AISB'05: Social Intelligence and Interaction in Animals, Robots and Agents

Proceedings of the Third International Symposium on Imitation in Animals and Artifacts

Î

12 - 15 April 2005

University of Hertfordshire, Hatfield, UK

SSAISB 2005 Convention





AISB'05 Convention

Social Intelligence and Interaction in Animals, Robots and Agents

12-15 April 2005 University of Hertfordshire, Hatfield, UK

Proceedings of the

Third International Symposium on Imitation in Animals and Artifacts

Published by



The Society for the Study of Artificial Intelligence and the Simulation of Behaviour www.aisb.org.uk

Printed by



The University of Hertfordshire, Hatfield, AL10 9AB UK www.herts.ac.uk

Cover Design by Sue Attwood

ISBN 1 902956 42 5

AISB'05 Hosted by



The Adaptive Systems Research Group adapsys.feis.herts.ac.uk

The AISB'05 Convention is partially supported by:



EPSRC Engineering and Physical Sciences Research Council

The proceedings of the ten symposia in the AISB'05 Convention are available from SSAISB:

Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05) 1 902956 40 9

Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action 1 902956 41 7

Third International Symposium on Imitation in Animals and Artifacts 1 902956 42 5

Robotics, Mechatronics and Animatronics in the Creative and Entertainment Industries and Arts 1 902956 43 3

Robot Companions: Hard Problems and Open Challenges in Robot-Human Interaction 1 902956 44 1

Conversational Informatics for Supporting Social Intelligence and Interaction - Situational and Environmental Information Enforcing Involvement in Conversation 1 902956 45 X

Next Generation approaches to Machine Consciousness: Imagination, Development, Intersubjectivity, and Embodiment 1 902956 46 8

Normative Multi-Agent Systems 1 902956 47 6

Socially Inspired Computing Joint Symposium (Memetic theory in artificial systems & societies, Emerging Artificial Societies, and Engineering with Social Metaphors) 1 902956 48 4

Virtual Social Agents Joint Symposium (Social presence cues for virtual humanoids, Empathic Interaction with Synthetic Characters, Mind-minding Agents) 1 902956 49 2

Table of Contents

The AISB'05 Convention - Social Intelligence and Interaction in Animals, Robots and Agents i <i>K. Dautenhahn</i>
Symposium Preface - Third International Symposium on Imitation in Animals and Artifacts iv <i>Yiannis Demiris, Kerstin Dautenhahn, Chrystopher Nehaniv</i>
Robot Imitation from Human Body Movements
Achieving Corresponding Effects on Multiple Robotic Platforms: Imitating in Context Using Different Effect Metrics
Acquisition of a "mirror" system for speechreading
Using Visual Velocity Detection to Achieve Synchronization in Imitation
Learning Discretely: Behaviour and Organisation in Social Learning
Motor resonance and MOSAIC
Interpersonal Maps and the Body Correspondence Problem
Teaching a Robot to Behave Like a Cockroach
Goal level imitation with emergent repertoires
Hierarchies of Coupled Inverse and Forward Models for Abstraction in Robot Action Planning, Recognition and Imitation
A Developmental Roadmap for Task Learning by Imitation in Humanoid Robots
Joint Attention Development in Infant-like Robot Based on Head Movement Imitation
Grey Parrots Do Not Always "Parrot": Roles of Imitation and Phonological Awareness in the Creation of New Labels from Existing Vocalisations
Human and Robotic Action Observation Elicit Automatic Imitation
An Examination of the Static to Dynamic Imitation Spectrum

View Sensitive Cells as a Neural Basis for the Representation of Others in a Self-Centered	
Frame of Reference	9
Eric L. Sauser and Aude G. Billard	

The AISB'05 Convention Social Intelligence and Interaction in Animals, Robots and Agents

Above all, the human animal is social. For an artificially intelligent system, how could it be otherwise?

We stated in our Call for Participation "The AISB'05 convention with the theme *Social Intelligence and Interaction in Animals, Robots and Agents* aims to facilitate the synthesis of new ideas, encourage new insights as well as novel applications, mediate new collaborations, and provide a context for lively and stimulating discussions in this exciting, truly interdisciplinary, and quickly growing research area that touches upon many deep issues regarding the nature of intelligence in human and other animals, and its potential application to robots and other artefacts".

Why is the theme of Social Intelligence and Interaction interesting to an Artificial Intelligence and Robotics community? We know that intelligence in humans and other animals has many facets and is expressed in a variety of ways in how the individual in its lifetime - or a population on an evolutionary timescale - deals with, adapts to, and co-evolves with the environment. Traditionally, social or emotional intelligence have been considered different from a more problem-solving, often called "rational", oriented view of human intelligence. However, more and more evidence from a variety of different research fields highlights the important role of social, emotional intelligence and interaction across all facets of intelligence in humans.

The Convention theme *Social Intelligence and Interaction in Animals, Robots and Agents* reflects a current trend towards increasingly interdisciplinary approaches that are pushing the boundaries of traditional science and are necessary in order to answer deep questions regarding the social nature of intelligence in humans and other animals, as well as to address the challenge of synthesizing computational agents or robotic artifacts that show aspects of biological social intelligence. Exciting new developments are emerging from collaborations among computer scientists, roboticists, psychologists, sociologists, cognitive scientists, primatologists, ethologists and researchers from other disciplines, e.g. leading to increasingly sophisticated simulation models of socially intelligent agents, or to a new generation of robots that are able to learn from and socially interact with each other or with people. Such interdisciplinary work advances our understanding of social intelligence in nature, and leads to new theories, models, architectures and designs in the domain of Artificial Intelligence and other sciences of the artificial.

New advancements in computer and robotic technology facilitate the emergence of multi-modal "natural" interfaces between computers or robots and people, including embodied conversational agents or robotic pets/assistants/companions that we are increasingly sharing our home and work space with. People tend to create certain relationships with such socially intelligent artifacts, and are even willing to accept them as helpers in healthcare, therapy or rehabilitation. Thus, socially intelligent artifacts are becoming part of our lives, including many desirable as well as possibly undesirable effects, and Artificial Intelligence and Cognitive Science research can play an important role in addressing many of the huge scientific challenges involved. Keeping an open mind towards other disciplines, embracing work from a variety of disciplines studying humans as well as non-human animals, might help us to create artifacts that might not only do their job, but that do their job right.

Thus, the convention hopes to provide a home for state-of-the-art research as well as a discussion forum for innovative ideas and approaches, pushing the frontiers of what is possible and/or desirable in this exciting, growing area.

The feedback to the initial Call for Symposia Proposals was overwhelming. Ten symposia were accepted (ranging from one-day to three-day events), organized by UK, European as well as international experts in the field of Social Intelligence and Interaction.

- Second International Symposium on the Emergence and Evolution of Linguistic Communication (EELC'05)
- Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action
- Third International Symposium on Imitation in Animals and Artifacts
- Robotics, Mechatronics and Animatronics in the Creative and Entertainment Industries and Arts
- Robot Companions: Hard Problems and Open Challenges in Robot-Human Interaction
- Conversational Informatics for Supporting Social Intelligence and Interaction Situational and Environmental Information Enforcing Involvement in Conversation
- Next Generation Approaches to Machine Consciousness: Imagination, Development, Intersubjectivity, and Embodiment
- Normative Multi-Agent Systems
- Socially Inspired Computing Joint Symposium (consisting of three themes: Memetic Theory in Artificial Systems & Societies, Emerging Artificial Societies, and Engineering with Social Metaphors)
- Virtual Social Agents Joint Symposium (consisting of three themes: Social Presence Cues for Virtual Humanoids, Empathic Interaction with Synthetic Characters, Mind-minding Agents)

I would like to thank the symposium organizers for their efforts in helping to put together an excellent scientific programme.

In order to complement the programme, five speakers known for pioneering work relevant to the convention theme accepted invitations to present plenary lectures at the convention: Prof. Nigel Gilbert (University of Surrey, UK), Prof. Hiroshi Ishiguro (Osaka University, Japan), Dr. Alison Jolly (University of Sussex, UK), Prof. Luc Steels (VUB, Belgium and Sony, France), and Prof. Jacqueline Nadel (National Centre of Scientific Research, France).

A number of people and groups helped to make this convention possible. First, I would like to thank SSAISB for the opportunity to host the convention under the special theme of Social Intelligence and Interaction in Animals, Robots and Agents. The AISB'05 convention is supported in part by a UK EPSRC grant to Prof. Kerstin Dautenhahn and Prof. C. L. Nehaniv. Further support was provided by Prof. Jill Hewitt and the School of Computer Science, as well as the Adaptive Systems Research Group at University of Hertfordshire. I would like to thank the Convention's Vice Chair Prof. Chrystopher L. Nehaniv for his invaluable continuous support during the planning and organization of the convention. Many thanks to the local organizing committee including Dr. René te Boekhorst, Dr. Lola Cañamero and Dr. Daniel Polani. I would like to single out two people who took over major roles in the local organization: Firstly, Johanna Hunt, Research Assistant in the School of Computer Science, who efficiently dealt primarily with the registration process, the AISB'05 website, and the coordination of ten proceedings. The number of convention registrants as well as different symposia by far exceeded our expectations and made this a major effort. Secondly, Bob Guscott, Research Administrator in the Adaptive Systems Research Group, competently and with great enthusiasm dealt with arrangements ranging from room bookings, catering, the organization of the banquet, and many other important elements in the convention. Thanks to Sue Attwood for the beautiful frontcover design. Also, a number of student helpers supported the convention. A great team made this convention possible!

I wish all participants of the AISB'05 convention an enjoyable and very productive time. On returning home, I hope you will take with you some new ideas or inspirations regarding our common goal of understanding social intelligence, and synthesizing artificially intelligent robots and agents. Progress in the field depends on scientific exchange, dialogue and critical evaluations by our peers and the research community, including senior members as well as students who bring in fresh viewpoints. For social animals such as humans, the construction of scientific knowledge can't be otherwise.



Dedication:

I am very confident that the future will bring us increasingly many instances of socially intelligent agents. I am similarly confident that we will see more and more socially intelligent robots sharing our lives. However, I would like to dedicate this convention to those people who fight for the survival of socially intelligent animals and their fellow creatures. What would 'life as it could be' be without 'life as we know it'?

Beppu, Japan.

Kerstin Dautenhahn

Professor of Artificial Intelligence, General Chair, AISB'05 Convention Social Intelligence and Interaction in Animals, Robots and Agents

University of Hertfordshire College Lane Hatfield, Herts, AL10 9AB United Kingdom

Symposium Preface Third International Symposium on Imitation in Animals and Artifacts

SYMPOSIUM OVERVIEW

We are pleased to introduce the proceedings of the Third International Symposium on Imitation in Animals and Artifacts held at the University of Hertfordshire, UK, during 12-14 April 2005.

The first Imitation Symposium was held from 7-9 April 1999 as part of the AISB'99 Convention at the Edinburgh College of Arts & Division of Informatics, University of Edinburgh, Scotland. The second symposium in the series ran from 7 - 11 April 2003 at University of Wales, Aberystwyth, United Kingdom as part of the AISB 2003 convention. The aim of the imitation symposium series remains constant in trying to bring together researchers from robotics, computer science, psychology, animal ethology, neuroscience, brain imaging, pathology, and other areas to present and exchange ideas addressing the important issue of imitation. Species from rats to birds to humans have been observed to turn to their peers for efficient learning of useful knowledge, and imitation plays a prominent role in this learning process. However, explaining the mechanisms underlying the imitative abilities of humans and other animals has proved to be a complex and challenging subject. The mechanisms are not well-understood, and their connections to sociality, communication, development, and learning are deep, as research from various disciplines has started to reveal.

The proceedings contain 17 papers, ranging from computational implementations of imitation mechanisms and their robotic applications, to imitation experiments with parrots and humans. This variety of approaches, indicative of the interdisciplinary interest to imitation, promises to stimulate exciting discussions during the 3 days of the symposium.

We would like to thank the authors for submitting high-quality research papers for consideration to this symposium, and we would like to thank the following members of the program committee, who worked hard to provide us with high quality reviews:

Andrew Meltzoff, University of Washington, USA Aris Alissandrakis, University of Hertfordshire, UK Aude Billard, EPFL, Switzerland Auke Jan Ijspeert, EPFL, Switzerland Cecilia Heyes, UCL, UK Chrystopher Nehaniv, Adaptive Systems Research Group, Hertfordshire, UK Erhan Oztop, ATR, Japan Geoffrey Bird, UCL, UK Joanna Bryson, University of Bath, UK Gillian Hayes, IPAB, University of Edinburgh, UK Giorgio Metta, LIRA, University of Genoa, Italy Giulio Sandini, LIRA, University of Genoa, Italy Gordon Cheng, ATR, Japan Harold Bekkering, University of Nijmegen, Netherlands Hideki Kozima, CRL, Japan Jose Santos-Victor, ISR, Technical University of Lisbon, Portugal Kerstin Dautenhahn, Adaptive Systems Research Group, Hertfordshire, UK Martin Giese, University Clinic Tubingen, Germany Meredith Gattis, University of Cardiff, UK Minoru Asada, Osaka University, Japan Philippe Gaussier, ENSEA, France Robert Mitchell, Eastern Kentucky University, USA

Sethu Vijayakumar, IPAB, Edinburgh, UK.
Stefan Wermter, University of Sunderland, UK
Irene Pepperberg, Radcliffe Institute, Harvard, USA
Jacqueline Nadel, CNRS, France
Luc Berthouze, Neuroscience Institute, AIST Japan
Luciano Fadiga, Faculty of Medicine, University of Ferrara, Italy
Marco Iacoboni, UCLA, USA
Rajesh Rao, University of Washington, USA
Vittorio Gallese, Universita di Parma, Italy.
Wolfgang Prinz, Max Planck Institute for Psychological Research, Germany
Yiannis Demiris, BioART, EEE, Imperial College London, UK

We hope you will enjoy the symposium; we hope its discussions will stimulate you with new ideas, and increase the level of interdisciplinary collaboration in tackling the fascinating issue of imitation.

Yiannis Demiris, Kerstin Dautenhahn, Chrystopher Nehaniv January 2005

SYMPOSIUM CHAIR

Dr. Yiannis Demiris BioART, Intelligent Systems and Networks Group Department of Electrical and Electronic Engineering, Imperial College London, SW7 2BT, London, UK http://www.iis.ee.ic.ac.uk/yiannis

PROGRAMME CO-CHAIRS

Professor Kerstin Dautenhahn Adaptive Systems Research Group University of Hertfordshire, School of Computer Science College Lane, Hatfield, Hertfordshire AL10 9AB, United Kingdom URL: http://homepages.feis.herts.ac.uk/~comqkd

Professor Chrystopher L. Nehaniv Adaptive Systems & Algorithms Research Groups University of Hertfordshire, School of Computer Science College Lane, Hatfield, Hertfordshire AL10 9AB, United Kingdom URL: http://homepages.feis.herts.ac.uk/~nehaniv/welcome.html

Robot Imitation from Human Body Movements

Carlos A. Acosta Calderon* and Huosheng Hu*

*Department of Computer Science, University of Essex Wivenhoe Park, Colchester CO4 3SQ, United Kingdom caacos@essex.ac.uk, hhu@essex.ac.uk

Abstract

Imitation represents a useful and promising alternative to programming robots. The approach presented here is based on two functional elements used by humans to understand and perform actions. These elements are: the body schema and the body percept. The first one is a representation of the body containing information of the body's capabilities. The body percept is a snapshot of the body and its relation with the environment at a given instant. These elements are believed to interact between each other generating among other abilities, the ability to imitate. This paper presents our approach to robot imitation and experimental results, where a robot is able to imitate the movements of a human demonstrator via its visual observations.

1 Introduction

Today, many robotics applications are being investigated, including space exploration, hazardous environments, service robotics, cleaning, transportation, emergency handling, house building, elderly assistance, and so forth (Liu and Wu, 2001; Fong et al., 2002). These novel applications involve interaction between humans and robots where the robots must coordinate their efforts with their human owners. Therefore, robots must be able to recognize the observable actions of the other teammates in order to understand the goals of the actions (Breazeal et al., submitted for publication). In addition, some robots must also learn new observable actions in order to be able to exchange roles with their teammates. Nevertheless, the introduction of robots in places where humans live or work requires safety, functionality, and effective human-robot interaction (Zollo et al., 2003).

It is here where imitation arises as a very promising approach. Imitation, the ability to recognize, learn and copy the actions of others, rises as *a very promising alternative solution to the programming of robots*. It remains a challenge for roboticists to develop the abilities that a robot needs to perform a task while interacting intelligently with the environment (Bakker and Kuniyoshi, 1996; Acosta-Calderon and Hu, 2003b). Traditional approaches to this issue, such as programming and learning strategies, have been demonstrated to be complex, slow and restricted in knowledge.

Imitation could equip robots with abilities to per-

form *efficient human-robot interaction*, eventually helping humans in personal tasks (Acosta-Calderon and Hu, 2003b; Dautenhahn and Nehaniv, 2002; Becker et al., 1999). It also seems that imitation could be a tool to acquire new behaviors and to adapt these within new contexts (Acosta-Calderon and Hu, 2003a).

Imitation has several advantages that can be transmitted from humans to robots. In humans, this ability permits one to treat the other as a conspecific (Meltzoff and Brooks, 2001) by *perceiving similarities between oneself and other*. This sort of perspective shift may help us to predict actions; enabling us to infer the goal enacted by one another's behaviors (Breazeal et al., submitted for publication).

Our approach to robot imitation is based on how humans acquired the necessary information to understand and execute action (Acosta-Calderon and Hu, 2004a). In humans, the information required to perform an action is obtained from two sources: the body schema, which contains the relations of the body parts and its physical constraints; and the body percept, which refers to a particular body position perceived in an instant (Acosta-Calderon and Hu, 2004b). The body schema and the body percept give us the insight into recognizing actions and thereby performing these actions, therefore, The understanding of other people's actions would lead to imitation (Oztop and Arbib, 2002). We use these fundamental parts and describe their relation throughout four developmental stages used to describe the imitative abilities in humans. This paper describes our approach to addressing imitation of body movements. Results of experiments with a robotic platform implementing mentioned approach are also described. Related works on imitation using robotics arms focus on reproducing the exact gesture, which means to minimize the discrepancy for each joint (Ilg et al., 2003; Zollo et al., 2003; Schaal et al., 2003). The work described here uses a different approach: to focus only on the target and to allow the imitator to obtain the rest of the body configuration. This approach is valid when the imitator and the demonstrator do not share the same body structure.

The rest of the paper is organized as follows. Section II presents the background theory that has inspired our work on imitation. Section III describes briefly the body configuration. In Section IV we present our mechanism for imitation of body movements and implementation issues for the robotic platform. Experiment results are presented in Section V. Finally, Section VI concludes the paper.

2 Background

Humans can perform actions that are feasible with their bodies. To achieve those actions humans use the information derived from two sources (Reed, 2002):

- *The body schema* is the long-term representation between the spatial relations among body parts and the knowledge about the actions that they can and cannot perform.
- *The body percept* refers to a particular body position perceived in an instant. It is built by instantly merging information from sensory input and proprioception; with the body schema. It is the awareness of a body's position at any given moment.

The body schema presents two significant functions, which use the knowledge of the feasible actions that every part of the body can perform:

- *Direct action.* When an action is performed from a current position, a new one is produced.
- *Inverse action*. When an action that satisfies a goal position is selected.

The interaction of both functions allows one to *simulate another person's actions* (Goldman, 2001). When a goal state is identified, then the inverse action generates the motor commands that would achieve

the goal. Those motor commands are sent to the direct action which will predict the next state. This predicted state is compared with the target goal to take further decisions.

These two functions share the idea that has been used in motor control but they are known as controllers and predictors. Demiris and Johnson (2003) used functions with the same principle but called them inverse and forward models.

When we observe someone performing a particular action, one can easily determine how one would accomplish the same task using one's own body. This means that it is possible to *recognize the action* that someone else is performing. The body schema provides the basis to *understand similar bodies and perform the same actions* (Meltzoff and Moore, 1994). This idea is essential in imitation. In order to imitate, it is first necessary to identify the observed actions, and then to be able to perform those actions. Thus, in order to achieve a perceived action a mental simulation is performed constrainting/restraining the movements to those that are physically possible.

There are different approaches to describe the way that humans develop the ability to imitate. One attempt to explain the development of imitation is given by Rao and Meltzoff (2003), who had introduced a four-stage progression of the imitative abilities. Details of those four stages are presented below:

- *Body babbling.* This is the process of learning how specific muscle movements achieve various elementary body configurations. Thus, such movements are learned through an early experimental process, e.g. random trial-and-error learning. Thus, Body babbling is related to the task of building up the body schema (the physics of the system and its constraints).
- *Imitation of body movements.* This demonstrates that a specific body part can be identified i.e. organ identification (Meltzoff and Moore, 1992). Here, the body schema interacts with the body percept to achieve the same movements, once these are identified.
- *Imitation of actions on objects.* This stage starts underlying mental stages about others' behaviour and oneself. This also represents flexibility to adapt actions to new contexts.
- *Imitation based on inferring intentions of actions.* This requires the ability to read beyond the perceived behaviour to infer the underlying goals and intentions.

These four developmental stages serve as a guideline for our progress in research. This paper reports mainly our experiences accomplishing imitation of body movements with a robotic system. We also describe briefly our work on body babbling.

3 Body configuration

Body babbling endows us with the elementary configuration to control our body movements by generating a map. This map contains the relation of all the body parts and their physical limitations. In other words, this map is the body schema.

As humans grow and their bodies change, the body schema is constantly updated by means of input from the body percept. The body percept, in turn, gathers its information from sensory and proprioception information. If there is an inconsistency between the body schema and the body percept, then the body schema is updated.

In robotics, since the bodies of robots are changeless in size and weight, body babbling is simplified by endowing the robot with a control mechanism. Such a mechanism must permit the robot to know its physical abilities and limitations. Therefore, for the experiments with the robotic platform we use the kinematic analysis as the mechanism of position control. The *forward kinematic analysis* calculates the position and orientation of the *gripper* of the robot. In similar way, to determine the values of the robot's joints to produce a desired position and orientation, we use the *Resolve Motion Rate Control (RMRC)*. Further details of these methods and implementation issues can be found in (Acosta-Calderon and Hu, 2004a,b)

4 Imitation of Body Movements

4.1 Identification of a body part

The first step towards imitation is the recognition of the action to imitate. Hence, the imitator must be able to differ among the demonstrator's body parts to identify those those to imitate. The approach described here uses key ideas of the mirror neurons.

These particular neurons have been found in macaque monkeys. These neurons fire when the monkey observes movements executed by another monkey or human demonstrator, as well as when the monkey executes similar goal-oriented movements(Oztop and Arbib, 2002). Neuropsychological experiments in humans described in (Buccino et al.,

2001; Charminade et al., 2002) have revealed brain regions that present similar activity to the one presented by mirror neurons, for both perception and execution of action.

One interesting feature is that, mirror neurons only fire when they perceive similar body parts to the monkey's (mechanical devices do not activate them). Hence, the detection of the similar body parts tends to release the mirror neurons' activity.

Psychologists propose a innate observationexecution pathway in humans (Meltzoff and Moore, 1992; Charminade et al., 2002), here, mirror neurons give a good insight into understanding this idea. Therefore, we can use the same idea of mirror neurons to identify a body part. However, an interesting question arises: do we need to implement a mirror neuron model to every single part of the body? If so, the model would be extremely complicated due to the number of possible combinations of body parts.

The solution could be in an insight of how human beings focus attention on body parts. When humans observe a body movement they do not focus their attention on every single body part. Instead, humans focus their attention on the "end-effector", discarding the position of the other body parts (Mataric, 2002; Mataric and Pomplun, 1998). The body schema finds the necessary body configuration for the rest of our body's parts thereby satisfying the target position for the end-effector.

The implementation of the identification model is done within the body schema module. Here, the end-effector of the demonstrator is marked in distinct color, which can be easily extracted from the image. For our purposes is sufficient to use this simple approach.

Therefore, it is important to remark *the level of imitation* Billard et al. (2004); Dautenhahn and Nehaniv (2002); Nehaniv and Dautenhahn (2002) used in this work. The level of imitation utilized here is *the reproduction of the path followed by the target*, where the imitator will only focus to follow the path described by the end-effector of the demonstrator. The level of *reproduction of the exact gesture* was not chosen due to our approach allows the body schema to find the body configuration satisfying the target position. The discrepancy among the bodies of the imitator and the demonstrator supports the validation of the level of imitation selected. Nevertheless, this discrepancy of bodies arises a problem: the correspondence problem.

4.2 The Correspondence Problem

A successful imitation requires that the imitator be able to recognize structural congruence between oneself and the demonstrator (Meltzoff and Brooks, 2001). When both the demonstrator and the imitator have a common body representation, the body schema of the imitator is then, by itself, capable enough to understand the demonstrator's body. Nevertheless, in a situation where the demonstrator's body differs from the imitator's body schema, there must be a way that the imitator can overcome this so called correspondence problem (Nehaniv and Dautenhahn, 2001, 2002). For our implementation, this correspondence problem is worked out by providing the representation of the body of the demonstrator and a way to relate this representation (Acosta-Calderon and Hu, 2004a).



Figure 1: The correspondence between the bodies of the robot (left) and the demonstrator (right). Two joints, the shoulder and the wrist, have correspondence in both bodies.

Figure 1 presents the correspondence between the body of the demonstrator and that of the imitator. Here, a transformation is used to relate both representations. This transformation is based on the knowledge that in the set of joints of the demonstrator there are three points that represent an arm (shoulder, elbow, and wrist). The remaining two points (the head and the neck) are used just as a reference. The reference points are used to keep a relation among the distances in the demonstrator model. This information about the representation of the demonstrator is extracted by means of color segmentation.

The transformation relates the demonstrator's body to the robot's body. The demonstrator's shoulder is used as the origin of the workspace of the robot. Hence, the shoulder of the demonstrator is treated as the reference point for the calculation of the remaining two points of the demonstrator's arm. Note that only the position of the demonstrator's end-effector (wrist) is then converted and fitted into the workspace of the robot.

Each new position of the end-effector identified in the workspace of the robot triggers the body schema to fulfill it. Since the robot only cares about the position of the end-effector, it uses the body schema (the control method) to obtain the rest of the body configuration (Acosta-Calderon and Hu, 2004a).



Figure 2: The architecture used to imitate the body movements. The information about the demonstrator is extracted and then converted to the robot's workspace. This information represents the new position to be imitated.

The mechanism implemented for the imitation of body movements is depicted in Fig. 2. Hence, to satisfy a new position of the end-effector the body schema employs the *inverse action* function (Resolve Motion Rate Control - RMRC). This function obtains the new values for the body parts to satisfy the desired position. The body configuration obtained leads to a controllable motion preventing the joints from moving too fast whilst the kinetic energy is minimized; just like humans do when we imitate *the path described by the target* and not *the exact gesture*. Further details of the RMRC implemented can be seen at Acosta-Calderon and Hu (2004a).

Although, the body configuration obtained for the robot, might not be similar to the one presented by the demonstrator. Instead of copying the extract posture, the level of imitation that we are addressing is to reproduce the same goal position. This is mainly because the robot and the demonstrator do not share the same body structure. This can avoid the situation where one body configuration can not be achieved by physical constraints. Here, the body schema plays a crucial role minimizing the motion between positions, while considering the physical constraints, and selecting the more efficient body configuration.

Once a body configuration has been found this can either be sent to the actuators and executed, or inhibit the output to the actuators and send it to the *direct action* function (Forward Kinematics). The direct action will simulate the action of sending those values to the actuators and return the achieved position for that particular set of values. The new reached position is used to generate the current body percept, as the new position, in other words, a mental rehearsal of the observed action.

4.3 Movements

The movements imitated are represented as paths consisting of a set of points. Each point represents the demonstrator's end-effector both the position (defined in Cartesian coordinates by x, y, and z) and the orientation (defined by the *roll*, *pitch*, and *yaw* angles) (Acosta-Calderon and Hu, 2004a,b).

Each new position in the movement of the demonstrator is smoothened by using *cubic spline curves*. These kinds of curves have the feature that they can be interrupted at any point and fit smoothly to another different path. More points can be added to the curve without increasing the complexity of the calculation. Using spline curves reduces the noise in the data from the color segmentation.

The identification of a movement is a complex process. In the process of identification it is necessary to find, if there is, a matching movement from the previously learnt movements in the library.

The matching process consists of comparing a movement with those already stored in the library, and selects the one with the minimal error defined by (1)

$$\varphi_k = \arg_i \min\left(\overrightarrow{f} - \overrightarrow{f_i}\right)^2$$
 (1)

where φ_k is the minimal error for the movement in the library with the index k. $\overrightarrow{f_i}$ is the function that represents the featuring vector of the movement with index i, as shown in (2). The minimal error obtained from the elements in the library does not guarantee the new element corresponds to a similar class of movements. Hence, the minimal error φ_k is compared with a threshold. When φ_k is less than the threshold, it is assumed that the observed movement is close enough to the one represented by the best match k. Thus, the movement k in the library is updated using interpolation with the observed movement. On the other hand, when φ_k is greater than the threshold, the observed movement would be treated as a new movement and finally added to the library. This process can be seen from Fig. 3

$$\overrightarrow{f_i} = (f_1, f_2, \dots, f_N) \tag{2}$$

The extraction of the features for the movement i is performed by using a *grid-based extraction* as described by Shen and Hu (2004). This method divides an image into a fixed number of cells N defined by the number of *columns* and *rows*. The next step is to visit each cell j in the grid and the number of Relevant Features RF counted. Finally, this value is normalized by the total number of Relevant Features of the movement i via (3).

$$f_j = \frac{RF_j}{\sum RF_j} \tag{3}$$

After visiting all the cells, all the feature values f_j are collected into the featuring vector $\overrightarrow{f_i}$. The values contained in the featuring vector are relative values, which are robust to variations in the slope of the movement. A variation in the slope of a sub-area of the movement does not represent a significant variation in the featuring vector.



Figure 3: Interpolation of the library movement (a) with a new movement (b) the result is movement (c).

Figure 4 presents two movements divided into subareas by the grid. In order to compare both movements they must have the same scale, the same number of columns and rows, and of course, the same number of pixels in each sub-area of the grid.

5 Experimental results

To investigate the abilities of the approach presented, we described our experience with experiments of imitation of actions on objects. In our set-up we used



Figure 4: Two movements are divided into cells and to be compared.

the robot United4, as the imitator, which faced a human demonstrator. The robot observed the movements performed by the demonstrator in order to imitate them later. The experiments were conducted in two phases for all the cases:

- *Learning phase*, in which the robot was observing the demonstrator's movements, while identifying and recording them to be executed later.
- *Execution phase*, here the robot performs the movements learnt in the previous phase.

The robotic platform used is a mobile robot Pioneer 2-DX with a Pioneer Arm and a camera, namely United4. The robot is a small, differential-drive mobile robot intended for indoors. The robot is endowed with the basic components for sensing and navigation in a real-world environment. It is also equipped with a color tracking system. United4 has a Pioneer Arm, which is a robotic arm with five *degrees-of-freedom*, the end-effector is a gripper with fingers allowing for the grasping and manipulation of objects.

The experiments have been conducted in our Brooker laboratory. The relevant objects in the environment (demonstrator's joints) were marked with different colors to simplify the feature extraction. The less cluttered background permits the robot to focus only on the significant information. We also consider only planar motions in order to validate our approach.

Our first set of experiments of movements of body parts involved movements describing different paths. In Figure 5, we present one path used in the experiments.

In Fig. 5, (a), (c), and (e) show the demonstrator performing a path from up to down with his right hand. While (b), (d), and (f) present the robot imitating such movement. In addition, We can observe that the robot presented *the mirror effect*. Hence, if the demonstrator, located in front of the robot, moves its



Figure 5: Movements performed by the demonstrator and imitated by the robot.



Figure 6: The movement of the demonstrator (solid line) and the performance of the robot (dotted line), extracted from the movements in Fig. 5.

left arm, then the imitator would move its arm toward the right, acting as a mirror.

In Fig. 6, the solid line is the path extracted from the movements performed by the demonstrator in Fig. 5. The dotted line represents the robot's performance. The path was extracted and adjusted in order to be performed by the robot since the size and shape of the workspace for the model and the robot were not the same.

The second set of experiments on imitation of body movements involved movements writing different letters, e.g. e, s. The robot observed the demon-



Figure 7: During the learning phase, shown in (a) and (b) the demonstrator is writing the letters "e" and "s". During the execution phase, shown in (c) and (d) the robot is writing those letters.

strator performing the handwriting while, by means of the colored markers that the demonstrator wears, the body representation of the demonstrator was extracted. This representation was related with the robot's representation by the body schema. Therefore, the robot could understand the new position of the demonstrator's end-effector within its workspace. The configuration needed to reach this desired position was eventually calculated by means of the kinematics methods. Finally, the path described by the end-effector was recorded and ready to be executed.



Figure 8: Letter "e." The solid line is the performance of the robot (from Figure 7.c), and the dotted line is the path that the robot generated after observing the demonstrator's performance (from Figure 7.a).

Figure 7 presents the letters "e" and "s." The learning phase is presented in (a) and (b), where the demonstrator has written these letters. When the demonstrator was describing the path of these letters, the robot was observing and relating those movements to its own. In the execution phase, (c) and (d), the robot is performing the paths described by the letters.



Figure 9: Letter "s." The Performance of the robot in the solid line (from fig. 7.d), and the dotted line is the path that the robot generated by observing the demonstrator performance(from fig. 7.b).

Each path is extracted and adjusted in order to be performed by the robot since the size and shape of the workspace for the model and the robot are not the same. To minimize the noise in the path, we smooth the path by using *cubic spline curves*.

6 Conclusions and future work

Roboticists have begun to focus their attention on imitation. Since the capability to obtain new abilities by observation presents considerable advantages in contrast with traditional learning approaches. Finally, imitation might equip robots with the abilities for an efficient human-robot interaction.

The presented approach is based on *the body schema* and *the body percept*, which are used by humans to understand how other people perform actions. Since the knowledge of feasible actions and physical constraints is implicit in the body schema, it is possible to do a mental rehearsal of other peoples' actions and gather the results of those actions at particular body percepts for the body schema. It is believed that these two key-parts play a crucial role in achieving imitation.

We used an approach of four developmental stages of imitation in humans, to prove the key-role of these two components. The scope of this paper describes our progress mainly on imitation of body movements. In this stage, we used the idea to focus on the endeffector as humans do and to allow the body schema to obtain the rest of the configuration. We have also described our experiments with a robot as the imitator, imitating the movements of a human demonstrator. Our experiments show the feasibility of the proposed approach at this stage of imitation. Our future work involves extending the experiments to the next stage, imitation of action on objects.

References

- C. A. Acosta-Calderon and H. Hu. Goals and actions: Learning by imitation. In *Proc. AISB03 Second Int. Symposium on Imitation in Animals and Artifacts*, pages 179–182, Aberystwyth, Wales, 2003a.
- C. A. Acosta-Calderon and H. Hu. Robotic societies: Elements of learning by imitation. In *Proc. 21st IASTED Int. Conf. on Applied Informatics*, pages 315–320, Innsbruck, Austria, 2003b.
- C. A. Acosta-Calderon and H. Hu. Imitation towards service robotics. In *Int. Conf. on Intelligent Robots and Systems IROS 2004*, pages 3726–3731, Sendai, Japan, 2004a.
- C. A. Acosta-Calderon and H. Hu. Robot imitation: A matter of body representation. In *Int. Symposium on Robotics and Automation ISRA 2004*, pages 137–144, Queretaro, Mexico, 2004b.
- P. Bakker and Y. Kuniyoshi. Robot see robot do: An overview of robot imitation. In AISB96 Workshop on Learning in Robots and Animals, pages 3–11, Brighton, England, 1996.
- M. Becker, E. Kefalea, E. Mael, C. V. D. Malsburg, M. Pagel, J. Triesch, J. C. Vorbruggen, R. P. Wurtz, and S. Zadel. Gripsee: A gesture-controlled robot for object perception and manipulation. *Autonomous Robots*, (6):203–211, 1999.
- A. Billard, Y. Epars, S. Calinon, S. Schaal, and G. Cheng. Discovering optimal imitation strategies. *Robotics and Autonomous Systems*, (47):69– 77, 2004.
- C. Breazeal, D. Buchsbaum, J. Gray, D. Gatenby, and B. Blumberg. Learning from and about others: Towards using imitation to bootstrap the social understanding of others by robots. *Artificial Life*, submitted for publication.
- G. Buccino, F. Binkofski, G. R. Fink, L. Fadiga, L. Fogassi, V. Gallese, R. J. Seitz, K. Zilles, G. Rizzolatti, and H. J. Freund. Action observation activities premotor and parietal areas in a somatotopic

manner: An fmri study. *European Journal of Neuroscience*, (13):400–404, 2001.

- T. Charminade, A. N. Meltzoff, and J. Decety. Does the end justify the means? a pet exploration of the mechanisms involved in human imitation. *Neuro Image*, (15):318–328, 2002.
- K. Dautenhahn and C. L. Nehaniv. *Imitation in Animals and Artefacts*, chapter The Agent-Based Perspective on Imitation. The MIT Press, Cambridge, MA, 2002.
- Y. Demiris and M. Johnson. Distributed, predictive perception of actions: a biologically inspired robotics architecture for imitation and learning. *Connection Science*, 15(4):231–243, 2003.
- T. Fong, I. Nourbakshsh, and K. Dautenhahn. A survey of social robots. *Robotics and Autonomous Systems*, (42):143–166, 2002.
- A. Goldman. *Intention and Intentionality*, chapter Desire, Intention, and the Simulation Theory, pages 207–224. The MIT Press, Cambridge, MA, 2001.
- W. Ilg, G. H. Bakir, M. O. Franz, and M. A. Giese. Hierarchical spatio-temporal morphable models for representation of complex movements for imitation learning. In *Proc. of the 11th IEEE Int. Conf. on Advanced Robotics*, pages 453–458, Coimbra, 2003.
- J. Liu and J. Wu. *Multi-Agent Robotic Systems*. CRC Press, London, 2001.
- M. J. Mataric. Imitation in Animals and Artefacts, chapter Sensory-Motor Primitives as a Basis for Imitation: Linking Perception to Action and Biology to Robotics, pages 392–422. The MIT Press, Cambridge, MA, 2002.
- M. J. Mataric and M. Pomplun. Fixation behavior in observation and imitation of human movement. *Cognitive Brain Research*, 7(2):191–202, 1998.
- A. N. Meltzoff and R. Brooks. *Intention and Intentionality*, chapter Like Me as a Building Block for Understanding Other Minds: Bodily Acts, Attention, and Intention, pages 171–191. The MIT Press, Cambridge, MA, 2001.
- A. N. Meltzoff and M. K. Moore. Early imitation within a functional framework: The important of person identity, movement, and development. *Infant Behaviour and Development*, (15):479–505, 1992.

- A. N. Meltzoff and M. K. Moore. Imitation, memory, and representation of persons. *Infant Behaviour* and Development, (17):83–99, 1994.
- C. L. Nehaniv and K. Dautenhahn. Like me? measures of correspondence and imitation. *Cybernetics and Systems*, 32(1):11–51, 2001.
- C. L. Nehaniv and K. Dautenhahn. *Imitation in Animals and Artefacts*, chapter The Correspondence Problem. The MIT Press, Cambridge, MA, 2002.
- E. Oztop and M. A. Arbib. Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87(2):116–140, 2002.
- R. P. N. Rao and A. N. Meltzoff. Imitation learning ininfants and robots: Towards probabilistic computational models. In *Proc. AISB03 Second Int. Symposium on Imitation in Animals and Artifacts*, pages 4–14, Aberystwyth, Wales, 2003.

- C. L. Reed. *The Imitative Mind*, chapter What is the body schema?, pages 233–243. Cambridge University Press, Cambridge, 2002.
- S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Philosophical Transaction of the Royal Society of London: Series B, Biological Sciences*, (358):537– 547, 2003.
- J. Shen and H. Hu. Mobile robot navigation through digital landmarks. In *Proc. of the 10th Chinese Automation and Computing Conf.*, pages 151–156, Liverpool, England, 2004.
- L. Zollo, B. Siciliano, C. Laschi, G. Teti, and P. Dario. An experimental study on compliance control for a redundant personal robot arm. *Robotics and Autonomous Systems*, (44):101–129, 2003.

Achieving Corresponding Effects on Multiple Robotic Platforms: Imitating in Context Using Different Effect Metrics

Aris Alissandrakis, Chrystopher L. Nehaniv, Kerstin Dautenhahn, and Joe Saunders*

*Adaptive Systems Research Group School of Computer Science, University of Hertfordshire College Lane, Hatfield, Hertfordshire, AL10 9AB, UK a.alissandrakis@herts.ac.uk

Abstract

One of the fundamental problems in imitation is the *correspondence problem*, how to map between the actions, states and effects of the model and imitator agents, when the embodiment of the agents is dissimilar. In our approach, the matching is according to different metrics and granularity. This paper presents JABBERWOCKY, a system that uses captured data from a human demonstrator to generate appropriate action commands, addressing the correspondence problem in imitation. Towards a characterization of the *space of effect metrics*, we are exploring absolute/relative angle and displacement aspects and focus on the overall arrangement and trajectory of manipulated objects. Using as an example a captured demonstration from a human, the system produces a correspondence solution given a selection of *effect metrics* and starting from dissimilar initial object positions, producing action commands that are then executed by two imitator target platforms (in simulation) to successfully imitate.

1 Introduction

Imitation is a powerful learning tool that can be used by robotic agents to socially learn new skills and tasks. One of the fundamental problems in imitation is the *correspondence problem*, how to map between the actions, states and effects of the model and imitator agents, matching according to different metrics and granularity, when the embodiment of the agents is dissimilar (Nehaniv and Dautenhahn (1998)). The following statement of the *correspondence problem* (Nehaniv and Dautenhahn (2000, 2001, 2002)) draws attention to the fact that the model and imitator agents may not necessary share the same morphology or may not have the same affordances:

Given an observed behaviour of the model, which from a given starting state leads the model through a sequence (or hierarchy [or program]) of *sub-goals* in *states*, *action* and/or *effects*, one must find and execute a sequence of actions using one's own (possibly dissimilar) embodiment, which from a corresponding starting state, leads through corresponding sub-goals - in corresponding states, actions, and/or effects, while possibly responding to corresponding events.

In this approach, a solution to the correspondence problem can be used to generate a *recipe* (a loose plan) through which an imitator can map sequences of observed actions of the model agent to its own repertoire of actions as constrained by its own embodiment and by context (Nehaniv and Dautenhahn (2000, 2001, 2002)). Qualitatively different kinds of social learning result from matching different combinations of matching actions, states and effects at different levels of granularity (Nehaniv (2003)). The sub-goals define the granularity to match and vice versa.

Artificial agents that have the ability to imitate may use (perhaps more than one) metric to compare the imitator agent's own actions, states and effects with the model's actions, states and effects, in order to evaluate the imitation attempts and discover corresponding actions that they can perform to achieve a similar behaviour. The choice of metrics used is therefore very important as it will have an impact on the quality and character of the imitation. Many interesting and important aspects of the model behaviour



Figure 1: **The JABBERWOCKY system architecture.** Using data captured from a human and given appropriate metrics and sub-goal granularity, the multi-target system can produce action command sequences that when executed by a software or hardware agent can achieve corresponding actions, states and/or effects. The corresponding actions, states and effects as demonstrated by the imitator can also be captured and used as a demonstration for another imitating agent. Differently embodied and constrained target systems in various contexts need to be supported.

need to be considered, as the metrics capture the notion of the salient differences between performed and desired actions and also the difference between attained and desired states and effects (Nehaniv and Dautenhahn (2001, 2002)). The choice of metric determines, in part, what will be imitated, whereas solving the correspondence problem concerns how to imitate (Dautenhahn and Nehaniv (2002)). In general, aspects of action, state and effect as well as the level of granularity (what to imitate) do all play roles in the choice of metric for solving the problem of how to imitate (Nehaniv and Dautenhahn (2001); Alissandrakis et al. (2002); Billard et al. (2004)). On-going research is thus addressing the complementary problem of how to extract sub-goals and derive suitable metrics automatically from observation (Nehaniv and Dautenhahn (2001); Nehaniv (2003); Billard et al. (2004); Calinon and Billard (2004)).

In our setting, it will be desirable to have different kinds of agents in the learning process, i.e. humans and robots interacting socially. Focusing on object manipulation and arrangement demonstrated by a human, this paper presents a system that uses different metrics and granularity to produce action command sequences that when executed by an imitating agent can achieve corresponding effects (manipulandum absolute/relative position, displacement, rotation and orientation). Depending on the particular metrics and granularity used, the corresponding effects will differ (shown in an example), making the appropriate choice of metrics and granularity depend on the task and context.

The work presented in this paper is motivated by the EU Integrated project COGNIRON ("The Cognitive Robot Companion") and addresses the problem of how to teach a robot new complex tasks through human demonstration. The learning algorithms to be developed should be general and address fundamental questions of imitation learning, applied to manipulation tasks. For example a robotic companion at home could acquire knowledge of e.g. laying out a table or drawing on a canvas from observing its human owner. Acquiring such skills socially requires matching different aspects of the effects that the human actions have on objects in the environment. Also the context within which a skill is replicated might require its generalization to various settings and to other types and shapes of objects.

2 The JABBERWOCKY System

This section presents the JABBERWOCKY system developed for the COGNIRON project, that uses captured data from a human demonstrator to generate appropriate action commands (see Figure 1), addressing the correspondence problem in imitation. The action commands can be targeted for various software and hardware platforms. These actions will allow the imitating agent to achieve corresponding actions, states and effects, depending on the given (relevant to the demonstrated task and context) metrics and granularity (provided by a *what to imitate* and *sub-goal extraction* module), embodiment restrictions and constraints (imposed by the targeted imitator platform), and possibly different initial state of the objects in the environment.

The design of the JABBERWOCKY system is informed by the ALICE (Action Learning via Imitating Corresponding Embodiments), a generic framework



Figure 2: Captured demonstration (left), and the extracted critical points (right). The colors (red, green and blue) indicate the three different objects. The dotted outlines indicate the initial position and orientation of the objects, while the solid thick outline the final position. For the demonstration data, the intermediate object's position and orientation is shown with solid thin outlines, linearly scaled (at intervals equal to one tenth of the overall trajectory only, for clarity) to indicate the direction of the movement. For the critical points, each object's position and orientation is shown at every critical point, again linearly scaled.

for solving the correspondence problem (see Alissandrakis et al. (2002, 2004)). The ALICE framework builds up a library of actions from the repertoire of an imitator agent that can be executed to achieve corresponding actions, states and effects to those of a model agent (according to given metrics and granularity). The ALICE framework provides a functional architecture that informs the design of robotic systems that can learn socially from a human demonstrator.

The system bears some similarity to the one presented in (Kuniyoshi et al. (1994)), but with the main differences that it can use *any* given metric and granularity and that it is designed to be able to generate action commands targeted for a variety of platforms, both in software and hardware to match different behaviour aspects and achieve various types of social learning.

2.1 Demonstrator (Model Agent)

The system uses captured data from a human demonstrator. The demonstrated behaviour is captured using motion sensors (Polhemus LIBERTYTM motion capture system). By attaching the motion sensors on the arms, hands and torso of the human, as well as on the objects that the demonstrator is manipulating, we can obtain the *actions* (e.g. hand movements, gestures),



Figure 3: An example of dissimilar initial object positions. The dotted outlines indicate the initial position and orientation of the objects in the demonstrator's workspace (from the demonstration shown in Figure 2, left) and the solid outlines the (dissimilar) initial configuration of the objects in the imitator's workspace. The blue object has the same initial position.

states (e.g. arm and body postures) and *effects* (e.g. positioning, displacement, rotation of objects in the workspace) of the demonstrator.

In example shown in Figure 2, the demonstrated task consists of three block objects (colored red, green and blue) arranged in a 2D workspace surface by a human who acts as the demonstrator. The workspace is a square grid 50 cm by 50 cm, and the sizes of the objects are: 10 cm by 8 cm (red) and 8 cm by 5 cm (green and blue). As the manipulations occur only in a 2D plane, only the XZ dimensions are given here (and shown in the figures) omitting the Y dimension (height).

The current work focuses on the *effects* aspect of the demonstrated behaviour, so only the position and orientation of the objects as they are manipulated by the demonstrator are captured, omitting the demonstrator's actions (arm movements) and states (body posture). The choice of initially concentrating on effects for this work is guided by the assumption that the manipulation of objects will be the most important aspect of the demonstrated behaviours that users would like a robotic companion to imitate in a home environment (e.g. fetching objects or arranging them in particular ways).

In ongoing work, three (or more) additional sensors will be used, one attached to the human torso and one at each hand/arm, providing additional information about the demonstrator's states and actions. Taking into account the *states* aspect would help the JAB-BERWOCKY system solve possible ambiguities when producing the corresponding actions for imitation. For example, a humanoid robot imitator, considering the states of the demonstrator would obtain possibly useful information e.g. which hand to use (left or right) to reach an object from its current configuration, based on the choice of hand used by the demonstrator.

2.2 What to Imitate Module

The character of the imitation will depend on the metrics and granularity (Alissandrakis et al. (2002, 2004); Alissandrakis (2003)). The *what to imitate* module will use the captured demonstration data to extract appropriate sub-goals (granularity) and also discover what metrics must be used to capture the appropriate aspects of the particular demonstration.

In the current implementation of the JABBER-WOCKY system the metrics and the sub-goal granularity are given, instead of being discovered by the imitator agents based on the observed demonstrated task. The *what to imitate* module provides a choice of metrics and granularity based on the task and context of the demonstration, although there might not always be a unique, "correct", choice. Here, the various possible metrics and granularity have been selected in advance. It can be shown that the character of the resulting matched effects can be very different, depending on the choice of metrics and granularity used.

The sub-goal granularity is given by finding the *critical points* in the trajectories of the manipulated objects. A critical point occurs when the direction of the captured trajectory and/or the orientation of an object changes by more than a certain threshold.

Several different *effect* metrics have been defined (see section 2.3) that are used in the experiment presented in this paper. In the future the work will be extended to consider also the *state* and *action* aspects of a demonstration.

2.3 Metrics

Towards a characterization of the *space of effect metrics*, we are exploring absolute/relative angle and displacement aspects and focus on the overall arrangement and trajectory of manipulated objects. Looking at how objects can be manipulated (in a nondestructive and combining way), there are two different perspectives: how the object was displaced and how it was rotated. The displacement can be either relative or absolute related to the final position in the workspace, or relative to the other objects within the workspace. The rotation can be also be relative or absolute related to the final orientation of the object. To fully describe the manipulation of an object, both displacement and angular effect aspects must be considered. We consider these aspects in a two-dimensional workspace, such as a table surface.

If the initial configuration of the (same or corresponding) objects is the 'same' for both the model and the imitator agents, then there is no observable distinction between using either the absolute and relative displacement/rotation or the relative position (if the objects are manipulated in the same order). But if the agents are active in a different workspace starting from a different initial configuration of objects, or the timing and the order of the manipulations is not the same, it will be impossible to satisfy simultaneously all the aspects. Therefore choosing to satisfy one particular aspect will result in a qualitatively different effect than if another one was chosen, but still satisfy those similarity quantitative criteria.

2.3.1 Displacement Effect Metrics

The model is moving an object from position X_M to position X'_M on the workspace, achieving an object displacement $\Delta X_M = X'_M - X_M$, where $X_M = \begin{bmatrix} x_M \\ y_M \end{bmatrix}$, $X'_M = \begin{bmatrix} x'_M \\ y'_M \end{bmatrix}$, and $\Delta X_M = X'_M - X_M = \begin{bmatrix} x'_M - x_M \\ y'_M - y_M \end{bmatrix}$. The imitator should move the same (or corresponding) object from position X_I to position X'_I on the workspace, with a displacement $\Delta X_I = X'_I - X_I$, such that a displacement metric is minimised (see Fig. 4, left). **Relative Displacement Effect Metric** is minimized if $\Delta X_I = \Delta X_M$ and $X'_I = X_I + \Delta X_M = \begin{bmatrix} x_I \\ y_I \end{bmatrix} + \begin{bmatrix} x'_M - x_M \\ y'_M - y_M \end{bmatrix} = \begin{bmatrix} x_I + x'_M - x_M \\ y_I + y'_M - y_M \end{bmatrix}$. **Absolute Displacement Effect Metric** is minimized if $X'_I = X'_M$ and $\Delta X_I = X'_M - X_I = \begin{bmatrix} x_M - x_M \\ y'_M - y_M \end{bmatrix}$

$$\begin{bmatrix} X_I \\ x'_M - x_I \end{bmatrix}$$

 $\begin{bmatrix} y'_M - y_I \end{bmatrix}^T$ **Relative Position Effect Metric** is minimized if the object is moved to a similar position relative to other objects in the workspace. The *relative position* effect metric is defined here for three objects in the workspace.

The center of the manipulated object is defined as $A = \begin{bmatrix} x_A \\ y_A \end{bmatrix}$, and the centers of the other two objects as $B = \begin{bmatrix} x_B \\ y_B \end{bmatrix}$ and $C = \begin{bmatrix} x_C \\ y_C \end{bmatrix}$. The imitator must move the same (or corresponding) object to form a triangle *ABC* so that it is the "same" as the triangle formed by the model, i.e. the angles $C\hat{A}B$,



Figure 4: A selection of *displacement* (left) and *angular* metrics (right). To evaluate the similarity between object displacements, the *relative displacement*, *absolute position* and *relative position* effect metrics can be used. To evaluate the similarity between object rotations, the *rotation* and *orientation* effect metrics can be used. The second row shows the way the corresponding object (in a different workspace) needs to be moved or rotated by an imitator to match the corresponding effects. The grey triangles are superimposed to show that for the *relative position* effect metric, the relative final positions of the objects are the same.

 $A\hat{B}C$ and $B\hat{C}A$ are equal. The triangle sides $A\bar{B}$, $\bar{B}C$ and $\bar{C}A$ can be equal only if the objects start from the same initial configuration for both agents and are manipulated in the same order, so only the equality of the angles can be used in general¹.

The relative position effect metric is minimized if

$$X'_I = A \text{ and } \Delta X_I = A - X_I = \begin{bmatrix} x_A - x_I \\ y_A - y_I \end{bmatrix}.$$

2.3.2 Angular Effect Metrics

The model is rotating an object from orientation θ_M to orientation θ'_M on the workspace, with a rotation $\Delta \theta_M = \theta'_M - \theta_M$. The imitator should rotate the same (or corresponding) object from orientation θ_I to orientation θ'_I on the workspace, with a rotation $\Delta \theta_I = \theta'_I - \theta_I$, such that a displacement metric is minimised (see Fig. 4, right).

Rotation Effect Metric is minimized if $\Delta \theta_I = \Delta \theta_M$ and $\theta'_I = \theta_I + \Delta \theta_M$.

Orientation Effect Metric is minimized if $\theta'_I = \theta'_M$ and $\Delta \theta_I = \theta'_M - \theta_I$.

2.3.3 Other Effect Metrics

Depending on the initial configuration of the corresponding objects in the imitator's workspace, or the particular task that the imitator would like to achieve, it might be desirable to use also other metrics that take into account mirror symmetry, both positional and angular, to features of the environment or other agents. For example: Mirror Displacement Effect Metric is minimized if

$$\Delta X_I = -\Delta X_M \text{ and } X'_I = X_I - \Delta X_M = \begin{bmatrix} x_I \\ y_I \end{bmatrix} - \begin{bmatrix} x'_M - x_M \\ y'_M - y_M \end{bmatrix} = \begin{bmatrix} x_I - x'_M + x_M \\ y_I - y'_M + y_M \end{bmatrix}.$$
Mirror Botation Effect Metric is minimized if

Mirror Rotation Effect Metric is minimized if $\Delta \theta_I = -\Delta \theta_M$ and $\theta'_I = \theta_I - \Delta \theta_M$.

Parallel Orientation Effect Metric is minimized if $\theta'_I = \vartheta$ and $\Delta \theta_I = \vartheta - \theta_I$, where ϑ is the orientation of a feature in the environment (e.g. one edge of the table). If the features in the workspace of the imitator are the same as the model's, then $\vartheta \equiv \theta'_M$ and this metric becomes equivalent to the *orientation effect metric*.

2.4 Combinations of Effect Metrics

To evaluate both the movement and the orientation of an object, both metric types must be used. To match the observed effect, the (corresponding) object needs to be moved on the workspace according to the displacement given by the displacement effect metric and rotated according to the angular effect metric used.

A weighted combination of more than one displacement metric can also be used, by averaging the displacement vectors that minimise each metric. For example, if $\Delta X_i = \begin{bmatrix} \Delta x_i \\ \Delta y_i \end{bmatrix}$ is the displacement that minimises a displacement effect metric *i*, and $\omega_1, \ldots, \omega_n$ are the weights of the *n* displacement effect metrics to be combined, the displacement that



Figure 5: The *three robots as objects* imitator platform.

minimizes this composite metric is then given by

$$\Delta X = \begin{bmatrix} |\Delta X| \times \cos(\phi) \\ |\Delta X| \times \sin(\phi) \end{bmatrix}, \text{ where } |\Delta X| = \omega_1 \times \sqrt{\Delta x_1^2 + \Delta y_1^2} + \dots + \omega_n \times \sqrt{\Delta x_n^2 + \Delta y_n^2} \text{ and } \phi = \omega_1 \times \tan^{-1} \left(\frac{\Delta y_1}{\Delta x_1}\right) + \dots + \omega_n \times \tan^{-1} \left(\frac{\Delta y_n}{\Delta x_n}\right).$$

2.5 Imitator

The system is addressing the correspondence problem for dissimilarly embodied imitators, so the *how to imitate* module must produce action commands that can be used by multiple different target platforms as imitator agents, both in simulation (software) and hardware (robots).

Each particular target platform will pose different embodiment restrictions and constraints to the actions, states and effects it can achieve, and eventually to the quality and character of the imitation.

The demonstrator and the imitator might share the same workspace or they might operate in different ones. Even in the same workspace, unless the objects and agents positions are arranged back into the same initial configuration before the imitative behaviour, the context will be different and the imitator therefore has to take that into consideration when imitating.

Two targeted platforms are used in the current realization of the system, both implemented using the WebotsTM robot simulation software.

2.5.1 Three Robots As Objects

In the first imitator platform, the imitator's workspace contains no objects. Instead, the imitator is 'embodied' as three mobile robots, each corresponding to one of the objects manipulated by the demonstrator (see Figure 5). Each robot is square 4cm by 4cm (so in this case, besides dissimilar demonstrator-imitator



Figure 6: The manipulator and three objects imitator platform.

embodiments, there is also dissimilar object correspondence, mapping the objects to mobile robots). The robots can follow the individual trajectories of the objects as arranged by the demonstrator, but cannot match the orientation (while moving) because they are differential wheel robots. Therefore the *angular* effect aspect will be ignored when they imitate, matching only the *displacement* effect aspect.

In the simulation, as the robots move around the workspace, they leave behind a colored trail (of same color as themselves and their corresponding objects) to help visualize the imitated trajectories.

2.5.2 Manipulator and Three Objects

In the second targeted imitator platform, the imitator's workspace contains three objects, of the same size and color as the corresponding objects in the demonstrator's workspace (in this case). The imitator is embodied as a single arm manipulator, positioned above the workspace and able to pick-up, move and rotate the three objects (see Figure 6). This embodiment, although dissimilar to the one of the human demonstrator, is nevertheless able to match both *displacement* and *angular* effect aspects of the demonstration.

As the objects are moved (and rotated) around the workspace by the manipulator in the simulation, they leave behind a colored trail (of same color as themselves) to help visualize the imitated trajectories. The manipulator is shown as a vertical yellow cylinder mounted at the end of a bar positioned above the workspace.

2.6 How to Imitate Module

The *how to imitate* module uses the captured data from the demonstration, the metrics and the sub-goal

granularity discovered by the *what to imitate* module to produce a sequence of action commands for an agent to execute and imitate. These action commands are made target specific by taking into account the particular embodiment, affordances and restrictions of the target imitator agent, and also contextual information (including the initial state) for both the agent and the environment. In the current system implementation both the metrics and the sub-goal granularity (critical points) are given.

Concentrating on the effects aspect of the demonstrated behaviour to be imitated only, an embodiment-independent solution to the correspondence problem can be found, taking into account the effect metrics and the sub-goal granularity. For example consider a human opening a cupboard, removing an object, closing the cupboard and placing the object on a table. This sequence of events can be achieved by agents of varying embodiments, ignoring state aspects like e.g. which hand was used to open the cupboard or how the object was held (or grasped) or even action aspects e.g. the way the human walked (gait) across the room. Any agent that can open the cupboard, transport the object and place it on the table can potentially imitate the effects of this particular demonstration. But for this solution to be useful to an imitating robotic companion, it must be converted to action commands that take into account its embodiment and also the context (e.g. the cupboard is already open, the object is located on a different shelf in the cupboard, the table is in another room), so that the imitator uses its motors and actuators to achieve the desired effects of the task.

The *how to imitate* module considers the given effect metrics and sub-goal granularity, together with the (possible dissimilar) initial configuration of the objects in the imitator's workspace (also given) to produce initially an embodiment-independent correspondence solution (since only the effects behaviour aspects are considered).

To discover this correspondence, the JABBER-WOCKY system currently uses a simple simulation of the 2D workspace that can handle various 'block' objects moving and rotating around, accounting for object collisions and workspace confines. This simulation can replay the captured model data at a given granularity, displaying the trajectory and orientation of the objects as they move and rotate on the workspace, from the initial configuration to the final captured frame. In parallel, starting from a different initial configuration of the same (or different) corresponding objects on the imitator's workspace, the simulation produces a sequence of changes to displacement and rotation for each object, that minimize the given effect metrics.

For example if the effect metric used is the relative displacement effect metric, and the demonstrator moved an object 10 cm to the right, then in order to minimize the metric, the corresponding object in the imitator's workspace must be also moved 10 cm to the right. But some displacements or rotations, although minimizing the metric, might be invalid because the path or final position is occupied by other objects or agents, e.g. if the corresponding object is less than 10 cm away from the right edge of the workspace (because the initial position was different), the entire move cannot be performed. The how to imitate module will then have to discover an alternative way in the given context (including other agents, static or dynamic obstacles) to achieve the same effects according to the metric. In this case it might be acceptable to move the object up to the right edge and then continue the rest of the imitative behaviour. In another context, it might be preferable not to move the object at all. This contextual information should be ideally provided by the what to imitate module, based on observations of the currently demonstrated task and not pre-defined. In the current JABBERWOCKY implementation, the system attempts to move (or rotate) the objects until they reach an obstacle (based on simple 2D object collision detection), and then stop, instead of considering another path to reach the position (and/or achieve the orientation) that minimizes (if possible) the metric used.

To imitate and achieve similar *effects* as the model, an imitator agent will have to adopt this (largely) embodiment-independent correspondence solution to move and rotate the objects, using a generated sequence of action command instructions. These action commands will be targeted to multiple imitator platforms, taking into account the embodiment constraints and restrictions of imitator embodiments.

Figure 7 (left) shows an example correspondence converted to action commands for the *three robots as objects* target platform. Each robot is given a sequence of way-points, depending on its corresponding object. For each of these way-points, the robot must use its differential wheel embodiment to move in a straight line up to that position in the workspace, and after reaching the target position, move on to the next. Figure 7 (right) shows the resulting captured imitative behaviour.

Figure 8 (left) shows an example correspondence converted to action commands for the *manipulator and three objects* target platform. The action sequence consists of a continuous (closed) path, with



Figure 7: An example of corresponding action commands for the three robots as objects imitator platform (left) and the resulting imitative behaviour (right). Using the critical points shown in Figure 2, starting from the initial positions shown in Figure 3, and minimizing the absolute displacement (red object), relative displacement (green) and relative position (blue) effect metrics, each of the robots must move along the way-points shown (left). The initial (dotted outline) and final (solid outline) positions are shown as circles, indicating that the orientation of the robots is not considered (the actual robots are square, but of equivalent size). Each way-point is indicated as a dot. The robots then perform an imitative behaviour (in Webots) and the captured results from the simulation are shown in the right plot.

way-points above the current (and future) positions of the objects. When the manipulator is above an object that must be moved, the manipulator will pick it up, then move (together with the object) to the target position and place the object down (while also, if required, rotating it), before continuing to the next object. To match the effects at each critical point, the order the manipulator approaches the objects is the same (red object, green, blue). If no displacement or rotation is required for an object during each of these turns, that object is ignored, simplifying the manipulator's path. Figure 8 (right) shows the resulting captured imitative behaviour.

3 Conclusions and Discussion.

The experiments shown in Figures 7 and 8 illustrate the diverse character of different successful imitative behaviours optimized to match particular aspects of the effects of demonstrated human manipulation of objects. Aspects captured by metrics for *absolute displacement, relative displacement, relative position, rotation* and *orientation* could all successfully be matched. The results illustrate the multiplatform targetability of the JABBERWOCKY system



Figure 8: An example of corresponding action commands for the manipulator and three objects imitator platform (left) and the resulting imitative behaviour (right). Using the critical points shown in Figure 2, starting from the initial positions shown in Figure 3, and minimizing the absolute displacement (red object), relative displacement (green) and relative position (blue) effect metrics, the manipulator must follow the continuous closed path (starting and ending at the left top corner of the workspace) shown as a dotted line (left). Since the human demonstrator did not rotate the objects, no angular effect metrics were used. The line in drawn using a gray to black color gradient to indicate the direction of the path. When reaching an object, the orientation that the object must be rotated to is shown by a small arrow. The manipulator then performs an imitative behaviour (in Webots) and the captured results from the simulation are shown in the right plot.

to map human demonstrated manipulations to matching robotics manipulations (in simulation), generalizing to different initial object configurations.

From the examples shown it becomes apparent that the relative/absolute position and rotation of objects are important aspects of a demonstrated task to match (or not) according to effect metrics, depending on the state of the objects in the environment and the context. The exploratory characterization of the space of effect metrics reveals that matching of "results" is a more sophisticated issue that generally acknowledged. This wide range of possible effect metrics illustrates that even the effect aspect of the correspondence problem for human-robot interaction by itself is already quite complex. Goal extraction in terms of effect metrics and granularity may have many different solutions that might not all be appropriate according to the desired results or context. Depending on the constraints of the imitator embodiment, a 'many-to-one' or 'one-to-many' correspondence between imitator and model sub-goals may be required for specific parts of the task. It is also possible that an

imitating agent has to switch metrics and granularity during the imitation attempt. This has not been emphasized at all in the literature so far (but see Alissandrakis et al. (2004)). This creates particular problems and challenges for sub-goal and metric extraction systems that can be used in programming robots by demonstration. The use of repeated demonstrations (Billard et al. (2004)), saliency detection (Scassellati (1999)) and goal-marking via deixis and non-verbal signaling by humans (Butterworth (2003); Call and Carpenter (2002); Bekkering and Prinz (2002)) may help contribute solutions to these problems. Other research questions yet to be addressed include the importance of order effects in manipulation and establishing object-object correspondence.

Acknowledgements

The work described in this paper was conducted within the EU Integrated Project COGNIRON ("The Cognitive Robot Companion") and funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020.

References

- A. Alissandrakis. Imitation and Solving the Correspondence Problem for Dissimilar Embodiments - A Generic Framework. PhD thesis, University of Hertfordshire, 2003.
- A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Imitation with ALICE: Learning to imitate corresponding actions across dissimilar embodiments. *IEEE Trans. Systems, Man & Cybernetics: Part A*, 32(4):482–496, 2002.
- A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Towards robot cultures? – Learning to imitate in a robotic arm test-bed with dissimilar embodied agents. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 5(1):3–44, 2004.
- Harold Bekkering and Wolfgang Prinz. Goal representations in imitative actions. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 555–572. MIT Press, 2002.
- A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal. Discovering optimal imitation strategies. *Robotics and Autonomous Systems*, 47:2-3, 2004.
- George Butterworth. Pointing is the royal road to language for babies. In Sotaro Kita, editor, *Pointing: Where Language, Culture, and Cognition Meet*, pages 9–26. Lawrence Erlbaum Assoc Inc, 2003.

- S. Calinon and A. Billard. Stochastic gesture production and recognition model for a humanoid robot. In *IEEE/RSJ Intl Conference on Intelligent Robots and Systems (IROS)*, 2004.
- Josep Call and Malinda Carpenter. Three sources of information in social learning. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- K. Dautenhahn and C. L. Nehaniv. An agent-based perspective on imitation. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 1–40. MIT Press, 2002.
- Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observations of human performance. *IEEE Trans. Robot. Automat.*, 10:799–822, November 1994.
- C. L. Nehaniv. Nine billion correspondence problems and some methods for solving them. In *Proc. Second International Symposium on Imitation in Animals and Artifacts – Aberystwyth, Wales, 7-11 April 2003*, pages 93– 95. Society for the Study of Artificial Intelligence and Simulation of Behaviour, 2003.
- C. L. Nehaniv and K. Dautenhahn. Mapping between dissimilar bodies: Affordances and the algebraic foundations of imitation. In John Demiris and Andreas Birk, editors, *Proceedings European Workshop on Learning Robots 1998 (EWLR-7), Edinburgh, 20 July 1998*, pages 64–72, 1998.
- C. L. Nehaniv and K. Dautenhahn. Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications. In J. Demiris and A. Birk, editors, *Interdisciplinary Approaches to Robot Learning*, pages 136–161. World Scientific Series in Robotics and Intelligent Systems, 2000.
- C. L. Nehaniv and K. Dautenhahn. Like me? measures of correspondence and imitation. *Cybernetics and Systems*, 32(1-2):11–51, 2001.
- C. L. Nehaniv and K. Dautenhahn. The correspondence problem. In K. Dautenhahn and C. L. Nehaniv, editors, *Imitation in Animals and Artifacts*, pages 41–61. MIT Press, 2002.
- Brian Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills. In C. L. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 176–195. Springer Lecture Notes in Artificial Intelligence, Volume 1562, 1999.

Notes

¹ Given $C\hat{A}B_M$, $A\hat{B}C_M$, $B\hat{C}A_M$ and \overline{BC}_I , we can find the other two sides $\overline{AC}_I = \sqrt{\frac{(1-\cos^2(A\hat{B}C_M))\times \overline{BC}_I^2}{(1-\cos^2(C\hat{A}B_M))}}$ and $\overline{AB}_I = \sqrt{\frac{(1-\cos^2(A\hat{B}C_M))\times \overline{BC}_I^2}{(1-\cos^2(C\hat{A}B_M))}}$

 $\sqrt{\frac{(1-\cos^2(B\hat{C}A_M))\times\overline{BC}_I^2}{(1-\cos^2(C\hat{A}B_M))}}, \text{ to satisfy the equalities } C\hat{A}B_I = C\hat{A}B_M, A\hat{B}C_I = A\hat{B}C_M \text{ and } B\hat{C}A_I = B\hat{C}A_M.$ Assuming that side \overline{BC}_I lies on the $(0, +\infty)$ x-axis with points $\mathcal{B} = \begin{bmatrix} 0\\0 \end{bmatrix}$ and $\mathcal{C} = \begin{bmatrix} |\overline{BC}_I|\\0 \end{bmatrix}$ corresponding to B_I and C_I , we can then find a point $\mathcal{A} = \begin{bmatrix} \frac{\overline{BC}_I^2 - \overline{AC}_I^2 + \overline{AB}_I^2}{2\times\overline{BC}_I}\\ \sqrt{(-\overline{BC}_I + \overline{AC}_I - \overline{AB}_I)\times(-\overline{BC}_I - \overline{AC}_I + \overline{AB}_I)\times(-\overline{BC}_I + \overline{AC}_I + \overline{AB}_I)\times(\overline{BC}_I + \overline{AC}_I + \overline{AB}_I)}\\ \sqrt{(-\overline{BC}_I + \overline{AC}_I - \overline{AB}_I)\times(-\overline{BC}_I - \overline{AC}_I + \overline{AB}_I)\times(-\overline{BC}_I - \overline{AC}_I + \overline{AB}_I)\times(\overline{BC}_I + \overline{AC}_I + \overline{AB}_I)} \end{bmatrix}$ corresponding to D_I corresponding to

 A_I , such that the equalities $\overline{\mathcal{AB}} = \overline{\mathcal{AB}}_I$, $\overline{\mathcal{BC}} = \overline{\mathcal{BC}}_I$ and $\overline{\mathcal{CA}} = \overline{\mathcal{CA}}_I$ To find A_I we need to rotate and translate \mathcal{A} in respect to the actual co-ordinates of $B_I = \begin{bmatrix} x_B \\ y_B \end{bmatrix}$ and $C_I = \begin{bmatrix} x_C \\ y_C \end{bmatrix}$ in the imitator's workspace: $A = \begin{bmatrix} x_A \\ y_A \end{bmatrix} = \begin{bmatrix} \cos\phi & \sin\phi \\ -\sin\phi & \cos\phi \end{bmatrix} \times \mathcal{A} + \begin{bmatrix} x_B \\ y_B \end{bmatrix}$, where $\phi = \tan^{-1}\left(\frac{y_C - y_B}{x_C - x_B}\right)$.

Acquisition of a "mirror" system for speechreading

Luc Berthouze*

*Neuroscience Research Institute (AIST) Tsukuba Central 2, Umezono 1-1-1, Tsukuba 305-8568, Japan Luc.Berthouze@aist.go.jp

Abstract

Articulatory mimicry is a spontaneous feature in both deaf and hearing infants. We discuss the role of this activity in the perception of visual speech, and speculate on how it shapes the underlying neural circuitry. We argue that in the early stages, speechreading involves an active phase of selection and sequencing of motor plans corresponding to representations of visible articulators acquired during articulatory mimicry. This sequencing activity results in activation of lateral and medial premotor areas (BA6) which we observed in our fMRI study of speechreading in naive subjects. As the repertoire of visual-motor associations expands, the automatic recognition of the visual stimulus (and the retrieval of the corresponding motor plan) becomes possible, consistent with the activation of the left inferior frontal gyrus (putative locus of the human mirror system) reported in studies of speechreading of trained stimuli. We conclude by outlining a computational model, and reporting on simple experiments of deferred head imitation.

1 Introduction

speechreading is the ability to perceive speech by (a) watching the movement of a speaker's mouth, (b) observing all other visible cues including facial expressions and gestures, (c) using the context of the message and the situation, and (d) exploiting the knowledge of the speaker's particular ways to articulate. In the deaf, speechreading is not formally taught but naturally occurs as a result of exposure to hearing teachers and/or parents. The deaf infant is not required to articulate, as long as it properly comprehends speech. In the oralist tradition, speechreading is the primary means of communication. In the hearing-impaired, or in Prillwitz's holistic view of deaf education, it augments communication. In the hearing, speechreading also occurs, as evidenced by the McGurk effect (McGurk and MacDonald, 1976) and recent neuroimaging studies on visible speech (e.g. Calvert and Campbell, 2003).

Our particular interest in speechreading stems from the fact that articulatory mimicry is a spontaneous feature in both deaf and hearing infants, even though the lack of auditory feedback in the deaf would suggest that an alternative route would be used. The discovery of mirror neurons (Gallese et al., 1996), which are now commonly seen as providing a link between language and gestures, seem to offer a reasonable explanation, especially in light of the revised theory of speech of Liberman and Mattingly (1985). Yet, it remains to be seen whether articulatory mimicry can be explained by mirror neurons as found in the monkey. In its early stages, i.e., before it becomes a linguistic competence, articulatory mimicry shares a lot with facial imitation. The infant can neither see nor hear the consequences of its own facial movements, not can it feel the muscle activities of the faces it imitates (Studdert-Kennedy, 2002). Yet, at least in the monkey, mirror neurons only code actions that are already known to the animal, and those neurons do not seem adapted to serve imitation of new, never seen, never executed actions (Fadiga, personal communication). Thus, we are left with the question of whether this early imitation involves a different circuitry (e.g., very low-level matching), or a specialized mirror system - a reasonable assumption from an evolutionary perspective.

Meltzoff and Moore (1997), who reported imitation of facial gestures within hours of birth, proposed an inter-modal matching mechanism (AIM). This mechanism translates visual perceived stimuli from an external coordinate frame of reference to a viewer-centered representation that can be used along with the viewer's proprioceptive state to drive the matching process. The fact that, for apical segments (when the utterance has the most characteristic and distinctive phonological structure), visible image properties can often be sufficient to identify speech sounds (Calvert and Campbell, 2003) would support such extension. However, the number of such visible articulators is, in reality, very limited. Over two-thirds of English speech sounds, for example, are either invisible or visually indistinguishable from one another, and skilled speech-readers barely score 25% in separate phoneme visual recognition. Finally, two critical differences between facial imitation and articulatory imitation make the idea of a simple extension of the AIM mechanism less likely. First, self-produced articulations play a preponderant role in what is being mimicked. Wihman (2002), for example, reported that "the experience of frequently produced CV [consonant-vowel] syllables sensitizes infants to similar patterns in the input speech stream", and various studies pointed that children choose words that match their available articulatory routines (Studdert-Kennedy, 2002). Secondly, articulatory activity follows a developmental trajectory, from pre-linguistic mouthing to purposive phonetic act. This, in turns, involves a transition from recognizing discrete patterns (elementary gestures, or movement of speech articulators) to recognizing continuous patterns (coordinative structure of gestures). Indeed, the utterance of words requires an accurate timing of each gesture itself and accurate phasing of gestures with respect to one other (Studdert-Kennedy, 2002).

2 Working hypothesis and relation with existing studies

The above discussion leads us to formulate the following three design principles to model the role of articulatory mimicry as a precursor to speechreading:

- (a) Acquisition of a basic repertoire of face action pattern mappings from motor babbling and contingent imitation by the caregiver. Such developmental mechanism has already been suggested by Yoshikawa et al. (2003) for modeling infants' acquisition of vowels. Since infants are observed to execute consonants with their precise, categorical loci of constriction more accurately than the less precise, continuously variable vowels (Studdert-Kennedy, 2002), those articulations are more likely to elicit a response by the caregiver and, as a result, specific articulatory patterns can be mapped to visible mouth movements (or visemes).
- (b) Words (coordinated structures of articulations) are then defined as continuous trajectories in the

(discrete) viseme space. An obvious corollary of this definition is that, since confusion occurs often between consonants belonging to a viseme, there are multiple possible readings of a single utterance. This is plausible. In skilled speechreaders, word recognition is only made possible by context modulation (given a sufficient linguistic competence).

(c) speechreading proceeds from a motor simulation process. The idea of a motor simulation process is not novel per se (e.g., Demiris, 2002; see also Miall, 2003 for review). However, in existing models, the rehearsal of motor plans relies on the ability of the forward controllers to predict the next state of the system given a motor command. In our context, however, such forward controllers are not necessarily available (at least, not in the initial stages) and a generative process is therefore necessary. Thus, we hypothesize that perception of novel visual speech involves an active phase of generation, selection and sequencing of actions, biased by already acquired visual-motor associations. As such, our proposal has conceptual similarities with the ASL (Associative Sequence Learning) hypothesis of Heyes (2001), in particular, the idea that the mechanism is highly experiencedependent, and that it involves bidirectional excitatory links between sensory and motor representations of movement units rather than an innate supramodal mechanism.

This last hypothesis leads to the prediction that, at least in the early stages, speechreading should activate premotor areas typically involved in motor response selection and sequencing, rather than Broca's area (the putative locus of mirror neurons in the human brain). Existing studies, however, do not show such pattern. Campbell et al. (2001), for example, reported "extensive activation in posterior-inferior regions, bilaterally. These included the middle occipital and fusiform gyri and the posterior part of the inferior temporal gyrus. Further activation was evident in the superior temporal gyrus, with large clusters of activation showing peak foci in the STS bilaterally, and in the inferior frontal gyrus, more extensively in the left than right hemisphere." Noting that in that study (and others), subjects were shown examples of the stimuli prior to scanning so that they were able to speechread the stimuli with high accuracy, we carried out our own study (Berthouze et al., 2004) in which we measured cortical activity during speechreading of novel stimuli in subjects with no prior formal experience in speechreading.

3 Brief summary of fMRI study

Eighteen Japanese hearing right-handed university students (10 males and 8 females, aged 20-29) participated after providing informed consent according to AIST safety and ethics guidelines. The subjects were not exposed to the stimuli before scanning and were not informed of the nature of the stimuli (i.e., language used). The subjects were instructed to covertly speechread (silent reading) 40 muted video-clips of isolated low-frequency / low-visibility Japanese and English words articulated by a Japanese face.

First-level analysis (fixed effects) showed activities related to baseline to overlap largely with results obtained in other speechreading studies. However, significant (corrected for family-wise error, p < .05) activation of the right-hemisphere BA6 area (middle frontal gyrus, ke=219, x=57, y=8, z=38) was also observed, which was not previously reported. This activity was also significant (corrected for falsediscovery rate, p < .05) in the random effects analysis (i.e., inference at the population level).

With respect to our hypothesis, this result is significant because covert speech has been widely shown to elicit a rather exclusive left lateralization of the precentral gyrus activation (Wildgruber et al., 1996; see also Dogil et al., 2002; Nixon et al., 2004). Thus, if this right-hemisphere activation is not accounted for by covert speech, it may then be related to our hypothesized motor sequencing activity. In fact, studies on motor sequence learning actually support that view. Rushworth et al. (1998), for example, provided evidence that the lateral premotor cortex is concerned with the learning of both sequences of sensory guided movement responses and with the learning of single responses instructed by arbitrary sensory cues.

4 Outline of model and results

To validate our hypothesis, and provide a platform with which to make further predictions, we constructed an integrated experimental system that implements the specifics of speechreading. The system consists of three major modules modeling the critical components of the hypothesized speechreading circuitry:

Motor apparatus A three-dimensional facial simulator was implemented ¹ that can produce articulatory sequences visually consistent with those produced by a human speaker. The simulator's smooth skin surface is supported by a threedimensional wireframe structure (see Figure 1, left) and 18 muscles organized according to the FACS (Facial Action Coding System) of Ekman and Friesen (1977). The contractions or relaxations of each muscle result in a motion field in the skin structure, including the lips. A jaw mechanism enables the mouthing actions needed for articulation, with a specific mouth muscle (sphincter) controlling the roundness of the lips. Although the facial simulator currently lacks a tongue, it was successfully used to implement 10 out of 13 visible articulators (visemes) common to both English and Japanese language (see Figure 1, right for example).



Figure 1: (Left) Skin surface of the facial simulator and the underlying musculoskeletal structure. (Right) Appropriate control synergies between jaw articulation, mouth sphincter, and facial muscles can implement visible articulators such as e/a (top) and w/r (bottom).

Visual apparatus The visual apparatus consists of a distributed network of feature detectors that respond selectively to apical segments of articulations. Each detector is implemented in the form of a cascade of boosted classifiers working with haar-like features (see Figure 2) according to Lienhart and Maydt (2002)'s method. These detectors are trained with sample views of a particular object (e.g., the mouth) called positive examples, that are scaled to the same size, and negative examples, arbitrary images of the same size. They are designed so that they can be easily "resized" in order to find objects of interest at different sizes. In the experiment described

¹The simulator is an extension of the facial simulator developed

at Imperial College under the supervision of Y. Demiris.

in this paper, face orientation selective detectors were constructed using the same principle.



Figure 2: Haar-like features for fast object recognition, from Lienhart and Maydt (2002)

Sequence learning module This module consists of sequence learning networks that seamlessly combine the learning and the prediction of arbitrary sequences of patterns into a single generative process (Berthouze and Tijsseling, in review). The sequence learning neural network (see Figure 3) was constructed according to design principles derived from neuroscience and existing work on recurrent network models. It utilizes sigmoid-pulse generating spiking neurons to extract timing information from the input stream and modifies its weights using an adaptive learning rule with synaptic noise. Combined with coincidence detection and an internal feedback mechanism, it implements a learning process that is driven by dynamic adjustments of the learning rate. This gives the network the ability to not only adjust incorrectly recalled parts of a sequence but also to reinforce and stabilize the recall of previously acquired sequences. Separate instances of these networks are used to encode visual input stream (a time-series in the viseme space), and articulatory motor sequences (the motor patterns required to implement a given viseme). Hebbian learning is used to establish connections between visual and motor networks so that resonant coupling can be achieved through correlated experience of observation and execution (Heyes et al., in press).

At this stage of the project, the integration of the three components was only tested on a simplified task: the deferred imitation of complex sequences of head panning movements. Although that particular task has



Figure 3: The input layer is a placeholder for each pattern in a presented sequence, while the context layer receives both external contextual information as well as feedback information from the predicted context module (i.e. the context that the network has learned to associate with the current sequence). Input and context information as well as feedback from the output module is propagated to the central module that contains a variant of spiking neurons. This module is responsible for extracting the variety in timing information from the input. All learning occurs in the connections from the central module to the output and the predicted context modules. The output module is also connected to the coincidence detector module, which regulates the learning rate by calculating the familiarity of the current output based on a history of previous states.

already been investigated, in particular using a model of inter-modal matching (Demiris et al., 1997), our focus was not on the task itself, but on the application, and validation, of the hypothesized mechanism. Five detectors were trained off-line to detect five discrete head orientations (the simplified equivalent of the apical segments of a visible articulator). The corresponding (discrete) visual-motor mappings were obtained as a result of contingent imitation (by


Figure 4: Learning curve for a novel stimulus as a function of the number of presentations. The horizontal line denotes the timing of the actual visual stimulus (i.e., optimal performance). The curve was obtained from 10 trials.



Figure 5: Relationship between perceived orientation (vertical axis) and actual orientation (horizontal axis). The blue line denotes the identity function. The red line denotes a fit by logistic regression. The effect of the discrete encoding of the head orientation (-90,-45,0,45,90) is noticeable in the acquired representation.

the human agent) of the simulator's motor babbling. During subsequent interaction, the incoming visual stimuli (head panning movements) were processed by each detector in parallel, and a predictive filter was used to filter out noise-induced errors in the set of detectors. With the system initially without any established visual-motor mapping, but the ones described above, the feeding of the time-series of detector activities to the sequence learning module resulted in the system generating head movements, with a bias on the repertoire of motor actions which had elicited contingent imitation (such a bias has been observed in young infants). As a result of learning, a continuous visual-motor mapping was acquired (see Figure 5). Successful acquisition required only a relatively low number of presentations (see Figure 4), after which the sequencing activity was reduced to a minimum.

5 Conclusions

Because the implementation of the model is still in its early stages, it is difficult to draw any conclusion as to how well the model can account for behavioral data obtained with human subjects. Nonetheless, there is supporting evidence for the three design principles used in the model. Studies showing that cells in the superior temporal sulcus (STS) are sensitive to discrete features of biological motion provide plausibility to our thesis that infants could construct detectors for apical segments of articulations. The fact that displaying such segments during perception of (time-varying) speech results in McGurk effects (Calvert and Campbell, 2003) justifies our idea that articulations are trajectories in the viseme space. This, in turn, could well explain why infants proceed from prosodic to segmental imitation. Indeed, a limited articulatory behavior of the child may result in its inability to detect continuous changes in the incoming visual patterns, and thus puts the focus on the duration (rhythm) of each discrete (visible) pattern. As the repertoire extends, segmental imitation becomes possible, through resonant coupling between external events and internally (motor-based) representations. A future focus of this research will be to investigate the origins of the differences observed in the neural substrate of speechreading in the hearing and in the deaf. Since we considered a single model to account for both deaf and hearing articulation mimicry, it will be interesting to see if the above differences can be explained by feedback modality, rather than by functional differences.

Acknowledgements

This research was funded by the Advanced and Innovational Research program in Life Sciences from the Ministry of Education, Culture, Sports, Science and Technology, the Japanese Government. The author wishes to thank Torea Foissotte for his programming contribution, and Naomi Nakagawa for her help with the human experiments. The author also thanks two anonymous reviewers for their comments.

References

- L. Berthouze, S. Phillips, O. Terasaki, and K. Kawano. Role of precentral gyrus (BA6) in speechreading: An fMRI study on hearing japanese subjects. *NeuroImage*, 22 (suppl.1):S36, 2004.
- L. Berthouze and A. Tijsseling. A neural network architecture for the categorization of temporal information. *Neural Networks*, in review.
- G.A. Calvert and R. Campbell. Reading speech from still and moving faces: The neural substrates of visible speech. *Journal of Cognitive Neuroscience*, 15(1):57–70, 2003.
- R. Campbell, M. MacSweeney, S. Surguladze, G. Calvert, P. McGuire, J. Suckling, M.J. Brammer, and A.S. David. Cortical substrates for the perception of face actions: an fmri study of the specificity of activation for seen speech and for meaningless lower-face acts (gurning). *Cognitive Brain Research*, 12:233–243, 2001.
- J. Demiris, S. Rougeaux, G. Hayes, L. Berthouze, and Y. Kuniyoshi. Deferred imitation of human head movements by an active stereo vision head. *IEEE International Workshop on Robot and Hu*man Communication (RoMan'97), Sendai, Japan, pages 88–93, Sept. 29-Oct. 1st, 1997.
- Y. Demiris. *Imitation in Animals and Artifacts*, chapter Imitation as a dual-route process featuring predictive and learning components: a biologicallyplausible computational model, pages 327–361. Cambridge, MA: MIT Press, 2002.
- G. Dogil, H. Ackermann, W. Grodd, H. Haider, H. Kamp, J. Mayer, A. Riecker, and D. Wildgruber. The speaking brain: a tutorial introduction to fmri experiments in the production of speech, prosody and syntax. *Journal of Neurolinguistics*, 15(1):59– 90, 2002.
- P. Ekman and W.V. Friesen. *Manual for the Facial Action Coding System*. Palo Alto: Consulting Psychologists Press, 1977.
- V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti. Action recognition in the premotor cortex. *Cognitive Brain Research*, 119:593–609, 1996.
- C. Heyes. Causes and consequences of imitation. *Trends in Cognitive Sciences*, 5(6):253–261, 2001.

- C. Heyes, G. Bird, H. Johnson, and P. Haggard. Experience modulates automatic imitation. *Cognitive Brain Research*, in press.
- A.M. Liberman and I.G. Mattingly. The motor theory of speech perception revised. *Cognition*, 21:1–36, 1985.
- R. Lienhart and J. Maydt. An extended set of haarlike features for rapid object detection. In *Proceedings of IEEE International Conference on Image Processing*, pages 900–903, 2002.
- H. McGurk and J. MacDonald. Hearing lips and seeing voices. *Nature*, 264:746–748, 1976.
- A.N. Meltzoff and M.K. Moore. Explaining facial imitation: A theoretical model. *Early Development* and Parenting, 61:179–192, 1997.
- R.C. Miall. Connecting mirror neurons and forward models. *Neuroreport*, 14(17):2135–2137, 2003.
- P. Nixon, J. Lazarova, I. Hodinott-Hill, P. Gouch, and R. Passingham. The inferior frontal gyrus and phonological processing: An investigation using rtms. *Journal of Cognitive Neuroscience*, 16:289– 300, 2004.
- M.F.S. Rushworth, P.D. Nixon, D.T. Wade, S. Renowden, and R.E. Passingham. The left hemisphere and the selection of learned actions. *Neuropsychologia*, 36(1):11–24, 1998.
- M. Studdert-Kennedy. *Mirror Neurons and the Evolution of Brain and Language*, chapter Mirror neurons, vocal imitation, and the evolution of particulate speech, pages 207–227. Amsterdam/Philadelphia: John Benjamins, 2002.
- M.M. Wihman. *Mirror Neurons and the Evolution* of Brain and Language, chapter The role of mirror neurons in the ontogeny of speech, pages 305–314. Amsterdam/Philadelphia: John Benjamins, 2002.
- D. Wildgruber, H. Ackermann, U. Klose, B. Kardatzki, and W. Grodd. Functional lateralization of speech production at primary motor cortex: A fMRI study. *Neuroreport*, 7:2791–2795, 1996.
- Y. Yoshikawa, M. Asada, K. Hosoda, and J. Koga. A constructivist approach to infants' vowel acquisition through mother-infant interaction. *Connection Science*, 15(4):245–258, 2003.

Using Visual Velocity Detection to Achieve Synchronization in Imitation

Arnaud J. Blanchard*

A.J.Blanchard@herts.ac.uk

*Adaptive System Research Group School of Computer Science University of Hertfordshire College Lane, Hatfield, Herts AL10 9AB Lola Cañamero†

L.Canamero@herts.ac.uk

[†]Adaptive System Research Group School of Computer Science University of Hertfordshire College Lane, Hatfield, Herts AL10 9AB

Abstract

Synchronization and coordination are important mechanisms involved in imitation and social interaction. In this paper, we study different methods to improve the reactivity of agents to changes in their environment in different coordination tasks. In a robot synchronization task, we compare the differences between using only position detection or velocity detection. We first test an existing position detection approach, and then we compare the results with those obtained using a novel method that takes advantage of visual detection of velocity. We test and discuss the applicability of these two methods in several coordination scenarios, to conclude by seeing how to combine the advantages of both methods.

1 Introduction

Synchronization and coordination are important mechanisms involved in imitation and social interaction. As put forward by psychological studies, e.g. (Hatfield et al., 1994), people often synchronize with their interaction partners using different methods, for example they synchronize their movements and rhythm. However, achieving good coordination is a very challenging problem in robotics. In this study, we take a first step to develop suitable mechanisms to this end.

In imitation and synchronization problems, the agent that is imitating (the "subject" agent) needs some inputs to know what the agent that is imitated (the "object" agent) is doing. A property often used as input information for imitation is the position of the object agent. Using position information, the subject agent can learn to reproduce or copy a trajectory. Position information can also be used to achieve synchronization—while dancing, for example.

In their studies of imitation tasks using robots, Andry et al. (2002) use the quantity of movement (temporal luminosity variation) to perceive the target position. This technique is efficient and simple as it does not need complex visual tasks such as object recognition. However, a problem with this mode of imitation in robotics is that there is always a delay between the object agent and the subject one. In fact, the subject agent can start to move only after the object agent is in a new position. Even if such delay is not always a problem when following trajectory, it usually poses a problem for synchronization tasks.

In this paper, we propose a velocity detection system to synchronize the movements of two robots avoiding the delay problem. This system is applicable not only in the case of precise reproduction of movements (e.g., when mirroring a movement) but also in cases in which imitation does not need to be precise but must be very well timed at the same rhythm, such as when dancing. Our experimental results show how this system outperforms other systems based on position detection in different synchronization tasks.

Finally, to conclude the paper, we discuss the limitations of using only velocity detection in other imitation tasks and we see how we can combine position and velocity detection to improve performance.

2 Problem Addressed

In the context of an autonomous mobile robot that has to interact with other robots in its environment, the problem that we have addressed in this study aims at achieving natural and fast, adapted reactions of the robot to changes detected in its environment. Minimizing the reaction time to respond to environmental changes is very important, in particular when the limited (perceptual and computational) resources of the agent impose severe constraints. This was made possible by our biologically plausible, bottom-up approach, following which we have adopted a minimal architecture that we have built using a neural network.

We have therefore designed an architecture to make a robot follow a target or to be synchronized with the target movement. We have developed four methods for this, two of them based on position detection and two based on velocity detection: 1) position detection with Winner-Take-All (WTA), 2) position detection without WTA, 3) velocity detection with focalization, and 4) velocity detection without focalization. We have implemented this architecture in a Hemisson robot (our "subject" robot) fitted with a video camera. The target is composed of two vertical strips or a pattern of strips drawn on a white paper attached to an object Koala robot, as shown in Fig. 1.



Figure 1: Experimental setup. On the left the Koala robot (object) moves the target observed by a Hemisson robot (subject) on the right.

2.1 Position detection

The basic principle is the one we can see in (Gaussier et al., 1998).

The area where the object is moving corresponds to the area of maximum luminosity difference. We first use a temporal smoothing in order to keep a small signal when the target stops moving for a short time. Then we use a WTA to set the position with the maximum quantity of movement among all the positions of the visual field. Once this position has been set, the subject agent only has to follow this position (method 1).

In fact with our bottom-up approach, we always try to build the system as simple as possible to realize the task and to take advantage of the side-effects that can be useful (Steels, 1994). In the present architecture, we can simplify and remove the WTA (method 2). The new resulting behavior of the robot is not the same but is still interesting: now, the subject robot reaction not only depends on the target position, but also on its contrast and activity. The problem is that the subject robot does not move if the target has a small activity whatever its position.

2.2 Velocity detection

In order to increase the reactivity of the agent to changes perceived in its environment, we put forward the idea of using velocity of the target as input information to use for synchronization. This velocity detection method, proposed by Johnston et al. (1999), is based on the hypothesis that each object's point has constant luminosity. Therefore, the luminosity variation of an image is due only to the movement of its objects. By considering v_x the velocity of one point in x, k a constant coefficient that essentially depends on the distance to the object, and i the light intensity, we use (1).

$$v_x = k \times (\partial i/\partial t)/(\partial i/\partial x)_{(x)} \tag{1}$$

Dividing the variation of luminosity $(\partial i/\partial t)$ by the contrast $(\partial i/\partial x)$ is a problem when the contrast is almost null. This is not surprising since without contrast we cannot estimate the movement of an object. To solve this problem we use a threshold for the contrast: a low value of contrast (i.e., below the threshold) will produce null velocity.

We can be interested either in focusing our attention on a small part of the visual field (method 3), or in the global velocity of the entire visual field, often due to the self movement of the robot (method 4). We can use the system of position detection to focus on the target (Fig. 2).

3 Experiments

3.1 Setup

In all the experiments, we have used the Hemisson as subject robot and the Koala as object that carries the target stimulus, and we measure the velocity order that the Hemisson would send to its wheels. Since it is impossible to know the exact position of a Hemisson robot (it has no odometer sensor), we had to design our experiments taking account of this constraint: all the computations are carried out normally to produce the motor command that the subject robot should execute to follow the target but self-motion of the robot is inhibited.



Figure 2: Architecture to detect the velocity of a focussed object (method 3). The gray part could be replaced by a large static gaussian and the architecture now only takes care of the global overview velocity (method 4). On this scheme, the curves are the result of real data.

With the first three methods, we use the same setup (see Fig. 1): the Koala robot moves right and left at a sinusoidal velocity with two vertical strips drawn on the target, while the subject Hemisson observes (without moving) the target at a distance of a floppy disk (3.5 inch). To test the last method (4) we use a very similar setup but this time the target is a wide pattern of vertical strips.

3.2 Results

We present one experimental result from a dozen with similar results in Fig. 3. The first graph shows the results of the synchronization task using position detection with WTA (method 1) and the second one without WTA (method 2). The two right graphs show the results of the synchronization task using the velocity detection, with a target's focus (method 3) on the third graph, and without focalization, but with a wide target covering all the visual field (method 4) on the last graph.

On each graph, the singularities observed over the first two iterations have no meaning. The dash line corresponds to the velocity of the object agent and the solid line corresponds to the velocity of the subject agent. Each iteration carried on for around 100 ms.

3.3 Discussion

All the methods that we have presented here have some interesting properties, depending on the task, when we want agent interactions, notably in imitation and synchronization.

The first method, which uses position detection, is very useful to follow the target trajectory. Nevertheless, the delay that it produces is not very convenient for synchronization tasks or when we have a situation that changes often. The second method, which uses a simpler version of the same principle, is suited to follow a target position even with a small embedded system (little calculus power is needed) but also for some specific behaviors.

The third method uses focalization on the object agent defined using the detection position system. The reaction is fast and proportional to the stimulus velocity since only the area of the target is considered. This is the ideal method for synchronization in dance.

The last method, which integrates each pixel's velocity without focalization, allows us to do pure synchronization. The target position does not matter and all the visual field is considered. Therefore, if the object agent is moving in the visual field, the subject agent moves in the same direction but not with a proportional velocity since the background is considered. This method is very useful when all the visual field is moving—e.g. when the camera itself is moving. We can use this to stabilize the agent's own movements, in the same way as a fly does (Holst and Mittelstaedt, 1950). We have been able to reproduce the fly phenomenon with our robot. We put the robot in a drum with black and white strips and, when we move the drum, the robot turns with the same velocity in the same direction. The robot thus stays relative to the drum at the same place.

We can see that we have two kinds of methods (position detection or velocity detection) that have advantages and disadvantages. The first category does not produce a drift but is not very reactive. The second category is very reactive but has a drift that does not permit a prolonged interaction since the target becomes lost. To drive a system it is possible to use either the position (first order) with a stable but slow system, or the velocity (second order) with a fast but unstable system. The best results are obtained by combining both methods and this leads us to think that we should do the same.



Figure 3: Results of the four methods tested: method 1 on the far left, method 4 on the far right.

4 Conclusion

We have presented different methods that allow us to increase the level of interaction (synchronizationimitation) thanks to biologically plausible processes. These processes are simple and easy to implement.

If we want to synchronize a dance, velocity detection is very useful. However, the detection of position is more useful to follow a moving target. We see also that velocity detection can help to anticipate the target tracking by anticipating. The robot could learn to anticipate the position using velocity perception for best tracking. Studies such as Hofsten and Rosander (1996) and Richards and Holley (1999) investigate how babies develop the capacity of smooth tracking with the same kind of protocol. Since we have access to the velocity and not only to the area of movement, we should be able to make the robot learn what is associated with its own movement. Hofsten and Rosander (1996) also show that babies progressively develop a better coordination between the movement of the eyes and the head. We could use this work as inspiration to reproduce this phenomenon with robots.

Further work could try to make this architecture more biologically realistic, allowing the robot to integrate or to predict the consequences of its own movement and apply this method to the synchronization and coordination problem. Therefore, we will focus our work on the learning of the perception-action mapping inspired by the psychology studies of Prinz (1997), which seem to fit well our robotics approach.

Acknowledgments

Arnaud Blanchard is supported by a research scholarship of the University of Hertfordshire. This research is partly supported by the EU FP6 Network of Excellence HUMAINE (IST-FP6-507422).

References

- P. Andry, P. Gaussier, and J. Nadel. From sensorimotor development to low-level imitation. In 2nd Intl. Wksp. on Epigenetic Robotics, march 2002.
- P. Gaussier, S. Moga, JP. Banquet, and M. Quoy. From perception-action loops to imitation processes. *Applied Artificial Intelligence*, 1(7), 1998.
- E. Hatfield, J. Cacioppo, and R. Rapson. *Emotional* contagion. Cambridge university press, 1994.
- C. von Hofsten and K. Rosander. The development of gaze control and predictive tracking in young infants. *Vision Research*, 36, 1996.
- E. von Holst and H. Mittelstaedt. Das reafferenzprinzip. wechselwirkungen zwischen zentralnervensystem und peripherie. *Naturwissenschaften*, 37, 1950.
- A. Johnston, C.P. Benton, and M.J. Morgan. Concurrent measurement of perceived speed and speed discrimination threshold using the method of single stimuli. *Vision Research*, 39, 1999.
- W. Prinz. Perception and action planning. *European journal of cognitive psychology*, 9(2), 1997.
- J.E. Richards and F.B. Holley. Infant attention and the development of smooth pursuit tracking. *Developmental Psychology*, 35, 1999.
- L. Steels. Mathematical analysis of behavior systems. In IEEE Computer Society Press, editor, *From Perception to Action Conference*, Lausanne, Switzerland, 1994.

Learning Discretely: Behaviour and Organisation in Social Learning

Joanna J. Bryson

Mark A. Wood

*Artificial models of natural Intelligence (AmonI) University of Bath, Department of Computer Science Bath, BA2 7AY, UK {jjb,cspmaw}@cs.bath.ac.uk

Abstract

This paper describes what is required to learn new tasks in general, then applies this knowledge to understanding imitation learning in specific. We make some reference to the neurological literature, including dual-speed hippocampal / neo-cortical learning systems. We suggest that this model solves the problem of discrete replicants in memetics. We also describe some very preliminary work in implementing and testing our ideas through social learning in a computer game context.

1 Introduction

Human-like intelligence requires an enormous amount of knowledge --- solutions to the hard problems of survival and reproduction, which for our species have come to involve complex social and technological manipulations. Some of these solutions are passed to us genetically, and some are learned by an individual during their lifetime through trial-and-error experience. For humans, one key source of knowledge is culture. By culture here we mean any knowledge an agent has derived from conspecifics by non-genetic means. In order for such knowledge to be acquired efficiently, the process of acquiring it must be significantly less time consuming (at least for the individual) than individual trial-and-error learning.

In this paper we discuss first how such learning may be accumulated socially by a culture, and then relate this to what we know about learning in individuals. We propose a model for task learning in general, which is clearly facilitated by social information. We then briefly describe our preliminary attempts to build and exploit such a model of learning.

2 Discretion in Memetics

Dawkins (1976b) proposes that knowledge and behaviour can be viewed as developing through a process of evolution, just as biological life has. Ideas or behaviours are propagated if they survive intact long enough to be reproduced. Reproductive success requires replication beyond a single host behaving agent. While some behaviours are known explicitly and transmitted deliberately (by teaching), there is evidence that our species may have evolved the ability to take advantage of this powerful mechanism for increasing knowledge and fitness before we were capable of such explicit mechanisms, and that indeed we still implicitly learn complex multimodal behaviours from our conspecifics. This allows us to build and transmit knowledge that our cultures have not yet developed words or theories to describe or deliberately represent. This theory of cumulative knowledge generation is called *memetics*.

Dawkins (2000) describes a fundamental problem with the theory of memetics. Memetics is based on the concept of a *meme* which is meant to be analogous to a gene. Some theorists have claimed that this analogy is invalid, on the grounds that genes are discrete, but memes are not. This claim is itself suspect, since to this day the term *gene* still does not describe a well-defined entity (Dennett, 2002), but is based on the fact that the DNA molecule ultimately encodes information in terms of discrete patterns of four possible chemical chains.

The underlying representation for a meme, though still completely unknown, is suspected not to be discrete, and therefore to be open to corruption. To describe the problem, Dawkins (2000) proposes a thought experiment where a child is shown a drawing of an unfamiliar type of boat and asked to copy it; then the process is repeated with another child who sees only the new drawing. Dawkins believes the boat would rapidly become as unrecognisable as a phrase whispered by children playing a game of 'telephone' [Chinese Whispers U.K.]. Dawkins proposes a solution to this problem, which is that one learns not gross behaviours, but instructions as to how to behave. He proposes an alternate thought experiment, whereby children learn to build a boat by origami, an art based on folding paper. Here small mistakes produced by one child will be corrected by the next, because the second child is able to deduce the intention of the first (or of the designer) because they understand the nature of the operations. In other words, because a process of origami consists of a relatively short list of well-defined operations, Dawkins claims it can be replicated more robustly than a process of drawing.

We believe that Dawkins' requirement that memes must be instructions is over-specific, though correct in principle. We think individuals learn in terms of *skills*, not instructions. There are two differences:

- skills are not necessarily known or communicated explicitly¹, and
- skills are developed by the individual, and thus open to individual variation.

This hypothesis has several interesting ramifications, mostly having to do with the consequences of having variations of granularity in memetic representation. For example, consider some teacher J who starts with relatively few mathematical skills, but has by a slow laborious process managed to learn a technique for writing back-propagation networks. Her representation might be a long string of relatively simple arithmetic and trigonometric operators. If she has a student, M, with more mathematical skills (for example, calculus), and he observed her coding a network, he might be able to form a new representation which would create exactly the same sort of system. But M's representation of the system would be quite different from J's, consisting of a smaller number of larger-granularity operators. Note too that the situation could be reversed — if J only knows trigonometry but observes M coding an algorithm, she might well be able to imitate the algorithm herself, however again she would perceive and remember the algorithm at a different level of granularity than that with which M was generating the code.

This sort of model could explain the results of Whiten (2000). Whiten presents various species of primates (including children) with complicated puzzle boxes which require one of a number of sequences of actions to get open. Subjects are generally able to open these boxes if they have first observed a demonstrator, but they will not necessarily go through all the same steps in the same order as the demonstrator. However, if the demonstrator demonstrates repeatedly, on the second or third try the subjects will often perfectly replicate the demonstrator's model, at least in terms of the sequence of affordances used. Subjects may still choose to pull out a pin using their teeth rather than their fingers, for example.

Our explanation would be that initially the subjects are imitating only the goal and perhaps some other simple attributes of the solution (e.g. knowing which knobs on the box need attending to.) However, as they develop skills by opening the box themselves, the difficulty of performing a perfect replication is reduced, because it becomes a relatively short sequence of relatively large-grain actions rather than a long sequence of basic motor commands.

3 Learning in Brains

The hypothesis described above ties in neatly to another hypothesis in learning — this one about how brains can learn from experience.

There are two ways to learn from experience. First, we can learn very slowly, taking a large number of examples to build up a model of how the world seems to be working, or at least what the right thing is to do in a particular context. The second way is to learn very quickly. The problem with learning very quickly is that we may be overly influenced by a very improbable event, taking it to mean more than it should. Learning from a large number of experiences very quickly / perfectly also runs the risk of over-fitting. General-purpose knowledge is usually considered to derive from compiling large amounts of knowledge into a few general rules or policies (Mitchell, 1997, for a summary), although in some relatively deterministic domains it can be derived by extrapolating over a set of exemplars (Poggio, 1990; Atkeson et al., 1997).

Generally speaking, our skills seem to be built up slowly through practise over time. But any such slowlearning system that builds its knowledge from experience faces a problem. The problem is, experience happens quickly. Consequently, what is needed is a second, quick system for jotting down salient events as they happen. McClelland et al. (1995) build a model of such a system, and using the neuroscience literature, tie down their model to particular regions

¹Though quite probably Dawkins didn't mean to limit memetics to explicit knowledge and was using the term *instruction* in some kind of loose computational metaphor.

of the brain. Slow learning, they say, happens in the neocortex — fast learning happens in the hippocampus (see also Treves and Rolls, 1994)

Another problem with fast learning is that it requires learning a large number of things - particularly if the system needs to hold each learned thing around long enough to allow a slow-learning system to process it. If two different things are learned that happen to be similarly indexed (by whatever category mechanism has emerged in a largely unsupervised system), they might overwrite each other. If accommodation of new information is not done systematically (which is generally seen as the purpose of a slow learning system (McClelland et al., 1995; Mitchell, 1997)) there's no reason to expect two such similarly-indexed events to be neatly, compatibly catalogued together. One way to reduce the probability of such 'collisions' (information about multiple events overwritten into the same locations) is to make sure that the information is encoded in a very sparse way. That is, to use relatively few changes in memory in order to represent the full event. And indeed, this seems to be what the hippocampus does (Rolls, 1996).

In order for a few changes to represent a complex event, each change must be highly salient ---it must represent a relatively broad chunk of semantics, a complex concept. As McClelland et al. (1995) point out, this strategy is very compatible with the hippocampal memory indexing theory (Teyler and Discenna, 1986). However, that theory was originally motivated as the use of the hippocampus for a compact, almost symbolic type of representation that would be useful for certain kinds of complex processing. For example, animals without a hippocampus can learn a new map, but they can't learn how to learn a map if they've never learned one before (Bannerman et al., 1995). Similarly, animals without a hippocampus can learn to associate actions with stimuli, but they can't learn to prioritise these actions (Alvarado and Bachevalier, 2000; Wood et al., 2004; Buckmaster et al., 2004). Whichever purpose might have originally driven the evolution of a hippocampus, the sparse representation is clearly useful enough to be necessary for at least some sorts of long-term memory storage (Squire et al., 2001), though it's possible that another similar region, perhaps the entorhinal cortex, performs some of the quick-learning roles that McClelland et al. propose for the hippocampus.

We believe that this indexical learning may be based on dynamic categories. That is, the representation of a newly observed behaviour is determined by the 'granularity' of the indexing in the fast-learning system, which is in turn driven by a set of skills learned or formed in the slower learning system. We already know that representations in the hippocampus are highly dynamic and vary by context (Wiener, 1996; Kobayashi et al., 1997). And clearly learned experience is itself a form of context. Thus the hypothesis that what (and how) we can learn with this system changes over time and experience is not excessively radical, although it does have interesting implications for the veracity of recall.

4 A Model of Task Learning



weights for percept-action pairs

Figure 1: Task learning requires learning four types of things: relevant categories of actions, relevant categories of perceptual contexts, associations between these, and a prioritized ordering of the pairings. Assuming there is no more than one action per perceptual class, ordering the perceptual classes is sufficient to order the pairs. See text for details.

In short, we believe there are at least *four* separate types of things that are learned in the process of learning a task (see Figure 1):

- 1. *perceptual classes:* What contexts are relevant to selecting appropriate actions.
- 2. *salient actions:* What sort of actions are likely to solve a problem.

- 3. *perception/action pairings:* Which actions are appropriate in which salient contexts.
- 4. *ordering of pairings:* It is possible that more than one salient perceptual class is present at the same time. In this case, an agent needs to know which one is most important to attend to in order to select the next appropriate action.

With respect to perception/action pairings, our current work indicates that there should only be one action possible per salient perceptual context, but there may be many perceptual contexts in which a particular action may be relevant, particularly if the object of the action is coded diectically (Wood et al., 2004; Bryson and Leong, 2005). Also not that although we mention perceptual contexts, we obviously do not mean the full context of all sensory information from a moment in time. Such a representation leads to overfitting / failure to generalize, besides generally being computationally intractable to process. Rather, detailed perception at any particular moment tends to be focussed on a few salient cues which will hopefully help disambiguate the current action-selection problem (Rensink, 2000).

Researchers familiar with Behaviour-Based AI may think of these four sub-problems in a different way. The first three items contribute to forming behaviour modules --- tight couplings of perception and action, while the last contributes to forming behaviour arbitration (Bryson, 2000a; Bryson and Stein, 2001). Researchers familiar with Cognitive Modelling may realise that what we describe is quite similar to ACT-R (Anderson and Matessa, 1998) except with extra emphasis on the forming of categories for sensing and action. However, ACT-R has a relatively simplistic ordering system which cannot account for all animal data on even relatively constrained tasks (Wood et al., 2004). ACT-R learns relatively simple 'utility values' for each perception/action paring, but complex tasks may require hierarchy and/or some other powerful sequencelearning representation such as POMDPs (Kaelbling et al., 1998; Bryson, 2000b).

Clearly solving four problems simultaneously makes learning new skills a very hard problem, but equally it motivates social learning. In a social context, sensing and action categories can be recognised by their co-occurrence (Roy, 1999). In all probability, sequential and hierarchical ordering may also be induced (Dawkins, 1976a).

5 Learning in Practice

In previous work we have shown successful models of solitary primate (including human) task learning where the salient actions and perceptions were already fixed, but the pairings between actions and perceptions and the prioritizations between these varied (Bryson, 2005; Bryson and Leong, 2005). We have also shown that one can create a complete set of possible mappings between perceptual and action classes and then simply prioritize all of these, since only the highest priority item for any perceptual category will be chosen (Wood et al., 2004).

In our current work, we are looking at the role of social learning in perceptual category formation. We are also hoping to explore more complex hierarchical representations. A complete agent needs to be able to move between many different tasks, and indeed determining when one is in a new task context is clearly a part of the problem for determining salient actions and perceptions.

5.1 A Working Model

As in our previous work, we are again not attempting to learn all four categories simultaneously. We have made the following simplifications / assumptions in our preliminary experiments:

- The imitator is initially able to recognise some actions that are key to learning the task.
- Only *one* perceptual class applies to the imitator at any one time.

The second assumption means that, for the time being, we are not worrying about learning priortizations, but merely perceptual classes and their pairings to actions.

Our perceptual classes are defined by boundaries in *n*-dimensional sensor space (*n* is the number of sensors providing a reading at any given time). Thus far we have kept *n*, and the *n* operative sensors themselves, constant throughout, although having different sets of sensors operating in parallel is one possible way of introducing parallel perceptual classes. The cardinality of each dimension of sensor space differs depending upon the sensor type. For example, a sensor which measures the presence of an object would return a discrete reading $\in \{true, false\}$, whereas a sensor which measures distance would return a continuous reading $\in \mathbb{R}^+$.

The actions we have made recognisable by the imitator are simply discrete. In some sense, actions *must* be discrete (e.g. in categories like turn, move and shout), but they could also be defined by parameters. These in turn can be either absolute (turnTo *north*) or deictic (turnTo *nearest_actor*) discrete values, but they can also be in terms of continuous values (turn 42.6°).

Since exactly one perceptual class applies to any given context, and only one action should be associated with any given perceptual class, the problem of perception/action pairing is in this case equivalent to partitioning sensor space and mapping each partition to an action. Given that there is no need for the prioritization of these pairings, this completes our simplified version of the model.

5.2 Domain

For both this initial exploration of perception/action map generation, and future more complex studies, we are carrying out experiments in the domain of virtualreality computer games. VR games are an excellent platform for experiments involving learning from human subjects because they are real-time, provide a common sensing and action framework for both artificial and human agents, and require many elements of human and animal intelligence, including navigation, reacting to complex, dynamic environments, planning and cooperation (Laird and van Lent, 2001).

We are currently working with two games: *Robocode* (Nelson, 2002) and *Unreal Tournament* (Digital Extremes, 1999).

Robocode is designed to be both a game and a Java teaching tool, provided for free download from IBM alphaWorks. Users have no direct control over their agents, but must provide Java code to drive them. The agents themselves are robotic tanks armed with a single cannon, a few basic sensors, and enough action commands to navigate the map and 'interact' with the other agents therein. The map is a simple 2-D rectangle surrounded by walls, without any obstacles that are not opponents.

Unreal Tournament (UT) is a commercially released, multi-player 'First Person Shooter'. As the term suggests, the user has an agent's-eye view of the game and direct, real-time control of an avatar's actions. UT also supports remote control of agents by sending commands to the game server over a network. This provides a framework for allowing external programs to direct an agents' actions. Such AI-controlled agents are commonly known as 'bots' in the literature and gaming community. The game server, in turn, sends two categories of sensor data back to the client. The first is synchronous: at regular intervals the client is informed of the agent's status (e.g. level of health, amount of ammunition, currently wielded weapon, etc). The other is asynchronous: for example whenever a wall is bumped, a footstep is heard or damage is taken.

5.3 Preliminary Experiments

Our earliest social experiments were conducted in Robocode, because we believed it would be simpler since it was two dimensional (2D) and came with pre-coded sample opponents. However, many aspects of Robocode control and sensing proved inaccessible, presumably to keep competitors from 'cheating' by affecting the code of other robots. Subsequently we have switched to a simple, effectively 2D Unreal Tournament environment.

Our work in UT is still in early stages, but we have had agents successfully learn simple plans from observation. In addition to providing a basic proof of concept, these experiments also point to representational issues which lie ahead. These will be discussed below.



Figure 2: The experimental arena

The experiment consisted of two actors (bots) moving within a single cuboid room (see Figure 2). World co-ordinates are given in three dimensions by the game engine, but since the bots only moved on the floor plane, the problem is well-defined in two dimensions. Similarly, the bots have three rotational degrees of freedom, but only one is used here (2D heading).

The **model** bot (labelled M in the figure) executes the following behaviour: move forward if not too

close to a wall; otherwise turn away from the nearest wall and then move forward. The actual distance at which the proximity sensor is triggered is determined by a setting in the sensor module. The region that this state applies to is represented by the shaded area in the figure. The angle through which the bot turns is calculated randomly, constrained by the fact that the bot must then head away from the nearest wall.

The goal of the **imitator** bot (labelled *I* in the figure) is to locate a model, and then remain a fixed distance behind it and record observations (after Billard and Dautenhahn (2000)). In this toy environment, it probably would have been sufficient to have a stationary imitator, but for larger and more complex environments and model behaviours, the imitator would need to stay close to its model in order to observe as closely as possible what the model observes. The imitator needs to be aware of when the model initiates a new action, so that it can record the sensor state at that instant and use it later to construct a perception/action mapping (see Section 5.1). We have tried two types of cue for this purpose:

- 1. The model acts explicitly as a teacher, informing the imitator of its decisions as and when they are made. The imitator only records an observation when this cue is given.
- 2. The model is passive, forcing the imitator to take snapshots of the sensor space at some predetermined regular interval.

The former simulates the training of a team-mate, i.e. where the goal of the model is for the imitator to learn as efficiently as possible. This method could not, however, be used to learn behaviour from 'unhelpful' agents (such as opponents). The latter could be used in this way, but risks missing the decision instant if the observation frequency is too low. There is also a risk of storing redundant data if thresholds between motions are not accurately detected. Nevertheless, either of these problems should be addressable given sufficient learning opportunities and a robust probabilistic representation.

Whichever cue is used, we endow the imitator with the ability to recognise the actions *move forward* and *turn*. The first set of sensors we gave to the imitator detected the x- and y-position respectively of the model in World co-ordinates, resulting in a 2D sensor space. At first glance, the partition would seem to be obvious; in fact directly analogous to the plan shown in Figure 2. The problem is that *move forward* decisions are taken both in the white zone, and in the shaded zone immediately after the robot has finished turning. If the shaded area cannot be mapped to a unique action, then the partition it generates is unsuitable. In fact, there is no suitable partition of this sensor space. Even if we take a more powerful representation and give the imitator a sensor that detects the distance of the model from the nearest wall, the problem remains.

There are (at least) two ways to solve this problem. The first is to give the imitator a sensor which detects the past (commonly known as memory) or, more specifically, detects the previously recorded action. If we use this in tandem with the distance sensor, we can create the following map: if close to a wall and the previously recorded action was move forward then turn; otherwise move forward. This makes sense, as there is an implicit two-item sequence present in the behaviour of the model. The second is to add to the distance sensor another which can detect whether or not the model is facing the nearest wall. The resulting map is equivalent to the one above: if close to a wall and facing it then turn; otherwise move forward. This also makes sense, as the model's behaviour contains a piece of state indicating whether or not it is facing the nearest wall. What is noteworthy is the two different ways of solving the same problem; one temporal and one atemporal.

5.4 Discussion and Future Work

Given this ambiguity, our next task is to investigate whether harder tasks are better solved by a greater number of 'immediate' sensors, or by the introduction of temporal dependencies. We expect POMDPs (Kaelbling et al., 1998), which we also mentioned in Section 4, will provide a way of more naturally modelling temporal systems, as well as the latent variables which are bound to be components of more complex behaviour.

Also, as we alluded to in Section 5.1, we conjecture that grouping sensors into modules that compete probabilistically for saliency in a particular action context will create naturally competing perceptual classes which will in turn need prioritization (see part 4 of the task learning model in Section 4). Using the scenario in Section 5.3, as an example suppose we created several sensor modules each containing one sensor as follows:

- 1. Distance of model from nearest wall.
- 2. Distance of model from second nearest wall.
- 3. Distance from the North wall.

After repeated observations it would become clear that module 1 influences the decision process of the model with a far greater probability than modules 2 or 3; that perceptual class should be given a higher priority at least in the context of generating turns. On the other hand, if there is a door in the North wall, for some other tasks the absolute location may be more salient. In general, we expect an agent will need to actively maintain modules which utilise different viewpoints (e.g. World (absolute / allocentric) view, model / imitator (egocentric) view, teammateand opponent-oriented views, etc.) to see which provide the most easily interpretable behaviour data in different contexts.

Currently however we are still working with relatively simple representational issues, including that of discreteness. Many actions are not easy to discretize: a bot that is turning may make one decision to turn in a long, continuous arc, or many consecutive decisions to turn in a series of smaller arcs. As we said in Section 2, it may not be important that the imitator forms the same perceptual categories as the model is using. In particular, since our models are using a radically different action-selection mechanism than our imitators, it is actually quite likely that the optimal behavioural categories may be different.

To reiterate our hypothesis, we assume that, some of these discriminations will be informed by skills the learning agent has already accumulated, whether through individual learning, previous imitation learning, or by 'innate' predisposition.

6 Summary

If we can build a model of task learning in the games domain, then it will be fairly simple to test how much social learning of a task can accelerate task learning by individual agents, as we can easily create experimental subjects that do or don't attend to other agents in the room. Also, if we allow for both individual and social learning at the same time, we believe we will quite naturally demonstrate agents with similar expressed behaviour, but with different internal representations. Finally, assuming we acquire learners with different perceptual categories (either through learning as just described or by programming) we will be able to test in what circumstances successful behaviours can be propagated through multiple 'generations' across multiple learning agents.

In this paper, we have described one of the key problems in memetics, the problem of discreteness in the representation of behaviour observed in conspecifics. We have suggested that the units of memetics may in fact vary between individuals based on their skills, knowledge and random factors in the self-organization of the underlying neurological representations of these. This may not be a bad thing, in fact it may account for why some observers are able to exceed the performance of their models. We have proposed a framework for representing this sort of learning and described preliminary experiments in building and using such a framework for social learning in the context of real-time multi-player computer games. In the future, we hope to radically expand our experiments and in the process continue to refine our model.

References

- Maria C. Alvarado and Jocelyne Bachevalier. Revisiting the maturation of medial temporal lobe memory functions in primates. *Learning and Memory*, 7(5):244–256, 2000.
- J. R. Anderson and M. Matessa. The rational analysis of categorization and the ACT-R architecture. In M. Oaksford and N. Chater, editors, *Rational Models of Cognition*. Oxford University Press, 1998.
- Christopher G. Atkeson, Andrwe W. Moore, and Stefan Schaal. Locally weighted learning for control. *Artificial Intelligence Review*, 11:75–113, 1997.
- D. M. Bannerman, M. A. Good, S. P. Butcher, M. Ramsay, and R. G. M. Morris. Distinct components of spatial learning revealed by prior training and NDMA receptor blockade. *Nature*, 378: 182–186, 1995.
- Aude Billard and Kerstin Dautenhahn. Experiments in social robotics: grounding and use of communication in autonomous agents. *Adaptive Behavior*, 7(3/4), 2000.
- Joanna J. Bryson. Cross-paradigm analysis of autonomous agent architecture. Journal of Experimental and Theoretical Artificial Intelligence, 12 (2):165–190, 2000a.
- Joanna J. Bryson. Hierarchy and sequence vs. full parallelism in reactive action selection architectures. In *From Animals to Animats 6 (SAB00)*, pages 147–156, Cambridge, MA, 2000b. MIT Press.
- Joanna J. Bryson. Evidence of modularity from primate errors during task learning. In *Modelling Language, Cognition, and Action: The Nineth*

Neural Computation and Psychology Workshop (*NCPW9*), 2005.

- Joanna J. Bryson and Jonathan C. S. Leong. Primate errors in transitive 'inference': Representations underlying transitive choice in humans and other primates. submitted for review, 2005.
- Joanna J. Bryson and Lynn Andrea Stein. Modularity and design in reactive intelligence. In *Proceedings* of the 17th International Joint Conference on Artificial Intelligence, pages 1115–1120, Seattle, August 2001. Morgan Kaufmann.
- Cindy A. Buckmaster, Howard Eichenbaum, David G. Amaral, Wendy A. Suzuki, and Peter R. Rapp. Entorhinal cortex lesions disrupt the relational organization of memory in monkeys. *The Journal of Neuroscience*, 24(44):9811–9825, November 3 2004.
- Richard Dawkins. Hierarchical organisation: A candidate principle for ethology. In P. P. G. Bateson and R. A. Hinde, editors, *Growing Points in Ethology*, pages 7–54. Cambridge University Press, Cambridge, 1976a.
- Richard Dawkins. *The Selfish Gene*. Oxford University Press, 1976b.
- Richard Dawkins. Foreword to Susan Blackmore's *Meme Machine*, pages *vii–xvii*. Oxford University Press, second edition, 2000.
- Daniel C. Dennett. The new replicators. In Mark Pagel, editor, *The Encyclopedia of Evolution*, volume 1, pages E83–E92. Oxford University Press, Oxford, 2002.
- Digital Extremes. Unreal Tournament, 1999. Epic Games, Inc.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998.
- T. Kobayashi, H. Nishijo, M. Fukuda, J. Bures, and T. Ono. Task-dependent representations in rat hippocampal place neurons. *JOURNAL OF NEURO-PHYSIOLOGY*, 78(2):597–613, 1997.
- John E. Laird and Michale van Lent. HumanlevelAI's killer application: Interactive computer games. *AI Magazine*, 22(2):15–25, 2001.

- James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995.
- Tom Mitchell. *Machine Learning*. McGraw Hill, New York, NY, 1997.
- Mathew Nelson. Robocode, 2002. IBM alphaWorks.
- Tomaso Poggio. A theory of how the brain might work. Memo 1253, Massachusetts Institute of TechnologyArtificial Intelligence Lab, Cambridge, MA, December 1990.
- Ronald A. Rensink. The dynamic representation of scenes. *Visual Cognition*, 7:17–42, 2000.
- Edmund T. Rolls. A theory of hippocampal function in memory. *Hippocampus*, 6(6):601–620, 1996.
- Deb Kumar Roy. *Learning from Sights and Sounds:* A Computational Model. PhD thesis, MIT, Media Laboratory, sep 1999.
- Larry R. Squire, Robert E. Clark, and Barbara J. Knowlton. Retrograde amnesia. *Hippocampus*, 11 (1):61–71, 2001.
- T. J. Teyler and P. Discenna. The hippocampal memory indexing theory. *Behavioral Neuroscience*, 100:147–154, 1986.
- Alessandro Treves and Edmund T. Rolls. A computational analysis of the role of the hippocampus in memory. *Hippocampus*, 4:374–391, 1994.
- Andrew Whiten. Primate culture and social learning. Cognitive Science, 24(3):477–508, 2000.
- S. I. Wiener. Spatial, behavioral and sensory correlates of hippocampal CA1 complex spike cell activity: Implications for information processing functions. *Progress in Neurobiology*, 49(4):335, 1996.
- Mark A. Wood, Jonathan C. S. Leong, and Joanna J. Bryson. ACT-R is *almost* a model of primate task learning: Experiments in modelling transitive inference. In *The* 26th Annual Meeting of the Cognitive Science Society (CogSci 2004), pages 1470– 1475, Chicago, 2004. Lawrence Erlbaum Associates.

Motor resonance and MOSAIC

Thierry Chaminade ATR Computational Neuroscience Laboratory 2-2-2 Keihanna Science City, Soraku-gun, Kyoto, 619-0288, Japan tchamina@atr.jp

Abstract

Motor resonance is the automatic involvement of motor control systems during perception of actions. Behaviourally, it is evident in behaviours like motor priming, the facilitation of the execution of an action by seeing it done, and believed to be used imitation. Description of motor resonance properties based on computational models has recently been proposed, and is extended here using MOSAIC. Inverse models are described as goal-directed aspects of actions, and this function is ascribed to the mirror neurons of the ventral premotor cortex. Forward models correspond to sensory-motor aspects of actions, found in the mirror neurons of the inferior parietal lobule. The parieto-premotor connectivity underlies the pairing between forward and inverse models in MOSAIC. Finally, this model argues in favour of the hypothesis that lower aspects of motor resonance, like motor priming, and higher-aspects, like emulation, have interlaced but separable underlying principles.

1 Introduction

Motor theories of social behaviours have flourished in the scientific literature (Blakemore and Decety, 2001; Gallese, 2003; Gallese et al., 2004; Rizzolatti et al., 2001). The foundation of these theories is that the same neural structures - neurons in monkey neurophysiology and functional brain areas in human brain imaging - show an increase of activity both when executing a given action and when observing another individual executing the same action, a process we will refer to as motor resonance¹. For example, a "mirror neuron" is a neuron in the monkey's premotor F5 region which discharges during goal-directed actions as well as when the monkey observes another individual perform the action encoded by this neuron (Rizzolatti and Craighero, 2004; Rizzolatti et al., 1996). This "mirror system" has been proposed to underlie a number of social behaviours such as imitation (Rizzolatti et al., 2001) but also mind reading and empathy (Gallese, 2003). This overstatement of the role of one phenomenon and, in its most extreme versions, of one functional brain area or type of neurons, to describe human social cognition, is arguable (Jacob and Jeannerod, 2005). On the other hand, current descriptions of the physiological processes underlying motor resonance mechanisms rely on the concepts such as "motor words" (Fadiga *et al.*, 2000), reminiscent of action representation in humans (Jeannerod, 1997). This phenomenology of motor resonance fells short of explaining the underlying processes.

Descriptions of motor resonance properties based on computational models have recently been proposed (Miall, 2003). We will extend this proposition considering how forward and inverse models could provide insights into the different phenomena thought to result from motor resonance. This extension will be particularly interesting to separate substrates for lower and higher aspects of motor resonance. A first section will describe observations of human behaviours that favour the hypothesis that observation of action automatically interferes with the execution of action, with the attempt to separate lower aspects, such as motor priming, and higher aspects, such as emulation. Then a computational model for motor control, MOSAIC, will be introduced and the possibility to describe lower and higher aspects of motor resonance based on this framework will be discussed. We will also propose a cortical implementation for the different aspects of this framework, focusing on two regions found in

¹ To resonate has three meanings: matching (between perceptual and motor representations of actions); automatic (intrinsic property of the neuronal system); resounding (the effect lasts and increases by repetition). Other uses of this term can be found in relation to mirror neurons (Fadiga *et al.*, 2000; Rizzolatti and Craighero, 2004).

humans and monkey, the ventral premotor and the inferior parietal cortex, and reinterpret their role in motor control, action observation and imitation at the light of the MOSAIC framework. A conclusion will discuss the scope and the limits of this framework, emphasizing that it should be considered as an automatic and unconscious perceptual process, responsible for some primitive aspects of social cognition but that additional components, such as agency and a system to represent mental states, are needed to explain higher cognitive behaviours.

2 Behavioural exploration of motor resonance

2.1 Motor interference

Motor interference relates to the influence the perception of another individual action has on the execution of action. When asked to raise their fingers in response either to a symbolic cue appearing on a nail or to a movement of the finger of a hand presented visually (Brass et al., 2000), it was found that the finger movement influences the response to the cue -measured as reaction time-, but the reverse effect is very small. In comparison to responding to the cross alone, lifting the index finger in response to a cross appearing on the index fingernail takes more time if the middle finger of the target is lifted at the same time, and less time if its index finger is moving. In other word, when responding to a symbolic cue, the response is hindered by the observation of an incompatible action and facilitated by a compatible one. The same effect is seen when the gesture is a hand posture and the cue the colour of the stimulus hand (Sturmer et al., 2000). Similar experimental paradigms were used to show that this phenomenon was complementary to other stimulus-response compatibility effects such as spatial compatibility (Brass et al., 2001a), emphasizing the specificity of the effect of compatibility between the observed and the executed action. Thus producing an action similar to an observed action is a prepotent response that requires to be inhibited to execute the correct response².

The variance of horizontal and vertical arm movements is significantly increased when watching a human, but not an industrial robot, perform a spatially incongruent movement (Kilner *et al.*, 2003). As in the previous experiments, it must be emphasized that the effect cannot only be explained as direct mapping of perceived action onto an execution system because the spatial congruency implies a mirror effect. My left to right arm movement is congruent with your right to left arm movement if we are facing each other. Thus both spatial and action compatibility are in play. We performed a related experiment to determine which features of the agent you interact with are involved in motor interference (Oztop et al., 2004). We found that a humanoid robot which movements actually reproduce human movements causes an attenuated but reliable interference effect in comparison to a human. Since the robot is obviously not perceived as human, this result implies that some characteristics of any agent, in term of aspect and motion, define its ability to cause motor interference.

Another result favouring a mapping between perceived and executed action is the motor priming effect observed in prehensile action. A human or a robotic hand performed a grasping action on a small or a large target prior to subjects performing a grasping action themselves. Congruence of the presented movement kinematics with the response movement kinematics, obtained when the model and subject grasp objects of the same size, had a significant priming effect when the model was a human but not when it was a robotic hand. This result implies that the primed movement kinematics influence the execution of a grasping movement (Gallese et al., 2002). A subsequent experiment showed that in addition to the effect of the action observation, object affordance also plays a role in the priming of the grasping action (Edwards et al., 2003).

2.2 The chameleon effect

Does the motor priming described in a laboratory environment have any reality in everyday life? The chameleon effect was introduced to describe the unconscious reproduction of "postures, mannerisms, facial expressions and other behaviours of one's interacting partner" (Chartrand and Bargh, 1999). This effect can easily be experienced in face-to-face interactions, when one crosses his arms or legs to see his partner swiftly adopt the same posture. Subjects unaware of the purpose of the experiment interacted with an experimenter performing one of two target postures, rubbing the face or shaking the foot. Analysis of the behaviour showed a significant increase of the tendency to engage in the same action. In another study, experi-

 $^{^2}$ Discussions on the frontal origin of this inhibition are out of the scope of the present report, but can be found in patient (De Renzi *et al.*, 1996) and neuroimaging (Brass *et al.*, 2001b) studies.

menters mimicking the subjects were rated as more likable. Subjects were unaware of the mimicking manipulation of the experiment. This emphasizes that the chameleon effect is an unconscious process improving sociability of an interaction by which parts of an individual behaviour is transmitted to another individual during interaction. This automatic trigger of social behaviours by the mere observation of the same behaviours in others is referred to as the "perception-behaviour expressway". It should be noted that mimicry has been described as a source of empathy (Decety and Chaminade, 2003). In contrast, imitation is intentional.

2.3 Imitation

There have been numerous claims that the mirror neurons underlie imitation. This attractive hypothesis suffers from several drawbacks. The main problem is the definition of imitation, which depending on the stance of the author can extend to a very large set of behaviours having little in common (Byrne and Russon, 1998). Another debated issue is the existence and the forms of imitative behaviours in monkeys, the only species in which mirror neurons were directly investigated. Since imitation is scarce in monkey, there is no recording of "mirror neuron" during imitation, imposing caution when linking such a high-level behaviour to the neurophysiology.

2.3.1 Proto-imitation

Infants between 12 and 21 days of age can imitate both facial and manual gestures (Meltzoff and Moore, 1977). After being presented visually with facial movement (tongue-protrusion, lipprotrusion, mouth opening) or sequential finger movements, newborns responses were recorded and their gestures were categorized. Results showed a significant increase of the target stimulus in comparison to the other gestures. A similar result was obtained with younger newborns, ruling out an effect of early social experience and thus favouring an innate capacity, proto-imitation. In six-week-old newborns, a tongue protrusion to the side led the execution of a similar tongue protrusion after correcting earlier approximations (Meltzoff and Moore, 1994).

These results imply that newborns can innately equate their own unseen behaviours with gestures they see others perform, as described in the "active intermodal mapping" model (Meltzoff and Moore, 1997). Using an initial organ identification system and trial-and-error movements, infants attempt to match a relation between organs they see in the adult with the relation between organs they feel when performing the action themselves. This matching mechanism relies heavily on a representational system that allows infants to interlace felt transformations of their body transformations – somatosensory input- and seen transformations of someone else's body –visual input.

2.3.2 Goal-directed imitation

In contrast to proto-imitation, it was recently proposed that imitation is primarily directed to the goal of the observed action, and that goaldirected models better explain imitation than direct matching models (Wohlschlager et al., 2003). Arguments are derived from a set of experiments in children and adults, which demonstrated that when asked to imitate, subjects reproduce the most salient goal sometimes at the expense of the reproduction of the given action. For example, one key experiment (Wohlschlager and Bekkering, 2002) was a modification of the finger movement paradigm presented in 2.1 (Brass et al., 2000). Subjects had to imitate a downward finger movement with one of the two hands, which was made ipsi- or controlaterally. A dot was present on the table in the goaldirected conditions. Reaction time showed a clear facilitation for ipsilateral movements and a negative effect on the number of errors for controlateral movements when the dot was present. Thus the goal has a decisive influence on the imitation behaviour.

In the goal-directed imitation model, the imitator decomposes the observed action into its separate aspects and reproduces the most salient one(s). In an extensive review of the literature on learning by imitation, Byrne and Russon (1998) made a number of useful distinction between behaviours that could be described as imitation, emphasizing their differences in a phenomenological perspective. Particularly relevant to the present discussion is the distinction between "action level" and "program level" imitation. Action level imitation requires that style and minor details should match between mimic and model. According to Byrne and Russon, it is likely to involve cognitively simple kinaestheticvisual and sensorimotor matching. Its involvement would concern social functioning rather than learning of behaviour, because of the role of contingencies in interactions. Proto-imitation, motor priming, and the chameleon effect could be related to action level resonance. In program level imitation, animals and humans reproduce the hierarchical organization of behaviour; actions are understood as hierarchically organized subgoals reproduced idiosyncratically. The goal-directed imitation model describes some form of program level imitation. One extreme form of program level imitation is emulation, in which knowledge about the relationships between objects and goals are acquired by observation of a conspecific using these objects.

Analysis of behaviour, in particular the case of imitation, indicates two levels of resonance, sensory-motor and goal-directed. In the next part we will argue that describing perception and imitation of action using the MOSAIC model offers explanations for these two types of resonance.

3 Computational model of motor control



Figure 1: Simplified version of the MOSAIC (MOdular Selection And Identification for Control) model for motor control (Wolpert *et al.*, 2003).

3.1 Internal models in the control of action

3.1.1 Forward models or predictors

Very early in the history of psychology was it anticipated that the motor and sensory aspects of actions were tightly coupled. "An anticipatory image [...] of the sensorial consequences of a movement [...] is the only psychic state which introspection lets us discern as the forerunner of our voluntary acts" (James, 1890). William James' "efferent discharge" seems a premonition of the "corollary discharge" or "efference copy" involved in optimal motor control. "Corollary discharges" are copies of the motor command (\mathbf{x}_t in Figure 1) used by internal models of the body, the forward models (FM), to predict the sensory consequences of the actions (Kawato, 1999). The prediction $(\hat{\mathbf{x}}_t)$ can then be compared to the actual sensory consequences of the action, (\mathbf{y}_t) and used to optimize motor control. The forward model can also be used to filter reafferent

sensory information: the sensory consequences of an action correctly predicted by the system are attenuated making non predicted sensory inputs more salient (Frith *et al.*, 2000). This could explain why we are not aware of our voice when speaking, of touching an object during a visually guided grasp, or of self-tickling.

3.1.2 Inverse models or controllers

Inverse models can calculate the feedforward motor commands \mathbf{x}_t from a desired trajectory information \mathbf{x}^*_t (Kawato, 1999). When experiencing a new object or context, the motor system sends a feed-forward signal and gets feedback signals that allow him to correct its internal representation of the relation between the desired state and the motor commands. With training and generalization, the model can then act as a controller of action providing the motor command adapted to the new object or context. Inverse models can thus be thought of as mappings between goals and contexts on the one side, and motor commands on the other side.

3.2 MOSAIC

3.2.1 Action control

In MOSAIC (Figure 1), predicted states can be used to identify the current context within which control is being attempted. Its uses multiple parallel modules, each comprising paired forward and inverse models dedicated to action control. In each module a forward model uses a copy of the motor command \mathbf{x}_t to predict the next state of the system; this prediction is compared to the sensory feedback, the actual state of the system \mathbf{y}_t , to produce a prediction error. The error of all modules are combined to estimate the responsibility signal of each module $\boldsymbol{\lambda}_i$, and the model resulting in a minimum error is given the higher responsibility in describing the current sensory-motor context.

This weighting is also applied to calculate the contribution of the inverse model from the i-th module in the final motor command x_t . Therefore the inverse model most adapted to the current sensory-motor context, according to the forward model, is selected to control the motor system.

3.2.2 Action perception

The possibility of the MOSAIC model to be utilized in action perception and imitation has been described by Wolpert *et al.* (2003). During perception of someone else's behaviour the motor system produces no action. In this context, all modules are initially equiprobable and the output of each inverse model is used as the input of the paired forward model. Each model provides a prediction of the next state of the observed action, which is compared with the observed next state of the action. Results from this comparison are used to calculate a responsibility, which is this case describes the efficacy a given module has in predicting the partner's behaviour.

The selection of the module yielding the highest responsibility score equates to making more salient a pair of a forward and an inverse model. The higher activity of one module increases its output, the motor command \mathbf{x}_{t} , and is a possible foundation for lower level aspects of motor resonance (chameleon effect, motor interference). On the other hand a read-out of the relative activations of the inverse models would provide an inference of the desired state underlying the observed action, the motor intention of the partner, and could explain higher aspects of motor resonance (goal-directed imitation, emulation, action understanding). In conclusion, the selection of a given module resulting from the comparison of the predictions of each module with the observed action can explain both lower and higher aspects of motor resonance, though the second necessitates an extra 'read-out' process.

4 From motor control to motor resonance

4.1 Putative neural bases of MOSAIC

Figure 2 highlights regions involved in action perception and imitation, which functions can be described within the MOSAIC framework.



Figure 2: Lateral render of macaque and human brains showing the putatively homologous brain regions involved in perception and imitation of action. STS: Superior temporal sulcus; IPL: Inferior Parietal Lobule; Crblm: Cerebellum.

4.1.1 Inverse models

Inverse models calculate the motor command given an intended state of the body, or goal. Premotor neurons in general, and mirror neurons in particular, are suited to implement this type of transformation.

Rostral to the motor cortex and caudal to the prefrontal cortex, to which they are connected, the role of premotor areas of the prefrontal cortex is the control of actions. Recent advances from monkey electrophysiology led to the description of a parcelled cortex, where different parcels are involved in different aspects of action control (Rizzolatti et al., 1998). Premotor areas can be subdivided in dorsal and ventral areas. Interestingly, the ventral parts -areas F4 and F5- would mostly code the motor control for peripersonnal space (part of the body/part of the body, part of the body/object, and part of the body/object/part of the body interactions), while dorsal parts -areas F2 and F7- would mostly code learnt artificial sensory-motor association. In contrast, the primary motor cortex -Brodmann area 4 in humans or area F1 in monkey'splays a major role in controlling more elementary features of movement control necessary to achieve a given action.

All premotor subregions have sensory as well as motor properties, especially interesting in area F5, where two types of premotor neurons with different visual properties have been described. Premotor neurons in area F5 encode object-directed behaviour performed with the hand and the mouth, and their connection with primary motor area offers a way to control action directly. "Canonical neurons" respond to the perception of objects, and "mirror neurons" to the perception of another individual performing an action similar to that encoded by the neuron. Mirror neurons firing is also elicited by the observation of object-related actions, even when the final part of the action is hidden but can be inferred (Umilta et al., 2001), showing that they encode the relation between the desired end-point -or goal- and the effector kinematics. Though mainly based on inference, the human brain area homologous to the monkey F5 is believed to be Broca's area, which is a large region of the inferior frontal cortex.

In humans, Broca's area activity has been associated with imitation. Neuroimaging experiments in humans have investigated the cerebral network underlying imitation. One set of results using fMRI (Iacoboni *et al.*, 1999) and MEG (Nishitani and Hari, 2000) showed an area in the left inferior frontal gyrus, known as Broca's area, to demonstrate activation patterns similar to those expected from the neuronal substrate of a human "mirror system". It is activated in both execution and observation conditions, and more activated in the imitation condition. This experiment was refined to test whether the presence of a goal, characterized by dots tapped by the fingers on the support of the hand, reproduce the activity of Broca's area (Koski *et al.*, 2002). The presence of goal increased the response in Broca's area, and caused additional bilateral dorsal premotor activities.

Altogether properties of inferior premotor cortices are in line with the idea that this area participates to internal models for the control of action, and the goal directedness favours an involvement at the level of inverse models.

4.1.2 Forward models

Forward models compute the expected next sensory feedback given an internal copy of the action motor command and the current state of the body, and the comparison of this prediction with the actual feedback is used to update the actual state of the body and correct the movement. Both monkey (Rushworth *et al.*, 1997) and human (Desmurget *et al.*, 2001) neurophysiology have demonstrated that the posterior parietal cortex is active during on-line movement control.

The posterior parietal cortex has also been implied in other functions of forward models in action control, for example in updating the internal model of the body based on sensory information (Blakemore et al., 1998; Wolpert et al., 1998). It is notable that it the posterior parietal cortex integrates sensory signals from many modalities (e.g. visual, proprioceptive, auditory and vestibular), as well as efference copy from motor structures (Andersen et al., 1997). With the exception of the neuropsychological data, in which the patient showed a large lesion of the parietal cortex, the inferior parietal lobule (IPL) is found in all studies. Accordingly, the "ventral-dorsal" visual stream introduced by (Rizzolatti and Matelli, 2003), which covers the inferior parietal cortex area PF in the monkey, is crucial to action organization. Altogether, these examples illustrate that the IPL could be involved in forward models based on the integration of sensory and motor signals.

In accordance with the proposed role of forward models in lower aspects of action imitation, activity in the left IPL can be found in neuromaging studies of human imitation. For example we found using PET that the IPL is activated when reproducing (Decety *et al.*, 2002) but also when simply tracking (Chaminade and Decety, 2002) another individual action. An fMRI investigation of the neural substrate of body-part and movement parameter coding during imitation of intransitive action (Chaminade *et* *al.*, 2005) showed that the IPL encodes the bodypart, while the superior parietal lobule encodes the spatial aspect. Similarly, reproduction of hand and finger static postures activate the inferior parietal cortex in the absence of any perceived movement (Tanaka and Inui, 2002), which may be related to the motion implied by the presentation of static images of body postures.

4.1.3 Responsibility estimates

Together with the parietal cortex, the cerebellum is another candidate for the localisation forward models (Blakemore and Sirigu, 2003; Kawato, 1999). The cerebellum participates to motor control with premotor and parietal cortices (Haslinger *et al.*, 2002). It is crucially involved in many mechanisms underlying internal models for motor control - prediction of the sensory consequences of action (Blakemore *et al.*, 2001), learning of internal models (Imamizu *et al.*, 2000). But until recently, there was no satisfactory description of its involvement in internal models for motor control (Miall, 2003).

Imamizu, Kawato and colleagues (Imamizu *et al.*, 2004) investigated the brain networks involved in representing inverse and forward models as well as regions responsible for the switch in the MOSAIC framework. An interpretation of the results was that premotor, anterior parietal and cerebellar regions contained internal models, but that switching of internal models relied on the cerebellar cortices. One possible explanation is that the cerebellum computes and compares prediction errors and responsibility, being able to work in a default mode with familiar actions (e.g. walking), and that inverse and forward models are represented in the cortical premotor and parietal areas respectively.

4.2 Explaining motor resonance

4.2.1 Parieto-premotor loops and action control

In monkeys, it is believed that parietopremotor networks form the core of the motor control system. Taking into account the connections between parietal and premotor areas, a series of segregated parieto-premotor functional circuits can be distinguished, each being involved in a specific sensory-motor transformation for action (Rizzolatti *et al.*, 1998). The ventral premotor area F5 and the ventral parietal PF, two areas were mirror neurons were found and which correspond roughly to Broca's area and the IPL in humans, are anatomically and functionally connected for the control of object-directed actions.

The present proposal is that direct or/and indirect (e.g. through the cerebellum) connectivity between areas F5 and PF in monkey sustains modular pairing of inverse and forward models, with an emphasis on its goal aspects in area F5 (as in inverse models) and on its sensory aspects in area PF (as in forward models). In this framework, motor resonance corresponds to the pairing of computations performed by the inferior parietal cortex and the ventral premotor cortex in action representations, interlacing sensory and intentional aspects for control and for perception of action.

In action control, the premotor region is in charge of programming the action, and the parietal cortex provides on-line action control. The present model proposes that the motor command is mainly based on the output of the inverse model from one module, and a copy of the resulting motor command, an efference copy, is sent to the parietal cortex area in charge of the forward model of the same module. The comparison between the actual and predicted sensory consequences of action is responsible for on-line control and correction of the action.

4.2.2 The case of action perception

The most important aspect is that the motor resonance system can be activated by the mere observation of action in the absence of any motor output. Evidence abounds in favour of this aspect (Chaminade and Decety, 2001), which actually motivate current research to understand the motor resonance process.

In the present model, perception of someone else's action, mediated by the superior temporal sulcus (STS³), is used to select the modules whose forward model is best suited to reproduce the sensory aspect of the observed action. This leads to an increase of activity: in premotor areas in relation to the activity of the inverse models providing input to the forward models; in the parietal cortex in relation to the computation of the predicted sensory inputs and their comparison with the actual input; and finally in the cerebellum in relation to estimating responsibility and selecting the module(s) with the inverse model most likely to give rise the same sensory consequences when performed by the self. Results are in line with these predictions.

First mirror neurons, found in the monkey's ventral premotor area F5 (Rizzolatti et al., 1996) and parietal area PF (Gallese et al., 2002) are activated both when the monkey perform a given action and when it sees an individual performing a similar action. A related result in human show that observation of action activates the premotor and parietal cortex in a somatotopic manner (Buccino et al., 2001). A meta-analysis of neuroimaging studies on action execution, observation and simulation reveals that parietal and premotor regions are activated by the three tasks (Grezes and Decety, 2001). Finally, parietal and premotor activity when predicting the outcome of dots kinematics representing simple movements, writing and pointing, is found in the cortical regions known to be involved in the execution of the same task (Chaminade et al., 2001). Parietal and premotor cortices controlling action are thus also activated, in an effector- and actionspecific way, by the observation of action.

4.2.3 The case of imitation

The ventral parieto-premotor network, together with the STS, is particularly important for imitation (Chaminade *et al.*, 2002; Chaminade *et al.*, 2005; Decety *et al.*, 2002; Miall, 2003; Rizzolatti *et al.*, 2001). The parietal cortex has been involved in lower aspects of imitation like body-parts coding and cinematic aspects (see 4.1.2). Accordingly in monkey, the PF mirror neuron system is believed to code for the kinaesthetic and somatosensory components of actions. In the present view, their activity reflects the process of selecting the forward model that can best reproduce the outcomes of the observed action that has been encoded by the STS.

Mirror neurons in F5 encode goal-directed actions, and accordingly the human homologous region, Broca's area, is activated by goal directed aspects of imitation (see 4.1.1). Some results even suggest that it is not specific to imitation per se, but can be activated by goal-directed aspects of an action when it is triggered by symbolic means instead of imitation, and when it is on-line or delayed (Makuuchi, 2004). Its involvement provides the system with a way to read out the intention of the observed action.

No experimental data to date is able to differentiate between the present and other models describing the involvement of motor cortices in action understanding. The present model allows for testable predictions. For example, it predicts an increase of activity of the cerebellum when under-

³ The STS is involved in transforming a visual input into a code that can be used by the resonance system. Its function is purely visual and out of the scope of the present discussion (see Allison *et al.*, 2000).

standing an observed behaviour, related to the switching between the different paired inverse/forward models. It also predicts that contributions of parietal cortex and premotor cortex in motor resonance are different, being involved in forward and inverse internal models of action.

5 Conclusions

Motor resonance is the concomitance of the activation of sensory aspects and goal-directed aspects of motor control, believed to be located in the inferior parietal and ventral premotor cortices respectively, during execution and perception of action. Using MOSAIC, it is possible to describe the sensory aspects as forward models, and the goaldirected aspects as inverse models. Parieto-premotor connectivity offers a way to pair inverse and forward models within action modules. This framework is coherent with the current knowledge on the brain bases of motor resonance, and offers testable predictions.

Yet it leaves several key points unexplained. How does this system differentiate between observed actions and actions produced by the self if both use the same code? Investigations on the sense of agency provides answers that could be transposed in the present framework (see for example Decety and Chaminade, 2003). Another example is the representation of higher-order intentions. When observing other behaviour, we have access not only to motor intention, but also to social intentions, taking contextual information into account.

Motor resonance could be a foundation for social understanding, by transforming a sensory code into an action code which interferes with action production (lower aspects of motor resonance), and by providing a mean to understand observed behaviour (higher aspects of motor resonance). The need to investigate the current proposal is reinforced by the possibility it offers to model social interactions.

Acknowledgement

The author wishes to thank Mitsuo Kawato and Erhan Oztop for helpful discussions of the ideas discussed in this manuscript, as well as an anonymous reviewer for his comments on an earlier version.

References

Allison T., Puce A. and McCarthy G. Social perception from visual cues: Role of the STS region. *Trends in Cognitive Sciences*, 4(7):267-278, 2000.

Andersen R. A., Snyder L. H., Bradley D. C. and Xing J. Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, 20:303-330, 1997.

Blakemore S. J., Wolpert D. M. and Frith C. D. Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1(7):635-640, 1998.

Blakemore S. J. and Decety J. From the perception of action to the understanding of intention. *Nature Reviews Neuroscience*, 2(8):561-567, 2001.

Blakemore S. J., Frith C. D. and Wolpert D. M. The cerebellum is involved in predicting the sensory consequences of action. *Neuroreport*, 12(9):1879-1884, 2001.

Blakemore S. J. and Sirigu A. Action prediction in the cerebellum and in the parietal lobe. *Experimental Brain Research*, 153(2):239-245, 2003.

Brass M., Bekkering H., Wohlschlager A. and Prinz W. Compatibility between observed and executed finger movements: comparing symbolic, spatial, and imitative cues. *Brain and Cognition*, 44(2):124-143, 2000.

Brass M., Bekkering H. and Prinz W. Movement observation affects movement execution in a simple response task. *Acta Psychologica*, 106(1-2):3-22, 2001a.

Brass M., Zysset S. and von Cramon D. Y. The inhibition of imitative response tendencies. *Neuroimage*, 14(6):1416-1423, 2001b.

Buccino G., Binkofski F., Fink G. R., Fadiga L., Fogassi L., Gallese V., Seitz R. J., Zilles K., Rizzolatti G. and Freund H. J. Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *European Journal of Neuroscience*, 13(2):400-404, 2001.

Byrne R. W. and Russon A. E. Learning by imitation: a hierarchical approach. *Behavioral and Brain Sciences*, 21(5):667-684, 1998.

Chaminade T. and Decety J. A common framework for perception and action: neuroimaging evidence.

Behavioral and Brain Sciences, 24(5):879-882, 2001.

Chaminade T., Meary D., Orliaguet J. P. and Decety J. Is perceptual anticipation a motor simulation? A PET study. *Neuroreport*, 12(17):3669-3674, 2001.

Chaminade T. and Decety J. Leader or follower? Involvement of the inferior parietal lobule in agency. *Neuroreport*, 13(15):1975-1978, 2002.

Chaminade T., Meltzoff A. N. and Decety J. Does the End Justify the Means? A PET Exploration of the Mechanisms Involved in Human Imitation. *Neuroimage*, 15(2):318-328, 2002.

Chaminade T., Meltzoff A. N. and Decety J. An fMRI study of imitation: action representation and body schema. *Neuropsychologia*, 43(1):115-127, 2005.

Chartrand T. L. and Bargh J. A. The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6):893-910, 1999.

De Renzi E., Cavalleri F. and Facchini S. Imitation and utilisation behaviour. *Journal of Neurology, Neurosurgery and Psychiatry*, 61(4):396-400, 1996.

Decety J., Chaminade T., Grezes J. and Meltzoff A. N. A PET Exploration of the Neural Mechanisms Involved in Reciprocal Imitation. *Neuroimage*, 15(1):265-272, 2002.

Decety J. and Chaminade T. When the self represents the other: A new cognitive neuroscience view on psychological identification. *Consciousness and Cognition*, 12(4):577-596, 2003.

Desmurget M., Grea H., Grethe J. S., Prablanc C., Alexander G. E. and Grafton S. T. Functional anatomy of nonvisual feedback loops during reaching: a positron emission tomography study. *Journal of Neuroscience*, 21(8):2919-28., 2001.

Edwards M. G., Humphreys G. W. and Castiello U. Motor facilitation following action observation: A behavioural study in prehensile action. *Brain and Cognition*, 53(3):495-502, 2003.

Fadiga L., Fogassi L., Gallese V. and Rizzolatti G. Visuomotor neurons: ambiguity of the discharge or 'motor' perception? *International Journal of Psy-chophysiology*, 35(2-3):165-177, 2000.

Frith C. D., Blakemore S. J. and Wolpert D. M. Ab-

normalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London: B Biological Sciences*, 355(1404):1771-1788, 2000.

Gallese V., Fadiga L., Fogassi L. and Rizzolatti G. (2002). Action representation and the inferior parietal lobule. In W. Prinz and B. Hommel (Eds.), *Common mechanisms of perception and action* (Vol. XIX., pp. 334-355). Oxford: Oxford University Press.

Gallese V. The manifold nature of interpersonal relations: the quest for a common mechanism. *Philosophical Transactions of the Royal Society of London: B Biological Sciences*, 358(1431):517-528, 2003.

Gallese V., Keysers C. and Rizzolatti G. A unifying view of the basis of social cognition. *Trends in Cognitive Sciences*, 8(9):396-403, 2004.

Grezes J. and Decety J. Functional anatomy of execution, mental simulation, observation, and verb generation of actions: a meta-analysis. *Human Brain Mapping*, 12(1):1-19, 2001.

Haslinger B., Erhard P., Weilke F., Ceballos-Baumann A. O., Bartenstein P., Grafin von Einsiedel H., Schwaiger M., Conrad B. and Boecker H. The role of lateral premotor-cerebellar-parietal circuits in motor sequence control: a parametric fMRI study. *Cognitive Brain Research*, 13(2):159-168, 2002.

Iacoboni M., Woods R. P., Brass M., Bekkering H., Mazziotta J. C. and Rizzolatti G. Cortical mechanisms of human imitation. *Science*, 286(5449):2526-2528., 1999.

Imamizu H., Miyauchi S., Tamada T., Sasaki Y., Takino R., Pütz B., Yoshioka T. and Kawato M. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature*, 403(6766):192-195, 2000.

Imamizu H., Kuroda T., Yoshioka T. and Kawato M. Functional magnetic resonance imaging examination of two modular architectures for switching multiple internal models. *Journal of Neuroscience*, 24(5):1173-81, 2004.

Jacob P. and Jeannerod M. The motor theory of social cognition: a critique. *Trends in Cognitive Sciences*, 9(1):21-25, 2005.

James W. (1890). Principles of psychology. New-

York: Holt.

Jeannerod M. (1997). *The Cognitive Neuroscience of Action*. Oxford: Blackwell.

Kawato M. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6):718-727, 1999.

Kilner J. M., Paulignan Y. and Blakemore S. J. An interference effect of observed biological movement on action. *Current Biology*, 13(6):522-525, 2003.

Koski L., Wohlschlager A., Bekkering H., Woods R. P., Dubeau M. C., Mazziotta J. C. and Iacoboni M. Modulation of Motor and Premotor Activity during Imitation of Target- directed Actions. *Cerebral Cortex*, 12(8):847-855., 2002.

Makuuchi M. Is Broca's Area Crucial for Imitation? *Cerebral Cortex*:Epub ahead of print, 2004.

Meltzoff A. N. and Moore M. K. Imitation of facial and manual gestures by human neonates. *Science*, 198(4312):74-8, 1977.

Meltzoff A. N. and Moore M. K. Imitation, memory, and the representation of persons. *Infant Behavior and Development*, 17(1):83-99, 1994.

Meltzoff A. N. and Moore M. K. Explaining facial imitation: A theoretical model. *Early Development and Parenting*, 6(3/4):179-192, 1997.

Miall R. C. Connecting mirror neurons and forward models. *Neuroreport*, 14(17):2135-7, 2003.

Nishitani N. and Hari R. Temporal dynamics of cortical representation for action. *Proceedings of the National Academy of Science U S A*, 97(2):913-918, 2000.

Oztop E., Franklin D. and Chaminade T. (2004). *Human-humanoid interaction: is a humanoid robot perceived as a human.* Paper presented at the Humanoids 2004, Los Ageles, CA, USA.

Rizzolatti G., Fadiga L., Gallese V. and Fogassi L. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2):131-141, 1996.

Rizzolatti G., Luppino G. and Matelli M. The organization of the cortical motor system: new concepts. *Electroencephalography and Clinical Neurophysiology*, 106(4):283-296, 1998. Rizzolatti G., Fogassi L. and Gallese V. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2(9):661-670, 2001.

Rizzolatti G. and Matelli M. Two different streams form the dorsal visual system: anatomy and functions. *Experimental Brain Research*, 153(2):146-57, 2003.

Rizzolatti G. and Craighero L. The mirror-neuron system. *Annual Review of Neuroscience*, 27:169-92, 2004.

Rushworth M. F., Nixon P. D. and Passingham R. E. Parietal cortex and movement. I. Movement selection and reaching. *Experimental Brain Research*, 117(2):292-310, 1997.

Sturmer B., Aschersleben G. and Prinz W. Correspondence effects with manual gestures and postures: a study of imitation. *Journal of Experimental Psychology: Human Perception and Performance*, 26(6):1746-1759, 2000.

Tanaka S. and Inui T. Cortical involvement for action imitation of hand/arm postures versus finger configurations: an fMRI study. *Neuroreport*, 13(13):1599-1602, 2002.

Umilta M. A., Kohler E., Gallese V., Fogassi L., Fadiga L., Keysers C. and Rizzolatti G. I know what you are doing. a neurophysiological study. *Neuron*, 31(1):155-165, 2001.

Wohlschlager A. and Bekkering H. Is human imitation based on a mirror-neurone system? Some behavioural evidence. *Experimental Brain Research*, 143(3):335-341, 2002.

Wohlschlager A., Gattis M. and Bekkering H. Action generation and action perception in imitation: an instance of the ideomotor principle. *Philosophical Transactions of the Royal Society of London: B Biological Sciences*, 358(1431):501-515, 2003.

Wolpert D. M., Goodbody S. J. and Husain M. Maintaining internal representations: the role of the human superior parietal lobe. *Nature Neuroscience*, 1(6):529-533, 1998.

Wolpert D. M., Doya K. and Kawato M. A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society of London: B Biological Sciences*, 358(1431):593-602, 2003.

Interpersonal Maps and the Body Correspondence Problem

Verena V. Hafner and Frédéric Kaplan Sony Computer Science Laboratory, 6 rue Amyot, Paris, France hafner@csl.sony.fr, kaplan@csl.sony.fr www.csl.sony.fr

Abstract

In this paper, we introduce the concept of "interpersonal maps". They realize a representation of one's own body to include the body of one's peers. In cases of strong couplings between agents, a "we-centric" space can emerge in which the agent's body structure can be directly mapped onto the structure of an observed body. Based on a set of robotic experiments, we argue that this unified representation can help to elucidate both the formation of a body schema and the body correspondence problem.

1 Introduction

The establishment of the self-other identity is a crucial milestone towards the development of more sophisticated forms of social interaction. It serves as a basis for developing intentional understanding, joint attention and imitative capabilities. Matching and discriminating between oneself and others results certainly from the interplay of several developmental dynamics. In this paper, we focus on a subset of this complex issue by considering the links between the formation of the body schema and the body correspondence problem (Nehaniv and Dautenhahn (2002)). We introduce the concept of "interpersonal maps", realizing a representation of one's own body as well as the body of peers.

This idea is related to several existing concepts. To account for early imitation, Meltzoff and Moore argue for the existence of an intermodal mapping establishing equivalence relations between different modalities such as vision or motor actions (Meltzoff and Gopnick (1993); Moore and Corkum (1994)). Such a model suggests that both perceived and observed behaviour could be represented in a shared neural format. Similarly, Gallese has argued that since the beginning of our life we inhabit a shared multidimensional interpersonal space. When we observe other individuals, "a meaningful embodied interpersonal link is established". Gallese refers to this form of intersubjectivity as the shared manifold space. Furthermore, his theory predicts the existence of "somatosensory mirror neurons" giving the capacity to map different body locations during the observation of the bodies of others (Gallese (2004)).

However, few models try to give a precise ac-

count on how such interpersonal or intermodal mappings could be developed. We believe that research in developmental robotics can play a relevant role to progress in understanding the development of such mappings. Designing algorithms addressing the body correspondence problem and the constitution of the body schema is one of the major challenges of this domain (Kaplan and Hafner (2004)). These issues have been investigated in separate manners (e.g. Yoshikawa et al. (2002, 2004) for the body scheme and Nehaniv and Dautenhahn (2002) for approaches of the correspondence problem). Our model results in a preliminary investigation in trying to address both problems in a unified framework.

2 Maps Based on Information Distances

2.1 Definition

Our approach takes inspiration from research carried out by Olsson et al. (2004) concerning the use of information distances between sensors. This research shows that maps can be built as metric projections showing informational relationships between sensors. It is based on the methods by Pierce and Kuipers (1997) on map learning. In such maps, sensors that are informationally related are close to each other. A related approach was investigated by Kuniyoshi et al. (2004). They argued that such information maps could appropriately be related to "somatosensory maps" such as the ones known to exist in the cortex (Penfield and Rasmussen (1950)).

Such a map can be built in the following way:

Computation of the information distance matrix

Let us assume that the robot R_X is equipped with n sensors (proprioceptive and distance sensors). At any time t its sensory state can be captured by the vector X(t)

$$X(t) = (X_1(t), X_2(t), \dots, X_n(t))$$
(1)

For any sensor X_i the entropy $H(X_i)$ can be calculated as

$$H(X_i) = -\sum_{x_i} p(x_i) \log_2 p(x_i)$$

where $p(x_i)$ is the probability mass function over all possible discretised values x_i . To calculate it, the histogram of X_i has to be calculated with a careful choice of the number of bins (see Schreiber (2000)).

The conditional entropy for two sensors X_i and X_j can be calculate as

$$H(X_j|X_i) = -\sum_{x_i} \sum_{x_j} p(x_i, x_j) \log_2 p(x_j|x_i)$$

where $p(x_j|x_i) = p(x_j, x_i)/p(x_i)$.

We chose to use

$$d(X_i, X_i) = H(X_i|X_i) + H(X_i|X_i)$$

as the distance used in the distance matrix since it has several advantages compared to the mutual information (Crutchfield (1990)). d is a metric for the space of information sources. This means that it has the three properties of symmetry, equivalence and triangle inequality.

- d(X, Y) = d(Y, X) follows directly from the symmetry of the definition
- d(X,Y) = 0 if and only if X and Y are recoding-equivalent (in the sense defined by Crutchfield Crutchfield (1990)).
- $d(X,Z) \le d(X,Y) + d(Y,Z)$

Two-dimensional metric projection

A two-dimensional projection is ideal for visualisation of the data. In order to create a two-dimensional body map from the sensor data, we apply a relaxation algorithm. The algorithm is an iterative procedure of positioning the sensors in a two-dimensional space in such a way that the metric distance between two sensors in this map is as close as possible¹ to the distance in the *n*-dimensional information space.

Different algorithms have been suggested (Hafner (2000); Duckett et al. (2002); Pierce (1995)) which convert an *n*-dimensional input into an *m*-dimensional map (m < n). Here, the algorithm of Pierce (1995) is used since it does not require any information about the relative orientation of connections between sensor nodes.

The algorithm used in this paper consists of an iteration of two simple steps:

First, each sensor X_i is randomly assigned to a point $\mathbf{p_i}$ on a two-dimensional plane.

1. The force f_i on each point $\mathbf{p_i}$ is computed as:

$$f_i = \sum f_{ij}$$

where

$$f_{ij} = (||\mathbf{p_i} - \mathbf{p_j}|| - d(X_i, X_j))(\mathbf{p_j} - \mathbf{p_i})/||\mathbf{p_j} - \mathbf{p_i}||$$

2. Each point **p**_i is moved according to the force *f_i*:

$$\mathbf{p_i} = \mathbf{p_i} + \eta f_i$$

where $\eta = 1/n$.

The advantage of using the relaxation algorithm is that it only requires the distances, and not the actual positions, which are not available in our case. A Kohonen self-organising map would therefore not be applicable on this data (Kohonen (2001)).

2.2 Example

Sensory data have been collected from an AIBO robot performing a slow walk while moving its head continuously from side to side. The recorded sensors are:

- 1-3 distance sensors
- 4-6 head (proprioceptive sensors)
- 7-9 right front leg
- 10-12 right hind leg
- 13-15 left front leg
- 16-18 left hind leg

During the walk, 1000 sensor values have been collected for each of these 18 sensors. Figure 1 shows an example of the development of the distance matrices and the maps using the sensor measurements

¹a perfect mapping given the $n \times n$ information distance matrix is possible in an (n-1)-dimensional space.



Figure 1: Development of distance matrices and corresponding body maps over time. Left: 10 measurements, centre: 100 measurements, right: 1000 measurements. The values in the matrices range from zero (dark blue) to high (red). In the body map on the right, the mapping from the sensors to the position of the sensors on the robot's body is already clearly visible.

of the AIBO robot after 10, 100 and 1000 steps. The 18×18 information distance matrix D is symmetrical with zeros in the diagonal, since $d(X_i, X_i) = 0$ and $d(X_i, X_i) = d(X_i, X_i)$.

In the map of figure 1 right, the arrangement of the sensors in the body map already corresponds roughly to the sensor distribution on the body of the robot. Distance and head sensors are arranged in the upper right half of the map, the knee joints of all four legs on the lower right of the map and all other leg sensors on the left side. The exact map depends on the random initial conditions which are different for each run of the relaxation algorithm, but the maps have comparable structures.

The particular emergent organisation of the map results from the body structure of the robot as well as from the behavioural patterns it conducts in a particular environment. In that sense, such maps can be interpreted as a body image.

3 Interpersonal Maps

3.1 Definition

The concept of a map can be extended to include not only internal proprioceptive sensors but also external sensors such as visual information. This permits to relate in the same format information about the robot's own body with information about other robots perceived through sensors. Let us define the state of the robot R_Y by a vector of size m:

$$Y(t) = (Y_1(t), Y_2(t), \dots, Y_m(t))$$
(2)

A possible formalisation of this situation can be obtained by supposing that the behaviour of the other robot R_Y is perceived through k new sensors in addition to the ones dedicated to proprioception. The new vector X(t) of size n + k can be expressed as below, where g is a potentially complex function linking the state of R_Y (dimension m) to the perceived state of R_X (dimension k).

$$X(t) = (X_1(t), \dots, X_n(t), g_1(Y(t)), \dots, g_k(Y(t)))$$
(3)

In such conditions, a map can be built using the same method as the one described in the previous section. In general, the sensors corresponding to the perceived state of R_Y will not be correlated with the activity of R_X , but they should show separated intracorrelated patterns. In such a case, the body schemas of R_X and R_Y should appear as two distinct clusters in the maps. However in some cases, some intercorrelations could be found between the two sets of sensors. This could be in particular the case when the two robots interact in a closely coupled manner, for instance during a direct imitation task. Such maps can be seen as conceptual signatures for the body correspondence problem. We will now show examples of these two situations.

For the sake of simplicity, we assume in the following examples that g offers a linear mapping linking the sensory states of the observed robot to the states perceived by the observing robot. We will discuss this assumption in the next section.

3.2 Example 1: No Intercorrelation

In this example, we used the sensors recorded from the walking robot together with the sensors of another robot it could have observed. The other robot was sitting and stretching its legs and neck. Altogether, this results in a recording of 36 sensors during 1000 time steps.

Since there is no interaction between the two robots, the two sensor groups are not directly correlated. This results in a higher information distance on average between two sensors of the same robot than between two sensors of different robots. The interpersonal body map in figure 2 therefore shows two clusters. The first cluster can be seen on the lower part of the body map with sensor indices from 1 to 18 printed in black, the second cluster can be seen above the first one with sensor indices from 19 to 36 printed in red. The body schemas within the two clusters are more distorted than the one in figure 1 right due to the interplay of the sensors, but a concentration of the head and distance sensors towards the centre of the map is still visible.

3.3 Example 2: Intercorrelation

This example studies the sensory information of one robot imitating the behaviour of the other. In this case, the robots were walking. The experiment has been performed with imitation with a time delay of 10 recordings which corresponds to about half a second (figure 3). In this case, the interpersonal body map does not show two clusters anymore but shows a mapping between sensors of a similar type. Sensors with indices i and i + 18 are very close to each other on the body map and are plotted in the same colour (e.g. X_1 and X_{19} on the upper right side).

4 Discussion

Our model makes a series of assumptions that can be discussed. The first one is to separate sensors related to proprioception with sensors related to external perception. In practice, such a clear distinction cannot be obtained. Our embodied perception merges both internal and external stimuli without a priori discrimination. However, presenting the model this way helps clarifying the mechanism we describe.

More importantly, we assume that R_X 's perception of the behaviour of robot R_Y can be modelled using a function g mapping the state of R_Y to R_X 's perceptual state. This is a reasonable assumption in the sense that in some way or another the observation of the behaviour of R_Y can be related to its internal state. The fact that relevant information about R_Y 's state can be reconstructed after this function has been applied is potentially more questionable. In our context, what counts is that some intercorrelation between Y and X can still be discovered. For instance if g is a linear transformation, such kind of information will be entirely conserved.

But it is likely that g is a much more complex function. Even in that case intercorrelations could potentially be discovered in several circumstances. One possibility is that R_Y scaffolds the interaction to make its perceived behaviour more tuned to its own internal state. It has been well studied that adults adapt to children in order to make their overt behaviour more easily analysed (Schaffer (1977); Kaye (1982)).

Another possibility is that the biases of g are evaluated by a separated mechanism. More generally, the progressive awareness of self and others is likely to be linked with several other developmental processes. Other embodied developmental models suggest for instance that discrimination based on levels of predictability could play a key role in development of the animate/inanimate distinction and the self/other discrimination (Kaplan and Oudeyer (2005)).

5 Conclusion

Interpersonal maps may offer a possible unified framework accounting for the structure of the agent's body schema as well as a representation of the observed behaviour of another agent. In cases of strong couplings between agents, a "we-centric" space can emerge in which the agent's body structure can be directly mapped onto the structure of an observed body. We strongly believe that the dynamics responsible for self-other distinction are tightly related with



Figure 2: Information distance matrix and interpersonal body map for a robot observing another robot behaving independently.

the ones accounting for the construction of the body schema and that both processes must be studied together. Our future research in developmental robotics will investigate further the conditions for the emergence of this interpersonal space and the possible usage of this information representation in the larger context of robotic control architecture. We also wish to address more precisely the relevance of this mechanism for the development of the self-other matching and discrimination as observed during children's early development.

6 Acknowledgements

Research funded by Sony CSL Paris with additional support from the ECAGENTS project founded by the Future and Emerging Technologies programme (IST-FET) of the European Community under EU R&D contract IST-2003-1940.

References

- J. P. Crutchfield. Information and its metric. In L. Lam and H. C. Morris, editors, *Nonlinear Structures in Physical Systems – Pattern Formation, Chaos, and Waves*, pages 119–130. Springer Verlag, 1990.
- T. Duckett, S. Marsland, and J. Shapiro. Fast, on-line learning of globally consistent maps. *Autonomous Robots*, 12:297–300, 2002.

- V. Gallese. The manifold nature of interpersonal relations: the quest for a common mechanism. In C. Frith and D. Wolpert, editors, *The Neuroscience* of Social Interaction, chapter 7, pages 159–182. Oxford University Press, 2004.
- V. V. Hafner. Cognitive maps for navigation in open environments. In Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS-6), pages 801–808, 2000.
- F. Kaplan and V.V. Hafner. The challenges of joint attention. In L. Berthouze, H. Kozima, C. Prince, G. Sandini, G. Stojanov, G. Metta, and C. Balkenius, editors, *Proceedings of the 4th International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic System*, pages 67– 74. Lund University Cognitive Studies 117, 2004.
- F. Kaplan and P-Y. Oudeyer. The progress-drive hypothesis: an interpretation of early imitation. In K. Dautenhahn and C. Nehaniv, editors, *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions*. Cambridge University Press, 2005. to appear.
- K. Kaye. *The mental and social life of babies*. University of Chicago Press, Chicago, 1982.
- T. Kohonen. Self-Organizing Maps. Extended edition. Springer, Berlin, 2001.
- Y. Kuniyoshi, Y. Yorozu, Y. Ohmura, K. Terada, T. Otani, A. Nagakubo, and T. Yamamoto. From



Figure 3: Information distance matrix and interpersonal body map for a robot being imitated by another robot with a time delay of half a second.

humanoid embodiment to theory of mind. In *Embodied Artificial Intelligence*, pages 202–218. Springer Verlag, 2004.

- A. Meltzoff and A. Gopnick. The role of imitation in understanding persons and developing a theory of mind. In H. Tager-Flusberg S. Baron-Cohen and D.Cohen, editors, *Understanding other minds*, pages 335–366. Oxford University Press, Oxford, England, 1993.
- C. Moore and V. Corkum. Social understanding at the end of the first year of life. *Developmental Review*, 14:349–372, 1994.
- C.L. Nehaniv and K. Dautenhahn. The correspondence problem. In K. Dautenhahn and C.L. Nehaniv, editors, *Imitation in animals and artifacts*, pages 41–61. MIT Press, 2002.
- L. Olsson, C.L. Nehaniv, and D. Polani. The effects on visual information in a robot in environments with oriented contours. In *Proceedings of the Fourth International Workshop on Epigenetic Robotics*, pages 83–88, 2004.
- W. Penfield and T. Rasmussen. *The Cerebral Cortex* of *Man*. Macmillan, New York, 1950.
- D. Pierce and B. Kuipers. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*, 92:169–227, 1997.
- D. M. Pierce. *Map Learning with Uninterpreted Sensors and Effectors*. PhD thesis, The University of Texas at Austin, 1995.

- H. Schaffer. Early interactive development in studies of mother-infant interaction. In *Proceedings of Loch Lomonds Symposium*, pages 3–18, New York, 1977. Academic Press.
- T. Schreiber. Measuring information transfer. *Physical Review Letters*, 85(2):461–464, 2000.
- Y. Yoshikawa, K. Hosoda, and M. Asada. Binding tactile and visual sensations via unique association by cross-anchoring between double-touching and self-occlusion. In *Proceedings of the Fourth International Workshop on Epigenetic Robotics*, pages 135–138, 2004.
- Y. Yoshikawa, H. Kawanishi, M. Asada, and K. Hosoda. Body scheme acquisition by cross map learning among tactile, image, and proprioceptive spaces. In *Proceedings of the Second International Workshop on Epigenetic Robotics*, 2002.

Teaching a Robot to Behave like a Cockroach

THOMAS HELLSTRÖM*

*Department of Computing Science Umeå University Sweden thomash@cs.umu.se

Abstract

Techniques for learning reactive robot behaviors have been an active field of research in robotics for many years. In this paper the method for representing behaviors is based on association rules. Learning the association rules is accomplished by recording training data for a manually programmed controller. The data is then used to generate a set of association rules that replaces the manually programmed controller, and manages to reproduce the demonstrated behavior. Reactive behaviors have obvious limitations, caused by the reactivity itself. Sequences of behaviors are hard to model, unless the switch between behaviors is synchronous with changes in the sensor data. Two ways to get around this limitation are discussed, and the method is demonstrated with examples: one road sign problem with a mix of two wall-following behaviors, and a more complex sequenced light-avoiding cockroach behavior. The results show that association rules are a powerful and practical way to implement rule-based controllers for reactive and semi-reactive robots.

1 Introduction

The work reported in this paper addresses the wellstudied problem of making robots learn reactive behaviors from demonstrations. In general, the process is divided into three steps: 1. The robot is controlled, either by remote control by a human operator, or by a manually coded software controller to perform a certain task. The sensor data S(t) at time t is recorded along with the commanded response signal R(t). 2. The recorded data is used in a modeling, where a control law $B: S(t) \rightarrow R(t)$ is created. 3. The controller B is implemented in the robot, which hopefully manages to perform the demonstrated task autonomously. The various approaches to this general setup can be distinguished by the machine-learning technique chosen to arrive at the control law.

Reinforcement learning (RL) is a commonly used methodology (Lin (1991); Carreras et al. (2002)), which maps the state of the environment to an action that in turn maximizes the accumulated future rewards. The main advantage of RL is that it does not require all data to be available at the same time, as do most other machine-learning techniques. As a result, RL is suitable for online robot learning. The main disadvantages are long learning time, and problems with continuous variables (Carreras et al. (2002)). Artificial neural nets that have also been applied to

the problem of finding reactive behaviors from data. Martin and Nehmzow (1995) use simple single layer perceptrons to represent behaviors for obstacle avoidance, wall following, cleaning, and route learning. Fuzzy rule bases have also been widely chosen to represent learned reactive behaviors. Ward et al. (2000) uses training data for a remote controlled robot to generate a fuzzy rule base capable of reproducing behaviors, such as wall following, corridor following, and docking. Evolutionary techniques have been combined with fuzzy rule bases to find optimal rules, e.g. in Hoffmann and Pfister (1997). In our approach, B is represented by a rule base with association rules. Association rules (Agrawal et al. (1993)) have been successfully used for data mining, where the goal is to explore complex databases to find patterns that might prove useful for various purposes. However, association rules have so far not been extensively used in robotics. One advantage with this machine-learning technique is the handling of uneven distribution of training examples. Most other techniques have a tendency to focus on the most common examples, and learn less from the scarce examples. For example, this is a well known problem when using neural nets for the learning process (Ward et al. (2000)).

The concept of association rules and how they are used to represent reactive behaviors is introduced in Section 2. The general method for building a controller is described in Section 3. Results of practical experiments are presented in Section 4, including a road-sign-following behavior and a more complex cockroach-hide behavior. Section 5 concludes the paper with a summary and conclusions.

2 Behaviors as Association Rules

Association rules are a way of expressing dependencies between items in databases. Association rules have the general form $X \Rightarrow Y$, where both X and Y are sets of items. Given transactions $T \in D$, where D is a database and each transaction is a set of items, the rule $X \Rightarrow Y$ expresses a statistical correlation between X and Y. The rules can be constructed according to different quality measures for different purposes. The *coverage* of the rule $X \Rightarrow Y$ is defined as

$$coverage(X \Rightarrow Y) = cover(X),$$

where cover(X) is defined as the number of transactions containing all items in X, divided by the size of the database. I.e., the *coverage* is the fraction of transactions in the database that contain all items in the left-hand side X of the rule. The *support* measures the fraction of transactions that contain all items in both X and Y:

$$support(X \Rightarrow Y) = cover(X \cup Y).$$

For some applications, the statistical correctness of the correlation is critical. The important measure for this quality is called *strength*. The *strength* (sometimes also called *confidence*) of an association rule $X \Rightarrow Y$ is the proportion of the transactions that contain X that also contain Y. It can be computed as

$$strength(X \Rightarrow Y) = \frac{support(X \Rightarrow Y)}{coverage(X \Rightarrow Y)}$$

Coverage and support are of interest when estimating the significance of the strength, since they quantify on how many observations of X and Y the computation of strength is based. For more information about these and related measures see Hellström (2003a).

2.1 Behavior Representation

The robotics framework in this paper is basically reactive, and each behavior is defined by a control law $B: S(t) \rightarrow R(t)$, where S is the vector of stimuli available at time t (a purely reactive scheme involves only stimuli from the current time t,) and R(t) is the response vector issued at time t. B is implemented as a rule base of rules of the form $S \Rightarrow R$. S is a conjunction of boolean expressions $s_i = v_i$, where s_i is a discretized sensor variable or derived expressions thereof and v_i is an integer value. R has the form y = a, where y is a discretized response variable and a is an integer value. With this notation, a rule has the general form

$$s_i = v_i \wedge s_j = v_j \dots \wedge s_k = v_k \Rightarrow y = a.$$
(1)

In our experiment we have a Khepera robot with 8 infrared sensors $IR_0, IR_1, ..., IR_7$ to measure the distance to the closest obstacle. Each sensor delivers an integer between 0 (corresponding to a distance larger than the sensor range which is about 4 cm.) and 1023 (corresponding to a distance less than about 1 cm.). For experiment 1, each sensor readout IR_i is split into 3 ranges 0, 1, 2. For experiment 2, 4 ranges are used and represented by a discrete variable ir_i according to

$$ir_{i} = \begin{cases} 0 \text{ if } 0 \leq IR_{i} < 100 & (\text{long distance}) \\ 1 \text{ if } 100 \leq IR_{i} < 600 & (\text{medium dist.}) \\ 2 \text{ if } 600 \leq IR_{i} < 900 & (\text{short dist.}) \\ 3 \text{ if } 900 \leq IR_{i} \leq 1023 & (\text{very short} \\ & \text{distance}) \end{cases}$$
(2)

The robot has two wheels with independent motor control, so both robot speed and turning radius are controlled by setting the left and right speed values v_l and v_r . v_l and v_r can be set to integer values in the range [-127,127]. The response y in our experiments is a coded combination of v_l and v_r according to:

y	v_l	v_r	Action
9	0	0	stop
3	-5	5	anti clockwise on the spot
2	0	5	anti clockwise around left wheel
1	2	5	soft anti clockwise
0	5	5	straight ahead
-1	5	2	soft clockwise
-2	5	0	clockwise around right wheel
-3	5	-5	clockwise on the spot
			(3)

As an example, a rule for a left-wall-following behavior may look like this:

$$ir_1 = 0 \land ir_2 = 1 \Rightarrow y = 1.$$

The rule should be interpreted as follows:

if
$$0 \le IR_1 < 100 \land 100 \le IR_2 < 600$$
 then
 $v_l = 2$ and $v_r = 5$.

In plain English this reads as:

if IR_1 senses a long distance and IR_2 senses a medium distance, then turn soft anti clockwise.

3 Building a Controller

The rule base to control the robot is generated from data recorded from a manually programmed controller, demonstrating the required behavior. In this way we obtain a set of stimuli/response pairs that can be used to automatically generate a rule base. This rule base then replaces the controller, and hopefully produces the same behavior as the manually programmed controller. Each sample has the form

$$ir_0, ir_1, \dots, ir_7, y$$
 (4)

where each ir_i is an infrared sensor readout and y is the commanded velocity signals from the manually programmed controller. The rules we want to find have the form defined in (1), where each term is an attribute-value pair of the form s = v, where s is a discretized sensor variable and v is an integer value. Algorithms that efficiently search large databases for association rules have been previously developed (e.g. Agrawal et al. (1993)).

The generated rules are implemented as a controller in the robot. During execution, the sensed data is matched with the left-hand side of the rules. A rule, for which all terms $s_i = v_i$ in the left-hand side match the sensed data, is said to fire. Three cases can occur: 1. Exactly one rule fires. The right-hand side y = a of the rule is used to control the robot. 2. More than one rule fires. The one with the highest strength is chosen. 3. No rule fires. The task of finding a rule for sensor data that lies outside all defined rules can be viewed as a classification problem: to which rule does the sample belong? We have successfully designed and implemented a method called k-nearest rules, based on the classification technique k-nearest neighbors (kNN). For more information, refer to Hellström (2003a).

4 Experiments

We present results from two experiments demonstrating the power of using association rules to model reactive behaviors in the way described in the previous section. The experiments also show how non-reactive behaviors can be tweaked into the reactive framework by pre-processing the sensor data.

4.1 Experiment 1

This experiment deals with the Road sign problem (Linåker and Jacobsson (2001)), in which the robot has to act on a road sign it had passed earlier. It is impossible to achieve this in a purely reactive manner, since the robot has to choose between a left and a right turn, depending on past stimuli. The situation is illustrated in Figure 1.

Our approach is to let the robot act on preprocessed sensor data with a perceptual decay (Werger (1999)). The perception of a road sign remains even after the stimuli have disappeared and slowly fades out with time. In this way the behavior can still be purely reactive, since the memory is hidden in the robot's perception. This is indeed a simplification of the original road sign problem, but it serves our purpose well. The purpose of the experiment is to see how a complex behavior can be modeled by the rule base of automatically generated association rules. The idea with perceptual decay is illustrated in Figure 2. The original stimuli as a function of time are shown in the lowermost pane. The perceptual decay in the middle pane shows how the perception remains and gradually decays after the original stimuli has disappeared. The uppermost pane shows another processing of the stimuli used in experiment 2.

The demonstrated behavior is manually coded as a switching between two controllers, a left-wall follower and a right-wall follower. The switching occurs when the robot encounters a road sign, describing the recommended way to go in the upcoming junction. The road signs are constructed of small bulbs attached to the walls of the robot's maze. The bulbs on the wall are sensed by the ambient light sensors on the Khepera robot. The sensors for left and right bulb detection are denoted AL_l and AL_r respectively. To enable the robot to act on a road sign that appears and disappears before a junction, a virtual road sign sensor RSis defined as:

$$RS = \begin{cases} 2 \text{ if } decay(AL_l) > decay(AL_r) \\ 0 & otherwise \end{cases}$$
 (5)

The decay function computes a perceptual decay of the sensed road sign signal, and serves to make the robot gradually forget about road signs as time passes after the road sign has disappeared out of the robot's sight. The RS sensor is a binary signal with the value 2 if the last seen road sign was a left sign, and 0 otherwise. The perceptual decay is a slight side step from a pure reactive design, but is a neat way of stretching the borders of the reactive paradigm when the robot's action has to depend on "old" sensor data. In our example, the RS signal is added to the 8 infrared sen

sors $ir_0, ir_1, ..., ir_7$ as an additional input, and serves as a switch between the two wall-followers in the learning mode. The manually programmed controller performs a left/right wall-following task as described by the pseudo code below:

where *left-wall follower* and *right-wall follower* are simple rule-based controllers described in Hellström (2003b). In step 2 of the basic learning process (see Section 1,) RS is made available as an extra input in the search for association rules, and should then (automatically) be added as a high-level condition that groups the generated rules in two categories: left-wall following and right-wall following. Of course, rules common to both behaviors may be unaffected by the value of the RS input.

The controller is run with a cycle time of 0.1 seconds for 100 seconds. This results in 1000 samples of training data, each sample consisting of 8 discretized sensor read-outs ir_{0-7} , RS and one discretized action y. The sensor data is discretized in 3 ranges, and the actions in 8 categories as described in (3) (the stop action is never used in this example.) The training data is then used to automatically generate association rules, which in turn are used to construct a robot controller.

Table 1 shows performance for a number of different controllers with different numbers of rules. The number of rules is set by giving a lower limit to the strength value. Each controller is evaluated on one row in the table. The rules are applied to two data sets, the 1000 samples big training data set, which was used to generate the rules, and a test data set separately generated. The e_{tr} and e_{te} are the fractions of samples that give incorrect action when compared to the manually programmed controller. By demanding a strength value equal to 1.0, 31 rules are selected. The column labeled 0rule% is the fraction of samples, for which no matching rule can be found in the controller's database. The 31 rule controller leaves 9.6 % of the samples not matched by any rule. The 1-nearest rule developed in Hellström (2003a) handles this reasonably well with 5.0 % incorrect actions on the test data set. The column labeled 1rule% is the fraction of samples covered by exactly one rule. The rightmost 3 columns are the fractions of samples covered by 2, 3 and more than 3 rules respectively. In 23.9~% of the cases, two or more rules fire at the same time. This is resolved by majority voting among the rules that fire. It is clear from the table that the best controller is achieved by a controller with the 43 rules with *strength* ≥ 0.95 . These rules give minimum error on both training and test data sets. Furthermore, the number of cases where no rule fires is reduced to zero when these 43 rules are used.

A comparison between the training set error e_{tr} and test set error e_{te} exhibits a difference that would normally be diagnosed as overfitting. This concept is largely ignored in the association rule community (Freitas (2000)), while it is very common in other areas of machine learning. However, acting on rules with very low *strength* or *support* corresponds to adding more nodes to a neural net, or adding higher-degree terms to a polynomial model. Simple techniques, such as computing performance for both training data and previously unseen test data should therefore be a standard procedure when using association rules for prediction or induction, in particular with noisy data, such as robot applications.

Table 2 lists a few of the generated rules and shows that not all rules responsible for turning contain the RS variable as condition on the left-hand side of the rule. However, this is not necessarily incorrect, since turning may occur not only when performing a turn in a junction, but also for wall-avoidance, which could be handled uniformly, regardless of the road sign condition. When the rules are installed as a controller on a real Khepera robot, the robot successfully manages to switch between left and right-wall following depending on road marks placed along the route in the maze. For a more detailed analysis of the road sign experiment, see Hellström (2003a) and Hellström (2003b).

4.2 Experiment 2

This experiment aims at developing a rule base capable of mimicking the behavior of an imagined lightavoiding cockroach. A program performing the following robot behaviors is first developed (refer to Figure 3):

- If the light is switched off, explore the surroundings while avoiding obstacles
- If the light is switched on, perform the following sequence: 1. Turn around 180 degrees. 2. Move in a straight line to a wall. 3. Follow the wall until a hiding place is found. 4. Turn around and stop until the light is switched off.

This is a fairly challenging task even for a human robot programmer. In reality it took many days to construct a program able to successfully perform all the described steps with the Khepera robot. It is clear that a pure reactive approach is not enough to achieve step 1 above. For this reason pre-processing of the ambient light sensor is introduced. The layout of this "habituation" function is illustrated in the topmost pane of Figure 2. The initial response follows the actual stimuli (the ambient light sensor) but falls off after a fixed time. In this way it is possible to model a time limited response in a semi-reactive fashion. The actual behavior is purely reactive but the pre-processing is not. It should be noted that the time for the response to fall off is tailored to match the time it takes for the robot to turn 180 degrees, i.e. to complete sub-behavior 1 above. This is necessary to make a reactive modelling possible, but may at first look like cheating. However, both animals and humans exhibit such tailored perception for various behavioral support. And after all, already the choice of sensors for a robot dictates which behaviors are feasible for the robot. The rest of the behavior described above can be programmed in a purely reactive fashion, using the infrared sensors to identify a hiding place and to turn around. The programmed controller is run with a cycle time of 0.1 seconds for 1000 seconds. This results in 10000 samples of training data. The data is then used to generate a rule base, which manages to reproduce the entire cockroach-like behavior using 100 association rules. In particular, the method manages to reproduce the time-limited turn in step 1 by including the pre-processed ambient-light sensor in the rules controlling the rotation.

5 Summary

We have demonstrated how association rules can be used by intelligent robot controllers for learning reactive and semi reactive behaviors. Two techniques to extend the reactive paradigm in this context have been presented. One road-sign-following task uses perceptual decay to achieve a memory of the type of the latest road sign. Another pre-processing of sensor data introduces habituation and makes it possible to implement a sequenced light-avoiding cockroach behavior.

References

Rakesh Agrawal, Tomasz Imielinski, and Arun N. Swami. Mining association rules between sets of items in large databases. In Peter Buneman and Sushil Jajodia, editors, *Proceedings of the 1993* ACM SIGMOD International Conference on Management of Data, pages 207–216,1993.

- M. Carreras, P. Ridao, J. Batlle, and T. Nicosevici. Efficient learning of reactive robot behaviors with a neural-q learning approach. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002.
- Alex Alves Freitas. Understanding the crucial differences between classification and discovery of association rules - a position paper. *SIGKDD Explorations*, 2(1):65–69, 2000.
- T. Hellström. Association Rules for Learning Behavioral Mappings in Robotics. Technical Report UMINF-03.12 ISSN-0348-0542, Department of Computing Science Umeå University, Umeå Sweden, 2003a.
- T. Hellström. Learning robotic behaviors with association rules. In Nikos Mastorakis, editor, *WSEAS transactions on systems*, 2003b. ISBN 1109-2777.
- F. Hoffmann and G. Pfister. Evolutionary design of a fuzzy knowledge base for a mobile robot. *Int. J. of Approximate Reasoning*, 17(4):447–469, 1997.
- Long Ji Lin. Programming robots using reinforcement learning and teaching. In AAAI, pages 781– 786, 1991.
- Fredrik Linåker and Henrik Jacobsson. Mobile robot learning of delayed response tasks through event extraction: A solution to the road sign problem and beyond. *International Journal of Computational Intelligence and Applications*, 4(1):413– 426, 2001.
- Paul Martin and Ulrich Nehmzow. Programming by teaching: Neural network control in the manchester mobile robot. In *Proceedings Intelligent Autonomous Vehicles*, 1995.
- Koren Ward, Alexander Zelinsky, and Phillip McKerrow. Learning robot behaviours by extracting fuzzy rules from demonstrated actions. *Australian Journal of Intelligent Information Processing Systems*, 2000.
- Barry Brian Werger. Cooperation without deliberation: A minimal behavior-based approach to multirobot teams. *Artificial Intelligence*, 110(2):293– 320, 1999.

Table 1: Performance for road sign controller. Majority voting is used when more than one rule fires. The error rate is much higher than for the simple wall following task. The difference between the training error e_{tr} and test error e_{te} is an indication of overfitting.

Strength	#rules	\mathbf{e}_{tr} %	\mathbf{e}_{te} %	0rule%	1rule%	2rules%	3rules%	>3rules%
1.00	31	1.0	5.0	9.6	66.5	18.9	3.2	1.8
0.98	33	0.9	4.9	6.5	66.8	20.7	4.0	1.9
0.95	43	0.5	2.4	0.0	26.8	46.9	10.2	16.1
0.90	56	3.1	4.7	0.0	7.7	44.4	21.1	26.8
0.85	66	4.7	6.0	0.0	7.1	26.8	31.1	35.0
0.80	76	4.2	6.2	0.0	6.2	23.2	2.8	67.8
0.75	87	7.1	9.9	0.0	5.9	0.6	9.1	84.4
0.70	97	6.0	9.1	0.0	5.9	0.3	6.4	87.4



Figure 1: The road sign problem, adapted from Linåker and Jacobsson (2001), in which the robot has to decide on a left or right turn in each junction, depending on the past stimulus from the road signs. Our approach is to add a perceptual decay to the road sign perception. The robot switches between a left- and right-wall following behavior to perform the turnings in the crossings.
Rule	No.	Coverage	Support	Strength
$ir_1 = 2 \wedge ir_2 = 2 \Rightarrow y = -3$	1	4	4	1.00
$ir_0 = 1 \wedge ir_1 = 2 \Rightarrow y = -3$	2	4	4	1.00
$ir_1 = 2 \land RS = 2 \Rightarrow y = -3$	3	6	6	1.00
$ir_0 = 1 \wedge ir_2 = 1 \wedge ir_6 = 1 \Rightarrow y = -3$	4	3	3	1.00
$ir_1 = 1 \land RS = 2 \Rightarrow y = -3$	5	67	67	1.00
$ir_2 = 2 \wedge ir_3 = 0 \Rightarrow y = -3$	6	2	2	1.00
$ir_2 = 1 \land ir_6 = 2 \land ir_7 = 1 \Rightarrow y = 3$	7	2	2	1.00
$ir_4 = 2 \wedge ir_6 = 1 \Rightarrow y = 3$	8	7	7	1.00
$ir_4 = 2 \wedge ir_7 = 1 \Rightarrow y = 3$	9	2	2	1.00
$ir_2 = 1 \land ir_4 = 1 \Rightarrow y = 3$	10	26	26	1.00
$ir_1 = 1 \land ir_4 = 2 \Rightarrow y = 3$	11	4	4	1.00
$ir_0 = 1 \land ir_4 = 2 \Rightarrow y = 3$	12	2	2	1.00
$ir_4 = 2 \land RS = 0 \Rightarrow y = 3$	13	25	25	1.00
$ir_3 = 2 \land RS = 0 \Rightarrow y = 3$	14	32	32	1.00
$ir_0 = 2 \wedge ir_1 = 0 \wedge RS = 2 \Rightarrow y = 0$	15	94	94	1.00
$ir_4 = 0 \land ir_5 = 2 \land RS = 0 \Rightarrow y = 0$	16	148	148	1.00
$ir_1 = 2 \land RS = 0 \Rightarrow y = -1$	17	8	8	1.00
$ir_1 = 2 \wedge ir_2 = 0 \Rightarrow y = -1$	18	6	6	1.00
$ir_0 = 1 \wedge ir_5 = 1 \wedge ir_6 = 1 \Rightarrow y = -1$	19	3	3	1.00
$ir_0 = 1 \land ir_6 = 1 \land RS = 0 \Rightarrow y = -1$	20	3	3	1.00
$ir_0 = 1 \land ir_7 = 1 \land RS = 0 \Rightarrow y = -1$	21	7	7	1.00
$ir_0 = 2 \land RS = 0 \Rightarrow y = -1$	22	13	13	1.00
$ir_3 = 0 \land ir_5 = 1 \land ir_6 = 0 \land RS = 0 \Rightarrow y = -1$	23	116	116	1.00
$ir_4 = 0 \land ir_5 = 1 \land RS = 0 \Rightarrow y = -1$	24	142	142	1.00
$ir_4 = 1 \land RS = 2 \Rightarrow y = 1$	25	13	13	1.00
$ir_1 = 0 \land ir_2 = 2 \land RS = 2 \Rightarrow y = 1$	26	101	101	1.00

Table 2: Part of generated rule base for road sign controller. The binary RS variable controls left- and right-wall following.



Figure 2: Two ways of introducing non reactivity by pre-processing of sensor data. The perceptual decay enables extended response to a stimulus. The habituation enables a sequence of two behaviors as a response to a stimulus.



Figure 3: A robot emulating a cockroach's light-avoiding behavior. Between 1. and 2. the light is off and the robot moves around randomly, while avoiding obstacles. At 2 the light is switched on and the robot turns 180 degrees, moves until it hits a wall, which it follows until it reaches a hiding place, where it turns around and stops.

Goal level imitation with emergent repertoires.

Bart Jansen*

*Vrije Universiteit Brussel Pleinlaan 2, 1050 Brussels, Belgium bartj@arti.vub.ac.be

Abstract

In this document we present computer simulation experiments on goal level imitation. We investigate how the study of imitation at the population level can be extended from action level to goal level imitation. We provide an imitation game in which agents learn to detect intentions in the actions of other agents. Agents then imitate the goals of other agents behaviour rather than their exact actions.

1 Introduction

Imagine a mother painting a wall using a paint roller. When her youngest daughter observes her mother paint, the child wanders around and sees her doll on the ground. She takes the doll and starts to rub the doll on the wall. The other daughter enters the room and sees mother and daughter "painting" the wall. She joins them, wanders around and finds a brush on the ground. Using the brush, she starts painting as well.

Both children imitate their mother's painting. The youngest child imitates the exact actions she sees her mother perform, although she uses the wrong tools. The other child does not imitate the exact actions, which is not possible since she's using a brush. However, her actions will have the same effect as her mother's actions.

This example¹ clearly illustrates the difference between action level and goal level imitation (Byrne and Russon, 1998). Goal level imitation is similar to effect level imitation (Demiris and Hayes, 1997), in which the effect on the play board rather than the intention is imitated.

In the field of robotics, it was soon recognized that imitation is a powerful mechanism to learn robots to execute a wide range of tasks (Billard and Hayes, 1997; Dautenhahn and Nehaniv, 2002; Kuniyoshi et al., 1994). An agent based perspective on imitation was proposed in order to be able to answer 5 essential questions on imitation (Dautenhahn and Nehaniv, 2002): Who, when, what and how to imitate and what makes a successful imitation. Often, those five issues were studied in interactions between a single teacher and a single student, in which the teacher starts with a set of skills and transfers them to the student. Later on, the importance of a population approach was stressed. It was shown how students can act as teachers, once matured enough. Even cyclic interaction patterns between a set of agents were studied (Alissandrakis et al., 2004).

We have argued in previous work (Jansen, 2003; Jansen et al., 2003, 2004) that in a true population approach, roles can not be fixed, nor can the interaction pattern be fixed. In other words, at any time any agent randomly selected from the population can imitate any other agent randomly selected from the same population.

In this document, we follow the same rationale to investigate how agents can learn to imitate intentional behavior. Little work has been done on goal level imitation in robotics. We believe that ideas on goal level imitation from developmental psychology might enrich the learning by demonstration research in robotics.

Detecting intentions in other agents' actions is a hard problem. Since there is a many-to-many relation between actions and intentions, the actions themselves do not provide enough information to extract intentions from it. Therefore, other (external) information must be taken into account by the agent in order to detect the intention of the behaviour. External information might include prior knowledge about the agent performing the action and cues in the context (Meltzoff, 1995; Searle, 1984). These bits of extra information constrain the huge search space of possible intentions of the agents actions. Supporting evidence was found in experiments with 9 months old infants. Even at this very young age they seem to use information external to the observed behaviour to interpret the intentionality of the behaviour (Woodward et al., 2001).

¹adapted from the well known example in (Nehaniv and Dautenhahn, 2000).

In this document a computer simulation experiment is proposed in which agents learn to imitate the intentions of the behaviour of other agents. The computer simulations serve as a proof of concept before transferring the goal level imitation to physical robots. We have successfully deployed this approach before in our research on action level imitation. Now the intentions, rather than the precise actions are imitated.

2 Imitating goals

Instead of considering a single teacher and a single student and the imitative process by which the student learns from the teacher, we consider a population of agents in which agents do not have fixed roles. The smallest population consists of only two agents. In previous work, we have reported successful imitation in populations containing up to 50 agents. The experimental setup for this work differs from previous experiments on action level imitation in that agents now can manipulate objects. In order to allow them to do so, every agent has its own play board it can observe and manipulate. Agents can also observe other agents perform actions on their boards.

		Obj-2	
Obj-1			
	Obj-3		

Figure 1: An example of the play board the agents can interact with.

An agent's play board is a simple two dimensional blocks world. Manipulation is done by a limited set of action primitives. These primitives include for instance moving a block one cell to the left, one to the right, ... Every agent has its own set of primitives, thus allowing for heterogeneous populations. In the work presented here, we assume homogeneous populations. We do not study the emergence of (possible shared) action primitives, so action primitives are built in or acquired earlier in this work.

Agents can observe the manipulations on each other's play board. They can recognize simple relations between objects like *left-of?* and *above?*. We assume that all agents have the same built in set of

relations they can detect. Agents can categorize their play board by concatenating such relations.

As the goals the agents have in this experiment are expressed as properties of the play board, goals can be represented by a concatenation of these relations as well. For an agent, a goal represents the desire to obtain a play board satisfying certain conditions. For instance, suppose that the goal of the agent is

```
(AND (above? obj-1 obj-2)
```

```
(left-of? obj-3 obj-2)
(left-of? obj-1 obj-3))
```

and that the play board is the one depicted in figure 1, then the agent could for instance execute the following sequence of actions

```
(AND (up obj-1) (down obj-2))
```

in order to obtain a play board in which the specified goal holds.

It is important to notice that every agent acts upon its individual play board. The experimental setup was designed deliberately in that manner to illustrate that the mere copying of actions would not result in successful imitation. Thus, if agents simply copy the actions the other agents perform, the imitative success is always very low as it is very unlikely to obtain the same goal with the same set of actions, simply because the initial configurations of the blocks are different. Having separate play boards for all agents is not crucial and does not hinder social learning: goals are learnt by observing the other agents' actions in their own play board.

2.1 Interactions

The interactions between agents are defined by imitation games. An imitation game is played with only two players. These are selected randomly from the population. Many games are played in the population, such that every agent interacts with many others.

The imitation game starts by randomly selecting two agents from the population. One agent will take the role of *initiator*, while the other agent takes the role of *initiator*. The initiator randomly selects a goal from its repertoire and builds a plan for reaching this goal. Actions are performed according to this plan. The imitator observes this sequence of actions and selects from its own repertoire the goal that best matches the perceived actions. The imitator now builds a plan for this goal and acts accordingly. The initiator observes this action sequence and verifies whether its initial goal holds in the resulting play



Figure 2: A cycle in the imitation game: the initiator selects a goal, builds a plan for it and executes appropriate actions. The imitator observes the action sequence and categorizes it in terms of its own goals. The imitator builds a plan for the categorized goal and acts accordingly.

board of the imitator. If that is the case, the game succeeds, otherwise it fails. If any of the previous steps of the game failed, the game fails as well. At the end of the game the initiator sends a single bit of feedback to the imitator, such that both agents know the outcome of the game. Both agents update their repertoires, depending on the outcome of the game. The overview of the game is shown in figure 2. Pseudo code for this interaction pattern is given in table 1. In the pseudo code E stands for the current configuration of the play board, C is the set of competing goals, while G is the repertoire of learnt goals.

2.2 Components

The imitation game requires the agents to have multiple capabilities. Agents need to have a repertoire to store goals, they need a planner, the ability to recognize the goal of the other agent's actions, a learning mechanism and the ability to observe and perform actions. For each of these skills, we have implemented a very simple approach allowing us to investigate a minimalistic set up in which a simple kind of goal level imitation can be demonstrated.

2.2.1 Repertoires of goals

The agent maintains a repertoire of goals. Initially, this repertoire is empty. During the imitation game, the agent can add new goals to this repertoire and delete unsuccessful goals. Therefore, success and usage counters are associated with every goal.



Figure 3: Interpretation of the observed actions as intentional behaviour. Every box represents a goal from the repertoire, together with its score.

2.2.2 Planner

This component constructs a plan for achieving a given goal, starting from a given play board. In our implementation, we use A^* search with a simple heuristic function, including checks for repeated states. If the search takes too long, it is aborted and the planning process fails. Even in this simple blocks world, the search space is huge. Suppose an m-by-nblocks world, containing k blocks. If an agent has a repertoire of l action primitives, every block can be moved by *l* actions, if they don't move the block outside the play board. The bigger the play board, the less likely this is. With k blocks, this results in a branching factor of at most kl. If the agents for instance can move a block in any of the four directions, the branching factor is kl - 1 (since one action always restores the previous configuration) with $l = 4 - \frac{2}{m} - \frac{2}{n}.$

2.2.3 Goal recognition

It is obvious that recognizing goals in behaviour requires the agents to be able to interpret the actions of the others as intentional behavior. Just as humans seem to use external information in this process, the agents do so as well. They assume that the other agents' cognitive capabilities are the same as their own. For instance, agents perform only relevant actions and do not deliberately confuse other agents with their behaviour. They explicitly assume that other agents do that as well. If an agent for instance wants to put a block A on top of a block B, no other blocks than A and B are manipulated unless this is required in order to be able to put A on top of B.

Moreover, agents maintain a repertoire of known goals. Any observed behavior is categorized as one of the stored goals. As the repertoire of known goals is only a very limited subset of all possible goals in the agents' play board, this severely constrains the interpretation process. Learning is done by extending and restricting the repertoire of known goals, or by modifying goals in the repertoire.

The agent observes the actions performed by the other agent as a sequence of play boards. In order to recognize the goal of these actions, the agent selects the best matching goal from its own repertoire of goals. For every goal in its repertoire the agent verifies whether the goal holds in the last play board of the action sequence and does not hold in the first play board of the action sequence. If such goals exist, the most successful goal is selected from this subset. All other goals in this subset are called "competing goals". This process is illustrated in figure 3.

2.2.4 Learning mechanism

Learning consists of several components: addition of new goals, deletion of unsuccessful goals and updating of the scores of goals. At the beginning of every game agents add new random goals to their repertoires with a small probability. When the imitator fails to detect the goal of the initiator's action, the game fails. In such case, the imitator constructs a random goal that is fulfilled in the observed play board and adds it to its repertoire. Goals that were not successful in the past are removed from the repertoire.

A usage counter and a score are associated with every goal. The usage counter of a goal is increased whenever an agent tries to fulfil this goal during the imitation game. Maintaining scores for every goal serves two purposes. It allows us to discriminate successful from unsuccessful goals, such that unsuccessful goals can be removed from the repertoire. Secondly, the score is used in the categorization of observed actions into goals. From all goals that hold for a particular sequence of play boards, the most successful is selected as the category of the observed action. Thus, by influencing the scores of the goals, categorization is influenced directly. If imitation was successful, we want to reinforce the goals that were used, thus we increase their success scores. This will increase the probability that those goals that were successful in previous games are used in later categorization. By decreasing the scores of all competing goals, this effect is further enhanced. When the game fails, the opposite reasoning is followed. Since the goal that was used does not lead to successful imitation, the probability of using it in similar cases should decrease; therefore its score is decreased, while the scores of the competing goals are increased. A threshold function is used to ensure that the scores of all goals remain between 0 and 1. Pseudo-code of these processes is listed in table 2.

Initiator	Imitator	
$E_1 \leftarrow random play board$	$E_2 \leftarrow random \ play \ board$	
$C_1 \leftarrow \emptyset$	$C_2 \leftarrow \emptyset$	
if $G_1 = \emptyset$ new-goal (G_1)	if $G_2 = \emptyset$ new-goal (G_2)	
$g \leftarrow \text{random from } G_1$		
$p_1 \leftarrow \text{build plan for } g$		
$E_1 \leftarrow \mathbf{execute} \ p_1 \ \mathrm{on} \ E_1$		
	observe action sequence A_2	
	$g_{rec} \leftarrow \text{goal from } G_2 \text{ that}$	
	best matches A_2 for E_1	
	$C_2 \leftarrow \text{all other matching}$	
	goals	
	$p_2 \leftarrow \text{build plan for } g_{rec}$	
	$E_2 \leftarrow \mathbf{execute} \ p_2 \ \mathrm{on} \ E_2$	
observe play board E_2		
if g holds in E_2		
send feedb. "success"		
else		
send feedb. "failure"		
$update(g, feedb., G_1, C_1)$	$update(g_{rec}, feedb., G_2, C_2)$	
do-other-updates()	do-other-updates()	

Table 1: Pseudo code of the imitation game for goal level imitation. The pseudo code for the update functions and some auxilaries is given in table 2.

3 Experimental setup and results

The experiments reported in this paper were performed in simulation. The block worlds of the agents consisted of a grid of 5 by 5 cells. Three objects were present in the world, placed initially at random positions. Agents can categorize the world they observe by using the predicates left-of?(obj-1,obj-2) and above?(obj1,obj2). All agents in this first experiment have the same four action primitives: *move-up(obj)*, *move-down(obj)*, *move-left(obj)* and *move-right(obj)*. An experiment with only two relations in a 5-by-5 grid world with only three objects and with only four actions might seem too much a toy problem. Although the total number of representable goals with length less than L is finite in this set up, it is huge. The number of goals that are actually learnt by the agents is only a fraction of this search space. In every game the probability of adding a random new goal (*addprobability*) is 0.05. Goals that were used more than 7 times with a score of less than 0.3 were removed from the repertoires. Δ was set to 0.05.

The experiments were repeated ten times. There-

update(g, signal, G, Competing)		
$g.usage \leftarrow g.usage + 1$		
if signal = "success"		
$g.score \leftarrow \sigma(g.score + \Delta)$		
$\forall g' \in \text{Competing } \mathbf{do} \ g'.score = \sigma(g'.score - \Delta)$		
else		
$g.score \leftarrow \sigma(g.score - \Delta)$		
$\forall g' \in \text{Competing do } g'.score = \sigma(g'.score + \Delta)$		
new-goal(G)		
$g \leftarrow random goal$		
$G \leftarrow G \cup g$		
$\left(\begin{array}{cc} 0 & \text{if } x < 0 \\ 0 & \text{if } x < 0 \end{array}\right)$		
$\sigma(x) = \begin{cases} 1 & \text{if } x > 1 \\ x & \text{if } 0 < x < 1 \end{cases}$		
$\begin{array}{c} x & 110 \leq x \leq 1 \\ \hline \end{array}$		
do-other-updates(G)		
$\forall g \in G \ \mathbf{do}$		
if g.score < *throwawaythreshold* and		
g.usage > *minimumuses*		
$G \leftarrow G \setminus g$		
with probability *addprobability * $\operatorname{\mathbf{do}}\operatorname{\mathbf{new-goal}}(G)$		

Table 2: Update procedures in the intentional imitation game.

fore results show average values and 95% CI. The graphs in figure 4 shows the imitative success and the number of goals that a population of only two agents developed during 10000 imitation games. In figure 5, the same results are shown for a population consisting of ten agents. The imitative success is calculated as the fraction of successful imitation games over the last 100 games. Results show how the repertoires of goals steadily increase, while the imitative success slowly increases up to 70%.

Whether the population contains two or ten agents, imitative success remains about 70%, suggesting that this type of imitation game scales well with larger population sizes. The number of learnt goals however in populations with ten agents is only one-fifth of the number of goals in populations with two agents, simply because in the larger population on every agent only participates in one-fifth of the games on average.

Results should be compared to a baseline experiment in which every agent in the population starts with a pre-programmed and fixed repertoire of random goals. In that case, the imitative success is 50% (verified experimentally). The clearly above random success ratio combined with the fact that a repertoire of goals steadily emerges, proves that the agents are developing a repertoire of goals that enables them to successfully imitate intentional behaviour.

However, there are two important issues that can not be neglected. The success ratio of 70% is much



Figure 4: Imitative success and number of goals developed in an imitation game between two agents over 10000 games.



Figure 5: Imitative success and number of goals developed in an imitation game between ten agents over 10000 games.

below the success ratio of 95% observed in imitation of actions, as shown in previous work. The lower success ratio is not a surprise, since the problem is much harder than learning to imitate simple actions. Moreover, the learning mechanism that is used is very simple. Learning consists of three components in this case: adding random goals, deleting unsuccessful scores and keeping a scoring mechanism. However, once created, a goal is never modified.

Secondly, the number of learnt goals seems to increase without bond. Relating the results back to previous work on emergent repertoires of actions, this can easily be explained. In learning actions, the agents learn to discretise the continuous action space by inventing a repertoire of categories for the actions. In that case, the number of learnt actions tends to increase without bound as well. However, by introducing the concept of a merging operator, which merges actions that resemble each other too much, the number of actions stabilizes after a while. A good merging operator was however not yet established for the current work since there is no continuous space anymore with an associated distance function. Moreover, noise on the perception and execution of actions resulted in actions that could not be distinguished from each other, even though they are different. In the simulation experiments presented here, no noise is introduced in the simulation.

4 Discussion

The imitation game as described in this document provides a computational simulation of goal level imitation, i.e. not the actions but the goals of the actions of the agents are imitated. One could however argue that effects instead of goals are imitated in this work. It is true that goals are represented as relations among objects; however they are not purely effects. Agents categorize the play boards they observe in terms of known goals. This entails that the agents will not perform the same actions, neither will they imitate by obtaining the same blocks configuration, i.e. the same effect. Any configuration of blocks for which the goal holds, will result in success.

As this is work in progress, results are preliminary and both the precise dynamics of the game and the learning mechanism are subject to change. However, this document shows how a population of agentsstarting with empty repertoires-succeeds in constructing a repertoire of shared goals. The repertoires are constructed by local interactions only; agents are not endowed with telepathic capabilities. Sharedness of the repertoires of goals, which is thus an emergent property, means that the repertoires of the different agents are sufficiently similar to enable successful imitation. By constructing these shared repertoires of goals, the agents gradually learn to successfully recognize goals in the other agent's behavior. Since we provide the agents with a built-in planning module, agents can successfully imitate goals as soon as they can recognize them.

5 Conclusion

In this document we have presented preliminary results on simulation experiments of imitation of intentional behaviour. Through imitation, a population of agents learns to recognize the intentions in the actions of the others. Gradually, they develop a repertoire of goals they can detect and pursue. The repertoire grows by creating random goals and by creating goals when imitative attempts fail. By maintaining a scoring mechanism and by removing unsuccessful goals, only those goals that can be observed, pursued and discriminated remain. Every agent has its own different play board it acts on, such that simple action level imitation is guaranteed to fail.

Acknowledgments

Bart Jansen is sponsored by a grant from the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT).

The work described in this paper was partially conducted within the EU Integrated Project COGNIRON ("The Cognitive Companion") and funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020

References

- Aris Alissandrakis, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. Towards robot cultures? - learning to imitate in a robotic arm test-bed with dissimilar embodied agents. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, 5(1):3–34, 2004.
- Aude Billard and Gillian Hayes. Transmitting communication skills through imitation in autonomous robots. In Andreas Birk and Yiannis Demiris, editors, *Proceedings of EWLR97, Sixth European Workshop on Learning Robots, Lecture Notes on Artificial Intelligence*, volume 1545, pages 79–95. Springer, 1997. Brighton, UK, July 1997.
- Richard W. Byrne and A. E. Russon. Learning by imitation: A hierarchical approach. *Behavioral and Brain Sciences*, 21:667 – 684, 1998.
- Kerstin Dautenhahn and Chrystopher Nehaniv. The agent-based perspective on imitation. In Kerstin Dautenhahn and Chrystopher Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, Cambridge, Mass., USA, 2002.
- J. Demiris and G. Hayes. Do robots ape? In *Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents.*, pages 28–31, MIT, Cambridge, Mass., 1997.
- Bart Jansen. An imitation game for emerging action categories. In Wolfgang Banzhaf, Thomas

Christaller, Peter Dittrich, Jan Kim, and Jens Ziegler, editors, *Proceedings of the 7th European Conference on Artificial Life, Lecture Notes in Artificial Intelligence*, pages 800–809, Berlin, 2003. Springer.

- Bart Jansen, Bart de Vylder, Bart de Boer, and Tony Belpaeme. Emerging shared action categories in robotic agents through imitation. In *Proceedings of the Second International Symposium on Imitation in Animals and Artifacts.*, pages 145–152, 2003.
- Bart Jansen, Tom Ten Thij, Tony Belpaeme, Bart De Vylder, and Bart de Boer. (to appear, 2004) imitation in embodied agents results in self-organization of behavior. In *Proceedings of The Eighth International Conference on the Simulation of Adaptive Behavior*, Los Angeles, USA, 2004. The MIT Press, Cambridge.
- Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *T*-*RA*, 10:799–822, 1994.
- A. N. Meltzoff. Understanding the intentions of others: Re-enactment of intended acts by 18-monthold children. volume 31, pages 838–850, 1995.
- C. Nehaniv and K. Dautenhahn. Of hummingbirds and helicopters: An algebraic framework for interdisciplinary studies of imitation and its applications. In J. Demiris and Andreas Birk, editors, *Interdisciplinary approaches to robot learning, World Scientific Series in robotics and intelligent systems, volume 24*, pages 136–161. World Scientific Press, 2000.
- J.R. Searle. *Minds, Brains, and Science*. Harvard University Press, 1984.
- Amanda L. Woodward, Jessica A. Sommerville, and Jose J. Guajardo. How infants make sense of intentional actions. In Bertram F. Malle, Louis J. Moses, and Dare A. Baldwin, editors, *Intentions and Intentionality; Foundations of Social Cognition*, pages 149–170. MIT Press, Cambridge, Mass., USA, 2001.

Hierarchies of Coupled Inverse and Forward Models for Abstraction in Robot Action Planning, Recognition and Imitation

Matthew Johnson & Yiannis Demiris

*Imperial College London Exhibition Road, London, SW7 2BT {matthew.johnson,y.demiris}@imperial.ac.uk

Abstract

Coupling internal inverse and forward models gives rise to on-line simulation processes that may be used as a common computational substrate for action execution, planning, recognition, imitation and learning. In this paper, *multiple* coupled internal inverse and forward models are arranged in a *hier-archical* fashion, with each level of the hierarchy interacting with other levels through top-down and bottom-up processes. Through experiments involving imitation of a human demonstrator performing object manipulation tasks, this architecture is shown to equip a robot with a multi-level *motor abstraction* capability. This is then used to solve the correspondence problem in action recognition. The architecture is inspired by biological evidence.

1 Introduction

Research has shown the direct involvement of the human motor system when observing, as well as imitating, actions performed by others (Meltzoff and Decety, 2003). This, along with the neuroscientific discovery of mirror neurons in area F5 of the macaque monkey premotor cortex, which respond when both performing and observing the same action (Rizzolatti et al., 2002), has led to the proposition of a mirror system underlying the recognition and understanding of behaviour (Fadiga and Craighero, 2003). This system is compatible with the simulation theory of mind-reading (Gordon, 1999), and connections have been made between the two (Gallese and Goldman, 1998). Much progress has been made in building artifical models of the mirror system, particularly using internal models (Demiris, 1999). Such models have been deployed onto robots, so as to investigate the practical aspects of using the simulation theory to understand and imitate the behaviour of other robots and of humans (Demiris and Johnson, 2003, 2004). Experiments with this approach have demonstrated that recognising the actions of a human requires a robot to apply a motor abstraction capability to observed actions, otherwise the recognition is impossible due to differences in human-robot morphology, and the much greater size of the human action space compared to that of the robot (Johnson and Demiris, 2004). In this paper that abstraction is achieved through modeling the motor system as a *hierarchy* of multiple coupled internal inverse and forward models.

2 Background

2.1 Inverse Models

Inverse models represent functionally specialised units for generating actions to achieve certain goals. The generic inverse model takes as input the *current state* of a system, a *goal state* that is the system's desired state, and produces as output the action required to move the system from its current state to the goal state (Narendra and Balakrishnan, 1997; Wada and Kawato, 1993). In the control literature, the inverse model is known as a *controller* and its outputs are control signals; when applied to robotics, the current state is the state of the robot and its environment, and the outputs are motor commands.

In the architecture described in this paper there are multiple inverse models, used at different levels of a hierarchical action execution and recognition system. When using multiple inverse models, each inverse model is considered *valid* for a specific goal or set of goals; that is, it can be used to achieve those goals. Thus, the *purpose* of an inverse model can be defined in general terms by the region of the goal space for which it is valid, and in specific terms at a single point in time by a particular goal taken from within that region. For example there may be an inverse model "grasp object", whose purpose is to be able to grasp a variety of possible objects. The further specification of the goal, such as specifying which object is to be grasped, may be supplied to the inverse model as a *goal parameter*.

There are situations in which an inverse model may or may not generate output. These situations are represented in the inverse model by the following states:

- If an inverse model is producing output from a current state and set of goal parameters, then it is in the state of *executing*.
- If, through comparison, the inverse model calculates that the current state is sufficiently close to the specified goal state, then no action is required. In this situation, the inverse model is *complete*.
- The inverse model may be presented with a current state that renders it unusable, as regards its purpose. The inverse model is then *ineligible*. An example would be a "Place object on table" inverse model, when there is no object.
- Although the current state may make the inverse model eligible for use, there may be a specified goal parameter for which the inverse model cannot produce any action that will result in it becoming complete. In this case, the inverse model is *not applicable*. An example would be the "Place object on table" inverse model when the object placement location has been obstructed.

The inverse model states defined above are considered binary states.

The inverse models described in this paper are not equipped with explicit initial knowledge as to the region of the goal space for which they are applicable. Instead, the inverse models determine whether or not they are capable of achieving a specific goal through an ongoing, active, simulation process, which performs action planning and results in action generation. This *simulation planning* requires the use of a *forward model*.

2.2 Forward Models

The generic forward model takes as input the current state of the system and the control signals acting on the system, and offers as output a prediction as to the *next state* of the system (Jordan and Rumelhart, 1992). In this architecture, multiple forward models are *coupled* to inverse models to create a simulation

process. This approach is similar to that used in other internal model-based systems (Wolpert and Kawato, 1998; Wolpert et al., 2003). When coupled to an inverse model, a forward model receives the action output from the inverse model through an *efference copy*. The forward model then generates a prediction of the state that would result, if the action was to be performed. This prediction can then be used for action planning and action recognition, as described in section 3 below.

2.3 Abstraction in Recognition

The architecture described here achieves action recognition by matching internally generated actions to observed external actions. In doing so, it is solving the correspondence problem (Nehaniv and Dautenhahn, 2002; Alissandrakis et al., 2002). When using robots to recognise and imitate actions performed by a human, solving the correspondence problem is made more complicated by the difference in morphology. This difference can lead to considerable disparities between the actions the robot would use to accomplish a task, and the actions the human uses to accomplish the same task in a demonstration. If the difference in morphology is small, i.e. if the robot is humanoid but with fewer degrees of freedom, then the robot can be equipped with a human motion model for action generation, which will bring the robot's actions closer in nature to that of the human demonstrator (Simmons and Demiris, 2004).

However, if the robot's morphology is so dissimilar to that of a human that it cannot produce humanlike actions, then this is a direct problem for using simulation theory for action recognition in robots. To address this issue, the motor system is developed as a hierarchical architecture, in which actions are prepared before execution using inherently more abstract simulation processes at higher levels of the hierarchy, a strategy similar to that used in (Haruno et al., 2003). Motor abstraction for successful recognition of observed human actions is then accomplished by using the higher levels of the hierarchy in a simulation theory approach.

3 The Hierarchical Architecture

3.1 Overview

The hierarchy is constructed using multiple coupled inverse and forward models. Figure 1 gives an overview of a hierarchy of K levels. The lowest level of the hierarchy contains a set of *primitive inverse*

models I_p , which generate motor commands M_t at each timestep to directly activate motor units (Bentivegna and Atkeson, 2002). The forward model in this level is a forward kinematics model of the robot, and thus offers predictions as to the *trajectory* that results from executing a motor command.

Higher-level inverse models generate actions that are sent down to the lower levels of the hierarchy for further interpretation and elaboration. Actions at higher levels are thus a more abstract representation of the eventual motor behaviour of the robot. The higher-level forward models offer predictions as to the *outcomes* and internal states of the inverse models in the lower levels that would result from the action, when it is interpreted in the level below. For example, an inverse model "grasp object" will have an outcome state "holding object = true", and the "gripper close" inverse model will then become *ineligible* for use. Thus the coupling of the high level forward and inverse models provides a simulation capability that is abstract over spatial and temporal trajectory, and which can be used for abstraction in action planning and recognition.

3.2 Action Representation

At the lowest level of the hierarchy, the primitive inverse models generate actions that are *motor commands*, meaning that they directly stimulate their intended motor units in order to realise the given action. At higher levels, inverse models generate actions that are represented by *action graphs* and *goal parameter vectors*. These actions require further elaboration at lower levels to enable final execution.

Action graphs are constructed as directed acyclic graphs, in which the nodes are inverse models and the edges specify the sequence of inverse model execution. These inverse models may produce actions that are themselves constructed as action graphs and goal parameter vectors, which are then passed on to the lower level of the hierarchy.

The recursive formulation of the action graph for action representation allows for a *multi-level* hierarchy of inverse models in action generation. An action is performed by traversing the action graph. The inverse models encountered are executed in the lower levels with the goal parameters supplied by the goal parameter vector until they are complete, and then the traversal continues.

An action graph is represented throughout the architecture by its *adjacency matrix*, denoted ψ (Jain and Krishna, 2003). To construct ψ , the N inverse models in the lower level of the hierarchy are enumerated $1, \ldots, n$, so as to index the rows and columns of ψ during its construction. To demark the beginning and end of an action, and to facilitate computation and processing, the marker nodes start and end are introduced. ψ then becomes an $(N+2) \times (N+2)$ matrix. The adjacency matrix is constructed such that if there is a directed edge from node *i* to node $j (i \rightarrow j)$ then the matrix element in the i^{th} column and j^{th} row of $\psi(\psi_{ij})$ equals one $(\psi_{ij} = 1)$, otherwise it is set to zero $(\psi_{ij} = 0)$. Thus, when parsing the matrix, an entry of "1" indicates that there is an edge from the node specifying the column to the node specifying the row, and an entry of "0" indicates no connection. When executing an action using ψ , the matrix is interpreted in a breadth-first manner, so that all the inverse model nodes leading to a single node must be completed before moving on to executing that subsequent node. This allows an action to be comprised of many parallel-executing components. An example of an action graph is given in Figure 2(A), and an example of ψ is given in Figure 2(C).

The goal parameter vector, denoted λ , has an entry for each of the N inverse models enumerated as for the action graph. If a particular inverse model requires no goal parameters, then its respective entry in λ remains zero.

3.3 Efferent Signals

When a higher-level inverse model generates an action, that action is sent in the form of an adjacency matrix and goal parameter vector as an *efferent* signal, to the level beneath in the hierarchy. The subsequent evaluation of the ensemble $\{\psi, \lambda, \mathbf{I}\}$ of the adjacency matrix ψ , the goal parameter vector λ , and the set \mathbf{I} of inverse models, results in the generation of more specific actions, and those actions are propagated all the way down the hierarchy, until the action becomes a motor command M_t and is eventually realised in the motor units.

3.4 Afferent Signals

Proprioceptive information for joint configurations, and exteroceptive information regarding objects in the environment, are continually provided by sensor units. This information is arranged into the current state vector S_t and is sent up through the hierarchy as an *afferent* signal. Every level of the hierarchy receives this signal. For higher levels, the state information is supplemented by the status of the inverse models in the previous level, i.e. whether those inverse models are *complete*, *eligible*, *applicable*, or *executing*. Along with the efferent signals from the



Figure 1: A K-level hierarchy of coupled inverse and forward models. The same architecture is used for both performing an action and recognising actions. The lowest level contains a set of primitive inverse models denoted I_p . F_p is the forward model for these primitives. $S_{t,r}$ is the state of the robot at time t, and \hat{S}_t is the predicted state at time t. The D in each level indicates a time delay, which is used to bring the prediction temporally in-line with the current state for meaningful comparison.

level above, this afferent flow of status information provides for reciprocal connections between the levels of the motor system.

3.5 Simulation Processes

The dashed lines in Figure 1 mark the feedback generated from the closing of two *simulation loops*. These simulation loops may be used for action execution, planning, recognition, and learning, depending on the requirements of the robot.

3.5.1 Inner Loop

The *inner* simulation loop is used for planning and modulating an on-going action during action generation. The inverse model generates multiple *action hypotheses* that it postulates will achieve the specified goal parameters. The action hypotheses are tested on the forward model, resulting in *predicted states* that are sent back to the inverse model. The inverse model can then use these predicted states in *substitution* for the current state, creating a simulation process that allows it to plan actions into the future, by searching the possible action space. Through comparison with the goal parameters, the inverse model converges to an *action solution*. There may be many potential action solutions that accomplish a given goal. The most appropriate solution at any given time is selected by a winner-takes-all mechanism, on the basis of the smallest action-graph depth, and sent to the level below. All the levels perform the same simulation process continually, and in parallel. The result is a distributed on-line hierachical control model that directly and indirectly modulates an action as it unfolds.

If, through the inner-loop simulation process described above, an inverse model determines that it is *unable* to achieve its goal, then this "not applicable" state is signalled as part of the overall state of the inverse models in that level (other states are *complete*, *eligible*, and *executing*). The afferented robot state information is supplemented by this inverse model state information as it reaches each higher level. The combined state information is then used in the *outer simulation loop*.



Figure 2: An example of an (A) action graph. 'S' and 'E' are start and end nodes, respectively. (B) enumerated inverse model set I and (C) adjacency matrix ψ , for an action "Grasp Object" generated by a high-level inverse model. In this example, nodes 1, 2, and 3 execute in parallel, and each must become complete before 4 can be executed.

3.5.2 Outer Loop

The outer loop is a prediction-comparison process. The forward model produces a prediction \hat{S}_t as to the result of the supplied action solution, and this is *buffered* by the delay component D, before comparison with the *actual resulting state*, S_t . The resulting *prediction error* P_e may be used both for action generation and learning of forward and inverse models when the supplied current state is the agent's own (Haruno et al., 2003), or action recognition and imitation learning, when the current state is that of an observed actor (Demiris and Hayes, 2002; Demiris and Johnson, 2003). In this architecture, the prediction error is calculated as being the sum over the nstate elements, of the absolute difference between the predicted state and the actual state:

$$P_{e} = \sum_{i=1}^{n} \left| S_{t,i} - \hat{S}_{t,i} \right|$$
(1)

3.6 Recognition and Imitation

The same arrangement of structures, as shown in Figure 1, is used for action recognition as well as execution. In recognition, the state input to the architecture is not taken from the robot, but is derived from visual observation of the demonstrator. All the inverse models in every level of the hierarchy that are "eligible" for execution, and not "complete", are then executed in parallel. The inverse models in a particular level compete with the other inverse models in that level for *confidence*, which is awarded at each time step to inverse models that match well with the perceived action. A winner-takes-all selects the inverse model with the highest confidence at any point in time as being the recognised action. The robot's motor hardware is taken off-line to prevent physical "mirroring" of the perceived action, by inhibiting the motor commands generated by the primitive inverse models in the lowest level of the hierarchy. When recognition is complete, imitation may proceed by executing the observed action.

3.7 Confidence Calculation

3.7.1 Lowest Level

The inverse models compete for confidence. At each timestep, the inverse model with the lowest prediction error P_e is *rewarded*, and the rest of the inverse models are *punished*. The inverse model with the lowest prediction error has its confidence C_t rewarded as follows:

$$C_{t} = \begin{cases} C_{t-1} + \frac{1}{\epsilon} & \text{if } P_{e} < \epsilon \\ C_{t-1} + \frac{1}{P_{e}} & \text{otherwise} \end{cases}$$
(2)

The other inverse models have their confidences punished, according to:

$$C_t = \frac{C_{t-1}}{2} \tag{3}$$

Initial confidences are zero for all inverse models. In the following experiments, ϵ was chosen to be 0.04.

3.7.2 Higher Levels

The forward models predict the outcomes and internal states of the lower-level inverse models that are the components of the action input. Thus, the higherlevel inverse models are rewarded when the prediction error P_e is less than ϵ , and their confidences are reset when they become complete:

$$C_t = \begin{cases} C_{t-1} + \alpha & \text{if } P_e < \epsilon \\ 0 & \text{if inverse model is complete} \end{cases}$$
(4)



Figure 3: Confidence levels of primitive inverse models in the lowest level of the hierarchy during a demonstration of picking up an object and placing it back on the table. The sequence of movements is: move to object \rightarrow move object away from table \rightarrow move object to table \rightarrow move away from object. The confidence values have been normalised at each time step.

As the prediction error is less than ϵ only at specific times, the confidence is never punished and the inverse models do not compete for confidence. Initial confidences are zero for all inverse models. In the experiments that follow, α was chosen to be 10.

4 Implementation

To demonstrate the architecture, it was implemented in a two-level hierarchy on a robot in an experimental scenario involving the recognition and imitation of object manipulation actions performed by a human demonstrator.

The lower level of the hierarchy was populated with six primitive inverse models, "gripper open", "gripper close", "move to object", "move away from object", "move object to table", and "move object away from table". The higher level was equipped with the inverse models "grasp object" and "place object", both of which accomplished their goals by combining the low-level primitives into action graphs. To simplify the implementation, only one object and one table were used, restricting the goal parameter space.

4.1 Robot Platform

The Peoplebot is equipped with a Canon VCC4 pantilt-zoom (PTZ) camera, two degrees of freedom gripper, and sonar and infra-red sensors. In these experiments, the camera was used as the main tracking and range-finding sensor. The sonar and the infra-red sensors were not used. All processing was done in real-time, with one full iteration of the architecture's mainloop executing in 0.5 seconds. The software was written in C++ for an AMD Athlon 64, which controlled the robot remotely over a wireless ethernet link.

4.2 Visual Systems

The visual tracking of the object and the hand was accomplished using the CAMShift algorithm (Bradski, 1998), working on a hue and saturation histogram back-projection of camera images taken at a pixel resolution of 640×480 and at 2 frames per second. The low frame rate was deliberately chosen to reduce noise in the visual signal. The ARToolkit (Billinghurst et al., 2001) was used to determine the robot's position relative to the table, as stereo vision was not available on the robot. Depth information was thus obtained by affixing an 8 cm \times 8 cm marker to the table's midpoint.

5 Experiments

The object manipulation actions chosen for the experiments were the common tasks of picking an object up from a table, and placing an object onto a table. These behaviours are well suited to the robot used, an ActivMedia Peoplebot, with its mobile platform and gripper assembly.

The robot was positioned facing a table, upon which was placed an object that was readily manipulable by both the robot and the human demonstrator. In these experiments, the object used was a tub. The initial robot-table distance was 1 m, sufficient for the robot's camera to view the entire scene, including the table, object, and the hand of the demonstrator as she moved to place or pick up the object. The demonstrator was unfamiliar with the operational details of



Figure 4: Confidence levels of inverse models in the higher level of the hierarchy during a demonstrations of grasping an object and placing an object.

the architecture, and was instructed when to start the demonstration. If the robot recognised the demonstrated action then it performed the action for itself, completing the cycle of imitation.

6 Results

Figures 3 and 4 show typical results from the experimental trials. Figure 3 shows the confidence levels of four primitive inverse models in the lowest level of the hierarchy during a demonstration of picking up and placing the object (the primitive inverse models shown in this graph are "Move to object", "Move object away from table", "Move object to table", and "Move away from object"). The architecture achieves successful recognition, ascribing high confidence levels to the primitive inverse models that generate trajectories that match with the observed actions. The progression of the confidence values shows the competition between inverse models during transitional stages of the action, where one inverse model builds up confidence at the expense of the others (iterations 12-14, 24-26, and 30-31). The duration of a recognised action can be seen as the length of time that the confidence level for a particular primitive remains at 1.

Figure 4 shows the confidence levels of the two inverse models in the higher level of the hierarchy. The peaks in the confidence clearly demark the "grasp object" action and the subsequent "place object" action. The higher-level inverse models do not match on action trajectory, but on subgoals during an action, resulting in a more abstract recognition that clearly distinguishes different observed behaviours.

7 Discussion

For large numbers of inverse models in any given level of the hierarchy, an adjacency matrix becomes a memory-inefficient means of action representation. The computational cost of adding inverse and forward models is therefore less than the overall memory cost. However, due to the directed nature of the action graphs, the matrix ψ is sparse, and can be efficiently managed through the use of look-up tables. It is expected that on modern computers the system could handle up to and beyond a hundred inverse and forward models.

Although the abstraction architecture is capable of recognising actions performed in different ways, the visual system is sensitive to the speed at which the actions are performed. This results in situations where recognition may not be successful. If the demonstrator moves too slowly, then noise in the visual system overcomes the movement signal and lower-level recognition fails, although higher-level recognition may succeed. Recognition at all levels fail if the human performs the movement too fast for the architecture to extract a reasonable signal.

Acknowledgments

This research was funded by the UK's Engineering and Physical Sciences Research Council (EPSRC) project GR/S11305/01, and the Royal Society. The authors would like to thank the BioART team, and also Mercedes Lahnstein and Tania Katzouraki for their invaluable support and assistance.

References

- A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Imitating with ALICE: Learning to imitate corresponding actions across dissimilar embodiments. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, 32(4):482– 496, 2002.
- D. Bentivegna and C. Atkeson. A framework for learning from observation using primitives. Presented at the Symposium of Robocup 2002, Fukuoka, Japan, June 2002.
- M. Billinghurst, H. Kato, and I. Poupyrev. The magicbook: A transitional AR interface. *Computers* and Graphics, pages 745–753, November 2001.
- G. R. Bradski. Computer vision face tracking for use in a perceptual user interface. *Intel Technology Journal*, Q2:1–15, 1998.
- J. Demiris. *Movement imitation mechanisms in robots and humans*. PhD thesis, University of Edinburgh, Scotland, UK, 1999.
- Y. Demiris and G. Hayes. Imitation as a dual route process featuring predictive and learning components: a biologically-plausible computational model. In K. Dautenhahn and C. Nehaniv, editors, *Imitation in Animals and Artifacts*. MIT Press, 2002.
- Y. Demiris and M. Johnson. Distributed, predictive perception of actions: A biologically inspired robotics architecture for imitation and learning. *Connection Science*, pages 231–243, 2003.
- Y. Demiris and M. Johnson. Simulation theory for understanding others: A robotics perspective. In *Imitation and Social Learning in Robots, Humans and Animals: Behavioural Social and Communicative Dimensions*. Cambridge University Press, 2004. to appear.
- L. Fadiga and L. Craighero. New insights on sensorimotor integration: From hand action to speech perception. *Brain and Cognition*, 53:514–524, 2003.
- V. Gallese and A. Goldman. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2(12):493–501, December 1998.
- R. M. Gordon. Simulation vs theory-theory. In R. A. Wilson and F. Keil, editors, *The MIT Encyclopædia of the Cognitive Sciences*, pages 765–766. MIT Press, 1999.

- M. Haruno, D. M. Wolpert, and M. Kawato. Hierarchical MOSAIC for movement generation. *International Congress Series*, 1250:575–590, 2003.
- S. Jain and S. Krishna. Directed graphs and their adjacency matrices. In S. Bornholdt and H. G. Schuster, editors, *Handbook of Graphs and Networks: From the Genome to the Internet*, chapter 16.2.1, pages 357–363. Wiley-VCH, 2003.
- M. Johnson and Y. Demiris. Abstraction in recognition to solve the correspondence problem for robot imitation. In *Proceedings of TAROS*, pages 63–70, 2004.
- M. Jordan and D. Rumelhart. Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16:307–354, 1992.
- A. N. Meltzoff and J. Decety. What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience. *Phil. Trans. of the Royal Society of London B*, 358:491–500, 2003.
- K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Transactions on Automatic Control*, 42(2):171–187, February 1997.
- C. L. Nehaniv and K. Dautenhahn. *Imitation in Animals and Artifacts*, chapter 2: "The Correspondence Problem", pages 41–61. MIT Press, 2002.
- G. Rizzolatti, L. Fogassi, and V. Gallese. Motor and cognitive functions of the ventral premotor cortex. *Current Opinion in Neurobiology*, 12:149– 154, 2002.
- G. Simmons and Y. Demiris. Imitation of human demonstration using a biologically inspired modular optimal control scheme, 2004. To appear at IEEE-RAS/RSJ International Conference on Humanoid Robots (Humanoids 2004).
- Y. Wada and M. Kawato. A neural network model for arm trajectory formation using forward and inverse dynamics models. *Neural Networks*, 6:919–932, 1993.
- D. M. Wolpert, K. Doya, and M. Kawato. A unifying computational framework for motor control and social interaction. *Phil. Trans. of the Royal Society of London B*, 358:593–602, 2003.
- D. M. Wolpert and M. Kawato. Multiple paired forward and inverse models for motor control. *Neural Networks*, 11:1317–1329, 1998.

A Developmental Roadmap for Task Learning by Imitation in Humanoid Robots

Baltazar's Story

Manuel Lopes^{†,‡}

Alexandre Bernardino[†]

José Santos-Victor[†]

[†]Instituto de Sistemas e Robótica, Instituto Superior Técnico, Lisboa, Portugal [‡]Escola Superior de Tecnologia, Instituto Politécnico de Setúbal, Setúbal, Portugal http://vislab.isr.ist.utl.pt {macl,jasv}@isr.ist.utl.pt

Abstract

We present a strategy whereby a robot follows a developmental pathway to (i) explore its own visuomotor capabilities, (ii) understand its surrounding environment, (iii) become aware of people acting in the environment and finally imitate observed actions. We describe some results of the different developmental stages, involving perceptual and motor skills, implemented in our humanoid robot, Baltazar. In addition to the overall system, another important contribution is the use of a two-phase, uncalibrated algorithm for object grasping. The last phase is driven under closed-loop vision based control, where the Jacobian is learned online.

1 Introduction

"Friendly" and social interaction between robots and humans is a grand challenge for robotics. Due to the diversity of actions/tasks to be performed and the range of possible interactions with objects and humans, it would be impractical (if not impossible) to explicitly pre-program a robot with such capabilities. Instead, such systems must be able to learn by themselves what tasks to execute and how they should be performed, which requires sophisticated motor, perceptual and cognitive skills.

To address these challenges, we (as well as other researchers) adopt two fundamental metaphors: (i) learning by imitation as a powerful means to teach a complex humanoid-like (social) robot and (ii) a developmental approach that can balance the complexity of the system at the various levels of functional performance.

Learning by imitation is likely to become the primary form of teaching such social, cognitive robots. Let us consider a system able to learn how to solve some tasks by imitation, e.g. by observing a human manipulating a set of objects. This problem of skill transfer has three major difficulties: (i) how to gather task-relevant information? (ii) how to convert the data that are valid for a human for a robot? and (iii) how to infer the important parts of the demonstration (e.g. "understand" the task).

Several approaches have been adopted to gather the



Figure 1: Baltazar. A 14 degrees of freedom humanoid torso.

information for imitation. Schaal et al. (2003) use an exoskeleton to capture kinematic data. Oztop and Arbib (2002) rely on some marks to get visual features for hand detection and grasping, in the context of imitation and modeling of the Mirror neurons. Lopes and Santos-Victor (2003) exploit task-contextual bias to modulate the information extraction process. Imitation and skill transfer between systems with different bodies (kinematics, dynamics and skills) was addressed by Nehaniv and Dautenhahn (1998) using an algebraic formulation (bodies with different skills were considered). For the case of a humanoid robot, Nakaoka et al. (2003) introduce adaptation of the trajectories to be able to guarantee the correct balance during task execution.

Kuniyoshi et al. (1994) proposed one of the first

works in imitation, a system able to learn how to imitate an assembly task by extracting a hierarchical description of the task. Billard et al. (2004) address the problem of inferring the important parts of the task by casting it into an optimization framework. Zöllner and Dillmann (2003) present a system where two hand tasks are imitated, using information about the functionality of each object and handling temporal task restrictions, in a symbolically manner.

Even if imitation can allow a robot to learn an extremely large variety of tasks, it is clear that it requires the robot to have several sophisticated motor, perceptual and cognitive capabilities. Hence, building such complex skills can become an overwhelming task in itself. For learning one particular skill, many other systems may need to be present and their inter-connections properly established.

The developmental perspective, as proposed by e.g. Weng (1998), is a new paradigm aiming at overcoming this complexity problem, of learning and properly integrating many perceptual, motor or cognitive skills.

The robot should "start" with a minimal subset of core capabilities (as newborns do) to bootstrap learning mechanisms that, through self-experimentation and interaction with the environment and other humans, would progressively lead to the acquisition of new skills, adapted to particular contexts, and having the system integrating all the learning methods internally. Metta (1999) used a developmental approach for a robot that successively acquired vergence, saccade and vestibular control, as well as head-arm coordination.

Amongst the capabilities required to interacting with objects, understand their spatial configuration and learning by imitation, perception is perhaps the most fundamental. They allow gathering (task or contextually relevant) information and training samples for all other forms of learning. Then, some motor capabilities need to be in place before the robot can start interacting with the world and providing "calibration" information for other modules (e.g. relating depth information from vergence with arm length).

The development of imitation capabilities requires an appropriate definition of the sequence of learning steps to reach that goal, as well as adequate performance evaluation methods to decide when to switch to higher developmental levels. In other words, it is important to define the overall hierarchy of developmental stages and the skills that must be acquired at each level. Table 1 shows the structure we adopt for the main developmental stages the robot (or a human infant Arsénio (2004); Natale (2004)) will go through: (i) Learning about the self; (ii) Learning about objects and the world and (iii) Learning about others and imitation.

For each stage in this "developmental pathway", we show the set of skills to be acquired, and the time line explaining the restrictions governing the system. We do not distinguish between innate versus learned behavior in biological systems ("the nature versus nurture" question). Instead, we just request all the modules to be present before the system can develop to the next level.

Table 1. Developmen	nai paniway ioi me reiceptuai		
and Motor capabilities (in <i>italic</i> the modules that are			
learnt by the robot)			
Time line	Percentual/Motor Canabilities		

Table 1. Developmental nothway for the Dercentual

Time line	Perceptual/Motor Capabilities
	eye vergence
	random movements
↓ self-awareness	Arm-head coordination
	near-space mapping
	near-space mapping
↓ world-awareness	visually initiated reaching
	visual control of grasp
\downarrow imitation	detection of other's actions
	imitation of tasks

In the first developmental level, the robot acquires very simple and yet crucial capabilities: vergence control and object foveation/tracking. Then, by executing random arm movements, in a self exploratory mode, it begins to coordinate head and arm configurations, by creating a arm-head map. This map is accurate enough to allow for reaching and grasping objects in easy positions.

In the second developmental stage, the robot builds a map of the surrounding area (object positions and identification). Driven by attentional cues, the robot engages in more challenging grasping tasks, for which the previously learned arm-head map is not sufficiently accurate. For that reason we propose a novel method for visually controlled grasping, which improves over time and ensures the necessary robustness.

Previous approaches for object grasping were either completely visual controlled Kragic et al. (2002), with problems in guaranteeing the presence of the hand in the visual field, or completely open-loop Natale (2004) with no capability of error correction. Instead, we combine the two modalities, with an openloop phase moving the effector to the field of view followed by a closed-loop method with the precision necessary to put the effector in contact with the object. At the final developmental stage, the presence of a demonstrator will elicit a task imitation behavior, that will decompose the actions and them replicate with a given metric. For this purpose, the system must be able to decompose the observed action into the relevant key elementary actions that must be executed for performing a task.

To conclude, our main contributions are two-fold. On one hand we present a developmental strategy for humanoid robots, according to widely accepted stages in developmental psychology. On the other hand, we propose a visually guided grasping process, where learning is driven by the motivation to precisely grasp objects, that continuously adapts over time (open ended learning).

All experiments in the paper were implemented in our humanoid robot, Baltazar, equipped with a 4 dofbinocular head, a 6 dof arm and an 11 dof (underactuated by 4 motors) hand. The robot is shown in Figure 1 and described in detail in Lopes et al. (2004).

In Section 2 we present the development of self awareness. Section 3 deals with the understanding of the world and the interaction with objects. Imitation learning of tasks is presented in Section 4. Conclusions and future work are drawn in Section 5.

2 Self-Awareness

Humans take a long time before becoming selfsufficient. Knowing how to walk, how to recognize objects, understanding how to solve a task, interacting with objects, are very difficult tasks, and only after several years all mechanism necessary are available and reliable. Becoming aware of its own body and then start to coordinate it is the first step to survival. Infants have several mechanism guiding its development.

For the case of the head-eye system, voluntary control appears very early. Some reflexive movements are evident from birth (head-righting reflex Payne and Isaacs (1999)) but voluntary control becomes apparent only at the end of the first month. A five-month old child already shows a good control. This control of the head will enable the tuning of the vision system to start looking at (and understanding) the environment. In van der Meer et al. (1995) there is a discussion about the significance of neonate's, apparently, random arm movements.

Several reflexes allow newborns to look to their hand. The "Asymmetric Tonic Neck Reflex" can be elicited when the baby is prone or supine. When the head is turned to one side or the other, the limbs on the face side extend while the limbs on the opposite side flex. This reflex is believed to facilitate the development of an awareness of both sides of the body as well as help develop eye-hand coordination.

The interaction between eyes and hand is very important. This interaction will allow the newborn to tune its eyes, distinguish depth and recognize touchable objects. For a baby exploring the hand, how it moves and how it looks will be the most interesting thing in the first few months. Learning to make it do what it wants to do, will be a very complex learning task. In the end the reward will be tremendous: being able to predict hand movements and to touch objects.

In this section we present the capabilities allowing the system to be aware of its own body and to learn how to coordinate it.

2.1 Near-Space perception

The near-space contains the touchable objects and our own body. Being able to understand what happens there is fundamental. Bernardino and Santos-Victor (2002) suggest a method where the disparity between images is used, together with some neuronal-based filters, to segment objects at different depths. The head can be moved to look toward the hand using disparity as a feedback signal to control it. Figure 3 shows a result of verging on an object. This same mechanism will later be used to map object positions.

2.2 Arm-Head Coordination

Many tasks need a very fine Arm-head coordination. Object manipulation is only possible with precise visual control of the hand. In order to coordinate Arm positions with Head position, we are going to create an *Arm-Head Map*. This map is bidirectional. If the head position is fixed moving the arm to the mapped position puts the hand in the fovea of the two eyes. If the arm is fixed, we can visually locate it by moving the head to the mapped position.

Several approaches can be used to learn this map. In Lopes and Santos-Victor (2003), a neural network was used to map from arm feature points to joint angles. In D'Souza et al. (2001), a very powerful method is used to learn inverse kinematics of a humanoid. Vijayakumar and Schaal. (2000) created a method, *Locally Weighted Projection Regression*, that will be used for learning the map. This method is linear with the number of samples and every new sample can be added easily. As it is not capable of extrapolating, the working space must be well covered in the training set.

The data set is gathered by self-observation. The arm is moved around in the space, while the hand is

tracked and foveated. Figure 2 shows the hand being moved to the front of the eyes by using the *Head-Arm map*. The quality of this map is good enough to guarantee that the hand is always in the image but not in the fovea. In our experiments, the average error is about 5cm, corresponding to 15% of the image.



Figure 2: Head-Hand Coordination

This map will enable the system to reach and, in special cases, grasp objects. This will be very motivating and in the next level object grasping will develop further.

2.3 Attention Mechanism

When looking around us some objects attract more our attention than others. This is related to the urgency of each task and to reduce the amount of information to process. Context may influence the attention drawn by some objects, (e.g. food when hungry). In our approach, the attentional process depends on the developmental stage. In the beginning, the hand is the main focus of attention, facilitating the learning of the Head-Arm Map. Also salient objects in the scene attract the system's attention in a bottomup process. Later on, in the second stage, Baltazar will search around him, pay attention to all objects, one at a time and create a map of the nearby area. In the final stage, attention will be driven toward the person doing the demonstration and the manipulated objects. In these later stages, attention becomes gradually more driven in a top-down, context and task dependent manner.

3 World Awareness

As the robot gains control over its own perceptual and motor capabilities, it gets more and more interested in exploring its surrounding world. This exploratory motivation will call for the development of more advanced manipulative capabilities as opposed to the rudimentary skills available during phase one.

For object grasping, it is necessary to have several motor programs: the arm must be able to approach the object (reaching) before finally grasping it, the hand must be able to have a stable grasp and pre-shaping can be necessary for faster movements or moving objects. However at this stage all the robot can do is to fixate at salient objects and approach them in a primitive form of grasping. The development path will require the following new skills:

- 1. detect object's positions in the nearby space and store this information in some sort of representation (near space map).
- 2. learn how to reach objects in a controlled manner, using visual feedback, and grasp them.

This section describes algorithms that solve all these steps, allowing a robot to move on to the next developmental level, where it gains awareness of others (humans or robots) and the actions they perform. In addition to the reaching step based on the *Head-Arm* Map presented in Section 2, we propose a new algorithm to grasp objects based on visual servoing techniques estimated online.



Figure 3: Verging on an object. Left (\triangle) and right (+) eye.

3.1 Near-Space (Objects) Mapping

There is neurological evidence of spatial aware neurons that are active when movement or objects are present near the skin Rizzolatti et al. (1977). It is also known in developmental psychology that infants became aware of the near and far space very early. It is very useful to know where an object is and whether it can be grasped or not. After all the time spent interacting with its own hand, the system can already

distinguish objects at different depths and search for the desired one.

By this exploratory behaviour, we create a map of the localization of objects around us - the peripersonal map - through various steps:

- 1. Find an object in the visual space
- 2. Foveate on this object
- 3. Memorize the object position in head (proprioceptive) coordinates (Θ_{Head}).

Through exploration, the robot thus creates a mental image of the surrounding space. The position of the objects are memorized in terms of head (proprioceptive) coordinates. In Figure 4 Baltazar is searching for "fruits" around him where different objects are assumed to have different colours.



Figure 4: Mapping object positions in head coordinates.

3.2 Object Grasping - a two step approach

Infants start reaching objects without any visual feedback. The movement is only initiated with vision but not guided throughout the entire action. In case of failure, the movement restarts from the beginning.

At the first stage of development, the estimated *Arm-Head map* allows the system to (crudely) move the hand towards an object. Hence, if a simple trajectory is followed, the hand may well succeed in touching the object. The problem with this (open-loop) approach is the absence of a mechanism for error correction. This is the reason why babies in this phase restart the grasp quite often, instead of correcting it Payne and Isaacs (1999).

The second stage of object reaching relies on visual feedback, coping with the problem of error correction. The *Head-Arm map* is used to move the hand to the objects vicinity. Then, accurate positioning is achieved by visual guidance in closed loop. With this phase, it is possible to grasp objects in a reflex type manner, the hand closing after touch.

The method presented in D'Souza et al. (2001) could be used here. Their approach consists in mapping motor positions and velocities to image velocities, using a very strong statistical learning approach, yielding good results. The disadvantages arise from the lack of extrapolation capabilities and by not having an explicit Jacobian estimation, thus needing more time to gather the information, and preventing the use of well studied visual servoing control algorithms.

We adopted a visual servoing perspective, described by e.g. Hutchinson et al. (1996). However, although it is possible to solve the problem with an algebraic formulation, we adopted a model-less way, as it allows the system to learn and develop from its own experience. A particularly useful method for online estimation of visual motor relations is presented by Jaegersand (1996). The image *Jacobian* (*J*) relating image changes (Δy) caused by motor movements ($\Delta \theta$), can be interactively estimated by:

$$\hat{J}(t+1) = \hat{J}(t) + \alpha \frac{\left(\Delta \mathbf{y} - \hat{J}(t)\Delta \theta\right)\Delta \theta^{T}}{\Delta \theta^{T}\Delta \theta}$$

where α denotes the Jacobian update rate. To move the system to the desired image position y^* , we apply the following control law:

$$\Delta \theta = g \left(J^+ \left(\mathbf{y}^* - \mathbf{y} \right) \right)$$

where J^+ represents the pseudo-inverse of J and the function g(.) can be chosen to have a exponential, linear or any other type of convergence.

In order to deal with a larger workspace and to incorporate some open-loop movements, we had to improve the existing algorithm. More details can be found in the Appendix A. Figure 5 shows the resulting behavior of the system while grasping objects. The hand is closed after sensing the contact with the object. The capability of pre-shaping the hand will only develop at a later stage. For small grasping velocities, this type of movements can be sufficient, but bigger velocities will require learning some form of pre-shaping and predicting the time of contact with the object.

At this point in development, the system can not only control its own body and perceptual abilities but also perform relatively complex manipulation tasks, memorize objects spatial configurations, search for objects, etc. It is then ready to start looking at humans or other robots and the tasks they perform.



Figure 5: Several frames in the sequence from the initial position resulting from the *Head-Arm Map*, then the visual guided part and finally the object grasping.

4 Imitation

Figure 6 shows an example of a task being executed. It consists of picking up some objects and moving them around. To imitate this task, the robot will first need to understand the spatial relations of objects around the demonstrator (understand the far space). Then, understanding the near space becomes fundamental to establish correspondence between the demonstrator perspective its own (self) viewpoint (i.e. the blue object is on the left of the demonstrator, but it is in front of me). After the observation of the demonstration movements, the important task moments must be extracted and segmented. Finally the task is repeated by the robot, using the task description and all the modules previously learned. The following sections will provide details on the different modules developed at this stage.



Figure 6: Several frames of the task demonstration.

4.1 Far-Space Interpretation

Understanding events and object's localizations at far distances (i.e. more than the arm can reach) is different from mapping our surrounding space. The frame of reference will no longer be our own body, instead we describe object's positions relative to another person, this is specially useful for imitation learning. Object's position will be codified in terms of allocoordinates. Some simplifying assumptions can be made about depth in order to reduce the complexity of scene reconstruction.

4.2 Task Segmentation

The actions and movements of the demonstrator must be segmented and codified in a way useful for imitation. We developed a method consisting in a multiple object tracking and a task point detector. When doing manipulation our hand will occlude objects very frequently. Grasping and releasing can be very difficult to detect. Being the hand the only actuator enables the usage of information to deal with occlusions. Every object can have three movement models: static, moving and being moved. When an object is moving its velocity profile can be predicted with Newtonian dynamics, when being moved is has the same velocity as the hand. The algorithm will mark every point in the trajectories of the objects that satisfy the following constraints: all object are static, the hand is not moving and the hand is not occluding any object.

The task is then codified by having objects with their physical properties (shape and color) and their spatial relations (A between B and C;A right of B or A left of B).

The complete sequence shown in Figure 6 has 234 frames, this sequence was processed online and the task points, shown in Figure 7, were automatically extracted. We can see that the system succeeds in detecting what frames are important to describe the task.

4.3 Imitation

As mentioned in Gergely et al. (2002), imitation goals are not always very clear. In our case the imitation task will proceed in order to have the same spatial relations. In case the demonstrator has made a movement and there is no difference in the ordering of objects (Figure 7), the robot will mimic the absolute spatial positions. We can see that all the modules developed until this point were essential to be able to replicate the task at hand.

5 Conclusions/Future Work

We have presented a developmental route for a humanoid robot¹ to acquire increasingly more complex skills.

The robot first learns about it's own body and surrounding environment. All information is gathered by self-exploration. The quality of the Arm-head coordination achieved in this phase is sufficiently good to ensure that the hand always remains in the image and that objects can be grasped in simple cases. In a second phase, motivated to further interact with objects, the system develops a closed-loop control behavior capable of precise grasping. It also creates a map of the interesting objects in the surrounding space. In the final developmental phase, people acting in the environment are the major source of information. The observed tasks are segmented in special points in order to finally imitate the task.

The developmental pathway allows the robot to acquire new skills on top of the existing (learned) capabilities. We described results of the various developmental stages of the system: the vergence and object tracking system, the learning of the Armhead map, the visually initiated object grasping system and a new solution to visually guide grasping. The method consists in two phases: an open-loop controller putting the hand close to the object, and a closed-loop vision-based controller for precisely touching the object. This method does not need calibration and can be learned on-line in a very efficient way. In the future, we will focus our efforts on the aspects of learning the interaction between people and objects.

A Visual Grasp

In this section we present a generalization of the method suggested by Jaegersand (1996), to be used to visually control the arm. The image *Jacobian* (*J*) relating image changes (Δy) caused by motor movements ($\Delta \theta$), can be iteratively estimated by:

$$\hat{J}(t+1) = \hat{J}(t) + \alpha \frac{\left(\Delta \mathbf{y} - \hat{J}(t)\Delta\theta\right)\Delta\theta^T}{\Delta\theta^T\Delta\theta}$$

where α is the Jacobian update rate. To move the system to the desired image position y^* , we can apply the following control law:

$$\Delta \theta = g \left(J^+ \left(\mathbf{y}^* - \mathbf{y} \right) \right)$$

where J^+ represents the pseudo-inverse of J and g(.) can be chosen to have a exponential, linear or any other kind of convergence.

When the working volume is very large the Jacobian can no longer be accurately estimated with only one linear model. To solve this we propose a new method. With only one linear model the update mechanism must be fast enough to have an accurate model for each region. In the case of openloop movements the system can no longer update the model and a specific model for the new region must

¹see http://vislab.isr.ist.utl.pt/baltazar for videos showing the experiments in this work



Figure 7: Segmentation of a task. Notice that from the third to the fourth image there is no difference in the ordering of the object, just their absolute distances. These relevant points where extract online from a video sequence with 234 frames.

already be present. The workspace should be partitioned in several regions, R_i , i = 1...N. At each instant the distance c is measure between the current position and all the regions, the selected Jacobian Jis the one corresponding to the nearest area R_i . We use a Mahalanobis distance with covariance D. The covariance can be updated online to reduce the number of regions and to better adjust the linear model to the non-linear system. Trying to update the regions center creates problems by overlapping regions and with region transitions.

The Jacobian update rate (α) should be larger when the model is inaccurate and then reduced to improve convergence. One measure to access the model quality (*mq*) can be:

$$mq(t) = mq(t-1) + \gamma < \Delta y, J_k \Delta \theta >$$

 γ is a decaying factor and $\langle . \rangle$ represents internal product. mq is positive when the observed movements has a direction error less than 90 degrees.

The regions centers x_i may correspond to motor features $x = \theta$, visual features x = y or a combination of them. With visual features there is the possibility of doing planning in visual space but there are different motor positions that give the same visual features and should have different linearizations.

Table 2 presents the complete algorithm for doing the visual controlled grasp.

 J^+ must be carefully implemented. As some directions are not observed, the Jacobian inversion will be very unstable. To solve this problem the pseudo

Table 2: Uncalibrated Visual Servoing Algorithm

To move the system to the desired image position y^*

1. Choose the region R_i corresponding to the actual state x:

$$c_i = (x - x_i)^T D_i (x - x_i)$$
$$R_i : \min_i c_i$$

if max $c_i < C$ create a new area l with $x_l = x$, $D_l = D$ and $J_l = J_i$. Choose $R_i = R_l$.

2. apply the control law:

$$\boldsymbol{\Delta}\boldsymbol{\theta} = K_i \frac{J_i^+ \left(\mathbf{y}^* - \mathbf{y} \right)}{\left\| J_i^+ \left(\mathbf{y}^* - \mathbf{y} \right) \right\|}$$

- 3. observe image changes Δy
- 4. make the update to the model *i* corresponding to position *x* with:

$$\hat{J}_i = \hat{J}_i + \alpha_i \frac{\left(\Delta \mathbf{y} - \hat{J}_i \Delta \theta\right) \Delta \theta^T}{\Delta \theta^T \Delta \theta}$$

5. if $|y^* - y| > E$ goto 1

inverse is implemented with a SVD method and any singular values less than 10% of the larger are treated as zero.

Chaumette (1998) show some problems present in Visual servoing methods. Our method solves the problem of the Jacobian derivation and the calibration of the robot and cameras. In general these methods are sensitive to initial positions, being prone to fall in local minima but, in our approach, the system always starts near the final position due to the *Head-Arm map*, thus making convergence easier.

We made several experiments to access the quality of the resulting algorithm. Our system measures a specific dot in the hand with two cameras giving an image position of the hand (u_l, v_l) for the left eye and (u_r, v_r) for the right eye. The features are calculated as follows:

$$\mathbf{y} = \begin{bmatrix} \frac{u_l + u_r}{2} \\ \frac{v_l + v_r}{2} \\ u_l - u_r \end{bmatrix}$$

This gives position and distance information estimation of the hand related to the head. The head was maintained fixed and four arm joints were used. The distance between the central point of each zone was 10 *degrees*. The Jacobian update rate was equal in all regions and choosen as $\alpha = 0.1$ while mq < 0and $\alpha = 0.01$ while mq > 0.

Figure 5 shows some quantitative results of the grasp sequence shown in Figure 5 using our proposed algorithm. The hand was positioned near the object using the *Head-Arm map*. The resulting error corresponds to about 8 *cm*. The associated image error is corrected in the final phase (visually controlled) with a linear convergence rate.

Acknowledgments

Work partially supported by: EU Proj. (IST-004370) RobotCub http://www.robotcub.net and by the FCT POSI-Program in the frame of QCA III.

References

- Artur Arsénio. Cognitive-Developmental Learning for a Humanoid Robot: A Caregiver's Gift. PhD thesis, Massachusetts Institute of Technology, Boston, USA, Sept 2004.
- Alexandre Bernardino and José Santos-Victor. A binocular stereo algorithm for log-polar foveated systems. In *Biological Motivated Computer Vision*, Tuebingen, Germany, 2002.



Figure 8: Servoing results for object grasping.

- A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal. Discovering optimal imitation strategies. *Robotics and Autonomous Systems*, 47:2-3, 2004.
- F. Chaumette. Potential problems of stability and convergence in image-based and position-based visual servoing. In D. Kriegman, G . Hager, and A.S. Morse, editors, *The Confluence of Vision and Control*, pages 66–78. LNCIS Series, No 237, Springer-Verlag, 1998.
- A. D'Souza, S. Vijayakumar, and S. Schaal. Learning inverse kinematics. In *IEEE International Conference on Intelligent Robots and Systems*, Maui, USA, 2001.
- György Gergely, Harold Bekkering, and Ildikó Király. Rational imitation in preverbal infants. *Nature*, 415:755, Feb 2002.
- S. Hutchinson, G. Hager, and P. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, October 1996.
- M. Jaegersand. Visual servoing using trust region methods and estimation of the full coupled visual-motor jacobian. In *IASTED Applications of Control and Robotics*, Orlando, EUA, 1996.
- D. Kragic, L. Petersson, and H. I. Christensen. Visually guided manipulation tasks. *Robotics and Autonomous Systems*, 40(2-3):193–203, August 2002.
- Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *Trans. on Robotics and Automation*, 10(6):799– 822, 1994.
- M. Lopes, R. Beira, M. Praça, and J. Santos-Victor. An anthropomorphic robot torso for imitation: design and experiments. In *International Conference on Intelligent Robots and Systems*, Sendai, Japan, 2004.
- M. Lopes and J. Santos-Victor. Visual transformations in gesture imitation: What you see is what you do. In *International Conference on Robotics and Automation*, Taipei, Taiwan, 2003.
- Giorgio Metta. *Babybot: A Study on Sensori-motor Development*. PhD thesis, University of Genova, 1999.

- S. Nakaoka, A. Nakazawa, K. Yokoi, H. Hirukawa, and K. Ikeuchi. Generating whole body motions for a biped humanoid robot from captured human dances,. In *ICRA*, Taipei, Taiwan, 2003.
- Lorenzo Natale. Linking Action to Perception in a Humanoid robot: a Developmental Approach to Grasping. PhD thesis, University of Genova, 2004.
- C. Nehaniv and K. Dautenhahn. Mapping between dissimilar bodies: Affordances and the algebraic foundations of imitation. In *European Workshop* on Learning Robots, Edinburgh, Scotland, 1998.
- Erhan Oztop and Michael A. Arbib. Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87: 116–140, 2002.
- V. Gregory Payne and Larry D. Isaacs. *Human Mo*tor Development: a Lifespan Approach. Mayfield Publishing Company, California, USA, 4th edition, 1999.
- G. Rizzolatti, L. Fadiga, L. Fogassi, and V. Gallese. The space around us. *Science*, 277:190–191, 1977.
- S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Phil. Trans. of the Royal Society of London: Series B, Biological Sciences*, 358:537–547, 2003.
- A. van der Meer, F. van der Weel, and D. Lee. The functional significance of arm movements in neonates. *Science*, 267(5198):693–695, Feb 1995.
- S. Vijayakumar and S. Schaal. Locally weighted projection regression: An o(n) algorithm for incremental real time learning in high dimensional spaces. In *ICML*, Stanford, USA, 2000.
- Juyang Weng. The developmental approach to intelligent robots. In AAAI Spring Symposium Series, Integrating Robotic Research: Taking The Next Leap, Stanford, USA, Mar 1998.
- R. Zöllner and R. Dillmann. Using multiple probabilistic hypothesis for programming one and two hand manipulation by demonstration. In *IROS*, Las Vegas, USA, 2003.

Joint Attention Development in Infant-like Robot based on Head Movement Imitation

Yukie Nagai

National Institute of Information and Communications Technology 3-5 Hikaridai, Seika-cho, Soraku-gun, Kyoto, 619-0289 Japan yukie@nict.go.jp

Abstract

The ability to imitate others enables human infants to acquire various social and cognitive capabilities. Joint attention is regarded as a behavior that can be derived from imitation. In this paper, the developmental relationship between imitation and joint attention, and the role of motion information in the development are investigated from a viewpoint of cognitive developmental robotics. It is supposed in my developmental model that an infant-like robot first has the experiences of visually tracking a human face based on the ability to preferentially look at salient visual stimuli. The experiences allow the robot to acquire the ability to imitate head movement by finding an equivalence between the human's head movement and the robot's when tracking the human who is turning his/her head. Then, the robot changes its gaze from tracking the human face to looking at an object at which the human has also looked at based on the abilities to imitate a head turning and gaze at a salient object. Through the experiences, the robot comes to learn joint attention behavior based on the contingency between the head movement and the object appearance. The movement information which the robot perceive plays an important role in facilitating the development of imitation and joint attention because it gives an easily understandable sensorimotor relationship. The developmental model is examined in learning experiments focusing on evaluating the role of movement in joint attention. Experimental results show that the acquired sensorimotor coordination for joint attention involves the equivalence between the human's head movement and the robot's, which can be a basis for head movement imitation.

1 Introduction

Neonatal imitation is a remarkable capability in human development. Such behavior might tell us that infants can associate their own action with others' action they see. The ability to imitate enables infants to acquire social identification and further social and cognitive capabilities (Meltzoff and Moore, 1997). Through the experiences of reproducing others' action, infants come to be able to understand the meaning of the action and the others' intention. Joint attention (Scaife and Bruner, 1975; Butterworth and Jarrett, 1991; Moore and Dunham, 1995) is one of the capabilities that can be derived from imitation. It is defined as a behavior to look where someone else is looking by following his/her gaze. In other words, joint attention is regarded as a type of imitative behavior that one turns one's own head and eyes towards the same side as another turns his/hers.

In this paper, the developmental relationship between imitation and joint attention, and the role of movement information in the development are discussed. Many researchers in cognitive science and developmental psychology have been investigating the capabilities of imitation and joint attention as the basis for infant development (Moore and Dunham, 1995). However, it is difficult to find the study in which the developmental relationship between the two abilities was examined. As described above, joint attention is an imitative behavior to copy others' head and eyes turning. It can emerge in infant-caregiver interactions when either of them, mostly a caregiver, introduces an object into their dyadic interactions based on the imitation. Considering the developmental progress from the dyadic to the triadic interaction, i.e. joint attention, is important for understanding the social and cognitive development in infants. This paper presents the developmental progress by which an infant-like robot incrementally learns to imitate and establish joint attention through interactions with a human caregiver. It is discussed from a standpoint of cognitive developmental robotics (Asada et al., 2001)

what capabilities a robot should be equipped with for interacting with an environment and learning the experiences, and how a caregiver should encourage and support the robot's development. As a key for the consecutive development from imitation to joint attention, a robot employs movement as its perceptual information. It is known in infant development that the motion information facilitates the development of the two abilities, e.g. (Vinter, 1986; Moore et al., 1997). Infants are more able to imitate others' action and comprehend others' gaze when they are presented with the behavior with the movement rather than without the movement. On the basis of the knowledge, a learning model by which a robot acquires joint attention ability through the experiences of head movement imitation by using motion information is proposed.

The rest of the paper is organized as follows. First, the findings about imitation and joint attention in infants are referred, in which the role of movement in the development is suggested. The finding of head movement imitation is also indicated, which is considered as a basis for joint attention development. Then, the current robotics models of imitation and joint attention are reported. Various models have been proposed with the aim of investigating infant development and/or constructing intelligent robots. The problems that the models did not deal with the developmental progress between imitation and joint attention and did not utilize motion information are pointed out. Next, a developmental model by which a robot learns joint attention based on head movement imitation is proposed. By utilizing motion information, a robot incrementally learns to imitate head movement and achieve joint attention without any a priori or symbolic representation for perceptual information given by a designer. Experiments that examined the validity of the model by using an infantlike robot are then described. Finally, discussion and ongoing work are given.

2 Related work on imitation and joint attention

2.1 Findings from infant studies

Meltzoff and Moore (1977, 1989) investigated the ability to imitate in infants at a few days or a few weeks of age. They found that infants were able to imitate facial and manual gestures and head movements demonstrated by an adult. On the basis of the finding, Meltzoff and his colleagues (Meltzoff and Moore, 1997; Rao and Meltzoff, 2003) proposed an active intermodal mapping model as the mechanism for early facial imitation. According to their model, infants can imitate an action by evaluating the equivalence between the action they see and their own action in a supra-modal representational space. In contrast, Jacobson (1979) suggested that facial and manual gestures of infants could be elicited by the presentation of a moving object. She showed that a moving pen and a ball were as effective as the tongue model of an adult in eliciting tongue protrusion by infants, and that a dangling ring elicited as much hand opening and closing as the adult hand model. This finding suggests that the motion information which infants perceive plays an important role in their early imitation. Vinter (1986) also indicated the significance of motion information in infant imitation. She showed that infants were more likely to imitate facial and manual gestures when they were presented with the gestures with the movement rather than without the movement. The reason was conjectured that the movement which infants perceive is effective in encoding their perceptual information.

Joint attention development has also been suggested to be facilitated by motion information. Moore et al. (1997) compared the infants' ability to learn gaze following when infants were presented with the final static state of an adult's head turning and the ability when infants were presented with the head turning with the movement. Their comparison showed that only infants presented with the movement were able to learn to establish gaze following. Lempers (1979) studied the developmental change in the ability to comprehend deictic gestures of infants at 9 to 14 months of age. His observational results showed that motion information helped younger infants to understand others' pointing and gaze. Corkum and Moore (1998) investigated the origin of joint attention and found that infants have a developmental stage at which they respond sensitively to the movement of an adults' gaze shift. They also examined the learning performance of joint attention in infants by presenting the infants with unnatural situations in which an interesting target appeared in the opposite side to the direction of an adult's head turning. Their examination showed that infants did not acquire the behavior to look at the object by turning to the opposite side of the head turning, but acquired the behavior to follow the adult's head turning although they could not find any object. This means that the learning mechanism of joint attention is not only based on the contingency between the adult's head turning and the object activation but also facilitated by the physical characteristics of the adult action, i.e. the direction of the head movement. I suppose from the result that infants learn the relationship between their own action and others' action before learning to find an object based on the others' cue.

2.2 Computational and robotic models

In order to investigate infant development and/or construct intelligent robots, computational and robotic models of imitation and joint attention have been proposed based on the findings from infant studies. Demiris and his colleagues (Demiris and Hayes, 1996; Demiris et al., 1997) constructed a model of head movement imitation based on the scheme of the active intermodal mapping proposed by (Meltzoff and Moore, 1997). Their model enabled a robot to imitate a human's head movement by establishing an equivalence between the human's head posture, which was estimated from the movement detected as an optical flow, and the robot's posture, which was given as encoder values. Scassellati (1999) built a humanoid robot that could imitate yes/no nods of a human. In his model, a robot recognized the yes/no nods by detecting the cumulative displacement of a human face in the robot's vision and then drove the fixed-action patterns for moving the robot's head as an imitative behavior.

The author (Nagai et al., 2002, 2003) proposed developmental models by which a robot learned joint attention through interactions with a human caregiver. I investigated how a robot with limited and immature capabilities could acquire the joint attention ability based on the evaluation from a caregiver or based on the robot's ability to autonomously find a sensorimotor contingency through its experiences. Triesch and his colleagues (Carlson and Triesch, 2003; Lau and Triesch, 2004) introduced the scheme of reward-based learning for a computational developmental model of gaze following. They suggested that the infant abilities of preferential looking, habituation, and reward-based learning, and an environmental setup in which a caregiver looks at an object that an infant prefers to look at can be a basic set for the emergence of gaze following. Shon et al. (2004a,b) constructed a model by which a robot acquired the ability to establish joint attention based on the imitation of a human's head movement. In their model, the imitation was achieved based on the scheme of the intermodal equivalence mapping (Meltzoff and Moore, 1997). In other words, a robot could imitate a head movement by turning its head to the same posture as that of the human, which was estimated from an image pattern of the human head. Then, the imitation of the head movement enabled a robot to achieve joint attention by finding an object at which the human was looking based on a probabilistic model.

However, these models of robotic imitation and joint attention have problems that they did not utilize motion information detected from visual perception and that they learned the mechanism to estimate the posture of a human head by using the exact posture which could not be detected by a robot. The following section presents a developmental model by which a robot consecutively learns to imitate and establish joint attention by utilizing both static and motion information detected by itself.

3 Joint attention development based on head movement imitation

3.1 Developmental progress

The developmental progress of joint attention via head movement imitation is shown in Figure 1. The development is based on the infant abilities to interact with an environment and learn the experiences and encouragement by a caregiver.

An infant is supposed to have the capability to preferentially look at salient visual stimuli, such as a bright colored object and a human face. This basic capability enables an infant to interact with an envi-



Figure 1: The developmental progress of joint attention via head movement imitation.



Figure 2: A learning model of joint attention based on head movement imitation. The visual attention controller enables a robot to have experiences of preferentially looking at a human face and a salient object. Through the experiences, the robot learns the sensorimotor coordination to imitate a head movement and achieve joint attention through the lower three modules.

ronment and have experiences for learning to imitate and establish joint attention. In early development, an infant often has dyadic interactions with a caregiver because of the caregiver's encouragement. A caregiver attempts to involve an infant in face-to-face interactions and emotionally communicate with the infant by showing facial expressions and head movements. The caregiver's movement drives the infant to visually track the caregiver's face as an interesting target, which provides experiences for learning to imitate head movement. In other words, when the caregiver turns his/her head vertically or laterally, the infant also turns his/her head to almost the same direction by tracking the caregiver's face. As the result, the infant finds an equivalence between the movement of the caregiver's head and that of the infant's and consequently acquires the ability to imitate head movements.

In parallel with or following the learning of head movement imitation, an infant starts to learn to achieve joint attention. A caregiver introduces an object, at which an infant prefers to look, into their dyadic interactions by presenting the infant with the object near the line of the infant gaze. The caregiver attempts to control the infant attention by moving the object and shifting the caregiver's own gaze to the object. The caregiver's encouragement drives the infant to change his/her attention target. The infant shifts his/her gaze from looking at the caregiver to looking at the object based on the abilities to imitate the caregiver's head movement and preferentially look at a salient object. This provides an experience for learning joint attention. The infant can acquire the sensorimotor coordination for joint attention by finding a contingency between the caregiver's gaze shift and the appearance of the object.

3.2 Learning model of joint attention based on head movement imitation

Figure 2 shows a proposed model by which a robot incrementally learns to imitate head movements and

establish joint attention. The model consists of four modules: a visual attention controller, an image feature detector, a learning module, and a coordinator. The visual attention controller enables a robot to have experiences of looking at salient visual stimuli. The latter three modules enable the robot to learn the sensorimotor coordination for imitation and joint attention through the above experiences.

3.2.1 Visual attention controller

The visual attention controller enables a robot to have fundamental experiences for the development. This module enables a robot to preferentially look at salient visual stimuli, such as a human face and a bright colored object, in an environment. A human face and a salient object are respectively detected by template matching and using color information from a peripheral camera image. Figure 3 (a) shows an example of the peripheral image, in which a human face and a yellow object are indicated by rectangles. In this case, the robot is controlling its gaze to look at the human face at the center of the image. A motor command to look at the object can be generated by multiplying the horizontal and vertical displacement between the object and the center of the image by scalar values.

3.2.2 Image feature detector

The image feature detector extracts visual information needed to achieve imitation and joint attention. The detector extracts the edge image E of a human face and the optical flow F of the human's gaze shift from foveal camera images I_{t-1} , I_t . An example of the detected features is shown in Figure 3 (b)-(d), in which (b) shows the foveal camera image when the robot is gazing at the human face as shown in (a), and (c) and (d) show the edge image and the optical flow detected from the center area (168 × 168 pixels) enclosed by a rectangle in (b). The position of the enclosed area is fixed at the center of the foveal image. The foveal and peripheral cameras are mechanically fixed and controlled to gaze at a visual target at the center of the peripheral image.

The edge image E is generated by orientation selective filters. Four filters that are selective with respect to four orientations $(e_1, e_2, e_3, e_4) =$ $(-, \, \, \)$ extract edge images E_n , where n = $1, \ldots, 4$, each of which includes one oriented edge. The value of each pixel $E_n(x, y)$ is calculated as

$$E_n(x, y) = \begin{cases} 1 & \text{if } \epsilon_n(x, y) > \epsilon \text{threshold} \\ 0 & \text{otherwise,} \end{cases}$$



(a) peripheral camera image (b) foveal camera image: I_t



(c) edge image: E (d) optical flow: F (e) motor output: $o_{e'f}$

Figure 3: An example of input-output datasets, in which (a) and (b) show a peripheral and a foveal camera image when the robot is looking at the human; (c) and (d) show the edge image and the optical flow detected from the center area in (b); (e) shows motor output to follow the human gaze, which is encoded in motion direction selective neurons.

where

$$\epsilon_n(x, y) = \left| \sum_{i=-1}^{1} \sum_{j=-1}^{1} \alpha_n(i, j) I(x+i, y+j) \right| \\ - \left| \sum_{i=-1}^{1} \sum_{j=-1}^{1} \beta_n(i, j) I(x+i, y+j) \right|.$$
(1)

(x, y) indicate a position in a camera image, and the coefficients, $\alpha_n(i, j)$ and $\beta_n(i, j)$, are given as

$$\boldsymbol{\alpha}_1 = \boldsymbol{\beta}_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & -1 & -1 \end{bmatrix}, \ \boldsymbol{\beta}_1 = \boldsymbol{\alpha}_3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & 1 & -1 \end{bmatrix},$$
$$\boldsymbol{\alpha}_2 = \boldsymbol{\beta}_4 = \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}, \ \boldsymbol{\beta}_2 = \boldsymbol{\alpha}_4 = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 0 \end{bmatrix},$$

where

$$\boldsymbol{\alpha}_{n} = \begin{bmatrix} \alpha_{n}(-1,-1) & \alpha_{n}(0,-1) & \alpha_{n}(1,-1) \\ \alpha_{n}(-1, 0) & \alpha_{n}(0, 0) & \alpha_{n}(1, 0) \\ \alpha_{n}(-1, 1) & \alpha_{n}(0, 1) & \alpha_{n}(1, 1) \end{bmatrix}.$$
(2)

Figure 3 (c) shows the edge image E combining E_n (n = 1, ..., 4), in which edges with one of the four orientations, -, \setminus , |, and /, are colored red, cyan, blue, and green, respectively. The edge image provides information to estimate the static direction of the human head and allows the robot to acquire the accurate sensorimotor coordination to achieve head movement imitation and joint attention.

The image feature detector also extracts the optical flow F. The center area of the foveal image is divided into small areas called receptive fields (24 × 24 pixels). The optical flow F^k in the *k*-th receptive field is calculated as the cumulative displacement of the image feature in the receptive field over ten image frames:

$$\boldsymbol{F}^{k} = \begin{bmatrix} \sum^{10 \text{frames}} (x_{k} - px) \\ \sum^{10 \text{frames}} (y_{k} - py) \end{bmatrix}, \quad (3)$$

where (x_k, y_k) and (px, py) are the center position of the k-th receptive field in I_t and that of the corresponding image area detected by template matching in I_{t-1} , respectively. Figure 3 (d) shows the optical flow detected when the human changes her gaze from looking straight at the robot's camera to looking at the yellow object shown in (a). Like the edges, the flows are drawn with four colors. Although the optical flow cannot provide enough information to infer the exact direction of the human head compared with the edge information, it gives a rough but easily understandable relationship with the movement direction of the human's head turning. Therefore, the flow information should enable the robot to quickly acquire rough sensorimotor coordination for head movement imitation and joint attention.

In addition, the flow information is utilized as a cue for the robot to control the timing of its own head turning. The temporal change in the amount of the optical flow indicates the start and end of a human's head turning. In other words, when the flow becomes zero after exceeding an upper threshold, this means that a human has shifted his/her head direction from looking at one location to looking at another and is gazing at a certain location. Based on this mechanism, the robot obtains the input data of the optical flow when the flow has a maximum value and the edge image when the flow becomes zero. This enables the robot to immediately follow a human's head turning without any explicit cue.

3.2.3 Learning module

This module learns the sensorimotor coordination between the edge input and motor output and between



(a) the encoding of edge input (b) the encoding of flow input

Figure 4: The encoding of detected image features into the input neurons, in which (a) and (b) show the encoding of edge and flow inputs into the four orientation selective neurons and the eight directions selective neurons, respectively. The length of a line in each circle denotes the activity of the neuron.

the optical flow and motor output through two independent neural networks (see Figure 2). The neural network for the edge input (the edge-NN) consists of three layers: input, hidden, and output layers, because edge information is difficult to interpret into the human's head direction. In contrast, the neural network for the optical flow input (the flow-NN) has two layers: input and output layers, because flow information gives an easily understandable relationship with the motor output to imitate the human's head movement and achieve joint attention.

Input to the edge-NN is represented as activities of four kinds of neurons that are selective to four orientations. Figure 4 (a) shows edge input encoding into the selective neurons. The activities of the four neurons $a_{e_n}^k$ (n = 1, ..., 4) in the *k*-th receptive field are calculated as

$$a_{e_n}^{k} = E_n^{k} / \max_k \sum_{m=1}^{4} E_m^{k},$$

where $E_n^{k} = \sum_{x_k} \sum_{y_k} E_n(x, y).$ (4)

 $E_n(x, y)$ is given by (1), and E_n^k means the amount of the edge e_n in the k-th receptive field. In the bottom of Figure 4 (a), the length of a line in each circle shows the activity of each neuron. No line means that the activity is zero.

Like the encoding of edge input, the optical flow

is encoded in eight kinds of neurons that are selective to eight directions $(f_1, f_2, \ldots, f_8) = (\leftarrow, \checkmark, \ldots, \checkmark)$ as shown in Figure 4 (b). The activities of the eight neurons $a_{f_n}^k$ $(n = 1, 2, \ldots, 8)$ in the *k*-th receptive field are calculated as

$$a_{f_n}^k = \begin{cases} \mathbf{F}^k \cdot \mathbf{u}_n / \max_k \|\mathbf{F}^k\| & \text{if } \mathbf{F}^k \cdot \mathbf{u}_n \ge 0\\ 0 & \text{otherwise,} \end{cases}$$
(5)

where F^k is given by (3), and u_n are unit vectors in eight directions. The activities of the eight neurons are also drawn as the length of the arrows as shown in Figure 4 (b). The methodology of coding edge and flow information is based on physiological evidence that the visual cortex in some animates has orientation selective neurons (Hubel and Wiesel, 1959) and motion direction selective neurons (Barlow and Hill, 1963). The similarity in the representation of edge and flow inputs leads to the possibility that the robot translates a well-acquired sensorimotor coordination in the edge-NN or the flow-NN into the other.

Outputs from the edge- and flow-NNs are represented as the activities of eight neurons, $o_{e'_n}$ and o_{f_n} (n = 1, ..., 8), which are selective to eight motion directions $(e'_1, ..., e'_8) = (f_1, ..., f_8) = (\leftarrow$ $,..., \checkmark)$, respectively. Figure 3 (e) shows an example of the activities of the output neurons. The representation is similar to that of encoded optical flow data. The activities of the output neurons are decoded into a motor command $\Delta \theta$ to move the robot's head by the coordinator described in the next section.

3.2.4 Coordinator

This module coordinates motor outputs from the edge- and flow-NNs. In the experiments discussed here, the robot uses a simple method that generates a motor command $\Delta\theta$ by decoding the mean value of the two outputs:

$$\boldsymbol{\Delta\theta} = \begin{bmatrix} \Delta\theta_{pan} \\ \Delta\theta_{tilt} \end{bmatrix} = \begin{bmatrix} g_{pan} \sum_{n} u_{nx} o_{e'f_n} \\ g_{tilt} \sum_{n} u_{ny} o_{e'f_n} \end{bmatrix}, \quad (6)$$

where g_{pan} and g_{tilt} are scalar gains; u_{n_x} and u_{n_y} are the horizontal and vertical components in u_n ; $o_{e'f_n}$ is the mean value of $o_{e'_n}$ and o_{f_n} . A motor command to move the robot's head is represented as displacement angles in the pan and tilt directions.

3.3 Learning processing

Employing the model, a robot has two-staged learning. First, a robot learns the sensorimotor coordination to imitate head movements. As a human turns his/her head vertically and laterally in front of the robot, the robot also turns its head to almost the same direction by tracking the human face based on the visual attention controller. Through the experiences, when the robot detects simultaneous activation of the input and output neurons that are selective to the same directions in the flow-NN, it learns the equivalence of the movement by multiplying the connecting weights between the neurons. This leads to the ability to imitate head movements. Next, the robot learns the sensorimotor coordination for joint attention. The human starts to introduce an object into the human-robot dyadic interactions. When the human shifts his/her gaze direction from the robot to the object by turning his/her head, the robot first imitates the head movement based on the acquired sensorimotor coordination and then changes its gaze from looking at the human to looking at the object based on the visual attention controller. This provides a sensorimotor experience of joint attention. The robot learns the sensorimotor coordination in the edge- and flow-NNs by back propagation based on the input-output dataset obtained in the above process and consequently acquires joint attention ability.

4 Preliminary experiment

4.1 Experimental setup

As a preliminary experiment, the model was evaluated with a focus on the role of movement in learning joint attention. The model was implemented into an infant-like robot, called *Infanoid* (Kozima, 2002), shown in Figure 5, which was developed by our group



Figure 5: Human-robot joint attention, in which an infant-like robot, called *Infanoid*, is looking at the stuffed toy that the human is looking at.

as a tool for investigating the cognitive development in human infants. Infanoid has a stereo-vision head with three degrees of freedom (DOFs) in its neck (one for the pan and two for the tilt directions) and three DOFs in its eyes (two for the each pan and one for the common tilt directions). Each eye has two color CCD cameras: a peripheral camera and a foveal camera, and the two camera images from the left eye were used in the experiment. The three DOFs in the neck were used to move the robot's head while the three DOFs in the eyes were fixed at the center positions. The displacement angle $\Delta \theta_{tilt}$ derived from (6) was equally divided into the two tilt DOFs in the neck.

A human sat face to face with Infanoid and interacted with it by using a salient object. In every trial, the human replaced the object at random positions and then changed her gaze from looking at the robot to looking at the object by turning her head. The human always looked at the object in front of her face.

4.2 Evaluating the role of movement in learning joint attention

The role of motion information in learning joint attention was evaluated. In this experiment, Infanoid learned to establish joint attention without learning to imitate. In other words, the robot learned a contingency between the human's head turning and the object appearance to acquire the sensorimotor coordination for joint attention through the edge- and flow-NNs without using any pre-acquired sensorimotor coordination to imitate head movements.

Figure 6 shows the changes in joint attention performance over the learning period, in which the horizontal and vertical axes respectively denote the learning step and the success rate of joint attention. The success of joint attention means that the robot looks at the object at which the human is looking within ± 8 degrees of error. The learning experiment was conducted off-line by repeatedly using 200 input-output datasets acquired beforehand, and the sensorimotor coordination acquired through learning was evaluated in joint attention experiments every 200 learning steps. The red line plots the result when the model used both the edge and flow inputs. The blue and green lines plot the results when the model used only the edge or the flow input, respectively. The graph shows the mean result of fifty experiments with different initial conditions and its standard deviation. Comparing the results for when the robot used either the edge or the flow input, it is confirmed that the flow input accelerated the start-up time of learning while the edge input gradually improved the task perfor-



Figure 6: The change in the task performance of joint attention over the learning period. The red, blue, and green lines indicate the results when the model utilized both the edge and flow inputs, only the edge input, and the flow input, respectively.

mance. This complementary result can be expected from the characteristics of the two inputs. By using both the edge and flow inputs, the model enabled the robot to quickly acquire the high performance of joint attention by combining the advantages of the two inputs.

4.3 Joint attention experiment after learning

The acquired sensorimotor coordination was evaluated in joint attention experiments. Figure 7 (a) and (b) show the two cases of input-output datasets when the robot attempted to achieve joint attention based on the acquired NNs. In case (a), the human shifted her gaze from looking at the robot to looking at an object in the outer left side of the foveal image. In case (b), the human shifted her gaze direction from the robot to an object in the outer lower right of the foveal image. The upper side of each figure shows the change in the foveal image when the robot shifted its head direction based on the output from the coordinator shown in the lower side. From these results, we can see that both the edge-NN and the flow-NN generated appropriate output to achieve joint attention. In these two cases, the robot was able to find the object at which the human was looking and establish joint attention. The success rate of joint attention with the same person as in the learning experiment was 90% (18/20 trials). In addition, we can confirm from this result that the flow-NN acquired one-to-one correspondence between the activities of the input and output neurons. The direction of the motor output from the flow-NN is



(a) In the case that the human shifted her gaze direction from the robot to an object in the outer left side of the foveal image.



(b) In the case that the human shifted her gaze direction from the robot to an object in the outer lower right of the foveal image.

Figure 7: The input-output datasets when the robot attempted to achieve joint attention based on the acquired NNs. The robot was able to establish joint attention in these two cases. clearly corresponding with the same direction of the optical flow. This means that the sensorimotor equivalence, which should be acquired through imitation learning, is also utilized in joint attention.

5 Discussion and ongoing work

This paper has presented a developmental model by which a robot learns joint attention based on head movement imitation. The preliminary experiments showed that the model accelerated the learning of joint attention by using movement information and that the equivalence of self and other movement was utilized to achieve joint attention. This result supports the idea that joint attention emerges through the experiences of head movement imitation. Ongoing work is to examine that learning to imitate head movements facilitates the development of joint attention. This is expected to lead to the possibility to reveal the role of other neonatal imitation, such as tongue protrusion and hand opening-closing, in the development of social and cognitive capabilities of infants. Another issue to be solved is to develop a mechanism that enables a robot to recognize not only head directions but also gaze directions. It was assumed in the experiments that a human shifted his/her gaze by turning his/her head and looked at an object in front of his/her face. This assumption is likely in joint attention by infants. However, infants can acquire the ability to recognize gaze directions. To solve the problem, I will apply a mechanism that changes the resolution of the receptive fields in the NNs as learning proceeds. Such mechanism will increase the resolution around the image area including important facial features, e.g. eyes and mouth, and consequently enable a robot to acquire the ability to recognize gaze directions.

References

- Minoru Asada, Karl F. MacDorman, Hiroshi Ishiguro, and Yasuo Kuniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37:185–193, 2001.
- H. B. Barlow and R. M. Hill. Selective sensitivity to direction of movement in ganglion cells of the retina. *Science*, 139:412–414, 1963.
- George Butterworth and Nicholas Jarrett. What minds have in common is space: Spatial mechanisms serving joint visual attention in infancy.
British Journal of Developmental Psychology, 9: 55–72, 1991.

- Eric Carlson and Jochen Triesch. A computational model of the emergence of gaze following. In *Proceedings of the 8th Neural Computation and Psychology Workshop*, 2003.
- Valerie Corkum and Chris Moore. The origins of joint visual attention in infants. *Developmental Psychology*, 34(1):28–38, 1998.
- J. Demiris, S. Rougeaux, G. M. Hayes, L. Berthouze, and Y. Kuniyoshi. Deferred imitation of human head movements by an active stereo vision head. In *Proceedings of the 6th IEEE International Workshop on Robot and Human Communication*, pages 88–93, 1997.
- John Demiris and Gillian Hayes. Imitative learning mechanisms in robots and humans. In *Proceedings* of the 5th European Workshop on Learning Robot, pages 9–16, 1996.
- D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *Journal* of *Physiology*, 148:574–591, 1959.
- Sandra W. Jacobson. Matching behavior in the young infant. *Child Development*, 50:425–430, 1979.
- Hideki Kozima. Infanoid: A babybot that explores the social environment. In K. Dautenhahn, A. H. Bond, L. Canamero, and B. Edmonds, editors, *Socially Intelligent Agents: Creating Relationships* with Computers and Robots, chapter 19, pages 157–164. Amsterdam: Kluwer Academic Publishers, 2002.
- Boris Lau and Jochen Triesch. Learning gaze following in space: a computational model. In *Proceedings of the Third International Conference on Development and Learning*, 2004.
- Jacques D. Lempers. Young children's production and comprehension of nonverbal deictic behaviors. *The Journal of Genetic Psychology*, 135:93–102, 1979.
- Andrew N. Meltzoff and M. Keith Moore. Imitation of facial and manual gestures by human neonates. *Science*, 198:75–78, 1977.
- Andrew N. Meltzoff and M. Keith Moore. Imitation in newborn infants: Exploring the range of gestures imitated and the underlying mechanisms. *Developmental Psychology*, 25(6):954–962, 1989.

- Andrew N. Meltzoff and M. Keith Moore. Explaining facial imitation: A theoretical model. *Early Devel*opment and Parenting, 6:179–192, 1997.
- Chris Moore, Maria Angelopoulos, and Paula Bennett. The role of movement in the development of joint visual attention. *Infant Behavior and Development*, 20(1):83–92, 1997.
- Chris Moore and Philip J. Dunham, editors. *Joint Attention: Its Origins and Role in Development*. Lawrence Erlbaum Associates, 1995.
- Yukie Nagai, Minoru Asada, and Koh Hosoda. Developmental learning model for joint attention. In Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 932–937, 2002.
- Yukie Nagai, Koh Hosoda, Akio Morita, and Minoru Asada. A constructive model for the development of joint attention. *Connection Science*, 15(4):211–229, 2003.
- Rajesh P. N. Rao and Andrew N. Meltzoff. Imitation learning in infants and robots: Towards probabilistic computational models. In *Proceeding of Artificial Intelligence and Simulation of Behaviour: Cognition in Machines and Animals*, 2003.
- M. Scaife and J. S. Bruner. The capacity for joint visual attention in the infant. *Nature*, 253:265–266, 1975.
- Brian Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills on a humanoid robot. In C. Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, volume 1562 of LNCS, pages 176–195. Springer-Verlag, 1999.
- Aaron P. Shon, David B. Grimes, Chris L. Baker, Matthew W. Hoffman, Shengli Zhou, and Rajesh P. N. Rao. Probabilistic gaze imitation and saliency learning in a robotic head. Technical report, University of Washington CSE, 2004a.
- Aaron P. Shon, David B. Grimes, Chris L. Baker, and Rajesh P. N. Rao. A probabilistic framework for model-based imitation learning. In *Proceedings of* the 26th Annual Meeting of the Cognitive Science Society, 2004b.
- Annie Vinter. The role of movement in eliciting early imitations. *Child Development*, 57:66–71, 1986.

Grey Parrots Do Not Always "Parrot": Roles of Imitation and Phonological Awareness in the Creation of New Labels from Existing Vocalizations

Irene M. Pepperberg

Radclffe Instititue for Advanced Studies Cambridge, MA 02138 impepper@media.mit.edu Brandeis University, Department of Psychology Waltham, MA 02454

Abstract

Evidence exists for a form of imitation, vocal segmentation, by a Grey parrot: Data show that the bird understands that his labels are comprised of individual units that can be recombined in novel ways to create a novel referential vocalization; that is, a novel act. Previous data suggested, but could not substantiate, this behaviour. Such evidence implies that a parrot not only has phonological awareness but also demonstrates true imitation, and has implications for programming speech.

1 Introduction

Given that imitation primarily involves the intentional copying of an otherwise improbable, novel act (Thorpe, 1963), the intentional, referential reproduction of a novel English vocalization by a Grey parrot (Psittacus erithacus) is a likely candidate for imitative behavior. And, given that imitation can also be seen as the integration of a number of familiar actions in novel ways to produce that novel act (e.g., Piaget, 1962; review in Arbib, 2002), of particular interest is what happens when the targeted novel vocalization can be constructed from related elements already in the parrot's repertoire. This particular type of combinatory behavior is actually a form of vocal segmentation. Successful segmentation shows that the bird understands that his *existent* labels are comprised of *individual units* that can be recombined in novel ways to create novel vocalizations. Previous data suggested, but could not substantiate, this behavior; current data does just that. Moreover, such evidence implies that a parrot has phonological awareness.

Demonstrating segmentation by Grey parrot would be an important milestone in comparing animal and human cognitive and communicative abilities. Although Grey parrots already use elements of English speech referentially (Pepperberg, 1999),¹ these birds are still sometimes regarded as mindless mimics. At least two reasons exist for that belief. One reason, that avian imitation of English speech does not involve intentional, accurate reproduction of human articulatory acts (as far as is possible with parrot anatomy), has been countered previously (e.g., Patterson and Pepperberg, 1994, 1998; Pepperberg, 2002). Another reason involves segmentation: Only limited evidence exists that parrots, or any animal taught a human communication code, can indeed segment the human code, that is, recombine existing labels intentionally either to describe novel situations or, for example, to produce a phrase to request novel items-rather than simply produce several referential labels that, by chance, appropriately apply to the situation (Fouts & Rigby, 1977; Pepperberg, 1999; Savage-Rumbaugh et al., 1993). Such intentional creativity is, in contrast, common in the earliest stages of normal human language acquisition (de Boysson-Bardies, 1999; Greenfield, 1991; Marschark, Everhart, Martin, & West, 1987; Tomasello, 2003). Another form of segmentation, the intentional recombination of existing phonemes (parts of words) or their approximations to create or reproduce what is for the subject a novel targeted utterance (Greenfield, 1991; Peperkamp, 2003), has not previously been reported in animals; it is not only considered basic to human language development (Carroll, Snowling, Hulme, & Stevenson, 2003), but also a uniquely human trait.

Such phonetic awareness, which requires understanding that words are made up of a finite number of sounds that can be recombined into an almost

¹ No claim is made that Alex's speech is isomorphic with human language (e.g., Alex cannot discourse about the weather), only that the elements that he does produce have been documented as being used referentially; labels are both understood and used in contexts that differ from and extend beyond training conditions.

infinite number of patterns (limited only by the constraints of a given language)-the parsing of a complex entity into pieces that are then integrated into a new schema that represents the imitated act (Arbib, 2002)-has additionally been considered a trait that is acquired over time. Children, for example, apparently shift from recognizing and producing words holistically (a simple form of imitation, Studdert-Kennedy, 2002; Arbib, in press) to recognizing words as being constructed via a rule-based phonology sometime around three years of age or later (Carroll et al., 2003; Vihman, 1996); furthermore, manipulation of individual parts of words is presumed to require development of an internal representation of phonological structure (Byrne & Liberman, 1999). That is, in order to sound outi.e., to imitate, rather than mimic-a novel label, the child must segment the stream of sound into discrete elements, recognize a match between those elements and elements (or close approximations) that exist in its own repertoire, and then recombine these elements in an appropriate sequence (see Gathercole & Baddeley, 1990; Treiman, 1995; Arbib, in press). Moreover, children's ability to focus on the sounds of words and sound elements of words rather than solely on word meaning appears to be assisted by training in sound-letter associations (Carroll et al., 2003; Mann & Foy, 2003). Most animals, lacking speech, are never exposed to, nor trained nor tested on, such issues of phonological awareness or imitation, nor are they expected to have internal representations of phonemes.²

Evidence now exists for this form of imitation (vocal segmentation, phonological awareness) by a Grey parrot: Here I show that my oldest speechtrained subject, Alex, understands that his labels are made of individual phonological units that can be recombined in novel ways to create novel vocalizations. Such evidence implies that parrots not only use English labels referentially, but also understand how such labels are created from independent sound patterns. I also suggest that the bird's ability is indeed a learned behaviour, is not uniquely human, and is dependent upon having considerable experience with English speech and sound-letter training. My younger birds, lacking such training, do not engage in such behaviour. Of specific interest is that this behaviour occurred in contrast to my parrots' customary patterns of label acquisition and demonstrates the steps the bird goes through in producing the imitated label.

2 Method

2.1 Subjects

Subjects were two Grey parrots (*Psittacus erithacus*). Alex, 27 years old, had had 26 years of intense training in interspecies communication (Pepperberg, 1999); Arthur, 3½ years old, had had the equivalent of about a year of comparable interspecies communication training. The birds live in a laboratory setting at all times. Their housing and day-to-day care have been previously described (e.g., Pepperberg & Wilkes, 2004). Using the training technique described below, Alex had previously learned to identify, request, refuse, categorize, and quantify a large number (>100) of objects referentially using English speech sounds (Pepperberg, 1999), and Arthur had already acquired four referential labels (Pepperberg & Wilkes, 2004).

2.2 General Procedures

Grey parrots in my laboratory generally learn referential English speech (e.g., to comprehend and produce labels for objects, colors, shapes; to answer questions about concepts of number, category, relative size, absence, same/different) via training with the Model/Rival (M/R) procedure. (Pepperberg, 1999). This procedure, introduced by Todt (1975) and adapted by Pepperberg (1981), involves threeway interactions between two human speakers and the avian student. While the bird watches, two humans handle an object in which the bird has demonstrated interest; one (the trainer) then questions the other (the parrot's model and rival for the trainer's attention) by using phrases like "What's here?", "What toy?", "What do you want?" etc. The trainer rewards correct responses with the object, thus demonstrating a label's referentiality (the connection between the label and the object to which it refers) and functionality (e.g., showing how the label can be used, as in a request). Humans model errors (e.g., poor pronunciation or other errors similar to the bird's at the time) and demonstrate the consequences of erring by having the trainer say "No!", look away, and briefly remove the object from sight; a correction procedure then follows. The parrot is also questioned and rewarded for an attempt or scolded for an error. Humans exchange roles of trainer and model on a regular basis, thus showing that one individual is not always the questioner and the other the respondent; the bird thus learns to interact with all the humans. Without such role reversal, a bird interacts with only the human acting as trainer (Todt, 1975). By involving the parrot in these interactions, the humans can adjust training to the bird's level (e.g., successively requiring better pronunciation).

This technique was used to train both Arthur and Alex on the label "spool". Arthur was trained first.

 $^{^2}$ Nonhuman primates have been trained and tested on their ability to segment human speech sounds (e.g., Newport et al., 2004), but not on sound-letter associations or on productive recombination of speech elements.

After Arthur's training, Alex began to show interest in the object, which he had previously ignored. Subsequent to Arthur's training, Alex, when given the item, began to chew it apart or roll it around his play stand. We therefore decided to initiate M/R training on the object for Alex.

2.3 Phoneme Training

For several years, Alex had received M/R training to associate Arabic letters B, CH, I, K, N, OR, S, SH, T with their corresponding appropriate phonological sounds (e.g., /bi/ for BI), with the plastic or wooden labels as his reward. Although his accuracy was greater than chance (generally about 50%, p<.01, chance of 1/9), it was never high enough (i.e., ~80%) to claim he had mastered the task.

2.4 Taping and Sonagraphic Analysis

Birds were taped with an AKG-70 microphone directly into an IBM T20 computer; wav. files were edited with *Audacity* and made into sonagrams by Dr. Diana Reiss (Wildlife Conservation Society, NYC) and Dr. Donald Kroodsma (UMass-Amherst) using *Raven* (Cornell Laboratory of Ornithology). An additional sonagram of Alex's and my formants was made by Dr. Ofer Tchernichovski (CCNY) with *Sound Analysis Pro* (http://ofer.sci.ccny.cuny.edu).

3 Results

The present data for my oldest subject, Alex, were obtained after our youngest subject, Arthur, had acquired the label *spool* to refer to wooden or plastic bobbins. The birds' labels usually appear in sessions initially as rudimentary patterns—first a vocal contour, then with vowels, finally with consonants (Patterson & Pepperberg, 1994, 1998) and Arthur's production followed the customary acquisition pattern, that is, beginning with /u/ ("ooo") and ending with a distinct, fully-formed "spool" (/spul/; see Figure 1A;



Figure 1: Sonagram of (A) Arthur's "spool": (B) Pepperberg's "spool"

the IPA transcription is approximate; Arthur used

more of a whistle than an actual /p/ and vowel sound). Although Arthur occasionally mis-identified the object as "wool" or "wood" (some of his other labels at the time, Pepperberg & Wilkes, 2004), with "wood" sometimes being a correct response for the object's material and "wool" being a reasonable phonological error, he did not consistently use such labels during training and was correct 87.5% on testing (Pepperberg & Wilkes, 2004).

Unlike Arthur, and unlike his usual form of acquisition, Alex, during training after watching Arthur playing with the object, began using a combination of existing phonemes and labels to identify the object: /s/ (trained independently in conjunction with the Arabic letter, S) and *wool*, to form "s" (pause) "wool" ("s-wool"; /s-pause-wUl/; Fig. 2).



Figure 2: Sonagram of Alex's "s-wool"

Note that no labels existed in Alex's repertoire that contained /sp/, nor did he have the labels "pool" or "pull", or any other label that included /Ul/; he did have labels such as "paper", "peach", "parrot", "pick", etc (Pepperberg, 1999). He retained this "s-wool" formulation for almost a year of M/R training, with no change whatsoever in the form of his production, although normally only about 25 M/R sessions (at most, several weeks of training) are sufficient for learning a new label (Pepperberg, 1999). The third parrot in the lab, Griffin, who was just beginning training on phonemes, and who heard exactly the same information from Arthur's sessions, did not exhibit this behaviour.

At the end of this year-long period, Alex spontaneously produced "spool", perfectly formed (/spul/; see Fig. 3), when I rewarded Arthur with the spool for producing the label. Thus, Alex added the /p/ phoneme and also shifted the vowel toward the appropriate /u/ sound. (Interestingly, both Alex's and my /u/'s are dipthongs, differing slightly from standard American English productions; Patterson & Pepperberg, 1994); note Alex's vowel changes from Fig. 2 to Fig. 3). Because the label "spool" appeared without any intermediary form from that of "swool", no statistical or other analysis of the process of change was possible.³



Figure 3: Sonagram of Alex's "spool"

Alex's and Arthur's productions differ significantly in auditory and sonagraphic patterns (see Figs. 1A, 3; .wav files on request), so that Alex did not simply learn to mimic Arthur's production; Arthur's utterance incorporated a avian whistle-like quality whereas Alex's utterance sounded distinctly human. Alex's vocal pattern more closely resembles mine (see Fig. 1B), even though I did less than onetenth of the M/R training. I had, however, done the majority of training on wool almost 20 years earlier (Pepperberg, 1999). In general, Alex's formant structures closely approach, although are not identical to, my own (Fig 4; see Patterson & Pepperberg, 1994, 1998 for detailed analyses of the similarities and differences between Alex's and my speech acts; Alex's patterns all closely approach mine, although identity is impossible by virtue of the difference in vocal tract sizes and Alex's lack of lips).⁴



Figure 4 A,B: Closeup of (A) Alex's formants for vowel part of "s-wool", (B) Alex's formants for the vowel part of "spool"



Figure 4C: Pepperberg's formants for the /p/ and vowel part of "spool".

4 Discussion

As noted above, parrots usually acquire labels by building the sound patterns gradually, beginning with vowels (Patterson & Pepperberg, 1994, 1998). Such behaviour may simply reflect the relative ease with which sounds that are more tonal can be produced relative to those that require, for example, plosive qualities in a subject lacking lips. Nevertheless, completely-formed new labels did occasionally materialize after minimal training and without overt practice (Pepperberg, 1983). In these cases, however (e.g., Alex's production of "carrot" the day after asking us what we were eating, or of the novel label "banerry" to refer to an apple), such utterances appeared fully-formed, with immediacy and no overt practice (Pepperberg, 1999). Even though the label generally contained sounds already in the repertoire (e.g., for "carrot", the /k/ from key, the remainder from parrot; "banerry" was derived from "banana"-"cherry"), neither my students nor I could convincingly argue that Alex had deliberately parsed labels in his repertoire to match a targeted utterance or to form novel vocalizations. A related argument could be made for Alex's abilities to referentially produce words that form minimal pairs (Patterson & Pepperberg, 1998). So, although Alex could state "Want corn" versus "Want cork", or "Want tea" versus "Want pea" (and refuse the alternatives), which suggests an ability to segment phonemes from the speech stream (somewhat like nonhuman primates; see Newport et al., 2004), we could not claim that he deliberately parsed these labels when learning to produce them.

That is, we could not claim that he acoustically represented labels as do humans with respect to phonetic categories and understood that his labels are made of individual elements that can be recombined in various ways to produce new ones. Possibly production of "carrot" was potentiated by his already being able to manipulate his vocal tract to produce such sounds, or the new labels were simply created from phonotactically probable sequences involving beginnings and ends of existent labels (Storkel, 2001), or, in the case of "banerry", from semantic relations. Moreover, data (Patterson & Pepperberg, 1994, 1998) demonstrating that he (a)

³ For Wav. forms, contact <u>impepper@media.mit.edu</u>. Production of "spool" first occurred during a taping session with the BBC, and after production of "s-wool" during the same session; thus the vocalization was evaluated by observers who had had no possible previous connection with training sessions.

⁴ See Beckers, Nelson, and Suthers (2004) for discussion of tongue placement and formation of true formants in parrot vo-calizations.

recognizes small phonetic differences ("tea" vs "pea") as meaningful, (b) produces initial phonemes differently depending upon subsequent ones (/k/ in "key" vs "cork"), and (c) consistently recombines parts of labels according to their order in existent labels (i.e., combines beginnings of one label with the ends of others—after analyzing over 22000 vocalizations, we never observed backwards combinations such as "percup" instead of "cupper/copper"; Pepperberg, Brese, & Harris, 1991), merely imply but did not prove that he engages in such *top-down* processing (Ladefoged, 1982).

Even the two closest behaviour patterns previously reported that suggest some form of label parsing (Pepperberg, 1990; Pepperberg et al., 1991), which both involve solitary sound play, differ from the bootstrapping described here. In one behaviour, Alex produced strings such as mail chail benail in private practice before producing the targeted, trained label *nail* (Pepperberg et al., 1991). The *nail* situation differs from bootstrapping in that the combinations of phonemes did not seem to be a deliberate attempt to create a new label from specific sound patterns that resembled the target, but rather to be deliberate play within a range of existent patterns in an attempt to hit on a correct pairing that matched some remembered template. That is, Alex's behaviour demonstrated an understanding of the combinatory nature of his utterances, but did not show that he understood how to segment the novel targeted vocalization exactly, then match its components to those in his repertoire in order to create trained label. In the second behaviour (Pepperberg, 1990), Alex babbled strings such as grape, grain, chain, cane in the absence of specific objects but in the presence of his trainers. Although these labels could quickly be referentially mapped onto physical objects, we had no reason to believe that production of such babbled strings was intentional, other than to gain the attention of trainers. Note that with the exception of *grape*, such labels would not have been used in the laboratory. The rhyme awareness demonstrated in these behaviour patterns, although separate from phoneme awareness, is still considered closely aligned to children's language skills (see Mann & Foy, 2003), and again supports the argument that Alex views his labels as being constructed from individual sound patterns.

The current data, when taken in combination with previous evidence, however, suggests that at least one parrot, much like a child, can actually apply a phonological rule derived from knowledge of its repertoire: recognize that sounds such as "car" and "pet" can be recombined for use in identifying and creating a totally distinct object—*carpet*— whose label has no referential correlation to the original utterances. That is, Alex appeared to form the closest match based on segmentation and on-

set+nucleus+rhyme (Storkel, 2002). Arguably, the data presented here could be considered stronger if Alex had had the labels *pull* or *pool* in his repertoire and had initially produced either "s (pause) pull" or "s (pause) pool". Note, however, that /p/ is particularly difficult for a parrot, lacking lips, to produce (Patterson & Pepperberg, 1998); Alex's first attempts at "peach", for example, sounded like "cheech" (Neal, 1996), and Patterson and I (1998) have suggested that he may be using a form of esophageal speech for /p/. I thus believe that his production of "s-wool" is actually more important, because, not having exact matches, he took the closest, readily-available sounds in his repertoire (i.e., "wool" is the only label out of approximately 50 documented in his repertoire that resembles "spool") to form the initial attempts at a novel vocalization, and by so doing, made the process transparent to his human trainers.

Another interesting issue exists concerning Alex's behaviour, and addresses the issue of whether he has simply shown a sensitivity to sound similarity. Exactly because of the difficulty of producing /p/, Alex may have used "s-pause-wool" as a way of initiating the vocalization such that two known utterances provided the overall structure and the pause was a place filler, somewhat like that occasionally used by young children, until he could learn how to insert the /p/ and adapt the vowel. Specifically, Peters (2001) suggests that children use certain sounds as fillers (a "holding tank") to preserve the number of syllables or the prosodic rhythm of the target vocalization until the standard form is learned (note also Leonard, 2001). Even though Alex used a pause, rather than another phoneme, his behaviour suggests (but, of course, does not prove), that he had an awareness of the need for something additional and somewhat different to complete the vocalization. Simply omitting or closing the gapand responding on the basis of sound similaritlywould have produced /swUl/ ("swull"), not /swul/ ("swoool").

One might, of course, question whether this single instance of combinatory behaviour qualifies as evidence of phonological awareness in a nonhuman. Arguably, Alex may have applied a phonological rule for combining utterances without truly understanding the basis for the rule. Such an argument could, indeed, be made for the labels that he produced in the absence of referents (i.e., in apparent sound play, Pepperberg, 1990), but the specificity and consistent use of the "s-wool" combination argues against such an alternative explanation, as well as against that of "babble-luck" (a fortuitously correct but accidental combination, Thorndike, 1943). Here, Alex had to have discriminated and extracted the appropriate speech sounds of the target label "spool", generalized these to the closest related items in his repertoire, fit the existent sounds together—apparently including a pause to maintain spacing for an absent sound—in a particular serial order so as to add to his lexicon, and additionally had to link the novel phonology referentially with a specific item. The uniqueness of the behaviour (i.e., the single reported instance) is also not remarkable: For the past several years, Alex's training has concentrated on concept, rather than label, acquisition, and he has had few opportunities to engage in novel label learning. Interestingly, in recent training on the label "seven", his first attempts have been "ss....nnn", which also suggest some phonological awareness.

It is now possible to return to the initial argument, that Alex's vocal segmentation provides evidence for true imitation, rather than mimicry. 'Mere' mimicry can be defined as the purposeless duplication of an act (for a bird, it would be rote reproduction of human speech without referential content), behaviour that lacks cognitive complexity and intentionality. But if an act is performed because the imitator understands its purpose-to reach a goal, be it an object or intentional communication, otherwise impossible to obtain-then the act is intentional, complex, likely indicates cognitive processing, and provides evidence for true imitation. As stated in the introduction, Alex's data demonstrate that he understands that his existent labels are comprised of individual units that can intentionally be recombined in novel ways to create referential, novel vocalizations.

Whether such data can be used to argue for a parrot's understanding of a phonetic 'grammar' (e.g., Fitch & Hauser, 2004) is unclear: Although the data suggest that Alex can generate novel meaningful labels from a finite set of elements, the rule system he demonstrated was relatively limited. Nevertheless, the data presented here add another intriguing parallel between Alex's and young children's early label acquisition (Pepperberg, 1999). For children, manipulation of individual parts of a word implies the existence of internal representations of words as divisible units, and normal children proceed in a fairly standard manner from babbling to full language. Alex does not have and will not likely reach the level of any young child-that is, in terms of grammar go beyond use of simple sentence frames such as "I want X" and "I wanna go Y", where X and Y are appropriate object or location labels-but any strides that a bird makes toward language-like ability-such as, for example, comprehending recursive conjunctive sentences⁵ or

demonstrating the kind of vocal segmentation described here—helps us understand the similarities and differences between humans and non-humans.

These findings may also be of use in two ways for computer scientists who are trying to develop speech skills in their atavars and robots. First, as Patterson and Pepperberg (1994, 1998) have demonstrated, Alex produces most of his utterances with little variation in his first formant, and most of the variation in his second (and possibly third). Thus speech modelling may be simplified if based on avian, rather than human, productions (note Schwartz, Boë, & Bessière, 2001). Second, Alex's pattern of acquisition might suggest how approximations and iterations can be used for the construction of novel speech sounds from existent programmed vocalizations, providing additional assistance to existent algorithms (see Higashimoto and Sawada, 2002; Nishikawa et al., 2002; Yoshikawa et al., 2003a, b).

In sum, I suggest that Alex's training on both referential labeling and sound-letter association has engendered levels of phonological awareness, vocal segmentation, and imitation that need to be addressed when arguing for (a) human uniqueness and (b) the exclusive use of humans as the bases for computational models.

Acknowledgements

I thank Drs. Diana Reiss (Wildlife Conservation Society, NYC) and Donald Kroodsma (UMass-Amherst) for creating the sonagrams on *Raven*, Dr. Ofer Tchernichovski (CCNY) for the sonagram using *Sound Analysis Pro*, and Dr. Diane Patterson for alerting me to the issue of 'filler' phonemes.

References

- Michael A. Arbib. The mirror system, imitation, and the evolution of language. In Kerstin Dautenhahn and Chrystopher L. Nehaniv (Eds.), *Imitation in Animals and Artifacts* (pp. 229-279). Cambridge, MA: MIT Press, 2002.
- Michael A. Arbib. From monkey-like action recognition to human language: An evolutionary framework for neurolinguistics. *Behavior and Brain Sciences*, X, xx-yy (in press).
- Gabriël J.L. Beckers, Brian S. Nelson, and Roderick A. Suthers. Vocal-tract filtering by lingual articulation in a parrot. *Current Biology*, 14(17):1592-1597, 2004.
- Brian Byrne and Alvin M. Liberman. Meaninglessness, productivity and reading: Some observa-

⁵ For example, given various trays each holding seven objects of several colors, shapes, and materials, Alex can respond to queries of "What object/material is color-A and shape-B?" versus "What shape is color-A and object/material-C?" versus "What color is shape-B and object/material-C?" (Pepperberg, 1992).

tions about the relation between the alphabet and speech. In Jane Oakhill and Roger Beard (Eds.), *Reading Development and the Teaching of Reading* (pp. 157-174). Oxford: Blackwell, 1999.

- Julia M.Carroll, Margaret J. Snowling, Charles Hulme, and Jim Stevenson. The development of phonological awareness in preschool children. *Developmental Psychology*, 39:913-923, 2003
- Benedicte de Boysson-Bardies. *How Language Comes To Children*. Cambridge, MA: MIT Press, 1999
- W. Tecumseh Fitch and Marc D. Hauser. Computational constraints on syntactic processing in a nonhuman primate. *Science*, 303:377-380, 2004.
- Roger Fouts and Randall Rigby. Man-chimpanzee communication. In Thomas Sebeok (Ed.), *How Animals Communicate* (pp. 1034-1054). Bloomington, IN: Indiana University Press, 1977.
- Susan E. Gathercole and Alan D. Baddeley. The role of phonological memory in vocabulary acquisition: A study of young children learning new names. *British Journal of Psychology*, 81:439-454, 1990.
- Patricia Greenfield. Language, tools and brain: The ontogeny and phylogeny of hierarchically organized sequential behavior. *Behavior & Brain Sciences*, 14:531-595, 1991.
- T. Higashimoto and H. Sawada. Speech production by a mechanical model: construction of a vocal tract and its control by neural network. In *Proceedings of the 2002 IEEE International Conference on Robotics & Automation*, 3858–3863, 2002.
- Peter Ladefoged. *A Course in Phonetics*. San Diego: Harcourt Brace Jovanovitch, 1982.
- Laurence B. Leonard. Fillers across languages and language abilities. *Journal of Child Language*, 28:257-261, 2001.
- Virginia A. Mann and Judith G. Foy. Phonological awareness, speech development, and letter knowledge in preschool children. *Annals of Dyslexia*, 53:149-173, 2003.
- Marc Marschark, Victoria S. Everhart, Jan Martin, and Sue A. West. Identifying linguistic creativity in deaf and hearing children. *Metaphor & Symbolic Activity*, 2:281-306, 1987.
- Karter B. Neal. The development of a vocalization in an African Grey parrot (*Psittacus erithacus*). Unpublished senior thesis, University of Arizona,

1996.

- K. Nishikawa, A. Imai, T. Ogawara, H. Takanobu, T. Mochida, and A. Takanishi. Speech planning of an anthropomorphic talking robot for consonant sounds production. *Proceedings of the* 20002 IEEE International Conference on Robotics and Automation, 1830-1835, 2002.
- Dianne K. Patterson and Irene M. Pepperberg. A comparative study of human and parrot phonation: acoustic and articulatory correlates of vowels. *Journal of the Acoustical Society of America*, 96:634-648, 1994.
- Dianne K. Patterson and Irene M. Pepperberg. Acoustic and articulatory correlates of stop consonants in a parrot and a human subject. *Journal of the Acoustical Society of America*, 103:2197-2215, 1998.
- Sharon Peperkamp. Phonological acquisition: Recent attainments and new challenges. *Language* & *Speech*, 46:87-113, 2003.
- Irene M. Pepperberg. Functional vocalizations by an African Grey Parrot (*Psittacus erithacus*). *Zeitschrift für Tierpsychologie*, 55:139-160, 1981.
- Irene M. Pepperberg. Cognition in the African Grey parrot: preliminary evidence for auditory/vocal comprehension of the class concept. *Animal Learning & Behavior*, 11:179-185, 1983.
- Irene M. Pepperberg. Referential mapping: A technique for attaching functional significance to the innovative utterances of an African Grey parrot (*Psittacus erithacus*). Applied Psycholinguistics, 11:23-44, 1990.
- Irene M. Pepperberg. Proficient performance of a conjunctive, recursive task by an African Grey parrot (*Psittacus erithacus*). Journal of Comparative Psychology, 106:295-305, 1992.
- Irene M. Pepperberg. *The Alex Studies*, Cambridge, MA: Harvard University Press, 1999.
- Irene M. Pepperberg. Allospecific referential speech acquisition in Grey parrots (*Psittacus erithacus*): evidence for multiple levels of avian vocal imitation. In Kerstin Dautenhahn and Chrystopher L. Nehaniv (Eds.), *Imitation in Animals and Artifacts* (pp. 109-131). Cambridge, MA: MIT Press, 2002.
- Irene M. Pepperberg, Katherine J. Brese, and Barbara J. Harris. Solitary sound play during acquisition of English vocalizations by an African Grey parrot (*Psittacus erithacus*): Possible parallels

with children's monologue speech. *Applied Psycholinguistics*, 12:151-178, 1991.

- Irene M. Pepperberg and Steven R. Wilkes. Lack of referential vocal learning from LCD video by Grey Parrots (*Psittacus erithacus*). Interaction Studies, 5:75-97, 2004.
- Ann M. Peters. Filler syllables: What is their status in emerging grammar? *Journal of Child Language*, 28:229-242, 2001.
- Jean Piaget. Play, Dreams and Imitation in Childhood. New York: Norton, 1962
- Sue Savage-Rumbaugh, Jeannine Murphy, Rose A. Sevcik, Karen E. Brakke, Shelly L. Williams, and Duane M. Rumbaugh. Language comprehension in ape and child. *Monographs of the Society for Research in Child Development*, 233:1-258, 1993.
- Jean-Luc Schwarz, Louis-Jean Boë, and Pierre Bessière. Speech robotics, a computational framework for the study and modelling of speech acquisition and evolution. <u>www.icp.inpg.fr/OHLL/</u> lesPubliRapports/article_Steeles_2001.pdf
- Holly Storkel. Learning new words II: Phonotactic probability in verb learning. *Journal of Speech, Language, & Hearing Research*, 44:1312-1323, 2001.
- Holly Storkel. Restructuring of similarity neighbourhoods in the developing mental lexicon. *Journal of Child Language*, 29:251-274, 2002.
- Michael Studdert-Kennedy. Mirror neurons, vocal imitation, and the evolution of particulate speech. In Maxim I. Stamenov and Vittorio Gallese (Eds.), *Mirror neurons and the evolution of brain and language* (pp. 207-227). Amsterdam, Netherlands: John Benjamins, 2002.
- Edward L. Thorndike. *Man and His Works*. Cambridge, MA: Harvard University Press, 1943.
- Dietmar Todt. Social learning of vocal patterns and modes of their applications in Grey parrots. *Zeitschrift für Tierpsychologie*, 39:178-188, 1975.
- Michael Tomasello. *Constructing a Language*. Cambridge, MA: Harvard University Press, 2003.
- Rebecca Treiman. Errors in short-term memory for speech: A developmental study. *Journal of Experimental Psychology: Learning, Memory, & Cognition,* 21:1197-1208, 1995.
- Marilyn M. Vihman. Phonological Development:

The Origins of Language in the Child. Malden, MA: Blackwell, 1996.

- Y. Yoshikawa, M. Asada, K. Hosoda, and J. Koga. A constructivist approach to infants' vowel acquisition through mother–infant interaction. *Connection Science*, 15(4):245-258, 2003a.
- Y. Yoshikawa, J. Koga, M. Asada, and K. Hosoda. Primary vowel imitation between agents with different articulation parameters by parrot-like teaching. *Proceedings of the 2003 International Conference on Intelligent Robots and Systems*, 2003b.⁶

⁶ Note that 'parrot-like teaching' for these authors is entirely different from our M/R technique, and does not at all involve parrots.

Human and Robotic Action Observation Elicit Automatic Imitation

Clare Press Dept of Psychology University College London 26 Bedford Way, London. WC1H OAP. c.press@ucl.ac.uk Geoffrey Bird Institute of Cognitive Neuroscience University College London 17 Queen Square, London. WC1N 3AR. g.bird@ucl.ac.uk Rüdiger Flach Dept of Psychology University College London 26 Bedford Way, London. WC1H OAP. r.flach@ucl.ac.uk

Cecilia Heyes Dept of Psychology University College London 26 Bedford Way, London. WC1H OAP. c.heyes@ucl.ac.uk

Abstract

Recent behavioural and neuroimaging studies found that observation of biological action, but not of robotic action, elicits imitation and activates the 'mirror neuron system' in the premotor cortex (Kilner, Paulignan, and Blakemore, 2003; Castiello, Lusher, Mari, Edwards, and Humphreys, 2002; Meltzoff, 1995; Tai, Scherfler, Brooks, Sawamoto, and Castiello, 2004). This implies that the actions of other people and of mechanical devices are processed in categorically different ways. However, if the mirror system develops through learning (Heyes, 2001), generalisation should result in some activation when observing robotic action. We asked subjects to perform a prespecified action on presentation of a human hand or a robotic device in the final posture of the same action or the opposite action (Heyes, Bird, Johnson, and Haggard, 2004; Stürmer, Ascherschleben, and Prinz, 2000). Both the human and the robotic stimuli elicited automatic imitation: the prespecified action was initiated faster when it was cued by the same action than when it was cued by the opposite action. However, even when the human and robotic stimuli were of comparable size, colour and brightness, the human hand had a stronger effect on performance. These results point to the shape of the human hand as a source of features distinguishing human from robotic action. They also suggest, as one would expect if the mirror neuron system develops through learning, that to varying degrees both human and robotic action can be 'simulated' by the premotor cortex (Gallese and Goldman, 1998).

1 Introduction

A number of studies have shown that action perception can influence action production. For example, in a reaction time (RT) paradigm Brass, Bekkering, and Prinz (2001) asked participants to execute a prespecified action (moving their index finger up or down) as soon as they saw another person's index finger begin to move up or down. An 'automatic imitation' effect was obtained such that upward movements were executed faster in response to upward movements than to downward movements, and vice versa for the execution of downward movements. Thus, even when the executed movement is simple, and has been prepared in advance, action perception can influence action production.

Interactions between action perception and production are thought to be mediated by structures

in the premotor and parietal cortices. The most widely-cited evidence in support of this view comes from electrophysiological studies of 'mirror neurons' in the premotor cortex (e.g. Gallese, Fadiga, Fogassi, and Rizzolatti, 1996; Rizzolatti, Fadiga, Gallese, and Fogassi, 1996) and inferior parietal lobule (Fogassi, Gallese, Fadiga, and Rizzolatti, 1998; Gallese, Fogassi, Fadiga, and Rizzolatti, 2002) of the macaque monkey. These cells fire both when the monkey performs an action and when it watches another monkey perform the Functional magnetic resonance same action. imaging (fMRI) has indicated areas with similar properties in the human premotor cortex and parietal lobes (e.g. Iacoboni et al., 1999; Buccino et al., 2001).

Evidence is accumulating that activation in 'mirror neuron' circuits, and behavioural phenomena like automatic imitation, occur when the stimulus action is biological, but not when it is robotic. For example, Castiello (2002) required participants to reach out and grasp an object after observing a human or a robot hand reaching for and grasping a similar object. When the stimulus hand was human, but not when it was robotic, the size of the object grasped by the stimulus hand influenced aspects of participants' action such as maximum grip aperture and time to reach peak velocity. Furthermore, using positron emission tomography (PET), Tai et al. (2004) found significant activation in the left premotor cortex when participants observed grasping actions performed by a human model, but not when the same actions were performed by a robotic model.

These results can be interpreted in at least two ways. First, they may indicate that the actions of other people and of mechanical devices are processed in categorically different ways. If this hypothesis is correct, one would not expect observation of robotic action to give rise to automatic imitation even when the robotic stimuli are as perceptually salient as human action stimuli. Second, results such as those of Castiello (2002) may indicate that, whereas both human and robotic movement stimuli give rise to motor activation, human movement stimuli typically receive more motor processing than robotic movement stimuli. According to this hypothesis, the difference between the two stimulus types is quantitative rather than qualitative. If it is correct, equally salient human and robotic movement stimuli should both elicit automatic imitation, and the human stimuli should have a stronger effect on performance.

The Associative Sequence Learning (ASL) model of imitation supports the second, quantitative

hypothesis over the first, qualitative hypothesis. It suggests that the capacity to imitate is learned in a Hebbian fashion; through experience which causes concurrent activation of visual and motor representations of the same action. Hand movements are perceptually transparent (Heyes and Ray, 2000), and therefore self-observation is likely to provide much of the experience contributing to hand movement imitation. However, insofar as robotic hands are visually similar to human hands, one would expect them to benefit from generalization of the 'training' received during selfobservation.

The present study aimed to determine whether human and robotic stimulus hands, matched on a range of physical dimensions, would both elicit automatic imitation, but to different degrees.

2 **Procedure and Results**

We presented participants with four hand types; human naturalistic, human schematic, robotic naturalistic and robotic schematic (see Figure 1). In order to control kinematic variables we used static rather than moving stimuli. In addition, schematic stimuli were matched for size, luminance and colour, and only differed in shape. Naturalistic stimuli were matched as far as possible on these dimensions.



Figure 1. Experimental stimuli: A human naturalistic, B human schematic, C robotic naturalistic, D robotic schematic. Images depict hand in a neutral posture (warning stimulus).

Within a block, participants made the same response (opening or closing) in every trial. They were instructed to execute this movement as soon as a hand in a neutral posture on the computer screen (the warning stimulus) was replaced by an opened or closed hand (the imperative stimulus). On compatible trials, the posture of the stimulus hand matched the end-point of the participant's response, and on incompatible trials, the stimulus hand was presented in the alternative posture. To control for spatial compatibility effects, the orientation of the participant's responding hand was orthogonal to that of the stimulus hand. Reaction times were recorded using EMG from the first dorsal interosseus muscle of the right hand.

The results showed that responding was faster on trials where stimulus movement type was compatible with response movement type, supporting previous findings of automatic imitation (e.g. Brass, Bekkering, and Prinz, 2001). There was a larger effect of automatic imitation with human stimuli than with robotic stimuli. This finding supported previous research suggesting human stimuli activate mirror systems to a greater extent than robotic stimuli (e.g. Tai et al., 2004). As some of our human and robotic stimuli were matched on all physical dimensions other than shape, the shape of a hand seems to be sufficient to modulate automatic imitation.

However, we still observed some automatic imitation with robotic hand stimuli. This implies both human and robotic action can be 'simulated' by the premotor cortex to varying degrees (Gallese et al., 1998), and is consistent with what one would expect if the mirror neuron system develops through learning (Heyes, 2001).

References

- M. Brass, H. Bekkering, & W. Prinz. Movement observation affects movement execution in a simple response task. *Acta Psychologica*, 106: 3-22, 2001.
- G. Buccino, F. Binkofski, G.R. Fink, L. Fadiga, L. Fogassi, V. Gallese et al. Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI Study. *European Journal of Neuroscience*, 13: 400-404, 2001.
- U. Castiello, D. Lusher, M. Mari, M. Edwards, & G.W. Humphreys. Observing a human and a robotic hand grasping an object: differential motor priming effects. In W.Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action, Attention and performance XIX.* New York: Oxford University Press, 2002.

- L. Fogassi, V. Gallese, L. Fadiga & G. Rizzolatti. Neurons responding to the sight of goal-directed hand/arm actions in the parietal area PF (7b) of the macaque monkey. *Society of Neuroscience Abstracts*, 24: 154, 1998.
- V. Gallese, L. Fadiga, L. Fogassi, & G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119: 593-609, 1996.
- V. Gallese, L. Fogassi, L. Fadiga, & G. Rizzolatti. Action representation and the inferior parietal lobule. In W.Prinz & B. Hommel (Eds.), *Attention and performance XIX. Common mechanisms in perception and action*. Oxford: Oxford University Press, 2002.
- V. Gallese & A. Goldman. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, 2: 493-501, 1998.
- C.M. Heyes. Causes and consequences of imitation. *Trends in Cognitive Sciences*, 5: 253-261, 2001.
- C.M. Heyes, G. Bird, H. Johnson, & P. Haggard. Experience modulates automatic imitation. *Cognitive Brain Research*, in press.
- C.M. Heyes & E. Ray. What is the significance of imitation in animals? *Advances in the study of* behaviour, 29: 215-245, 2000.
- M. Iacoboni, R.P. Woods, M. Brass, H. Bekkering, J.C. Mazziota & G. Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286: 2526-2528, 1999.
- J.M. Kilner, Y. Paulignan, & S.J. Blakemore. An interference effect of observed biological movement on action. *Current Biology*, 13: 522-525, 2003.
- A.N. Meltzoff. Understanding the intentions of others: Re-enactment of intended acts by 18month-old children. *Developmental Psychology*, 31: 838-850, 1995.
- G. Rizzolatti, L. Fadiga, V. Gallese, & L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3: 131-141, 1996.
- B. Stürmer, G. Ascherschleben, & W. Prinz. Correspondence effects with manual gestures and postures: A study of imitation. *Journal of Experimental Psychology*, 26: 1746-1759, 2000.

Y.F. Tai, C. Scherfler, D.J. Brooks, N. Sawamoto,& U. Castiello. The human premotor cortex is

'mirror' only for biological actions. *Current Biology*, 14: 117-120, 2004.

An Examination of the Static to Dynamic Imitation Spectrum

Joe Saunders, Chrystopher L. Nehaniv and Kerstin Dautenhahn

Adaptive Systems Research Group School of Computer Science

University of Hertfordshire

College Lane, Hatfield, Herts AL10 9AB, United Kingdom {J.2.Saunders|C.L.Nehaniv|K.Dautenhahn}@herts.ac.uk

Abstract

We consider the issues that arise from an examination of the continuum between two social learning paradigms that are widely used in robotics research: (i) *following* or *matched-dependent behaviour* and (ii) *static observational learning*. We use physical robots with minimal sensory capabilities and exploit controllers using neural network based methods for agent-centred perception of model angle and distance. The robot is first trained to perceive the dynamic movement of a robot model carrying a light source, then the robot learns by observing the model demonstrate a behaviour and finally it attempts to re-enact the learnt behaviour. Our results indicate that a dynamic observation using rotation performs significantly better than static observation. However given the embodiment of the robot a dynamic strategy using both rotational and translational movement becomes more problematic. We give reasons for this, discuss lessons learned for combining these types of social learning and make suggestions for requirements for imitator robots using dynamic observation.

1 Introduction

In this paper we build on our previous research (Saunders et al., 2004) in considering the issues that arise from an examination of the continuum between two social learning paradigms that are widely used in robotics research: (i) following or matcheddependent behaviour and (ii) static observational *learning*. Our motivation in examining these issues is the belief that an understanding of the mechanisms underlying social learning should be considered as a prerequisite for building adaptive and intelligent robots. We believe that social learning leads to an acceleration of the acquisition of intelligent behaviour (Zentall, 2001; Galef and Heyes, 1996; Dautenhahn and Nehaniv, 2002) with the promise of easier robot task acquisition, increased behavioural complexity and ultimately some form of cultural transmission (Alissandrakis et al., 2003). In this respect we focus on the mechanisms supporting Imitation¹ with experiments with physical robots in an attempt to simplify and focus on key aspects of imitative processes. The background of this paper is an ongoing investigation of social learning and the interaction between both human/robot and robot/robot pairs to understand the social dimension of imitative behaviour. The perspective of both the imitator and the imitatee and the problems of perception and action encountered by both are considered. Our starting point is the different imitator perspectives which are widely applied in paradigms used in robotics imitation research, namely *following* behaviour (Hayes and Demiris, 1994; Billard and Dautenhahn, 1997; Dautenhahn, 1994) and *static observation* behaviour (Kuniyoshi et al., 1994; Gaussier et al., 1997; Bentivegna and Atkeson, 2002; Schaal, 1997; Matarić et al., 1998; Alissandrakis et al., 2003).

From a psychological/ethological viewpoint *following* is more rightly considered as *matched-dependent behaviour* (Zentall, 2001). The imitator observes and immediately matches the behaviour of the model as it is being performed, staying close to the model. For example rats can be trained to follow a lead rat through a maze which they then learn to navigate (Miller and Dollard, 1941). The rats may have no idea of intentionality of the lead rat and can be trained to follow other salient (including non-animal) stimuli, this behaviour is sometimes called *discriminated following*.

¹We take Thorndike's 1898 classical definition of imitation (Thorndike, 1898) as "learning how to do something by seeing it done" but extended to include non-biological agents (Mitchell, 1987).

Likewise, *static observation* by the imitator who stays at a fixed location is related to the ethological/psychological notion of *observational learning*. Here the behaviour of the demonstrator is copied after it is observed carrying it out. Typically the demonstrator and imitator operate within a shared context but at a distance from one another. For example Norway Rats apparently develop food preferences by smelling the breath of a conspecific (Galef and Heyes, 1996), without reference as to whether the demonstrating rat becomes ill or dies. These examples hint at some interesting but not widely researched features of imitative behaviour in the relationship between static observation of, and active participation in, an event to be imitated.

In our previous research (Saunders et al., 2004) we considered the extremes of a purely reactive following behaviour and contrasted that against a static observational behaviour using some simple experiments with Khepera miniature robots. Two controllers were designed to allow either a reactive following behaviour or a static observation behaviour. Each robot either followed or statically observed another robot making various geometric shapes over varying terrain. In both cases the robot could learn the observed behaviour and attempt to re-enact it. The model was perceptable by the imitator due to placement a small light bulb on top of the model. No explicit communication was permitted between the model and imitator; in fact the sensory information was basically the perceived brightness of the moving light bulb. The research results from these experiments identified trade-offs that are summarised in the spectrum table shown in figure 1.

The results indicated that there was a clear tradeoff between positional accuracy obtained from static observation and the advantages of direct perceptionaction coupling available from following. This lack of precision during following we called impersistence to reflect the fact that the robot is always reacting to the latest sensor reading and not persisting to meet the goal signalled by the previous reading. We believed that the accuracy available from static observation was unsurprising, given that static observation allows the design of the robot controller to concentrate exclusively on angle and distance perception and apply more complex and engineered methods to this task. We believed that similar complexity in observational systems were also engineered into most other social learning robotics experiments.

The relative simplicity of the following paradigm also hid some key advantages, in that the robot was





Figure 1: The table summarises the key aspects revealed by the previous research experiments (Saunders et al., 2004) with extremes of each aspect shown. Comparative costs are shown in the boxes. The current research considers mixed approaches which might allow the balance of these costs and benefits.

able to directly map its perceptions against its motor actions. It was thus able to learn much about the environment directly and relatively cheaply. However to achieve positional accuracy, more complex observational algorithms were required, but observation alone was insufficient to completely assess the physical complexities of the environment. We said that there may be an argument for suggesting that observation could be most effective after a following episode, i.e. observation could fine-tune already stored movement patterns. Similarly there may be an appropriate time to 'see' (observe from a distance) as opposed to 'feel' (follow, experiencing the same context) in social learning. A mixed approach may be valuable, this approach corresponding to intermediate positions or switching in the spectrum table shown. One could imagine for example cases where the observation is less static e.g. several follow-observe-follow cycles, or where a series of static observations are made prior to each episode of following behaviour.

Dynamic Observation. In this paper we consider in more detail some of the effects of allowing a more dynamic observational approach. We study the quality of the imitation attempt from the imitator's perspective in two experiments using either an 'observe and rotate' or an 'observe and move' strategy to match the movement patterns of the model. These successively augment static observation, respectively, by adding orienting rotational changes to allow the imitator to track the model (observe and rotate) or by adding rotation and translation.

2 Experimental Overview

For our experiments we use a controller previously designed (Saunders et al., 2004) to investigate the

imitation of movements using static observation and extend it to provide a mechanism to investigate dynamic imitation. All experiments are carried out in real-time on physical robots (i.e. simulation is not used) on a desktop in a typical busy academic environment with light levels varying during the day.



Figure 2: The picture shows the experimental platform. The Khepera acting as a model has a small bulb placed on top of it. The imitator is shown tracking the model which is tracing out a triangle.

We examine the behaviour of the imitator when imitating various geometric shapes made by the model. We consider intermediate positions in the spectrum table in two experiments to examine the effects of a mixed observation and movement execution strategy. Both experiments involve dynamic observation which combines both observation and movement.

Observe and Rotate. The first experiment extends the static observational perspective by allowing the robot to alter its orientation so as the better exploit it sensory facilities. The embodiment of the Khepera robot is such that the majority of the light sensors are in front of the wheels, with two sensors at the back. The estimation of distance is therefore more accurate when the robot is able to employ all of its front facing sensors as it is receiving more information from the environment. To ensure that these sensors are in an optimal position we program the circular Khepera robot to rotate in place orienting toward the model.

The rotation is such that the imitator will attempt to directly face the model if the model's angle with respect to the imitator exceeds a given threshold. However, if the imitator has to rotate to achieve this then all subsequent observations must be converted back to the original reference frame in order to replay the imitation. To achieve this conversion, accuracy in measuring how far the robot has turned is critical to this process. We tested threshold angles of 0, 30, 60 and 90 degrees. In both this and the experiment described below the model was preprogrammed to make 4 geometric shapes. The first was a 10cm radius circle around the imitator, the second a 10cm circle 5cm in front of the imitator. The third and fourth a triangle and T-Shape 5cm in front of the imitator.

Observe and Move. The second experiment allows the robot to record a sequence of observations of the model and then attempt to use a given subset of these observations to imitate the model's movement sequence. Once the imitator has completed this part of the imitation it recommences observing.

In a two-dimensional parameterisation of the spectrum, different social learning mechanisms are given by varying both the number n of *observations* and the number m of *movements* made by the imitator. A single observation is an estimation by the imitator of the model's angle and distance from the imitator. A movement is the transformation and execution by the imitator of observations to motor-commands in order to achieve the same effect.

These mechanisms however present a series of challenges due to the fact that after each movement sequence the robot's memory of previous observations will be from a different perspective from the current observation set. This is because the imitator, after partially replaying the imitation (by transforming a subset of the observed vectors) will find that the remaining observations need to take account of the new observation position. Furthermore, the new observation position may not be optimal for accurate readings, therefore a rotation (as in experiment 1) will be necessary. To then replay the next part of the imitation the effect of the rotation must be reversed and subsequently a transformation of the observations reperformed.

3 Controller

The controller used in both experiments relies on computing the distance and angle from the imitator to the moving model and storing these observation points as a list of two element vectors. Prior to observing, the robot first learns how to measure angles and distance.

Learning to Measure Angles. The robot is first trained to accurately compute the angle of the light from the centre of the imitator. A number of methods were evaluated including using a light compass (Nehmzow, 1993), or computing the angle by using *vector summation* of the inputs to each of the light sensors (Arkin, 1998). However both of these methods were not accurate and suffered from incorrect readings especially when none of the robot's sensors were directly facing the light. A new method, which we call environmental sampling, was grounded in sensory experience and is to some extent nearer to a biological solution: the robot is allowed to learn about light angles simply by observing them. As the Khepera is a circular robot it rotates in a circle in the presence of the model. It detects when the circle is complete by polling its wheel encoders and stopping when the appropriate value has been exceeded. (During the turn it reads its light sensors every 200ms. A robot turning at 8mm/s would typically poll it sensors 65 times.) As the speed of the turn is constant the time interval between readings can thus be converted to an angle. Each of the sensor readings are then normalised. This has two effects, firstly that of making distant readings of angle equivalent to closer readings, and secondly allowing these values to be loaded directly as weights into a neural network (a counterpropagation network (Hecht-Nielson, 1988)). This is a fully connected feed-forward three layer network. The first layer takes the normalised input of the 8 light sensors.

the number of middle layer neurons is set to the number of times the robot was able to poll its sensors and the final layer used to output the conversion of these values to angles. Using this technique has a number of advantages. Firstly that the network can be built as the environment is observed, secondly there are no additional training steps i.e. there is no further training of the neural network, thirdly the size of the network is directly related to the internal rotation speed, sensor modality and sensor polling time of this particular robot and finally that the method is partially resilient to sensor failure. There are some biological observations which may show similar (though not equivalent) mechanisms in animals. For example young bees appear to record the image of their hive from many angles and positions around it: they fly in and out of the hive varying their circular flight path each time (Murphy, 2000).

Learning to Measure Distance. For distance measurements various mechanisms were also assessed. A first approach was to use *triangulation*, exploiting the fact that accurate angle measurement was now possible. The approach measured the light angle from the model, moved the imitator a fixed distance and then read the new angle. This allows the computation of the original distance using the two angles and the travelled distance. However this mechanism was unreliable for two reasons, firstly that, over small movement distances (which minimised errors in the odometry readings from the wheel encoders), the derived angle would be small and tiny errors in the angle measurement would result in an amplified error in the distance computation, secondly if the model was moving, the measurements/movement combination of the imitator could never be fast enough to resolve the position of the model accurately. An alternative method based on environmental sampling was used for the angle computation, the light sensors being summed as vectors as the robot turned. This exploited the fact that sensors directly facing the light would have a larger effect on the vector magnitude than those further away. The robot was trained by rotating at increasing 1cm distances from the light source. The vector magnitude was then held in a lookup table indexed by angle and distance. Using this method gave a reasonable distance accuracy to about 25 cm from the robot at an angle between approximately 30° to 150° in front of the robot. However, outside these parameters the distance accuracy was very poor.

Following these procedures the robot can compute both angle and distance without further training.

Observing Angles. After the learning phase is complete the network operates by feeding a normalised sensor vector to the input layer and receiving the angle from the output layer. The network is thus operating as a pattern matching mechanism. Automatic interpolation between observed values is achieved by setting the middle layer 'winning nodes' to a value greater than 1.

Observing Distance. During the observation phase the angle is computed, followed by magnitude of the vector summation², the two values providing the key to the lookup table to yield distance.

Altering the Angle of Observation. In both experiments the robot collects a set of angles/distances from itself to the model whilst the model is moving.

²Refer to (Saunders et al., 2004) for details.

The imitator cannot poll its sensors when it itself is moving. Thus in a fixed time period the number of possible observations when the imitator is not moving will be higher than when the imitator is moving. In the first experiment the imitator can either not move or rotate to face the model once a threshold angle has been exceeded (see figure 3). The lower the threshold angle the greater the rotational movement of the robot to face the imitator when the angle is exceeded. The higher the angle, the smaller the rotational movement, but the robot will move more often. In our previous research we had fixed the imitator position and allowed it to observe the moving model. The model was at all times in front of the imitator and therefore within range of angle/distance computation mechanism. We now allowed the model to be both in front of the imitator and at any angle around the imitator. By varying the rotation threshold we can then examine both the effect of rotational movement size and the effect of frequency of movement on the imitation attempt.

Time Averaging and Way Points. In both experiments the recorded observations are smoothed using a simple moving average. The smoothed trajectory is then thresholded to yield a set of way points. This procedure is necessary for two reasons. Firstly to eliminate the effect of noisy observations and secondly to avoid two observation points being too close to one another - this closeness causing large and potentially damaging changes in the robot's motor systems if replayed directly. The imitator uses the derived way points to then imitate the model's trajectories. In the second experiment this procedure is only applied when computing the required movement. Any unused observations (which result from the movement index being less than the observation index - see experiment 2 below), remain unmodified as these may be subject to geometric transformation following the actual movement of the imitator.

4 Experimental Results

In our experiments we compared imitation behaviour on four simple patterns. These were a triangle, a circle enclosing the imitator, a circle observed ahead of the imitator and the letter T. The triangle was chosen because of the sharp changes of direction at each vertex, the circle because of its continuous shape and the letter T because of the need to reverse direction and remap the shape. We emphasise that our goal was not to design robots that perfectly imitate geometric shapes but rather investigate relevant aspects of the imitation attempt using a more dynamic approach in observational learning.

4.1 Experiment 1 - Dynamic Observation with Rotation Only

Details of Set-up. In each case the imitator is placed at the centre of the experimental platform (shown as point 0,0 on the graphs in figure 4) facing forward (at 90° along the positive Y-axis). The model is pre-programmed to move according to the prescribed shape. A threshold rotation angle is then set and the imitation run commenced for a fixed period. The threshold supplies a range of values around the front of the robot. For example, setting a threshold of say 60° means that if the imitator perceives the model within a forward range of $60 - 120^{\circ}$ (see figure 3) no rotation will be applied. If however the model



Figure 3: Rotation Threshold. In this example the imitator will not move at values between 60° and 120°. Between 61° and 121° the imitator will rotate to face the model.

moves to, say, 50° the imitator will rotate so that the model is directly in front of it, and thus be, from the imitator's new perspective, at 90° . The higher the threshold angle (to the limit of 90°) the more often the imitator will move to match the model but it will rotate by a smaller amount. If the threshold is set to zero, then the imitator will only move when the model is outside the range $0 - 180^{\circ}$, however the robot will then rotate by at least one quarter of its circumference.

Results. Figure 4 shows the results from a test with the enclosing circle. The robot is placed facing forward along the positive Y-axis. After the run the



Figure 4: *Dynamic Observation with Rotation Only. Imitative Behaviours for an enclosing circle. The first diagram shows the result with no rotation, the final four graphs show rotation at* 0°, 30°, 60° *and* 90° *degrees thresholds. The continuous line shows the path of the model, dotted line the imitator observations, crosses the smoothed observation and large dots the way points which are replayed by the imitator.*

imitator robot attempts to re-enact the observed behaviour. The large dots on the graphs show the way points, these being the path that the imitator will take when replaying the imitation. The first graph shows the imitation when no rotation has been applied and thus where only static observation is taking place. As expected at angles outside the angle/distance range the imitation is poor. The second graph shows the first example of a dynamic observation with the imitator moving only when the model moves outside the range $0 - 180^{\circ}$. Two extreme observation points are shown reflecting the inability of the distance/angle sensor to correctly measure the distance. However, once the 180° or 0° angle is exceeded the robot turns and starts again to make reasonably accurate readings. On imitation replay the outlying readings are smoothed away. The situation is further improved at 30° when the sensory apparatus is always in range but



Figure 5: *Dynamic Observation with Rotation Only. Imitative Behaviours for T-Shape The two graphs show results with rotation at* 60° *and* 90° *thresholds.*

the number of moves small. However at 60° and 90° the situation is ambiguous. We would suggest that the imitation is slightly less accurate. This may be due to the increasing effect of odometry errors as the number of moves increases. This is especially true at 90° where there would be a small movement for every 1° change on the model's position.

Results for the forward circle and triangle (not shown) were less marked, however the robot was subject to less movement due to the constrained angles presented by both shapes. The T-Shape however is more interesting (see figure 5). The nature of the shape meant that at 0° , 30° and 60° the imitated trajectories were broadly similar, however at 90° the robot was affected again by odometry drift and a similar worsening of readings ensued.

Analysis. These effects show some of the advantages and disadvantages of a tracking mechanism described above. Observing whilst not moving (static observation) has the key advantages of being fast and thus able to make more observations in a given time period (given the sequential nature of the observe/move scenario presented here). There are no odometry concerns as the imitator is not moving and the energy required would be lower than for a moving imitator. The major disadvantage is of course that the model can move into imitator blind spots. The advantages of the tracking imitator (dynamic observation with rotation only) is that blind spots can now be seen, however this is offset by the disadvantages of increasing odometry errors as more movement is carried out, a higher energy cost, and more complex computation as reference frame adjustments are continuously required. However at a particular movement/rotation ratio, which for this robot appears to be around 30° , there appears to be a point where accuracy is optimised. This suggests an clear strategy - expend energy and computational costs by moving only when



Figure 6: Analysis of Observation and Movement Index. In the analysis the model describes a square pattern. Here we see the imitator using a observation:movement index of 2:1 and successfully matching it. Similar successful matching will always occur when the movement index is set to 1 regardless of the observation ratio.

not to do so would give incorrect results. Or more simply - keep still until movement is almost necessary (in our case when the model goes beyond the 30° threshold into 'peripheral' vision).

4.2 Experiment 2 - Dynamic Observations: Varying Observation and **Movement Cycles**

Theoretical Results and Detailed Set-up. This experiment explored how movement and observation might be intermixed. This was attempted by varying the number of look-ahead observations against an equal or

smaller number of moves. Thus the robot would first make an initial observation ³ and then subsequently observe for n cycles and then move, based on these observations, m times. This procedure iterated throughout the imitation attempt. Prior analysis of this method, using the imitator and model represented as points (see figures 6 and 7) suggested that accurate imitation may only be possible if the number of moves were set to 1. To simplify the analysis we assumed that the imitator and model moved at approximately the same constant speed.





Imitator observes Model at 1

Figure 7: Analysis of Observation and Movement Index. This example shows the imitator an index of 2:2. When the movement index is set above 1 the imitator always fails to match the pattern.

Additionally, due to the control system of the robot, observation and movement execution are not possible in the same time step. We then imagined three scenarios cyclically alternating n observations (o) of model transitions with m moves (x) by the imitator (note: it is not possible to imitate further than our observed sequence, and therefore n is always larger or equal to m). The first scenario was of n observations to 1 move e.g. 1:1 o-o x o x o x o $x \dots$, secondly a scenario where there are an equal number of observations and moves but where both are greater than 1, e.g. 2:2 o-o-o x-x o-o x-x o-o x-x ... and finally where n is greater than m and both are greater than 1 e.g. 3:2 o-o-o-o x-x o-o-o x-x o-o-o x-x Figure 7 shows an example of the failure to correctly match the movement pattern when the move index is set higher than 1. This occurs because the imitator has failed to observe one or more critical points in the model's move sequence. The effect is similar to the *impersistence* problem we noted when analysing 'following' behaviour (Saunders et al., 2004), however rather than failure to complete or persist in its goal, as was the case for following, here the problem is one of 'inattentiveness'. The imitator is blind to the moves of the model. This problem occurs at all values of n and m which are larger than 1. Figure 7 shows an example of this when n:m is set to 2:2.

Results. The robot was tested on a series of index

³For each move two observation vectors are required, therefore at the start of the run one additional observation is made.



Figure 8: Dynamic Observation: varying Observation and Movement cycles. The results show the inability of the imitator to correctly replay the model's path. In this case a triangle shape.

values on each geometric shape presented by the model. Figure 8 shows an example of the physical robot using a 5:1 n:m index on the triangle shape. The imitator fails to match the model. Similar failures occurred in all attempts with the physical robot on all shapes. This was initially surprising, however the difficulty became clear once the actual imitator movement was considered.

Analysis. The simplicity of the point analysis above hides some crucial implementation issues. For example the robot can only move in the direction of its fixed wheels (i.e. it cannot arbitrarily move sideways), therefore a rotation may be necessary to orient the robot to the correct movement vector signalled from the model. Also, the Khepera has a fixed placement of sensors around its circular wheelbase. In order to correctly 'focus' on the model the robot must be in the appropriate sensor range. Thus the rotation mechanism described in experiment 1 was employed. Therefore in addition to the move or moves calculated from the observations we may have up to 2 additional moves: one to focus the sensors on the model and the other to orient the imitator for its move. Whilst these moves are being carried out the problem of 'inattentiveness' is compounded. Two further issues were also apparent. Firstly, each move is accompanied by a small odometry error. The total error therefore increases as the number of moves increases. Secondly, the smoothing effect of time averaging has little or no effect when attempting to model a small number of moves. This means that unsmoothed noisy observations are replayed leading subsequently to a poor imitation attempt.

We believe that the failure of the imitation is due primarily to the constraints imposed by the embodiment of this particular robot and might be obviated if the sensory apparatus was independent of the actuator mechanism e.g. distance/angle sensors which rotated and were focusable independently of movements of the main robot body. Such a mechanism is in fact used by (Gaussier et al., 1997) in their experiments. In fact there are no known imitating animals whose observation sensors cannot focus at least somewhat independently of the orientation of their bodies.

5 Discussion

In this research we have started to examine some of the practical issues which face an imitator when trying to use a dynamic observational behaviour. Here the problems of perception, perspective and action must be considered. We have greatly simplified the problem domain by restricting the imitative actions to that of replaying geometric shapes and have used simple robots with fixed sensor embodiments and with limited perceptual capabilities. We have previously suggested that a 'following' behaviour, although limited in its imitative accuracy, has the major advantage of computational simplicity and the added value of direct interaction with the environment through proprioceptive polling of its actuators whilst moving. We do not suggest that this opportunity to 'feel' the environment is exclusive to a following strategy and accept that there are alternative and probably better ways to proprioceptively explore the environment. However this strategy has the straightforward merits described above. It is also true that both a follower and an static observer are necessarily out of phase with the model and for this reason it seems that the follower's sensory cues may not be more appropriate than an observer's, however work by (Billard and Dautenhahn, 1997) showed that the these cues are dependent on the distance between a follower and the model and within a critical distance the follower's sensory cues become very relevant. What we describe here is an initial attempt to provide a movement mechanism to an observer in order to combine the advantages of observational accuracy with the feedback obtained from actively exploring the environment. Clearly a simple and modular solution to this task would be to keep to the 'extreme' behaviours and simply apply each strategy in turn e.g. follow-statically observe-follow. One of the aims of this research has been to explore the challenges faced in combining these strategies whilst retaining the positive aspects of both. The experiments themselves are clearly limited as we are constrained both by the sensor embodiment of the robot and its internal control system, but we believe valuable lessons still emerge.

Suggestions for Imitators Dynamically Observing from a Fixed Location. Our first experiment showed that dynamic observation with rotation was successful in that it allowed the model to pass out of view of the imitator and be reacquired. It was superior to static observation alone in this respect and it appeared that the benefit of tracking accuracy could be balanced against the cost of rotation frequency and rotational movement based on a turn threshold. Thus to retain observational accuracy, rotational movement should be limited so that odometry errors are minimised in their effect on the geometric transforms required to replay the imitation. Thresholds near the periphery of vision balance these factors. In robotics the issue of errors from odometry drift is clearly not new, however the literature on robotic observational imitation seem rarely to cite it as being a problem for a moving imitator.

Suggestions for Imitators that Observe and Move.

Our second experiment showed that with this particular robot, dynamic observation with movement of the imitator was extremely difficult and failed to replicate with reasonable accuracy the model's path. Our theoretical analysis suggested that the 'inattentiveness' problem may be soluble for a dynamic observational imitator where the movement value is set to unity. This region in our spectrum corresponds with the methods of other research (Wit, 2000) where a single solution to this issue is considered. However the need to make additional movements over and above those required to track the model means that the movement value can never be unity for an embodiment where the sensor orientation is completely fixed for a fixed body orientation. Thus the imitation will be poor.

Possible Solutions. A solution to this might be *in-dependent sensing and actuator mechanisms*. We envisage that such a system would additionally employ independent computation facilities for both mechanisms to allow continuous and parallel calculation of model position. Thus appropriate movement vectors could be sent to the actuators reducing unnecessary movements and the associated additional odometry drift. The sequential nature of the move-sense cycle on our robot may mean that accurate dynamic observation is very difficult, however other control systems employing a *parallel cycle* may provide solutions. There may also be simpler alternatives, for example the model may *repeat the pattern* and the im-

itator might manage to fill the gaps caused by earlier inattentiveness, or the model might simply *wait for the imitator*.

Even in our own human experience it appears much harder to both partially replay an imitation and observe the model before the model has finished its actions. Animals in fact may have obviated this issue by evolving alternative mechanisms. In this respect the recent neurological evidence of 'mirror neurons' in primates and humans (Gallese et al., 1996) and their role in action perception may play a considerable role in static observational learning with the imitator experiencing perhaps as good a corrrelation to its own behavioural patterns whilst statically observing as when attemping to match movements directly.

Acknowledgements

Many thanks to the reviewers of the extended abstract of this paper for their useful and constructive comments.

The work described in this paper was partially conducted within the EU Integrated Project COGNIRON ("The Cognitive Companion") and funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020.

References

- A. Alissandrakis, C. L. Nehaniv, and K. Dautenhahn. Synchrony and perception in robotic imitation across embodiments. In *Proc. IEEE CIRA* '03, pages 923–930, 2003.
- Ronald C. Arkin. *Behavior-based robotics*. MIT Press, Cambridge, Mass., London, England, 1998.
- Darrin C. Bentivegna and Christopher G. Atkeson. Learning how to behave from observing others. In Proc. SAB'02-Workshop on Motor Control in Humans and Robots: on the interplay of real brains and artificial devices, Edinburgh, UK, August, 2002, 2002.
- A. Billard and K. Dautenhahn. Grounding communication in situated, social robots. In *Proc. TIMR*, 1997.
- K. Dautenhahn. Trying to imitate a step towards releasing robots from social isolation. In P. Gaussier and J.-D. Nicoud, editors, *Proc. From Perception to Action Conference, Lausanne, Switzerland*,

pages 290–301. IEEE Computer Society Press, 1994.

- K. Dautenhahn and C. L. Nehaniv, editors. *Imitation in Animals and Artifacts*. MIT Press, 2002.
- B. G. Galef and C. M. Heyes, editors. *Social Learning in Animals: The Roots of Culture*. Academic Press, 1996.
- V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119:593–609, 1996.
- P. Gaussier, S. Moga, J. P. Banquet, and M. Quoy. From perception-action loops to imitation processes: A bottom-up approach of learning by imitation. In *Socially Intelligent Agents*, pages 49–54. AAAI Press, Technical report FS-97-02, 1997.
- Gillian Hayes and John Demiris. A robot controller using learning by imitation. In *Proc. International Symposium on Intelligent Robotic Systems, Grenoble*, pages 198–204, 1994.
- Robert Hecht-Nielson. Applications of counterpropagation networks. *Neural Networks*, 1:131–139, 1988.
- Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: Extracting reusable task knowledge from visual observations of human performance. *IEEE Trans. Robot. Automat.*, 10:799–822, November 1994.
- M. J. Matarić, V. Zordan, and Z. Mason. Movement control methods for complex, dynamically simulated agents: Adonis dances the macarena. In *Proc. Autonomous Agents '98*, pages 317–324. ACM Press, 1998.
- N. E. Miller and J. Dollard. Social Learning and Imitation. Yale University Press, 1941.
- R. W. Mitchell. A comparative-developmental approach to understanding imitation. In P. P. G. Bateson and P. H. Klopfer, editors, *Perspectives in Ethology 7: Alternatives*, pages 183–215. Plenum Press, 1987.
- Robin R. Murphy. *Introduction to AI Robotics*. The MIT Press, Cambridge, 2000. p.76.
- Ulrich Nehmzow. Animal and robot navigation. In Luc Steels, editor, *The Biology and Technology of Intelligent Autonomous Agents*. Springer Verlag, 1993.

- J. Saunders, C. L. Nehaniv, and K. Dautenhahn. An experimental comparison of imitation paradigms used in social robotics. In *Proc. IEEE Robot and Human Interactive Communication (ROMAN '04)*, pages 691–696. IEEE, September 2004.
- Stefan Schaal. Learning from demonstration. In Michael C. Mozer, Michael I. Jordan, and Thomas Petsche, editors, Advances in Neural Information Processing Systems volume 9, pages 1040–1046. MIT Press, 1997.
- E. L. Thorndike. Animal intelligence: An experimental study of the associative process in animals. *Psychol. Rev. Monogr.*, 2:551–553, 1898.
- J.S. Wit. Vector Pursuit Path Tracking for Autonomous Vehicles. PhD thesis, University of Florida, 2000.
- T. R. Zentall. Imitation in animals: Evidence, function and mechanisms. *Cybernetics and Systems*, 32 (1-2):53–96, 2001.

View Sensitive Cells as a Neural Basis for the Representation of Others in a Self-Centered Frame of Reference

Eric L. Sauser*

*Ecole Polytechnique Fédérale de Lausanne Swiss Federal Institute of Technology, Lausanne Autonomous Systems Laboratory CH-1015 Lausanne eric.sauser@epfl.ch Aude G. Billard[†]

[†]Ecole Polytechnique Fédérale de Lausanne Swiss Federal Institute of Technology, Lausanne Autonomous Systems Laboratory CH-1015 Lausanne aude.billard@epfl.ch

Abstract

This work follows from a research project, in which we investigate the underlying mechanisms of human imitation and develop a neural model of its core neural circuits. The present paper presents a model of a neural mechanism by which an imitator agent can map movements of the end effector performed by other agents onto its own frame of reference. The model mechanism is validated in simulation and in a humanoid robot to perform a simple task, in which the robot imitates movements performed a human demonstrator.

1 Introduction

Imitation is the ability to recognize, learn and reproduce others' actions. This powerful cognitive mechanism is fundamental for the transmission of knowledge and skills within the same species and across species. It is also at the basis of primates' social communication. One can distinguish the numerous forms of imitation behavior displayed in nature according to levels of complexity. In its simplest form, imitation can be reduced to a sensorimotor mapping that transforms sensory information, usually visual, into corresponding motor commands. Such basic imitation would be displayed as a form of "emulation" or of "social facilitation" (Heyes, 2001). Moreover, Bekkering et al. (2000) have shown that imitation is generally goal-directed, that is children and adults tend more to reproduce the goal of an demonstrative act, rather that the exact sequence of movements leading to it. Indeed, compared to "mimicry", this mechanism doesn't require any body correspondence. What to imitate preponderates on how to imitate. Then, in its most complex form, imitation leads to or requires more complex cognitive capabilities, such as the recognition of conspecifics and the attribution of others' intentions or states of mind (Billard, 2002). It is often referred to as "true imitation" i.e. the ability to reproduce and learn new motor skills which are not part of the imitator's current motor repertoire. In true imitation, the imitator must be capable to extract the purpose of a given sequence of movements, namely

to be capable of action understanding¹.

In this paper, we aim at exploring the mechanisms underlying mimicry. Despite not being directly involved in the most common imitation mechanism that is goal-directed imitation, it is important to note, as mentioned by Wholschläger et al. (2003) that

it seems that if the goal is clear (or absent), then the course of the movement plays a more central role in imitation. One might also say therefore, that the movement itself becomes the goal.

Thus, the question we will develop here is how one can map motions performed by others onto his/her own perspective, and more precisely while considering the end effector trajectory, i.e. the hand of the demonstrator. Indeed, as simple as it appears to be to Ethologists, mimicry remains complex in terms of the basic cognitive capabilities it requires, such as the capacity to perform arbitrary frames of references transformations and to generate coherent sensorimotor mappings. Such cognitive processes are fundamental and necessary for more complex forms of imitation. They remain, however, ill-understood. We argue that a better understanding of the brain mechanisms underlying mimicry is necessary to provide the stages for understanding and modeling the leap from simple to complex forms of imitation in animals.

While the behavioral processes of imitation have been the focus of studies in Ethology and develop-

¹An action is understood here as a goal-directed sequence of movements.

mental Psychology for centuries, evidence of correlated neural processes is much more recent. Studies of brain lesions resulting in degenerate imitative behaviors (apraxia or echopraxia) were the first to give some insight into the brain areas responsible for imitation, pointing to generic areas in the frontal and parietal cortices (Lhermite et al., 1986; Shimomura and Mori, 1998). The field revived a new life with the discovery in 1992 of the mirror neurons system, direct-mapping mechanism between visual and motor systems. For recall, mirror cells respond both when the animal performs and sees a goal-directed sequence of movements, hence, suggesting that a direct-mapping mechanism between visual and motor system exists for the purpose of linking conspecifics' or humans' action observation with self motor execution. The mirror neurons were first detected in the macaque monkey premotor cortex (PM), posterior parietal cortex (PPC) and superior temporal sulcus (STS) (Fogassi and Gallese, 2002; Rizzolatti et al., 1996). Later, brain imaging studies of the human brain highlighted numerous areas, such as STS, PM and Broca (Nishitani and Hari, 2000; Iacoboni et al., 1999; Decety et al., 2002). While the discovery of the mirror neuron system is certainly a key step toward a better understanding of the brain mechanisms underlying primates' ability to perform various forms of imitation, one has yet to clearly spell out the role of the mirror neuron system as part of the general neural processes for imitation.

Mirror neurons are relatively far from the brain areas receiving primary sensory information. They react, thus, to highly processed stimuli, represented in a goal-centered frame of reference (FR). A proposal by Burnod et al. (1999) suggested that the series of FRs, required for transferring information in retina-based FR into a body-centered FR, is encoded by different cells along the visual pathway, following a sensory gradient of increasing complexity. Indeed, along the visual pathway (the "what stream"), the information flows from the primary visual cortex (V1) to the temporal lobes, including the inferior temporal area (IT) and the superior temporal sulcus(STS). IT contains populations of neurons that separately exhibit sensitivity to a variety of objects. Some of these populations are sensitive to the size and orientation relative to an viewer-centered FR, whereas others react in an object-centered FR (Booth and Rolls, 1998). Similarly, neurons in macaques' STS, have been found to respond to specific human body parts and correlate with various quantities such as the position, rotation and translation of limbs, hands, faces, eyes; as well as with complex motions such as walking. Perrett et al. (1989) showed that the FRs in which these neurons seem to react are multiple. Moreover, there is a body of evidence that spatial visual properties such as direction, orientation and size of objects are also encoded in PPC (Sakata et al., 1999). Finally, concerning the distance of the target objects and observed bodies, neurons activities in the ventral pathway and parietal cortex have been shown to correlate this parameter, firing differently for close or far stimuli in a modulatory fashion (Dobbins et al., 1998; Sakata et al., 1980). All these regions are tightly coupled and form a complex network (Wise et al., 1997) that plays a fundamental role in primates ability to reproduce movements and goal-directed actions, such as transforming viewer-centered information into an other-centered representation.

In former work (Arbib et al., 2000; Billard and Matarić, 2001), we started developing computational models of the complete visuomotor pathway underlying imitative behavior. In this paper, we present a neural model that accounts for the ability to perform arbitrary frame of reference transformations and to display mimicry of hand motions. The model attempts, once the goal has been clearly identified, to explain the core circuits underlying the ability to map goal-directed motion performed by others into a frame of reference located onto one's own body. Such basic imitative behavior is displayed both by monkeys and humans.

2 A Mimicry Task

The mimicry task we consider in this paper is illustrated in Figure 1. It consists of the following: An imitator and a demonstrator face one another. The demonstrator produces various movements with his right hand. The imitator tries to reproduce the demonstrator's actions simultaneously (immediate imitation). The imitator attempts to reach to the same location as the demonstrator's hand in its own frame of reference. For instance, when the demonstrator's hand performs a circular trajectory on his left side, the imitator has to perform a similar hand motion on his own left, independently on the demonstrator's orientation. Indeed, the imitator could face the observer, be on a profile view or even be turned upside-down. It must be able to still perform the correct frame of reference transformation.

However simple this task appears to be, it is non trivial to model the neural processes that underly it. Thus, we describe a distributed neural model, inspired from neurophysiological evidence of population vector coding, that is able to perform such



Figure 1: Illustration of the frames of reference transformation required to transfer the target from the demonstrator's view point to that of the imitator.

transformations. As an illustration, we implement this mechanism into a robotic platform using a minihumanoid robot shown in Figure 8, to perform immediate imitation of hand drawings produced by a human demonstrator.

2.1 The Frames of Reference Problem

Consider the core problem tackled in this paper: How does the central nervous system perform frames of references transformation in order to build a bodycentered or object-centered representation?

In mathematical terms, as illustrated on Figure 1, the question is how can we transform a vector \vec{v} given in a referential R into \vec{v}' in R', knowing the vector $\vec{v}_{\rm T}$ across the origins of the two referentials, and the axes of the referential R' itself, expressed in R. We assume that R and R' are given by

$$R = \{O, \vec{e}_1, \vec{e}_2, \vec{e}_3\}$$
$$R' = \{O', \vec{e}'_1, \vec{e}'_2, \vec{e}'_3\}$$

where $OO' = \vec{v}_{\rm T}$, and \vec{e}'_i , \vec{e}_i , $\forall i \in \{1..3\}$ correspond to the principal axes, as unit vectors, of the demonstrator's body and of the observator's body, respectively. These axes correspond to the right-left, feet-head and back-front axes, respectively. The orientation of R' with respect to R is given by the rotation matrix $M_{R'}$:

$$M_{R'} = \vec{e}'_1 | \vec{e}'_2 | \vec{e}'_3 . \tag{1}$$

By writing down the classical transformations across referentials and considering M_R as an identity matrix, we get the following forward and inverse equations:

$$\vec{v} = M_{R'}\vec{v}' + \vec{v}_{T}$$

$$\Leftrightarrow$$

$$\vec{v}' = M_{R'}^{-1}(\vec{v} \quad \vec{v}_{T}).$$
(2)

If we consider now that $M_{R'}$ is orthonormal, we know that $M_{R'}^{1} = M_{R'}^{T}$. This allows us to rewrite the previous equation using the dot product and we find:

$$\vec{v}' = \sum_{i \in \{1..3\}} \vec{e}'_i \cdot (\vec{v} \quad \vec{v}_{\rm T}) \vec{e}_i.$$
 (3)

Such transformation can thus be reduced to a combination of relatively simple (from a neurophysiological point of view) vectorial operations, consisting of sums, dot products, and unitary vector scaling.

This way, the vector \vec{v}' pointing to the target in the demonstrator's referential can be directly mapped into the imitator-centered referential, so that the demonstrator's target is considered as the imitator's one.

2.1.1 Population Vector Coding

We use the *population vector coding* paradigm as a neurophysiological substrate for representing each of the vectors of our referentials.

In this paper, we define a population as an ensemble of neurons whose distributed firing activities are correlated to a single macroscopic quantity that is a vector \vec{v} in a given frame of reference R. In such populations, each neuron is tuned to a preferred direction \vec{r} , i.e. its firing activity is maximal when \vec{v} and \vec{r} are collinear and point to the same direction, and decrease as \vec{v} diverges from \vec{r} . Then, in order to extract the information from a populations of neurons, as originally proposed by Georgopoulos (1996), we use the population vector. Considering that each neuron votes for its preferred direction proportionally to its firing activity, by taking the average of all these votes, we obtain the vector encoded by this population, i.e. the population vector.

2.1.2 View Sensitive Cells Defining Referentials

As mentioned in Section 1, we know from neurophysiology that neurons in STS and IT are sensitive to different orientations or views of bodies and objects, respectively. These neurons firing activities have also been shown to be correlated to different frames of references, mainly in a viewer, object or goal centered reference frame. Moreover, these populations of neurons exhibit a large range of preferred directions, tending to cover uniformly the possible orientation given rotations around the three principal axis.

In order to model these cells, we assume that there are three distinct populations of neurons encoding separately the three principal axis of a observed body or object. This principle might be consistent with neurophysiological data despite not being completely experimentally proved. Indeed, to our knowledge, there are no neurophysiological experiments that have systematically tested the response of orientation sensitive cells to the complete ensemble of possible orientation. Indeed, usually sole the classical rotations along the three principal axis were tested.

2.2 The Model

This section presents a summary of a neural model for frames of reference transformations that we have, in its major parts, already proposed in Sauser and Billard (2005). For more information on the mathematical development and on the implementation details, please refer to this paper. The novelty here, concerns the parallel use of three principal axes determining a frame of reference, rather than a set of angles that code for a series of rotations that are performed serially.

2.2.1 An attractor network

Let us consider Ω , a continuous population of neurons where each unit participating in the population is characterized by its preferred direction \vec{r} . In this paper, the preferred directions are assumed to be uniformly distributed along a 3 dimensional subspace $\Gamma = {\vec{r} \in \mathbb{R}^3 \mid ||\vec{r}|| = 1}$, that corresponds to the surface of a unitary sphere. The response of the whole population, the population vector, is given in a continuous form by

$$\vec{P} = \frac{1}{\kappa} \oint_{\Gamma} f \ u_{\vec{r}} \ \vec{r} \,\mathrm{d}\vec{r} \tag{4}$$

where $\kappa = \frac{2\pi}{3}$ is a normalization factor, $u_{\vec{r}}$ the neuron's membrane potential with preferred direction \vec{r} , and $f(u_{\vec{r}})$ its firing activity. f is a non-linear function equal to $f(x) = \max(0, x)$.

Let us now consider an attractor network (Salinas and Abbott, 1996) made of a fully connected population of neurons whose dynamics is governed by

$$\tau \dot{u}_{\vec{r}} = u_{\vec{r}} + \oint_{\Gamma} w_{\vec{r}' \to \vec{r}} f(u_{\vec{r}'}) \, \mathrm{d}\vec{r}' + x_{\vec{r}}$$
$$w_{\vec{r}' \to \vec{r}} = \gamma(\eta) \left(\vec{r}' \cdot \vec{r}\right) \tag{5}$$



Figure 2: On the left, architecture of the two layers neural network producing a non-linear composition of its inputs. On the right, the symbolic illustration of this network as will be used further in the paper.

where $w_{\vec{r}' \to \vec{r}}$ are the lateral weights that exhibit symmetric, rotation invariant, and center surround excitation inhibition characteristics, $x_{\vec{r}}$ is the external synaptic input, and $\gamma(\eta)$ is a scaling factor depending on the network parameter $\eta \in]0, 1[$ that controls the influence of the lateral weights². Assuming that the network input $x_{\vec{r}}$ is composed of a vectorial and a constant homogeneous input of the form given by

$$\begin{aligned} x_{\vec{r}} &= \vec{r} \cdot \vec{v} + h \\ &= \beta_v \left(\vec{r} \cdot \vec{r}_v \right) + h \end{aligned} \tag{6}$$

where $\beta_v = \|\vec{v}\|$ and $\vec{r}_v = \frac{\vec{v}}{\|\vec{v}\|}$. We have shown that the activity profile of this network converges toward a stable state that can be approximated by

$$u_{\vec{r}}^{\star} \approx h + \frac{1}{\chi(\eta)} \left(\vec{r} \cdot \vec{v} \right) + \frac{1}{\eta} h \left(\vec{r} \cdot \vec{r}_{v} \right)$$
(7)

where $\chi(\eta) = 1 - \gamma(\eta)\frac{\pi}{3}$. We can see that the approximation of this activity profile reflects both the vectorial and constant inputs, plus a modulatory term, which is the result of the interactions of the recurrent connectivity.

2.2.2 A Two Layers Neural Network

As seen on Equation 7, the current attractor network produces, as an output, a sum of vectorial and constant terms. In order to strictly keep the multiplicative term and thus have a network capable of producing a non-linear composition of its two input sources, we build a two layers neural architecture as illustrated in Figure 2. The first layer consists of the attractor network. The second layer is composed of another population Ω_{O} , o for output, without lateral weights. It receives projections from the recurrent population

²In the present case of populations representing 3D vectors, $\gamma(\eta) = \left(\frac{\pi}{3}\left(2+3\eta - \eta^3\right)\right)^{-1}$, such that Equation (5) has a non trivial state of convergence.



Figure 3: The architecture and connectivities of the gain field. The spheres containing a referential and a vector correspond to populations of neurons coding for a vector in a given referential.

using one to one synapses, and inhibitory inputs corresponding to the vectorial and constant inputs of the recurrent population with an appropriate scaling, such that

$$x_{\vec{r}}^{\rm O} = \eta \ f(u_{\vec{r}}) \quad h \quad \frac{1}{\chi(\eta)} (\vec{r} \cdot \vec{v}) \ .$$
 (8)

Considering that the neurons of $\Omega_{\rm O}$ support an immediate integration mechanism such that $u_{\vec{r}}^{\rm O} = x_{\vec{r}}^{\rm O}$, we obtain, after a substitution in the previous equation using (7), that the firing rate converges toward

$$f(u_{\vec{r}}^{O}) = f\left(\eta \ f(u_{\vec{r}}) \quad h \quad \frac{1}{\chi(\eta)}\beta_{v}(\vec{r}\cdot\vec{r}_{v})\right)$$

$$\approx \begin{cases} h\left(\vec{r}\cdot\vec{r}_{v}\right) & \vec{r}\cdot\vec{r}_{v}>0, h>0\\ 0 & \text{otherwise.} \end{cases}$$
(9)

This network is capable of encoding independently two separate quantities, that are the direction \vec{r}_v and the amplitude h, regardless of the intensity of the directional input β_v . In vectorial terms, this means that given a vector \vec{v} and a scalar h, the output population vector will tend toward $h \frac{\vec{v}}{\|\vec{v}\|}$. Therefore, this model can be used to form a vectorial basis, inspired from classical linear algebra. Moreover, it will also be the building block of a bigger network, the gain field.

2.2.3 The Gain Field

In order to combine two different sources of vectorial information, we propose a model of gain field that follows an architecture and connectivity shown in Figure 3. It consists in an assembly of building blocks, described in Section 2.2.2, that define a new dimension denoted by $\vec{s} \in \Gamma$. The external inputs come from two different vectorial sources represented by a modulatory population Ω_{mod} and an vectorial population Ω_{v} , that encode the vectors \vec{v}_{mod} and \vec{v} , respectively. They are separately applied to each dimension of the gain field, \vec{r} and \vec{s} , respectively. Hence, the input for each neuron \vec{r} of each layer \vec{s} in the gain field Ω_{GF} is defined by

$$x_{(\vec{r},\vec{s})}^{\text{GF}} = \oint w_{\vec{r}' \to \vec{r}}^{\text{mod} \to \text{GF}} f(u_{\vec{r}'}^{\text{mod}}) \, \mathrm{d}\vec{r}' + \oint w_{\vec{r}' \to \vec{s}}^{\text{v} \to \text{GF}} f(u_{\vec{r}'}^{\text{v}}) \, \mathrm{d}\vec{r}' \\ = (\vec{r} \cdot \vec{v}_{\text{mod}}) + (\vec{s} \cdot \vec{v})$$
(10)

Then, if we substitute this equation into (9), we obtain that the gain field output firing activity converges toward

$$\begin{aligned} f(u_{(\vec{r},\vec{s})}^{\text{GFO}}) &\approx & \beta_v(\vec{s}\cdot\vec{r}_v)(\vec{r}\cdot\vec{r}_{v_{\text{mod}}}) \\ &\approx & (\vec{s}\cdot\vec{v})(\vec{r}\cdot\vec{r}_{v_{\text{mod}}}) \end{aligned}$$
(11)

From this, we can see that the activity profile of the gain field output is symmetric and that the peak is located at the intersection of the directions currently encoded by both source populations. Moreover, considering the amplitude of its activity, sole the amplitude of \vec{v} is taken into consideration in this network. This property allows transformations that guarantee that the amplitude of the transformed vectorial quantity is preserved (Sauser and Billard, 2005).

2.2.4 Projections on Principal Axis and Others Centered Frame of Reference

The final step, and the new part of our model, is to show how our neural network model can perform arbitrary frames of reference transformations by applying the principles mentioned in Section 2.1 (see Equ. (3)). As shown in Figure 4, we consider five sources of information arising from five populations of neurons that encode $\vec{e}'_i, i \in \{1..3\}, \vec{v}$ and $\vec{v}_{\rm T}$. In order to compute the dot product, we need three gain fields whose modulatory inputs are connected to the populations coding for the principal axis $\vec{e}'_{i \in \{1..3\}}$, while their vectorial inputs are linked to the difference between populations coding for \vec{v} and $\vec{v}_{\rm T}$ that are connected using excitatory and inhibitory synapses, respectively. These gain fields project then to another population that will receive the result of the transformation: the vector \vec{v}' in a body or object centered frame of reference using the following synaptic weights, $\forall i \in \{1..3\}$

$$w_{(\vec{r},\vec{s})\rightarrow\vec{r}'}^{\text{GFO}_i\rightarrow\mathbf{v}'} = \frac{1}{\kappa^2} (\vec{r}\cdot\vec{s}) \ (\vec{r}'\cdot\vec{e}_i). \tag{12}$$



Figure 4: The architecture and connectivities of the model that can perform frames of reference transformations.

Then, using the activity profile of the gain fields described by Equation (11), each neuron of the final population receives a synaptic input equal to

$$x_{\vec{r}'}^{\mathbf{v}'} = \sum_{i \in \{1..3\}} \oint f \quad u_{(\vec{r},\vec{s})}^{\text{GFO}_i} \quad w_{(\vec{r},\vec{s}) \to \vec{r}'}^{\text{GFO}_i \to \mathbf{v}'} \, \mathrm{d}\vec{r} \, \mathrm{d}\vec{s}$$
$$= \sum_{i \in \{1..3\}} (\vec{v} \quad \vec{v}_{\mathrm{T}}) \cdot \vec{e}'_i \ (\vec{r}' \cdot \vec{e}_i)$$
$$= \vec{r}' \cdot \vec{v}' \tag{13}$$

This equation means that this population is now encoding \vec{v}' in a body or object centered frame of reference.

2.3 Experimental Setup

We implemented this system in a kinematic simulator of a pair of demonstrator - imitator humanoid avatars



Figure 5: Overview of the system implementation on a robotic platform. The surrounding dotted rectangle indicates the parts used the simulation.

(see Fig. 6) and in a humanoid robot, as shown in Figure 8. An overview of the overall system architecture is illustrated in Figure 5. The visual system consists in two webcams connected to a color-based stereo vision software that allow the tracking of specific colors marks in 3D space. The human demonstrator is placed in front of the cameras, with three different color marks on the left and right of his torso, and on his hand. Assuming that he is always in a standing posture the two marks on the body are sufficient to uniquely determine the demonstrator's principal axis, that are \vec{e}'_i , $i \in \{1..3\}$. The visual system also provides the body and hand position in a viewer centered frame of reference, $\vec{v}_{\rm T}$ and \vec{v} , respectively. These information are fed into our neural network in order to compute the target location in the demonstrator's body centered reference frame. It is directly applied to a self-centered frame of reference that gives the imitator its own target. In order to allow the robot to reach the target, this position of the target with respect to the imitator is fed to an inverse kinematic algorithm adapted from Wang and Verriest (1998), that provides the sequence of joint angles to the robot.

3 Results

3.1 Mimicry of hand gestures

We conducted simulations, in which the demonstrator avatar draw 8 different figures. Figure 7 shows superimposed the trajectories performed by the demonstrator and the imitator. Demonstrated and imitated



Figure 6: One avatar is drawing a figure while the other imitates the demonstrated trajectory.

movements show a high qualitative resemblance. However, one can observe a systematic shift in space and a slight deformation of the figure. This is an artifact resulting from the non-uniformity of the distribution of preferred directions in our neural population³. In other words, the neural populations produce a nonuniform map of their inputs, resulting in a slight deformation of the three dimensional representation of the target vectors.

We, then, conducted experiments, in which a humanoid robot imitated 4 trajectories produced by a human demonstrator. Figure 8 shows, superimposed, the trajectories of the demonstrator's and imitator's hand for the four examples. We can observe that the results are similar to those obtained in simulation. The imitation is qualitatively good. However, it suffers from a systematic shift in space and rescaling in amplitude. In addition, the use of a stereovision system for recording demonstrated and imitated trajectories creates a new source of errors.



Figure 7: Eight trajectories followed by the demonstrator's hand (dotted lines) and by the imitator's hand (plain lines) in simulation.



Figure 8: Top Figure: Hoap-2 a mini-humanoid robot built by Fujitsu, provided with 25 degrees of freedom, including 4 on each arm. The robot imitates a human trajectory forming an "S", while tracking the demonstrator's gesture using a pair of fixed cameras. Bottom Figure: Four trajectories followed by the demonstrator's hand (dotted lines) and by the robot's hand (plain lines) in simulation.



Figure 9: Error recorded during a simulation batch where populations with different parameters were given random input vectors and referentials.

3.2 Error measures

In addition to the errors that appear by discretizing continuous equations, the approximation we made in our mathematical development (see Equ. (7)) is a source of systematic errors between the theoretical resulting vector, denoted by $\vec{v}^{\prime\star}$, computed with classical algebraic equations, and the result \vec{v}^{\prime} produced by our network. To quantify them, we define E_{β} , the error on the amplitude, and E_{θ} , the error on the direction, by

$$E_{\beta}(\vec{v}\,',\vec{v}\,'^{\star}) = \frac{|\,\|\vec{v}\,'\|\,\,\|\vec{v}\,'^{\star}\|\,|}{\|\vec{v}\,'^{\star}\|} \tag{14}$$

$$E_{\theta}(\vec{v}\,',\vec{v}\,'^{\star}) = \operatorname{acos}\left(\frac{\vec{v}\,'\cdot\vec{v}\,'^{\star}}{\|\vec{v}\,'\|\,\|\vec{v}\,'^{\star}\|}\right) \quad (15)$$

that correspond to the relative difference between their norms, and to the angle they form, respectively. Figure 9 shows the errors E_{β} and E_{θ} that were

³In Sauser and Billard (2005), we showed that only a "quasi" uniform distribution of preferred directions can be obtained using iterative algorithms.

measured during a simulation batch where random vectors where transformed into random referentials. The different curves correspond to different network sizes. First, we can see as expected that the bigger is a population, the smaller are the errors. Second, consistently with our previous work (Sauser and Billard, 2005), the parameter η has a ambivalent influence on the network. On the one hand, small values increase the importance of the recurrent connections, hence increasing the errors due to an imperfect distribution of preferred directions. On the other hand, big values induce more errors due by our mathematical approximations (see Equ. (7)). These two properties explain why, on the left of the figure, an optimum can be observed.

4 Discussion and Conclusion

The model, presented in this paper, provides an example of neural mechanism for the representation of others in a self-centered frame of reference. As such, it is an important step toward a full-scale imitation model. Indeed, as illustrated in this paper, a model for solving the frames of reference transformation problem provides us automatically with a simple imitation mechanism. Note that the present model does not yet explain the tendency humans have to perform imitation in mirror fashion when reproducing meaningless gesture, and when demonstrator and imitator face each other (Wholschläger et al., 2003). It only shows a solution to the frame of reference problem. Note that the model could be extended to address this issue. The preference for mirror imitation could simply be an effect of early visual processing, occurring prior to the frame of reference transformation, that would represent the demonstrator's body in a referential that reflect best the natural symmetries of the human body; presenting motions perceived visually on the left handside of the imitator by corresponding motion on the left handside of the imitator.

Another important aspect not yet addressed by our model is how the rescaling of the demonstrator's motions to the imitator's body is performed. In the present implementation, rescaling is done by hand, providing a vector of an appropriate size to the network, so that the resulting vector after convergence lies within the robot's range of motion. The model could be extended to encapsulate explicitly the rescaling aspect, by exploiting the multiplicative nature of the network. Moreover, such a neural representation would be in accordance with biological evidence that neurons located in the visual cortex fire in response to the size of an object, regardless of the distance to the object (Dobbins et al., 1998; Sakata et al., 1980).

There is as yet no evidence to support our model's hypothesis that orientation sensitive cells in the visual areas STS and IT are grouped in populations that encode the principal axes of the demonstrator's body. If evidence of such an encoding was to be found, this would suggest that such groups of neurons may form a basis (in the vectorial sense) of a goal centered representation of hand motion. Unfortunately, to our knowledge, no systematic experiment have shown a complete description of single cell sensitivity to all possible orientations. Note that if these cells were to encode the three principal axes, this would offer a highly redundant representation of motion. One could consider less redundant forms of encoding 3D frames of reference. However, as discussed by Marr (1982), such representations are difficult to determine and the three axes representation remain the most natural representation for 3D frames of reference. Furthermore, Deneve and Pouget (2003) proposed a model that deals with a two dimensional object-centered representation using basis functions. The authors argue that a redundant neural substrate is well-suited to reduce neural noise and to simplify the complexity of single cell computation.

The time required for the model to perform a FR transformation is independent on the orientation of the two frames of reference. Such a result is in contradiction with the observation that humans produce a longer reaction time, when required to perform mental rotations in an "unusual" orientation, such as shifting an image upside-down. One could, however, imagine that another mechanism is at play. In absence of visual input, such a mechanism would set the principal axes of the demonstrator's referential to a default state (i.e. setting the preferred directions of the network in our model to a default value), expressing the imitator's expectation that the demonstration would stand vertically and would face him. In this case, the network's state in our model will take more time to match unusual visual orientations; hence, reproducing the expected observation. Note that the delay could also be due to a longer processing phase during preprocessing of the visual field, for recognizing the body features (used then to set the landmarks for determining the axes).

Finally, we showed that frame of reference transformation performed by the model result in qualitative discrepancies between demonstrated and imitated trajectories, while ensuring a high qualitative resemblance across demonstrated and imitated motions. Note that humans show also imprecision in their imitation, if other constraints, such as an alignment to landmarks, are not specified or absent. In future work, we will compare the imprecisions made by the model to those done by human imitators.

The model's implementation we presented in this paper focused on a body-centered frames of reference transformation. The model is, however, quite general and could, also, be applied to object-centered representations. The later representation being crucial to performing several daily tasks. In future work, this model will be adapted to form both object- and goalcentered representations in order to provide context dependent information for goal-directed imitation.

Acknowledgments This work was supported by the Swiss National Science Foundation, through grant no 620-066127 of the SNF Professorships program.

References

- M. Arbib, A. Billard, M. Iacoboni, and E. Oztop. Mirror neurons, imitation and (synthetic) brain imaging. In *Neural Networks*, volume 13 (8/9), pages 953–973, 2000.
- H. Bekkering, A. Wohlschläger, and M. Gattis. Imitation of gestures in children is goal-directed. *Quarterly Journal of Experimental Psychology*, 53(1):153–164, 2000.
- A. Billard. Imitation: A review. In *The Handbook of Brain Theory* and Neural Network. 2nd Edition. Michael A. Arbib (editor), pages 566–569. MIT Press, 2002.
- A. Billard and M. Matarić. Learning human arm movements by imitation: Evaluation of a biologically-inspired connectionist architecture. *Robotics & Autonomous Systems*, 941:1–16, 2001.
- M.C.A Booth and E.T. Rolls. View-invariant representations of familiar objects by neurons in the inferior temporal cortex. *Cerebral Cortex*, 8:510–523, 1998.
- Yves Burnod, Pierre Baraduc, Alexandra Battaglia-Mayer, Emmanuel Guigon, Etienne Koechlin, Stefano Ferraina, Francesco Laquaniti, and Roberto Caminiti. Parieta-frontal coding of reaching: an integrated framework. *Experimental Brain Re*search, 129:325–346, 1999.
- J Decety, T. Chaminade, J Grezes, and A. N. Meltzoff. A PET Exploration of the Neural Mechanisms Involved in Reciprocal Imitation. *Neuroimage*, 15:265–272, 2002.
- S. Deneve and A. Pouget. Basis function for object-centered representations. *neuron*, 37:347–359, 2003.
- A.C. Dobbins, R.M. Jeo, J. Fiser, and J.M. Allman. Distance modulation of neural activity in the visual cortex. *Science*, 281: 552–555, 1998.
- L. Fogassi and V. Gallese. The neural correlates of action understanding in non-human primates. In Mirror Neurons and the Evolution of Brain and Language, Advances in Consciousness Research, John Benjamins Publishing, pages 13–35, 2002.

- A.P. Georgopoulos. On the translation of directional motor cortical commands to activation of muscles via spinal interneuronal systems. *Cognitive Brain Research*, 3:151–155, 1996.
- C. Heyes. Causes and conscequence of imitation. Trends in Cognitive Sciences, pages 253–261, 2001.
- M. Iacoboni, R.P. Woods, M. Brass, H. Bekkering, J.C. Mazziotta, and G Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286:2526–2528, 1999.
- F. Lhermite, B. Pillon, and M. Serdaru. Human autonomy and the frontal lobes, par i: imitation and utilization behavior: a neuropsychological study of 75 patients. *Annual Neurology*, 19:326–34, 1986.
- D. Marr. Vision. Cambridge MA: MIT Press, 1982.
- N. Nishitani and R. Hari. Temporal dynamics of cortical representation for action. In *Proceedings of the National Academy of Sciences of the United States of America*, volume 97(2), pages 913–918, 2000.
- D. I. Perrett, M Harries, A.J. Mistlin, and A.J. Chitty. Three stages in the classification of body movements by visual neurons. In H.B et al. Barlow, editor, *Images and Understanding*, pages 94– 107. Cambridge University Press, 1989.
- G. Rizzolatti, L. Fadiga, V. Gallese, and L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3(2):131–141, 1996.
- H. Sakata, H. Shibutani, and K. Kawano. Spatial properties of visual fixation neurons in posterior parietal association cortex of the monkey. *Journal of Neurophysiology*, 43:1654–1672, 1980.
- H. Sakata, M. Taira, M. Kusunoki, A. Murata, K. Tsutsui, Y. Tanaka, W.N. Shein, and Y. Miyashita. Neural representation of three-dimensional features of manipulation objects with stereopsis. *Experimental Brain Research*, 128:160–169, 1999.
- E. Salinas and L.F. Abbott. A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences USA*, 93:11956–11961, 1996.
- E. Sauser and A. Billard. Three dimensional frames of references transformations using recurrent populations of neurons. *Neurocomputing, In press*, 2005.
- T. Shimomura and E. Mori. Obstinate imitation behaviour in differentiation of frontotemporal dementia from alzheimer's disease. *Lancet*, 352(9128):623–4, 1998.
- X.G. Wang and J.P. Verriest. A geometric algorithm to predict the arm reach posture for computer-aided ergonomic evaluation. *Journal of Visualization and Computer Amination*, 9:33– 47, 1998.
- A. Wholschläger, M. Gattis, and H. Bekkering. Action generation and action perception in imitation: an instance of the ideomotor principle. *Phil. Trans. R. Soc. Lond.*, 358:501–515, 2003.
- S.P. Wise, D. Broussaoud, P.B. Johnson, and R. Caminiti. Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annual Reviews in Neuroscience*, 20:25–42, 1997.

Nonconscious Imitation has Consequences That Go Beyond the Dyad

Rick van Baaren^{*} *University of Nijmegen R.vanbaaren@psych.ru.nl Tanya L. Chartrand[†] [†]Duke University TLC10@duke.edu

Abstract

People often nonconsciously imitate other people and imitation has positive consequences for the interaction. We argue that imitation not only has consequences for the way in which an imitated person feels towards the imitator, but that imitation also changes the way in which the imitated person feels towards other people in general. In two studies participants were unobtrusively imitated by a confederate and the effects on interpersonal closeness were measured. Experiment 1 showed that imitated participants feel closer to non-specified other people *in general* compared to non-imitated participants. Experiment 2 replicated this result using a seating distance measure. Together, these studies reveal that imitation has consequences that go beyond the dyad. Now the challenge is to look for a system, which can explain this implicit imitation recognition

1 Introduction

Imitation is often nonconsciously used as a tool to influence people (Cheng & Chartrand, 2003; Lakin & Chartrand, 2003). When imitation goes unnoticed, several beneficial consequences arise. For one, imitation increases liking and rapport in interactions (Bailenson & Yee, 2004; Chartrand & Bargh, 1999; Suzuki, Takeuchi, Ishii, & Okada, 2003). It also increases pro-social behavior (Van Baaren, Holland, Kawakami, & van Knippenberg, 2004). In the present studies, we want to start to investigate how this happens. That is, why do we like imitators more than non-imitators, and why do we behave more pro-socially towards them?

1.1 The Positive Consequences of Imitation

In many commercial books on influence and making friends, imitation is offered as one of the means to create a good impression on or rapport with others (e.g Lieberman, 2000). There is now experimental evidence that this occurs. Several studies that manipulated imitation by having one individual either mimic another person or not found positive consequences of subtle imitation. The developmental psychology literature documents evidence that infants react more favorably towards adults who imitate them than adults who do not (Asendorpf, Warkentin, & Baudonniere,1996; Meltzoff, 1990). In these studies, however, it is unknown whether these infants are aware or unaware of the imitation. Humans have a predisposition to unwittingly and automatically mimic the behaviors of others (Chartrand & Bargh, 1999; Prinz, 1990). It usually occurs nonconsciously and remains unnoticed.

Positive consequences have been observed for this nonconscious mimicry of body movements and speech variables. In a typical experiment, a participant and a confederate work on an irrelevant task. During that task, the confederate mimics (or not) the posture, mannerisms, and behaviors of the participant after a short delay. These can be gestures or movements such as face-rubbing, foot-shaking, playing with a pen, orientation of the body (avoiding movements that indicate power or status), or speech variables such as using the same phrases of speech. After this imitation manipulation, the dependent variable is assessed, which is often an evaluation of or behavior towards the confederate.

Chartrand and Bargh (1999) found that participants who were subtly mimicked by a confederate liked that confederate more and had smoother interactions with that confederate. Interestingly, similar consequences have been observed in humancomputer interactions. Bailenson and Yee (2004) had a realistic interface agent (i.e., an avatar using virtual reality technology) either imitate the participant's head movements or perform different head movements. The imitating interface agents were rated as more likeable and more persuasive than the non-mimicking avatars. Similarly, Suzuki, Takeuchi, Ishii, and Okada (2003) found that mimicry of certain properties of a participant's voice by a computer agent led to more favorable evaluations of the computer agent. Thus, the evaluative consequences of imitation are not unique to human-human interactions, although it should be noted that the avatars were very lifelike and thus were treated as human.

Van Baaren et al. (2004, Experiment 1) found that being imitated not only influences evaluations such as liking or rapport, but also makes people behave in a more pro-social manner. In the first study, an experimenter unobtrusively imitated participants (or not). After the imitation manipulation the experimenter left the room and returned a short while later while carrying some pens and papers. Upon entering the room, the experimenter "accidentally" dropped the pens. The dependent variable was whether participants got off their chairs and started to help (a measure used by Macrae & Johnston, 1998). The results revealed that imitated participants were considerably more helpful than non-imitated participants.

What was confounded in the above study, and in several other studies of the consequences of imitation, is that the effects of imitation were measured vis-a-vis the imitator. This is important to note, because it could theoretically be possible that the effects of imitation are not restricted to the dyad and the imitator. Perhaps the effects extend beyond the relation between the imitator and the imitated. Perhaps it affects the imitated person in a more fundamental way. It is possible that imitation makes one more pro-social towards other people in general. This is exactly what was observed in two experiments (Van Baaren et al. 2004, Experiment 2 and 3). Imitated participants behaved more pro-socially towards a second experimenter and gave more money to a charity than non-imitated participants. Do the observed effects on people other than the imitator suggest that being mimicked makes one feel closer to other people in general?

1.2 The Present Studies

In the present studies we investigate whether imitation leads to increased interpersonal closeness, and more specifically, whether imitation makes people feel closer to other people *in general*. In two studies participants were imitated or not and their subjective connectedness towards undefined others (Experiment 1) or unknown others (Experiment 2) was assessed.

2 Experiment 1

2.1 Method

2.1.1 Overview

Participants enrolled in an "advertisement study", during which they rated 10 advertisements on some irrelevant dimensions. During that task, following the procedure by Chartrand and Bargh (1999), an experimenter imitated the posture, gestures, and mannerisms of half the participants. The other half of the participants was not imitated. After this task, participants filled in a questionnaire designed to measure how close participants felt to "people in general".

2.1.2 Participants and design

Twenty-six participants (17 women and 9 men) were paid \$2 for their participation in this study. The experiment had a single factor (behavior: imitation or no-imitation) between-subjects design.

2.1.3 Procedure

Upon arrival at the laboratory, participants were led into a room by the experimenter and seated behind a desk. The participant's chair half-faced the experimenter. The experimenter, who was blind to the hypothesis, seated himself behind a desk and explained that the experiment was a advertisement study that tested the reaction of people to certain types of ads. The task of the participant was to look at each of the 10 ads and take about 30 seconds to describe his or her feelings toward the specific ad. The experimenter wrote down the answers on the note-pad in front of him. During the task, the experimenter would imitate the behaviors of half the participants. Specifically, the orientation of the body (forward or backward) and the position of the arms and legs were imitated after a several-second delay. In addition, gestures such as touching one's face or hair were (contra-laterally) imitated. The other participants were not imitated, which meant that the experimenter had to actively avoid having the same posture and gestures. The experimenter was trained to mimic (and to anti-mimic), but was unaware of the hypothesis.

After the advertisement task, participants were given a modified Inclusion-of-the-Other in-the-Self-Scale (IOS-scale, Aron, Aron & Smollan, 1992) that was designed to measure the closeness they felt towards other people in general. The closeness task depicted six pairs of circles, which were increasingly overlapping with each other. The instructions explained that one of the circles represented the participant and the other circle represented "other people in general." The participant had to indicate how close he or she felt towards "other people in general" by selecting one of the six pairs of circles. Higher numbers indicate a smaller felt distance.



Figure 1. Inclusion-of-the-Other in-the-Self-Scale.

After this task, the participant was thanked, paid and debriefed. Importantly, no participant indicated awareness of the imitation.

2.1.4 Results and Discussion

To test the prediction that participants who were imitated would show a greater closeness to "other people in general", the scores on the closeness questionnaire were submitted to a 2 (Behavior: imitation or no-imitation) X 2 (Gender: male or female) between subjects analysis of variance. As expected, a main effect for Behavior was found, F(1,22) = 6.46, p < .02. Participants who had been imitated by the experimenter felt closer to people in general (M =4.3) than the participants that had not been imitated (M = 3.5). In addition, a main effect for gender was found, F(1,22) = 7.64, p < .02, confirming that women feel closer to other people (M = 4.2) than men (M = 3.2). No interaction was found (F < 1). During the debriefing, no participant indicated awareness of the imitation (or lack thereof) or of its effect on the dependent variable.

These results confirmed the hypothesis that imitation increases interpersonal closeness towards undefined others thereby extending previous findings that imitation increases liking, rapport and prosocial behavior. The present data show that, after being imitated, people also feel closer to others. It is important to note that it is an increased closeness to "others in general," not toward any specific person. This finding is consistent with the finding that imitation also stimulates pro-social behavior to people other than the imitator, indicating that its effects are diffuse and not specifically targeted at one single person (Van Baaren et al., 2002).

The results of Experiment 1 furthermore indicated that women feel closer to others in general than men. Following the reasoning by Cross and Madson (1997), this effect may be explained by different socialization process of men and women in Western societies. If women pay more attention to and are more concerned with relationships, it is likely that this interdependent construal of the self is associated with a reduced interpersonal distance and more empathy.

Instead of using an abstract measure of interpersonal distance, in Experiment 2, we examined the consequences of imitation on a concrete, behavioral level. Specifically, seating distance from the imitator was measured (see Macrae, Milne & Bodenhausen, 1998).

3 Experiment 2

3.1 Method

3.1.1 Overview

During an ostensible interview, half of the participants were imitated by an experimenter. Afterwards, all participants were asked to take a seat in an adjacent room, where a bag, a jacket, and some documents on one of the chairs indicated the presence of another person. The dependent variable was the distance (measured in number of chairs) between the occupied chair and the chair on which the participant chose to sit.

3.1.2 Participants and design

Fifty-eight undergraduates (35 women and 23 men) were paid \$2 for their participation in this study. The experiment had a single factor (Behavior: imitation or no-imitation) between-subjects design.

3.1.3 Procedure

Upon arrival at the laboratory, participants were led into a room by a male experimenter, who informed them that they would take part in two separate studies. In the first study, the experimenter interviewed the participant about their travelling behavior (with the help of a questionnaire). During the interview, he unobtrusively imitated (contra-laterally) the posture and behavioral mannerisms of the participants randomly assigned to the *imitation* condition. In the no-imitation condition, participants were treated likewise by the experimenter, except that they were not imitated. It is important to note that in the noimitation condition, the experimenter was actively avoiding imitation, which meant that he had to pay as much attention to the participant and the participant's behavior as in the imitation condition.

After the imitation manipulation (the bogus interview), the participants were thanked and asked to take a seat in an adjoining room while waiting for the second study. The experimenter made clear he would not be supervising the second study and that a different experimenter would pick up the participant. In this waiting room, there were five chairs placed side by side along one of the walls. On top of the leftmost chair, a bag, a jacket and some documents were placed, thereby indicating the presence of another (and unknown) person. The distance between the "occupied" chair and the chair on which the participant chose to sit was an implicit measure of interpersonal closeness, the dependent variable (Holland et al., 2004; Macrae et al., 1998). After a short wait, a new experimenter entered the waiting room to pick up the participant for the (irrelevant) second study, and wrote down on which chair the participant sat. Finally, the participant was thanked, paid, and debriefed. Importantly, no participant indicated awareness of the imitation.

3.1.4 Results and Discussion

To test the prediction that participants who were imitated would choose to sit closer to an unknown other, the distance between the participant's chair and the occupied chair was submitted to a 2 (Behavior: imitation or no-imitation) X 2 (Gender: male or female) between subjects analysis of variance. As expected, a main effect for behavior was found, F(1,54) = 6.68, p < .05. Participants who had been imitated by the experimenter sat closer to the occupied chair (M = 1.47) than the participants that had not been imitated (M = 1.96). No main effect of Gender or interaction between Gender and Behavior were obtained.

These results replicate Experiment 1 and provide further evidence that imitation makes people feel closer to other people in general. Interestingly, the observed main effect of gender in Experiment 1 was not replicated in Experiment 2, although other work has recently found such a gender effect on seating distance (Holland et al., 2004).

4 General Discussion

The present studies demonstrate that being imitated has consequences beyond the dyad. Although it may be the case that imitation creates a "special bond" between the person who is being imitated and the imitator, the effects are not restricted to that dyad. Being imitated affects people more profoundly; it makes them feel closer to others in general. How might this work?

Participants are not aware of being imitated, but still they are affected by it. Although admittedly speculative, it is possible that on a non-conscious level one registers that one's perception and action are in synchrony, an implicit form of imitation recognition (Nadel, 2002). This signal may function as a proprioceptive cue that subsequently influences the way in which we interact with our environment. A challenge to future research is to examine whether in fact such a signal exists. Does synchrony between one's own actions and the actions one perceives another person perform indeed result in observable activity? There is quite some evidence for a close link between perception and action (Prinz, 1990; Iacoboni et al., 1999) and this could provide an architecture from which the implicit sensation of being imitated may occur.

Of course, these findings do not mean that imitation *only* had consequences for the imitated person and not for the specific dyad. It seems more likely that these are additive effects. First, imitation has positive consequences for the specific relationship in which imitation occurs: relatively more liking and rapport. In addition, being imitated changes one's interaction with, and possibly perception of, the environment more fundamentally. More research is needed to investigate this. Hopefully this will bring about a fuller, more sophisticated understanding of human interaction.

A challenge to future research is to examine whether in fact such a signal exists. Does synchrony between one's own actions and the actions one perceives another person perform indeed result in observable activity? There is quite some evidence for a close link between perception and action (Prinz, 1990; Iacoboni et al., 1999) and this could provide an architecture from which the implicit sensation of being imitated may occur.

Of course, these findings do not mean that imitation only had consequences for the imitated person and not for the specific dyad. It seems more likely that these are additive effects. First, imitation has positive consequences for the specific relationship in which imitation occurs: relatively more liking and rapport. In addition, being imitated changes one's interaction with, and possibly perception of, the environment more fundamentally. More research is needed to investigate this. Hopefully this will bring about a fuller, more sophisticated understanding of human interaction.

Acknowledgements

Work on this article was supported by NWO grant 451-03-051 (VENI) awarded to the first author. The authors would like to thank Johan Karremans, Sebastian Kramer and Hanneke Ribberink for their suggestions and help.

References

Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of other in the self scale and the structure
of interpersonal closeness. Journal of Personality and Social Psychology, 63, 596-612.

- Asendorpf, J. B., Warkentin, V., & Baudonniere, P. (1996). Self-awareness and other-awareness. II: Mirror self-recognition, social contingency awareness, and synchronic imitation. *Devel*opmental Psychology, 32, 313-321.
- Bailenson, J., & Yee. N. (2004). Digital Chameleons: Automatic assimilation of nonverbal gestures in immersive virtual environments. Manuscript submitted for publication.
- Chartrand, T. L., & Bargh. J. A. (1999). The chamleon effect: The perception-behavior link and social interaction. Journal of Personality and Social Psychology, 76, 893 910.
- Cheng, C. M., & Chartrand, T. L. (2003). Selfmonitoring without awareness: Using mimicry as a nonconscious affiliation strategy. *Journal* of Personality and Social Psychology, 85, 1170-1179.
- Cross, S. E., and Madson, L. (1997). Models of the self: Self-construal theory and gender. *Psychological Bulletin, 122*, 5-37.
- Holland, R. W., Roeder, U., van Baaren, R. B., Brandt, A., & Hannover, B. (2004). Don't stand so close to me: self-construal and interpersonal closeness. *Psychological science*, 15, 237-242.
- Iacoboni, M., Woods, R., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. *Science*, 286, 2526-2528.
- Lakin, J. & Chartrand, T.L. (2003). Increasing nonconscious mimicry to achieve rapport. Psychological Science, 27, 145-162.
- Macrae, C. N., Bodenhausen, G. V., & Milne, A. B. (1998). Saying 'no' to unwanted thoughts: Self-focus and the regulation of mental life. *Journal of Personality and Social Psychology*, 74, 578-589.
- Macrae, C. N., & Johnston, L. (1998). Help, I need somebody: Automatic action and inaction. *Social Cognition*, 16, 400-417.
- Meltzoff, A. (1990). Foundations for developing a concept of self: The role of imitation in relating self to other and the value of social mirroring, social modeling, and self-practice in infancy. In D. Cicchetti & M. Beeghly (eds.),

The self in Transition (pp. 139-164). Chicago: University of Chicago Press.

- Nadel, J. (2002). Imitation and imitation recognition: functional use in preverbal infants and nonverbal children with autism. In A. Meltzoff et w. Prinz (Eds.), *The imitative mind* (pp. 42-62). Cambridge: Cambridge University Press.
- Prinz, W. (1990), A common coding approach to perception and action. In O. Neumann & W. Prinz, (eds.), *Relations between Perception* and Action (pp.167-201). Berlin: Springer-Verlag.
- Suzuki, N., Takeuchi, Y., Ishii, K., & Okada, M. (2003). Effects of echoic mimicry using hummed sounds on human-computer speech interaction. *Speech Communication*, 40, 559-573.
- Van Baaren, R. B., Holland, R. W., Kawakami, K., & van Knippenberg, A. (2004). Mimicry and pro-social behavior. *Psychological Science*, 15, 71-74.