Facilitating Play Between Children with Autism and an Autonomous Robot

Dorothée C. M. François

October 2008

Submitted to the University of Hertfordshire in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Adaptive Systems Research Group, School of Computer Science,
Faculty of Engineering and Information Sciences,
University of Hertfordshire

Facilitating Play Between Children with Autism and an Autonomous Robot

Dorothée C. M. François
University of Hertfordshire

PhD Thesis
October 2008

Abstract

This thesis is part of the Aurora project, an ongoing long-term project investigating the potential use of robots to help children with autism overcome some of their impairments in social interaction, communication and imagination. Autism is a spectrum disorder and children with autism have different abilities and needs. Related research has shown that robots can play the role of a mediator for social interaction in the context of autism. Robots can enable simple interactions, by initially providing a relatively predictable environment for play. Progressively, the complexity of the interaction can be increased.

The purpose of this thesis is to facilitate play between children with autism and an autonomous robot. Children with autism have a potential for play but often encounter obstacles to actualize this potential. Through play, children can develop multidisciplinary skills, involving social interaction, communication and imagination. Besides, play is a medium for self-expression. The purpose here is to enable children with autism to experience a large range of play situations, ranging from dyadic play with progressively better balanced interaction styles, to situations of triadic play with both the robot and the experimenter. These triadic play situations could also involve symbolic or pretend play.

This PhD work produced the following results:

- A new methodological approach of how to design, conduct and analyse robot-assisted play was developed and evaluated. This approach draws inspiration from non-directive play therapy where the child is the main leader for play and

the experimenter participates in the play sessions. I introduced a regulation process which enables the experimenter to intervene under precise conditions in order to: i) prevent the child from entering or staying in repetitive behaviours, ii) provide bootstrapping that helps the child reach a situation of play she is about to enter and iii) ask the child questions dealing with affect or reasoning about the robot. This method has been tested in a long-term study with six children with autism. Video recordings of the play sessions were analysed in detail according to three dimensions, namely Play, Reasoning and Affect. Results have shown the ability of this approach to meet each child's specific needs and abilities. Future work may develop this work towards a novel approach in autism therapy.

- A novel and generic computational method for the automatic recognition of human-robot interaction styles (specifically gentleness and frequency of touch interaction) in real time was developed and tested experimentally. This method, the Cascaded Information Bottleneck Method, is based on an information theoretic approach. It relies on the principle that the relevant information can be progressively extracted from a time series with a cascade of successive bottlenecks sharing the same cardinality of bottleneck states but trained successively. This method has been tested with data that had been generated with a physical robot a) during human-robot interactions in laboratory conditions and b) during child-robot interactions in school. The method shows a sound recognition of both short-term and mid-term time scale events. The recognition process only involves a very short delay. The Cascaded Information Bottleneck is a generic method that can potentially be applied to various applications of socially interactive robots.

- A proof-of-concept system of an adaptive robot was demonstrated that is responsive to different styles of interaction in human-robot interaction. Its impact was evaluated in a short-term study with seven children with autism. The recognition process relies on the Cascaded Information Bottleneck Method. The robot rewards well-balanced interaction styles. The study shows the potential of the adaptive robot i) to encourage children to engage more in the interaction and ii) to positively influence the children's play styles towards better balanced interaction styles.

It is hoped that this work is a step forward towards socially adaptive robots as well as robot-assisted play for children with autism.

Acknowledgements

This PhD research was supported by a research scholarship of the University of Hertfordshire. The work was partially conducted within the EU Integrated Project RobotCub (Robotic Open-architecture Technology for Cognition, Understanding and Behaviours) and was partially funded by the European Commission through the E5 Unit (Cognition) of FP6-IST under contract FP6-004370.

I would like to thank my three supervisors, Kerstin Dautenhahn, Daniel Polani and Stuart Powell, for their their guidance and support. Thank you to René te Boekhorst who introduced me to statistics for behavioural sciences, and to Chrystopher L. Nehaniv for his useful comments on the thesis.

I am grateful to my viva's examiners, Michael Hillman and Rod Adams, for their useful comments and suggestions during the viva.

The Adaptive Systems Research Group is a multidisciplinary team and it was a very good experience to be part of it. In particular, I would like to thank Megan Davis, Assif Mirza, Alexander Klyubin, Michael Walters, Arnaud Blanchard and Joe Saunders.

I am also grateful to everyone at the school where I conducted the play sessions. Thank you to the teachers and carers from the school who welcomed me and showed great interest in my research. My thanks also go to the parents for accepting their child to take part in the play sessions and for their enthusiasm for my work, and of course to the children who participated in the play sessions. The play sessions with the children were a rich and unforgettable personal experience.

Many thanks also go to all my friends. These several years in the UK were a great opportunity to share and learn about different cultures with my friends from Oxford and London. Thank you to Frances, Xavier, Ade, Lily, William, Lisa, Rickard, Maki, Ross, Rani, Ludovic and Sascha for the great moments we shared. Thank you to Gaëtane, too. And thank you to Sarah, who has always been very attentive despite the miles that separate us.

I would like to thank my family for their support. Thank you to my brother, Benjamin, my parents, Dominique and Hugues, and my stepparents, Inge and Guy, who encouraged me all along.

Finally, I would like to thank Mathieu, for the wonderful years we have together.

Contents

Chapter 1

Introduction

The research presented in this thesis is part of the Aurora project (Aurora, 2008), an ongoing long-term project investigating the potential use of robots to help children with autism overcome some of their impairments in social interactions (Dautenhahn and Werry, 2004, 2000). The Aurora project is constituted of two main streams of research. One stream focuses on the robot as an autonomous toy and, in particular, addresses the question of real-time recognition and adaptation to human-robot interaction styles (François et al., 2007, 2008b). The second one focuses on the potential role of the robot as a mediator (Davis et al., 2005; Robins et al., 2005a; François et al., 2009), i.e. as a salient object that helps children interact with other children or adults.

Children with autism have impairments in communication, social interaction and imagination skills (National Autistic Society, 2008; Powell, 2000). Autism is a spectrum disorder and children with autism have different needs, abilities and skills (Association, 1994). The advantage of enabling the child to interact with a robotic platform is to reduce the complexity of the interaction and to initially create a relatively predictable environment, so that it can be easier for the child with autism to feel at ease. It also aims at enabling the child to understand better the interactions taking place. Progressively, the complexity of the robot's behaviours can be increased, along with the child's progress in coping with more complex social interactions, involving, ideally, both the robot and other children or adults.

1.1 Motivation

Through play, children can develop skills in various fields such as social, communicative, imaginative, logical and abstract skills (Chaillé and Silvern, 1996; Piaget, 1945; Boucher, 1999). Play is also a medium for self-expression. Research has shown that children with autism have a potential for play (Boucher and Wolfberg, 2003), but, unfortunately, often encounter obstacles to actualize this potential (Dautenhahn and Werry, 2004; Fritz, 1989). A challenging goal is to find a way to facilitate their access to various play situations despite this difficulty. The work presented in this thesis focuses on facilitating play between children with autism and an autonomous robot. The ultimate goal is to enable children to experiment with a variety of play situations, ranging from solitary (dyadic play with the robot) to social situations of play in a triad, with both the robot and another person, and from basic tactile interaction to more symbolic situations of play, involving possibly pretend play. It is hoped that children would develop appropriate skills, while involved in specific play situations. For instance, children may progressively understand some notions of causality: e.g. 'when activating a sensor, the robot shows a reaction'. Another example is the notion of chronology that is present in several games, such as 'hide and seek' or 'story telling'. More generally, if a child can play symbolically, she may develop imagination skills. And if she manages to play in a triad with both the robot and another person, she may deal with communicative and social skills. An additional challenge towards this goal is to be able to adapt to each child's specific needs and abilities. Autism is a spectrum disorder and the abilities and needs can vary enormously. In addition, we should try to make play with the robot as enjoying and fun as possible; thus, the personalities of the children, their liking and disliking, should also be taken into account. Consequently, this challenging goal of facilitating play between children with autism and an autonomous robot should be addressed along three axes:

- *The design of the play sessions*: How should the experimental sessions be designed in order to adapt to the children's needs and abilities? What should be the role of the experimenter?
- *The recognition of the interaction styles*: How could the robot recognize the play styles of the children in real time, in order to autonomously adapt to them?
- *The adaptation of the robot*: How could the robot adapt to the children' s needs, abilities and preferences so that fun and enjoyment are favored, and so that the children learn from their interaction with the robot?

1.2 Methodology and Practical Effort

The research presented in the thesis is really multidisciplinary, involving important challenges, particularly in robot-assisted play and in pattern recognition. For this work, I had to read about various areas: child development, developmental robotics, developmental psychology, autism research, therapy, education, assistive technology, robot-assisted play, social robotics and machine learning.

Concerning the research on pattern recognition, little research had been done before in the recognition of interaction styles in real time. I investigated various techniques. The prototypes for the interaction styles had to be generated under controlled laboratory conditions, in interaction with the robot which was a very time consuming process[1]. This notably implied to take into account the possible artifacts due to real data. The testing of the techniques was done in two steps: first, I ran a test with data generated under laboratory conditions, and, when successful, I applied a further test with data generated in real situations of play between children with autism and the robot. In the end, I developed a novel computational method for time series analysis based upon the existing Information Bottleneck method.

Concerning the long-term experiments in school, I should mention that it requires lots of practical effort. Firstly, it takes time to make contacts with schools. Schools are very busy, and, in terms of logistics, it is not necessarily easy to be attributed a room where the experiments can happen every week. In the context of this work, the school already knew about the Aurora Project beforehand and it was therefore faster to get the agreement to conduct the experiments in the school. Besides, for each child, the parent's written consent form is required before starting the trials, as well as a CRB check plus ethics approval from the University Ethics Committee. Furthermore, a documentary was broadcasted on 25^{th} August 2008 on my research on the German channel 3SAT. For this filming, I had to get firstly the authorization from the school to film in the school and, secondly, specific consent forms from the parents, so that the children could be filmed by the journalists. In particular, the journalists filmed play sessions that I conducted with the children. The documentary is now available online, on the internet site www.3sat.de, Rubric 'Nano', with the title 'Roboter soll Kommunikation autistischer Kinder foerdern'.

In terms of experiments, I conducted play sessions for more than a year on a weekly basis (15 months, including holiday time). During the first four months, six children participated in the trials. After the four first months, three other children

[1] There was no preexisting standardized databases which could have been used.

joined the play sessions; thus in total, nine children were then involved in the play sessions. Each play session was video recorded for a later analysis in details, which represents a considerable amount of time. It should also be mentioned that the play sessions require additional skills than scientific or technical ones. Extreme concentration, organisation, adaptation, observation, empathy and listening are required at all time.

1.3 Contribution to Knowledge

The research presented in this thesis contributes to several areas:

Robot-Assisted Play: I propose and experimentally test a new methodological approach of how to design, conduct and analyse robot-assisted play.

This approach is inspired by non-directive play therapy. The experimenter participates in the play sessions. The child is the main leader for play. However, under specific conditions that are precisely defined, the experimenter intervenes proactively in the play situations. This intervention aims at i) preventing the child from entering or staying in repetitive behaviours; ii) providing bootstrapping that helps the child reach a situation of play she is about to enter and iii) asking the child questions dealing with reasoning or affect related to the robot.

This methodological approach focuses on three intertwined dimensions that are play, reasoning and affect. The analysis of the play sessions relies on these three dimensions with a qualitative analysis. Each dimension is analysed according to a list of precise criteria.

This approach is tested with a long-term study with six children with autism in school and proves capable to adapt to each child's needs and abilities. Each child makes progress in at least one dimension (Play, Reasoning or Affect). In particular, children experiment with various situations of play that address specific aspects of play such as the use of causality/reaction, social play, chronology, symbolic and pretend play.

Pattern Recognition: I design a novel and generic computational method for the automatic recognition of human-robot interaction styles in real time. This method is experimentally tested with data that have been generated under laboratory controlled conditions in real interactions with a robot and with data that have been produced during child-robot interaction in school where children were not instructed how to play. Results show the capability of the method to recognize short-term and mid-term time scale events correctly. The recognition is made with a very small delay. This method is entirely generic for application with socially interactive robots. It forms a step towards socially adaptive robots.

Human-Robot Interaction: I demonstrate a proof-of-concept system for an adaptive robot responsive to different styles of interaction in human-robot interaction. I test its impact through a study with seven children with autism.

The adaptive robot uses the Cascaded Information Bottleneck Method for the real-time recognition of the interaction styles. I design a schema for the robot's adaptation that is based on the principle of rewarding well-balanced styles of interaction. In addition, the robot attempts to encourage children to engage in interaction if they are disengaged. Results from the trials conducted show the positive impact of this adaptive robot on the children's play styles.

Developmental Robotics: I contribute to the understanding of social behaviour and adaptation which are key topics in developmental robotics, inspired by research on child development and autism therapy.

Autism Therapy: I conduct a study that potentially may be developed towards a new method in autism therapy. At the moment, a roboticist is needed to deal with the issues implied by the use of a robot. In future, play therapists could apply the new methodological approach in robot-assisted play that is presented in this thesis.

1.4 Publications resulting from this work

Several publications resulted from this work:

- One journal paper (to appear): François et al. (2009);
- Two conference papers: François et al. (2008b, 2007);
- Two technical reports: François et al. (2008a,c);
- An abstract for a talk: This talk will be given on December 1^{st} 2008, in Coventry University Technocentre, at the Conference RAatE 2008 (Recent Advances in Assistive Technology and Engineering). The abstract is entitled *'Robot Assisted Play: Detecting Interaction Styles of Children with Autism Playing with a Zoomorphic Robot'* (authors: François, D., Dautenhahn, K., and Polani, D.).

1.5 Outline of the thesis

Chapter 2 begins with background on Human-Robot Interaction, a fairly recent and broad area of research, including Child-Robot Interaction and Robot-Assisted Therapy and Education, for which the current state of research is presented.

Chapter 3 is dedicated to Play and Autism: Play is a vehicle for learning and for experimenting with a multidisciplinary range of skills. Children with autism have a potential for play but often encounter obstacles to actualize their potential. Several approaches in psychology coexist with respect to learning and cognitive development. More recent ones underlining the role of social interaction in the process of learning are particularly highlighted. This leads to the formulation of the research questions that structure this research on facilitating play between children with autism and an autonomous robot.

In Chapter 4, I introduce a new methodological approach for designing, conducting and analysing robot-assisted play. This method is inspired by Non-Directive Play Therapy and the role of the experimenter is clearly defined. In this approach, the child is the main leader for play, but the experimenter participates in the play sessions and can intervene under precise conditions in order to facilitate or bootstrap the access to higher levels of play. This method is tested through a long-term study with six children with autism. The results are analysed with a specific methodology that focuses on three dimensions: Play, Reasoning and Affect.

Chapter 5 addresses the recognition of the interaction styles in real time. This work was mainly conducted in parallel to Chapter 4. Chapter 5 starts with the definition of basics concepts and a presentation of the related work. Then, it presents an early approach for the online classification of human-robot interaction styles based on Self-Organizing Maps which shows a good accuracy, but requires important hand-tuning to obtain acceptable delay in the recognition process. Two other techniques are therefore successively investigated, the Fisher Linear Discriminant Analysis and Clustering by compression; however, both of them fail in separating the classes. It leads me to design a novel computational method for the recognition of human-robot interaction styles, the Cascaded Information Bottleneck Method, that is presented in a third section of the chapter. This method extends the existing Information Bottleneck Method developed by Tishby et al. (1999) by providing a cascade of bottlenecks trained successively. Each bottleneck has the same amount of bottleneck variables and a measure to extrapolate cases that have not been seen during the training phase is introduced. The method is evaluated with both data generated in laboratory conditions (training set and cross-validation) and data obtained from interactions in the school, between children and the robot.

Chapter 6 was informed by Chapter 5 and introduces the notion of an adaptive robot. The main purpose of the real-time recognition of the interaction styles is actually to enable the robot to adapt its behaviour according to the children's play styles. Such a robot is called an adaptive robot (in comparison with a 'reactive' robot which reacts the same way whatever the play styles of the children are). Firstly, the architecture underlying the adaptation process is detailed and, secondly, results from a study investigating the impact of the adaptive robot on the play styles of children with autism are reported.

Chapter 7 discusses the contribution of this thesis and concludes it. Chapter 8 draws directions for future work.

Several Appendices are enclosed in this thesis. Appendix A provides additional figures for Chapter 6. Appendix B contains details about the children, in terms of their age and their level of autism. Appendix C explains the methodology to tailor the behaviours of the robot according to the children's needs, abilities and preferences. Appendix D shows the social story used for one child, in order to help this child understand better how the play sessions proceeded. Appendix E lists the publications resulting from this research. Appendix F presents the media coverage on this work. Finally, implementation details can be found on the CD attached.

Chapter 2

Robot-Assisted Therapy and Education

2.1 Human-Robot Interaction

Robot-Assisted Therapy and Education is one of the various domains of application of Human-Robot Interaction (HRI), a multidisciplinary research area, which requires the collaboration of researchers from different fields of expertise such as: psychology, social sciences, cognitive science, linguistic, artificial intelligence, mathematics, computer science, robotics, engineering and human-computer interaction (Dautenhahn, 2007a; Goodrich and Schultz, 2007).

HRI is a growing novel research field. It started to emerge in the mid 1990's, and numerous HRI studies have been conducted since, addressing a great diversity of applications, ranging from space applications to assistive robotics. Recently, Goodrich and Schultz (2007) presented a review on HRI, focusing particularly on presenting the main challenges in the various application domains and on extracting the first accepted practices that govern the field. They define an HRI problem as the one of "[understanding] and [shaping] the interactions between one or more humans and one or more robots". The goal is to make those interactions beneficial in some sense. They also identified five main attributes on which the designer can have an impact, which are the following: 1) level and behaviour of social autonomy, 2) nature of information exchange, 3) structure of the team, 4) adaptation, learning and training of people and the robot, 5) shape of the task.

HRI studies can adopt various approaches, which are not mutually exclusive, namely a robot-centred approach, a robot cognition-centred approach or a human-centred approach (Dautenhahn, 2007b). A robot-centred approach favors the view of the robot as an autonomous entity that is pursuing its own goals and interacts with people in order to pursue a precise goal. A human-centred approach focuses very much on humans and on how the robot can fulfill a specific task in a way that is acceptable by humans. A robot cognition-centred approach underlines the capability of the robot to make decision on its own and solve problem autonomously.

Dautenhahn (2007b) emphasizes the necessity of focusing on social skills for robots, which requirements may vary depending on the domain of application. Dautenhahn (2007b) actually defines a spectrum of requirements for robots' social skills, ranging from none for a remotely controlled robot, or a robot operating in a spatially-temporally separated environment from humans, to essential, for robots in nursing care, rehabilitation or therapy (e.g. autism therapy) and for robot companions in the home. Besides, Dautenhahn (2007b) provides a conceptual space of HRI approaches, where she addresses the notion of social robots and discusses some of its related concepts, e.g. 'socially evocative' (Breazeal, 2002, 2003), 'socially situated' (Fong et al., 2003) 'sociable' (Breazeal, 2002, 2003), 'socially intelligent' (Dautenhahn, 1998), and 'socially interactive robots' (Fong et al., 2003)).

Importantly, Dautenhahn (2007b) provides an enlightening grid for categorizing HRI studies. This grid is based on four criteria, each of them being possibly derived in a spectrum of intensity:

- 'Contact with humans' (spectrum ranging from 'none', to repeated, long-term and physical)
- 'Robot Functionalities' (spectrum ranging from limited and clearly defined functionalities to open, adaptive and shaped by learning)
- 'Role of the robot' (spectrum ranging from a machine tool to an assistant, a companion or a partner)
- 'Requirements of social skills' (spectrum ranging from no social skills required to essential)

2.2 Child-Robot Interaction

Numerous studies have been conducted in child-robot interaction research. Important research questions addressed are whether and how a robot could contribute to the social and cognitive development of the child, and how, under certain conditions, it could be used as a medium for social interaction. The main challenges are to identify the natural means by which children interact with robots, and to encourage long-term interaction by an appropriate design of the robots' features and capabilities and relevant scenarios.

Tanaka et al. (2006, 2005) lead an ongoing long-term study on child-robot interaction with a focus on the context of dancing. The main purpose of this ongoing study, named "Ruby Project", is to find principles for realizing long-term interaction between children and a robot. In the first year of the project, typically developing children, from age 18 to 24 months, encountered the Sony humanoid robot QRIO at school, in the context of dancing. Off-line analysis of the interactions between the children and QRIO showed that the children tended to progressively adapt their behaviour to the robot's characteristics. Besides, a further analysis on 45 successive sequences of interaction of those children with QRIO spanning 5 months (Tanaka et al., 2007) showed that those children tended to progressively consider QRIO as their peer rather than as a toy: The way they touched the robot was reorganised so that, in the end, the distribution of their touch towards the robot was converging to the one observed when they were touching their peers. This study relies mainly on design by immersion, which means here that scientists, engineers and robots are present in the everyday life environment of those children while shaping both hardware and software and addressing scientific questions early in the development process (Movellan et al., 2007). For instance, this design by immersion has enabled Tanaka et al. (2006) to highlight some basic necessary units for long-term human-robot interaction, respectively "sympathy" between the human and the robot and "variation" within the interaction styles.

In a different study, the potential of communication robots for elementary schools has been investigated with the robot Robovie (Kanda and Ishiguro, 2005). The focus of this study was on two different aspects: the first one addressed the role of the robot as a tutor for children's learning (Kanda et al., 2004) and the second one dealt with how it could be possible to encourage long-term child robot interaction (Kanda et al., 2007). The first aspect was addressed through a two weeks trial where the robot demonstrated positive effects for motivating children to learn foreign languages at school. Children showed statistically significant improvements in their listening tests which were linked to their interaction patterns with Robovie. Nevertheless, children tended to get bored by the robot after a week of interaction. This illustrates that, in order to enable long-term human-robot interaction, the robot should have additional features and capabilities, including some novelty. Besides, Kanda et al. (2007)

reported on a long-term study they conducted in school with 37 children. Those children could play freely with the robot on a school-daily basis for about 30 minutes per session during 2 months. The robot was placed in a classroom and equipped with two social communication abilities:
1) The capability of forming long-term relationship through 3 main principles of behaviour designs, which were the following: i) the robot called children by their names (it uses RFID tags), ii) the robot adapted its behaviour to each child on the basis of a pseudo development analysis: the more a child played with Robovie, the higher the diversity of Robovie's behaviours was, iii) Robovie told its personal matters to the children who had interacted with it for a sufficiently high period of time.
2) The capability of evaluating the friendly relationship among the children: this evaluation relied on the principle that people who spontaneously behaved as a group were friend. Then, the robot estimated for each child in the group the relationship with the child and the other children in the group.
Results showed that three main successive phases could be identified: 1) great excitement, 2) stable interaction to satiation and 3) sorrow because the robot was soon leaving (Kanda and Ishiguro, 2005). These experiments illustrate very well the challenges of designing a robot and scenarios for enabling long-term human-robot interaction. Communication robots for elementary school is a particular application of the general research question on how robots can take part in human daily life by playing the role of peers, playmate or partners.

2.3 Robot-Assisted Therapy and Education

Rehabilitation Robotics (also called Assistive Robotics) is a main application of HRI. Rehabilitation robotics is the use of robots for people with special needs. It embraces physical and cognitive impairment and also the effect of aging.

2.3.1 Robotic devices as tools for physical rehabilitation

In this case, the robot is the mean by which the human can recover some mobility or agility, lost by the physical impairment. Two examples are a robotic wheelchair (Yanco, 1998, 2001; Bailey et al., 2007), and a robotic arm mounted on a wheelchair (Hillman, 2003; Hillman et al., 1999; Hagan et al., 1997; Kwee et al., 1989). The latter may assist a person with physical impairments in i) eating and drinking, ii) personal hygiene, iii) mobility and access (e.g. opening doors) and iv) tasks related to reaching and/or moving objects (Hillman et al., 2001).

A third example is the use of robotic platforms for visually impaired people (e.g. Kulyukin et al. (2006); Lacey et al. (1999); Lacey and Dawnson-Howe (1998)). Elderly people with visual impairments are potential users of such robotic platforms and the challenge of providing both i) a physical support for the person walking and ii) obstacles avoidance has been addressed in e.g. the PAM-AID project (Lacey and Dawnson-Howe, 1998; Lacey et al., 1999).

Another use of robots in rehabilitation is for post-stroke rehabilitation (Kahn et al., 2001; Kwakkel et al., 2008; Burgar et al., 2002; Mahoney et al., 2003; Matarić et al., 2007). Stroke is an important cause of severe disability and can lead, in particular, to difficulties in accomplishing everyday movements activities. During the critical post-stroke rehabilitation, it is possible to improve this loss of function which is called "learned disuse". Research in this specific application domain splits into two different perspectives: 1)"hands-on" rehabilitation, during which the robot actively helps the patient repeat prescribed movements of a specific limb (e.g. Kwakkel et al. (2008), Burgar et al. (2002) and Mahoney et al. (2003)); 2)"hands-off" rehabilitation (Gockley and Matarić, 2006; Eriksson et al., 2005), during which the robot plays a social role rather than a physical role, described by Matarić et al. (2007) as "socially assistive robotics". In hands-off rehabilitation, the patient performs the active exercises by himself/herself, without physical help from the robot; rather, the robot interacts socially with the patient and encourages him/her in the rehabilitation exercises. Note that '2)' shifts from the pure notion of 'tools' for physical rehabilitation and steps towards robot-mediated rehabilitation through human-robot social interaction.

2.3.2 The robot as a peer or a playmate for therapy and education

Long-term studies like some studies conducted with the seal robot Paro (Shibata et al., 2005; Marti et al., 2005) have shown that specific everyday life situations exist in which human-robot interaction can have a positive effect on well being of human beings. The robot can play the role of a peer or a playmate which stimulates cognitive and/or physical capacities and may be a medium for social interaction. Several main domains of application are currently very actively addressed, respectively, 1) the role of human-robot interaction in therapy for physical, cognitive, communicative or social impairments, e.g. cerebral palsy and autism, and 2) human-robot interaction for elderly people, either from a pure interactive point of view, to stimulate aged people to interact, or on an assistive approach, whereby the robot could possibly play the role of an assistant or a companion in order to enable aged people to live as long as possible independently in their homes. The following paragraphs illustrate successively those domains of application.

Recent studies have investigated whether and how a robot could be introduced in therapy protocols to enhance or accelerate the benefits of a therapy. A long-term study with the seal robot Paro has shown that specific situations of human-robot interaction can be a significant factor of performance in therapy. This study designed engaging rehabilitation activities that would combine physical and cognitive rehabilitation (Marti et al., 2005). The participant, a child with severe cognitive and physical delays, interacted with Paro on a weekly basis over three months as follows: Paro was introduced in the context of the Bobath protocol and played the role of a playmate. The Bobath protocol is a method used for the rehabilitation of physical functional skills (Bobath protocol, 2008; Knox and Evans, 2002). It consists in training the child to acquire basic behavioural primitives of movements and positioning such as head control, grasping or equilibrium control during a movement or in case of a fall. In the Bobath protocol, diverse toys are used to engage the child in the therapy process. Here, Paro robot was used alternately in a passive way (i.e. like a simple toy) and in an interactive mode. The activities were designed to be as similar as possible to the ones used for the Bobath protocol. Results showed that the introduction of Paro in the Bobath protocol may have strengthened the efficiency of the method for this specific child by facilitating his active engagement in the rehabilitation exercises.

Several research laboratories address the particular case of autism and how child-robot interaction could possibly help children with autism experiment with social skills (Dautenhahn and Werry, 2004, 2000; Dautenhahn, 2007b; Kozima et al., 2005); for instance, children with autism may experience non-autistic behaviours while engaged in play with a robot (Stanton et al., 2008). Chapter 3 is dedicated to autism and play, and Chapter 4 will provide more details on related work in robot-mediated therapy and education for children with autism (refer to Section 4.3).

The potential use of robots for elderly people has been particularly addressed with studies with the seal robot Paro. Shibata et al. (2005) conducted a long-term study whereby Paro was introduced on a daily basis into the everyday life of some elderly people in two different institutions, in one of them for a daily duration of 20 minutes over 6 weeks and in the second one for 1 hour over more than a year. Elderly people were free to interact with the robot. Results showed that, on average, interacting with Paro improved the mood state of the participants and made them more active and more communicative with each other as well as with the caregivers.

The role of the robot as a cognitive robot companion is addressed by the Cogniron project[1] (Cogniron, 2008; Syrdal et al., 2008; Otero et al., 2008). The purpose of this European project is 1) to design the cognitive functions of a robot which would be able to serve humans in their daily life, as assistant or companions, and 2) to study and develop the capability of a continuous learning and education scheme of the robot which would enable it to mature to a true companion. A direct application of this project could be to provide companions in homes of elderly people to enable them stay longer at home.

2.4 Methodology

A diversity of approaches coexists in this new field, which is still in its infancy. Qualitative and quantitative analysis are used and experiments can be short-term or long-term studies depending on the research questions that are addressed. At this stage of development it may be valuable to be open to such a diversity of approaches, allowing the exploration of many paths to address HRI challenges. However, as Dautenhahn (2007a) underlined, it might also play a role in the difficulty to reproduce results from experiments.

2.4.1 Safety

The safety of the participants is the main priority. Any trial must be first tested under laboratory conditions in order to test the safety and the reliability of the experimental conditions, including, in particular, the robot's functionalities. Besides, the experimenter must have Ethics approval. Here, the required ethics approval was obtained from the University of Hertfordshire Ethics Committee.

2.4.2 Short-term / Long-term studies

Depending on the research questions addressed, the trials can be run short-term or long-term. At the very extreme, a short-term interaction could be a zero-acquaintance (Dautenhahn, 2007a). It means that the participant, who has not encountered the robot before, interacts with the robot for one session only. At the far other extreme, participants have the opportunity to progressively build a relationship with the robot, or at least to get familiar with the robot. This may result in progressive changes in the way he/she interacts with it. Such long-term studies are necessary for the design

[1] The Cogniron project focuses on a large range of possible users, not only elderly people; it does consider the robot as a companion, that would assist in daily task in a social interactive way.

by immersion of robots enabling long-term interaction (thus trying to counteract the possibility of getting bored by limited functionalities of the robot) and are also particularly relevant to therapeutic and assistive robotics applications.

2.4.3 Behavioural based data analyses

A main technique used in HRI trial analysis is the decomposition of behaviours previously videotaped during the trial, according to a predefined grid of criteria. An inter-rater reliability test should be applied which, in total, makes this analysis time consuming. But still, it remains widely used in HRI. In particular, Kahn et al. (2003) have developed a manual for the coding of child-robot interaction with a robotic pet. They addressed both behavioural and reasoning issues. We will refer to this manual in Chapter 4.

Moreover, it is useful to develop objective quantitative measurements of the interaction, which would be complementary to video analysis and provide additional insights on the interaction. The second part of the thesis focuses on such automatic classification of the interaction, where related work will be detailed. In particular, Scassellati (2005a,b) investigated the development of objective measurement of typical factors of autism through a robotic platform, which would provide a complementary insight to the traditional technique of diagnosis.

2.4.4 Self-reporting through questionnaires or semi-structured interviews

Another useful input is the analysis of questionnaires or semi-structured interviews. However, it should be noted that for our precise study with children with autism, such techniques can not be directly used. Questionnaires can be given to the parents or the teachers but trials in themselves must be evaluated through videos or objective quantitative measurement of the robot's sensor data.

2.5 Summary

In this chapter, we have introduced some background on Human-Robot Interaction (HRI), a fairly recent and pretty broad research area. We have presented some core concepts and methodological aspects, such as the grid proposed by Dautenhahn (2007b), based on four criteria to characterize HRI studies. Each criterion, namely, 'Contact with Humans', 'Robot functionalities', 'Role of the robot' and 'Requirements of social skills', can be derived in a spectrum of intensity.

Further to this, we have presented the state of research in Child-Robot Interaction, which principally addresses the questions on whether and how, under certain conditions, a robot could be used as a medium for social interaction. The main challenges here are to identify the natural means by which children interact with robots, and to encourage long-term interaction by an appropriate design of the robot's features and capabilities and relevant scenarios.

We have then presented Robot-Assisted Therapy and Education, which is a domain of application of HRI and can be addressed in the following ways:

- *Developing robotic devices as tools for physical rehabilitation*: Four examples are a robotic wheelchair (Yanco, 2001; Bailey et al., 2007), a robotic arm mounted on a wheelchair (Hillman, 2003; Hillman et al., 1999), robotic mobile platforms for visually impaired people (Kulyukin et al., 2006; Lacey et al., 1999) and robotic platforms for post-stroke rehabilitation (Kahn et al., 2001; Kwakkel et al., 2008; Burgar et al., 2002; Mahoney et al., 2003; Matarić et al., 2007).

- *The introduction of the robot as a peer or a playmate for therapy and education*: It embraces physical, cognitive (Marti et al., 2005), communicative and social impairment (Aurora, 2008; Dautenhahn and Werry, 2004; Stanton et al., 2008; Kozima et al., 2005), as well as the use of robots for elderly people, both from a pure interactive perspective (Shibata et al., 2005) and from an approach of assistance in daily tasks (Cogniron, 2008).

Finally, this chapter reported on methodology in Human-Robot Interaction studies, addressing in particular safety issues and classical techniques used for the analysis of HRI trials.

Chapter 3

Autism and Play

3.1 Autism

Autistic Spectrum Disorders can appear at various degrees and refer to different skills and abilities (Powell, 2000; Jordan, 1999). Detailed diagnostic criteria for autistic spectrum disorders are provided in the Diagnostic and Statistical Manual of Mental Disorders (Association, 1994)[1].

3.1.1 Main impairments

The main impairments highlighted by the National Autistic Society (2008) are:

- **Impaired social interaction:** Difficulties to make sense of a relationship with others, difficulties to guess or even understand what the other's intentions, feelings and mental states are.
- **Impaired social communication:** Difficulties with verbal and or non verbal communication (for example, difficulties to understand facial gestures).
- **Impaired imagination:** e.g. Difficulties to have imaginative play.

As a consequence of the above impairments, children with autism often choose a world of repetitive patterns and, for instance, often engage in playing in a repetitive way. Different theories try to explain why those with autism prefer to live in a predictable world. One of them, the Theory of Mind (Baron-Cohen, 1997) explains that

[1] DSM-IV (Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition) was published in 1994 and is the last major revision of the DSM.

children with autism tend to have difficulties in identifying mental states of others, i.e. in having a representation of what others may think. More precisely, it concerns a full range of mental states (e.g. beliefs, desires, intentions, imagination, emotions) that cause action (for a description of some of the manifestations of this impairment, please refer to Baron-Cohen (2001)). Consequently, it can be very hard for them to understand social interactions. In addition to this theory, they often lack the capability to generalize (Baron-Cohen, 1997) and, as a consequence, to classify entities. For example, two mugs may often appear to them as two completely distinct and uncorrelated items: it can be very hard for some of them to extract common properties of objects and then categorize according to these properties. What happens to objects, in this way, also happens to experience in life. It is most of the time impossible for those with autism to generalize situations they encountered previously. Furthermore, children with autism can distinguish between a human and an object but, in the most severe cases of autism, their behaviour towards humans may have elements of how they treat objects (Hobson, 2002). Moreover, since human beings are very complex with all their essential expressiveness, they tend to prefer interacting with objects which are simpler. This could be partly explained by a theory focusing on the sensory dimension (Williams, 1996). The latter suggests that children with autism are very sensitive to stimuli and that they prefer remaining in predictable environments so that they can avoid being hurt by overpowering sensory stimuli (Gillingham, 1995).

3.1.2 Autobiographical accounts

Of course, autobiographical accounts do not replace research on autism, which is the proper source of theory and knowledge on autism to build our research upon. Nevertheless, autobiographical accounts are very inspiring. Below is a summary of particularly relevant points to my research. Note, both of the persons whose autobiography is mentioned here are high-functioning or have Asperger Syndrome.

Daniel Tammet's autobiography (Tammet, 2006). Daniel Tammet, 26 years old, has both Asperger Syndrom and synaesthesia. He managed to overcome some of his impairments and to live in the world in which typically developed people live. He practises a job in pedagogy where he provides e-learning for foreign languages. Moreover he has a partner and close friends. Daniel Tammet provides an autobiography, with in depth description of his specific aptitudes and his internal states or feelings. He also directly contributes to research in autism by agreeing to become a subject of study for research in autism.

In his autobiography, Daniel Tammet implicitly refers to weak central coherence theory by explaining the fact that as a child, he could not "read between the lines": for instance, he could not infer the implicit meaning of two logical correlated sentences; he provides a metaphor to describe this impairment: "It is like joining the dots in a children's coloring book and seeing every dot but not what they create when joined together". Another point is his wishing, already when he was a child, to encounter friendship at such a point that he invented himself an imaginary friend, an old woman called Anne. As he said "People with Asperger syndrome do want to make friends, but find it difficult to do so". And, referring to the imaginary friend: "Looking back, Anne was the personification of my feelings of loneliness and uncertainty". Moreover, Daniel Tammet had the feeling of being different from typically developing children. Again, through his imaginary friendship, he asked Anne why he was different: "Once I asked her why I was so different from the other children but she shook her head and said that she could not say. I worried that the answer was terrible and that she was trying to protect me, and so I didn't ask her again".

Temple Grandin's autobiography (Grandin, 1986). Like Daniel Tammet, Temple Grandin mentions the consciousness at some point of 'being different'. One trait of autism Temple Grandin was particularly suffering from is oversensitivity to stimuli and irregularity of the reactions to stimulations. For instance, she writes about her paradoxical reactions to sound, which made her, at some point, not react to people speaking directly to her but react to, or even be annoyed by other sound stimuli (e.g. bird's song). Moreover, she was oversensitive to tactile stimulation too: at a tactile stimulation, she could suddenly overreact. But it also appeared to her that she could decrease her oversensitivity and anxiety through other specific tactile stimulations (like for instance when she wrapped up in a blanket). Later, she designed a therapeutic device relying on the principle of such tactile pressures.

Temple Grandin is now a designer of livestock handling facilities and a Professor of Animal Science. She points out some persistent difficulties in the everyday life. She indeed still finds it hard to understand human interaction; therefore, she tends to rely on a library of experiences referring to various social situations and the corresponding reactions of people to be able to predict what may happen in a current interaction. In the same way, she tends to rely on explicit social rules concerning behaviours that are usually intuitive for adults who do not have autism. For instance, one of the rules is to smile back at someone smiling to yourself. Another rule is to remember to look interested when someone is talking to you.

In both autobiographies, the sustained effort made by the child, on the one hand, and, on the other hand, the sustained education and the continuous emotional support from the family, the friends, the teachers and the carers, emerged as important key factors of success in dealing with impairments. Moreover, Daniel Tammet's parents seemed to have always encouraged their child to encounter new social situations, e.g. through play with his brother and sisters, holiday with family of friends or activities at which he was particularly good at such as chess in a chess club.

3.2 Play

3.2.1 What is play?

There is no precise definition of play partly because play covers a wide range of activities (Chaillé and Silvern, 1996). There is no clear boundary between play and not play. In fact, many disciplines deal with play and investigate this field differently according to their specialty. Since play can be investigated from different perspectives, various classifications can be found in literature. Some focuses on the result of play. For example, play may take two different forms with reference to learning: one is active and the other is passive learning. By this, it is meant that the child can either play in a way that he/she will try to learn a lot or, on the contrary, just play without efforts for learning. The first class, active learning, requires the child's interest, attention and mental activity so that knowledge construction can take place (Chaillé and Silvern, 1996). For example, if a child is just manipulating an object without applying mental activity, this will be classified as not active learning. On the contrary, if the child manipulates the object, with the intention of trying to understand what it can be used for and how it can be used, then this will be part of active learning.

Another way of classifying play is trying to differentiate different kinds of activities in play itself, like purely motor skills or abstract representation. A famous taxonomy about play is given by Piaget (1945). He differentiates four kinds of play: Practice play, symbolic play, games with rules and constructions. A different classification, provided by Boucher (1999), distinguishes at least four classes which are: i) sensory motor play, ii) manipulative and exploratory play, iii) rough-and-tumble play and active physical play and iv) social play (see Fig. 3.1). This classification is particularly interesting for this research study because it mixes the notion of exploration with the idea of social interaction.

Kind of play	Specificities
Sensory motor play	“teaches young infant about their own bodies and about objects in the immediate environments”
Manipulative and exploratory play	“teaches the older infant more about objects and their properties, and about how we can influence the world around us”
Rough-and-tumble play and active physical play	“teaches the toddler and preschooler gross motor skills, and provides experience of whole body interaction with others and with objects in the environments”
Social play	“teaches children about social relationships and how to engage in them, as well as about cultural norms of the society the child is growing up in.”

Figure 3.1: A classification of Play given by Boucher, quotations from Boucher (1999).

3.2.2 Children with autism and play

A relative potential for play Children with autism are able to play but the nature of their play may be described as restricted. Indeed, according to the American Psychiatric Association, “a lack of varied, spontaneous make-believe play is a defining feature of autism” (Association, 1994). In other words, children with autism often play in a repetitive way, which means that there is a lack of generating new behaviors during play. This can be correlated to the fact that children with autism prefer predictable environments. The other main idea emphasized in the previous quotation points out the important difficulties children with autism encounter in symbolic play (Chaillé and Silvern, 1996): only a few of them are able to draw properly (within a symbolic mode) or to make models (again, that may be classed as symbolic). So, children with autism lack some abilities for playing in the way that normally developing children do though they can play in a certain (autistic) way. Moreover they sometimes have a wrong perception of what play is since for example, they often consider obsessional activities as play (Boucher and Wolfberg, 2003). This introduces another notion, which is that children with autism have some obstacles for actualizing their potential for play (Boucher and Wolfberg, 2003).

Obstacles for developing their potential for play Different possible obstacles have been identified. Among them are impairments in socioemotional intersubjectivity, impairment in joint attention and impairment in Theory of Mind (Baron-

Cohen, 1997). These impairments damage interaction in general and, more specifically, must imply a lack of spontaneous and social reciprocity during play. The disability in perceiving the coherence of categories and concepts can besides be a reason why children with autism perceive objects in their parts and not as the wholes which is part of a weak central coherence theory (Fritz, 1989; Dautenhahn and Werry, 2004). But today, causes for impaired play are still not very well understood. These causes can vary for different children, depending notably on the personality of the child and her past experience of play. One last point to underline is that it may be very useful to encourage and support the child with autism while he/she is playing though the way in which that encouragement and support is operationalised is a matter for debate. Certainly, the typical way in which caregivers interact with typically developing children may not be the most appropriate model in autism.

3.2.3 Why focus on play?

Active education Through certain kinds of play, children can construct some understandings. Here, understanding means active construction of meaning. Children can arrive at understanding by creating hypotheses about items and events that they find interesting. They test hypotheses as they actively interact with the material and events in their environment (Chaillé and Silvern, 1996). Piaget uses the notion of active education to speak about the intentional process of constructing understanding. Active education involves four elements: interest, play, genuine experimentation and cooperation. Chaillé and Silvern (1996) argue that play (and the context of play) already includes the three other components (i.e. interest, genuine experimentation and cooperation) which would mean that (in certain conditions at least) play leads directly to active education. In other words, play is a vehicle for learning in itself.

Multidisciplinary learning Play is a vehicle for learning in various fields. We shall give a few examples:

- *Logical memory and abstract thought*: For example, trying to understand relationships between entities is a basis of logical mathematical knowledge.
- *Communication skills*: Children learn communication skills by constructing a shared understanding of literal and non literal meaning. This suspension of believe, which enables a child to become another person virtually, facilitates the exploration of the language (oral and written). The child is actually not

stuck by the correction or the constraints from real life and can feel free to test new vocabulary, new parts of grammar and conjugation.

- *Social skills*: Through play, children can explore social roles and experiment with social issues of conflicts, trust or compromise. It is a very good place for the child to experiment with particularly difficult issues since he/she will not take it too affectively because children are just playing roles which have, a priori, no link with their real life.

It is important to note that the learning can really be interdisciplinary in the sense that the child is guided to make some connections between different areas. For example, by playing cards, the child has to be able to communicate and also to count. Thus, each time he/she is playing cards, he/she will deal and experiment with these two aspects and progressively develop skills.

Play for itself In the previous paragraph we justified the focus on play with the argument that play is a vehicle for learning. In this paragraph we will underline the fact that play is also an end in itself. Children usually enjoy playing (though this might not be the case in autism). Their pleasure and motivation seem to increase when they have the impression that they master the play situation (Boucher, 1999). Consequently, if we try to help children with autism master situations of play, they may have more fun playing which may contribute, even very modestly, to global happiness. Another argument is that play is a medium for self-expression (Boucher, 1999). When a child plays, she shows some parts of her personality and can also express personal feelings which she would maybe not show in ordinary life.

3.3 Approaches in learning in psychology

3.3.1 Behaviourist approach

The behaviourist approach is centred on the interaction between the human and his/her surrounding environment and focuses on the study of observable behaviours and the role of the environment as a factor influencing behaviours (Watson, 1913). The original behaviourist approach rejects the use of references to mental states (Watson, 1925), arguing that behaviours should be studied directly. It defends the idea that behaviours can be explained as the product of learning, which, in the classical behaviourist context, consists of conditioning. Conditioning is defined by Colman (2001) as "the process of learning through which the behaviour of organisms becomes dependent on environmental stimuli"[2]. It can take two forms: the classical conditioning, defended by Pavlov (1927), and the operant conditioning (which introduces the notion of 'reinforcement'), proposed by Skinner (1974). Skinner's ideas actually differed[3] from the original behaviourist approach established by Watson (1913, 1925). In particular, he redefined the notion of 'behaviour' in order to include everything that an organism does, which includes thinking, feeling and speaking (Skinner, 1974).

In contrast to the Watsonian pure behaviourist approach, the cognitive-behaviourist perspective, argues that mental processes are key factors in the behaviour (Tolman, 1932). Tolman (1932), who was one of the first contributors of this cognitive-behaviourist perspective, even defended the notion of goal-directed behaviour and used the expression "purposive behaviour". In this perspective, learning happens through meaningful behaviours.

3.3.2 Piaget's constructivist approach

The constructivist approach has been mainly driven by Piaget (Piaget, 1928), in reaction to the behaviourist approach. Piaget did not deny the fact that learning was fairly influenced by the environment. However, he defended the idea that learning was mainly due to mental processes. According to Piaget (1928), the environment actually plays an important role in the sense that it enables the child to experiment with new situations and thus develop new skills. But these skills can only be actualised at specific stages of the development of mental processes: Piaget (1928) indeed argued that children's cognitive development progressed through a series of stages that unfold in a definite sequence.

It should be underlined that Piaget mostly focused on the child's cognitive development and did not emphasize much the role of social interaction in the cognitive development (which may also be a reason why his classification of play focused on the lonely child behaviour and did not highlight social aspects of play (Piaget, 1945)). Piaget's schema of stages of development received a few critics. In particular, Isaacs (1930) was at first enthusiastic for Piaget's theories on the cognitive development of young children, but later criticised his schema[4] (Isaacs, 1930). She reproached him for using systematically the notion of 'maturation' without precautions, thus arriving to the point of explaining with the notion of maturation some phenomena which, in

[2] The conditioning can be considered as a form of associationism.

[3] Skinner branched off a new version of behaviourism, called radical behaviourism.

[4] These critics were formulated in (Isaacs, 1930) and Piaget answered to those critics in (Piaget, 1931).

fact, according to her, depend on experience (Isaacs, 1930). In addition, Isaacs (1930) criticized Piaget's tendency to rely on questionnaires, which, according to her, lead to stereotypical situations and interfere with the results. Instead, she argued in favor of her methodology, which relied on the observation of the children in their everyday life setting (i.e. Malting House School): according to her, the direct observation of the children and the cooperation between the children enabled a better objectivity in the observation of the thought of the child (Isaacs, 1930). Isaacs (1933)' approach to play is, besides, particularly relevant to our focus on robot-assisted play for children with autism. According to her, "Play is a child's life and the means by which he comes to understand the world he lives in" (Isaacs, 1933).

3.3.3 Vygotsky's influence (socio-constructivist approach)

Vygotsky introduced the importance of social interaction in child's development (Vygotsky, 1978). He stated that learning must take place within social interaction (Vygotsky, 1988). He defined the concept of zone of proximal development, as the zone of potential learning for an individual child at a given time. Concretely, it corresponds to what the child can possibly learn at a given time with the help of a peer or an adult. The helping approach offered to a child by an adult that is sensitive to that child's current zone of proximal development is sometimes referred to as the 'Vygotskian tutorial'. Unlike Piaget who stated that the child should wait to have reached a given stage of development to be able to develop new skills, Vygotsky believed that the most valuable learning for children was the one which was slightly in advance of their development. According to Vygotsky (1988), children actually need to learn in order to be motivated and this stretching of their possibilities is a boost.

3.3.4 Bruner's approach

Vygotsky (1896-1934) became internationally famous only in the 1960's and Bruner has been one of the first psychologists to bring some of Vygotskian's ideas in the United States. Bruner has contributed a lot in educational psychology and in particular, developed the notion of 'spiral curriculum', which is also of importance in play: spiral curriculum is the idea that children will revisit play materials and activities over the years, but then use them differently because their development has progressed. Bruner insisted on the importance of the medium of children's play, stating that the material to be learned is ideally the highest motivation for learning.

Bruner was besides involved in the creation of the cultural psychology. This approach considers that language, and by extension human thinking, come from the interaction between the individual and the culture in which he/she develops (Bruner, 1990). Bruner defines three modes of representation of the world; the first level is 'enactive', i.e. the action is linked to the manipulation of objects. The second level is 'iconic': the child uses pictures to make a representation of the environment. The third level, the 'symbolic representation', can be reached with the acquisition of language. At this level, Bruner argues that the culture will bring to the child the basis for his/her cognitive development. Unlike Piaget, Bruner considers that the environment and the culture play a preponderant role in the child's development (Bruner, 1983). Moreover, Bruner insists on the fact that education is an interactive activity between the child, the teacher and the environment and he insists on the role of the adults in the child's mastering of activities (Bruner, 1996).

3.4 Summary

In this chapter we have presented the main specificities of autism. Autism is a spectrum disorder which means that we should take into account the singularity of needs and abilities of each child with autism individually. Through play, children can experiment with a variety of skills from different fields. Particularly, they can develop social, communicative and imaginative skills, plus the ability to deal with more abstract concepts through symbolic or pretend play. Children with autism can play but often encounter obstacles to develop their potential. Through play, they may experiment with a multiplicity of skills, in particular, imaginative, communicative and social skills. Moreover, play is a medium for self-expression. We have then summarized different approaches with respect to cognitive development. We have shown that some approaches tend to focus more and more on the importance of the social interaction in the process of learning such as the Vygotskian approach which states that learning must take place within social interaction (Vygotsky, 1988). Unlike Piaget who stated that the child should wait to reach a given stage of development to be able to develop new skills, Vygotsky believed that the most valuable learning for children was the one which was slightly in advance of their development and took the approach that children need to learn in order to be motivated and this stretching of their possibilities is a boost. In the Vygotskian tutorial, the tutor (parents, carer, educator) can help the child to develop cognitive skills by extending its zone of proximal development.

3.5 Research Questions

In the preceding chapter, we have presented background on Human-Robot Interaction and Robot-Assisted Therapy and Education, showing, in particular, possible applications of the use of robots in rehabilitation. The present chapter has highlighted the main specificities of autism and the role of play in the cognitive and social development of children. We have shown that play is a vehicle for learning and that children may experience with a diversity of skills while engaging in play situations. Besides, typically developing children usually enjoy playing and, through play, children can also express themselves. Children with autism have a potential for play but often encounter obstacles to actualise this potential and we believe that, if we can facilitate them the access to diverse play situations, they would experience play skills and may develop some of them, in particular, communication, social and imagination skills. Moreover, because robots are much simpler than humans, we believe, in the context of the Aurora project, that robots can be a good medium for social interaction, in particular through play. Besides, in contrast to a stuffed animal, a robot can be embedded with specific behaviours, adapted to each child's needs and abilities, that may influence the children's responses during the interaction.

The insight of different approaches in psychology related to learning and cognitive development suggests a diversity of approaches that could be adopted in our specific context of play whereby we would like to help the child progressively reach higher levels of play, thus learn from play situations previously encountered. Vygotsky very much considered that social interaction is a key factor for learning (Vygotsky, 1988) and, according to him, the adult (a parent, a carer, and, in our context, the experimenter) plays a role and helps the child progress by extending its zone of proximal development. On the other hand, Piaget states that the child should anyway wait to reach a given stage of development to be able to develop new skills. In our context of facilitating play, we should carefully address the role of the experimenter: What should be the role of the experimenter in the play sessions? Should she stay apart during the play sessions? Should she, on the contrary, take part in the play sessions with the children? If we translate the Vygostkian approach in terms of play, the experimenter should actively help the child reach higher levels of play. In contrast, a Piagetian approach may not require such an intervention of the experimenter.

In addition to the role of the experimenter, the role of the robot should be defined. It should be investigated in which terms the robot can adapt to the specific needs and abilities of the children, in order to guide them towards more balanced interaction styles. What should be the robot's behaviour? Should the robot react differently according to specific play styles? How could it encourage children to engage in play?

Those questions related to the role and capabilities of the robot highlight further issues: if the robot must adapt to the play styles of the children, how can it identify those play styles in real time? What should be its schema of adaptation in order to influence the children's play styles positively and guide them towards more balanced tactile interactions, i.e. interactions that are neither too weak, nor too forceful, and that happen with an appropriate frequency (not too low, not too high)?

This leads to the following research questions:

- What approach for the play sessions could be adopted in robot-assisted play to enable each child with autism to progress according to his/her specific needs and abilities, that is, experiment with progressively higher levels of play and possibly develop play skills which could further help him/her cope with more complex situations of communication and social interaction, and develop imagination?
- How can a robot recognize the interaction styles of each child in real time?
- How could the robot best adapt to the children's needs and abilities? Can a robot that adapts to the play styles of the children in real time impact the behaviour of the children? Could it, in this way, help the children engage progressively in better balanced interactions?

Chapter 4

A Novel Approach in Robot-Assisted Play Inspired by Non-Directive Play Therapy

4.1 Introduction

Until now, research in robot-mediated therapy for children with autism has mainly explored the use of specific games, such as imitation (Robins et al., 2005b, 2004) or chasing games (Werry and Dautenhahn, 1999) and only recently started to involve the experimenter in the trials, qualifying his role as the one of a "passive participant". This chapter presents a different perspective on robot-mediated therapy, which is not "task-oriented" but rather draws inspiration from non-directive play therapy (Axline, 1946, 1947; Ryan, 1999; Josefi and Ryan, 2004) and which expands and clearly defines the role of the experimenter[1], who takes part in the trials and, whose role goes beyond the one of a 'passive participant'.

This method strongly encourages the child's proactivity and initiative-taking with respect to the choice of play, the rhythm of play and verbal communication. While a task-oriented approach might expect the child to complete a specific task, such as for instance performing imitation, this approach, here, inspired by non-directive play therapy, enables the child to proactively experience various situations of play, from simple exploration of the robot's features and capabilities to more complex

[1]At the moment a roboticist is needed to deal with the programming issues, but in the future, ideally, play therapists would be able to use this method as a new approach of robot-assisted play in the context of play therapy.

situations of play, possibly involving an understanding of the notion of causality as well as an ability to play symbolically, or take on a specific role in play. Furthermore, at any moment, the child can appeal to the experimenter's participation in play, thus enabling the child to experience triadic play.

Besides, beyond inspiration from non-directive play therapy, this approach introduces a regulation process (see Section 4.4.3.1). This process notably enables the experimenter to regulate the interaction in order to guide the child towards other play styles when needed, modify slightly the rhythm of play if she feels the child is 'standing still', or ask questions to the child about reasoning or affect related to the robot.

This chapter presents and explores the potential of this pioneering approach in robot-assisted play[2], through a long-term study with six children with autism. This study should be regarded as a preliminary exploration of the feasibility of such a technique in the context of robot-mediated therapy for children with autism. Several research questions are addressed:

- Does such an approach for robot-assisted play, inspired by non-directive play therapy, help the child experience higher levels of play and enable him/her to develop new play skills?
- Does this approach encourage the child to play socially?
- Might this approach be appropriate for children who play solitarily and speak mostly by onomatopoeia[3]? Might it help him/her experience social play? If not, what might be the additional requirements necessary for such experience?

The remainder of this chapter is structured as follows. Section 4.2 presents background on non-directive play therapy. Related work is presented in Section 4.3. Section 4.4 explains the method in terms of procedures and measures. Results are then provided in Section 4.5 and discussed in Section 4.6. Finally, a conclusion (Section 4.7) closes the chapter.

4.2 Non-directive Play Therapy

This section summarizes the core ideas of non-directive play therapy as mainly developed by Axline (1947) and explained and illustrated by case studies reported on

[2] Note that robot-assisted play is considered as a subfield of robot-mediated therapy

[3] Onomatopoeia is a word that imitates the sound(s) associated with objects or actions it refers to, e.g. 'buzz'.

by Ryan and Wilson (1996).

Non-directive play therapy has its roots in Rogerian client-centred therapy with adults (Rogers, 1976), adapted to child therapy with a focus on play as the principal medium of communication (in contrast to verbal exchange). Rogerian theory relies on the idea that all human beings have a drive for self-realisation; it means that any human being tends to develop towards maturity, independence and self-direction. The individual needs to completely accept himself/herself as well as be accepted by others.

In non-directive play therapy, the child, rather than the therapist, chooses the type of play and the activity in general in the playroom. This contrasts with other play interventions. We shall cite Axline (1947) who primarily developed the method of non-directive play therapy: "Non-directive play therapy is not meant to be a means of substituting one type of behaviour, that is considered more desirable by adult standards, for another 'less desirable'. It is not an attempt to impose upon the child the voice of authority that says 'You have a problem. I want you to correct it'." A few limitations in the behaviour of the child are set which refers to safety and security reasons.

A relationship is progressively built up between the child and the therapist. This relationship enables the child to share his/her inner world with the therapist and, "by sharing, (the child) extends the horizons of both their world" (Axline, 1947). Ryan and Wilson (1996) state that this relationship, with the help of the therapist, progressively facilitates the child to choose freely the feelings he/she wishes to focus on as well as the way how he/she wants to explore them. Three mediums may be used for communicating these feelings: action, language and play.

The therapist participates in the therapy. He/she observes, listens and answers to the child. The therapist is reflecting the child's feelings or emotionalized behaviours in order to help him/her build a better understanding of himself/herself. The therapist's role has been characterized by eight basic principles set out by Axline (1947), see Fig. 4.1.

It should be noted that in the study presented in this chapter, the experimenter was not trying to engage in therapy; the study only drew inspiration from non-directive play therapy, thus the context may be a therapeutic one, but the experimenter, a human-robot interaction researcher, was not behaving exactly like a therapist. The experimenter was not applying strictly the eight principles set out by Axline (1947), see Fig. 4.1. She very much drew inspiration from principles 1, 2, 3,

1. ``The therapist must develop a warm, friendly relationship with the child, in which good rapport is established as soon as possible."
2. ``The therapist accepts the child exactly as he is."
3. ``The therapist establishes a feeling of permissiveness in the relationship so that the child feels free to express his feelings completely."
4. ``The therapist is alert to recognize the *feelings* the child is expressing and reflects those feelings back to him in such a manner that he gains insight into his behavior."
5. ``The therapist maintains a deep respect for the child's ability to solve his own problems if given an opportunity to do so. The responsibility to make choices and to institute change is the child's."
6. ``The therapist does not attempt to direct the child's actions or conversation in any manner. The child leads the way; the therapist follows."
7. ``The therapist does not attempt to hurry the therapy along. It is a gradual process and is recognized as such by the therapist."
8. ``The therapist establishes only those limitations that are necessary to anchor the therapy to the world of reality and to make the child aware of his responsibility in the relationship."

Figure 4.1: Eight basic principles set out by Axline for practice of non-directive play therapy, quotations from Axline (1947).

5 and 8, but she was not dealing with the fourth one; and, concerning principles 6 and 7, she was considering these principles with more flexibility. It is worthy of note here that this study is a first step towards a proof-of-concept and required robotics expertise; in future, play therapists may use this approach.

4.3 Related Work

4.3.1 Non-directive play therapy for children with autism

Non-directive play therapy has been largely used for children and adolescents with a wide variety of emotional and behavioural problems (Ryan, 1999, 2001; Ryan and Needham, 2004). Only recently have researchers started to investigate the feasibility of such techniques with children with autism. A pioneering case study is presented in 2004 by Josefi and Ryan (2004). In that paper, Josefi and Ryan (2004) present a case study with a 6-year-old-boy with severe autism by using the non-directive play therapy technique. Before starting the experiments, the boy was mostly communicating non-verbally, and hardly controlled his sudden excess of energy. He was described as never playing with his brother and sisters and whenever he played, he only engaged in playing mechanically with toys. The child attended 16 non-directive play therapy sessions of an hour over a 5-month period in the child's special school. The room was empty except from specific materials selected for their "expressive, imaginative, relaxing and interactive properties". Results were analysed both qualitatively and quantitatively. The findings showed an increase in the child's autonomy and initiative-taking. Besides, the child developed attachment to the therapist. According to Josefi and Ryan (2004) it was shown that non-directive play therapy itself may provide children with autism with: "(i) emotional security and relaxation, (ii) an enhanced and attentive adult environment in which playing together is emphasized, and (iii) the acceptance by therapists of children's ability to instigate therapeutic change for themselves under favourable conditions". These conditions constitute the basis for therapeutic progress as written in play literature (Axline, 1947). Besides, the child's repertoire of play appeared to expand and the child managed to concentrate progressively longer during the sessions. During the last sessions the child proactively engaged in play requiring more joint attention and direct social interactions with the therapist. He started to become more and more interested in toys that have symbolic characteristics. He also communicated more and more verbally with the therapist. It is perhaps worthy of note here that the symbolizing capacities have similarities with, and may overlap capacities, to learn language during normal development; in return, it is very likely that learning a language requires some symbolizing capacities and processes. However, repetitive and obsessive behaviours were not considerably reduced. As a conclusion, Josefi and Ryan (2004) stated that non-directive play therapy with children with autism may be complementary to behaviour therapy, non-directive play therapy likely to be more efficient in the child's gaining autonomy, taking initiative, joining attention and developing social and symbolic play, while behaviour therapy would be more efficient in reducing ritualistic and obsessive behaviours.

4.3.2 Robot-mediated therapy in the context of autism

Within the Aurora Project, Robins et al. carried out several studies analyzing on the one hand the role of the robot as a mediator (Robins et al., 2005a) and on the other hand the role of the experimenter (Robins and Dautenhahn, 2006) in the trials. Robins and Dautenhahn (2006) describe the role of the experimenter as the one of a "passive participant" who responds to the children if they initiate interaction with him/her. In Robins et al.'s experiments, children interacted with a small robotic doll, Robota, by imitation of gestures, that is imitation of position or movement of arms and legs. In these trials, Robota was either simulating a dance or being controlled remotely by the experimenter. Thus, there was no autonomous reaction from the robot to the child's interactions in their study.

In different studies, Werry et al. focused on free-play with a mobile autonomous robotic platform, Labo-1, equipped with infrared and heat sensors (Werry and Dautenhahn, 1999; Dautenhahn et al., 2002; Werry et al., 2001). Its shape is rectangular (30cm wide by 40cm long), it weights 6.5kg and does not have pure tactile sensors. The play situations were mainly approach and avoidance games whereby turn-taking emerged from the child-robot interactions (Dautenhahn, 2007b). The experimenter did not take part in the games; they only responded to the child when the child initiated communication or interaction with them (Dautenhahn and Werry, 2002). The child played therefore in a relatively unconstrained environment on his/her own with the robot (Werry and Dautenhahn, 1999), or two children interacted at the same time with the robot (Werry et al., 2001).

Outside the Aurora project, Kozima et al. (2005) used a small dancing creature-like robot, Keepon, in a long-term study with children with autism, most of the time in partly unconstrained conditions. During these experiments, the small creature-like robot was manually controlled by the experimenter who was not part of the trials. Rather, carers were part of the trials with the child. The experiments highlight the role of Keepon as a pivot in triadic interaction by studying, in particular, the emergence of joint attention. This result reinforces the idea that child-robot interaction may be valuable for children with autism.

4.4 Method

4.4.1 Participants

All the children taking part in the experiments have a diagnosis of autism and are from the same school based in Hertfordshire, UK. This school welcomes children between 4 and 11 years old with moderate learning difficulties. In particular, an Autism Base provides extra care and a specific education program for children with autism to start within the school. When the child gets older or when he/she has made sufficient progress (especially if he/she improved social skills) he/she can be integrated in a more general class, which gathers children with specific needs and abilities but not only children with autism.

For clarity and simplicity purposes, a consistent naming of the children will be used in the whole thesis, starting with A and then, alphabetically, in order of appearance in the text.

Two boys from the Autism Base, Child A (seven years old) and Child B (eight years old) were invited to take part in the experiments. Both of them find it hard to express themselves verbally and their behaviour often includes onomatopoeia and repetitive gestures. According to the teachers, Child A often shows apprehension towards dogs and doors and Child B has a fascination for computers. Child C took part in the experiments too. She is a seven years old girl. During the experiments, she was part of the Autism Base but in the process of being integrated to another class with children with moderate learning difficulties but not only children with autism. She therefore started to follow part-time the education program of this class and the rest of the time stayed in the Autism Base. She masters verbal communication pretty well and teachers describe her behaviour as proactively social, as far as play at playtime is concerned.

Three older children took also part in the experiments. All of them are integrated in classes for general moderate learning difficulties. Child D, ten years old, is described by his teacher as a solitary child. In the classroom the position of his desk, fairly isolated from the others, gives him an 'own' space. Child D understands pretty well when one addresses him verbally but mostly speaks by onomatopeia. At school, he often uses the computer to do exercises, especially exercises on words and writing. Two other children, Child E, ten years old and Child F, nine years old, took also part in the study. They communicate verbally and are not described as solitary children.

The children's specific levels of autism are specified in Appendix B. The study was carried out with approval of the University of Hertfordshire Ethics Committee. The parents of all the children who took part in this study gave written consent, including permission to videotape the children and utilize photos in publications.

4.4.2 Artifact

The main artifact used in this study was a white robotic mobile autonomous dog, the Sony Aibo ERS-7, an off-the-shelf robot commercialised by Sony (Fig. 4.2). Aibo ERS-7 weights approximately 1.65kg and measures approximately 180(w) x 278(h) x 319(d) mm. It is equipped with a great variety of external sensors (e.g. infrared sensors, stereo microphones, tactile sensors). In our study, tactile sensors that the children can activate by stroking the robot played a major role. Those sensors are: the head sensor, the chin sensor and the three back sensors. Aibo's control programming is achieved using URBI (Universal Real-Time Behaviour Interface (Baillie, 2005)). URBI is a scripting interpreted language. It uses a client/server architecture, the connection between the server (robot) and the client being made through a wireless connection. The client can control the joints of the robot and access sensors and any accessible part of the robot. URBI can be used with various robotics platforms and various languages (e.g. C++, java, etc) on the client side to control (program) the robot. In particular, existing libraries (e.g. the C++ URBI library, liburbi-C++, or the java URBI library, liburbi-java) provide simple ways to program a client and make a powerful use of URBI's functionalities. In this work, I developed my own programs using liburbi-java.

Figure 4.2: Aibo ERS-7.

4.4.3 Procedures and Measures

4.4.3.1 Procedures

Experimental Setup. The experiments took place once a week, on Wednesday mornings, in the school. Each child took part in a maximum of ten sessions. Not everybody could take part in ten sessions because some of them may have been away for a day or on a trip with their class. Note, an exception was made for one child who showed some apprehension towards the robot: for this specific child, experiments were stopped after five sessions and only restarted on the last day of the experiments when he proactively came to the trial.

The rooms used for the experiments changed several times due to circumstances at the school. In each case, the child may encounter possible distractive objects, like toys or mirrors. Thus, these experiments took place in a context of possible distraction. The different rooms used for these experiments are described in Fig. 4.3 and a list of the rooms used for each session is provided in Fig. 4.4.

Each trial involved one child with autism, the experimenter (myself) and possibly another researcher from the Aurora project with whom the children were familiar. The latter helped the experimenter film the trials and occasionally took part in a

Room	Description	Dimensions	Furniture in the room	Objects in the room
R 1	Small room	Approx. 10feet * 8feet	-small longitudinal window on the very top (children can't see through it), -cupboard, -low rectangular table, -2 children's chairs, -decoration on the wall (a clown's head drawn on a paper board).	Regular objects: - game with individual letters to form words, reflective blue metallic support, - coloured cubes (25mm*25mm) - rectangular paperboard 3D decoration, 1m*30cm*20cm , vertically in a corner. On occasion: man's like face drawn on a paperboard that children could hold in front of their face.
R 2	Small room in the Autism Base	Approx. 10feet * 12feet	-big window on a wall, -second internal window (semi-transparent, semi-reflective) with view on another classroom; -vertical mirror, children can see their whole body by reflection -shelves on the very top, children can't access -table & small chairs (session8 only)	- games in open boxes on the shelves (e.g. a doll); children can see them but can't access them.
R 3	Large meeting Room: library, kitchen and living room corners. Experiments took place in the living room corner.	-room: Approx. 35feet * 40feet; -living room corner, approx. 10feet * 12feet	-Large windows on two walls -2 sofas made of joint comfortable chairs -4 comfortable additional chairs -rectangular dinner table, 6 chairs -2 low coffee tables -shelves (at the entrance) -kitchen corner	-magazines on the coffee table -on the shelves, objects such as cloth samples in open boxes -small calculator -small alarm clock
R 4	Classroom; experiments took place in the library corner	-room: Approx. 30feet * 30feet; -library corner: approx. 10feet * 7feet	Library corner: -2 shelves separating the library corner from the rest of the classroom -small children's bench	Library corner: -books

Figure 4.3: Description of the school's rooms used for the experiments.

Session	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
Room	R1	R1	R1	R1	R1	R3	- Child C : R3 - Other children : R4	- Child C: R3 - Other Children: R2	R2	R2

Figure 4.4: List of the school's room(s) used for each session.

verbal communication process by answering a child's question directly addressed to her.

The duration of the sessions was variable. The child was free to play as long as he/she wanted with the following restrictions: i) the upper limit of time was 40 minutes (so that the child did not miss too much of his/her courses at school); ii) if the child had an obligation due to his/her planning, the session was shortened.

The Aibo robot was programmed in order to show simple behaviours, tailored progressively by immersion according to each child's needs and abilities. Note that "tailored by immersion" means here that the repertoire of appropriate robot's behaviours with respect to each child specific needs, abilities, dislikes and preferences was progressively refined as the experiments progressed. The mapping between the sensors and the reactions of the robot (also called behaviour-mode) could therefore vary from one session to the other and also during a session in order to meet as close as possible the needs, abilities and demands of the child at a given moment. The robot reacted autonomously to the activation of its sensors, with respect to the specific behaviour-mode it had been endowed with. The switch between various behaviour-modes was done manually by the experimenter through a wireless connection with a laptop. The laptop was located in the same room as the children, and thus constituted an additional source of distraction for the children.

Methodology of the approach. During the session, the child was invited to play with the Sony robotic pet Aibo. The experimenter took part in the experiment. The child was the major leader for play: the child was free to choose the game to focus on, the pace of play and he/she could engage in free-play (unconstrained play) with the robot and/or the experimenter; he/she was also free to engage in communication with the experimenter whenever he/she wanted. If the child appealed to the experimenter's participation, then the experimenter did take part in the game. If the child initiated verbal or non-verbal (e.g. smile, eye gazing) communication with the experimenter then the experimenter answered appropriately. With respect to verbal communication, the experimenter tried to answer every question of the child and rewarded him/her verbally whenever appropriate. Note that this approach is mainly child-centred, relies strongly on the child capabilities of designing his/her own trajectory of progression and on total respect and consideration towards the child from the experimenter. In this sense, this approach draws inspiration from non-directive play therapy.

Beyond inspiration from non-directive play therapy, this approach adds a regulation process under specific circumstances which are detailed below:

a) *to prevent from or get rid of a repetitive behaviour*: If the child was starting or about to start a repetitive behaviour, the experimenter intervened and tried to help the child play a different game;

b) *to help the child engage in play*: if the child did not engage in interaction with the robot, then the experimenter encouraged him/her to play with the robot, verbally and/or non-verbally (e.g. by stroking the robot and encouraging verbally imitation);

c) *to give a better pace to the game if already experienced by the child*: If the game was "standing still" but the child already experienced this game and had shown he/she was capable to play this specific game, then the experimenter could intervene punctually to confer a better pace to the game;

d) *to bootstrap a higher level of play*: if the child was about to reach an higher level of play but still needed some bootstrapping (some light guidance), the experimenter could provide it;

e) *to proactively ask questions related to affect or reasoning*: the experimenter could proactively ask the child simple questions related to affect or reasoning such as: "Do you think Aibo is happy today?" or "Do you like playing with Aibo?".

Note that e) enables: i) to test the ability of the child to answer and/or ii) to show the child a specific point for reasoning. We shall give several examples within various levels of reasoning:

1. technical issue: show the child how to change the battery of the robot so that he/she can do it next time in a context of cooperative task;
2. ask the child if he/she thinks Aibo is happy;
3. help the child reason on causal effect: stimulation of a sensor implies a specific reaction of the robotic dog;
4. show the child that a reaction can be interpreted: e.g. if I press this specific button, then Aibo wags the tail; and wagging the tail can mean that Aibo is happy; thus if you press this button, you can show that Aibo is happy.

4.4.3.2 Measures

Each session was filmed unless the child explicitly asked for not being filmed which rarely happened. First, the experimenter viewed the video recordings and wrote down notes on the events constituting each session. These notes described the events in detail and contained as few interpretation as possible. As a second step, the experimenter analysed the data in terms of more abstract criteria that would enable her to identify, for each child, both the profile according to the three dimensions (Play, Reasoning and Affect) and the progresses made over the 10 sessions. This methodology allows to first gather as much information as possible before deciding on the specific criteria; it has the advantage of not restricting the analysis to pre-defined criteria which might reveal a posteriori not being the optimal ones to base the analysis upon. This is especially relevant in the case of an exploratory study. This procedure follows the procedure described by Schatzman and Strauss (1973), stating that: "the researcher requires recording tactics that will provide him with an ongoing developmental dialogue". Schatzman and Strauss (1973) underline the importance of recording observations from the very beginning of research. They also suggest taking notes separately, categorizing notes into three different packages: a) "observational notes" based on events, without interpretation; b) "theoretical notes" representing an attempt to confer or denote the meaning from an observational note; c) "methodological notes" dedicated to methodological comments.

Results of the experiments were analyzed according to three (intertwined) dimensions, respectively Play, Reasoning and Affect.

Play This study aims at testing the feasibility of this approach to encourage the child to learn new play skills and enable him/her to experience more and more complex play situations with respect to the following main criteria:

a) social aspect of play,

b) proportion of symbolic and/or pretend play,

c) understanding/use of causality,

d) ability to handle the pace of a specific play and possibly the chronology or the transitions between two logical segments of play.

That is why, concerning the dimension of Play, what particularly matters is 1) to extract information qualitatively about play situations that the child has experienced in each session, and 2) see if the child really experienced a large repertoire of play and more complex levels of play gradually over the sessions.

For this purpose, a Play Grid was built (after the play sessions) based on the children's plays objectively observed during the experiments. This grid is exhaustive with respect to the variety of play situations which took place at least once during the experiments for at least one of the children. Besides, the different play situations were classified into 6 sets, each set denoting a specific level of complexity of play (Level 1 being the lowest and then gradually incrementing the level of complexity until Level 6). The level of complexity is defined according to four criteria:

a) social play,

b) proportion of pretend and/or symbolic play,

c) exploration of the use of causality/reaction,

d) chronology and/or number of different phases in the play, e.g. a simple reaction to a sensor is constituted of two phases while a search and rescue game involves many phases to handle chronologically: i) initial situation, ii) search phase, iii) rescue phase, iv) final situation.

The level of complexity is then deduced from an average evaluation over the four components which explains that the same level may contain play situations with a predominant component of "d)" and others with a predominant component of "b)". Consequently, within a same level of complexity, the different play situations are not ordered since they may be very different in nature. Ideally, the child would experience higher levels of play over the time and, within a same level of complexity, different play situations in nature.

The systematic analysis with the grid for each child and each session shows the trajectory of each child (i.e. the profile of the child). Each cell in the grid is filled in if and only if it corresponds to a play situation experienced by the child at least once during that specific session; and the content depends on the play situation being acted proactively or reactively (i.e. the child was slightly guided towards this play situation by the experimenter).

However, this grid is much enlightening for children who manage to play socially and manage to diversify their play. For those who do not interact a lot with the robot and, when playing, tend to experience mainly solitary play through the exploration of the robot's features and behaviours, a more adapted tool to evaluate their progresses was used. That evaluation was quantitative and relied on measuring for the whole duration of each session:

1. the total time spent in interaction with the robot,
2. the duration for each single uninterrupted phase (period) of pure interaction (note that the total duration is the sum of the duration of each single uninterrupted phase of play),
3. the amount of gestures imitated by the child and the number of gestures explicitly asked by the experimenter to be imitated.

Reasoning Through play, children can notably construct some understanding of social situations and gain experience of some situations they encountered while playing. If a child can reason on abstract concepts, infer mental states and make a sense of social rapports, it will be easier for him/her to play symbolically. Reciprocally, while the child experiences symbolic play, he/she manipulates abstract concepts such as inferring an emotion or handling social rapports. Both play styles and reasoning are therefore intertwined and both views should therefore be used to analyse the results of the experiments carried out for this study. Note that with respect to 'Reasoning', what is particularly relevant are both questions and answers emerging from play situations. The context of play enables the use of imagination, whereby Aibo may be assigned a specific role by the child, and it allows the child to attribute specific capacities to the robot such has having mental states (e.g. it enables to imagine that Aibo is taking on a specific role and make further assumptions on his mental state or his social status). Consequently, the context of play enables the robotic pet to be attributed with mental states as well as a social role, and possibly moral standing. In this way, it is possible to explore quite largely the reasoning part of the coding manual developed by Kahn et al. (2003) for the analysis of children's conception of the Aibo robot, by exploring the four following categories used in Kahn et al. (2003): "Essence", "Mental States", "Social Rapport" and "Moral Standing". According to Kahn et al. (2003), those categories "reflect a "quadrology" of children's conceptions of Aibo and Shanti[4]".

For each of those four categories a list of related questions can be formulated (Kahn et al., 2003):

a) **"Essence"**: Does the child consider Aibo as an artefact or a biological entity?

b) **"Mental states"**: Does the child attribute mental states to Aibo? Does the child consider that the robot develops in terms of age for instance? Does the child consider Aibo has a personality? Does he consider Aibo could live autonomously?

c) **"Social rapport"**: How does the child position Aibo relatively to himself/herself;

d) **"Moral standing"**: Can Aibo be physically or morally hurt? Can he be held responsible for something? Can Aibo be punished when necessary? Could Aibo be praised?

Note that Kahn et al.'s coding manual has been developed in a different context than the one of this study: they targetted typically developing preschool children who only encountered Aibo once and afterwards immediately answered specific questions about "reasoning" (Kahn et al., 2003, 2006) - while answering questions, children could

[4]Shanti is the name of the stuffed dog that was used in Kahn et al. (2003)'s study as a basis for comparison.

however carry on interacting with the robot. Here, the context used in our study is different since the succession of sessions enables the child to progressively build some reasoning and understanding, along with the progressive building of a shared space of expressions and routine activities between the child and the experimenter. Therefore, the reasoning related to the robot can be enriched. Besides, 'reasoning' here is part of play in itself. In the study presented in this chapter, the context of play is actually used to enable the child to explore issues such as mental states or social rapports, and the robot in itself is a support for embodying such issues through the imaginary context that comes with play. Moreover, since the experimenter takes part in the experiments, not only social rapport between the child and the robot should be considered, but also the child's view on the notion of social rapport between the robot and the experimenter and between himself/herself and the experimenter. Consequently, here, the dimension of 'Reasoning' is analysed as follows:

1. The main features of the four categories ("Essence", "Mental States", "Social Rapport" and "Moral Standing") are extracted from Kahn et al.'s coding manual (Kahn et al., 2003);

2. The issue of whether and how the child addresses those features is investigated for each child, in a perspective of questioning through play rather than giving firm answers.

Note that since the experimenter is not a therapist, and since the behaviour of children with autism might sometimes be interpreted differently from typically developing children, in the analysis we only consider events which are objectively and reliably identifiable. Verbal events are particularly reliable events; they can be statements or questions arising from the child (major events) or answer to the experimenter's question (minor events). Below are some examples: a) Essence: "He's a robot, he is a robot dog", "He has short teeth, he doesn't bite. Robot dogs don't bite, do some do?"; b) Mental states: "Aibo is happy", "How old is Aibo", "Aibo, answer me, do you like toys?"; c) Social Rapport: "It is your robot"; d) Moral standing: the child accidently kicks the robot and apologized verbally to the robot directly. Besides, in many cases, as already explained, reasoning and play are intertwined; for instance, when the child and the robot's relative social position in an enacted situation of pretend play is well-defined by the child (e.g. a competition with two participants, the child and Aibo), the notion of social rapports is certainly addressed. Another example is a play situation of asking the robot about its mental states and answering with the activation of a sensor.

As a further step in reasoning, the child may tackle a more general issue related to his/her mental states for instance, or to social rapport, concerning himself/herself or even the experimenter. This is a relevant point for this study: it would show the potential reuse in another context of skills the child may develop or practise through reasoning about the robot during play.

1)	**Proactive (major) event related to affect:**
i)	Child's statement or question referring directly to himself/herself liking the robot or the robot liking him/her. No hug or kiss from the child to the robot. Examples: ``I like Aibo", ``Aibo likes me".
ii)	Child's verbal compliment to/concerning the robot. No hug or kiss from the child to the robot. Examples: ``good doggy", ``nice dog", "he is a nice dog".
iii)	Child's hug to the robot, clearly identifiable, accompanied by a kind word from the child to/concerning the robot or verbal statement qualifying the hug. Example: the child hugs the dog and asks the experimenter to hug the dog: ``Put your hands and hug, hug, hug!"
iv)	Child's kiss to the robot, clearly identifiable, accompanied by a kind word from the child to/concerning the robot. Example: the child gives a kiss to Aibo after saying ``Goodbye Aibo, have a good sleep"
2)	**Reactive (minor) event related to affect:**
i)	Child's answer to a question about himself/herself liking the robot or the robot liking the child. Example: the experimenter asks the child: ``Is it a nice robot?" and the child answers ``Yes".
ii)	Child's answer to a question about himself/herself being happy to play with the robot. Example: the experimenter asks the child: ``You are happy playing with the robot?" and the child answers ``Yes".
	Note, reactive events related to affect are considered very cautiously in this study; they are not considered as sufficient to make firm deductions about the child addressing the notion of ``Affect".

Figure 4.5: Criteria for coding events related to Affect. An event is related to 'Affect' if it corresponds to one of the items provided in the table; in some of the following figures, events related to affect are qualified by a corresponding code: the code of an event related to affect is given by its corresponding item's index, e.g. "I like Aibo" is [1i].

Affect The 'Affect' dimension represents any expression indicating whether the child likes the robot or not, or if the child makes an assumption on the robot liking him/her. Here, only obvious signs of like/dislike are considered, in order to ensure that events considered as related to affect are clearly identifiable (see Fig. 4.5 provides the table of criteria for the coding of events related to affect). For instance, a gentle stroke is not classified as an event related to affect in this study, neither a gesture such as a kiss or a hug, if it is not accompanied by an appropriate child's statement.

4.4.4 Coding and Reliability

Inter-rater reliability testing was carried out for each of the three dimensions, respectively, play, reasoning and affect. A second coder who was not familiar with the aims of the study re-coded part[5] of the data. Good reliability was shown: a) On play, 80.7% agreement ($13min50s$ of videos coded divided among two children, Child E and Child C); b) On reasoning, 80.3% agreement ($18min24s$ of videos coded divided among two children, Child E and Child F); c) On affect, 93.3% agreement ($22min$ of Child C's videos coded).

4.5 Results

Child A Child A showed some apprehension towards the robot and did not interact at all during the five first sessions. The experimenter therefore decided not to require the child to come for the following sessions and let the child proactively decide whether he wanted to take part in the further trials or not. In the last session (Session 10), Child A proactively came for the trial. In that session he engaged in an interaction with the robot with the help of the experimenter: one interaction event happened between the child and the robot, during which the experimenter showed the child how to stroke the robot and the child imitated (Fig. 4.6). Afterwards, the child both showed signs of light apprehension (he moved his body slightly backwards) and enjoyment (he smiled).

Child B Child B took part in 9 sessions (Fig. 4.7). Child B naturally showed attempts to play with the laptop rather than with the robot. It was a big challenge to get the child away from the laptop and get his attention focused on something else. The experimenter used a simple trick by hiding the laptop with a cloth. But for practicality reasons (e.g. to connect or reconnect Aibo during the session), the cloth had to be removed from times to times during the session thus introducing an important source of distraction for Child B. Progressively, the child seems to have understood that he was allowed to punctually have a look at the laptop (as part of his well-being) but that he should mostly engage in interactions with the robot. The table provided in Fig. 4.8 shows the average amount of time Child B spent engaging

[5]The recoded segments contained only high involvement of the children in interaction. High involvement is characterised by the fact that i) children do not stop interacting for a period longer than a few seconds, and ii) children experience many situations of play, reasoning or affect related to the robot. Therefore, the density of events to identify and code is very high in the recoded segments which makes the evaluation highly meticulous.

		1	2	3	4	5	6	7	8	9	10
L1	*Solitary* Exploration										
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face										
L2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										P
L3	*Social* exploration (social play)										
	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo										
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 4.6: Child A. Play Grid. The first column describes the corresponding level of play, the second column details the various play situations for each level that the child experienced at least once; the following columns refer to the sessions, ordered chronologically. The table is then completed according to the following rules: a) if the child did not experience the play situation during the specific session, leave the corresponding cell blank; b) if the child experienced the specific play situation at least once during the session, then write "P" (if the child experienced it proactively only – i.e. it was his/her own initiative). Write "r" if the child never experienced it proactively (only reactively: the experimenter guided the child towards the play situation). Write "B" if the child experienced this play situation many times , sometimes proactively and sometimes reactively. Note that Child A did not take part in the play sessions 6, 7, 8 and 9.

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration	P	B	B	P	r			B	P	B
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)								r	r	B
L 3	*Social* exploration (social play)										
	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo										
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 4.7: Child B. Play Grid. See Fig. 4.6 for a detailed caption. Note that Child B was away for Session 7.

	Total duration of play (min:sec)	Repartition of the play time in single phases of play (min:sec and + between 2 single phases)	Aspects of imitation: In each single phase of play, numbers of gestures: Imitated by the child	Explicitely asked by the experimenter to be imitated	Verbal expression involving either the word 'dog' or 'robot'
Session1	0:06	0:06	0	0	
Session2	1:30	1:00 + 0:30 (mostly looking attentively at Aibo)	0	0	
Session3	0:40	0:40	0	0	
Session4	Almost null	Almost null	0	0	'The little dog was easy'
Session5	0:15	0:15 the experimenter helps by holding the child's hand to show him	0	0	
Session6	0:00	0:00	0	0	
Session7	away				
Session8	1:05	1:05	1	2	
Session9	2:21	0:40 +1:16 +0:16	0 +1 +0	0 +2 +0	
Session10	5:24	0:20 +1:47 +0:18 +2:46	0 +3 +0 +3	0 +3 +0 +1	

Figure 4.8: Child B. Dimension of play: quantitative results: For each session, the following indicators are reported: a) total duration of play; b) duration for each specific single session of play ; c) aspects of imitation with respect to i) the occurrence of gestures (touch or stroke of the robot) that the child imitated and ii) the occurrence of gestures that the experimenter explicitly asked the child to imitate; d) verbal expressions including the word "dog" or "robot".

in play with the robot during each session. The tendency is clearly that the child played longer with the robot in the two last sessions than in the previous ones and almost doubled his play time between the 9^{th} and 10^{th} session. If we consider in detail the duration of single phases of play, i.e. uninterrupted periods of time when the child continuously played with the robot, then, again, this table shows that the child experienced longer uninterrupted periods of play with the robot during the last sessions. Typically, two uninterrupted periods of play are often separated by an attempt of the child to play with the laptop. This shows that the child progressively learnt to focus more and more on the robot and on engaging in play with the robot. Nevertheless, the experimenter also often intervened to help the child carry on playing and keep focusing his total attention to the robot; this intervention usually happened in two ways: a) encouraging and rewarding the child verbally, or b) showing an example, e.g. stroking the robot and asking for the child to do the same. In this context, 'b)' is very relevant indeed since the child does not speak verbally and encouraging imitation is favourable for both relaunching the child's engagement in play and bootstrapping social play. It should be noted that in this specific context, imitation is very rudimentary: the experimenter either touches a specific sensor or gently strokes the robot (e.g. on the head) and explicitly asks the child to do the same. The child is considered to imitate the experimenter's gesture if he initiates within 10 seconds the same nature of gesture, i.e. either a touch of a sensor or a stroke, and if the gesture is applied on the same part of the robot's body; for instance, i) the experimenter touches the head sensor and, within 10 seconds, the child presses the same sensor (with or without activation depending on the child's precision of touch) ; or ii) the experimenter gives a gentle stroke on the back of the robot and, within ten seconds, the child gives a stroke on the back of the robot. Results show that Child B progressively experienced more situations of imitation. Besides, they also reveal that during the last session he imitated some gestures proactively, i.e. without being explicitly asked by the experimenter to imitate.

Concerning the "Reasoning" dimension, Child B did not address the issue verbally. Thus, no firm conclusions should be drawn. However, the detailed study of the child's gestures shows that the exploration of the child became progressively richer and richer over the sessions. The child varied his position relative to the robot, from sitting to kneeing and lying, and thus looked at the robot from various viewpoints. Moreover, he progressively varied his way of touching the robot: during the first sessions, he progressively abandoned random-like touch to develop more targeted touch. Note that targeted touch can be, for instance, trying to touch a single sensor precisely or stroke the robot gently and then activate many sensors. Besides, during the last session, the child experienced proactively a combination of two previous sensor activations: first, he imitated the experimenter and stroke the back of the robot; second, he imitated the experimenter again and touched the head; third, his next behaviour was the simultaneous activation of back sensors and the head sensor.

Concerning the third dimension, "Affect", no event that was related to affect (with respect to Fig. 4.5) was recorded.

Child D Child D was away for Session 3 and Session 6 and therefore took part in 8 sessions in total. The analysis of the Play Grid in Fig. 4.9 shows that Child D played mostly solitarily. He engaged largely in exploratory play which became progressively more and more enriched. Two main aspects objectively illustrate the phenomenon a) a progressive change of position (from sitting orthogonal to the robot and not facing the experimenter to facing the robot and the experimenter) and b) a

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration	P	P		P			P	P	P	P
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face	P			P	P		P	P	P	P
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)				P	P		P	B		
L 3	*Social* exploration (social play)										
	Simple Bite/Save or Give/Food - no use of the sensors										
	Position or locomotion game – with verbal qualification of the game										
	Cooperative technical task: change the battery, or turn on/off Aibo				P	P		B	P	P	B
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"										
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor										
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor										
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story									P	P
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 4.9: Child D. Play Grid. The first column describes the corresponding level of play, the second column details the various play situations for each level that the child experienced at least once; the following columns refer to the sessions, ordered chronologically. The table is then completed according to the following rules: a) if the child did not experience the play situation during the specific session, leave the corresponding cell blank; b) if the child experienced the specific play situation at least once during the session, then write "P" (if the child experienced it proactively only – i.e. it was his/her own initiative). Write "r" if the child never experienced it proactively (only reactively: the experimenter guided the child towards the play situation). Write "B" if the child experienced this play situation many times , sometimes proactively and sometimes reactively. Note, Child D was away for Session 3 and Session 6.

more diversified way of touching the sensors. Moreover, the child practised “solitary mirror play” frequently. It consists in looking at one’s image in the robot’s reflecting face. Child D experienced situations of looking at his image with other reflecting surfaces too, such as a window, partially reflecting, or a mirror, perfectly reflecting (room R2 contained a mirror). All of these play situations, consisting in looking at one’s image, were often fascinating for Child D, and sometimes prevented him from engaging in other kinds of play situations. Besides, Child D did not experience play involving explicitly causal reactions, such as showing a specific reaction of the robot through the sensors’ activation.

However, progressively, Child D experienced situations with some components of social play. From a cooperative point of view, the child did take part, both reactively and proactively in cooperative technical tasks such as turning on the robot. Furthermore, Child D, who mostly speaks by onomatopeia did develop some ways of expressing himself, by dancing in front of the mirror and/or the robot and even probably telling a story by using not proper words but onomatopeia. The situation described below, that Child D experienced, may actually be interpreted, with caution, as a storytelling situation: Child D chronologically a) pressed the button to “wake up” Aibo (i.e. turn Aibo on), then b) stood in front of the wall mirror in the room, still watching Aibo “waking up”; c) once Aibo had woken up, the child started dancing and saying onomatopeia in front of the mirror. At some point, the robot disconnected. During the whole process the experimenter told Child D many times that she thought he was telling a story and asked him if she was right. She got no answer. When the robot disconnected the child stopped dancing and the experimenter reiterated her question: “Was it a story that you were telling me? Yes or no?” and the child answered “Yes”. Then she asked: “Can you tell me another story, yes or no?” and the child answered “yes”. Then the child repeated the same succession of behaviours ‘a)’, ‘b)’ and ‘c)’ and she asked: “Is it about a boy the story?” And he answered “Yes”. It is worthy of note here that the child might have simply repeated the word ‘yes’ after each question without giving a ‘real’ answer to the questions. Nonetheless, that example shows how the child may have progressively opened up to more communication with his surrounding social environment for play (notably the experimenter).

This storytelling situation took place in the last sessions while the child was starting to answer some questions about reasoning as well as using proactively verbal expressions to express intention. An in depth study of the verbal answers the child formulated shows that over the first sessions, the child almost only answered “yes” or “no”, whenever he answered. Then, progressively, the child answered some questions by repeating words from the question: e.g. in Session 4 the experimenter asked “Do you want to play with the robot or go back to the classroom?”. The child answered: “play with the robot”. And in the last two sessions, the child did use expressions to express his own intentions; for instance, the expression “sitting down” means that he wants to remain sitting down on the ground to carry on playing with the robot. In Session 9, the experimenter actually asked the child: “Do you want to go back to the classroom or play with him (the robot)?” and the child answered “play with him”. Then later in the session, the experimenter asked the question “Shall we go back to the classroom now?”. And the child answered: “Sitting down”. During the last session, the child reused exactly the same expression (“sitting down”) to answer the experimenter’s question: “Would you like to go back to the classroom soon?”.

Regarding the analysis of the reasoning dimension, the child answered reactively very basic questions about Aibo’s mental states, such as “Do you think Aibo is happy today?” or about his own mental state: “Do you like playing with the robot?” but there was no proactivity from the child with respect to mental states.

Concerning “Social rapport”, the child progressively grasped the fact that Aibo belonged to the experimenter. In the first sessions, the experimenter had to explain many times to the child that he could not take the robot with him back to the classroom. In contrast, at the end of the last session, the child hesitated a short time and gave the robot back to the experimenter proactively. Apart from that, the child did not explicitly show any reasoning on “Social rapports”. Neither did he on Aibo’s “Moral standing”.

The dimension of Affect has been mostly addressed indirectly (Fig. 4.10), through simple questions from the experimenter: in Session 4, the child answered affirmatively to the following questions: a) “Is it a nice robot?” and b) “You are happy playing with the robot?”. Later, in session 9, the child answered affirmatively to the question “Do you think Aibo likes you?” And in Session 10, the child answered affirmatively to the question “You like the robot?”. Note that since these inputs did not emerge proactively we should be careful with too much interpretation. Nonetheless, it should be underlined that most of the time the child said he preferred playing with the robot rather than going back to the classroom, which shows the child was having fun playing with the robot. It is perhaps worthy of note here that the experimenter is aware that

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	
S2	· **[2i] "Do you like it?"** (Experimenter); **"Yes"** (Child D)
S3	
S4	· **[2i] "Is it a nice robot?** (Experimenter); **"Yes"** (Child D); · **[2ii] "You are happy playing with the robot?** (Experimenter); **"Yes"** (Child D)
S5	
S6	
S7	
S8	
S9	· **[2i] "Do you think Aibo likes you?"** (Experimenter); **"Yes"** (Child D)
S10	· **[2i] "You like the robot?"** (Experimenter); **"Yes"** (Child D)

Figure 4.10: Child D. Events related to Affect. Events are separated by bullet points, and provided with their context (normal font) in the table. Events written in bold are coded according to Fig. 4.5 (the code is provided in brackets in front of the event); please note, that when the child answers a question, the event in itself is the child's answer, but, in this table, in order to make it clear to the reader, the question that the answers refers to is also written in bold.

the child may just have given a stereotypical answer[6].

Child C Child C was away for Session 7 and thus took part in 9 sessions in total (note that in Session 6 she had a very limited time of play, approximately 10 minutes, because of a class trip). The Play Grid in Fig. 4.12 shows that Child C experienced more and more complex levels of play during the sessions (see Fig. 4.11). She experienced in play situations involving the activation of a specific sensor to generate a precise reaction only a bit. She rather proactively experienced firstly play situations where "affect" is largely addressed (e.g. "Social Hug"). Secondly, she developed play situations where the robot embodied a character in a story she was telling. Finally, in a third and last phase, she initiated play situations where she was able to tackle issues on social rapports or mental states (Session 10: "look after Aibo and set up rules" and "search and rescue" play situations).

The "looking after Aibo" game dealt with deciding that she and the experimenter would take care of Aibo, and Child C proactively suggested that, as a consequence, she and the experimenter would have to define rules the robot would have to respect; and she enumerated the rules (among them, a detailed list of what the robot is not allowed to eat, and the statement: "dogs must go outside and must walk", followed by "I need to make him walk"). This game also gave rise to proactive inferences of state, the child even saying: "Look! He is smiling!" in the proper context. The social

[6]For instance, the experimenter did not ask the question: "Does the robot hate you?", which the child might have said "yes" to as well.

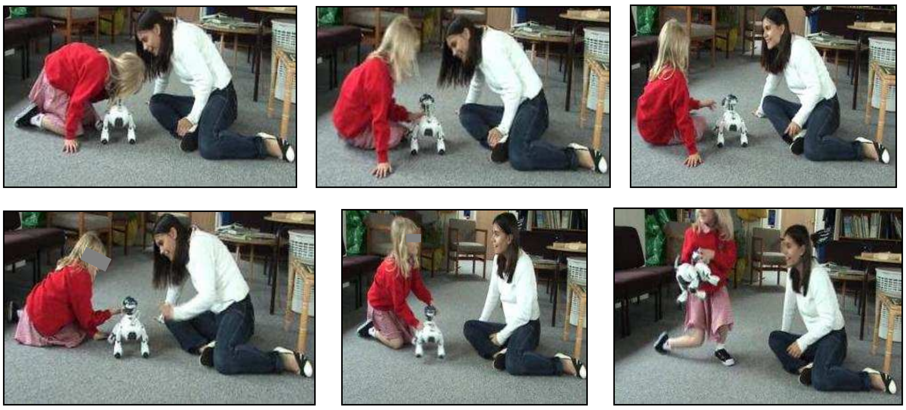

Figure 4.11: Child C involved in social play with the experimenter. 2 sequences are displayed, one on each line. Each sequence is organised chronologically; on the first line, picture on the right and on the second line, picture in the middle, Child C is making eye contact with the experimenter.

status that she took of taking care of Aibo led her to show the experimenter how to do specific things such as to make Aibo go forward: "You see, you must do like this, see".

Furthermore, this game was followed by a "search and rescue game" which was extremely rich in many ways:

a) The child led the rhythm, the pace, and the three steps of the play situation (chronologically):

- step 1: initial situation where Aibo is lost, the goal of finding Aibo is stated,
- step 2: the experimenter and the child are looking for the dog,
- step 3: final situation: the experimenter and the child find the dog.

b) The child slightly dilated step 2 over time so that she could deal with emotional states, particularly sadness: "You think we've lost him forever" said Child C; "Oh, that's sad" said the experimenter; and the child replied: "I think we're sad actually" thus conferring a socio-dramatic dimension to the current play situation.

c) During step 3, when the robot was found, the child introduced some reasoning about categories: she introduced the notion that it might be another robot than Aibo that she and the experimenter had found; she introduced this reasoning step by step and she might not have been really at ease with these concepts, but the point is that she practised them through experiencing them: Child C's reasoning started with "Oh no, there are two Aibos here" and, after several steps in the reasoning, she drew the following conclusion: "No there are two dogs, only one

		1	2	3	4	5	6	7	8	9	10
L1	*Solitary* Exploration										
	"Imitation" of robot's bark				P	P			P		
	Solitary mirror play – look at oneself in the robot's reflecting face										
L2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L3	*Social* exploration (social play)	P	P	P	P	P	P		P	P	P
	Simple Bite/Save or Give/Food - no use of the sensors						r				P
	Position or locomotion game – with verbal qualification of the game	P				P	P		P		
	Cooperative technical task: change the battery, or turn on/off Aibo		P	P	P		r		r	P	
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"		P	P	P				P	P	P
	Basic pretend & social play – imitate Aibo's snoring & verbal comment		P								
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo			P	P	P	P				
	Repeat after me - ask the experimenter to repeat verbal expressions										P
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)				P						
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)										
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo			P							
L4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors								B	B	P
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)										
	Speak directly to Aibo about Aibo's feeling (symbolic play)										
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor						P				
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor						r		P		r
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story				P		P		P	P	
L5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										P
	Pretend & symbolic & chronological play & social rapports: Search and rescue										P
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 4.12: Child C. Play Grid. See Fig. 4.9 for a detailed caption. Note, Child C was away for Session 7.

Figure 4.13: Child C's social hug to the robot. photos ordered chronologically. The child brings the robot to a second researcher (who helped out during this trial) while saying "Put your hands and hug, hug, hug" and both of them hug the dog. On the third picture from the left, Child C makes eye contact with the researcher.

Aibo. The clever one!" and she threw up her hands accompanied by a big smile. Again, here is illustrated that both "reasoning" and "play" dimensions are highly intertwined.

Concerning the notion of "Essence" for the Reasoning dimension, Child C mixed the use of artifacts and biological statements such as saying within the same session: "He's a robot, he's a robot dog" and "Nice dog", "He is a nice dog", "I love dogs", "A boy or a girl?" (Session 10).

Except in the last session, the notion of "Mental states", was addressed mostly reactively: the child answered to questions asked by the experimenter such as "Do you think Aibo is hungry" (which usually initiates the game "Give food/drink"). There were two exceptions: a) the child proactively said that the robot liked her, and b) the child could sometimes refer to mental states when telling stories she adapted from well-known children's books. During the last session, the child proactively referred to mental states of the robot as mentioned above in both "look after" and "search and rescue" play situations. During the "look after" play situation, she said: "We play, want to make the dog happy, make the dog feel pretty".

Moreover, as already mentioned above too, she experienced "Social rapports" a lot e.g. either simply by saying (in Session 9) "Look at Aibo, Aibo is your dog" or in taking on specific social roles in more elaborated play situations (e.g. in Session 10, during "look after" and "search and rescue" games).

Concerning "Moral standing", no objective event related to it happened.

The dimension of affect played an important role for the child (Fig. 4.14). In Session 1 already, she started saying "good doggy" with respect to the robot. Then, in Session 3 she introduced the notion of social hug (see Fig. 4.13), which consisted in asking the experimenter (or the second researcher present) to help her hug the dog:

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	· **[1ii] "Good doggy"** (Child C) while stroking the robot and looking at the experimenter (eye contact)
S2	
S3	· **[1iii] "Help me hug the dog: put your hands and hug, hug, hug"** (Child C) while bringing the robot near the assistant and showing how to hug · **[1ii] "Good doggy"** (Child C) · **[1i] "The dog really likes me"** (Child C). The experimenter answer "yes" · **[2i] "Do you like it?** (Experimenter). **"Yes"** (Child C)
S4	· **[1ii] "Good doggy"** (Child C), while stroking the robot · **[1i] "The dog really likes me"** (Child C) and she starts mimicking the noise that would do the dog by lapping her.
S5	· **[1ii] "Good doggy"** (Child C) and she looks at the experimenter; "yes very good doggy" (Experimenter).
S6	
S7	
S8	· **[1ii] "Good doggy"** (Child C) after the robot has "woken up" (i.e. is connected)
S9	· **[2i] Are you happy to see Aibo?** (Experimenter); **"Yes"** (Child C)
S10	· **[1ii] "Nice dog"** (Child C) · **[1i] "I love Aibo. I love Aibo"** (Child C) and she strokes the robot · **[1ii] "Good boy, good boy"** (Child C) and she strokes the robot · **[1i]** "Do you like the walk C, please tell me? (Experimenter); **"Yes, this is all about dogs like me"** (Child C) · **[2i] You like Aibo, right?** (Experimenter); **"Yes"** (Child C)

Figure 4.14: Child C. Events related to Affect. See caption of Fig. 4.10 for details.

"Put your hands and hug, hug, hug" Child C asked. Later in the same session, as well as in session 4, the child said, "The dog really likes me". Note that end of session 3 is the first time she answered to the question "Do you like it(Aibo)?" (she answered affirmatively). From that session onwards, the child confirmed several times the fact that Aibo liked her (e.g. session 4 "The dog really likes me") and that she liked Aibo (e.g. in session 10: "I love Aibo" and "Nice dog").

Child E. Child E took part in the 10 sessions of experiments. The Play Grid in Fig. 4.15 shows that Child E progressively experienced more and more complex levels of play over the sessions. During the first sessions, he attentively explored the reactions of the robot and in the following sessions, he experienced more and more simple causal reactions through the following games: a) "ask about a feeling, answer with a sensor", e.g. in Session 10 the child asked: "are you happy?" and pressed the head button which made the robot wave the mouth as to say "yes". b) "aim at a physical reaction, show it with sensors": e.g. the experimenter asked "Do you think

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration										
	"Imitation" of robot's bark										
	Solitary mirror play – look at oneself in the robot's reflecting face			P							
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L 3	*Social* exploration (social play)	P	P	P	P	P	P	P	P	P	P
	Simple Bite/Save or Give/Food - no use of the sensors						r	r			
	Position or locomotion game – with verbal qualification of the game	P		P					P		P
	Cooperative technical task: change the battery, or turn on/off Aibo	P	P	P	P	B	P	P	P	P	P
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"		P			P					
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo										
	Repeat after me - ask the experimenter to repeat verbal expressions										
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)										
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French										
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children										
	Simple play with accessory (symbolic play)							P			
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"		P	P		P	P	P	P		
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors						B	B	B	B	B
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors			P		B	r	P	P	P	P
	Complex turn off Aibo to sleep (symbolic play)							P			
	Speak directly to Aibo about Aibo's feeling (symbolic play)						P	P		P	P
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor					r		r		P	P
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor			r		B			r		r
	Cause-reaction play & basic pretend play, "caught on the act"										P
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter						P		P	P	
	Symbolic & pretend play Complex play with an accessory										
	Symbolic & pretend play Complex nap with Aibo										
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)										
	Causal composition of plays: Bite/Save & Give Food/Drink							P			r
	Causal composition of plays: Kiss & Bite/Save										
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors									P	
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor							P			

Figure 4.15: Child E. Play Grid. See Fig. 4.9 for a detailed caption.

Tornado (the name the child gave to the robot) can wag the tail today?" and Child E activated the right sensor at the first attempt and commented: "That's the tail one". Child E also proactively played the game of giving food or drink to the robot as well as a cooperative play situation of Bite/Save (see Fig. 4.16). Bite/Save play situation consisted of two chronologically steps: i) the robot bit the finger of either the child or the experimenter (through the use of the sensors) and ii) the person remaining (child or experimenter) saved the latter by freeing her/his finger: the freeing was done either by activating the sensor ("Complex Bite/Save") or by directly taking the finger out of the mouth of the robot ("Simple Bite/Save").

Furthermore, in Session 7, the child proactively combined 2 games, "Give food/-drink" and "Bite/save" and said: "He (the robot) is saying: give me a drink or I bite your fingers".

Another interesting play situation the child proactively experienced in Session 7 consisted of a competition between the robot and himself: both of them had to drink as fast as possible their invisible drink; the robot could only drink with the help of the experimenter (the experimenter was asked to activate the sensor linked to the opening of the mouth as fast as possible). At the end of the competition, Child E decided that the robot had won. Thus, in this play situation Child E experimented with:

a) dealing with rules of competition,

b) handling the temporal aspects of the game and the various chronological phases,

c) taking on the role of the participant (as a competitor) and the one of the organizer who announces the winner,

d) playing with abstract entities (invisible drink),

e) playing socially.

Concerning the reasoning dimension, it should be first noted that the child decided to rename the robot after the first session and call him "Tornado". Moreover, in the first sessions, most of his questions addressed the issue of the robot's technical capabilities and how to control the robot. In Session 2, for instance, the child said: "How is he doing that?" and "What's being on the head to make him walk?" (because when he touched the head and activated the head sensor, the robot walked). And later in the same session, while looking at the laptop he said "this must be the controller".

Figure 4.16: Child E. playing the game 'Bite/Save' with the experimenter. Chronological order of the photos: from left to right and top to bottom. First photo: the child activates the head sensor of the robot which make the robot open the mouth and enable the robot to 'bite' his finger. Second photo: the experimenter brings her hand close to the head of the robot in order to activate the head sensor. Third photo: the experimenter activates the robot's head sensor to make Aibo open the mouth in order to 'save' the child's finger; when the mouth opens, the child pull of his finger (third and fourth photos).

Furthermore, in Session 3, the child said: "I found how he might open his mouth"; the experimenter asked "is he moving the mouth?" and the child answered: "yes, when I stroke on the head, you see". This example illustrates that the child actively developed technical and causal reasoning about behaviours and capabilities of the robot. This questioning can be related to the category "Essence" and shows that the child considered primarily Aibo (Tornado) as a proper robot. It should be noted here that the child invented the concept of "invisible drink" as well as the way of calling it (very logically): "invisible robot drink". This illustrates the ability of the child to make links with real dog's life while adapting it correctly to the characteristics of robots.

The category "Mental state" was addressed during later sessions (from session 5 onwards). In session 5 the child actually said "he is wagging the tail"; the experimenter answered: "yes, that shows he is happy"; and the child replied "He likes me" and he stroked the robot. The experimenter reinforced the positive feeling: "yes, he likes you". That first step was expanded into the game "speak directly to Aibo about Aibo's feeling". In session 6 and onwards, the child addressed proactively the question of emotions but he tended to deal with a restricted repertoire of emotions only, such as " being scared" or "being terrified" (e.g. session 7 the child said: "You're scared Tornado, in fact you're terrified").

Child E dealt with "Moral standing" in session 5 when he accidentally kicked the robot and, in return, apologized to him directly ("Sorry Tornado") and comforted him by stroking him.

Finally, Child E addressed indirectly the question of "Social rapports" through play. For instance, in session 10, he conferred a specific role to the robot for the competition; the robot thus became his adversary, but on a very kind level, since the child decided at the end of the game that the robot had won the competition. Another example took place in Session 8 where the child asked directly questions to the robot (e.g. "Do you want to drink something Tornado?"). Then, he made the robot bark as an answer and the child "translated" the answer verbally for the experimenter: "He said yes". In this case, the child proactively played the social role of an intermediary position between the experimenter and the robot.

The dimension of affect (Fig. 4.17) appeared from Session 5 and onwards where the child proactively said "he (the robot) likes me". And the experimenter replied "Yes he likes you. You like him?" The child then answered "Yes". Then later, in Session 8, the child said "he (the robot) is very happy". The experimenter agreed with

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	
S2	
S3	
S4	
S5	· **[1i]** "Yes that shows he (the robot) is happy" (Experimenter); **"He likes me"** (Child E); "Yes he likes you" (Experimenter); · **[2i] "You like him** (the robot)**?"** (Experimenter); **"Yes"** (Child E)
S6	
S7	
S8	· **[1i]** "He (the robot) is very happy" (Child E) while making the robot bark; "Yes he is" (Experimenter)**,"Tornado likes me"** (Child E); "Yes he likes you" (Experimenter)
S9	· **[1ii] "Tornado is very friendly, isn't it?"**(Child E); "yes, he is"(Experimenter)
S10	

Figure 4.17: Child E. Events related to Affect. See caption of Fig. 4.10 for details.

him and then Child E added "Tornado likes me" and the experimenter reinforced the positive feeling: "Yes he likes you". In Session 9, Child E commented on the robot, qualifying him as 'friendly': "Tornado is very friendly, isn't it?" and the experimenter agreed verbally.

Child F. Child F was away for session 5. Thus he took part in 9 sessions. Note that on his explicit demand, session 7 and session 8 were not recorded (the experimenter had permission from the parents to videotape the child but she decided to value the child's request); thus information from sessions 7 and 8 is missing in the corresponding columns in the Play Grid. The Play Grid Fig. 4.18 shows that Child F engaged in social play almost all the time. He used verbal language a lot and progressively experienced some more complex levels of play notably pretend play with respect to "play with accessory". The first situations of "play with accessory" happened in Session 3. In this session, the child borrowed the mouse of the laptop and put it on the ground in front of Aibo at approximately 30 cm distance and asked the robot to touch the mouse with the paw. Then he activated the right sensor to make Aibo walk forward and approach the mouse. The child carried the robot for the 5 remaining centimetres separating the robot's paw from the mouse and finally the robot touched the mouse with his paw. Later, in session 4, the child experienced further situations of "play with accessory" in two successive steps. As a first step, he proactively played very simply with an accessory. For instance, Child F used the face of a character drawn on a piece of cardboard that he held in front of his face and told Aibo: "Stay

		1	2	3	4	5	6	7	8	9	10
L 1	*Solitary* Exploration										
	"Imitation" of robot's bark	P	P	P			P			P	
	Solitary mirror play – look at oneself in the robot's reflecting face										
L 2	*"Pre-social" or basic-social* exploration – stroke Aibo immediately after the experimenter (possibly basic imitation of the gesture)										
L 3	*Social* exploration (social play)	P	P	P	P		P			P	P
	Simple Bite/Save or Give/Food - no use of the sensors									P	P
	Position or locomotion game – with verbal qualification of the game	P			P		B			B	P
	Cooperative technical task: change the battery, or turn on/off Aibo	r		P	B		r			P	B
	Verbal order towards Aibo: e.g. "sit", "walk", "wake up"	P	P	P	P					B	P
	Basic pretend & social play – imitate Aibo's snoring & verbal comment										
	Basic play on affective gestures – give/receive a kiss and/or a lip to/from Aibo									P	P
	Repeat after me - ask the experimenter to repeat verbal expressions										P
	Look at Aibo through the camera (Possibly stroke Aibo & look at its reaction through the camera)			P	P		P			P	P
	Speak French with Aibo - e.g. "Hello" or "Bye-Bye" in French				r		B				r
	Show Aibo to other children (social play) Express verbally the willing/intention to show Aibo to the other children	P	P								
	Simple play with accessory (symbolic play)			P	P						
	Social Mirror play (social play) - look at oneself (and possibly at the experimenter) in the robot's reflecting face & express verbal comments, e.g. "Look at my arm!"										
	Social Hug – hug Aibo & ask the experimenter or the second researcher to hug Aibo										
L 4	*Complex* Give Food/Drink (cause-reaction play & symbolic play & social play) - use of sensors										
	Complex Bite/Save (cause-reaction play & pretend play & cooperative play) - use of sensors										
	Complex turn off Aibo to sleep (symbolic play)						P				P
	Speak directly to Aibo about Aibo's feeling (symbolic play)		P								
	Cause-reaction play & mental states: Ask a question to Aibo (e.g. identity, feeling), answer with a sensor		B	P	r		B				
	Cause-reaction play, Aim at a physical reaction of the robot, show it with a sensor		P	B	B		r			P	P
	Cause-reaction play & basic pretend play, "caught on the act"										
	Telling a story										
L 5	Cause-reaction play and explicit Social rapport: Ask a question to Aibo, answer with a sensor (e.g. press the sensor which opens the mouth), translate verbally the answer for the experimenter										
	Symbolic & pretend play Complex play with an accessory			P	P		P				
	Symbolic & pretend play Complex nap with Aibo				P						
	Symbolic & extrapolation play : "RobotCat" - Speak about the idea of a robotic cat (possibly imagine how one would play with it)									P	P
	Causal composition of plays: Bite/Save & Give Food/Drink										
	Causal composition of plays: Kiss & Bite/Save										P
	Pretend play & causal reaction & social rapports: Ask verbally Aibo to act a situation, use of sensors										
L 6	Pretend play & focus on Aibo's mental states: Mimic Aibo's cry, and explain Aibo is never crying but pretending to cry										P
	Pretend play & social rapports: Look after Aibo and set up rules										
	Pretend & symbolic & chronological play & social rapports: Search and rescue										
	Pretend & symbolic play & social rapport & cause-reaction play & chronological play: competition (drink fast) between the child or the experimenter and Aibo ; the non-competitor activates Aibo's sensor										

Figure 4.18: Child F. Play Grid. See Fig. 4.9 for a detailed caption. Note, Child F was away for Session 5 and, on his request, was not filmed during Sessions 7 and 8.

here Aivo, I've got something to show you". Note that the child slightly changed the pronunciation of the name of the robot and referred to Aibo as 'Aivo'. As a second step, later in the same session, the child proactively played a more complex accessory game with the robot, the "ghost dog". That play situation consisted in putting a cloth on top of Aibo and pretending Aibo was a ghost dog (Child F told Aibo: "You can be a ghost dog Aivo"); vocally, the child used classical onomatopeia mimicking ghost's "voice and presence". Moreover, in Session 6, the child decided to make the robot wear clothes and this game was expanded by:

a) a series of questions on inferring states of the robot with respect to like/dislike,

b) a direct communication with the robot to explain him what he was wearing (Child F told Aibo: "Look at you Aivo! You've got some paper on to be black");

c) a version of the game "aim at a physical reaction of the robot, show it with a sensor" (the experimenter asked "How do you make him walk with all these clothes?", the child replied "Walk?", and the child made the robot walk).

In addition to the accessory games, the child experimented with pretend play with the robot in a social context, e.g. pretending having a nap with the robot (in session 4) in a detailed (and complex) way resulting of:

1. using a cloth as a blanket to cover both of them,
2. deciding on the duration of sleep and asking for watching the clock to respect the time predefined for the nap,
3. pretending to snore,
4. both of them waking up again.

Besides, another way of tackling pretend play as well as robot's mental states happened in session 10 when the child imitated Aibo's crying, and then argued that Aibo was not crying but pretending to cry. And this notion of pretending to cry for the robot was reused many times during the last session (e.g. Child F said: "No, he's not crying, he is only pretending to cry.").

The reasoning dimension is really an important component of the profile of Child F. Child F principally addressed three of the four components, respectively, "Essence", "Mental States" and "Social Rapport", and, in minor importance, the issue of "Moral statement".

Concerning "Essence", the child really tackled the question of artefact or biological features, processes and categories. Categorywise, he often questioned about robot dogs boundaries, e.g. in Session 2: "Have you seen dogs that are not robot dogs, yes or no?" he asked the experimenter, and later in the same session: "He has short teeth, he doesn't bite. Robot dogs don't bite, do some do?"

The part on "mental states" is very rich since the child addressed all the aspects defined in the coding manual of Kahn et al. (2003) except probably the "autonomy" one. Actually, he attributed "intentions" to the robot in Sessions 1 and 2. He explicitly considered robot's "emotional states" in sessions 2, 4, 6 and 10. He also both tackled "emotional states" of the robot and his "personality" when he asked him questions about his likes/dislikes (e.g. Session 4: "Do you like toys Aivo, yes or no?"). Furthermore, he pretended the robot had some "cognitive abilities" and developed play upon it: in Session 4, for instance, he disguised himself with an accessory in order to "show" Aibo and thus presupposed -for the game- that Aibo could see. Later, in Session 6, again, the child presupposed for the game that the robot could see and told him: "Look at you Aivo. You've got some paper on to be black". The last aspect of "mental states" is the notion of "development" of the robot. Child F really questioned about it, from the very beginning of the sessions onwards. More than the notion of development, the child seems to have been willing to build a biography for the robot (i.e. the past of the robot) and therefore asked questions to the experimenter such as: a) in Session 1: "Where was this robot dog from?"; b) in Session 2: "Where was he born?" and "Has he travelled in a car?"; c) in Session 3: "Where did you get him from?", "Where does he live?", "How old is he?", etc.

Concerning the part on "Social rapports", the child really investigated the social links between the robot and the experimenter, who was considered by the child as being the "mum" of the robot (Child F told the experimenter "it's your dog son", meaning that Aibo is the experimenter's dog, and that the experimenter, in a way, is considered as being Aibo's 'mum'). Besides, he investigated the social links between the robot and himself, through situations of pretend play but also verbally. In Session 2 for instance, the child presupposed that there was a social rapport between the robot and himself since he told the robot: "When it is lunch time Aivo I got to go. And don't cry Aivo". Later, in Session 6, the child stated that the robot was his cousin: "Aivo is my cousin". And when the experimenter asked: "Aivo, do you like playing with F[7]? Can you tell me? Can you ask for his answer F?" then the child

[7]Child F is designed by F in the dialogue.

Session	**Events objectively related to Affect** (ordered chronologically with respect to first appearance, event only mentioned once per session)
S1	· **[1ii] "Ooh he is a nice dog"** (Child F) and he strokes the robot
S2	
S3	
S4	· **[1ii]** Child F brings a towel to put on the robot : **"Put this on Aivo, my dog, my friend, Aivo"** (Child F)
S5	
S6	· **[1i] "Aibo, do you like me?** You're my cousin. I'm your cousin Aivo**"** (Child F) · **[1iv]** Child F gives a **kiss** to the robot on the muzzle after saying **"OK, Goodbye Aivo, have a good sleep"**
S7	
S8	
S9	
S10	· **[1iv]** Child F has covered Aibo with a coat; he gives the robot **a kiss** on the forehead and says **"Goodnight Aivo"**

Figure 4.19: Child F. Events related to Affect. See caption of Fig. 4.10 for details.

told Aibo: "Aivo do you like me? You're my cousin. I'm your cousin Aivo". Besides, the child investigated beyond social rapports involving Aibo and, for instance asked the experimenter a few questions about her family: a) in Session 4, the child asked about the experimenter's French accent[8]: "What accent do you speak", which was further investigated in Session 6: "Why do you speak French?" and "Why were you born in France?"; b) in Session 6, he asked her about her family: "What are your parents' names?"; he investigated further questions on the experimenter's family in session 10.

On the affect level (Fig. 4.19), the child expressed himself a lot, both by gestures (e.g. giving a kiss to Aibo after saying "Goodbye Aivo, have a good sleep" in Session 6) and verbal expressions (e.g. in Session 4 when he dressed up Aibo: "Put this on, Aivo, my dog, my friend, Aivo"). It is perhaps worthy of note here that it might be the case that some gestures related to affect from a non-autistic perception (e.g. giving a kiss), do not have the same interpretation for a child with autism: for a child with autism, giving a kiss might, for instance, just be an imitated response. Concerning Child F, it might be the case that the child reproduced the gesture "giving a kiss" from a situation he had encountered or witnessed before; nonetheless it should be mentioned that his gesture was made proactively, with no previous reference from the experimenter to such a gesture.

[8]Child F masters some French vocabulary.

4.6 Discussion

Results from these experiments show that the children progressed differently, and that their profiles according to the three (intertwined) dimensions *Play* - *Reasoning* - *Affect* are unique. This highlights how the experimental approach presented in this study allows many trajectories for progressing and, more specifically, how it can meet the child's specific needs and abilities.
Furthermore, concerning the dimension of play, and, more precisely, concerning the children's progression with respect to solitary vs. social play, three groups can be highlighted. The first one, group 1, consists of children who mostly played solitarily and possibly encountered rudimentary situations of imitation, but no further components of social play. This group includes Child A who encountered imitation in session 10 and Child B. Note, both of them find it very hard to communicate verbally. For the children whose current play with the robot is mainly dyadic, it is particularly relevant to enable the robot to adapt automatically to their play styles in real time so that they can benefit from this dyadic play and progressively reach well balanced and potentially higher levels of play. This issue will be addressed in the next chapters of this thesis. The second group, group 2, consists of Child D who communicated mainly non-verbally yet progressively experienced situations of verbal communication and showed pre-social or basic social play during the last sessions. The third group, group 3, consists of Child C, E and F. Those children proactively played socially (i.e. in a triad including both the robot and the experimenter).
For those three groups, results shows that a) Child B (group 1) experienced progressively longer uninterrupted periods of play and engaged in basic imitation during the last sessions; b) children from group 3 tended to experience higher levels of play gradually over the sessions and constructed more and more reasoning about the robot (and sometimes experienced specific reasoning about real life situations as well). At a more basic stage, Child D (group 2) also experienced higher levels of play progressively. He started to reason about technical aspects of the robot as well, e.g. 'turning on/off' the robot and changing the battery. In the last sessions different elements suggested that he may also have experienced some reasoning about social rapport. Besides, the children's proactivity was encouraged, enabling them to take initiative and express intentions (cf. the proportion of proactive activities vs. reactive activities in the Play Grids).

These results are in agreement with Josefi and Ryan (2004)'s findings who have shown in the case study[9] they conducted that non-directive play therapy had encouraged the child's initiative-taking. Further to this, Josefi and Ryan (2004)'s study has shown that non-directive play therapy may encourage symbolic play, which is an important finding of our approach, too: In our study, children from group 3 progressively experienced situations of symbolic or pretend play. Note that, as already explained in Section 4.2, the study presented in this chapter took place in a therapeutic context but the experimenter was not behaving exactly like a therapist. Besides, we identify several advantages in introducing an autonomous robotic pet in the experimental setup (Josefi and Ryan (2004) used non robotic toys):

a) the use of a robot allows to simplify the interaction and to create a more predictable environment for play to begin with, thus facilitating the child's understanding of the interaction (e.g. by giving the robot a simple predictable behaviour to start with);

b) children tend to express interest in the robot, and occasionally affect towards Aibo, as our findings show;

c) here, one of the findings is that, in these experiments with this new approach, children tend to develop reasoning, and make comparisons to real dogs' lives through play with the robotic pet.

Thus, the robotic pet can be considered as a good medium for developing reasoning on mental states and social rapports upon, and for learning about basic causal reactions, too.

Davis et al. (2005) compared different robotic or computer platforms used in the Aurora project and compared their specific focus. She showed that mobile autonomous robots were adequate to unconstrained play situations, while the use of the humanoid robot Robota focused mostly on imitation of movements and gestures. However limited attention has been accorded so far to proper unconstrained play situations with an autonomous mobile robot and most experiments have been carried out using Robota and focusing on imitation. In Robins et al.'s studies with the non-mobile doll-like robot Robota (see Section 4.3), situations of child-robot interactions which actually happened were mostly restricted to situations of imitation of a gesture or a movement (Robins et al., 2004). Thus, even if the experiments may not have been qualified as such, they were in fact much more task-oriented, at least

[9] Josefi and Ryan (2004)'s experiments and results have been detailed in Section 4.3 of this chapter.

with respect to the dyadic child-robot interaction. Nadel et al. (1999) showed that imitation skills have a significant impact on the acquisition of social skills for children with autism. However focusing on imitation tasks only may not be sufficient when the child reaches some higher levels of play (cf. children from group 3 in the experiments presented in this study); Howlin and Rutter (1987) underlined the necessity of incorporating developmental aspects in pure behaviour principles.

Werry et al.'s trials (presented in Section 4.3) tended to encourage relatively unconstrained situations of play by using a mobile autonomous robotic platform (Werry et al., 2001; Werry and Dautenhahn, 1999). Shape, weight, sensors and the range of possible behaviours of the robot used (Labo-1, see more details in Section 4.3) are in contrast with Aibo's properties: Labo-1 is heavier, not pet-like, and has only a few sensors. Thus, interactions enabled by the robotic platform Labo-1 were very different in nature from the ones enabled by the use of Aibo. Besides, compared to Aibo's rich behaviour repertoire that we used in this study, the repertoire of behaviours of Labo-1 was fairly limited and situations of play were mainly approach and avoidance games. Note that in the present study, Child E most of the time asked for Aibo not to walk: the use of Aibo in trials enabled the child to play with Aibo either in a mobile or non-mobile mode (robot walking or non-walking mode), whatever the child prefers. Even when not walking, Aibo can still react in various ways (e.g. turning head, wagging the tail, barking etc.).

Moreover, in Werry et al.'s experiments, none of the experimenters participated in the experiments. The child played on his/her own with the robot (Werry and Dautenhahn, 1999), or two children interacted at the same time with the robot (Werry et al., 2001), but none of the experimenters did take part in the trials -they only responded to the child when the child initiated communication or interaction with them (Dautenhahn and Werry, 2002). In contrast, Robins and Dautenhahn (2006) started to investigate the role of the experimenter. Robins and Dautenhahn (2006) argued that the participation of the experimenter in the trials was necessary and described the experimenter's role as the one of a "passive participant" who responds to the children solicitation whenever they initiate interaction with him/her (Robins and Dautenhahn, 2006).

The study presented in this chapter goes beyond these previous experiments, since it provides the child with a relatively highly unconstrained environment of play: due to the mobile autonomous robotic pet, the child can engage in a larger repertoire of play situations (note that Robota is fixed in place) and notably experience causal reaction play and symbolic play. Imitation is used as a bootstrap to initiate more complex situations of interaction or to help the child reengage in the interaction. Besides, the experimenter is part of the trial and her role goes beyond the one described by Robins and Dautenhahn (2006) and is defined more precisely and formalized. In our method, the experimenter answers the child's solicitations and rewards him/her. In addition, her role is empowered under specific circumstances:

a) if the child is about to enter a repetitive behaviour, then the experimenter proactively intervenes to try to prevent the child from entering that repetitive behaviour or help the child change the game; note that "a)" aims at counterbalancing the fact that repetitive behaviours may not be considerably reduced by pure non-directive play therapy as stated in Josefi and Ryan (2004)'s study.

b) if the child does not engage in the interaction, then the experimenter encourages him/her to engage in playing with the robot,

c) if the game is "standing still" but the child has already experienced this play and has shown he/she is capable to play this specific game, then the experimenter can punctually intervene to give a better pace to the game;

d) if the child is about to reach a higher level of play but still needs some bootstrapping (or some guidance), then the experimenter can provide it;

e) the experimenter can proactively ask the child simple questions related to reasoning or affect such as: "do you think Aibo is happy today?" or "do you like playing with Aibo?".

Moreover, in this study, I have adopted a qualitative approach for the analysis of each dimension, Play, Reasoning and Affect. I was actually interested in the emergence and in the specificities of the play styles, questions or statements related to reasoning and events that could be objectively related to affect, rather than in the occurrences or the duration of each of them. In particular, two similar games might actually happen to be different in the way the child experiences them, such as for example, the fluency, the rhythm, the coherence etc. Consequently, unlike a quantitative analysis which often relies on micro-behaviour analyses[10] (Dautenhahn and Werry, 2002; Tardif et al., 1995), this qualitative analysis here focused on a bigger

[10]Micro-behaviour analysis is the analysis of videos based on the coding of low level behaviours such as eye gaze, eye contact, touch, etc.

scale, i.e. an intermediary scale[11]. This intermediary scale enabled us to consider events constituting a game as connected events and, in particular, to describe the structure of a specific play situation in possibly different (chronological) phases or identify in this play situation, the presence of social play, the proportion of symbolic or pretend play, and the use of causality.

The research presented in this chapter has provided novel insights into the methodology of using robots in robot-assisted play, going beyond previous work in this area. Results from trials with children with autism are very encouraging. Based on these results from an exploratory study, future research in this domain can extend and further develop and test this approach e.g. with larger user groups or specific control conditions that allow to specify in more detail the particular features of this approach that contributes to robot-assisted therapy for children with autism.

4.7 Conclusion

This chapter highlighted a new approach in the context of robot-mediated therapy with children with autism. This approach draws its inspiration from non-directive play therapy, notably encouraging the child's proactivity and initiative-taking. Beyond inspiration from non-directive play therapy, the approach introduces a regulation process. The experimenter, who takes part in the experiment, can indeed regulate the interaction under specific conditions detailed in Section 3; in brief:

a) to discourage repetitive behaviours,

b) to help the child engage in play,

c) to give a better pace to the game if it has already been experienced by the child,

d) to bootstrap a higher level of play,

e) to ask questions related to reasoning or affect.

A long-term study was carried out with six children which highlighted the capability of the method to adapt to the child's specific needs and abilities through a unique

[11] To make the parallel with the notion of micro-analysis used in (Tardif et al., 1995) that refers to the coding of micro-behaviours, one could qualify our approach here as a mesoscopic approach or a meso-analysis. The prefix 'meso' comes from the Greek word 'mesos', meaning middle. "Mesoscopic" is an intermediary scale between "microscopic" and "macroscopic". Those terms are commonly used in Physics and Chemistry, and can be transposed metaphorically to our context.

trajectory of progression with respect to the three dimensions, Play-Reasoning-Affect. In particular, each child made progress with respect to at least one of the three dimensions progressively over the sessions. Moreover, with respect to play, and, more precisely, solitary vs. social play, children could be categorized into three groups. Besides, the children who managed to play socially experienced progressively higher levels of play and constructed progressively more reasoning related to the robot; they also tended to express some interest towards the robot, including on occasions interest involving positive affect. This preliminary long-term study has therefore shown promising results for this new approach in robot-assisted play. It is a first study that potentially may be developed towards a new method in autism therapy.

Chapter 5

Real-Time Recognition of Human-Robot Interaction Styles

5.1 Introduction

5.1.1 Motivation

In the previous chapter, we have presented a novel approach for the play sessions, inspired by non-directive play therapy. We have shown through a long-term study how this approach could meet the specific needs and abilities of the children and encourage them experiment with progressively higher levels of play. Nevertheless, the study has highlighted that a few children (e.g. Child B) played a lot dyadically although they started to progressively experience some basic situations of social play, particularly in imitation games while stroking the robot. Besides, this dyadic interaction with the robot was mainly a tactile interaction.

It is therefore particularly relevant to enable those children to develop basic play skills through this tactile dyadic interaction, in order to help them reach progressively better balanced tactile interaction styles and higher levels of play. To this end, the robot should be able to appropriately adapt to the child's needs and abilities and to autonomously encourage the child's progress towards well-balanced and progressively more complex play styles. Combined with the novel approach in robot-assisted play presented in Chapter 4 where the experimenter takes part in the play sessions, such an 'adaptive' robot might also provide the child with additional opportunities to engage in triadic interaction with both the robot and the experimenter.

A first step towards this challenging goal of enabling the robot to guide the children towards well-balanced interaction styles is to enable the robot to recognize in real time the children's play styles. This is the main focus of this chapter.

5.1.2 Criteria to describe an interaction

An interaction can be characterised by various criteria. Here, in the contexts of autism and child-robot tactile interaction, the criteria should be very simple, in order to fit first levels of interaction.

Gentleness of the interaction This criterion refers to the forcefulness of an interaction. It may happen that a child touches the robot too forcefully. In this case, we want the child to learn to play less forcefully, which means, in a way, control the strength of its gesture towards other partners of interaction (a robot in this case).

An interaction is classified as 'gentle' if the participant strokes the robot gently, without signs of force. On the contrary, if the participant touches the robot with signs of force, then the interaction is classified as 'strong'.

Frequency of the interaction In our everyday life, we are involved in various interactions whose respective frequency can vary among a realistic spectrum. Here, we want the child to learn to play in a well-balanced frequency of interaction, i.e. not too low and not too high: if the frequency is too low, then the interaction is rare; at the other end, the higher the frequency is, the more difficult it may be, for a child, to understand the reaction to a specific stimulus. Thus, the frequency of interaction is categorised into four classes, defined by their typical periodicity of interaction:

- *very low* (S_0): the elapsed time between two tactile interactions is greater than 15 seconds. We will refer to it by saying that, for class S_0 the 'periodicity' is greater than 15 seconds.
- *middle inferior* (S_1): the elapsed time between two tactile interactions is lower or equal to 15 seconds and greater than 5 seconds. We will refer to it by saying that, for class S_1 the 'periodicity' is lower or equal to 15 seconds and greater than 5 seconds.
- *middle superior* (S_2): the elapsed time between two tactile interactions is lower or equal to 5 seconds and greater than 1 second. We will refer to it by saying that, for class S_2 the 'periodicity' is lower or equal to 5 seconds and greater than 1 second.

- *very high* (S_3): the elapsed time between two tactile interactions is lower or equal to 1 second. We will refer to it by saying that, for class S_3 the 'periodicity' is lower or equal to 1 second.

In this context S_1 and S_2 are considered as well-balanced frequencies of interaction, while S_0 corresponds to a rare interaction and S_3 to a very intense interaction.

In future work, other criteria, possibly more complex ones, could be investigated for the children who already master the first levels of interaction defined by these criteria.

5.1.3 Related Work

The role of tactile human-robot interaction in educational and therapeutic applications has been well highlighted by long-term studies with the seal robot Paro which have proven that specific everyday life situations exists in which human-robot interaction can have a positive effect on well-being of human beings (Shibata et al., 2005; Wada and Shibata, 2006) and even play a role in a therapeutic context of cognitive and physical rehabilitation (Marti et al., 2005). Tactile interaction is the primary means of interaction with the seal robot Paro, which is equipped with ubiquitous tactile sensors (Shibata, 2004). These sensors are sensible both to the pressure and the position on a flexible curved surface. Paro has internal states that influence its behaviour which can be proactive or reactive, i.e. in response to a sensor stimulation (Wada and Shibata, 2006). Moreover, the Huggable robot, a teddy-bear like robot, equipped with a full body sense of touch, has proven to be a promising support to investigate the quantitative characterisation of social affective content of touch (Stiehl et al., 2006).

Besides, offline characterisation of interaction styles in general, has been investigated recently with diverse approaches. Scassellati (2005b) focused on providing quantitative and objective measurements to assist in the diagnosis of autism. Measurements refer to the position in the room, vocal prosody and gaze pattern – whose characterisation relies on Linear Discriminant Analysis. Kanda et al. (2002) conducted a study that highlighted the feasibility to link quantitative robot's and human's data characterizing body movements with a subjective evaluation made by the participant. Later, Salter et al. (2006) showed the possibility, in the context of child-robot interaction, to reflect some traits of personality of the children with an offline clustering technique based on the empirical probability distribution of the activation of the sensors.

Concerning real-time classification of interaction styles, Salter et al. (2007) have presented a real-time simple recognition algorithm for four interaction styles ('alone', 'interacting', 'carrying' and 'spinning') using the robotic platform Roball. The algorithm is based on a decision tree whose conditions are set up manually, by visual inspection of sensor data. Moreover, Derakhshan et al. (2006) have developed a real-time classification algorithm of interaction styles for children playing on an adaptive playground that is made of tiles equipped with sensors. The algorithm relies on a multi-agent system approach of BDI (Belief—Desire—Intention) in combination with neural networks using supervised learning. It shall be further noted that in the slightly different context of gesture recognition, Hidden Markov Models have been used quite a lot for real-time recognition (e.g. Kim et al. (2007); Lee and Xu (1996); Calinon and Billard (2004)). An HMM is specified[1] by its number of hidden states, the initial state probability distribution and the two following probability matrices: the transition matrix, describing the conditional probability, given the state S at time step t, to be in the state S' at time step $t+1$, and the emission matrix, defining the conditional probability of emitting a signal O, given the state S. Current research on gesture recognition, e.g. as in Lee and Xu (1996), Calinon and Billard (2004) and Kim et al. (2007), usually implicitly refers to homogeneous HMMs, i.e. for a given HMM, the transition matrix does not change over time[2]. Classifying an observation with HMMs consists in finding, among all the different HMMs[3], the one which has the highest probability of emitting this observation (Lee and Xu, 1996).

5.1.4 Artifact and Sensors

Any robot or embodied agent situated and acting in an environment should have sensors through which it can receive information about itself and about its surrounding environment. Examples of sensors dedicated to the surrounding environment are: visual sensors, infra-red distance sensors, sonar sensors. Those sensing the internal state of the robot are, typically, motor position, internal temperature sensors and gyroscopic accelerometers. At the border between the external surrounding environment and the internal state of the robot, the tactile sensors play a major role, providing possibly information on the internal state of the robot. Also, more impor-

[1]For more detail on HMMs the reader can refer to Rabiner (1989).

[2]Note that homogeneous HMMs are usually simply referred to as HMMs. In contrast to homogeneous HMMs, with inhomogeneous HMMs, the transition matrix varies over time. One example of application of inhomogeneous HMMs is presented in Borodovsky and McIninch (1993) for gene finding.

[3]One HMM per class to distinguish.

tantly for our study here, some of them, which we call here external tactile sensors, can give information on the tactile interaction the robot is involved in as an embodied agent.

Artifact and external tactile sensors The robot used in the whole study is the AIBO ERS-7 commercialised by Sony. It is equipped with five external sensors, namely, the chin sensor, a boolean sensor, and four continuous sensors, the head sensor, the front back sensor, the middle back sensor and the rear back sensor. These sensors are the ones directly involved in tactile human-robot interaction.

Normalization of the sensor values Consider a sensor S, whose values s_t (continuous or discrete) are comprised between S_{min} and S_{max}, $S_{min} \leq S_{max}$. Then, s_t can be normalised as follows:

$$norm(s_t) = \frac{s_t - S_{min}}{S_{max} - S_{min}} \tag{5.1}$$

Global Variable From the individual external sensor data, a global variable G can be built. The global variable is obtained by summing, at a moment t_0, the sensor values that must have been either normalized or binned beforehand. G removes the spatial information on the sensor data but allows a simplification of the input data to analyse.

Quantitative Binning Binning the data principally enables one to reduce the complexity of the data to analyse by grouping them under specific constraints. The binning is a mapping from a discrete or a continuous space, to a subset of $\mathbb{N}$, $B_N = [0, 1, ...N - 1]$, where $N > 0$ is the cardinality of the binning (i.e. the number of bins). The binning can be done according to various criteria and the choice of these criteria is of great influence for a further analysis. For instance, bins can be defined according to the value, and in such a way that all bins have the same size, in terms of the range of possible values. In this case, a sensor s_t, with normalised value $norm(s_t)$, is mapped into a bin $b(s_t)$ as follows:

$$b(s_t) = IntegerPart(norm(s_t) * N) - 1 \tag{5.2}$$

where $IntegerPart$ is the function extracting the integer part of a real number.

The binning considered in this thesis is different from the one defined in Equation 5.2. In the present thesis, the data are still binned according to their value (quantitative binning) but the bins are defined in such a way that the probability to belong to a bin i, (except for one bin which represents the null value), is empirically equilikely. A bin i is described by its extreme values, (respectively the lowest $b_{min}(i)$ and the highest value $b_{max}(i)$). These values are defined by the empirical probability distribution of the variable that is binned[4]. Thus a normalised sensor value, $norm(s_t)$ belongs to the bin i if and only if $b_{min}(i) \leq norm(s_t) \leq b_{max}(i)$.

Time Series We consider any temporal variable X (typically a sensor or the global variable G in our context) observed from the moment t_0 over a temporal horizon h, $h > 0$ with the sequence of values $[x(t_0), x(t_0 + 1), ...x(t_0 + h - 1)]$. This sequence is a time series of the variable X.

5.1.5 Summary

This section has explained the motivations and rationale for enabling a robot to adapt to specific play styles of the children. A main motivation is to guide the children towards better balanced interaction styles (e.g. for the criterion 'Gentleness', neither too forceful nor too weak strokes) and, progressively, towards higher levels of play. We have then presented the two criteria of interaction which we will focus on in this thesis: the 'gentleness' and the 'frequency' of the interaction. Further to this, we have introduced the artifact and its sensors, in particular the external tactile sensors. The external tactile sensors are, in this study, the sensors involved in the characterisation of the interaction styles. Basic definitions useful for the preprocessing of the sensor data close the section.

5.2 Classification with Self-Organizing Maps

This section describes an early approach I adopted (François et al., 2007), using self-organizing maps (SOMs) for the classification of interaction styles. Because I used no a priori knowledge on the structure of the input data, this unsupervised non-linear mapping deemed to be a good first approach. In this section, I present a proof-of-concept of the method which led to a good recognition of the different styles for the criterion 'Gentleness'. However, the recognition was made with a fairly high delay that I then attempted to cut down.

[4]In our context, the extreme values describing the bins are determined by trials with the real robot: the empirical probability distribution of the sensor data is based on those experiments.

5.2.1 Preliminary Trials

Preliminary trials were conducted in order to generate, under controlled laboratory conditions, the prototypes for the interaction styles. During each trial, the participant interacted either on a gentle or on a strong mode exclusively with the robot. Five runs generated by two different adults were used for this study, three on the gentle mode and two on the strong mode of interaction. Each run lasted around five minutes. Every $32ms$ the robot Aibo sent an update of its sensors' values through a wireless LAN to the laptop. Finally, for each run, 9532 updates of the sensors' data were used for the further analysis, which means a total of 47660 updates of the sensors' values. Values were stored by runs and for each run sequentially (chronologically).

In this preliminary study, the five external sensor data were involved. At each time step, the value of each sensor data was binned and the five binned data were summed into a global variable G. Note, since the original focus was on information theory, binning was applied to the sensor data. This binning had not been removed for the first step of this study. In a second step where I optimised the delay for the recognition of the interaction styles, the binning had been removed, but this did not significantly change the results. I was interested to keep this binning because I had in mind to possibly go back to information theoretic techniques in a further step of this work.

Visual inspection of the variable G showed important differences of shape between time series originating from strong interactions and the ones originating from gentle interactions (see an example in Fig. 5.1 and Fig. 5.2). The ones originating from gentle interactions were characterised by typical blobs while those originating from strong interactions showed typical peaks.

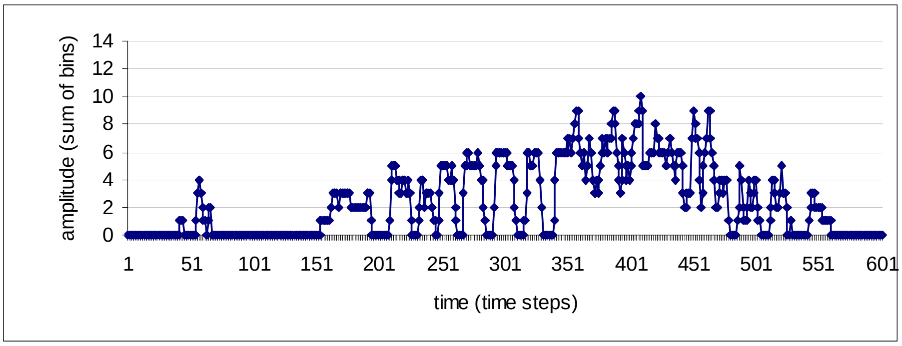

Figure 5.1: Typical time series from a gentle interaction.

This characterization by blobs and peaks in addition to the property of shift

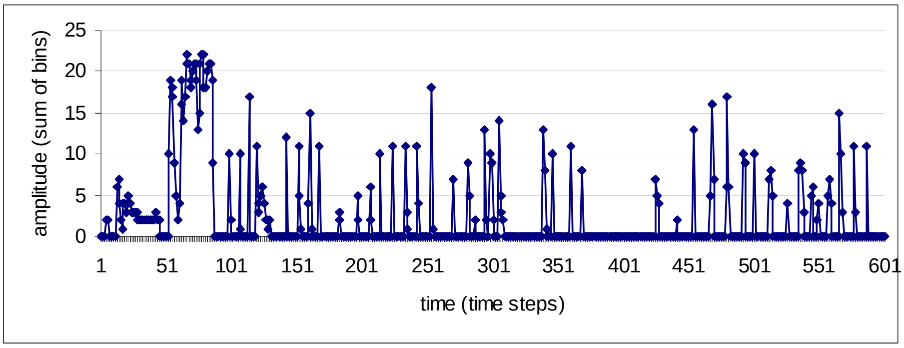

Figure 5.2: Typical time series from a strong interaction.

invariance led me to apply Fourier Transform. For each interaction style, the modulus of the Fast Fourier Transform (FFT) was calculated componentwise for a sliding window on G values. This preprocessing led to a new high-dimensional space to analyse (we call this space 'FFTInputSpace' and the length of the vectors resulting from this preprocessing is n). I used no other knowledge about the FFTInputSpace and sought a method to classify the data of this space according to the criterion gentleness[5]. Moreover, I wanted to get some insight on how the data were spatially organised in the FFTInputSpace. That is why I decided to use the Self-Organizing Maps (SOMs) which provide a non-linear projection from a high dimensional space to a lower dimensional space and is topology preserving (Kohonen, 2001).

5.2.2 Self-Organizing Maps

Self-Organizing Maps provide a non-linear projection from a high dimensional space to a lower dimensionality space and is topology preserving (Kohonen, 2001). In this study, SOMs were used as a classifier, i.e. as a method which classifies data from the high dimensional space (Wünstel et al., 2000).

SOMs are a specific class of Artificial Neural Networks, which rely on unsupervised, competitive learning. In a Self-Organizing Map, the neurons are placed at the nodes of a lattice that is usually one or two-dimensional. The neurons are selectively tuned to various input patterns in the course of a competitive learning process. A specific weight, from the same dimension as the input data, is attached to each neuron (node) of the lattice. Each node is connected to the adjacent ones according to a neighbourhood rule which derives from the topology of the map. The SOM is used

[5]The input vectors which originated from a 'strong' interaction should be separated from those which originated from a 'gentle' interaction.

in two phases: the training phase during which weights of the nodes are updated and the mapping phase, during which the classification or categorization of data can be made.

5.2.2.1 Training phase

First of all the network is initialized (either by random initialization, by initial samples, or through linear initialization). This process defines initial weight vectors, one for each node[6] of the network. Then, input data[7] are presented one by one to the network by a random selection. For each input data v, the distance between v and each weight of the network is measured according to a predefined metric[8] between the input vector and each node of the network. The node i^* minimizing the distance d is called the Best Matching Unit (BMU). Its corresponding weight $w_{i^*}(v)$ satisfies the following equation:

$$w_{i^*}(v) = \arg\min_{w_i} d(w_i, v) \tag{5.3}$$

Afterwards, the weights are updated according to the following equation:

$$w'_j = w_j + \epsilon(t) \cdot h_t(j, i^*) \cdot (v - w_j)\ ,\ j = 1, ..., \|\mathcal{K}\| \tag{5.4}$$

where:
w_j is the weight for the node j
w'_j is the updated weight for the node j
$\|\mathcal{K}\|$ is the size of neighbourhood $\mathcal{K}(i^*)$ for the winner node i^*.
$h_t(j, i^*) = exp(\frac{-d(j,i^*)^2)}{\sigma(t)^2})\ \forall j \in \mathcal{K}(i^*)$
ϵ and σ are monotonic decreasing functions of time.
To simplify, we can say that time t being static, the closer a node is from the BMU, the more it will learn; and globally, the network will learn less and less when time t is growing.

The whole process of finding the BMU, identifying the neighourhood and updating the selected nodes is often identified as a three step process:

1. *competitive process*: selection of the BMU,
2. *cooperative process*: the BMU is the center of a topological neighbourhood of cooperative neurons; the latter is defined by the choice of h_t,

[6] A node is also called a unit.
[7] These input data are the training data.
[8] The metric usually used is the Euclidean distance.

3. *adaptive process*: this is the update of the weight of the neurons according to equation 5.4.

The presentation of the entire set of input data constitutes what we call an epoch. A training phase results from the succession of many epochs. The learning-rate parameter ϵ, in the upper equation, is decreasing over-time and contributes to a convergence of the feature map. Kohonen (1982, 1997) actually described the existence of two successive phases in the adaptive process:

1. a *self-organizing phase also called ordering phase* during which the topological ordering of the weight vectors happens;
2. *the convergence phase* during which the feature map is finely tuned.

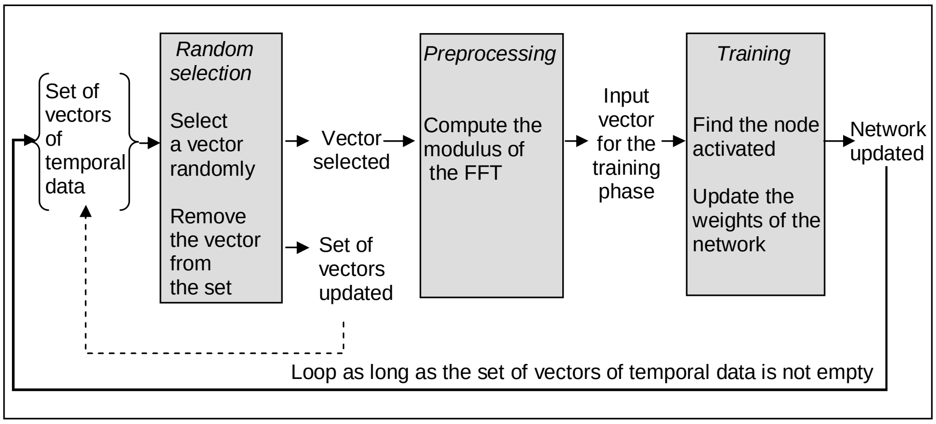

Figure 5.3: One Epoch of training.

Fig. 5.3 represents one epoch of training of the SOM in the context of this study.

5.2.2.2 Mapping phase

Once the network has been trained, it can be used for classifying (categorizing) data from the same space as the space where the input data used for the training phase come from. A data from this space is presented to the nodes successively. The node activated is the node corresponding to the BMU with respect to the same metric as the one used for the training phase (Eq. 5.3).

5.2.3 The recognition algorithm

Preprocessing for the SOM The temporal sensor data G was preprocessed in order to be used in the process of classification. The preprocessing resulted from the computation of the componentwise modulus of the FFT on a sliding window on G.

Characterisation of the nodes of the SOM feature map After the training phase of the SOM, a 'preliminary' mapping phase was conducted with the training set of data (Section 5.2.2.2) in order to characterize each node of the SOM feature map. The characterization was made according to the following criteria:

- A node activated by data originating from gentle interactions only was renamed *gentle node*.
- A node activated by data originating from strong interactions only was renamed *strong node*.
- A node activated by both types of data data was renamed *hybrid node*.
- A node never activated was renamed *null node*.

The training phase (Fig. 5.3) and the characterisation of the nodes of the SOM feature maps were made off-line.

Recognition of the interaction styles The algorithm operated on-line for the recognition of the interaction styles and the real-time adaptation of the robot. The whole process of classification of the interaction styles is summarized in Fig. 5.4. At each classification, the mapping phase of the SOM led to the activation of one specific node. According to its characterisation (gentle, strong, hybrid or null) the current state (gentle or strong) that determined the robot's behaviour was updated as follows:

- The initial state of the interaction was set to gentle
- If the node activated on the SOM feature map was a gentle node, then the current state was gentle;
- If the node activated on the SOM feature map was a strong node, then the current state was strong;
- If the node activated on the SOM feature map was a hybrid or a null node, it resulted in no change in the current state (as for the robot's behaviours, they did not change).

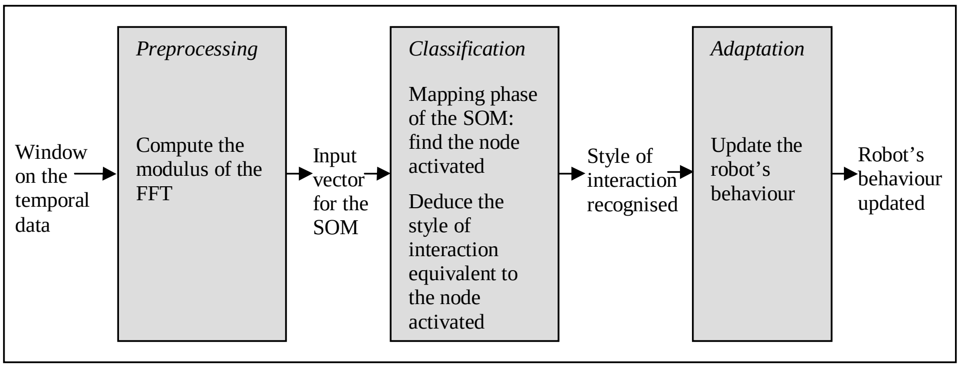

Figure 5.4: Unit process of classification and adaptation to the interaction style.

5.2.4 Implementation

5.2.4.1 Communication process

The robot used in this study was the Sony Aibo ERS-7. Its control programming was achieved using URBI, Universal Real-Time Behaviour Interface, (Baillie, 2005). Sensor/motor data were transmitted through a wireless LAN to a laptop. The Aibo sent current values of its sensors every 32ms. The laptop analyzed periodically the sensor data, classifying on-line the interaction correspondingly and sending the information back to the Aibo which then changed its behaviour accordingly. The process of classification of the interaction was written in Java.

5.2.4.2 Parameters for the process of classification

The parameters had to be finely tuned experimentally. Input vectors had to be sufficiently big to give enough information about the mode (gentle or strong) it originated from, so that it would be easier to 'separate' gentle and strong input vectors on the SOM Feature Map. This had to be well tuned with the size of the Feature Map, which, in a way, influenced the scale of the mapping. Different input vector's sizes and network's sizes were experimentally tested. The best results were obtained for an input window on temporal data from size 512 and a rectangular network of size 10*10. Each component of the input vector for the classification with the SOM was respectively the component of the modulus of the vector resulting from the FFT. The network was randomly initialized and trained (off-line) with 5 epoch.

Once the training phase had been finished, the behaviour classification was made on-line, the FFT algorithm being computed on-line as well as the activation of nodes for the SOM. However, since this process was time consuming, and since the modulus of the Fourier transform did not change significantly over a few time steps, I decided to set a frequency which would be more suitable. Experimentally it was found that updating the modulus on the FFT once in 120 updates of the sensor data was fast and precise. After each update of the interaction state through the classification process, the Aibo was informed of the result in order to adapt its own behaviour on-line.

5.2.5 Validation of the model

5.2.5.1 Validation of the topology of the SOM map

In order to have a precise and coherent classification of the interaction, a necessary condition was that the SOM map clearly distinguished topologically two regions, one corresponding to the gentle nodes and the second regrouping the strong nodes. Moreover, the proportion of hybrid and null nodes had to be very low compared to the proportion of gentle and strong nodes so that there were not two many cases in which the Aibo was not able to really ‘decide’ between strong and gentle interaction and relied on the previous state detected. Besides, hybrid nodes should rather be on the border or next to the border between gentle and strong regions (rather than in the inner part of the regions); this would suggest a smooth transition between the two regions.

Two SOM maps were successively trained. Both of them gave good results (see Fig. 5.5 which provides a graph of the first map). For each of them, the number of hybrid nodes was respectively 9 and 7 out of 100, while the number of null nodes was respectively 1 and 0. For the first map, all the hybrid nodes were on the border. For the second map, 3 hybrid nodes were not directly on the border but 2 of them were first neighbours of border nodes and the third one was second neighbour. This corresponded to a smoother transition between the two regions.

5.2.5.2 Validation of the real-time classification and adaptation

In order to evaluate the capability of the algorithm to recognize the interaction styles correctly within a short delay, four different trials were conducted whereby the algorithm operated online and the robot updated its behaviour accordingly in real time. The set of Aibo’s possible behaviors remained the same in the four experiments: Aibo was standing and waiting for at least one of its five external sensors to be activated. Whenever one of the latter sensors was activated, it a) wagged the tail if it had detected a gentle interaction, or b) barked if it had detected a strong interaction. We shall remind the reader that the activation of a hybrid or null node in the SOM feature

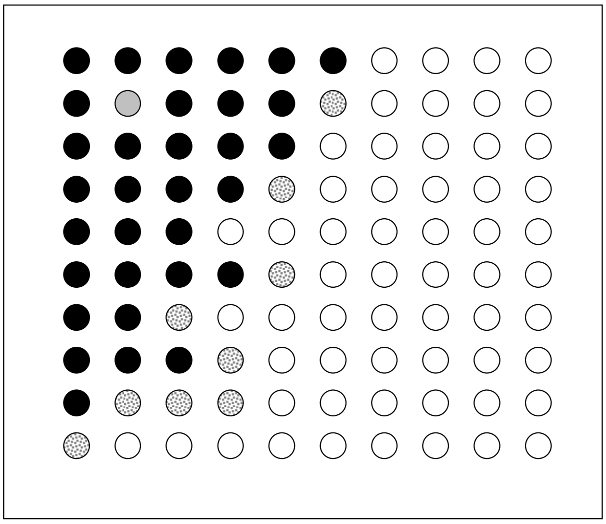

Figure 5.5: Map of a SOM. Legend: white for a 'gentle' node, black for a 'strong' node, stripes for a 'hybrid' node, blobs for a 'null' node.

map was interpreted as the state remaining the same as before (i.e. gentle if it was gentle, strong otherwise): in terms of robot's behaviour, the robot's current reaction to tactile stimuli remained the same. The robot's initial state corresponded to a gentle interaction. The rationale behind this choice was as follows: in child-robot play we want the robot to be able to maintain a well-balanced interaction style, i.e. for this criterion Gentleness, the interactions should be neither too forceful nor too weak. In this work, barking was used as a representative behaviour that might induce a human to back off, thus calming the interaction. Wagging the tail was used as an indicator to encourage interaction. Note that this proof-of-concept was carried out before I started the trials with the children with autism. Thus, those behaviours should be regarded as representative from the point of view of typically developed adults and are not representative of the specific interests, needs and abilities of children with autism. During the long-term experiments with the children I have conducted after this preliminary study, the concrete choice of the behaviours of the robot have been specifically tailored towards each child's interests and abilities (Appendix C).

The succession of interaction styles detected by the robot and the corresponding node activated on the SOM map were stored. According to the experiment, the participant had to play either gently or strongly, or alternating gentle and strong interactions. The participant had to maintain the same style of interaction until the Aibo had classified and adapted to this style. Each time the participant changed her way of interacting with the robot, the time at which it happened was stored as well as the time at which the Aibo adapted its behavior accordingly. Note, this study is a proof of concept; future work should cope with more frequent changes in play style, since child users will not be instructed how to play.

Experiment 1: This experiment tested the capability of the algorithm to recognize the style of interaction (i.e. gentle or strong) in the simple case where there was no transitions between gentle and strong strokes. This first trial therefore only tested the correctness of the recognition of the style and did not focus on the delay in the recognition process.
Two runs of three minutes each were conducted. In the first run, the participant interacted with the robot only gently. In the second run she interacted with the robot only in a strong style. Each run gave rise to 42 classifications computed by the algorithm which resulted in:

i) In the first run (gentle interaction), 39 activations of a gentle node, 3 activations of a hybrid or null node, no activation of strong nodes;

ii) In the second run (strong interaction), 41 activations of a strong node, 1 activation of a null node, no activation of gentle nodes.

We shall remind the reader that whenever a null node or a hybrid node was activated, the decision-making process led to the current state remaining the same as before (i.e. 'gentle' if the previous style recognised was gentle, 'strong' if the previous style recognized was strong). Consequently, in the first run, the current state was always gentle, and in the second run, the current state was always strong: in the two trials, there were no error in the recognition of the interaction styles and the robot adapted correctly to them.

Experiment 2: The purpose of this experiment was to test the Aibo's capability of adaptation over time involving all five tactile sensors. This experiment focused principally on the delay for recognizing the interaction style. The participant was asked to interact during eight minutes with the robot, alternating gentle and strong strokes. The participant tagged each stroke (i.e. she named the style of the stroke while interacting, e.g. 'gentle' or 'strong'). This process refers to the 'interaction subjectively evaluated' displayed in Fig. 5.6. Moreover, the participant was asked to vary the sensors she activated. Results showed that the robot adapted correctly to the interaction but with a fairly high delay which varied between $10s$ and $19s$. Fig. 5.6 shows an example of the dynamics of the robot's adaptation to the interaction for four transitions in the interaction style. It compares the transitions in Aibo's behaviour (as a consequence of adaptation) to the changes in the participant's behaviour scored subjectively.

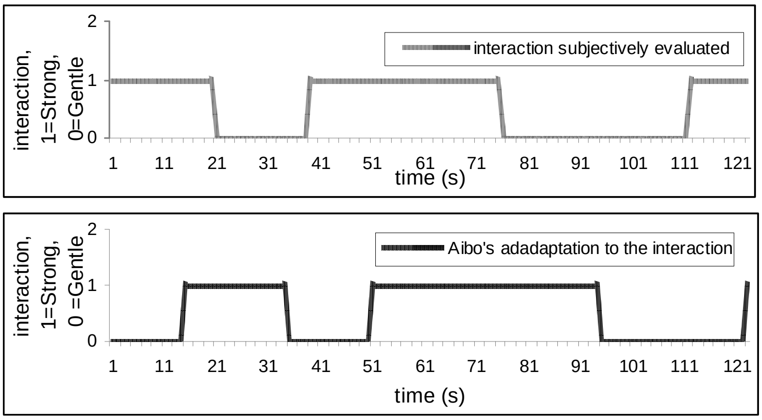

Figure 5.6: Example of the dynamics of the robot's adaptation to the interaction: The first graph represents the real interaction style over time; the second graph shows the robot's adaptation, which is accurate with a delay. On the y-axis, 1 stands for 'strong' and 0 for 'gentle.

Experiment 3: This experiment aimed at testing the impact of the nature of the sensor (continuous or boolean) on the recognition process. The two previous experiments have shown that, when the participant varied the sensors that she activated, the robot could adapt correctly to the interaction style although the delay to adapt was pretty high. However, because the sensors involved in tactile interactions could be from two kinds here (continuous or boolean) the algorithm had to be tested with each sensor separately. This setting was particularly inspired by situations of repetitive play, where children pursue the same action for a long time. Children with autism sometimes play in a repetitive way; thus, a situation where a child carries on touching the same sensor for a while was likely to happen in applications involving children with autism and therefore had to be tested beforehand in laboratory conditions.

Five runs were conducted. For each of them, the participant had to touch only one sensor, respectively the chin sensor, the head sensor and the three back sensors. For each run, the participant could vary the style of interaction (from gentle to strong, from strong to gentle) whenever she wanted. She tagged each stroke ('gentle' or 'strong'). Results showed that the Aibo adapted correctly to the interaction for trials involving the activation of a continuous sensor. However, the Aibo had diffi-

culties to adapt correctly to the interaction styles when the sensor was boolean. In the case of the boolean sensor, two types of errors indeed happened: a) the Aibo was not able to detect a gentle interaction within 1 minute (1 minute is a long time compared to the average time of adaptation to a new interaction style), or b) the Aibo had detected a gentle interaction for a very short time (around 4 seconds), the participant was keeping interacting subjectively gently but the Aibo started barking, which means that the algorithm had recognised a strong interaction style. This error in the recognition of the interaction style for the boolean sensor could be explained by the fact that its binned value could only be 0 or 9 and that the model used for classifying the data took mainly two factors into account: i) the relative modulus of the frequencies of the Fast Fourier Transform of one vector of sensor data indicates which frequencies are predominant; ii) the FFT respects the linear property.
This experiment showed therefore a limitation of this model. The relative percentage of activation of boolean sensors should be relatively small compared to the one of continuous sensors. Experiment 1 and 2 have shown that in the case of one boolean sensor in five where there was a fairly well-balanced repartition of the strokes among the different sensors, the recognition process worked correctly.

Experiment 4: In the present experiment, in order to avoid the risk of having errors induced by the use of a boolean sensor the participant had to respect the constraint of not touching the chin sensor, but she could touch all the four other sensors. She could change from one style of interaction to another (gentle, strong) whenever she wanted but she tried to vary the duration of time between the time the Aibo adapted to the current interaction and the time she changed the interaction afterwards. The idea was to check experimentally that the delay of adaptation was not directly influenced by the rhythm of changes in the subjective interaction. This idea is linked to the fact that we used a finite sliding window on the data to classify the interaction. This means that we took into account only a limited history of the interaction.
This experiment consisted in one run of seven minutes, whereby the participant alternated longer periods of changes in behaviour and shorter periods of changes. The longest duration of an interaction was 50 seconds, the shortest was 17 seconds. Fig. 5.7 represents on the x-axis the duration of a style of interaction and on the y-axis the delay of the adaptation to the next interaction style (e.g. length of gentle interaction and delay to adapt to the next kind of interaction which will be strong). The graph shows that there was no linear relationship between the period of changes

in behavior and the delay for adaptation.

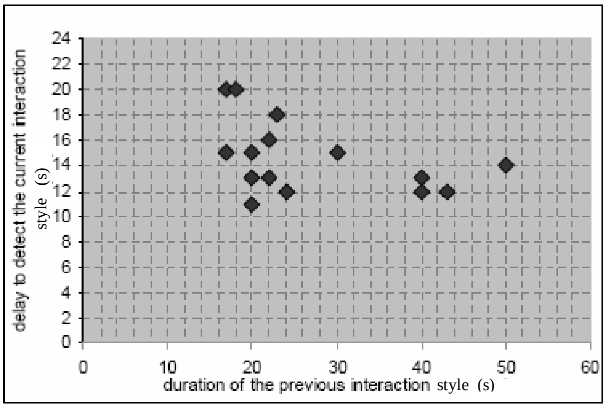

Figure 5.7: Experiment 4 : the graph represents the delay in the process of classification and adaptation (y-axis) and the corresponding duration of the previous interaction style (x-axis).

5.2.6 Discussion

These experiments have shown that the model was capable to classify pretty well the interaction styles, but with a remaining important delay. It is a main issue in human-robot interaction, because the robot should be able to give feedback and/or adapt rapidly enough in order to make the whole interaction process consistent with a human-centred perspective and worth maintaining.

I tried to optimise[9] the delay by reducing the size of the input window on the temporal data, thus limiting the amount of information the algorithm was dealing with to a shorter past (a window of size 512 corresponds to 16.4 seconds of interaction). However, while reducing its size, the SOM Feature Map was affected: the separation between zones for gentle and strong nodes became less and less obvious, the amount of null nodes and hybrid nodes increased. A null node is a node not activated during the training process. A hybrid node is a node activated by two types of input data used in the training phase, some originating from gentle interactions and others generated during strong interactions. In order to counteract this, a postprocessing was applied (first manually, and then automatically), in order to confer additional meaning to the map. This postprocessing was defined empirically under a probabilistic constraint. For this, additional runs with gentle and strong styles were generated

[9]For this optimisation I only considered the four continuous sensors. I tested it (separately) with binned sensor data and with normalised continuous data.

under controlled conditions. Null nodes that were activated with these new input data were tagged according to the following rule: a null node activated was renamed ‘gentle’ (respectively ‘strong’) if it was activated more often for gentle (respectively strong) behaviours than for strong (respectively gentle) behaviours. The delay could therefore be reduced to a few seconds (approx. 3-4 seconds). However, this method required important hand-tuning which made the solution very specific to one setting. This is an important issue for our main goal which was to find a method which enabled to classify not only strong and gentle behaviours, but could also be easily used to classify other criteria from tactile interaction, with the least ‘hand-tuning’ as possible.

5.2.7 Other methods

In order to find a more generic solution, I investigated other techniques using different approaches, i) the testing of the linear separability of the data with the Fisher Linear Discriminant Analysis (LDA) (Elizondo, 2006), and ii) the extraction of common features from the input data by compression (Cilibrasi and Vitanyi, 2005). In order to extend the domain of criteria, I also tested these methods with data generated for the specific context of the ‘frequency of the interaction’; data were generated in the same way as those for the criterion gentle/strong: runs with exclusively one class (very low, middle inferior, middle superior, very high) were conducted in order to collect data specific to each class for this criterion.

5.2.7.1 The Fisher Linear Discriminant Analysis

This method enables to test the linear separability of different classes. It relies on an informative projection of the data on an axis that satisfies the following condition: when the data from the classes are projected onto this axis (defined by its direction w^*), the ratio ‘distance between the classes projected’ to ‘distance of the projected cases within a class’ is maximized. This condition corresponds to a maximization of the Fisher Criterion (Equation 5.7).

For each class c, we consider the number of cases, N_c, the mean of these cases[10], μ_c , and its *within class scatter matrix*, S_W, defined by:

$$S_W = \sum_c \sum_{i \in c} (x_i - \mu_c)(x_i - \mu_c)^T \tag{5.5}$$

[10]The mean μ_c of the cases originated from a class c is defined by: $\mu_c = \frac{1}{N_c} \cdot \sum_{i \in c} x_i$.

The total number of cases over different classes is N and the mean of these cases[11] is μ. The *between class scatter matrix*, S_B, is defined by:

$$S_B = \sum_c N_c \cdot (\mu_c - \mu)(\mu_c - \mu)^T \tag{5.6}$$

The Fisher criterion is defined by the following equation:

$$J(w) = \frac{w^T \cdot S_B \cdot w}{w^T \cdot S_W \cdot w} \tag{5.7}$$

For two classes c_1 and c_2, with means, respectively, μ_1 and μ_2, a vector w^* that maximizes Equation 5.7 satisfies the following equation[12]:

$$\exists \alpha \in \mathbb{R} \ \ such\ as \ \ w^* = \alpha \cdot S_W^{-1} \cdot (\mu_1 - \mu_2) \tag{5.8}$$

Among the set of possible solutions of Eq. 5.8, we take w^* for which $\alpha = 1$. In order to evaluate whether the two classes are well separated with this projection, we now consider the hyperplane[13] defined by the normal vector w^* and the bias b. The bias b is determined in such a way that the hyperplane lies between the two means of the training data projected onto direction w^*. Thus, b satisfies the following equation[14]:

$$b = -\frac{1}{2} \cdot (< w^*, \mu_1 > + < w^*, \mu_2 >) \tag{5.10}$$

where $< w^*, \mu_c >$ is the scalar product between w^* and μ_c (which corresponds to the projection of the mean of the cases from class c onto direction w^*). If the two classes were separated by the LDA, then two decision regions would be separated by this hyperplane defined by its normal vector w^* and bias b, i.e. the hyperplane would separate the space in two regions, each one corresponding to a different class.

This method was applied successively for the criterion 'Gentleness' and for the criterion 'Frequency' of the interaction. For the criterion 'Gentleness', the LDA did not separate the two classes (gentle and strong). As for the criterion 'Frequency', attempts to separate contiguous classes two-by-two (i.e. firstly S_0 and S_1, secondly

[11] The mean μ of all the cases is defined by $\mu = \frac{1}{N} \cdot \sum_i x_i$, i.e $\mu = \frac{1}{N} \cdot \sum_c N_c * \mu_c$.

[12] For more details on how we obtain this equation the reader should refer to Haykin (1998).

[13] In an affine space of finite dimension n, a hyperplane is defined by its normal vector w and the bias b. A vector x from this space belong to this hyperplane if and only if it satisfies the following equation:

$$w^T \cdot x + b = 0 \tag{5.9}$$

[14] For an illustration of this method the reader can refer to Elizondo (2006).

S_1 and S_2 and finally S_2 and S_3) with the LDA did not work either. Therefore, the LDA was not further pursued.

5.2.7.2 Clustering by compression (Cilibrasi and Vitanyi, 2005)

This method takes inspiration from the Kolmogorov complexity and its related Normalized Information Distance (NID) which is non-computable in the Turing sense (Li et al., 2004; Li and Vitanyi, 2002). It provides a measure of similarity which is computable, the Normalised Compression Distance (NCD). Ususally, for an approximation of the NCD, standard compressors are used, like gzip. In this case, the NCD can be seen as the result of the approximation of the NID by a real compressor (Cilibrasi and Vitanyi, 2005). Clustering according to the NCD relies on grouping sequences that are similar according to features, but the knowledge of those features is not explicit (Cilibrasi and Vitanyi, 2005).

This technique was applied for each criterion of interaction (firstly the gentleness and secondly the frequency of the interaction). The test compared different compression rates. We shall illustrate this technique with a simple example. In this example, we consider two different classes S_0 and S_1 and two reference strings s_0 and s_1 (in our case the reference strings were windows on sensor data, each window originated from a different class, e.g. one from 'gentle', the other from 'strong'). We consider now an unknown string (a window on data) s, which belongs either to S_0 or to S_1. We concatenate successively s_0 and s_1 with s, ie. $concat(s_0, s)$ and $concat(s_1, s)$ and take the ratio of compression R of each concatenated string with a real compressor (in our case gzip), namely $R(concat(s_0, s))$ and $R(concat(s_1, s))$. If the classes are well separated by the clustering by compression, then the ratios $R(concat(s_0, s))$ and $R(concat(s_1, s))$ should be very different: if s originates from the source S_0, then $R(concat(s_0, s))$ should be much smaller than $R(concat(s_1, s))$; on the contrary, if s originates from the source S_1, then $R(concat(s_1, s))$ should be much smaller than $R(concat(s_0, s))$. It turned out that, in our application, this method was not really reliable for separating the classes (neither for the criterion gentle nor for the frequency of the interaction). Therefore, this method was not further pursued.

5.2.8 Summary

The purpose of this section was to present a proof-of-concept of the first model I used for the online recognition and adaptation to human-robot interaction styles. This method relied on Self-Organizing Maps, applied to the modulus of the Fast Fourier

Transform on windows on the input data. Results from a first level of testing have shown that the algorithm was able to classify pretty well the interaction styles for the criterion gentle/strong, but with a pretty high delay. It was possible to reduce this delay but it required important hand-tuning and made the solution very specific to a particular setting. Further to this I started investigating other techniques, firstly the Linear Discriminant Analysis which showed that the classes were not linearly separable and, secondly, Clustering by compression, which did not properly separate the classes. The next step of my research focused on the use of the Information Bottleneck Method, upon which I built a novel method for real-time classification of the interaction styles, which I present in the next section.

5.3 The Cascaded Information Bottleneck Method

5.3.1 Introduction

This section presents a novel method for time series analysis, the Cascaded Information Bottleneck Method, which I applied to the real-time recognition of human-robot interaction styles. This method, which enables time-filtering, is based on the concept of Information as introduced by Shannon (1949) and builds upon from the Information Bottleneck Method developed by Tishby et al. (1999).

Importantly, this work goes beyond prior work (see related work in Section 5.1.3) that either classified and characterized interactions off-line, i.e. after the interactions had taken place, or relied on explicit criteria tuned by hand (vs. automated training phase of the recognition algorithm). It also goes beyond the work I presented in Section 5.2 which enabled real-time recognition of interaction styles with respect to one criterion, the gentleness, using a different method, based on self-organizing maps (François et al., 2007). The Cascaded Information Bottleneck Method is entirely generic for applications with socially interactive robots.

5.3.2 Background: Information theory

This subsection summarizes basic notions of Information Theory[15] (Shannon, 1949).

5.3.2.1 Entropy

The entropy is a measure of the uncertainty of a random variable. It is a way to measure the amount of information required on the average to describe a random variable.

[15] For more details, the reader can refer to Cover and Thomas (1991); Crutchfield (1990)

Definition: Let X be a discrete random variable with alphabet $\mathcal{X}$ and probability mass function $p(x)$, $x \in \mathcal{X}$. The entropy H(X) of the random variable X is defined by

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x) \tag{5.11}$$

where the function log stands for the logarithm in base 2.

5.3.2.2 Joint Entropy and Conditional Entropy

In the previous subsection we have defined the entropy of a single random variable. Here we extend the definition to a pair of random variables.

Definition: Let X and Y be discrete random variables with alphabet, respectively $\mathcal{X}$ and $\mathcal{Y}$, and a joint distribution $p(x, y)$, whatever $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. The joint entropy $H(X, Y)$ of the pair of discrete random variables (X, Y) is defined as:

$$H(X, Y) = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(x, y) \log p(x, y) \tag{5.12}$$

where the function log stands for the logarithm in base 2.

Definition: Let X and Y be discrete random variables with alphabet, respectively $\mathcal{X}$ and $\mathcal{Y}$ and a joint distribution $p(x, y)$, whatever $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. The conditional entropy $H(Y|X)$ of the pair of discrete random variables (X, Y) is defined as:

$$H(Y|X) = \sum_{x \in \mathcal{X}} p(x) H(Y|X = x) \tag{5.13}$$

5.3.2.3 Relative Entropy

The relative entropy is a measure of the distance between two distributions.

Definition: The relative entropy or Kullback Leibler distance between two probability mass functions $p(x)$ and $q(x)$ is defined as:

$$D_{KL}(p||q) = \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)} \tag{5.14}$$

with the following conventions: a) $0 \log 0/q = 0$; b) $p \log p/0 = \infty$.

A very useful property of the relative entropy is that it is always non negative and is equal to zero if and only if $p = q$. However, the relative entropy is not a true distance between the distribution since it does not satisfy the property of symmetry.

5.3.2.4 Mutual Information

The mutual information is a measure of the amount of information that one random variable contains about another random variable. The mutual information between the random variable X and the random variable Y corresponds to the reduction in the uncertainty of one random variable (X or Y) due to the knowledge of the other one (respectively Y or X).

Definition: let X and Y be two random variables; let $p(x, y)$ be their joint probability mass function and $p(x)$ and $p(y)$ be respectively their marginal probability mass function. The mutual information $I(X; Y)$ is the relative entropy between the joint distribution and the product distribution $p(x)p(y)$:

$$I(X;Y) = \sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{Y}} p(x,y) \log \frac{p(x,y)}{p(x)p(y)} \tag{5.15}$$

The mutual information is symmetric and it is very easy to show that:

$$I(X;Y) = H(X) - H(X|Y) \tag{5.16}$$

Note that $I(X; X) = H(X)$.

5.3.3 Background: The Information Bottleneck Method

The Information Bottleneck Method (Tishby et al., 1999) is a clustering method based on an information theoretic approach (Shannon, 1949) whose purpose is to extract the relevant information[16] in a signal $x \in \mathcal{X}$ that is, extract features of a random variable (r.v.) X that are relevant to the prediction of Y. This problem is modeled by the following Bayesian network with Markov condition: $\tilde{X} \longleftarrow X \longleftarrow Y$ where $\tilde{X}$ is the variable that extracts information about Y through X.

This method provides an alternative to 'rate distortion theory' techniques which constitute a standard approach to lossy source compression. In the Information Bottleneck method, the relevance is not addressed through distortion but directly

[16] In this context, the relevant information is defined as the information that the signal $x \in \mathcal{X}$ provides about another signal $y \in \mathcal{Y}$.

through a new variational principle. The rationale is that the best trade-off between the compression of the signal and the preservation of the relevant information is the one that keeps a fixed amount of relevant information about the relevant signal Y while minimizing the number of bits from the original signal X, i.e. maximizing the compression. The optimal assignment can be found by minimizing the functional

$$\mathcal{L}[p(\tilde{x}|x)] = I(\tilde{X}; X) - \beta I(\tilde{X}; Y) \tag{5.17}$$

$I(X; Y)$ stands for the mutual information between X and Y. For β and the cardinal of $\tilde{X}$ fixed, an expression can be given which specifies implicitly the solution and leads to a fixed point iteration. β can be considered as the inverse of the temperature. This method uses a stochastic clustering top-down approach. The notion of stochastic refers here to the fact that the clustering is soft and that the input data are mapped to the different elements of $\tilde{\mathcal{X}}$ with a particular probability. For that information bottleneck setting, the Kullback-Leibler divergence $D_{KL}(p(y|x)||p(y|\tilde{x}))$ replaces the distortion function.

The Agglomerative Information Bottleneck algorithm (Slonim and Tishby, 1999) makes the assumption that β tends to ∞ in the Lagrangian equation (Eq. 5.17). When β goes to ∞, the first priority is to look for a solution that keeps all information about Y that X contains. The second priority is to remove all unnecessary information from X. In terms of mutual information, the mutual information between $\tilde{X}$ and Y is maximized and a hard partition of the data into subsets is induced, each subset corresponding to a bottleneck state $\tilde{x}$: for a fixed cardinal of $\tilde{X}$ (i.e. a fixed number of subsets - also called states - in the bottleneck), each member of the input signal $x \in \mathcal{X}$ belongs to one and only one subset $\tilde{x} \in \tilde{\mathcal{X}}$ and $\tilde{x}$ is the subset (the state) for which $p(y|\tilde{x})$ has the smallest $D_{KL}(p(y|x)||p(y|\tilde{x}))$. The hard partition can be soften afterwards, with reverse annealing. The pseudo-code of the algorithm can be found in Slonim and Tishby (1999).

5.3.4 The Cascaded Information Bottleneck Method

5.3.4.1 The principle

Based on the Information Bottleneck Method, I have developed a novel time-filtering method particularly adapted for pattern recognition in time series. Let x $\in X$ be the time series input signal of length l, $x = [x_0, ..., x_{l-1}]$. We take k and $S \in \mathbb{N}$, with $l = k * S$, such that x can be divided into S disjoined parts of cardinality k, X_s, $s = 0, ..., (S-1)$ in the following way:

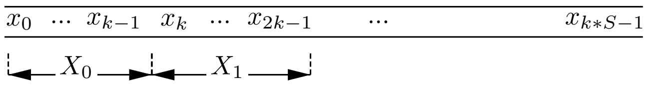

The Cascaded Information Bottleneck method relies on the principle that the relevant information can be progressively extracted from the time series with a cascade of successive bottlenecks sharing the same cardinality of bottleneck states but trained successively. The agglomerative information bottleneck algorithm is applied to each bottleneck successively, the first one being trained in the standard way while the next ones depend on the previous bottleneck states, as the following graph shows:

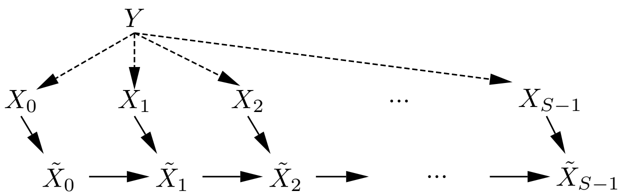

5.3.4.2 Extrapolation

The Cascaded Information Bottleneck method progressively extracts the relevant information from an input sample $X = [X_0, ..., X_{S-1}]$ by a recall on the successive components (X_0 for the first step of the cascade, $(\tilde{X}_{s-1}, X_s)$ for the other steps s). Each bottleneck is characterized by a hard mapping between: i) X_0 and $\tilde{X}_0$ for the first step, and ii) $(\tilde{X}_{s-1}, X_s)$ and $\tilde{X}_s$ for the other steps of the cascade. At each step s of the cascade, the algorithm looks for the equivalent $\tilde{x}_s$ given the input $(\tilde{x}_{s-1}, x_s)$ according to the hard mapping at step s: the equivalent $\tilde{x}_s$ satisfies the equation $p(\tilde{x}_s|(\tilde{x}_{s-1}, x_s)) = 1$. This means that if a pair has been observed during the training phase of the cascade, then there is only one outcome for the equivalent bottleneck state $\tilde{X}$. Note that the input $(\tilde{x}_{s-1}, x_s)$ corresponds to the input x in the original information bottleneck method.

It can happen that at a specific step s of the cascade, the pair $(\tilde{x}_{s-1}, x_s)$ for which we need to find the equivalent $\tilde{X}_s$ has never been encountered during the training process of this bottleneck. This pair is called an unseen pair. In the case of an unseen pair $(\tilde{x}_{s-1}, x_s)$ at step s, the cascade can a priori make no inference on $\tilde{X}_s$ because there is no preexisting default continuation of the cascade, due to the fact that the bottlenecks have been trained successively. In other words, for each pair $(\tilde{x}_{s-1}, x_s)$ which was not part of the training set data, $p(\tilde{x}_s|(\tilde{x}_{s-1}, x_s))$ is a priori undefined, whatever $\tilde{x}_s$ we take. For such cases, it is necessary to introduce a 'default' way leading from $\tilde{X}_{s-1}$ to $\tilde{X}_s$, i.e. we have to introduce an artificial identification of successive bottleneck states which consists in matching two bottleneck states (one at step $s-1$ and one at step s). Therefore I apply a reorganisation of the bottleneck states at each possible step s (i.e. a one-to-one mapping of the bottleneck states at step $s-1$ and the ones at step s which we call a permutation). For this purpose, I introduce the following measure $d_{(s-1,s)}$ allowing to directly compare the reorganised bottleneck states from step s with those from step $s-1$. Let $\tilde{\mathcal{X}}_{s-1}$ (respectively $\tilde{\mathcal{X}}_s$) be the set of bottleneck states $\tilde{x}_{s-1}$ (respectively $\tilde{x}_s$) and $p(\tilde{x}_{s-1})$ (respectively $p(\tilde{x}_s)$) the empirical probability; for each permutation r of the bottleneck states $\tilde{\mathcal{X}}_s$:

$$d_{(s-1,s)}(r) = -\sum_{\tilde{x}_{s-1}\in\tilde{\mathcal{X}}_{s-1}} p(\tilde{x}_{s-1}) \log \tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1}) \tag{5.18}$$

Note that if $\tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1})|\tilde{X}_{s-1} = \tilde{x}_{s-1}) = 0$ then, by convention, $d_{(s-1,s)}(r)$ is ∞. The logarithm measures the unpredictability of the next case (i.e. the unpredictability of $\tilde{X}_s$ given $\tilde{x}_{s-1}$). We want to choose r to minimize that unpredictability and weight for the probability that the state $\tilde{x}_{s-1}$ actually happens (because there is no sense in penalizing a deviation if the state does not happen.). We call this permutation $R(s-1,s)$.

The permutation of the bottleneck states that extracts the most similarity between bottleneck states at step $s-1$ and those at step s, $R(s-1,s)$, is given by:

$$R(s-1,s) = \arg\min{}_r d_{(s-1,s)}(r) \tag{5.19}$$

We consider $R(s-1,s)$ as the 'default' path between $\tilde{X}_{s-1}$ and $\tilde{X}_s$, i.e. as the criteria for extrapolating an unseen event at step s.

5.3.5 Implementation

In the following subsections we present an application of the Cascaded Information Bottleneck Method with real data which addresses the issue of interest in the present thesis: the automatic recognition of tactile interaction styles in the context of human-robot interaction. I conducted two series of trials, the first one under laboratory conditions and the second one in school where several children could interact (one child at a time) freely with the robot. In all experiments the robot used was the Sony Aibo ERS-7 and I focused on characterizing the tactile interactions according to two criteria, namely the *gentleness* and the *frequency of the interaction*. As defined in Section 5.1.2:

- An interaction is classified as 'gentle' (respectively 'strong') if the participant strokes the robot gently, without signs of force (respectively with signs of force).

- The frequency of interaction is categorized into four classes S_i, $i = 0...3$, defined by their typical periodicity of interaction T (in seconds): i) S_0: 'very low' ($T > 15$), ii) S_1: 'middle inferior' ($5 < T \leq 15$), iii) S_2: 'middle superior' ($1 < T \leq 5$), and iv) S_3: 'very high' ($T \leq 1$).

5.3.5.1 Implementation

Two cascades of bottlenecks were generated, one for the criterion *gentleness* and one for the criterion *frequency of the interaction*. The lengths of the cascades differed in these two different cases in order to meet the specificity of each criterion: the gentleness is a short-term time scale event while the frequency of the interaction is a mid-term time scale event, thus the length of the cascade for the criterion frequency was bigger than the one for the criterion gentleness. The whole list of parameters for each cascade of bottlenecks is provided in Fig. 5.8. The samples for the training of the cascade were generated during interactions with the Aibo ERS-7 in laboratory conditions within different runs. Each run contained one class exclusively, i.e. for the criterion gentleness, the samples generated within a same run contained only gentle or only strong styles of interaction (i.e. only gentle or only strong strokes were generated during a same run), and for the criterion frequency of the interaction, the samples generated within a same run contained only one type of frequency (i.e. S_0, S_1, S_2 or S_3 exclusively).

Criteria	**Classes**	**Length of the input vector (window size), l**	**Length of the individual subsequences, k**	**Length of the cascade, S**	**Number of bottleneck states, m**
Gentleness	2 classes: gentle/strong	50 (equivalent to 1.6 seconds)	2	25	4
Frequency	4 classes: S_0, S_1, S_2, S_3	472 (equivalent to 15.1 seconds)	2	236	6

Figure 5.8: Parameters for each cascade of bottlenecks.

Preprocessing Each criterion (*gentleness* and *frequency of the interaction*) is studied independently. In each case, the time series studied is the quantitatively binned sum of the normalized sensors values[17] involved in the type of interaction: the preprocessing normalizes each sensor data, sums these normalised values originated at

[17]The robot's sensor data are updated every 32ms.

the same time step, and bins this sum. Note that FFT is not applied contrary to the previous technique with SOMs presented in Section 5.2. For the criterion gentle/strong, the sensors involved are the four continuous external sensors, while the criterion frequency includes all five external sensors, i.e the four continuous sensors plus the boolean one[18].

Extra-conditions for the training

- for the criterion 'gentleness', the algorithm does not learn null samples (i.e. samples made of null events only),
- for the frequency of interaction, the system deals only with samples whose first component is not null.

In both cases, a sliding window proceeds on the sensor data time series. For the criterion gentleness, the window size is 50 while for the criterion frequency of the interaction, it is 472.

Postprocessing The postprocessing relies on a 'winner takes all' principle: The selected (winner state) is defined by $\arg\max_{y \in Y} p(y|\tilde{x}_{S-1})$.

5.3.5.2 Features of the trained cascade

The mutual information for the training set between the last bottleneck variable $\tilde{X}_{S-1}$ and Y is 0.8 bit for the criteria gentle/strong and 1.9 bits for the frequency of the interaction. The conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$ (Fig. 5.9) is globally decreasing over the cascade, pretty quickly, which suggests that a structure is progressively and rapidly emerging over the cascade: at the beginning of the cascade, a lot of new information is needed to deduce the next bottleneck state and then, when progressing in the cascade, less and less new information is needed. However, for the frequency of interaction, $H(\tilde{X}_{s+1}|\tilde{X}_s)$ has some small local peaks, both at the very beginning of the cascade and at the very end[19], which suggest that at these steps s, the input data X_s

[18] In this application, X is a window on the quantitatively binned sum of the normalised sensors data and Y is the class, i.e. the style of interaction, e.g. gentle or strong for the criterion gentleness and S_0, S_1, S_2 or S_3 for the criterion frequency of the interaction.

[19] Note that the small local peaks at the end of the cascade may reflect the importance of the last steps for distinguishing the classes S_0 and S_1. S_0 is defined by a periodicity greater to $15s$ while S_1 is defined by a periodicity greater to $10s$ and inferior or equal to $15s$. These small peaks in the end of the cascade appear at approximately $14s$ (the peaks appear just after $s = 220$, for $k = 2$ and the systems send updates of the sensor data every $32ms$, thus the duration is : $220 \times 2 \times 32ms = 14.08s$). I hypothesize that they reflect the typical periodicity of events from S_1.

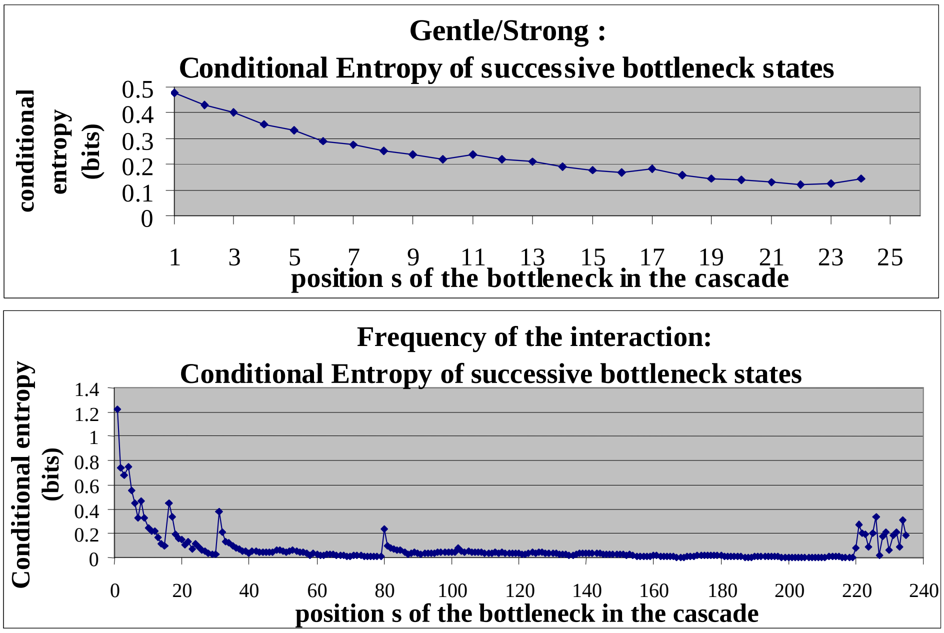

Figure 5.9: Conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$. There are four main parameters for the cascade: l (length of the input vector), k (length of the individual subsequences), S (length of the cascade), m (number of bottleneck states). For the frequency of interaction, $l = 472$ (equivalent to 15.1 seconds) , $k = 2$, $S = 236$, and $m = 6$. For the criterion gentle/strong, the corresponding parameters are: $l = 50$ (1.6 seconds), $k = 2$, $S = 25$ and $m = 4$.

may influence a bit more in the choice of next equivalent state $\tilde{X}_{s+1}$. This measure is correlated with the reorganisation measure for extrapolating $d_{s-1,s}(R(s-1,s))$ (equation (5.18) and equation (5.19)) which presents, respectively to each criterion of interaction, profiles similar to the conditional entropy with peaks positioned at the same place in the cascade (the mean of $d_{s-1,s}(R(s-1,s))$ is equal to, respectively, for Gentle/Strong, 0.037 bits, and, for the frequency of interaction 0.129 bits): when the distance becomes bigger, it means that there is less similarity between the successive bottleneck steps according to the measure $d_{s-1,s}(R(s-1,s))$. This is equivalent to say that more new information is needed to find the equivalent bottleneck state, which corresponds to a peak in Fig. 5.9. In the rest of the study, the algorithm will extrapolate between step 5 and 24 (respectively 5 and 216) of the cascade for the gentleness (respectively frequency of interaction).

5.3.6 Experiments

These experiments aim at assessing statistically:

- The soundness of the recognition of interaction styles by our algorithm, i.e., i) for the criterion 'gentleness', whether a behaviour that has been classified as gentle (respectively strong) by a human is indeed going to be classified as gentle (respectively strong) by our algorithm, or ii) for the frequency of interaction, whether a frequency of interaction that has been tagged by a human is indeed going to be correctly recognised by the algorithm.

- The delay for the recognition of local events.

Importantly, the criterion 'gentle/strong' characterizes local events, and the algorithm should be able to recognise each specific event 'gentle' or 'strong' within a short delay. In contrast, the criterion 'frequency of the interaction' requires the algorithm to classify mid-term time scale events. This study deliberately focuses on such different criteria in order to show the flexibility of the algorithm.

Experimental setup under laboratory conditions These trials are used as a first step in the statistical assessment of the soundness of the recognition of the interaction styles.
They involve one participant at a time who is asked to interact with the robot for a few minutes in a *predefined way* which is one of the following:

- for the 'Frequency of the interaction': only 'pure styles of interaction', i.e one class[20] exclusively.

- for the criterion 'Gentle/Strong': In a first step, it is pure styles exclusively[21]. In a second step, the participant is asked to alternate gentle and strong behaviour and, just before generating the first event of the new class, he/she must name the style (i.e. "gentle" or "strong"). All the sessions are video recorded and this tagging enables to determine very precisely the transitions for a further measure of the delay of the recognition process.

Experimental setup in school A further step in the validation of the algorithm is the testing with data obtained under natural situations of Human-Robot interaction. These experiments took place in a small classroom dedicated to the study, one child at a time being present in the room. Each child was invited to play freely for several minutes with the robot (the duration of play depended on the child's needs and abilities) in an unconstrained environment.

[20] very low (S_0), middle inferior (S_1), middle superior (S_2), or very high (S_3).
[21] gentle or strong only.

5.3.6.1 Measures

The experiments were all video-recorded and sensor data were stored. Note, the validation of the algorithm must be assessed offline but the recognition algorithm is designed to operate in real time. The evaluation follows two main successive steps: firstly the testing with the trained data and, secondly, cross-validation with, on the one hand, data generated under laboratory conditions and, on the other hand, data generated in school during interaction between children and the robot.

Samples excluding transitions from one class to another The profile of the classification made by the algorithm can be analysed with a confusion matrix which displays the probability distribution that events from class S_i are recognised by the algorithm as events of class S'_i ($i = 0$ or 1 for gentle/strong, $i = 0...3$ for the frequency of interaction).

Samples with transitions for the criterion gentle/strong These samples enable us to test the ability of the algorithm to recognise a transition and reach, after a short transition phase, a new equilibrium phase. One can model this process by a temporal curve that would indicate the state of the system for a transition happening at time t_0. Three typical domains can be identified: for $t < t_0$ the curve is constant, indicating a stable state; from $t = t_0$, the curve's value alternates to indicate an hesitation between the two possible states (thus identifying a change in the behaviour observed); from $t = t_0 + \tau$ the curve would keep the same value (the new state). Ideally, the second phase should be very short (i.e. τ is very small). I study three typical measures here: a) the number of transitions recognised by the algorithm; b) the time elapsed to reach the new equilibrium state, c) the ratio of errors made within this new equilibrium state. Note, a transition will be considered broadly as either a transition from a gentle (respectively strong) behaviour to a strong (respectively gentle) one, or from a state where no classification occurred (i.e. no interaction occurred during the past 1.6 seconds) to gentle or strong.

Samples with hybrid behaviours for the frequency of interaction Because this criterion is based on a mid-term time scale analysis, some samples generated in school can be hybrid, i.e contain a mix of features from different classes. In order to encapsulate hybrid behaviours, the human classifies the behaviours on a 'two choices' basis, i.e. he/she can select the two styles characterising the hybridity. In this case, the algorithm's classification is successful if it agrees with one of the two choices made by visual inspection.

Practically, the video and graphs of the temporal global variable are first manually tagged. In a second step, the classifications S_i resulting from the manual tagging are compared with the classifications S'_i made by the algorithm.

5.3.7 Results

In this subsection, we present the results for the two criteria of interaction successively, firstly the gentleness of the interaction and secondly the frequency of the interaction. Note that here we will refer to the samples of data that were classified without using the extrapolation[22], i.e. the samples that contained no unseen cases at any step of the cascade, as *samples classified without extrapolation*. In contrast, the samples of data that required an extrapolation at one or more steps of the cascade, i.e. the samples for which there were unseen cases to extrapolate [23] (i.e. cases that had not been encountered during the training phase of the algorithm), will be referred to as *samples classified with extrapolation*.

5.3.7.1 Results for the criterion 'Gentle/Strong'

In the four following paragraphs, we report on the results for the criterion 'gentleness of the interaction'. For this criterion, the algorithm was evaluated successively with a) the data from the training set, b) new samples of data generated under controlled conditions excluding transitions (cross-validation), c) new samples of data generated under controlled conditions including transitions (cross-validation) and d) samples of data generated in school by the children (cross-validation).

Training set of data: The 20,018 samples used for the training were classified by the algorithm with an overall success of 97.82% and, respectively, for gentle and strong, 96.83% and 98.81%.

Samples excluding transitions: They constitute 1 hour 2 minutes 49 seconds of interaction. 100,111 samples were classified with a ratio of success for correct classification of 0.948.

[22]Those samples were classified by recall according to the hard mapping defined during the training phase of the algorithm; since there were no unseen cases in these samples, those samples never used the extrapolation for their classification.

[23]Those samples were extrapolated according to the measure provided in Section 5.3.4.2

97.7% of samples were classified without extrapolation with 95.22% of success while the samples classified with extrapolation (3.3%) were well classified in 75.54% of cases which, considering that it results from an extrapolation, is quite a good result. Note that the parameters of the Cascaded Information Bottleneck Method were chosen in such a way to have a good balance between the extrapolation and the precision, which is reflected here in the low percentage of cases extrapolated.

Samples with transitions under laboratory conditions: The four runs constitute 19 minutes and 40 seconds of interaction to analyse. They contain 53,192 samples to classify and 0.01% of the samples were not classified because they could not be extrapolated by the algorithm[24]. 212 transitions were to be recognised, 99.1% of which were indeed well classified by the algorithm[25] with an average delay of 0.17 seconds. The cumulative probability distribution of the delay is displayed in Fig. 5.10. The curve grows very rapidly, thus showing that most of the delays are very small. Transitions recognised without any delay occur particularly in the case of a transition from no event to classify to any event to classify. The longest delay is 2.05 seconds, which I consider very acceptable for human-robot interaction kinesics. The average error ratio in the equilibrium phase is 0.02 and the cumulative probability distribution is displayed in Fig. 5.11. Here again, the curve grows rapidly and shows that the probability of the highest error ratio is very low and remains acceptable for real human-robot interaction.

Samples generated by the children in the school: Videos from five different children were analysed, which constitute 12 minutes and 52 seconds of interaction. These runs contain 6,660 samples to classify: 97.49% of these samples were classified by the algorithm. These samples contain 45 transitions. 91.1% of these transitions were indeed well classified by the algorithm within an average delay of 0.17 seconds. The cumulative probability distribution of the delay is represented in Fig. 5.10. The curve grows very rapidly, thus showing that most of the delays are very low. Transitions recognised without any delay occur, and, at the far end, the highest delay is 1.54 seconds, which is very acceptable for human-robot interaction kinesics. The mean

[24]These samples had to be extrapolated outside the range of steps considered for the extrapolation. The range of steps considered for the extrapolation is the range of steps where the algorithm is allowed to extrapolate unseen events. For the gentleness of the interaction, is it between step 5 and step 24; for the frequency of the interaction, the range of steps for extrapolation is between step 5 and step 216 of the cascade. If a sample had to be extrapolated outside this range, then it was not classified.

[25]A transition is considered as wrongly classified if the transition phase is very long compared to the new equilibrium phase.

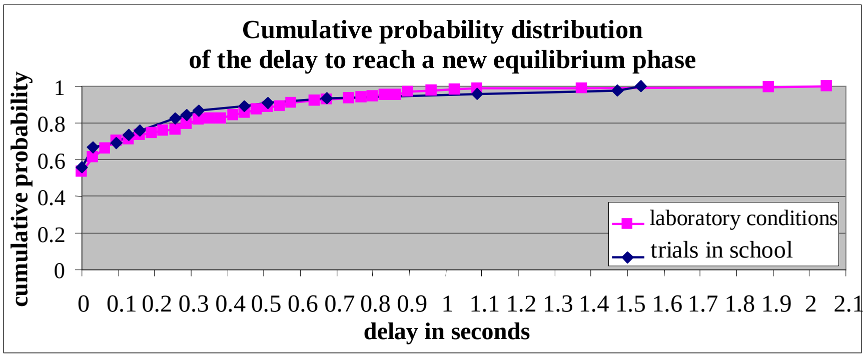

Figure 5.10: Cumulative probability distribution of the delay for recognising the transition. We display the cumulative probability, i.e. the probability that an event is recognised within (less or equal) n seconds for a given n. The delay corresponds to the length of the transition phase when a transition occurs.

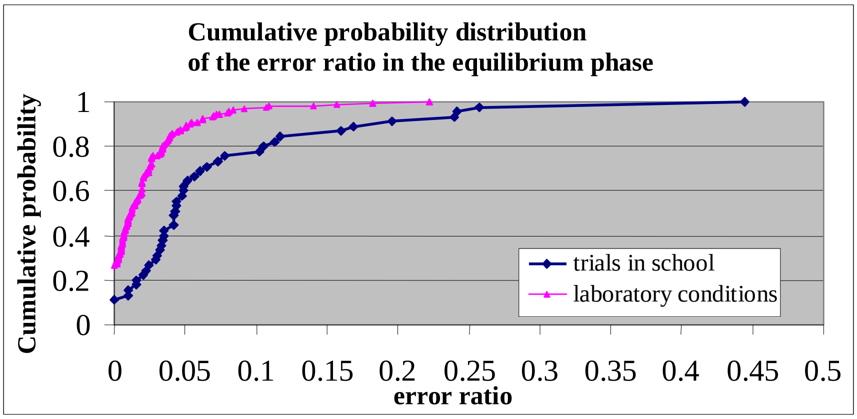

Figure 5.11: Cumulative probability distribution of the error ratio for the equilibrium phase. The ratio measures the number of errors of classification made during a phase of equilibrium divided by the number of samples to classify during this phase. The figure displayed gives, for a given r, the cumulative probability, i.e the probability that the error ratio is inferior or equal to r.

error ratio in the equilibrium phase is 0.1 and the cumulative probability distribution of this ratio is displayed in Fig. 5.11. Here again, the curve grows rapidly. It is worthy of note that the highest value obtained is 0.44 and the second one is much lower (0.26) which indicates that the first highest value can be seen as an extraordinary case. Looking at the sequential classification of the results, it appears that this highest error ratio was obtained while a child interacted in a very instable way that is, within 1.76 seconds three successive transitions were observed that are 1) no event to gentle (gentle phase lasted 1.37 seconds), 2) gentle to strong (the phase with strong style lasted only 0.26 seconds), 3) strong to gentle. It is the strong phase, after the transition from gentle to strong behaviour that was recognised with the highest error ratio (0.44), but it lasted for such a short time that it is not really a concern here (0.26 seconds is very low compared to the typical time for human-robot interaction which usually lasts a few seconds). Therefore, we can consider to omit this highest value 0.44 in the probability distribution and, looking at the resulting values, the results are good and comparable to the results obtained in the laboratory.

5.3.7.2 Results for the criterion 'Frequency of the interaction'

In the three following paragraphs, we report on the results for the criterion 'frequency of the interaction'. For this criterion, the algorithm was evaluated successively with a) the data from the training set, b) samples of data generated under controlled conditions excluding transition between classes (cross-validation), c) samples of data generated in school by the children (cross-validation).

	S'$_0$	S'$_1$	S'$_2$	S'$_3$
S$_0$	**1**	0	0	0
S$_1$	0.0008	**0.9992**	0	0
S$_2$	0	0	**1**	0
S$_3$	0	0	0	**1**

Figure 5.12: Confusion Matrix for the training set. The ratio is the one among events from type S_i. S_i represents the real class and S'_i the recognised class, $0 \leq i < 4$.

Training set of data: It constitutes 36 minutes 34 seconds of interaction and contains $4,865$ samples to classify (respectively, 450 for S_0, $1,208$ for S_1, $1,484$ for S_2 and $1,723$ for S_3). 99.98% of these samples were well classified; the ratio of success specific to each class is displayed in Fig. 5.12.

Samples generated under laboratory conditions: They constitute 51 minutes 44 seconds of interaction and contain $5,395$ samples to classify (respectively $1,017$ for S_0, 855 for S_1, $1,933$ for S_2 and $1,590$ for S_3) 91.16% of which were classified with an overall ratio of success of 0.922. 99.4% of the samples not extrapolated were well classified, and 76.41% of samples classified through extrapolation were well classified. Fig. 5.13 displays the confusion matrices.

Samples generated by the children in the school: Three runs of interaction were used for the validation of the frequency of interaction in a real situation, from three different children. They constitute 14 minutes 41 seconds of interaction and

No Extrapolation	S'_0	S'_1	S'_2	S'_3
S_0	**1**	0	0	0
S_1	0	**0.972**	0.028	0
S_2	0	0	**0.999**	0.001
S_3	0	0	0.006	**0.994**

Extrapolation	S'_0	S'_1	S'_2	S'_3
S_0	**1**	0	0	0
S_1	0.115	**0.864**	0.022	0
S_2	0.083	0.146	**0.768**	0.003
S_3	0	0	0.368	**0.632**

Figure 5.13: Confusion Matrices for pure sets of data (cross-validation) for, respectively, non extrapolated and extrapolated samples. Non extrapolated samples are samples which were classified without the need to use the extrapolation, because none of the cases were unseen cases (relatively to the training set samples). The results for those samples is provided in the table with mention 'No extrapolation'. On the contrary, extrapolated samples are samples that used the extrapolation at least once in the cascade (those samples contained at least one unseen case in the cascade, i.e. a case that had not been encountered during the training). The results for those samples are provided in the table with the mention 'Extrapolation'. See Fig. 5.12 for more details on the notion of confusion matrix.

contain $5,288$ samples to classify. 91% were classified (including 26.81% that had to be extrapolated) and 93% were classified correctly. Among samples classified with no extrapolation, the ratio of success for a sound classification was 0.96 while for samples classified with extrapolation, it was 0.84.

5.3.8 Discussion

The Cascaded Information Bottleneck Method has proven sound for the recognition of the two criteria of interaction. Concerning the criterion gentle/strong, results show that the two classes are well recognised and the delays very acceptable for human-robot interaction. The extrapolation works well, which shows the capability of the system to make a sound decision in case of unseen events. These results can be compared with the preliminary study presented in Section 5.2 where I used Self-Organizing Maps to classify this criterion of interaction (François et al., 2007), whereby the average delay to recognize transitions was much higher and the postprocessing required more effort.

Importantly, one might wish to define the styles slightly differently to the definition given here, such as, for instance, focusing on more details (in order to describe substyles for instance). This can be easily done by adjusting relevant parameters, mainly the number of bottleneck states, the binning and the training sets which condition the learning. One can also control how much information is being left by tuning the parameter β that is introduced in Eq. 5.17: in this thesis we used the limit 'β goes to ∞' but one could go back to a finite β which enables to control the quantity of information that X contains about Y that is being left.

The algorithm has also proved very capable of classifying real data over a mid-term time scale (cf. the criterion frequency of the interaction) which illustrates the ability of the method to make use of an existing temporal structure not only of short-term time scales but also mid-term ones. This ability is enabled by the use of different bottlenecks (thus different mappings) over the cascade. Besides, the process for extracting the information is transparent: we can say how much and which information is extracted at which step of the cascade, which gives a fine-grained control over what information is taken from the input data and where that information is taken in the cascade (at which step of the cascade). In contrast, with homogeneous HMMs, as used for gesture recognition in e.g. Lee and Xu (1996) and Calinon and Billard (2004), the mapping would be the same all over the time series[26], and, by trying to squeeze all temporal information into one flat transition structure, it might actually prevent homogeneous HMMs from a powerful exploitation of an existing temporal structure of the data. This hypothesis about homogeneous HMMs needs to be explored in future work which should include a comparison of the Cascaded Information Bottleneck Method with HMMs in these scenarios.

This method is designed for real-time use during natural human-robot interaction and little research had been done so far on real-time recognition of tactile interaction styles. Salter et al. (2007)'s adaptation algorithm was a first important step towards real adaptation. Yet, that system did not learn its own categorisation, which was completely described by a hand-tuned decision tree. In the present study, the recognition and the decision are made algorithmically, after a real learning phase and with a capacity to extrapolate unseen events, with very small delays. Furthermore, the method is very easy of use and can be tuned easily to adapt to other criteria of interaction. This method is entirely generic for different applications involving socially interactive (humanoid and non-humanoid) robots.

5.3.9 Summary

In this section, I have presented a novel method for time series analysis for detecting interaction styles in the context of Human-Robot Interaction. This method, namely the Cascaded Information Bottleneck Method, has its roots in the Information Bottleneck Method (Tishby et al., 1999) and relies on the principle that the relevant information can be progressively extracted from the time series with a cascade of successive bottlenecks sharing the same cardinality of bottleneck states but trained

[26]For more detail on homogeneous HMMs, the reader can refer to Section 5.1.3

successively. The first bottleneck is trained in the standard way while the next ones depend on the previous bottleneck states. This importantly contributes to enable the method to make a powerful exploitation of an existing temporal structure of the time series. Moreover, a structure progressively emerges through this cascade of bottlenecks, and I introduced a measure for extrapolating unseen cases, which are cases that have not been seen during the training phase of the algorithm. The Cascaded Information Bottleneck method is thus transparent and provides a fine-grained control over how much and what information is taken from the input data and where, in the cascade, this information is extracted.

I have applied this novel computational method to real-time recognition of human-robot interaction styles, in a detailed study, by implementing the algorithm for interactions with a real robot. The testing of the method had to be done offline, i.e. after the interactions had taken place, but the algorithm is designed to operate in real time in order to enable real-time adaptation of robots to the interaction styles.

I have shown the soundness of the method through extensive experiments, using successively samples of data generated under laboratory conditions and samples from natural situations of child-robot interaction in a school for children with autism. The algorithm was able to recognize short term events very well within an average delay of 0.17 seconds (the highest delay being 2.07 seconds). It was also able to recognise mid-term time scale events very well (during cross-validation, the percentage of events correctly classified was 92.2% under laboratory conditions and 93% with data from the child-robot interactions). This study has shown the soundness of the method for pattern recognition and illustrated its capability of time-filtering on real data. The method is entirely generic for applications with socially interactive robots.

The next chapter will focus on the application of the method in autism therapy where we find a strong need for socially adaptive robots. The ability of a robot to classify in real time human-robot interaction styles is a first step towards the challenging goal of enabling an autonomous robot to influence positively children's interaction styles to guide him/her progressively towards different therapeutically relevant levels of interaction.

Chapter 6

The Adaptive Robot in Robot-Assisted Play

This chapter addresses the role of the adaptive robot in robot-assisted play. A robot that is 'adaptive' can recognize interaction styles in real time and adapt to them appropriately. In other words, an adaptive robot reacts differently depending both on the origin of the stimulation (i.e. which sensor(s) is (are) activated) and on the styles of interaction recognised. In contrast, by 'reactive' robot, we mean here a robot that can only react differently depending on the origin of the stimulation, and which will not change its behaviour according to the interaction styles.

This chapter presents a proof-of-concept system of an adaptive robot responsive to different styles of interaction in human-robot interaction. The adaptive robot uses the Cascaded Information Bottleneck Method to recognize in real time the interaction styles and adapt its behaviour accordingly. The potential of this adaptive robot to influence the children's play styles is investigated experimentally through a short-term study with seven children with autism. The long-term goal of this study is to investigate whether an adaptive robot might help children with autism reach therapeutically relevant levels of interaction.

6.1 Schema of adaptation: the reward basis

The adaptive mode relies on a reward basis for well-balanced interaction styles: the child should get a positive feedback (also called reward) from the robot when he/she plays in an appropriate style of interaction[1]. The idea behind is the same all along this thesis: the child should always be encouraged and rewarded for every progress he/she made. With this approach, we hope to comfort the child in gaining self-confidence, enjoying himself/herself, and progressively acquiring a better understanding of the interactions he/she is involved in. It is hoped that the rewarding process can indirectly play the role of a trigger: the child wants to get the reward and therefore changes his/her behaviour until he/she actually gets it. Concretely, the robot should help regulate the interaction: if the child plays in a well-balanced interaction style, the robot reacts appropriately to the stimulation; on the contrary, if the interaction is, for instance, too strong, the robot does not show any reaction.

Moreover, the child should be encouraged engaging in the interaction if he/she does not. Therefore, the robot should be both rewarding and engaging[2].

6.1.1 Schema of adaptation

The reward is a physical reaction of the robot, which can be a gesture, a movement, a light or a sound. The concrete instantiation of these behaviours has been designed by immersion for each child beforehand (Appendix C), during long-term studies with each child, whereby the experimenter tested different robot's behaviours with each child in order to evaluate 1) whether the specific child liked it or not, 2) whether he/she conferred a specific meaning to the reaction and, particularly, whether the reaction had, in his/her view, a connotation of the robot being happy or sad. In this way, it appeared that the robot's barking was mostly interpreted by the children as the robot being very happy, which is contradictory to our *a priori* hypothesis that the robot's barking would induce a back off, thus calming the interaction (cf. Chapter 5.2).

We shall now detail the notion of reward: each time the child activates a sensor, the robot evaluates the interaction style in terms of gentleness and in terms of frequency and gives a reward, separately according to each criterion. If the interaction is gentle, then the robot shows a reaction to the child. The reaction depends on the sensor activated (there is a deterministic mapping between the sensors and the

[1] The feedback is designed specifically for each child and results in robot's specific behaviours. The robot's behaviours are chosen specifically for each child to ensure that each child reacts positively to them and thus, receives a positive feedback when he/she plays on an appropriate style of interaction; this positive feedback is viewed as a reward. Note that those behaviours have been designed during a long-term study described in Appendix C.

[2] The robot could try to 'trigger' or bootstrap the interaction if the child is not engaged in the interaction.

reactions of the robot for each child). If the stimulation takes place in a good overall frequency of interaction , i.e. a well-balanced frequency of interaction, two LEDs turn on on the robot's face (which is sometimes interpreted by the children as the robot's eyes). Note that a well-balanced frequency of interaction is a frequency not too low and not too high, represented in this thesis by the classes S_1 and S_2 as defined in Section 5.1.2. This model is totally generic and can be applied with different criteria of interactions.

Besides, future work could expand this model by focusing on a larger grid of criteria for the interaction styles: while the child progresses, the robot could increase the range of criteria it considers for characterizing the interaction and thus on which the decision is based for the reward for the child. In contrast, when the child encounters some difficulties, then the robot could simplify the range of criteria on which the reward for the child is based, so that the child can get a better understanding of the interactions happening. This progressive refinement in the adaptation process of the robot to the child's play styles could be linked, in some sense, to the notions of 'discrete development' and '(Alternate) Freezing and Freeing of Degrees of Freedom'[3] which has been widely used in developmental robotics (Berthouze and Lungarella, 2004; Lungarella and Berthouze, 2002; Gómez et al., 2004). This technique, typically applied for a system learning motor skills, can actually be transposed to a social system, constituted here by the child and the robot (see Fig. 6.1): this social system is freezing some complexity in the interaction to learn more efficiently how to deal with interaction in general. The adaptation in real time from the robot to the interaction styles is directly linked to the social potential of the robot. If the system has done enough progress with respect to the interaction styles then the system can release some degrees of freedom in the interaction i.e reach a more complex level of interaction. The process of reaching a more complex level of interaction could be enabled by the robot triggering more complex criteria of interaction and maybe also by the robot adopting more complex behaviours in response to the child's stimulations. A simple example of this discrete social development could be modeled by three levels (in terms of complexity, *level 0* $<$ *level 1* $<$ *level 2*):

- *level 0*: the robot only adapts to 'no interaction' by trying to engage the child in the interaction; whatever the stimulation from the child, it then gives feedback
- *level 1*: the robot adapts to the criterion gentle/strong and additionally rewards

[3] The notion of degrees of freedom can be used in different areas. In mechanics, it is defined as the set of independent displacements and rotations that describes completely the displaced or deformed position and orientation of a system.

for a good frequency of interaction

- *level 2*: the robot adapts to the criterion gentle/strong and to the frequency of the interaction (i.e both criteria have to be satisfied in order for the child to get a feedback from the robot).

It is beyond the scope of this thesis to test this model. We first need to show the potential role of an adaptive robot in robot-assisted play before starting longer and more complex trials that would necessitate the investigation of the potential of this model. But this should be done in future work, because this model seems *a priori* really valuable; in particular, it favors the flexibility potential of an adaptive robot (in comparison with a reactive robot), its capacity to embed a model of gradual (discrete) social development, and therefore, we hypothesize that the robot's behaviours determined by this model might be a relevant playmate to the children, which would be both entertaining and educating, by adapting the complexity of the interaction triggered to the specific needs and abilities of any child at a specific time.

System	Target	Internal Evaluation	Source of the increase in complexity
{Robot}	Motor skills	Locomotion	- Mechanical degrees of freedom (Lungarella & Berthouze, 2002) or - Concurrently (Gòmez et al., 2004): • mechanical degrees of freedom, • sensor resolution • neural capabilities
{Robot, Child}	Interaction level	Interaction	- Range of the interaction styles to which the robot adapts - Range of the robot's behaviours

Figure 6.1: Discrete Development: comparison of two systems. The first system is a robot learning motor skills; two examples are given. In the first one, the robot alternatively freezes and frees mechanical degrees of freedom (Lungarella and Berthouze, 2002). The second example focuses on the concurrent increase in complexity in the mechanical degrees of freedom, the sensor resolution and the neural structure (Gómez et al., 2004). The second system consists of the child and the robot; this systems 'learns' social interaction.

6.1.2 Adaptation according to two criteria

We shall now go back to our principal focus here, which is the adaptation of the robot according to two criteria, the gentleness and the frequency, where the child gets the feedback from the robot if he/she stroke the robot gently with an additional reward

(LEDs turning on and off) for a good frequency (thus corresponding to level 1 in the preceding model of discrete development).

6.1.2.1 Schema of adaptation for the two criteria

Fig. 6.2 presents the general Reward's schema for the two criteria.

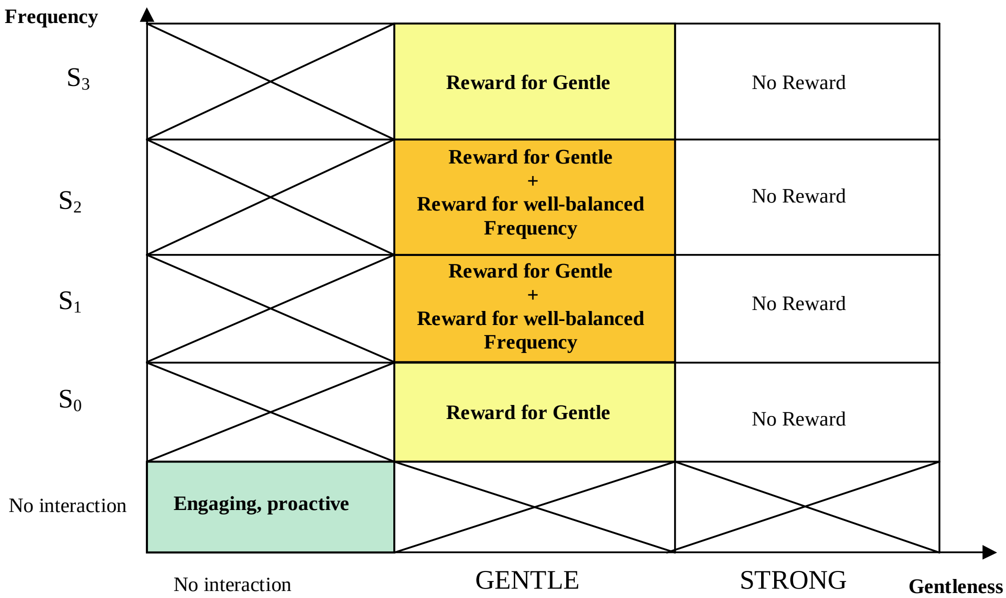

Figure 6.2: Reward Schema for the two criteria of interaction.

6.1.2.2 Architecture for Decision-Making on the interaction styles

The purpose here is to enable the robot:

- To give an appropriate feedback to the child after the child has activated a sensor of the robot, within a very short delay. In order to give the appropriate feedback, the robot must first detect the play style of the child, that is, i) whether the child's stroke is gentle or strong and ii) if the child's stroke took place in a context of a good frequency of interaction (i.e. the frequency of interaction is not too low and not too high);
- To detect whether the child is not interacting with the robot, i.e. not engaging in the interaction in order to try to encourage the child to engage in the interaction.

The first point is addressed as follows. The real-time recognition of the interaction styles uses the Cascaded Information Bottleneck Method. As we have shown in Chapter 5.3, the recognition process of the criterion Gentle/Strong is done with a possible small delay. We have modeled this delay by a curve with a transition phase, followed by an equilibrium phase. In order to completely automate the decision process, this transition phase must be taken into account and the decision-making process should be given enough time in order not to be altered by the transition phase. This wait is modeled by a pause in the decision-making process, that is a small latency (600ms) during which the algorithm ignores the current interaction style. After the sleep, the decision-making process looks at the successive classifications that are made by the Cascaded Information Bottleneck algorithm during a fixed short amount of time and counts the occurrences of Strong behaviours recognised. If it is above a fixed predefined threshold then the final choice (i.e. the decision) is that the behaviour is 'strong interaction style' and the child will not get the reaction from the robot to his/her stimulation. If it is under this threshold, then the decision is 'gentle interaction style' and the child gets the reaction from the robot corresponding to the sensor activated. Besides, the robot updates the criterion frequency of interaction with a 1 second periodicity according to the Cascaded Information Bottleneck method (different threads for the gentleness and the frequency of interaction running in parallel). If, when the child strokes the robot gently, the current frequency of interaction is S_1 or S_2, then the child will get the additional reward of the two lights illuminating on the robot's face, while the robot also shows the specific reaction correlated to the gentle stimulation.

Note that this decision-making process really reflects the variety of the interaction styles considered here, the criterion 'gentle/strong' corresponding to a short-term time scale event and the criterion 'frequency of the interaction' corresponding to a mid-term time scale event.

The second point (detecting an absence of engagement in tactile play from the child) is addressed as follows: we consider that the child should be encouraged to play with the robot if he/she has not stroked the robot for a specific time that we define here as just above 15 seconds (more exactly, the length of the window size for classifying the frequency of the interaction which is 472×32 ms). Thus the decision-making process here is straightforward: at each update of the frequency of the interaction, the algorithm checks if the input vector is null or not. If it is null, it means that the child has not engaged in tactile interaction with the robot for, at least, 15 seconds and the robot starts its engaging behaviour.

6.2 Can an adaptive robot influence children's play styles? A short-term study

6.2.1 Motivation

This study investigated whether an adaptive robot might impact the child's play styles in comparison with a reactive robot. Different research questions were addressed here:

- Does the adaptive robot encourage the children or, on the contrary, discourage the children from engaging in the interaction? Is their engagement in play similar to the one with a reactive robot?
- Does the child's play patterns differ when the robot is adaptive from when the robot is reactive? This question contains two subquestions as follows:
  i) Are the strokes qualitatively different (ideally more gentle) when the child plays with an adaptive robot?
  ii) Is the frequency of the interaction differently (ideally better) balanced when the child plays with the adaptive robot?

6.2.2 Method

6.2.2.1 Participants

Seven children with autism participated in the experiments which took place in the same school as the one described in Chapter 4. All these children had had the chance to play with the robot during several months beforehand: therefore, they were familiar with both the robot and the experimenter and had already experienced various situations of play with the robot and, possibly, the experimenter. In particular, they had already experimented with play sessions with the approach inspired by non-directive play therapy. Note that among the nine children involved in the play sessions on a regular basis in the school, only seven children could participate in the trials described in this chapter. This was due to the fact that it would have been too hard for those two remaining children to cope with the experimental setup: for them, the sessions would have been too long and too complicated.

6.2.2.2 Artifact

The robot was the Aibo ERS7. It behaved autonomously; the range of its behaviours depended on the specific child. Those behaviours have been designed by immersion for the children through a long-term study (Appendix C). Note that 'design by immersion' means progressively developing the robot's behaviour as the children play, so that the behaviours can correspond to the children's needs, abilities and preferences. This means observing the children's reaction when a precise robot's behaviour occurs, and identifying whether the child shows any hesitation towards it, or, on the contrary, whether he/she feels really at ease with that behaviour; note that, in particular, some children expressed verbally their reactions towards robot's behaviours, and other started smiling and laughing when specific robot's behaviours started.

The robot was either adaptive or reactive. In both cases the mapping between the sensors and the robot's reactions was the same except from the LEDs flashing for a good frequency of interaction, which was an additional feature for the adaptive robot, as well as wagging the tail when no interaction was detected. The behaviour mapping used for this specific study is detailed in Fig. 6.3.

Sensor	Corresponding behaviour
Chin sensor	Emit "bark" sound while opening-closing the mouth
Head sensor	Turn head (Head tilt)
Back front sensor	- Wag the tail (used for Child E) - Walk forward, turn right, stand, turn left, walk backwards (used for the other children)
Back middle sensor	Turn head (Head pan)
Back rear sensor	Emit "drum" sound while wagging the tail

Figure 6.3: Mapping between the external tactile sensors of the robot and its behaviours. In both modes (reactive and adaptive) the mapping between the sensors and robot's reactions is the same for a specific child. For child E, the walking has been removed and replaced by the robot's wagging the tail. The difference between the two modes is that in the adaptive mode, the robot's reaction happens only if the interaction style is gentle, and an additional reward is provided (flashing LEDs) if the stroke happens in a context of well-balanced frequency of interaction; plus, in the adaptive mode, the robot has an engaging behaviour, wagging the tail, when the child has not stroked the robot for more than 15 seconds.

6.2.2.3 Procedures and Measures

Procedures (experimental approach, methodology): Each child participated in two sessions and the experiments involved one child at a time. Each session consisted in three successive steps[4] (also called games or runs), each step being defined by the mode of the robot– reactive (R) or adaptive (A)– which alternated between two successive steps.

As a result, a session was defined by its setting which was either A-R-A or R-A-R. Each child experimented with both settings (each during a different session). Three children started with the setting A-R-A. The four remaining started with R-A-R (Fig. 6.4).

Child	**Setting 1**	**Setting 2**
Child A	A-R-A	R-A-R
Child G	R-A-R	A-R-A
Child H	A-R-A	R-A-R
Child C	R-A-R	A-R-A
Child E	R-A-R	A-R-A
Child F	A-R-A	R-A-R
Child D	R-A-R	A-R-A

Figure 6.4: Settings for the different children. Setting 1 corresponds to session 1 and setting 2 corresponds to session 2.

Each robot's mode was signaled to the child by a sticker with a specific geometrical form drawn on it (a triangle for adaptive and a circle for reactive); the right sticker was put on the back of the robot at the beginning of each step. At each step, the child was told which game he/she was now playing, i.e game 1 for step 1, game 2 for step 2 and game 3 for step 3 and the child could see the experimenter putting the sticker on the back of the robot. The sticker was used as a way to give a sign to the child that something could be different between situations where the robot has the triangle and those where the robot has circle. But the child had no information about the existence of adaptive and reactive modes; he/she could only possibly observe the difference in the reactions of the robot. The different stickers were used so that it was not too hard for the child to understand that the game was different.

During each game, the child could freely interact with the robot. Before the beginning of each game, the experimenter:

1. paused the algorithm (for game 2 and 3),

2. congratulated the child and told him/her that now he/she would move on to game 2 (respectively 3),

[4]A session resulted in three steps also called games, which are, successively, step 1 (game 1), step 2 (game 2) and step 3 (game 3).

3. put the corresponding sticker on,
4. sent the 'new robot's mode' through a wireless connection to the robot,
5. resumed the algorithm for the detection of play styles with the new robot's mode.

Each game lasted several minutes (depending on the children's specific needs and abilities); the minimum duration of each step was approximately 3 minutes. The experimenter did not touch the robot during the trials, except for putting on the sticker at the beginning of each step (but, of course, sensors data were not collected at this stage), neither did she try to influence the child's behaviour in any way. In this part, the experimenter did not take part in the interactions taking place with the robot in order not to interfere with the purpose of this study which had to focus on dyadic natural interactions between the child and the robot, in order to test the potential of an adaptive robot in terms of influencing the children's play styles[5].

Measures The experiments were video recorded. The sensor's data and the interaction styles detected with respect to the gentleness and the frequency of the interaction were recorded. These data were then analysed quantitatively. For the criterion gentle/strong, we actually looked at the overall proportion of sensor's activation and at the ratio of strong interaction styles. For the criterion 'frequency of the interaction', we took into account its evolution over time, which means here that we looked at the whole set of classifications, that is every 32 ms.

6.2.3 Results

6.2.3.1 Statistical analysis on the engagement in the interaction and on the gentleness of the strokes

Engagement in the interaction: In this paragraph, we study whether the adaptive robot may have a positive impact on the engagement of the children in play. Here, we do not consider the specificity of the strokes, i.e. whether they are gentle or strong. Instead, we are interested, for each child, in the total number of sensors' activations, that we want to compare for adaptive and reactive modes.

[5]Future work could consider introducing the adaptive robot in the approach inspired by non-directive play therapy where the experimenter participates in the experiments (Chapter 4).

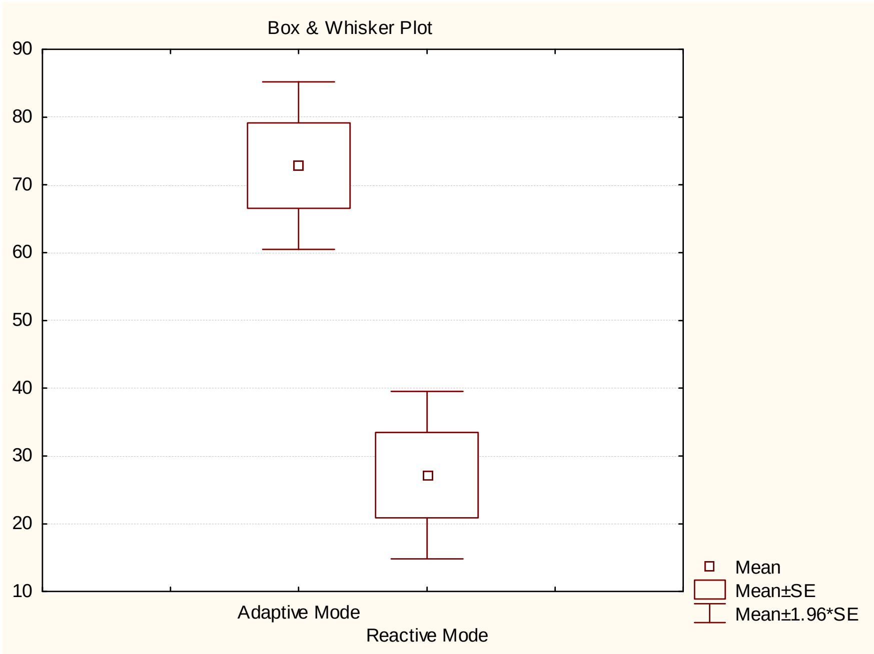

Figure 6.5: Mean, Standard Error of the Mean (SE) and Confidence Intervals for the sensors' activation on the two modes. The x-axis represents the two modes; the y-axis represents the repartition in percentage of the sensors's activation.

For each child and for each mode, we take the total number of times sensors were activated (each sensor[6] activated counts as one activation whatever the continuous external sensor it is), namely, $N(Reactive)$, for the reactive mode[7], and $N(Adaptive)$, for the adaptive mode[8]; for each child, we analyse the relative ratio of each mode[9], as follows:

$$r(Reactive) = \frac{N(Reactive)}{N(Reactive)+N(Adaptive)}$$

$$r(Adaptive) = \frac{N(Adaptive)}{N(Reactive)+N(Adaptive)}$$

The Wilcoxon test is applied to the data from the seven children for the two following variables (Fig 6.6 and Fig 6.7): $r(Adaptive)$, representing the adaptive mode, and $r(Reactive)$, representing the reactive mode. The test shows that there

[6]Here we consider the activation of any of the four continuous external sensors, that are: head sensor, back sensor front, back sensor middle and back sensor back.

[7]$N(Reactive)$ is the sum over all runs conducted in the reactive mode, for a specific child.

[8]$N(Adaptive)$ is the sum over all runs conducted in the adaptive mode, for a specific child.

[9]Some children will naturally interact a lot with the robot, while others may stroke the robot only a few time during a session, thus we prefer to look at relative ratios.

Child	**r(Adaptive) in percentage** Relative percentage of the number of activations on the adaptive mode (per child)	**r(Reactive) in percentage** Relative percentage of the number of activations on the reactive mode (per child)
Child A	72.22	27.78
Child G	87.84	12.16
Child H	44.74	55.26
Child C	82.01	17.99
Child E	57	43
Child F	75.56	24.44
Child D	90.48	9.52

Figure 6.6: Table of data for the Wilcoxon test to compare the engagement of the children on adaptive and reactive modes.

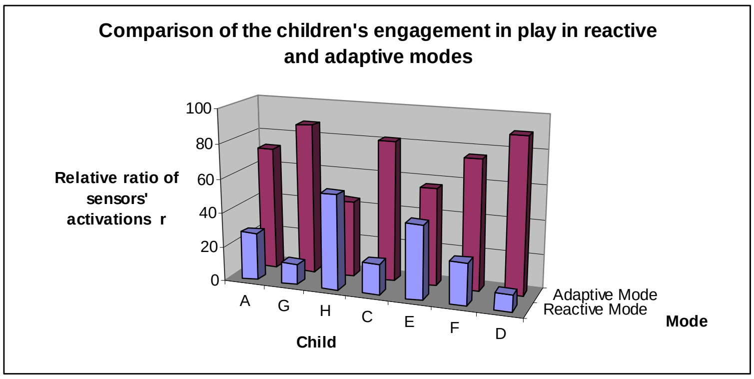

Figure 6.7: Graph showing the relative engagement of the children in adaptive and reactive modes.

is a significant effect of the experimental conditions adaptive versus reactive[10] (for $T = 1.000$, $p < 0.028$, with $N = 7$, Fig. 6.5). Thus, we can conclude that the children engage significantly more in the interaction[11] when the robot is adaptive (in comparison with reactive). It is worth noting that those results include all settings (A-R-A and R-A-R) and all steps[12] constituting a session: step 1, 2 and 3.

[10] In brief (more details on the Wilcoxon test can be found in (Siegel and Castellan, 1988)), for $T_{observed}$ (this is the T calculated, in this case, $T = 1.00$),

- If $p(T <= T_{observed}) < 0.05$: i) rejects H_0, ii) accept H_A: there is a significant difference between the conditions adaptive and reactive.
- If $p(T <= T_{observed}) >= 0.05$, i) accept H_0: there is no significant difference between the conditions adaptive and reactive, ii) reject H_A.

[11] In this context, we measure the engagement of the child in terms of how many times he/she activates sensors. Those sensors are the continuous external sensors. Thus, when we say that a child 'interacts more with the robot' or 'engages more in the interaction', it implicitly means that the child activates a higher number of times these sensors.

[12] We shall remind the reader that each step is defined by the mode of the robot, two successive steps are defined by different modes.

Gentleness of the interaction: We are interested in this paragraph in the nature of the activation in terms of gentleness, i.e whether an activation is gentle or strong. We therefore consider here the percentage of strong strokes (also called strong activations) among the total number of sensors' activations, per run and per child. For each child and for each mode, we take the average of this percentage over the runs from the two sessions[13] (Fig. 6.8).

Child	Average percentage of strong activations in the adaptive mode	Average percentage of strong activations in the reactive mode
Child A	20.52	71.97
Child G	2.08	12.50
Child H	5.56	9.09
Child C	3.53	11.75
Child E	15.23	15.79
Child F	17.51	67.74
Child D	60.58	33.33

Figure 6.8: Table providing the average percentage of strong strokes in each mode for each child.

The Wilcoxon test is applied to the data from the seven children for the two following variables (Fig 6.8): the average of the percentage of strong strokes in the adaptive mode and the average of the percentage of strong strokes in the reactive mode. The test shows that there is no significant effect of the experimental conditions on the gentleness of the strokes ($N = 7$ and, for $T = 5.00$, on gets $p < 0.128$): there is no significant difference in the amplitude of the average percentage of strong strokes between adaptive and reactive modes. However, the proportion of cases where this average is smaller in the adaptive mode is 6 cases out of 7. The probability of obtaining such a deviation (6 or more cases out of 7) from a fifty-fifty ratio is 0.016 (two-tailed probability in the binomial test[14]) which shows that, in the adaptive mode, the percentage of children who react less strongly in the adaptive mode deviates significantly from a fifty-fifty ratio.

6.2.3.2 Detailed analysis per child on the engagement in the interaction and on the gentleness of the strokes

In order to analyse results in more details, we conduct a detailed analysis per child. Each child's analysis starts with a description of the impact of the adaptive robot

[13]Concerning the criterion Gentleness, we want to encourage children to play more gently. Thus, we are looking at the ratio of strong activations and investigate whether this ratio is inferior when the robot is in the adaptive mode, compared with when the robot is in the reactive mode.

[14]A probability of deviation lower to 5% points to a significant outcome.

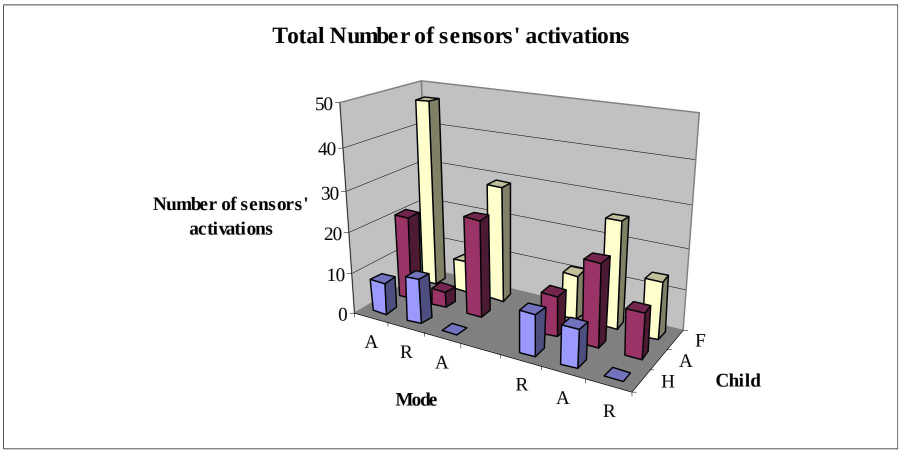

Figure 6.9: Total of activations for Children H, A and F. Setting A-R-A followed by R-A-R.

on his/her engagement in the interaction. Then, the nature of these activations is described in terms of gentleness, and how the adaptive mode might encourage a higher proportion of gentle strokes (in comparison with strong strokes). To this end, in this subsection, we analyse the changes not only in terms of average over the runs from the two sessions, but also in detail for each session. This detailed analysis enables to characterize the changes in the gentleness of the strokes between two successive runs of a same session, one in reactive mode and the other in adaptive mode. In other words, this detailed analysis characterizes, for each session, the differences in the proportion of gentle strokes when the robot was adaptive and when it was reactive. Such an analysis notably enables to get additional insight on the impact of the adaptive robot by reporting, for each child separately, on common tendencies or changes between Session 1 and Session 2.

Child A Fig 6.9 shows that Child A engaged much more in activating the continuous external sensors when the robot was in the adaptive mode. The video analysis shows, in particular, that the engaging behaviour of the robot ('wagging the tail') had a real impact on the child, since it managed to attract the child's attention, which was usually distracted by the presence of doors. The 'wagging the tail' of the robot, resulting from the fact that no interaction was detected, actually attracted the child's attention and often resulted in the child stroking the robot. Besides, in the adaptive mode, Child A also tended to reiterate the activation of the sensors until he actually got a feedback from the robot. Fig. 6.10 shows that, within each session, the

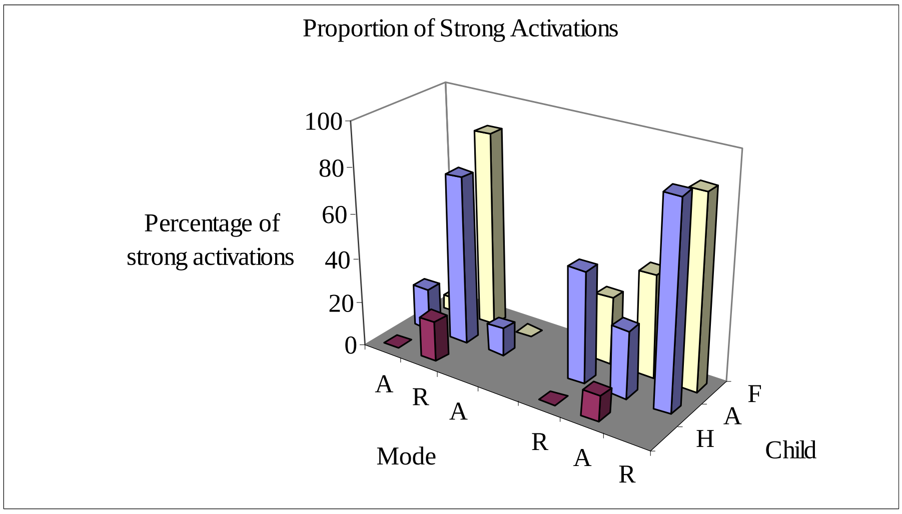

Figure 6.10: Percentage of Strong activations (versus Gentle activation) for Children H, A and F. A-R-A followed by R-A-R. The graph represents the percentage of strong activations among all the activations (the activations can be either gentle or strong). When a bar is absent, it means that the child did not activate any external sensor during the corresponding run. The precise values are displayed in Fig. A.2.

ratio of strong strokes was clearly higher when the robot was in the reactive mode (in comparison with the adaptive mode). This means that the child played much more gently while the robot was in the adaptive mode than when it was in the reactive mode, whatever the order (A-R, R-A) or the setting (A-R-A, R-A-R).

Child G Fig. 6.11 shows that Child G engaged much more in activating the continuous external sensors when the robot was in the adaptive mode. During the long-term study carried out beforehand to observe the child's play styles and tailor the robot's behaviours appropriately according to his needs and abilities[15] (Appendix C), Child G actually tended to interact in a way that could be qualified as a bit 'shy' as i) he tended to stroke so lightly the sensors they would not be activated[16], and ii) he regularly 'avoided' zones with sensors to stroke non sensing parts of the robot. Here, the adaptive robot has clearly encouraged the child to stroke and even activate the sensors of the robot. In particular, the video analysis shows that when the robot wagged the tail (as an engaging behaviour) the child tried to activate sensors[17]. Note

[15] During these long-term trials the robot was in the reactive mode.

[16] In order for a sensor to be considered as 'activated', its value must reach a predefined threshold.

[17] In brief, in the adaptive mode: 1) the stroke of any external sensor stopped the robot's engaging behaviour (i.e. it made the tail stop wagging) and 2) a gentle activation of any sensors gave rise to

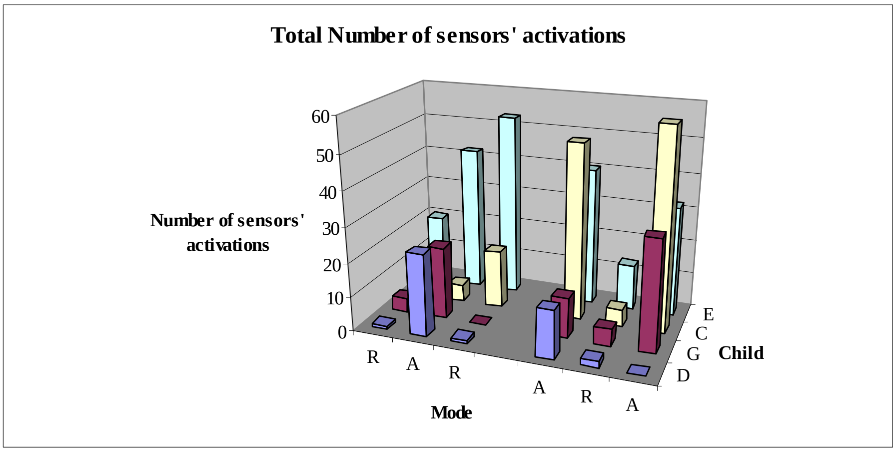

Figure 6.11: Total of activations for Children D, G, C and E. Setting R-A-R followed by A-R-A.

that the duration of a session might have been a bit too long for this child, which might partly explain that he did not interact at all with the robot on the last step of the first session when the robot was in the reactive mode.

Concerning the proportion of gentle (versus strong) strokes, Fig. 6.8 indicates that, on average, Child G interacted more gently with the robot in the adaptive mode. The analysis in detail of each session (Fig. 6.12) shows that in the first session (setting R-A-R) the child clearly interacted more gently in the first run in adaptive mode than in the run in reactive mode. As already reported on, in the last run, the child did not activate any sensor. During the second session (setting A-R-A), child G did not stroke the robot strongly, both in the first run (adaptive mode) as well as in the second run (reactive mode); he showed a few strong behaviours in the last step of the session (third run), when the robot was back in the adaptive mode.

Child H Child H tended to interact slightly more in the reactive mode than in the adaptive mode as show Fig. 6.7 and Fig 6.9. In each session, the child did not activate the robot's sensors during the last step of the session, which might suggest that the sessions were a bit too long for him. In order to understand better those results, it should be noted that during the preceding long-term study for the design of the robot's behaviours (Appendix C), Child H was always very interested in playing with the robot and he always showed lots of concentration on exploring the features and capabilities of the robot. Child H did not use verbal communication but rather used

an appropriate robot's behaviour as a positive feedback to the child.

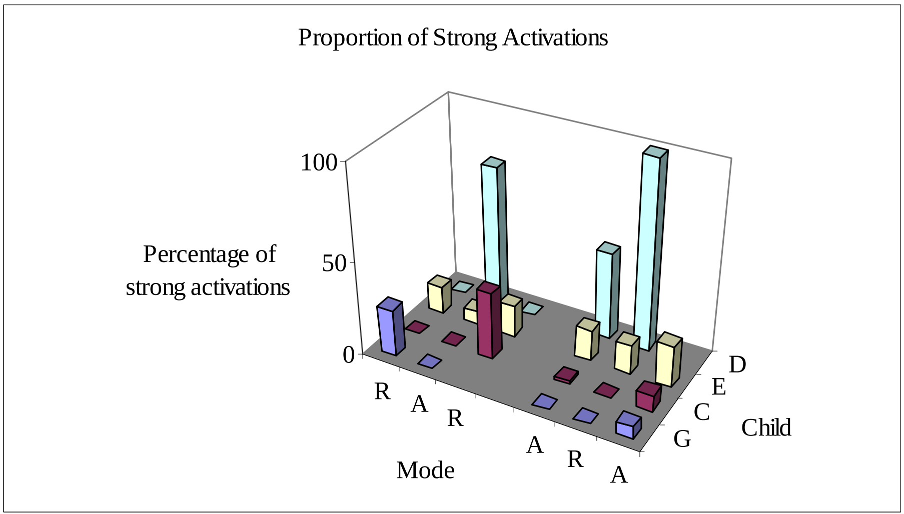

Figure 6.12: Percentage of Strong activations (versus Gentle activation) for Children D, G, C and E. R-A-R followed by A-R-A; for more details on the graph, refer to Fig 6.10. The precise values are displayed in Fig. A.3.

tactile interaction a lot, looking at the robot from various positions. The analysis of the video shows that the engaging behaviour of the robot (i.e. wagging the tail when the child did not interact) may have had the contrary effect to the one expected for Child H: he actually seems to have been discouraged (at least in the first -early- stage) to interact with the robot since he positioned himself rather as an observer than as an actor in the game. Child H may actually have understood that, during this game, where the robot was in the adaptive mode, he should rather look at the robot than stroke it, which is what he actually mainly did. This particular case shows, again, how difficult it is to design appropriate solutions that would possibly encompass the diversity of profiles, needs, abilities and personalities that recover autism; this is also a reason why statistical results must always be considered cautiously when analysing studies with those with autism.

Concerning the criterion gentle/strong, Fig. 6.8 indicates that, on average, Child H interacted more gently with the robot in the adaptive mode. The analysis of each session in detail for this child shows however that there is no clear tendency in the changes on the gentleness of the interaction between two successive runs of a same session, one when the robot is adaptive and one when the robot is reactive (Fig. 6.10). In the first session, the percentage of strong strokes was higher in the reactive mode while, in Session 2, this percentage was higher in the adaptive mode. It seems that child H tended to interact more strongly during second steps than during the first ones in both settings (setting one : A-R-A, setting 2: R-A-R), whatever the mode it corresponds to. Should we compare those two runs, the percentage of strong strokes from the adaptive mode is slightly smaller.

Child C Child C engaged on average a lot more in the interaction when the robot was adaptive than when it was reactive (Fig. 6.7). The analysis in detail per session shows however that Child C interacted much more in the second session with the robot (Fig. 6.11). The video analysis indicates that Child C showed at first some hesitation because the play session was not as ‘usually’. In particular, the experimenter did not join the games. This can be explained by the fact that Child C naturally used lots of verbal communication and, during the long-term study presented in Chapter 4, she showed very capable of playing socially. In the long-term study presented in Chapter 4, she actually experimented with more situations of social play with both the robot and the experimenter, than with dyadic situations of play with the robot only. Here, because of the purpose of this study on adaptation, the experimenter had to not take part in the trials and it might have taken some time to Child C to really engage in the dyadic interaction. It should be further noted that, during the long-term study presented in Chapter 4, Child C experienced many and various play situations involving pretend and symbolic play. But the time where she naturally engaged in pure tactile dyadic play were not very long compared to other situations of play. She tended to embed those tactile interactions in a story or in a game of hugging the robot. However, the fact that her engagement increased a lot between Session 1 and Session 2 suggests that Child C progressively coped with (and even maybe adapted to) this new game and actively engaged in it.

Looking now at the percentage of strong strokes (versus gentle strokes) per run, Fig. 6.8 indicates that, on average over the two sessions, Child C interacted more gently with the robot in the adaptive mode. Nonetheless, the analysis of each session in detail suggests different tendencies (Fig. 6.12): in the first session, the percentage of strong strokes was clearly lower in the adaptive mode while, in the second session, this percentage was slightly higher in the adaptive mode.

Child E Child E tended to interact more with the robot when the robot was in the adaptive mode than when it was in the reactive mode (except one case), especially in session 2, for the setting A-R-A (Fig. 6.11).
Concerning the percentage of strong strokes per run, it was slightly lower on average in the adaptive mode (Fig. 6.8). During the first session, with the setting R-A-R,

Child E interacted more gently in the adaptive mode than in the reactive mode (Fig. 6.12). However, during the second session, with the setting A-R-A, Child E interacted slightly more gently in the reactive mode (Fig A.3). In both settings Child E interacted more gently in step 2. For this child, further trials should be carried on to be able to refine the characterisation of the impact of the robot's mode on the gentleness of the interaction.

It is interesting to note that, in the previous play sessions conducted, like Child H, Child E was always very interested and concentrated on investigating the possible features and capabilities of the robot. He liked exploring tactile interaction with the robot. The difference with Child H is that Child E used verbal communication a lot, and, during the long term study described in chapter 4, he did experience with social play and symbolic and pretend play situations a lot. Child H did also take part in play sessions with the approach inspired by non-directive play therapy[18]. He did progress a lot, but, unlike Child E, did not experiment with symbolic or pretend play. He did play with accessories and he did play socially with both the robot and the experimenter, but the situations of social play were 'limited' to the game 'ask for a physical reaction, show it with a sensor'[19].

Child F Child F engaged much more in the interaction in the robot's adaptive mode, compared to the robot's reactive mode, whatever the setting (A-R-A or R-A-R) and the order (A-R or R-A), see Fig 6.9. Concerning the nature of the strokes, Fig. 6.8 shows that the proportion of strong strokes per run was, on average, fairly lower in the adaptive mode than in the reactive mode. The detailed analysis of each session indicates only one case in which Child N interacted more gently in the reactive mode than in the adaptive mode in a same session. This happened in Session 2 (Setting R-A-R): in the first run (reactive mode) the percentage of strong strokes was 30 while in the second run (adaptive mode), this percentage was 46.1 (Fig. 6.10).

Child D Child D engaged more in the interaction in the robot's adaptive mode than in the robot's reactive mode, except on the last step of session 2 where the child did not interact at all with the robot (Fig. 6.11). For child D, it might be that the session lasted a bit too long for ensuring him to be playing with the robot during the three steps of each session. Child D usually often needs to be encouraged to engage in

[18]These play sessions with Child H were conducted after the ones presented in Chapter 4.

[19]In the game 'ask for a physical reaction, show it with a sensor', the experimenter asks the child to show a physical reaction of the robot; the child then tries to activate the right robot's sensor that leads to that specific behaviour of the robot.

the interaction and the robot's engaging behaviour (wagging the tail) in the adaptive mode, seems to have encouraged him a lot playing with the robot.

Concerning the strokes, the adaptive mode did not encourage the child to play more gently, as Fig. 6.8 and Fig. 6.12 show. During playtime in school, Child D stays very much isolated instead of playing with other children (for more details, refer to Section 4.4.1). When using the approach inspired by non-directive play therapy, Child D progressed a lot, progressively opening to some basic situations of social play (Section 4.5). In later stages of the long-term study with the approach inspired by non-directive play therapy (later stages than the ones reported in Chapter 4), Child D even played a lot the game 'aim at a physical reaction, show it with a sensor'. Nevertheless, he may need more time and might need some additional slight guidance from the experimenter to become sensible to the actual quality of touch (gentle/strong). At his stage, he seems to rather focus on the spatial distribution of touches.

6.2.3.3 Impact of the adaptive robot on the frequency of interaction

To analyse the impact of the adaptive robot on the frequency of interaction, we look at the four classes S_0, S_1, S_2, S_3 and how their occurrence varies, in a same session, between a run in the robot's reactive mode and a run in the robot's adaptive mode.

We define R as the set of the three runs (steps) within a session for a specific child and $N_{S_i}(r)$ as the number of events from a class S_i for a specific run r. For each class S_i, each child, and each session, we define the relative ratio $\rho_{S_i}(r)$ for a given run r, defined as follows:

$$\rho_{S_i}(r) = \frac{N_{S_i}(r)}{\sum_{\tilde{r}\in R} N_{S_i}(\tilde{r})} \tag{6.1}$$

For each child, for each mode m (adaptive or reactive) and for each class S_i, the average relative ratio over the two sessions is called $Av_m(\rho_{S_i})$. For each child and for each mode m, the average relative ratio over the four classes is called $Av_m(\rho)$.

Note that the division by the factor $\sum_{\tilde{r}\in R} N_{S_i}(\tilde{r})$ in Eq. 6.1 enables:

- to normalize the data according to the average activity for each class during a session, which enables to compare the effect of the adaptive mode i) for a specific class S_i between different sessions (the average richness and engagement of a child may vary from one session to the other) and ii) for different classes (each child may play with a particular tendency to favor some frequencies of

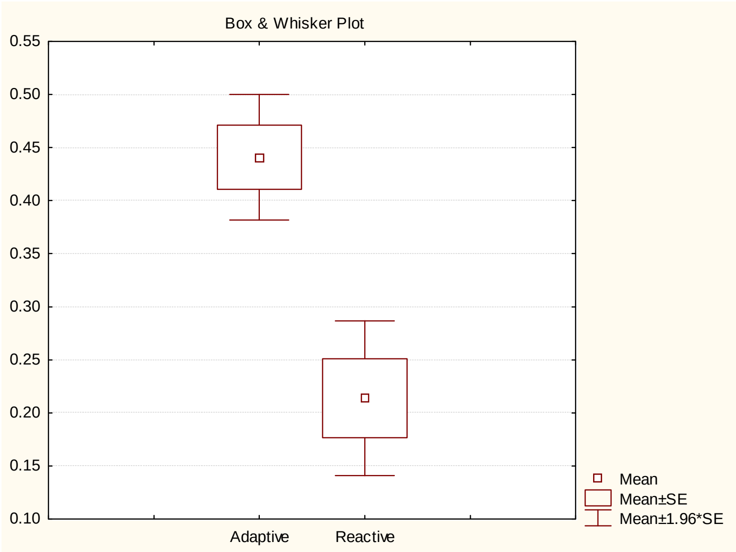

Figure 6.13: Mean, Standard Error of the Mean (SE) and Confidence Intervals. The richness of the interaction is measured in terms of occurrences of events happening in each class S_0, S_1, S_2, S_3.

interaction –i.e. some classes may be more present than others; the normalisation enables to look at variations between adaptive and reactive modes with a comparable scale for the different classes);

- to remove any possible artefact on a possible not equilikely probability distribution over the classes, which would be due to i) the inherent definition of the classes and ii) the Cascaded Information Bottleneck Algorithm, which, as implemented for these experiments, concerning the frequency of interaction, only classifies events (sensors data frames) starting with a non null value (see Section 5.3.5.1). Given the definition of the classes, this constraint may thus naturally lead, for a same duration of interaction in two different frequencies, to the classification of more events for higher frequencies than for lower ones.

Thus, we can then compare the different $Av_m(\rho_{S_i})$ for the different classes in order to determine which classes are particularly positively impacted by the robot's adaptive mode (in comparison with reactive).

The Wilcoxon test is firstly applied to the two following variables: $Av_{Adaptive}(\rho)$ (representing the adaptive mode) and $Av_{Reactive}(\rho)$ (representing the reactive mode). The test shows that there is a significant effect of the experimental conditions (adaptive versus reactive) since for $T = 0$, one has $p < 0.018$, with $N = 7$ (Fig. 6.13).

We can conclude that, in the adaptive mode, the interactions are significantly richer than in the reactive mode.

Secondly, the Wilcoxon test is applied for each class i separately, to the following variables: $Av_{Adaptive}(\rho_{S_i})$ (representing the adaptive mode) and $Av_{Reactive}(\rho_{S_i})$ (representing the reactive mode) with the following results:

- **class** S_0 **(Fig. 6.14)**: for $T = 5.000$, $p < 0.128$ ($N = 7$), thus there is no significant difference between the two experimental conditions (adaptive versus reactive) for the class S_0: there is no significant difference in the amplitude of the average relative ratios $Av_{Adaptive}(\rho_{S_0})$ and $Av_{Reactive}(\rho_{S_0})$. However, the proportion of cases where $Av_{Adaptive}(\rho_{S_0}) > Av_{Reactive}(\rho_{S_0})$ is 6 cases out of 7. The probability of obtaining such a deviation (6 or more cases out of 7) from a fifty-fifty ratio is 0.016 (two-tailed probability in the binomial test) which shows that the percentage of children for which there are more events related to S_0 in the adaptive mode than in the reactive mode deviates significantly from a fifty-fifty ratio.

Child	$Av_{Adaptive}(\rho_{S0})$	$Av_{Reactive}(\rho_{S0})$
Child A	0.430	0.237
Child G	0.230	0.437
Child H	0.393	0.274
Child C	0.461	0.205
Child E	0.369	0.298
Child F	0.406	0.261
Child D	0.531	0.136

Figure 6.14: InputData for the Wilcoxon test applied to S_0.

Child	$Av_{Adaptive}(\rho_{S1})$	$Av_{Reactive}(\rho_{S1})$
Child A	0.423	0.243
Child G	0.667	0
Child H	0.133	0.533
Child C	0.556	0.111
Child E	0.360	0.306
Child F	0.333	0.333
Child D	0.333	0

Figure 6.15: InputData for the Wilcoxon test applied to S_1.

- **class** S_1 **(Fig. 6.15)**: for $T = 4.000$, $p < 0.173$ ($N = 7$), thus there is no significant difference between the two experimental conditions (adaptive versus reactive) for the class S_1: there is no significant difference in the amplitude of

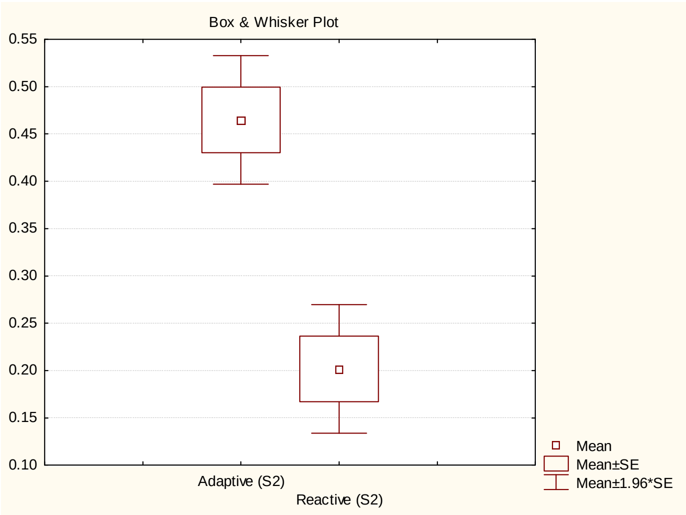

Figure 6.16: Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_2. The two variables are $Av_{Adaptive}(\rho_{S_2})$ and $Av_{Reactive}(\rho_{S_2})$.

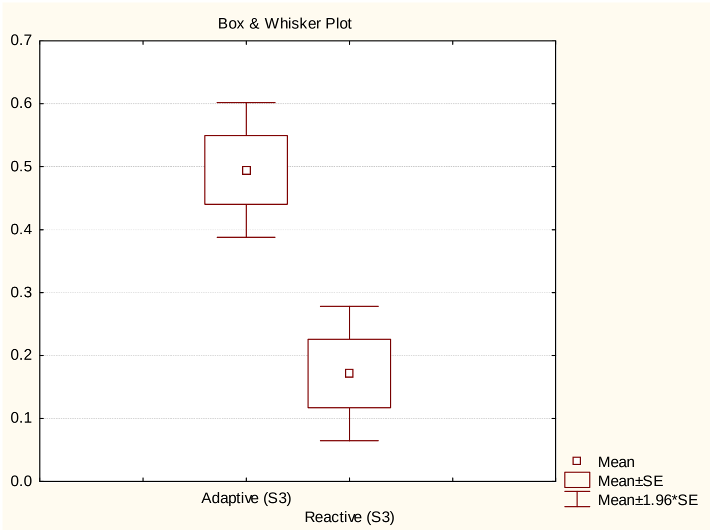

Figure 6.17: Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_3 The two variables are $Av_{Adaptive}(\rho_{S_3})$ and $Av_{Reactive}(\rho_{S_3})$..

the average relative ratios $Av_{Adaptive}(\rho_{S_1})$ and $Av_{Reactive}(\rho_{S_1})$. However, the proportion of cases where $Av_{Adaptive}(\rho_{S_1}) > Av_{Reactive}(\rho_{S_1})$ is 6 cases out of 7. The probability of obtaining such a deviation (6 or more cases out of 7) from a fifty-fifty ratio is 0.016 (two-tailed probability in the binomial test) which shows

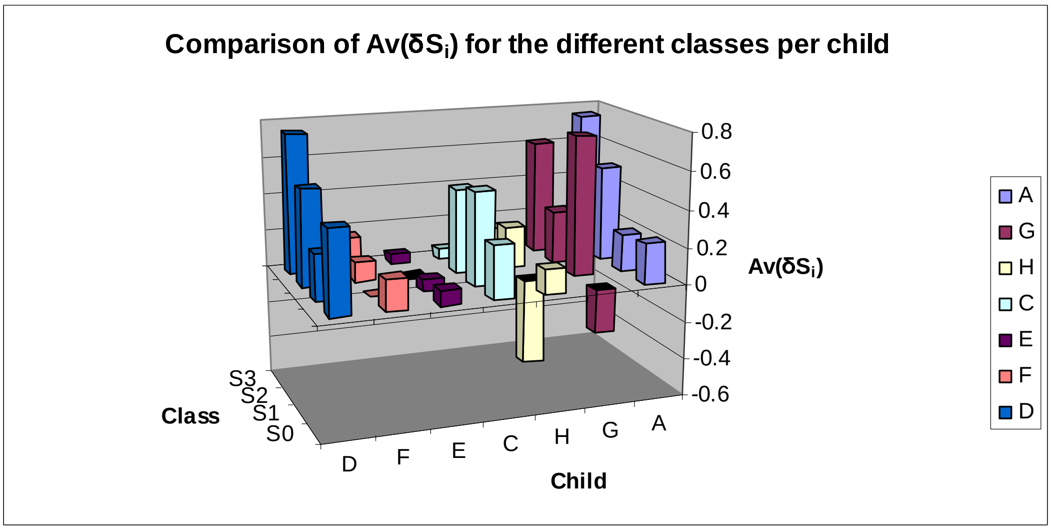

Figure 6.18: Average of δS_i over the two sessions per child. The precise values are provided in Fig. A.4.

that the percentage of children for which there are more events related to S_0 in the adaptive mode than in the reactive mode deviates significantly from a fifty-fifty ratio.

- **class** S_2 **(Fig. A.6)**: for $T = 1.000$, $p < 0.028$ ($N = 7$, Fig. 6.16) thus, there is a significant effect of the experimental conditions Adaptive and Reactive with respect to the class S_2: in the adaptive mode, there are significantly more events from class S_2 than in the reactive mode.

- **class** S_3 **(Fig. A.7)**: for $T = 0.000$, $p < 0.018$ ($N = 7$, Fig. 6.17) thus, there is a significant effect of the experimental conditions Adaptive and Reactive with respect to the class S_3: in the adaptive mode, there are significantly more events from class S_3 than in the reactive mode.

We shall now describe in detail the results for each child. For each child, each class and each session, we look at the parameter $\delta(S_i)$, which measures the difference between the adaptive mode and the reactive mode in terms of events from a class S_i for two successive runs in a same session. It is defined as follows:

$$\delta_{S_i} = \rho_{S_i}(Adaptive) - \rho_{S_i}(Reactive) \tag{6.2}$$

For each child and each session, we get two measures of δ_{S_i} (see the table on

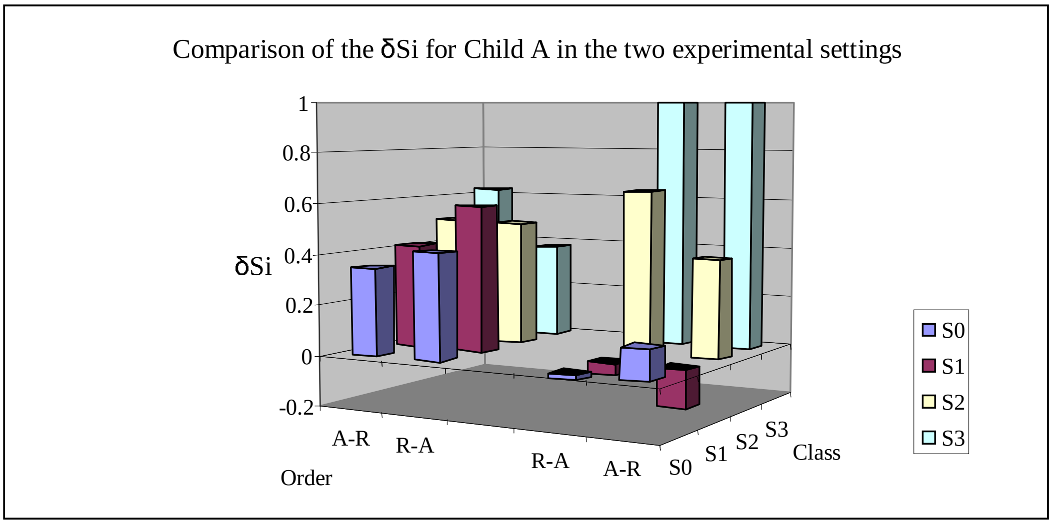

Figure 6.19: Comparison of δ_{S_i} for Child A. The precise values are provided in Fig. A.16 .

Fig. A.16 in Appendix A):

- for the setting A-R-A: δ_{S_i} between the first adaptive run and the reactive run, and δ_{S_i} between the last adaptive run and the reactive run;
- for the setting R-A-R: δ_{S_i} between the adaptive run and the first reactive run, and δ_{S_i} between the adaptive run and the last reactive run.

δS_i therefore compares the distribution of the events from a class S_i between two successive runs from a same session (one adaptive and one reactive): δS_i is positive if and only if there are more events from class S_i in the run in the adaptive mode than in the run in the reactive mode. The average of δS_i on the runs from the two sessions per child is called $Av(\delta S_i)$ and is provided in Fig 6.18.

Child A On average over the two sessions, for Child A, the adaptive robot encouraged principally the frequencies represented by the classes S_3 and S_2 (Fig 6.18). In the first session (setting A-R-A) the $\delta(S_i)$ of the four classes stayed quite close to each other (i.e. one group) while in the second session (setting R-A-R) S_2 and S_3 did increase a lot, while the two others (S_0 and S_1) decreased significantly compared to session 1 (Fig. 6.19). This suggests that the adaptive mode of the robot tended, for Child A, to trigger higher frequencies (first position for S_3 and second one for S_2).

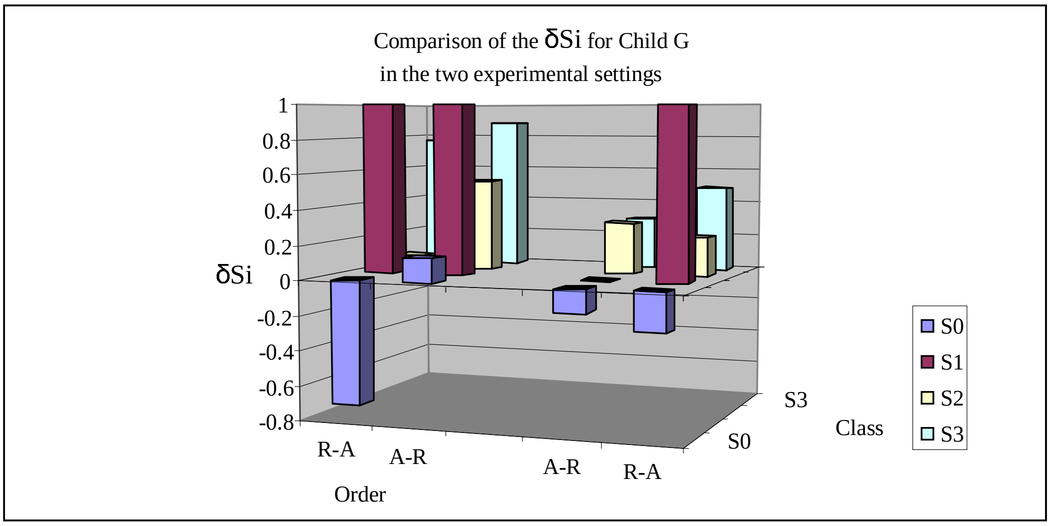

Figure 6.20: Comparison of δ_{S_i} for Child G. The precise values are provided in Fig. A.16 .

Given the profile of Child A[20] (see Section 4.4.1), this result highlights the potential of the adaptive robot's mode to keep longer the attention of the child, that is, an uninterrupted period of play with the robot (i.e. typically, a phase of play with the robot between two phases of looking at doors) tended to be richer in interactions when the robot is adaptive. Child A might also have progressed or learned between session 1 and 2 since the profiles of the classes fairly differ (Fig. 6.19). Note that this assumption should be further confirmed with a long-term study in future work. Here we carefully describe tendencies.

Child G The presence of S_1, S_2 and S_3 globally significantly increased (their respective δ_{S_i} is greater and sometimes even much greater than 0) in the adaptive mode while S_0 decreased (i.e. δ_{S_0} was negative), except from one run where it did slightly increase. δ_{S_3} is always bigger (or approximately equal) than δ_{S_2} (Fig. 6.20). In both sessions, the adaptive mode encouraged the apparition of events from class S_1 which did not happen during the reactive runs.

Child H On average, the biggest increase in terms of δ_{S_i} concerned the class S_2 (Fig 6.18). Concerning the detailed analysis for each session, the graph provided in Fig. 6.21 shows two different tendencies: in session 1 (setting A-R-A) S_0 is the leader,

[20]Child A is fascinated by doors and often looks at them, even during play sessions with the robot. He often alternates stroking the robot and looking at doors.

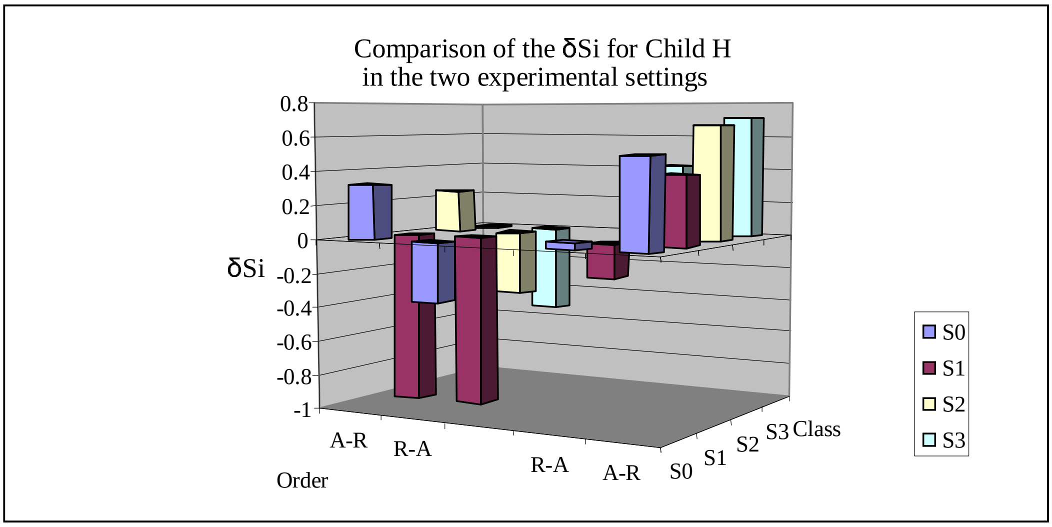

Figure 6.21: Comparison of δ_{S_i} for Child H. The precise values are provided in Fig. A.16 .

S_2 is just behind, in second position. S_3 is on third position. In session 2 (Setting R-A-R), S_3 is the leader and S_2 is just behind (second position). As described in the analysis on the criterion Gentle/Strong, during the first session, Child H tended to observe the robot rather than engage in play with it while the mode was adaptive. In particular, during the last run of the first session, Child H did not stroke the robot at all. Results from the second session show that the frequency of interaction has been pretty high in the adaptive mode compared with the reactive mode[21]. This indicates that the child's reaction to the adaptive robot may have changed between Session 1 and Session 2. Having observed the child during long-term studies and with the highlight of the video analysis, my hypothesis is that the child is naturally interacting a lot with the robot and that he might have been a bit surprised when the robot showed a reaction by itself without stimulations. As I described in the paragraph on the criterion gentle/strong, he might first have understood that the new game was in this case looking at the robot rather than stroking the robot, or he might simply have felt a bit hesitating. In session 2 he was far less hesitating with the adaptive mode and engaged longer in the interaction with the robot and with higher frequencies. This suggests that Child H might have progressively 'adapted to' the adaptive mode, and particularly to the engaging behaviour of the robot. However, since Child H is naturally engaging a lot in the interaction (he usually removes his

[21]In both sessions, the high absolute values of δ_{S_i} for the two last runs might also be explained by the fact that the sessions were a bit too long and that the child was less involved in the last run than in the two first ones.

attention from playing with the robot only when it is time for him to go back to the classroom), this engaging behaviour from the robot could certainly be removed for future experiments with Child H.

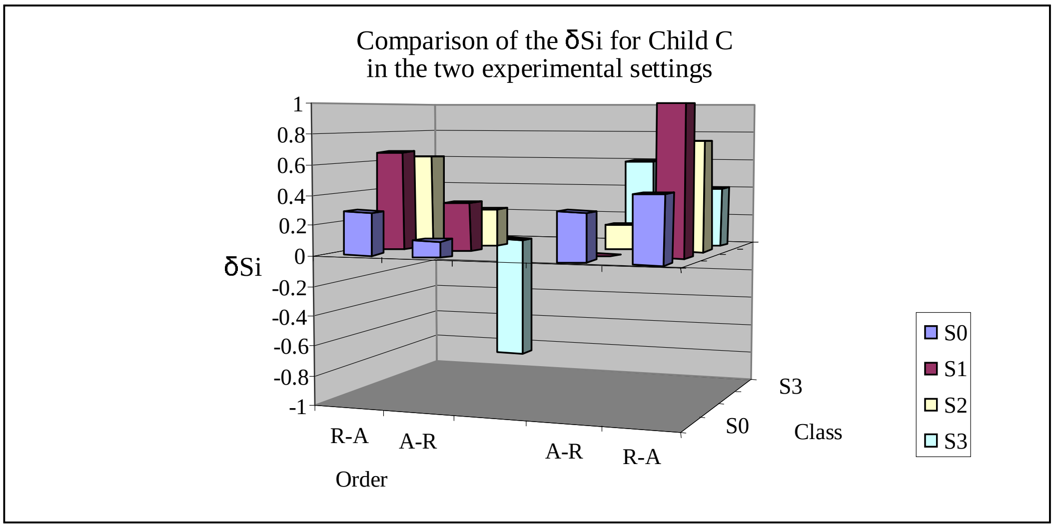

Figure 6.22: Comparison of δ_{S_i} for Child C. The precise values are provided in Fig. A.16 .

Child C On average, the adaptive robot encouraged principally frequencies from class S_1 and S_2, which correspond to a well-balanced frequency of interaction (Fig 6.18). A detailed analysis for each session shows (Fig. 6.22): a) in session 1 (setting R-A-R), S_1, closely followed by S_2, have the highest δS_i; b) in Session 2 (setting A-R-A), the first adaptive run has principally favoured S_3 with respect to δS_i while, in the second adaptive run, S_1, closely followed by S_2, had the highest δS_i.

Child E The detailed analysis of each session shows that variations in terms of frequencies between the adaptive and the reactive modes are small for Child E (Fig. 6.23), except S_1 whose δS_i moved from a negative value in Session 1 to a high positive value in session 2 (setting A-R-A). In the first session, the adaptive mode triggered, on average, mainly events from the class S_0, while, in the second session, it triggered mainly events from the class S_1. This suggests that Child E tended to slightly adjust the frequency of interaction in order to get the additional reward. This result must be linked with the analysis of the video which shows that Child E very much focused his attention on the additional reward he could get, that is the LED turning on and off (simulating the eyes of the robot): each time it happened,

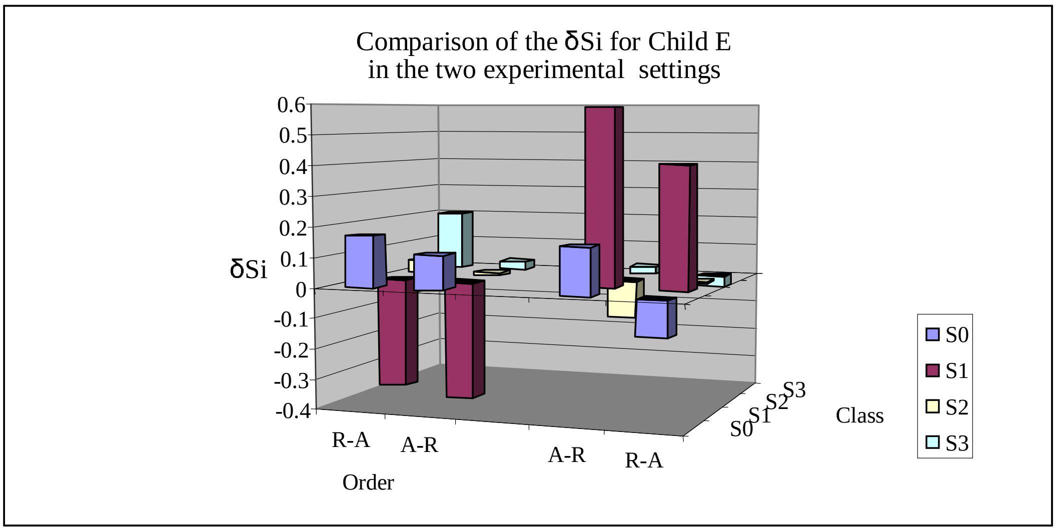

Figure 6.23: Comparison of δ_{S_i} for Child E. The precise values are provided in Fig. A.16 .

the child actually mentioned it to the experimenter. It is not clear whether the child explicitly understood that it was linked to a specific frequency of interaction. What is important here, is that the child tried to reproduce gestures that made him achieve the flashing LEDs (i.e. the reward for a good frequency), or adjust the strokes and persevere until he got the reward.

Child F In the first session (setting A-R-A) S_1, S_2 and S_3 increased a lot when the robot was in the adaptive mode, with δS_2 and δS_3 approximately equal to each other (Fig. 6.24). Unlike the other classes δS_0 decreased quite importantly between the order A-R and R-A. In the second session, the adaptive mode does not really appear to have been a facilitator factor because the δS_i are all negative or null for the comparison with the first reactive run, and only two of them (δS_0 and δS_3) are non negative in the comparison with the last reactive run. Note, results from the second session should be taken very cautiously, because during this trial, Child F seemed less receptive and less actively focused than usually.

Child D The different classes have a high δ_{S_i} (except from S_1 in the first session) in all runs but the last one in session 2, where the child did not interact at all (Fig. 6.25 and Fig 6.18). It is probably due to the fact that the session lasted a bit too long for him. Thus the last run is ignored for the analysis. There was no presence of S_1 in the first session (R-A-R). In session 2, the adaptive mode facilitated the occurrence of S_1

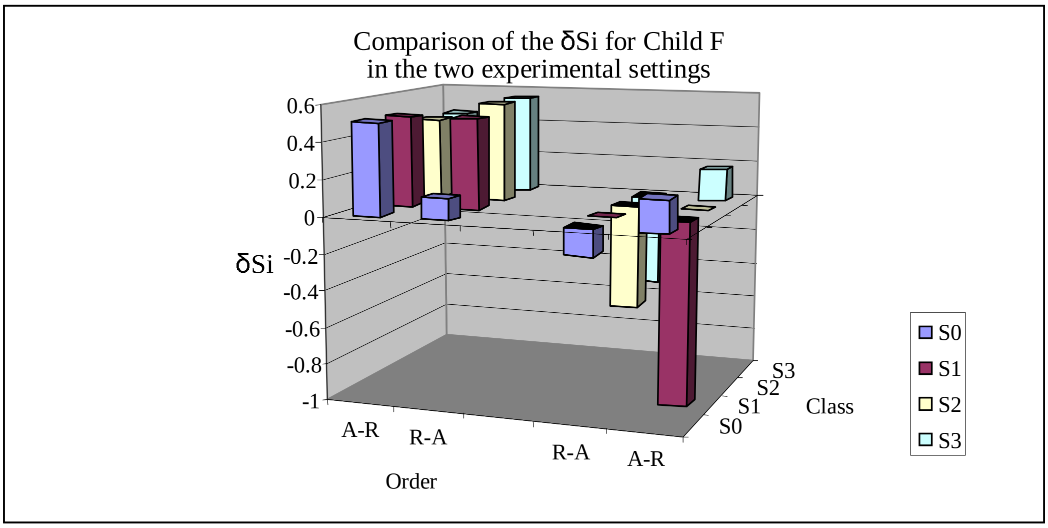

Figure 6.24: Comparison of δ_{S_i} for Child F. The precise values are provided in Fig. A.16 .

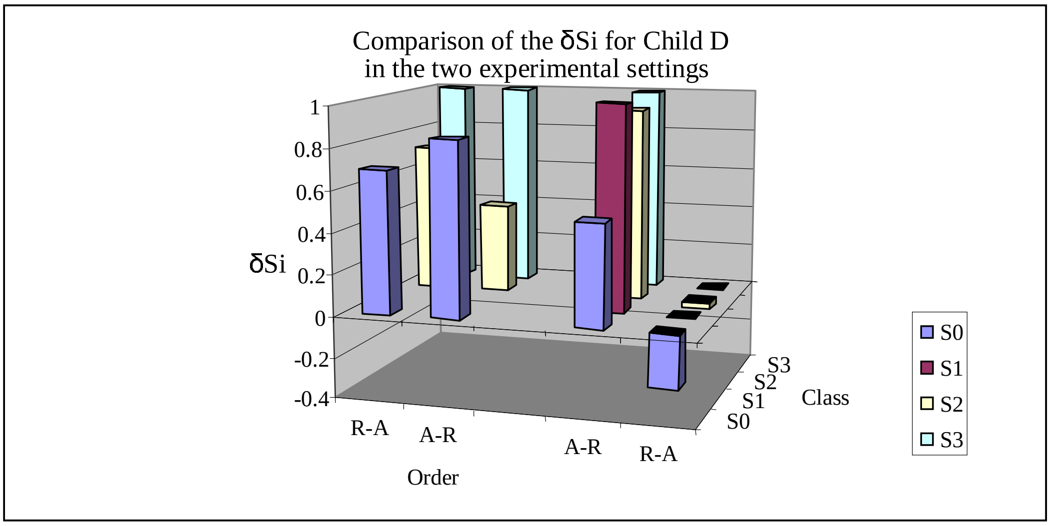

Figure 6.25: Comparison of δ_{S_i} for Child D. The precise values are provided in Fig. A.16 .

(S_1 occurred 20 times, see Fig. A.8). During this session, the δ_{S_0} decreased compared to Session 1, while all the others increased or kept the same value over time. In all cases, S_3 was absent from all runs in the reactive mode, and importantly present in the runs in adaptive mode (the number of occurrences on the first session was 182, and 22 on the second session, see Fig. A.8 in Appendix A).

6.2.4 Discussion

This study has shown that the children engaged significantly more in interaction with the robot when the robot was in the adaptive mode, in comparison with the reactive mode. Moreover, the adaptive mode had some positive effects on the nature of the interaction, with respect to the gentleness of the interaction. On average over the runs from the two sessions, all the children except one child interacted more gently with the robot when it was adaptive than when it was reactive.
Furthermore, the analysis in detail per child of the changes in the percentage of strong strokes for two successive runs from a same session (one run in the adaptive mode, the other in the reactive mode) shows that the tendency to play more gently for each run in the adaptive mode (compared with reactive mode) within each session was very clear for Child A. It was also the case for Child F (except one run), and, to a more basic extent for Child G. For several other children, the tendency is slightly less clear and further trials should be conducted to get a better idea of the impact on the gentleness of the interaction on a long-term basis.

Besides, the adaptive mode induced significant changing in the frequency of interaction of the children. The tactile interactions were significantly richer in the robot's adaptive mode (in comparison with reactive mode). The very high frequency (class S_3) and a well-balanced frequency (class S_2) were both significantly more present while the robot was in the adaptive mode than in the reactive one.
Further to this, a detailed analysis per child has enabled to highlight some individual tendencies. It is clear that, for Child A, the adaptive mode has encouraged higher frequencies of interaction (mainly the very high frequency S_3, and the well-balanced one S_2, Fig 6.18). Concerning Child G, the adaptive mode has encouraged higher frequencies and, on average, mainly a well-balanced frequency (S_1) followed by the very high one (S_3) (Fig 6.18). For Child C, the adaptive mode has mainly encouraged a well-balanced frequency of interaction (i.e. is S_1 and S_2). It is moreover interesting to underline that in some cases, the adaptive mode has triggered the apparition of classes that were absent in the reactive mode: for instance, the well-balanced frequency S_1 and the very high one (S_3) did happen for Child D in at least one run in the adaptive mode although they were absent in the reactive mode. Nevertheless, it seems that, while for some children the tendency is pretty clear (Child A, Child G, Child C and Child D), for some others, it would be useful to conduct further experiments to investigate some observations and hypothesis formulated here. Importantly, it also seems that, for some children, the frequency of interaction has already changed between the first and the second session (e.g. Child H and Child E). It would be very

interesting to observe further evolutions on a long-term study.

The research questions motivating this study can now be answered directly (see Section 6.2.1):

- The adaptive robot does encourage the children to engage in the interaction, and the children engage significantly more in the interaction with the adaptive robot, compared with the reactive robot.

- The children's play patterns differ when the robot is adaptive from when the robot is reactive:
  i) the strokes are qualitatively different and significantly more children play more gently with the robot in the adaptive mode; ii) the frequency of interaction differs: the tactile interactions are significantly richer in the adaptive mode and the well-balanced frequencies (S_2) and high frequencies (S_3) are significantly more present. The detailed analysis per child has enabled to identify several children for whom a clear tendency could be established, with respect to the impact of the adaptive robot on the frequency of the interaction. It has clearly a positive influence on Child A, Child G, Child C and Child D. Child H and Child E's frequency of interaction seems to have already changed and further changes should be investigated in a long-term study. Future work could also define the classes slightly differently, maybe looking at a shorter window frame for the frequency of interaction, so that this criterion might be more directly 'accessible' to the children[22].

Finally, it should be noted that, among the children (Child A, G, H and D) who, during the play sessions conducted with the approach described in Chapter 4, tended to engage mainly in tactile exploration games with the robot (and possibly engaged in the triadic play situation 'ask for a physical reaction show it with a sensor' which involves the direct use of tactile sensors to induce a precise reaction of the robot), the majority of those children (Child A, G and D) had their play styles importantly positively impacted by the robot.

[22]Looking at the periodicity over approximately 15 seconds is meaningful in this context. However, it might be sometimes difficult for the child to deal directly with this period of time which could be a bit long for some children.

6.3 Summary

In this chapter, we have presented the adaptive robot, which recognizes the play styles of the children with the Cascaded Information Bottleneck Method, and gives feedback according to those styles to the children, based on a schema for adaptation, that relies on a reward basis. We have briefly discussed the notion of discrete social potential of an adaptive robot, which relies on a transposition of the technique of 'Freezing and Freeing Degrees of Freedom' which is commonly used to learn motor skills, and which can be transposed for the context of human-robot social interaction in robot-assisted play. Further to this, we have presented a study investigating the potential role of an adaptive robot. This study was conducted in school with seven children with autism, over two sessions. We have shown that the adaptive robot did significantly positively impact the children's play styles, in terms of engagement in the interaction and in terms of the richness of the interactions generated. Those two aspects have been established with statistical techniques. The Wilcoxon test separately applied for S_2 and S_3 has besides shown that, events from class S_2 (i.e. well-balanced frequencies) and events from class S_3 (i.e. very high frequencies) happened significantly more in the adaptive mode than in the reactive mode. Moreover, on average over the two sessions, significantly more children interacted more gently with the robot in the adaptive mode (in comparison with the reactive mode).

In addition, a detailed analysis per child was conducted, which notably compared within each session the proportion of strong strokes in adaptive runs and in reactive runs. It showed that for several children, within each session, the adaptive robot clearly increased the ratio of gentle strokes. As for the frequency of the interaction, the adaptive robot triggered, for several children, frequencies that had not occurred in the reactive mode (typically S_1 and S_3 for Child D, S_1 for Child G).

This study is a step forward in the investigation of the potential role that adaptive robots can play in robot-assisted play for children with autism.

Chapter 7

Discussion and Conclusion

7.1 Limitations

This thesis has addressed a large range of issues in order to facilitate play between children with autism and an autonomous robot. Nevertheless, it presents some limitations that are exposed in this section.

Constraints on the input data for the Cascaded Information Bottleneck Method: The algorithm works on a restricted range of input windows. Concerning the criterion Gentleness of the interaction, it is perfectly fine because the restriction is limited to non null events only: only null events will not get a classification, which is perfectly fine since a null event is not categorised into gentle or strong anyway. In contrast, for the criterion 'Frequency of the interaction', the constraint is a bit bigger since it only classifies events starting with a non null value, i.e. the first value of the input data has to be non null. This constraint is nevertheless largely inherent to the nature of this criterion and, additionally, to the fact that, for each criterion, the algorithm considers a fixed window's length.

The criteria of the interaction: In this thesis, we only consider two criteria of interaction. It might be interesting to extend the analysis to a larger range of criteria. A strong point in our analysis here is that we encompass both short-term (criterion Gentleness) and mid-term (Frequency of the interaction) time scale analyses.

It might also be useful to add an additional triggering on the frequency of interaction that would guide the children more directly in the first stages of their progress. For instance, one could add a second classification of the frequency over the 5 to 7 last seconds of interaction.

The trials in school: The trials involved one experimenter only and a small number of children. The long-term study presented in Chapter 4 involved six children. Seven children participated in the trials with the adaptive robot (Chapter 6). Note that when I started the experiments in school, six children took part in the play sessions, and the trials reported in Chapter 4 involve all of them. Three other children joined the play sessions a few month later. In the study reported in Chapter 6, two children were not involved because these sessions would have been too long and too complicated for them.

The session's length in the trials on the adaptive robot (Chapter 6): For one or two children, it happened that the sessions may have been slightly too long since they did not interact during the last run, or only a bit. Finding the right balance for the session's length is challenging: on the one hand, the child's motivation to play should not be affected by a too long trial; on the other hand, enough data should be collected for the further analysis. This challenge was principally due to the fact that the three runs had to be conducted on the same day. This decision was taken because children's mood can importantly vary from one day to the other and, therefore, I thought it would be more consistent to be able to make direct comparisons between adaptive and reactive runs that happened on the same day[1].

The small number of sessions in the trials on the adaptive robot (Chapter 6): The experiments lasted only two sessions. Thus, it was possible to analyze the immediate impact of the adaptive robot that was compared, within a same session, with the reactive robot. It was also possible to see the first progress between the first and the second session. But it would be interesting to run a further long-term study that would focus on long-term changes in the play styles of the children with the adaptive robot (in comparison with the reactive robot), as well as study whether the impact of the adaptive robot might change over time, for instance, progressively while the children themselves get a better understanding of this mode[2].

[1] Moreover, the tendencies described in each session could then be compared in order to get additional insight on the possible stability or changes in the child's profile.

[2] A better understanding of the adaptive mode could, for instance, be to understand that the robot only reacts under specific strokes, and a second step would be that the child purposefully reproduces the strokes that bring the robot's reaction.

7.2 Contributions

In this thesis, I have addressed the issue of facilitating play between children with autism and an autonomous robot. I have adopted a multidisciplinary approach, which notably enabled to contribute to three main domains, the domain of robot-assisted play, the field of pattern recognition and the dimension of robots' adaptation to social contexts. We shall now go back to the research questions formulated in Section 3.5 and discuss how I answered them.

Research Question 1: *What approach for the play sessions could be adopted in robot-assisted play to enable each child with autism to progress according to his/her specific needs and abilities, that is, experiment with progressively higher levels of play and possibly develop play skills which could further help him/her cope with more complex situations of communication and social interaction, and develop imagination?*

I have designed a novel method inspired by non-directive play therapy. This method precisely describes which of the eight principles (Fig 4.1) proposed by Axline (1947), underlying the non-directive play therapy approach, are considered in our context of robot-mediated play for children with autism.

Beyond inspiration from non-directive play therapy, this method adds a regulation process which enables the experimenter to intervene under precise circumstances, in brief:

- to discourage repetitive behaviours
- to help the child engage in play
- to give a better pace to the game if it has already been experienced by the child
- to bootstrap a higher level of play
- to ask questions related to reasoning or affect

This method is a new step in robot-assisted play, which, traditionally tended to focus on a restricted repertoire of games for the trials, such as imitation with a remotely controlled robotic doll (Robins et al., 2004) or chasing games with a rectangular robotic platform (Werry and Dautenhahn, 1999), while, here, we address a large range of games with the autonomous robotic pet. Moreover, trials in robot-assisted play have, for a long time, kept the experimenter physically apart or not involved in the situations of interaction. Robins et al. started to introduce the role of the experimenter, qualifying her role as the one of a 'passive participant' (Robins and Dautenhahn, 2006). Our method goes beyond and proposes a precise definition of the role of the experimenter who can intervene and regulate the interaction under specific conditions that are detailed and formalized. A main goal here is to facilitate the access to higher levels of play and to reasoning related to the robot.

Results from a long-term experiment in school with six children with autism have been analysed with a specific qualitative method which enables to focus on three dimensions, Play, Reasoning and Affect. For each dimension, I proposed a methodology: i) I defined a Play Grid for the analysis of play situations; ii) For the analysis of reasoning about the robot I referred to four categories of the reasoning part of the coding manual developed by Kahn et al. (2003), namely "Essence", "Mental States", "Social Rapport" and "Moral Standing"; iii) I coded the 'Affect' dimension according to precise explicit criteria.

Results have shown that this method is capable to adapt to the children's specific needs and abilities since all the children progressed, and progressed differently, according to their needs, abilities and preferences. Moreover, with respect to play and more specifically solitary vs. social play, children could be categorized into three groups. The first group is constituted by children not playing or mostly engaged in dyadic play with the robot. The second group is constituted by those initially playing solitarily and communicating mostly non-verbally but progressively experiencing more complex situations of verbal play as well as few pre-social or basic social situations of play. The third group is constituted by the children who managed to play socially (i.e. play in a triad including both the robot and the experimenter). It was found that:

- Children from the first group tended to progressively experience longer periods of uninterrupted play with the robot and started engaging in basic imitation during the last sessions;
- Children from the third group and, at a more basic stage, those from the second group, tended to experience higher levels of play gradually over the sessions and constructed more and more reasoning related to the robot; they sometimes demonstrated specific reasoning on real life situations as well.

Last but not least, children from the second and third group tended to express verbally or physically some interest in the robot, including on occasion interest involving affect. Finally, it was globally found that this approach did encourage proactivity and initiative-taking.

Research Question 2: *How can a robot recognize the interaction styles of each child in real time?*

The real-time recognition of the interaction styles has been investigated with several techniques, firstly with the Self-Organizing Maps, which showed a good accuracy to classify strokes according to the gentleness. Attempts to reduce the delay led to substantial hand-tuning and I preferred a solution that would be more easily generalizable to other criteria of interaction. I therefore applied successively two other techniques. Firstly, the Linear Discriminant Analysis and secondly, the clustering by compression (Cilibrasi and Vitanyi, 2005) did not lead to a separation of the classes and therefore were not further pursued.

I designed a novel method for the real-time recognition of Human-Robot Interaction Styles, the Cascaded Information Bottleneck Method, which extends the existing Information Bottleneck Method (Tishby et al., 1999). It relies on a succession of bottlenecks, trained successively, with the same cardinality of bottleneck states. The first bottleneck is trained in the standard way (Tishby et al., 1999) while the next ones depend on the previous bottleneck states. This successive training of the bottlenecks notably favours a powerful exploitation of the temporal structure of the data. Further to this, I introduced a measure which evaluates the similarity between states from two successive bottlenecks in order to extrapolate events that have not been encountered during the training phase of the algorithm.

I have shown the soundness of this method through extensive testing, with both i) data generated under laboratory conditions (training data and cross-validation) during human-robot interactions with a physical robot and ii) samples from natural situations of child-robot interaction in a school for children with autism. The algorithm was able to recognize short term events very well within and average delay of 0.17 seconds (the highest delay being 2.07 seconds). It was also able to recognise mid-term time scale events very well (the percentage of events correctly classified was 92% under laboratory conditions and 93% with data from the child-robot interactions).

The method is entirely generic for applications with socially interactive (humanoid and non-humanoid) robots. The ability of a robot to classify in real time human-robot interaction styles is a first step towards the challenging goal of enabling an autonomous robot to influence positively children's interaction styles to guide them progressively towards different therapeutically relevant levels of interaction.

Research Question 3: *How could the robot best adapt to the children's needs and abilities? Can a robot that adapts to the play styles of the children in real time impact the behaviour of the children? Could it, in this way, help the children engage progressively in better balanced interactions?*

I investigated the role of an adaptive autonomous robot, which reacts differently according to the child's play styles (in comparison with a reactive autonomous robot). I designed a schema of adaptation which relies on a reward basis. The interaction styles are categorised in real time with the algorithm I developed, the Cascaded Information Bottleneck Method. The robot's behaviours have been tailored by immersion according to each child's specific needs and abilities (Appendix C).

I tested the impact of the robot in a study in school with seven children with autism, over two sessions. I analysed the results, firstly with nonparametric statistics, which showed that children engaged significantly more in the interaction and generated richer interaction when the robot was adaptive (in comparison with reactive). They interacted significantly more on higher frequencies (both very high frequencies represented by S_3 and well-balanced ones represented by S_2) in the adaptive mode. Besides, the binomial test was applied on the average relative ratio of strong strokes over the runs from the two sessions for respectively the adaptive and reactive modes. It showed that significantly more children played, on average over the two sessions, more gently with the adaptive robot. Furthermore, I conducted a detailed analysis for each child combining qualitative and quantitative analysis. Results have shown that, in terms of the gentleness of the strokes, the adaptive robot clearly (and importantly) impacted the play styles of several children who generated within a same session, a higher ratio of gentle strokes (in comparison with strong strokes) while the robot was adaptive. As for the frequency of the interaction, the detailed analysis per child showed that, for some children, the adaptive robot even triggered frequencies of interaction (e.g. S_1 and S_3) which did not happened in the reactive mode. The adaptive robot therefore positively impacted the children's play styles and, in particular, it was found that, for most of the children who belonged to the first and second groups as defined in the study with the method inspired by non-directive play (Section 4.6), this positive impact tended to be very important. It is a valuable result because, during the play sessions conducted with the novel approach described in Chapter 4, these children tended to naturally mostly engage in exploration games by stroking the robot. Some of them had experienced the game 'ask for a physical reaction, show it with a sensor', too, which implies interaction with both the robot and the experimenter. Nonetheless, those children did not use much of verbal communication which would facilitate them access to proper situations of symbolic play with both the robot and the experimenter. Therefore, for those children, at this stage, tactile play remains the main means by which they interact with the robot and, it is very important to have found that their play styles can be positively impacted by the adaptive robot. It means that this dyadic interaction may enable them to learn more balanced levels of interactions.

7.3 Conclusion

In this thesis, I have adopted a multidisciplinary approach to address the issue of facilitating play between children with autism and an autonomous robot. This led me firstly to develop a novel approach for the design of play sessions in robot-assisted play. This approach draws inspiration from non-directive play therapy and adds a regulation process that enables the experimenter to guide the child towards other play styles under specific conditions or ask questions on reasoning of affect related to the robot. The long-term study that I conducted showed that this method can adapt to the specific needs and abilities of the children and encourage them explore a diversity of play situations and, in particular, social play. Three groups were highlighted based on the capacity of the children to progressively play socially or not with both the robot and the experimenter. It was shown that for the children who play mostly dyadically with the robot, the interactions with the robot tended to last longer over the sessions and some situations of imitations happened, which constitute a very first step towards triadic interaction. For the children who played socially with both the robot and the experimenter, higher levels of play were progressively experienced as well as reasoning and possibly affect related to the robot. This preliminary long-term study has therefore shown promising results for this new approach in robot-assisted play. It is a first study that potentially may be developed towards a new method in autism therapy.

Secondly, I have tested different methods for the real-time recognition of human-robot interaction styles and proposed a new method. The first technique I tested, based on Self-Organizing Maps, showed capable of classifying the criterion 'gentle/strong'. However, in order to have a recognition made within a reasonable delay, important hand-tuning was required which made the solution very specific to that particular criterion and was time-consuming. Two other methods were then successively tested, firstly the Fisher Linear Discriminant Analysis and secondly Clustering by compression which did not enable a separation of the classes. Thus, those two methods were not further pursued and I developed a new method, the Cascaded Information Bottleneck Method. This method consists of a cascade of bottlenecks trained successively. I have shown that a structure over the cascade emerges and I have introduced a measure for extrapolating unseen events. This measure enables to control the degrees of freedom of the system and is a first way to prevent the system from over-learning. An additional way to control how much and what new information is taken at which step of the cascade would be to move back from the agglomerative setting to a finite β setting. This shows how the method is transparent and enables control over how much and what new information is taken at which step of the cascade. The method was evaluated with two criteria of interaction, the criterion 'gentle/strong' which corresponds to a short-term time scale event and the criterion 'frequency of the interaction' which corresponds to a mid-term time scale event, both with, successively, trained data and cross-validation. The algorithm showed sound for recognizing both of these criteria. The short-term time scale events were recognized with a very small delay and the method made a powerful exploitation of the existing temporal structure of mid-term time scale events. Note that the algorithm was also tested with data generated in school with children with autism, whereby they could play freely, i.e. they were not instructed how to play. This method is entirely generic to applications with socially interactive robots and is a step towards socially adaptive robots.

Thirdly, I investigated the role of the adaptive robot in robot-assisted play in a study with children with autism in school. I designed a schema of adaptation based on a reward which uses the Cascaded Information Bottleneck Method for the recognition of the interaction styles in real time. The study showed the positive impact of the adaptive robot on the children's play styles. The adaptive mode encouraged the children to engage significantly more in tactile interaction by activating the sensors more often within a session. Besides, on average over the two sessions, the proportion of gentle strokes increased when the robot was in the adaptive mode (except for one child). It was moreover found that for several children, this tendency was very clear within each session. Furthermore, the analysis of the criterion 'frequency of the interaction' showed that the interactions were significantly richer in the adaptive mode. In particular, higher frequencies were significantly more present in the adaptive mode, including notably frequencies that we qualified as well-balanced. For some children, the adaptive robot even triggered some frequencies that were absent from the reactive runs within a same session. This study is a step forward in the investigation of the potential role that adaptive robots can play in robot-assisted play for children with autism.

To summarize, this thesis contributes to a wide range of areas:

- **Robot Assisted Play**: I proposed and experimentally tested a new methodological approach of how to design, conduct and analyse robot-assisted play.
- **Machine Learning**: I proposed and experimentally tested a novel and generic computational method.
- **Human-Robot Interaction**: I demonstrated a proof-of-concept system of an adaptive robot responsive to different styles of interactions in human-robot interactions and tested its impact through a study with children with autism.
- **Developmental Robotics**: I contributed to the understanding of social behaviour and adaptation which are key topics in developmental robotics, inspired by research on child development and autism therapy.
- **Autism Therapy**: I conducted a study that potentially may be developed towards a new method in autism therapy.

Chapter 8

Future Work

In this chapter we draw some directions for future work.

Application of the Cascaded Information Bottleneck Method to other criteria of interaction: The real-time recognition of human-robot interaction styles is a step towards socially adaptive robots (Dautenhahn, 2007b, 1998). The Cascaded Information Bottleneck Method is entirely generic for applications with socially interactive robots. A further step could be to generalize this method to the recognition of additional criteria of interaction. Note that this method could also be tested and used in different contexts of pattern recognition than HRI, since the method is a powerful time filtering process that progressively extracts information from time series and makes a good exploitation of the temporal structure of the data with transparency and the possibility to control how much and what new information is taken at which step of the cascade.

Comparison of the Cascaded Information Bottleneck Method with other methods: Future work could include, in these scenarios of Human Robot interaction, a comparison of the Cascaded Information Bottleneck Method with other methods used for pattern recognition such as HMMs. My hypothesis is that classical homogeneous HMMs as used in e.g. Lee and Xu (1996) and Calinon and Billard (2004) might have difficulties to model an existing mid-term temporal structure of the data by trying to squeeze all temporal information into one flat transition structure[1]. On the contrary, the Cascaded Information Bottleneck Method relies on different

[1]In order to get more insight on how the HMMs would be used in this specific context, please refer to the detailed description provided in Section 5.1.3

bottlenecks trained successively (i.e. different mappings over a time series), thus enabling a powerful exploitation of the temporal structure of the data. The problem with a cascade of bottlenecks trained successively could be here that the system has too many degrees of freedom and could overlearn. The extrapolation with the measure that I have introduced is a first step in the control of the degrees of freedom of the system. In addition, the overlearning can be tightly controlled by penalizing the intake of novel information. For this, we would have to move from the agglomerative model (where β goes to ∞) to a model with a finite β that would control the information intake per step. This shows how the Cascaded Information Bottleneck method is transparent and gives more control over how much and what new information is taken at which step in the cascade. Since the method is so transparent and easy to control, there could be even further enhancements and improvements that use these properties.

A complementary long-term study with the adaptive robot: It would be enlightening to conduct a long-term study with the adaptive robot, complementary to the short-term study presented in Chapter 6 that would analyse on a long-term basis the impact of the adaptive robot, compared to the reactive robot as follows:

i) How do the children's play styles change over many sessions with the adaptive robot (in comparison with the reactive robot)?

ii) Does the impact of the adaptive robot change over time (in comparison with the reactive robot)?

During this long-term study, one might also wish to adapt slightly the reward schema of the robot for Child H who might have been a bit confused by the engaging behaviour of the robot[2]. This could result, for this specific child, in modifying the engaging behaviour of the robot, or even removing the engaging behaviour and only focusing on the robot's rewards for gentle strokes and good frequencies of interaction.

The impact of the familiarity with the robot on the children's play styles: An additional study could involve children who would have met the robot only during a few (several) sessions so that they have not had a chance to become too familiar with the robot yet. However, the number of sessions should reach a minimum threshold

[2]We should remind here that Child H usually interacts a lot with the robot. However, when the robot started the engaging behaviour, i.e. wagging the tail, Child H tended to rather look at the robot than stroke it. For this specific child, the engaging behaviour of the robot may thus have had the reverse effect, that is disengaging the child from the tactile interaction with the robot.

in order to enable the experimenter to get basic clues to tailor the robot's behaviours according to the children's specific needs and abilities (Appendix C). It would be interesting to see whether the impact of the robot on those children would be different to the one on the children of this present study who all participated, beforehand, in a long-term study with the method inspired by non-directive play and, through it, became familiar with the robot and experimented with play skills with both the robot and the experimenter. Note that between the long-term study with the method inspired by non-directive play and the trials investigating the impact of the adaptive robot on the children's play styles, those children had several play sessions during which the experimenter progressively decreased her participation in the games, in order to ensure a progressive transition with the study on the adaptive robot whereby the experimenter did not take part in the experiments[3].

Investigating the social potential of the robot: A further step both in robot-assisted play and towards socially adaptive robots would be the implementation of the different discrete levels of the social potential of the robot as described in Section 6.1.1 and its testing in the context of robot-assisted play: in this context the robot would be able to select its level of adaptation according to the child's progresses, needs and current abilities, by following the transposed principle of "Alternate Freezing and Freeing of Degrees of Freedom", as illustrated in Fig. 6.1. The robot would therefore adapt to the interaction styles but also select its level of adaptation according to the child's profile, needs, progresses and abilities. In this sense, one may imagine that, following the example of discrete social potential given in Section 6.1.1, the robot first adapts to the criterion Gentleness, and, when the child has shown capable to stroke the robot gently many times, it moves up to a higher level where it will classify both the gentleness and the frequency of interaction. But, if this level appears, at some point, to be too complicated for the child, the robot might go back to the easier level where it only triggers the Gentleness of the interaction.
In this context, one might also think about extending the social potential of the robot by adding more criteria of interaction for the recognition of the children's play styles. In Section 6.1.1, I suggested one possibility for the discrete social development of the robot that would follow three levels and take into account at most two criteria, the gentleness and the frequency of the interaction. One could actually imagine to

[3]The experimenter did not take part in these trials, i.e. she only responded to the children's questions. This was made in order not to interfere with the main purpose of this study which was to test the impact of the adaptive robot. In future work, the adaptive robot could be introduced in play sessions where the experimenter takes part in the play sessions.

extend the discrete social development of the robot to more levels that would include other criteria of interaction, such as, for instance, an additional classification of the frequency of the interaction based on a smaller range of frequencies with a 5 to 7 seconds analysis, and, the topological diversity of the interaction, that is the variety in terms of which sensors are activated.

Robot-assisted play at home and with play therapists: Those two approaches, on the one hand, the method inspired by non-directive play therapy, and, on the other hand, the adaptive robot, could, in future, be used by play therapists with children with autism. At the moment, a roboticist is needed to deal with all the technical issues addressed by the use of a robot. But the goal is to enable play therapists to apply these methods, in future, as a complementary approach to existing therapies in autism. Moreover, ideally, in future, children would be able to play in schools or even at home with the adaptive robot, which might be additionally equipped with a discrete social potential (if it appears to be positive for the children). Of course, such play sessions would necessarily be supervised by an adult, in order to check safety issues. Ideally, the child would meet a play therapist regularly, who could (re)adjust the robot's behaviours and its social potential for the following days, according to the child's needs, abilities, progresses and preferences.

This thesis has focused on facilitating play between children with autism and an autonomous robot and has addressed the issue with a multidisciplinary approach which led to a number of novel results and contributions. It has firstly enabled the design of a new methodological approach in robot-assisted play that was experimentally tested in school. Secondly, it led to the development of a novel and generic computational method for the real-time recognition of the interaction styles. Thirdly, it demonstrated a proof-of-concept system of an adaptive robot responsive to different styles of interaction in human-robot interaction. A study evaluated its impact on the play styles of children with autism. Taken together, I hope that these achievements represent a step forward in socially adaptive robots and in robot-assisted play for children with autism.

I conducted play sessions in school for more than a year with children with autism. Personally, those play sessions with the children were a wonderful and unforgettable experience. I would like to thank those children and I dedicate this thesis to them.

Appendix A

Short-term Study: Figures

Child	Average percentage of activations on the adaptive mode	Average percentage of activations on the reactive mode
Child A	72.22	27.78
Child G	87.84	12.16
Child H	44.74	55.26
Child C	82.01	17.99
Child E	57	43
Child F	75.56	24.44
Child D	90.48	9.52

Figure A.1: Table providing the average relative engagement of the children in adaptive and reactive modes.

Child	**Setting A-R-A**			**Setting R-A-R**		
	A	**R**	**A**	**R**	**A**	**R**
Child A	19.05	75.00	12.50	50.00	30.00	90.91
Child H	0	18.18		0	11.11	
Child F	6.38	87.50	0	30.00	46.15	85.71

Figure A.2: Table providing the percentage of strong strokes among all the strokes (the strokes can be gentle or strong) for the children A, H and F. A void cell means that, for the run corresponding to the cell, the child did not activate sensors.

Child	**Setting R-A-R**			**Setting A-R-A**		
	R	**A**	**R**	**A**	**R**	**A**
Child G	25.00	0		0	0	6.25
Child C	0	0	35.25	1.96	0	8.62
Child E	15.00	7.14	16.98	16.67	15.38	21.88
Child D	0	75.00	0	46.15	100	

Figure A.3: Table providing the percentage of strong strokes among all the strokes (the strokes can be gentle or strong) for the children G, C, E and D. A void cell means that, for the run corresponding to the cell, the child did not activate sensors.

Child	**Av(δ_{S0})**	**Av(δ_{S1})**	**Av(δ_{S2})**	**Av(δ_{S3})**
Child A	0.217	0.202	0.508	0.750
Child G	-0.233	0.750	0.284	0.614
Child H	0.133	-0.450	0.222	0.151
Child C	0.288	0.500	0.451	0.058
Child E	0.081	0.061	-0.017	0.054
Child F	0.164	0	0.108	0.170
Child D	0.444	0.250	0.516	0.750

Figure A.4: Table providing the average of δSi over the two sessions per child.

Child	$Av_{Adaptive}(\rho)$	$Av_{Reactive}(\rho)$
Child A	0.520	0.147
Child G	0.491	0.176
Child H	0.340	0.327
Child C	0.477	0.189
Child E	0.353	0.313
Child F	0.382	0.284
Child D	0.523	0.060

Figure A.5: Input Data for the Wilcoxon test to measure the richness of the interaction.

Child	$Av_{Adaptive}(\rho_{S2})$	$Av_{Reactive}(\rho_{S2})$
Child A	0.559	0.108
Child G	0.459	0.207
Child H	0.432	0.235
Child C	0.534	0.133
Child E	0.326	0.341
Child F	0.381	0.285
Child D	0.563	0.104

Figure A.6: InputData for the Wilcoxon test applied to S_2.

Child	$Av_{Adaptive}(\rho_{S3})$	$Av_{Reactive}(\rho_{S3})$
Child A	0.667	0
Child G	0.606	0.060
Child H	0.400	0.266
Child C	0.359	0.308
Child E	0.357	0.309
Child F	0.409	0.258
Child D	0.667	0

Figure A.7: InputData for the Wilcoxon test applied to S_3.

Child	Setting	Mode	N(S0)	N(S1)	N(S2)	N(S3)
Child A	A-R-A	A	193	15	178	198
		R	35	0	0	0
		A	226	21	177	121
	R-A-R	R	95	28	1	0
		A	90	24	21	125
		R	61	37	9	0
Child G	R-A-R	R	201	0	84	93
		A	33	1	97	735
		R	0	0	0	0
	A-R-A	A	75	0	412	226
		R	106	0	142	42
		A	54	22	346	345
Child H	A-R-A	A	153	0	85	75
		R	79	55	51	77
		A	0	0	0	0
	R-A-R	R	160	21	46	19
		A	172	14	94	46
		R	0	0	0	0
Child C	R-A-R	R	30	0	0	0
		A	78	16	94	45
		R	60	8	56	538
	A-R-A	A	138	0	69	204
		R	28	0	9	0
		A	181	17	273	139
Child E	R-A-R	R	59	33	391	416
		A	98	6	449	808
		R	72	35	436	753
	A-R-A	A	125	16	367	496
		R	84	0	557	460
		A	53	11	566	410
Child F	A-R-A	A	120	15	500	850
		R	24	0	0	0
		A	47	15	611	1041
	R-A-R	R	167	0	94	274
		A	118	0	19	86
		R	58	6	19	19
Child D	R-A-R	R	8	0	0	0
		A	44	0	28	182
		R	0	0	11	0
	A-R-A	A	91	20	64	22
		R	31	0	2	0
		A	0	0	0	0

Figure A.8: Table providing the number of occurences for each class Si, $N(Si)$.

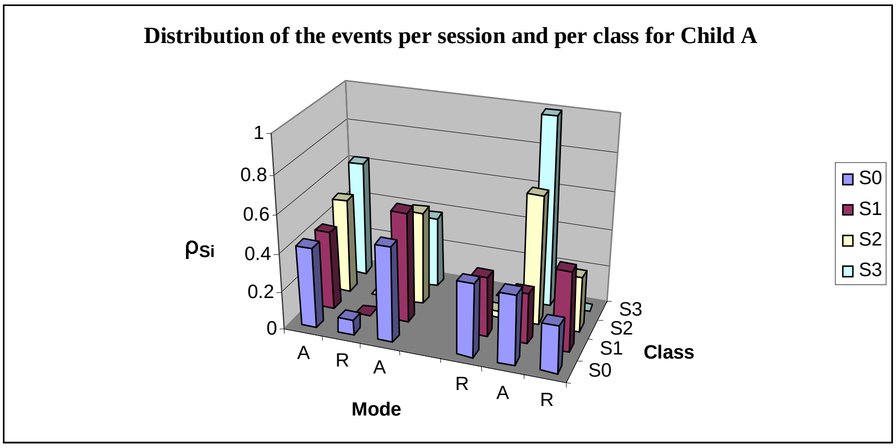

Figure A.9: Distribution of the events per session and per class for Child A.

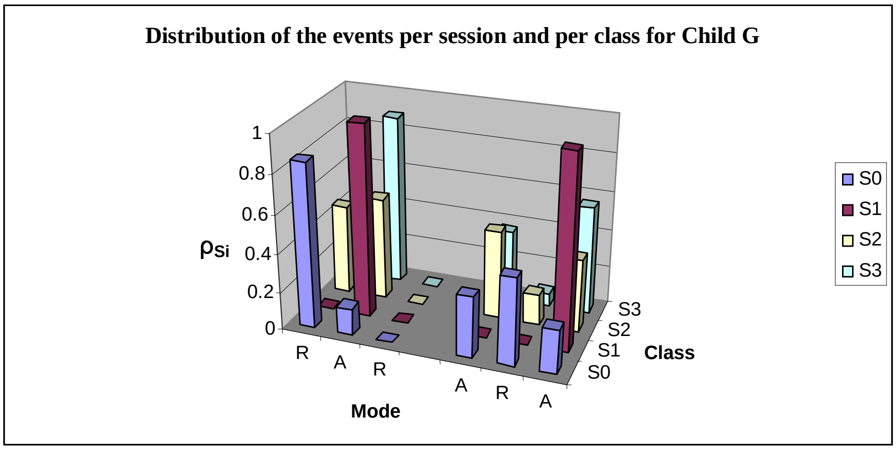

Figure A.10: Distribution of the events per session and per class for Child G.

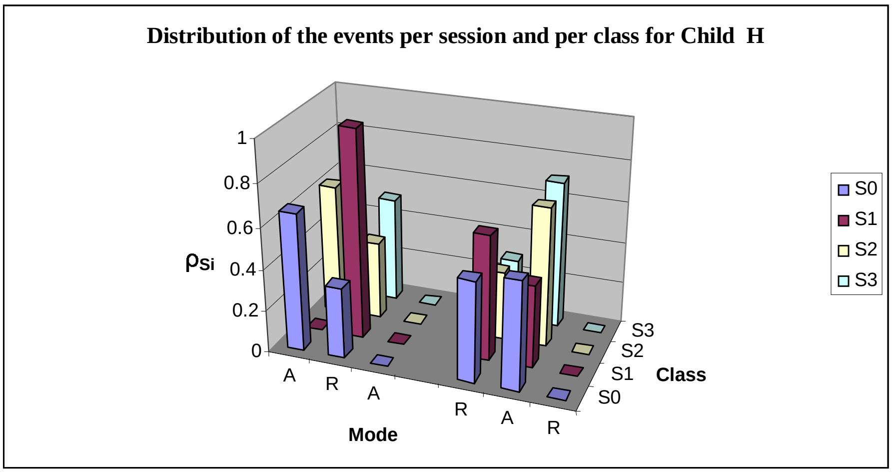

Figure A.11: Distribution of the events per session and per class for Child H.

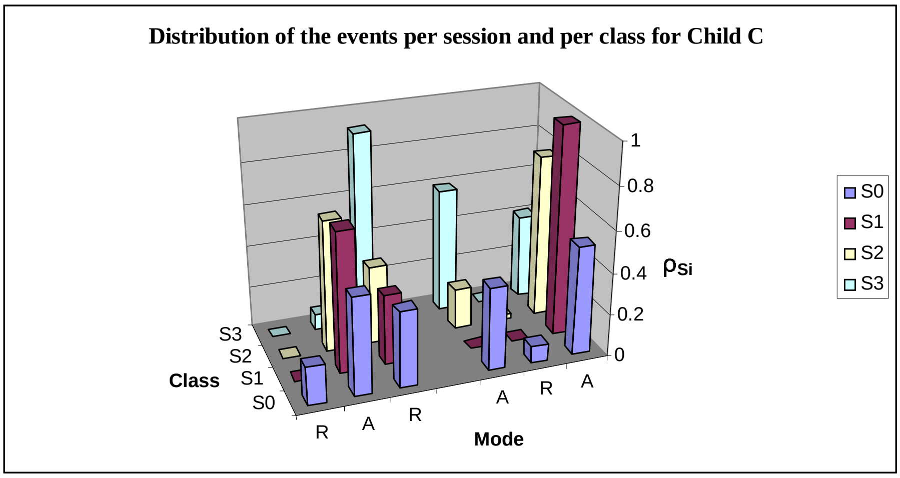

Figure A.12: Distribution of the events per session and per class for Child C.

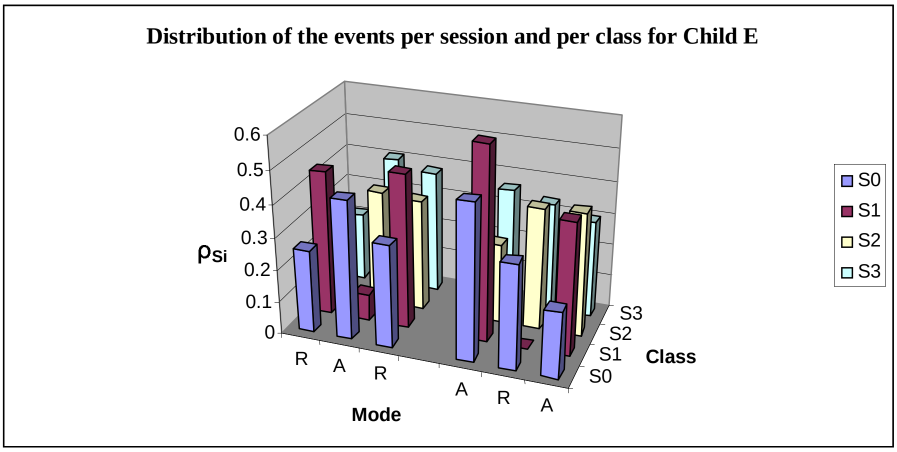

Figure A.13: Distribution of the events per session and per class for Child E.

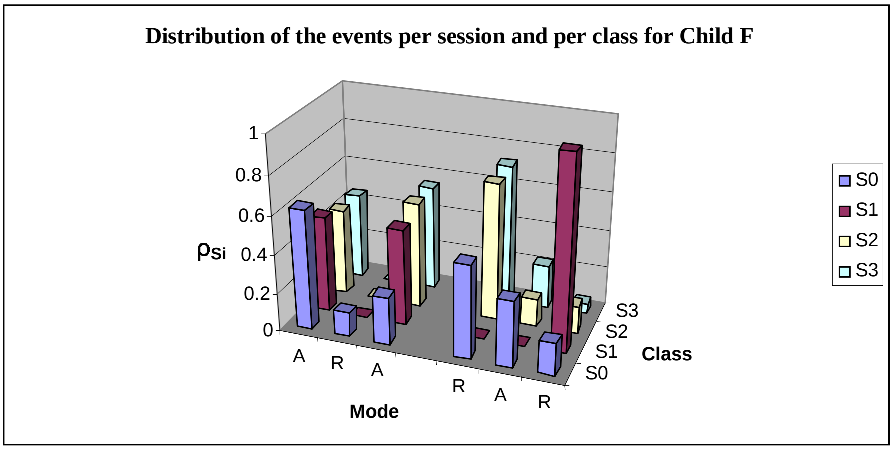

Figure A.14: Distribution of the events per session and per class for Child F.

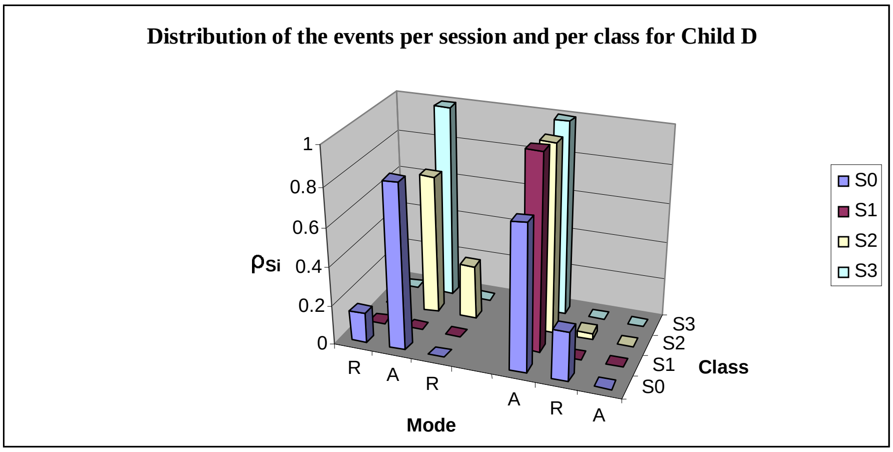

Figure A.15: Distribution of the events per session and per class for Child D.

Child	Setting	Order	δS0	δS1	δS2	δS3
Child A	A-R-A	A-R	0.348	0.417	0.501	0.621
		R-A	0.421	0.583	0.499	0.379
	R-A-R	R-A	-0.020	-0.045	0.645	1.000
		A-R	0.118	-0.146	0.387	1.000
Child G	R-A-R	R-A	-0.718	1.000	0.072	0.775
		A-R	0.141	1.000	0.536	0.888
	A-R-A	A-R	-0.132	0	0.300	0.300
		R-A	-0.221	1.000	0.227	0.494
Child H	A-R-A	A-R	0.319	-1.000	0.250	-0.013
		R-A	-0.341	-1.000	-0.375	-0.507
	R-A-R	R-A	0.036	-0.200	0.343	0.415
		A-R	0.518	0.400	0.671	0.708
Child C	R-A-R	R-A	0.286	0.666	0.626	0.077
		A-R	0.107	0.333	0.253	-0.846
	A-R-A	A-R	0.317	0	0.170	0.595
		R-A	0.441	1.000	0.752	0.405
Child E	R-A-R	R-A	0.170	-0.365	0.045	0.198
		A-R	0.114	-0.392	0.010	0.028
	A-R-A	A-R	0.156	0.593	-0.128	0.026
		R-A	-0.118	0.407	0.006	-0.037
Child F	A-R-A	A-R	0.503	0.500	0.450	0.449
		R-A	0.120	0.500	0.550	0.551
	R-A-R	R-A	-0.143	0	-0.568	-0.496
		A-R	0.175	-1.000	0	0.177
Child D	R-A-R	R-A	0.692		0.718	1.000
		A-R	0.846		0.436	1.000
	A-R-A	A-R	0.492	1.000	0.939	1.000
		R-A	-0.254	0	-0.030	0

Figure A.16: Table providing the δ_{Si} for each session for each child.

Appendix B

Children's age and level of autism

In the school where the play sessions were conducted, the level of autism of the children was evaluated with the Childhood Autism Rating Scale (CARS) (Schopler et al., 1980). It is one of the most widely used standardised instruments specifically designed to aid in the diagnosis of autism. It can be used for children from two-years old and more. The test is organised in 15 areas which are detailed in Fig. B.1.

1. Relating to people
2. Imitation
3. Emotional response
4. Body use
5. Object use
6. Adaptation to change
7. Visual response
8. Listening response
9. Taste, smell, and touch response and use
10. Fear and nervousness
11. Verbal communication
12. Nonverbal communication
13. Activity level
14. Level and consistency of intellectual response
15. General impressions

Figure B.1: The fifteen areas in the Childhood Autism Rating Scale.

For each area, the child gets a rate from 1 (which means 'normal for the child's age') to 4 (which means 'severely abnormal for the child's age'). The individual rates from these 15 areas are then summed up. The total indicates whether the child has autism or not, and whether the autism is mild-to-moderate or severe. A score greater or equal to 30 means that the child has autism. If the score is greater or equal to 37, then the child has severe autism.

Fig. B.2 provides, for each child involved in the experiments, the CARS score.

Child	**Chronological age at the beginning of the trials**	**CARS**
Child A	7 years old	51
Child B	8 years old	48
Child C	7 years old	35
Child D	10 years old	42
Child E	10 years old	35
Child F	9 years old	38
Child G	5 years old	41
Child H	5 years old	45
Child I	7 years old	NA

Figure B.2: Children's profile of autism according to the Childhood Autism Rating Scale (CARS).

Appendix C

Tailoring the robot's behaviours

The robot's behaviours have been tailored, for each child, by immersion. This means that the repertoire of appropriate robot's behaviours with respect to each child specific needs, abilities, dislikes and preferences was progressively designed and refined as the experiments progressed. During these trials, several sources of stimulation were successively tested: sound (e.g. barking) and movement (e.g. head turning, walking). The speed of the movements was also progressively tuned to fit each child's specific needs and preferences. The idea was to start very simply and, progressively, add some diversity and complexity in the robot's behaviours[1], in order to identify the reaction of the child and the possible interpretation he/she gave to a specific robot's behaviour.

If the child liked the behaviour, then the behaviour was adopted for the robot. For example, the first time the robot's walking was enabled, *Child I* stroked the robot, the robot walked and *Child I* suddenly laughed and smiled. During the whole session, whenever the robot's walk happened, she smiled and laughed again. *Child I* did not communicate verbally during the play sessions. Here, her laughing expressed a positive reaction to the robot's behaviour.

In contrast, if it was felt that a child showed some hesitation in front of a new behaviour of the robot, then this behaviour was removed. For instance, the robot's

[1]Firstly, 'robot's barking' as well as 'slowly moving head' and 'wagging tail' were tested. Then, the robot's walking was introduced. Further to this, the range of robot's gestures was expanded, as well as its walking, which could be forward and backwards. This phase notably included the testing of some behaviours that would last a bit longer (approximately 2-3 seconds) than the ones developed in further stages (0-2 seconds). Because the children were very interested in robot's emitting sounds, and because one child even asked for it, a new sound was finally introduced, in addition to the barking. In addition to the nature of the behaviour, different mappings were also tested, in order to i) make the child experiment with changes in the robot's reactions, and ii) adapt to the children's play styles, preferences and demands.

walking was removed for Child E.

Further to this, if a child was asking for a specific behaviour (among the range of realistic behaviours in this context of play), then the experimenter would update the robot's behaviour immediately to make the specific reaction happen[2]. A first example is about Child F liking a specific sound that the robot could play (this sound sounded like a drumming sound). When this behaviour arose the first time, Child F liked it very much and called it 'electric stroke' (when he was stroking the robot on one of the back sensors, this sound arose). In the following sessions, he always asked to have this behaviour on the robot. A second example is about correlated behaviours of the robot: during an advanced session, Child F asked why the robot was not opening the mouth when barking. The experimenter asked him whether he would like the robot to do that, and he said yes. Therefore, the robot was immediately programmed to bark and open the mouth at the same time. In later stages of this study, Child F had encountered different mappings[3] for the robot's behaviours. Depending on the play situation he was involved with (e.g. give food to the robot, make the robot 'walk in the air', explore the robot's features by stroking him, etc.) he sometimes asked for a specific[4] mapping.

---

[2] The behaviours were programmed with URBI (Baillie, 2005).

[3] By 'mapping' we mean here a one to one deterministic mapping between the robot's external sensors and its behaviour.

[4] In particular, when he gave food to the robot, he wanted the barking to be removed when the robot opened the mouth. But once he was finished, he wanted to have it back.

Appendix D

Social Story

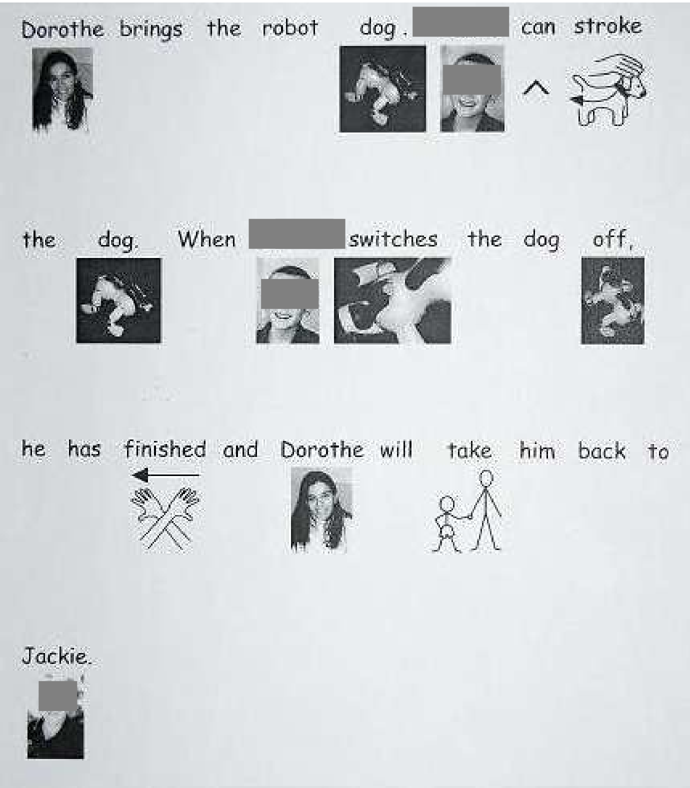

Figure D.1: Social Story used for Child A. In order to help Child A understand how the play sessions proceeded, a social story was made by the teachers of the autism base. The name of the child has been erased and both faces of the child and the teacher have been hidden.

Appendix E

Publication List

Several publications resulted from this research:

Journal Paper (to appear):

François, D., Powell, S., and Dautenhahn, K. (2009). A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative taking. *To appear in: Interaction Studies. Special Issue: Robots in the Wild: Exploring Human-Robot Interactions in Naturalistic Environments.*

Conference Papers:

François, D., Polani, D., and Dautenhahn, K. (2007). On-line behaviour classification and adaptation to human-robot interaction styles. *In Proc. 2nd ACM/IEEE International Conference on Human-robot Interaction (HRI 07)*, pages 295-302.

François, D., Polani, D., and Dautenhahn, K. (2008b). Towards socially adaptive robots: A novel method for real time recognition of human-robot interaction styles. *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids 08)*, pages 353–359.

Abstract for a talk:

François, D., Dautenhahn, K., and Polani, D. (2008). Robot Assisted Play: Detecting Interaction Styles of Children with Autism Playing with a Zoomorphic Robot. Abstract for talk to be given on December 1st 2008, in Conventry University Technocentre, at the Conference RAatE 2008 (Recent Advances in Assistive Technology and Engineering).

Technical Reports:

François, D., Polani, D., and Dautenhahn, K. (2008a). Real time recognition of human-robot interaction styles with cascaded information bottlenecks. Technical report 478, School of Computer Science, Faculty of Engineering and Information Sciences, University of Hertfordshire.

François, D., Powell, S., and Dautenhahn, K. (2008c). A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative taking. Technical report 477, School of Computer Science, Faculty of Engineering and Information Sciences, University of Hertfordshire.

Appendix F

Media

"Roboterhund hilft bei Autismus": Documentary broadcasted on 25^{th} August 2008 on the German Channel 3SAT (http://www.3sat.de/). Rebroadcasts on 3Sat and the partner channels MDR, SF, RBB and BRalpha.

This five minute documentary reported on my research. The filming took place both in the school (where the journalists filmed several play sessions with the children) and in the Science and Technology Research Institute, at the University of Hertfordshire.

Bibliography

Association, A. P. (1994). *Diagnostic and Statistical Manual of Mental Disorders DSM-IV, fourth Edition*. American Psychiatric Publishing.

Aurora (2008). Aurora project. http://www.aurora-project.com/.

Axline, V. (1946). *Dibs: in search for Self*. New York: Ballantine Books.

Axline, V. (1947). *Play therapy*. New York: Ballantine Books.

Bailey, M., Chanler, A., Maxwell, B., Micire, M., and Yanco, H. (2007). Development of stereo vision-based navigation for a robotic wheelchair. In *Proceedings of the International Conference on Rehabilitation Robotics (ICORR 2007)*, pages 951–957.

Baillie, J.-C. (2005). URBI: Towards a universal robotic low-level programming language. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 820–825.

Baron-Cohen, S. (1997). *Mindblindness, an essay on Autism and Theory of Mind*. The MIT Press, new edition.

Baron-Cohen, S. (2001). Theory of mind in normal development and autism. *Prisme*, 34:174–183.

Berthouze, L. and Lungarella, M. (2004). Motor skills acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behaviour*, 12(1):47 – 64.

Bobath protocol (2008). Bobath protocol. http://www.bobath.org.uk/.

Borodovsky, M. and McIninch, J. (1993). GENMARK: parallel gene recognition for both DNA strands. *Computers and Chemistry*, 17(2):123–133.

Boucher, J. (1999). Editorial: Interventions with children with autism - methods based on play. *Child language teaching and therapy*, 15:1–5.

Boucher, J. and Wolfberg, P. J. (2003). Editorial: Aims and design of the special issue. *Autism*, 7:339–346.

Breazeal, C. (2002). *Designing Sociable Robots*. MIT Press.

Breazeal, C. (2003). Towards sociable robots. *Robotics and Autonomous Systems*, 42(3-4):167–175.

Bruner, J. (1983). *Child's Talk: Learning to Use Language*. New York: Norton.

Bruner, J. (1990). *Acts of Meaning*. Cambridge, MA: Harvard University Press.

Bruner, J. (1996). *The Culture of Education*. Cambridge: Harvard University Press.

Burgar, C. G., Lum, P. S., Shor, P. C., and Van der Loos, H. F. M. (2002). Development of robots for rehabilitation therapy: The palo alto VA/Stanford experience. *Journal of Rehabilitation Research and Development*, 37(6):663–673.

Calinon, S. and Billard, A. (2004). Stochastic Gesture Production and Recognition Model for a humanoid robot. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems(IROS)*, pages 2769–2774.

Chaillé, C. and Silvern, S. B. (1996). *Understanding Through Play*. The British Library - "The World's Knowledge".

Cilibrasi, R. and Vitanyi, P. M. B. (2005). Clustering by compression. *IEEE Transactions on Information Theory*, 51(4):1523–1545.

Cogniron (2008). Cogniron project. http://www.cogniron.org/.

Colman, A. M. (2001). *Dictionary of Psychology*. Oxford University Press.

Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley-Interscience, New York.

Crutchfield, J. P. (1990). *Information and its Metric*. Nonlinear structures in physical systems- Pattern Formation, Chaos and Waves, L.Lam and H. C. Morris Eds. Springer Verlag, pp 119-130.

Dautenhahn, K. (1998). The art of designing socially intelligent agents: Science, fiction, and the human in the loop. *Applied Artificial Intelligence*, 12(8-9):573–617.

Dautenhahn, K. (2007a). Methodology and themes of human-robot interaction: A growing research field. *International Journal of Advanced Robotics Systems*, 4(1):103–108.

Dautenhahn, K. (2007b). Socially intelligent robots: Dimensions of human-robot interaction. *Philosophical Transactions of the Royal Society B*, 362(1480):679–704.

Dautenhahn, K. and Werry, I. (2000). Issues of robot-human interaction dynamics in the rehabilitation of children with autism. In *Proc. From Animals to Animats, 6th Int. Conference on the Simulation on Adaptive Behavior (SAB2000)*, pages 11–15.

Dautenhahn, K. and Werry, I. (2002). A quantitative technique for analysing robot-human interactions. In *Proc. IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, pages 1132–1138.

Dautenhahn, K. and Werry, I. (2004). Towards interactive robots in autism therapy: Background, motivation and challenges. *Pragmatics and Cognition*, 12(1):1–35.

Dautenhahn, K., Werry, I., Rae, J., Dickerson, P., Stribling, P., and Ogden, B. (2002). *Robotic Playmates: Analysing Interactive Competencies of Children with Autism Playing with a Mobile Robot.* K. Dautenhahn and A. Bond and L. Cañamero and B. Edmonds (Eds), Socially Intelligent Agents - Creating Relationships with Computers and Robots. Kluwer Academic Publishers, 117-124.

Davis, M., Robins, B., Dautenhahn, K., Nehaniv, C., and Powell, S. (2005). A comparison of interactive and robotic systems in therapy and education for children with autism. In *Assistive Technology: From Virtuality to Reality AAATE 2005*, pages 353–357.

Derakhshan, A., Hammer, F., Demazeau, Y., and Lund, H. H. (2006). Mapping children and playgrounds into multi-agent systems. In *Proc. 11th Int. Symposium on Artificial Life and Robotics (ISAROB)*.

Elizondo, D. (2006). The linear separability problem: some testing methods. *IEEE Transactions on Neural Networks*, 17(2):330–344.

Eriksson, J., Matarić, M., and Winstein, C. J. (2005). Hands-off assistive robotics for post-stroke arm rehabilitation. In *Proc. Int. Conference on Rehabilitation Robotics*, pages 21–24.

Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3-4):143–166.

François, D., Polani, D., and Dautenhahn, K. (2007). On-line behaviour classification and adaptation to human-robot interaction styles. In *Proc. 2nd ACM/IEEE International Conference on Human-robot Interaction (HRI07)*, pages 295–302.

François, D., Polani, D., and Dautenhahn, K. (2008a). Real time recognition of human-robot interaction styles with cascaded information bottlenecks. Technical report 478, School of Computer Science, Faculty of Engineering and Information Sciences, University of Hertfordshire.

François, D., Polani, D., and Dautenhahn, K. (2008b). Towards socially adaptive robots: A novel method for real time recognition of human-robot interaction styles. In *Proc. IEEE-RAS International Conference on Humanoid Robots (Humanoids08)*, pages 353–359.

François, D., Powell, S., and Dautenhahn, K. (2008c). A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative taking. Technical report 477, School of Computer Science, Faculty of Engineering and Information Sciences, University of Hertfordshire.

François, D., Powell, S., and Dautenhahn, K. (2009). A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative taking. *To appear in: Interaction Studies. Special Issue: 'Robots in the Wild: Exploring Human-Robot Interactions in Naturalistic Environments'.*

Fritz, U. (1989). *Autism: Explaining the Enigma.* Oxford, Blackwell.

Gillingham, G. (1995). *Autism Handle with care.* TX. future Education Inc.

Gockley, R. and Matarić, M. (2006). Encouraging physical therapy compliance with a hand-off mobile robot. In *Proc. Int. Conference on Human-Robot Interaction*, pages 150–155.

Gómez, G., Lungarella, M., Hotz, P. E., Matsushita, K., and Pfeifer, R. (2004). Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity. In *Proc. 4th Int. Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 119–122.

Goodrich, M. A. and Schultz, A. C. (2007). Human-robot interaction: A survey. *Foundations and Trends in Human-Computer Interaction*, 1(3):203–275.

Grandin, T. (1986). *Emergence: Labeled autistic*. Warner Books.

Hagan, K., Hagan, S., Hillman, M., and Jepson, J. (1997). Design of a wheelchair mounted robot. In *Proceedings of the International Conference on Rehabilitation Robotics (ICORR'97)*, pages 27–30.

Haykin, S. (1998). *Neural Networks: A Comprehensive Foundation.* Prentics Hall; 2nd Edition.

Hillman, M. (2003). Rehabilitation Robotics from past to present - a historical perspective. In *Proceedings of the 8th International Conference on Rehabilitation Robotics (ICORR).*

Hillman, M., Hagan, K., Hagan, S., Jepson, J., and Opwood, R. (2001). The Weston wheelchair mounted assistive robot - the design story. *Robotica*, 20:125–132.

Hillman, M., Hagan, K., Hagan, S., Jepson, J., and Orpwood, R. (1999). A wheelchair mounted assistive robot. In *Proceedings of the International Conference on Rehabilitation Robotics (ICORR'99)*, pages 86–91.

Hobson, P. (2002). *The Cradle of Thought.* Oxford University Press.

Howlin, P. and Rutter, M. (1987). *Treatment of Autistic Children.* John Willey and Sons.

Isaacs, S. (1930). *Intellectual growth in young children.* London: Routledge.

Isaacs, S. (1933). *Social Development in young children: A study of beginnings.* London: Routledge.

Jordan, R. (1999). *Autistic Spectrum Disorders, An Introductory Handbook for Practitioners.* David Fulton publishers Ltd.

Josefi, O. and Ryan, V. (2004). Non-directive play therapy for young children with autism: A case study. *Clinical child Psychology and Psychiatry*, 9(4):533–551.

Kahn, L. E., Averbuch, M., Rymer, W. Z., and Reinkensmeyer, D. J. (2001). Comparison of Robot-Assisted Reaching in Promoting Recovery from Chronic Stroke. In *Proc. Int. Conference on Rehabilitation Robotics*, pages 39–44.

Kahn, P. H., Friedman, B., Freier, N. G., and Severson, R. (2003). Coding manual for children's interactions with AIBO, the robotic dog - The preschool study. Uw cse technical report 03604603, Department of Computer Science and Engineering, University of Washington.

Kahn, P. H., Friedman, B., Pérez-Granados, D. R., and Freier, N. G. (2006). Robotic pets in the lives of preschool children. *Interaction Studies*, 7(3):405–436.

Kanda, T., Hirano, T., Eaton, D., and Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: a field trial. *Journal of Human Computer Interaction*, 19(1-2):61–84.

Kanda, T., Ishiguro, H., Ono, T., Imai, M., and Nakatsu, R. (2002). Development and evaluation of an interactive humanoid robot "Robovie". In *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA2002)*, pages 1848–1855.

Kanda, T. and Ishiguro, I. (2005). Communication robots for elementary schools. In *Workshop on Robot Companion at AISB 2005*.

Kanda, T., Sato, R., Saiwaki, N., and Ishiguro, H. (2007). A two-month field trial in an elementary school for long-term human-robot interaction. *IEEE Transactions on Robotics (Special issue on Human-Robot Interaction)*, 23(5):962–971.

Kim, D., Song, J., and Kim, D. (2007). Simultaneous gesture segmentation and recognition based on forward spotting accumulative HMMs. *Pattern Recognition*, 40(11):3012–3026.

Knox, V. and Evans, A. L. (2002). Evaluation of the functional effects of a course of bobath therapy in children with cerebral palsy: a preliminary study. *Developmental Medicine and Child Neurology*, 44(7):447–460.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.

Kohonen, T. (1997). Exploration of very large databases by self-organizing maps. In *Proceeding of the IEEE International Conference on Neural Networks*, pages 1–6.

Kohonen, T. (2001). *Self-Organizing Maps*. Springer, Third extended extension.

Kozima, H., Nakagawa, C., and Yasuda, Y. (2005). Interactive robots for communication-care: A case study in autism therapy. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, pages 341–346.

Kulyukin, V., Gharpure, C., Nicholson, J., and Osborne, G. (2006). Robot-assisted wayfinding for the visually impaired in structured indoor environments. *Autonomous Robots*, 21(1):29–41.

Kwakkel, G., Kollen, B. J., and Krebs, H. I. (2008). Effects of Robot-Assisted Therapy on Uper Limb Recovery After Stroke: A Systematic Review. *Neurorehabilitation and Neural Repair*, 22(2):111–121.

Kwee, H. H., Duimel, J. J., Smits, J. J., de Moed, A. A. T., and van Woerden, J. A. (1989). The Manus wheelchair-borne manipulator: System review and first results. In *Proceedings of the 2nd Workshop on Medical and Healthcare Robotics*.

Lacey, G. and Dawnson-Howe, K. (1998). The application of robotics to a mobility aid for the elderly blind. *Robotic and Autonomous Systems*, 23(4):245–252.

Lacey, G., MacNamara, S., Petrie, H., Hunter, H., Karlsson, M., Katevas, N., and Rundenschold, J. (1999). Adaptive control of a mobile robot for the frail visually impaired. In *Proc. Int. Conference on Rehabilitation Robotics ICORR*, pages 60–66.

Lee, C. and Xu, Y. (1996). Online, interactive learning of gestures for human/robot interfaces. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2982–2987.

Li, M., Chen, X., Li, X., Ma, B., and Vitanyi, P. M. B. (2004). The similarity metric. *IEEE Transactions on Information Theory*, 50(12):3250–3264.

Li, M. and Vitanyi, P. M. B. (2001/2002). *Algorithmic Complexity*. International Encyclopedia of the Social and Behavioral Sciences, N.J. Smelser and P.B. Baltes, Eds., Pergamon, Oxford.

Lungarella, M. and Berthouze, L. (2002). Adaptativity via alternate freeing and freezing of degrees of freedom. In *Proc. 9th Int. Conf. on Neural Information Processing*, pages 482–487.

Mahoney, R. M., Van Der Loos, H. M., Lum, P. S., and Burgar, C. (2003). Robotic stroke therapy assistant. *Robotica*, 21(1):33–44.

Marti, P., Fano, F., Palma, V., Pollini, A., Rullo, A., and Shibata, T. (2005). My Gym Robot. In *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, pages 64–73.

Matarić, M. J., Eriksson, J., Feil-Seifer, D. J., and Winstein, C. J. (2007). Socially assistive robotics for post-stroke rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 4(5).

Movellan, J. R., Tanaka, F., Fasel, I., Taylor, C., Ruvolo, P., and Eckhardt, M. (2007). The RUBY project: A progress report. In *2nd ACM/IEEE Int. Conference on Human-Robot Interaction (HRI07)*, pages 333–339.

Nadel, J., Guerini, C., Peze, A., and Rivert, C. (1999). *The evolving nature of imitation as a format for communication*, pages 209–234. Cambridge University Press.

National Autistic Society (2008). National autistic society. http://www.nas.org.uk/.

Otero, N. R., Saunders, J., Dautenhahn, K., and Nehaniv, C. L. (2008). Teaching robot companions: The role of scaffolding and event structuring. *Connection Science*, 20(2-3):111–134.

Pavlov, I. P. (1927). *Conditional Reflexes.* London Oxford University Press.

Piaget, J. (1928). La causalité chez l'enfant. *British Journal of psychology*, 18(3):276–301.

Piaget, J. (1931). Le développement intellectuel chez les jeunes enfants: étude critique. *Mind*, 40(158):137–160.

Piaget, J. (1945). *La formation du symbole chez l'enfant: Imitation, jeu et rêve, image et représentation.* Delachaux et Niestlé.

Powell, S. (2000). *Helping Children with Autism to Learn.* David Fulton Publish.

Rabiner, L. R. (1989). A tutorial on Hidden Markov Models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286.

Robins, B. and Dautenhahn, K. (2006). The role of the experimenter in HRI research - a case study evaluation of children with autism interacting with a robotic toy. In *Proc. 15th IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN 06)*, pages 646–651.

Robins, B., Dautenhahn, K., and Dubowski, J. (2005a). Robots as isolators or mediators for children with autism? A cautionary tale. In *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, pages 82–88.

Robins, B., Dautenhahn, K., te Boekhorst, R., and Billard, A. (2005b). Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society (UAIS)*, pages 105–120.

Robins, B., Dickerson, P., Stribling, P., and Dautenhahn, K. (2004). Robot-mediated joint attention in children with autism: A case study in robot-human interaction. *Interaction Studies*, 5(2):161–198.

Rogers, C. (1976). *Client-Centred Therapy.* Constable.

Ryan, V. (1999). Developmental delay, symbolic play and non-directive play therapy. *Clinical child Psychology and Psychiatry*, 4(2):167–185.

Ryan, V. (2001). Non-directive play therapy with children experiencing psychic trauma. *Clinical child Psychology and Psychiatry*, 6(3):437–453.

Ryan, V. and Needham, C. (2004). Non-directive play therapy with children experiencing psychic trauma. *Clinical child Psychology and Psychiatry*, 9(1):75–87.

Ryan, V. and Wilson, K. (1996). *Case-Studies in Non-Directive Play Therapy.* Baillière Tindall.

Salter, T., Dautenhahn, K., and te Boekhorst, R. (2006). Learning about natural human-robot interaction. *Robotics and Autonomous Systems*, 54(2):127–134.

Salter, T., Michaud, F., Lee, D., and Werry, I. P. (2007). Using proprioceptive sensors for categorizing human-robot interactions. In *Proc. of the 2nd Int. Conference on Human-Robot Interaction HRI07*, pages 105–112.

Scassellati, B. (2005a). How social robots will help us to diagnose, treat, and understand autism. In *Proc. 12th Int. Symposium of Robotics Research (ISSR)*, pages 552–563.

Scassellati, B. (2005b). Quantitative metrics of social response for autism diagnosis. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, pages 585–590.

Schatzman, L. and Strauss, A. (1973). *Strategy for recording.* Schatzmann, L., Strauss, A.L. Eds. Field Research. Strategies for a natural sociology. Prentice Hall.

Schopler, E., Reichler, R. J., Vellis, R. F. D., and Dally, K. (1980). Towards objective classification of childhood autism: Childhood autism rating scale (CARS). *Journal of Autism Developmental Disorders*, 10(1):91–103.

Shannon, C. E. (1949). *The mathematical theory of communication.* The University of Illinois Press, Urbana.

Shibata, T. (2004). Ubiquitous surface tactile sensor. In *Proc. 2004 1st IEEE TExCRA on Robotics and Automation*, pages 5–6.

Shibata, T., Wada, K., Saito, T., and Tanie, K. (2005). Human interactive robot for psychological enrichment and therapy. In *Proc. AISB'05 Symposium on Robot Companion Hard Problem and Open Challenges in Human-Robot Interaction*, pages 98–109.

Siegel, S. and Castellan, N. J. (1988). *Nonparametric Statistics for the behavioral sciences.* McGraw-Hill International Editions.

Skinner, B. F. (1974). *About Behaviourism.* New York: Vintage.

Slonim, N. and Tishby, N. (1999). Agglomerative information bottleneck. In *Proc. of Neural Information Processing Systems (NIPS 99)*, pages 617–623.

Stanton, C. M., Kahn, P. H., Severson, R. L., Ruckert, J. H., and Gill, B. T. (2008). Robotic animals might aid in the social development of children with autism. In *Proceedings of the 3rd International Conference on Human-Robot Interaction.*

Stiehl, W. D., Lieberman, J., Breazeal, C., Basel, L., Lalla, L., and Wolf, M. (2006). The design of the huggable: A therapeutic robotic companion for relational, affective touch. In *AAAI Fall Symposium on Caring Machines: AI in Eldercare.*

Syrdal, D. S., Walters, M. L., Koay, K. L., and Dautenhahn, K. (2008). The role of autonomy and interaction type on spatial comfort in an HRI scenario. In *Proceedings of 'Robotic Helpers: User Interaction, Interfaces and Companions in Assistive and Therapy Robotics', a workshop at ACM/IEEE HRI 2008*, pages 27–29.

Tammet, D. (2006). *Born on a blue day: A Memoir of Asperger and an Extraordinary Mind*. Hodder and Stoughton.

Tanaka, F., Cicourel, F., and Movellan, J. R. (2007). Socialization between toddlers and robots at an early childhood education center. In *Proc. of the National Academy of Sciences of the USA (PNAS) 104(46)*, pages 17954–17958.

Tanaka, F., Fortenberry, B., Aisaka, K., and Movellan, J. R. (2005). Developing dance interaction between qrio and toddlers in a classroom environment: Plans for the first steps. In *Proc. 14th IEEE Int. Workshop on Robot and Human Interactive Communication (RO-MAN)*, pages 223–228.

Tanaka, F., Movellan, J. R., Fortenberry, B., and Aisaka, K. (2006). Daily hri evaluation at a classroom environment: Reports from dance interaction experiments. In *Proc. 1st Annual Conf. on Human-Robot Interaction (HRI)*, pages 3–9.

Tardif, C., Plumet, M.-H., Beaudichon, J., Waller, D., Bouvard, M., and Leboyer, M. (1995). Micro-analysis of social interactions between autistic children and normal adults in semi-structured play situations. *International Journal of Behavioural Development*, 18(4):727–747.

Tishby, N., Pereira, F. C., and Bialek, W. (1999). The information bottleneck method. In *Proc. of the 37-th annual Allerton Conference on Communication, Control and Computing*, pages 368–377.

Tolman, E. C. (1932). *Purposive behavior in animals and men*. New York: Appleton.

Vygotsky, L. (1988). *The genesis of higher mental functions*. Richardson, K. and Sheldon S. (eds) Learning to think, London, Routledge/The Open University.

Vygotsky, L. S. (1978). *Mind in Society: Development of Higher Psychological Processes*. Harvard University Press.

Wada, K. and Shibata, T. (2006). Robot therapy in a care house - Its sociopsychological and physiological effects on the residents. In *Proc. 2006 IEEE International Conference on Robotics and Automation*, pages 3966–3971.

Watson, J. B. (1913). Psychology as the behaviourist views it. *Psychological Review*, 20:158–177.

Watson, J. B. (1925). *Behaviourism.* New York, Norton.

Werry, I. and Dautenhahn, K. (1999). Applying mobile robot technology to the rehabilitation of autistic children. In *Proc. SIRS99, 7th Int. Symposium on Intelligent Robotics Systems*, pages 265–272.

Werry, I., Dautenhahn, K., Ogden, B., and Harwin, W. (2001). Can social interaction skills be taught by a social agent? The role of a robotic mediator in autism therapy. In *Proc. 4th Int. Conference on Cognitive Technology: Instruments of Minds (CT2001)*, pages 57–74.

Williams, D. (1996). *Autism: an Inside-Out approach.* Jessika Kingsley Publishers.

Wünstel, M., Polani, D., Uthmann, T., and Perl, J. (2000). Behavior classification with self-organizing maps. In *Proc. 4th Int. RoboCup Workshop*, pages 179–188.

Yanco, H. (1998). Integrating robotic research: A survey of robotic wheelchair development. In *AAAI Spring Symposium on Integrating Robotic Research.*

Yanco, H. (2001). Development and testing of a robotic wheelchair system for outdoor navigation. In *Proceedings of the 2001 Conference of the Rehabilitation Engineering and Assistive Technology Society of North America.*