbol{x}_t$. The likelihood function L is formulated as follows:

$$c_{it} = \sum_j h(i,j)\, s_{jt}^2, \tag{2}$$

$$\log L(\boldsymbol{x}_1, \cdots, \boldsymbol{x}_n;\, \boldsymbol{w}_1, \cdots, \boldsymbol{w}_m) = \sum_{t=1}^{n} \sum_{i=1}^{m} G(c_{it}), \tag{3}$$

where $h(i,j)$ is binary filter function for selecting the elements that neighbor i-th components on topography. This filter makes sure that adjacent elements in the final map have similar weight vectors, allowing the map to have a topographical organization. $G(c_{it})$ denotes the probability density function of c_{it}, which we defined as:

$$G(c_{it}) = \log p(c_{it}) = -\sqrt{0.005 + c_{it}}. \tag{4}$$

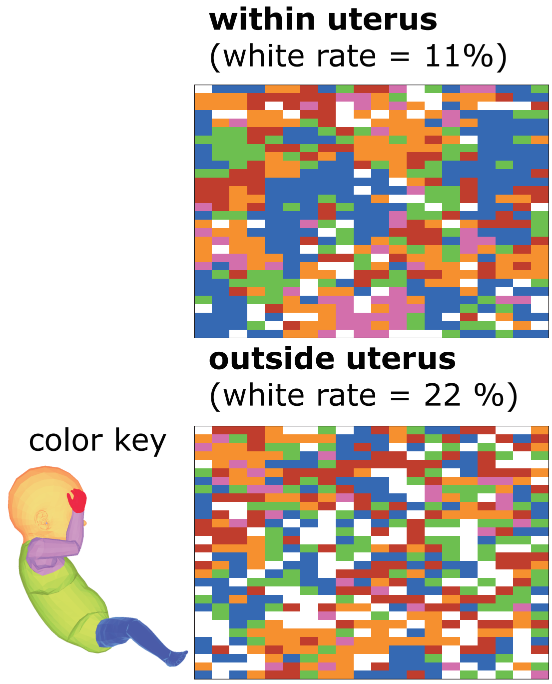

Fig. 2. Learned S1 maps. Colors represent each body parts, and white color represents somatosensory components which could not be categorized into any specific body part.

By defining the probability density function in the above fashion allowed the resulting map to be sparse. In this experiment, the dimensions of the resulting two-dimensional topographical map 30 $\times$ 20 elements ($m = 600$). The map had a torus configuration (opposite edes were connected) to avoid border effects.

III. Experiments

In order to investigate relationship between the uterine environment and organization of the S1 model, we conducted fetus simulations within and outside the uterus, and then built S1 maps as defined by tactile sensory information. Therefore, we set the time step of the simulation to 1 ms, and ran the simulation for 500 s. As for tactile sensors, we used the left-side of the body. We analyzed (1) whether each component in the S1 model represent specific five body parts: head, arm, hand, torso, leg and (2) whether the S1 map is organized so that adjacent components represent the neighboring body part.

First, we determined which body part was represented by each tactile component in S1 (Fig.2). If more than half of the strongest inputs to a given tactile sensor came from one specific body part, it was categorized as being dominantly represented by that body part. We calculated the percentage of components which could not be categorized into any specific body parts ("white rate" in Fig. 2). The percentages were 11% and 22% within and outside the uterus, respectively. Figure 3 shows the array of tactile sensors contributing to the body parts represented in S1. We confirmed that these sensors tended to be spatially localized to their respective body parts.

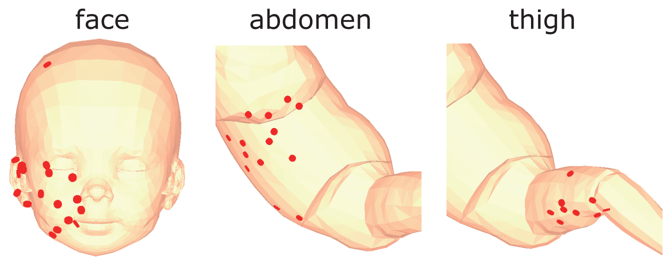

Fig. 3. Examples of cortical reprsentation in S1. Red circle is a tactile cell, which strongly inputs one component in S1 map.

Second, to evaluate the degree of topography in S1, we investigated the degree of clustering in S1. The number of tactile components which had neighboring components also categorized into the same body part were summed. Results showed a significant increase in the number of clustered components in S1 maps created within rather than outside the uterus. The results showed that such area within uterine environment significantly increased compared with those outside uterus (Mann-Whitney test, $p < 0.005$).

IV. Conclusion

Animals are dynamically coupled to their environments, with environment shaping the structure of sensory input, and sensory information determining neural dynamics. In this paper, we argue that interaction structured by the environment can guide the formation of somatosensory representations in human fetuses. To test our hypothesis, we conducted computer simulations using fetus model and compared the organization of such representations within and outside uterine environment. We found that S1 within the uterus had two times the number of localized body representations than outside the uterus. Furthermore, the fetus within the uterus is significantly larger than outside the uterus in terms of somatotopic organization. Our results suggest that uterine environment possesses rich regularities that structure tactile information and guide the organization of the S1 body map.

Acknowledgment

The work reported in this paper has been supported by MEXT Grant-in-Aid for Scientific Research on Innovative Areas "Constructive Developmental Science" (24119002).

References

[1] M. H. Johnson, *Developmental Cognitive Neuroscience*. Blackwell Publishing, 2005.

[2] R. M. Bradley and C. M. Mistretta, "Fetal sensory receptors," *Physiological Reviews*, vol. 55, no. 3, pp. 352–382, 1975.

[3] M. Milh, A. Kaminska, C. Huon, A. Lapillonne, Y. Ben-Ari, and R. Khazipov, "Rapid cortical oscillations and early motor activity in premature human neonate," *Cerebral Cortex*, vol. 17, no. 7, pp. 1582–1594, 2007.

[4] H. Mori and Y. Kuniyoshi, "A human fetus development simulation: Self-organization of behaviors through tactile sensation," in *IEEE 9th International Conference on Development and Learning*, 2010, pp. 82–87.

[5] Y. Yamada and Y. Kuniyoshi, "Emergent spontaneous movements based on embodiment: Toward a general principle for early development," in *Proceedings of the Post-Graduate Conference on Robotics and Development of Cognition*, J. Szufnarowska, Ed., September 2012, pp. 52–55.

[6] A. W. Freeman and K. O. Johnson, "A model accounting for effects of vibratory amplitude on responses of cutaneous mechanoreceptors in macaque monkey," *The Journal of Physiology*, vol. 323, no. 1, pp. 43–64, 1982.

[7] M. Hadders-Algra, "Putative neural substrate of normal and abnormal general movements," *Neuroscience & Biobehavioral Reviews*, vol. 31, no. 8, pp. 1181–1190, 2007.

[8] Y. Kuniyoshi and S. Sangawa, "Early motor development from partially ordered neural-body dynamics: experiments with a cortico-spinal-musculo-sleletal model," *Biological Cybernetics*, vol. 95, pp. 589–605, 2006.

[9] W. Penfield and E. Boldrey, "Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation," *Brain*, vol. 60, pp. 389–443, 1937.

[10] H. Terashima and M. Okada, "The topographic unsupervised learning of natural sounds in the auditory cortex," *Neural Information Processing Systems*, 2012 (to appear).

[11] A. Hyvärinen, P. Hoyer, and M. Inki, "Topographic independent component analysis," *Neural computation*, vol. 13, no. 7, pp. 1527–1558, 2001.

Ugur, E., Nagai, Y., Oztop, E., and Asada, M. (Eds) Proceedings of Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities? November 29, 2012. Osaka, Japan

Central tendency in space perception varies in development and is altered in human-robot interactive tasks.

Alessandra Sciutti, Andrea Del Prete, Lorenzo Natale, Giulio Sandini, Monica Gori
Robotics, Brain and Cognitive Sciences Dept.
Istituto Italiano di Tecnologia
Genoa, Italy

David Burr
Dipartimento di Psicologia
Università degli Studi di Firenze
Florence, Italy

I. Introduction

A fundamental principle in time perception is the so-called central tendency: reproductions of time intervals regress to the mean value of the previous stimuli distribution [1, 3-4]. Therefore, the estimate of sample duration differs depending on the distribution from which it is drawn, i.e. its statistical context. Central tendency is not just the consequence of a decrease of attention in a repetitive task, but rather optimizes temporal reproduction by minimizing the total error, which comprises both the accuracy and the variance of the responses [1, 4]. This strategy compensates for low sensory resolution by sacrificing veridicality, as it takes into account the statistics of previous stimuli rather than just the current stimulation. To explain the benefit of applying this strategy, we can propose an analogy with an everyday judgment. Consider the task of estimating the real dimension of an object (e.g. a car) by looking at its picture. If we were to evaluate the object size only on the basis of the current visual information, a car would be misperceived as a few cm long. On the contrary, usually our estimation is based on an internalized average measure, derived from a statistics of all the cars that we have seen in the past. The ability to take into account environmental statistics during perception could be advantageous also on a robotic platform. However, from a practical point of view, it is important to evaluate whether this optimization approach could be always beneficial and, if not, in what circumstances would it be more advantageous. In particular, it is interesting to investigate how this mechanism develops during childhood. Recent Bayesian models of the phenomenon in human adults [1, 4] have shown that the central-tendency strategy is beneficial only when perceptual judgment is imprecise and that the entity of the regression to the mean depends strongly on the precision with which such judgment can be made.

In this study we evaluated whether other factors can influence the feasibility of the central tendency in perception, to establish when it would be appropriate to endow the robot with this optimization mechanism. In particular we considered the relevance of three factors in determining the adoption of the central tendency in humans:

The quantity to be judged, moving from time to space perception. The idea of studying space rather than time perception derives from a hypothesis put forward by Hollingworth already in 1910 [3]. According to his view, the perceptual principle of central tendency should apply also to other sensory judgments rather than time, although recent studies have been focused only on the latter. Therefore, in this work we evaluated whether central tendency generalizes also to space perception.

The developmental phase, as different strategies could be beneficial at different phases of the development. Although the central tendency mechanism is a gold standard in perceptual judgments in adults, how this strategy develops with age is still unknown. On the one hand, in children sensory precision is usually lower than in adults. Therefore, a strategy aimed at minimizing the variance of sensory judgments could be particularly beneficial at younger ages. On the other hand, for the developing brain it could be fundamental to formulate estimations as veridical as possible, at the expenses of production variance, so that – by trial and the feedback of the error – children could develop the ability to produce accurate judgments. Hence, a Bayesian model aimed at minimizing the total error by reducing veridicality could be adopted only later in the development. This is what happens for instance in the development of sensory fusion: before adopting the Bayesian

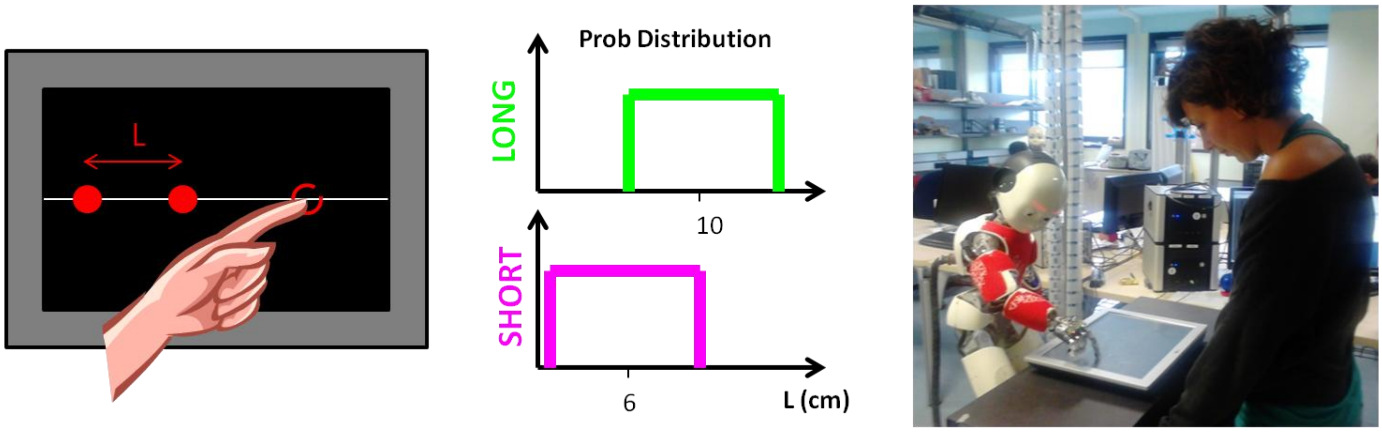

Fig. 1. Schematic representation of the task and of the stimuli distributions in the LONG and SHORT conditions. On the right, picture of the interactive task.

optimization of a multimodal judgment, the brain needs to calibrate the single modality estimation by using the more reliable sense [2]. Only when the calibration has occurred, the adult-like multimodal Bayesian integration occurs. In this study we evaluated the developmental trend of the central tendency mechanism in children between 7 and 14 years of age.

Interactive vs. not interactive context, as a judgment finalized to an interaction might need to obey to different constraints than a judgment per se. Even in adults, the central tendency mechanism could sometimes be not beneficial. In particular, when we move from perceptual tasks to interactive tasks, accuracy (or veridicality) could acquire a higher relevance than robustness to perceptual noise. Indeed, an inaccurate evaluation of the amplitude of the arm movement of another agent passing an object could imply a failure of the cooperative task. Hence, it could sometimes be inappropriate to sacrifice accuracy for minimizing the total error. We performed an interactive spatial task with the humanoid robot iCub as a co-actor to assess whether central tendency is normally adopted by adults in interactive scenarios too.

The results of this study can give insights on whether it is really relevant (and under which conditions it is appropriate) to implement the central tendency optimization mechanism in a robotic device.

II. METHODS

In this study we tested the central tendency strategy for space estimation in adults, in both a perceptual task (6 subjects) and an interactive task (7 subjects). Moreover, we tested children ranging from 7 to 14 years of age (a total of 77 subjects divided into five age groups) in the first task, to evaluate the development of the phenomenon. In the developmental study we used such a larger number of subjects, (in traditional studies with adults the sample is of 6 subjects [1, 4]) because children data are usually characterized by a high variability.

In the perceptual task, on each trial, subjects were presented with two subsequent flashes of light (red disks of 1 cm of diameter, each flash lasting 400 ms) positioned along a visible straight white line crossing the whole screen at its middle height. The first flash was located at a variable distance from the left border of the screen (0.5-3.5 cm, randomly selected). On its disappearance, a second disk appeared at a variable distance from it. Subjects were requested to touch a point on the straight line in order to reproduce, with respect to the second disk, the distance between the first and the second disk. After the touch, a red disk appeared to indicate where the subject had pressed the screen. No feedback was provided. Each new trial started after the experimenter’s button press, with the first light appearing after 500 ms. Each subject participated in two sessions: a SHORT condition, in which the spatial distance between the first two disks ranged from 2 cm to 10 cm, and a LONG condition, in which the presented distances ranged from 6 cm to 14 cm (see Fig.1). To avoid interference between the two contexts the two conditions were tested in two different days. The order of the sessions was randomized between subjects. Each condition was characterized by 11 different sample intervals (separated by 0.8 cm each), each of which was presented 8 times, yielding to a total of 88 trials per subject per condition.

In the interaction task the stimuli were similar, but the task was presented as a collaborative game. The humanoid robot iCub pressed the touch screen in two different points in sequence and the subjects had to complete its action by

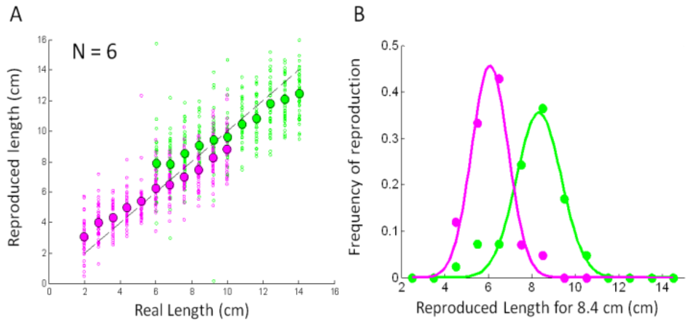

Fig. 2. A) Reproduced lengths as a function of stimulus length for two stimulus ranges (short – magenta, long – green). B) Reproduction distribution of visual stimuli for the stimulus 8.5 cm during sessions where stimuli were drawn from the short or the long contexts (same color coding as Fig. 2A).

touching a third point to reproduce the distance shown by the robot. iCub performed a human-like approximate minimum-jerk movement (6) and exhibited a naturalistic gaze behavior (5), with the eyes looking towards the pointing target before the movement completion and then fixating the human subject when his/her turn started. The stimuli presentation was slower than in the perceptual condition, with the mean robot velocity being constant across trials (average hand speed of about 0.1 m/s). The use of the robotic platform iCub guaranteed a complete control over the statistics and the timing of the stimuli presentation, to allow for a comparison with the computerized perceptual task. Before each experimental session, a calibration was performed to register the frames of reference of the robot and the touch screen. The distances presented by the robot were not sampled in 11 intervals but were drawn from the same uniform distributions that we used in the perceptual task, as in (4).

III. RESULTS

A. *Central tendency in the perception of space*

The results in Fig. 2 show that the phenomenon of central tendency is present in adults also for space perception. The average reproduced lengths (larger dots in Fig. 2A) do not correspond to the real stimulus amplitude (the data would lie on the identity line in such case), but tend to the mean value of the corresponding stimulus distribution (6 cm and 10 cm for the magenta and green data points, respectively), lying on flatter lines. The central tendency can be quantified by the regression index, i.e. the difference in slope between the best linear fit of the reproduced lengths and the identity line. This index varies from 0 (veridical performance) to 1 (complete regression to the mean) and was on average 0.35 ± 0.08 (Standard Error of the Mean, SEM). The central tendency is even more clearly depicted in Fig. 2B, where the distribution for the reproductions of the stimulus of 8.5 cm depends strongly on the sample range from which it was drawn, tending towards a shorter mean length for stimuli in the short range (magenta) and a longer mean for stimuli in the long range (green).

B. *The development of the central tendency mechanism*

As depicted in Figure 3A, during development a clear change is observed in the reproduction of spatial stimuli. Interestingly a developmental trend seems to be active until 13 years of age, with a progressive increase in children's accuracy shown by the decrease of the regression index. Indeed, the regression index decreases substantially with age, reaching adult-like values around 11-13 years of age (see Fig. 3B).

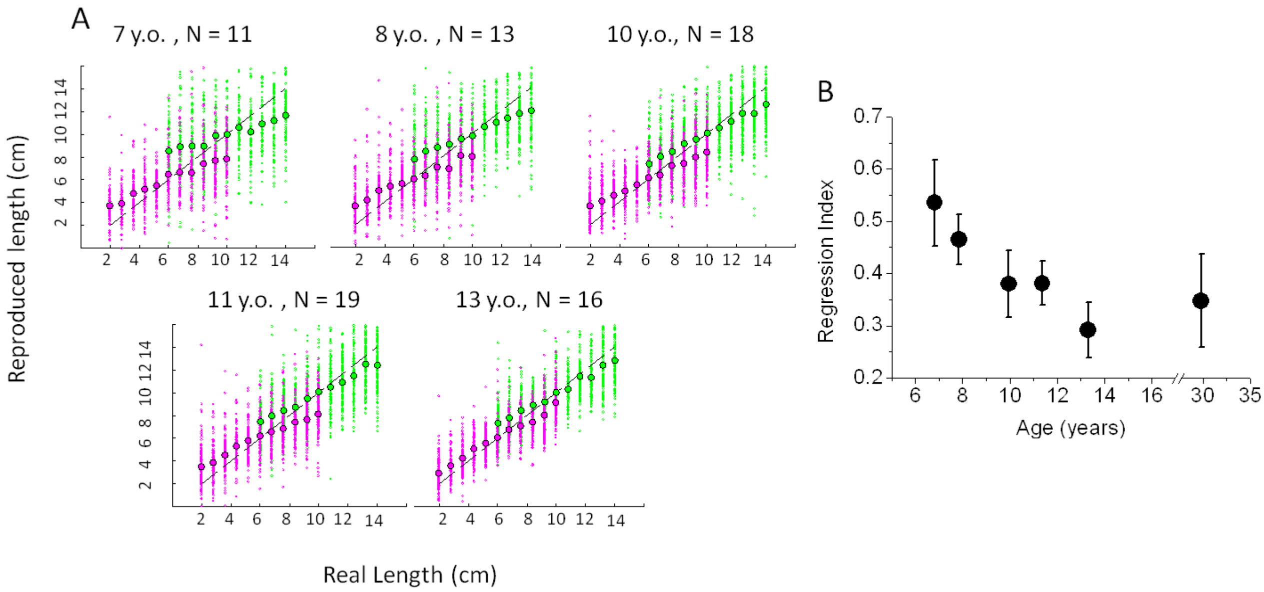

Fig. 3. A) Reproduced length as a function of stimulus length for the different age groups. Same color code as Fig. 2. B) Average regression index as a function of age. Error bars represent group standard errors

This trend follows the improvement in visual precision in size perception (see e.g. [2]), as the variability associated to perceptual estimates of size decreases progressively during childhood, at least until 10 years of age. Moreover, the tendency seems to indicate that during late childhood kids adopt even lower regression values than adults, as if giving more relevance to veridicality than adults

C. *Accuracy wins over central tendency in interactive scenarios*

In adults when space reproduction is inserted in an interactive framework, the central tendency almost disappears (see Fig. 4A). Indeed, subjects on average reproduce as accurately as possible (not considering an individual constant bias) the distance presented by the other agent. In fact, on average the regression index decreases significantly in the interactive condition with respect to the space reproduction task performed in solo (one-sided, two-sample t-test, $t(6.95)= 2.55324$, $p= 0.019$, see Fig. 4B).

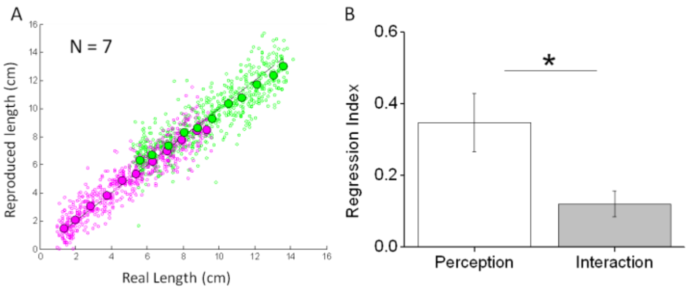

Fig. 4. A) Reproduced lengths as a function of stimulus length in the interactive task. Same color code as Fig. 2. B) Average regression index as a function of task type. Error bars represent standard errors. The star indicates significant difference in a t-test ($p < 0.05$.)

D. *General results*

Our results show that in adults the central tendency is present not only for time, but also for space perception. Younger children show a stronger regression to the stimulus mean. Adult-like regression level is attained only around 11-13 years of age, following the reduction in perceptual variance associated with development. A tendency for preferring accuracy over regression to the mean is observed at the later developmental stage, but it still needs to be investigated in more detail. Most interestingly, the central tendency is almost abandoned in interactive tasks, where the accuracy of the reproduction is maximized with almost no regression to the stimulus average.

IV. DISCUSSION

The central tendency mechanism is a fundamental principle of optimization adopted in human perception, not dissimilar to a filtering approach that takes into consideration a (moving) window of measures to obtain a better estimation of a certain quantity. An important question is under what circumstances its implementation could become advantageous also in robotics. Our study confirms that sacrificing accuracy for noise-robustness can be beneficial for perception of various quantities (e.g., time and space) when sensory precision is low, as during development. However, in interactive scenarios accuracy is preferred over noise-robustness, suggesting that the adoption of central tendency is task-dependent. These findings indicate that the implementation of the central tendency mechanism should not be ubiquitous, but would need to be considered as a function of the task and of the variance of the robot sensory inputs.

REFERENCES

[1] Cicchini, G. M.; Arrighi, R.; Cecchetti, L.; Giusti, M. & Burr, D. C., “Optimal encoding of interval timing in expert percussionists.”, J Neurosci 32(3),2012, pp. 1056 –1060.

[2] Gori, M.; Viva, M. D.; Sandini, G. & Burr, D. C., “Young children do not integrate visual and haptic form information”, Curr Biol 18(9), 2008, pp. 694–698.

[3] Hollingworth, H. L, “'The central tendency of judgements”, J Philos Psychol Sci Methods 7, 1910, pp. 461 – 469.

[4] Jazayeri, M. & Shadlen, M. N., “Temporal context calibrates interval timing. ”, Nat Neurosci 13(8), 2010, pp. 1020–1026.

[5] Pattacini, U. , “Modular cartesian controllers for humanoid robots: Design and implementation on the iCub”, Master's thesis, Ph. D. dissertation, RBCS, Istituto Italiano di Tecnologia, Genoa, 2010

[6] Pattacini, U.; Nori, F.; Natale, L.; Metta, G. & Sandini, G., “An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots”, IEEE/RSJ International Conference on Intelligent Robots and Systems', 2010.