Available online at www.sciencedirect.com

SCIENCE DIRECT®

Robotics and Autonomous Systems 54 (2006) 127–134

Robotics and Autonomous Systems

www.elsevier.com/locate/robot

Learning about natural human–robot interaction styles

Tamie Salter*,[1], Kerstin Dautenhahn, René te Boekhorst

Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, College Lane, Hatfield AL10 9AB, United Kingdom

Received 13 October 2004; accepted 13 September 2005

Abstract

If we are to achieve natural human–robot interaction, we may need to complement current vision and speech interfaces. Touch may provide us with an extra tool in this quest. In this paper we demonstrate the role of touch in interaction between a robot and a human. We show how infrared sensors located on robots can be easily used to detect and distinguish human interaction, in this case interaction with individual children. This application of infrared sensors potentially has many uses; for example, in entertainment or service robotics. This system could also benefit therapy or rehabilitation, where the observation and recording of movement and interaction is important. In the long term, this technique might enable robots to adapt to individuals or individual types of user.

Keywords: Natural human–robot interaction; Patterns of behaviour; Touch; Adaptation

1. Introduction

Human–Robot Interaction (HRI) is a highly interdisciplinary field where behavioural and psychological approaches towards understanding the *nature of human–robot interaction* complement robotics and engineering oriented work.[2] The work presented in this paper focuses on the nature of *interactions* between children and robots, investigating how a simple and robust technique, based on a robot's sensor readings, can be used to automatically detect and distinguish human contact.

Our work is part of the AuRoRA project [1], which is attempting to develop robots for use with children with autism in a therapeutic and educational context. Robots are already increasingly being used in assistive technology, rehabilitation, surgery and therapy; e.g. [2–6]. Other areas in which the use of robotics is growing, include service and entertainment domains; e.g. [7–9]. Methods that enable easy and effective communication between robots and humans are crucial in all of these areas. When investigating human–robot communication, people normally consider using cameras for visual recognition or audio for sound recognition. There are however other ways in which people communicate such as touch and through patterns of behaviour, which can be recorded by other means. Passive sensors, e.g. infrared (IR), are increasingly being used in the homes of the elderly to communicate to security systems [10]. We will show how through the use of IR sensors we have been able to capture and record natural touch from children interacting with a mobile robot. The records of interaction derive from the robot interacting with a single child. The robot functions as a mobile moving toy and registers the child's proximity and physical handling during play. We analyse the sensor input data from the infrared sensors to recognize patterns in the way children interact with the robot and compare the results with an analysis of behavioural observations of the children's play with the robot. In the discussion we will outline the limitations of this technique and suggest the application of IR sensors as a valuable channel in human–robot interaction.

1.1. Communicating with robots

For robots to be useful to humans they must be able to respond to our needs and demands and this will require some form of communication. Human–robot communication systems

* Corresponding author.
E-mail address: tamie.salter@usherbrooke.ca (T. Salter).

[1] Present address: LABORIUS, Department of Electrical Engineering and Computer Engineering, Local C1-5103, Faculté de génie, Université de Sherbrooke, 2500 boul. Université, Sherbrooke, Québec, Canada J1K 2R1.

[2] Recent workshop and conference proceedings on HRI topics are strongly reflecting this trend, e.g. Proc. IEEE Ro-man 2004, 13th IEEE International Workshop on Robot and Human Interactive Communication September 20–22, 2004 Kurashiki, Okayama, Japan.

0921-8890/$ - see front matter © 2006 Published by Elsevier B.V.
doi:10.1016/j.robot.2005.09.022

are becoming increasingly sophisticated but there is still a long way to go before natural communication will be achieved. Most current robot–human interfaces involve vision or speech recognition (e.g. [11–13]). Some systems make use of push down buttons to register touch or communication from people; e.g. [11,14,7]. However, communicating through natural touch or physical contact is an area which has not yet been fully explored by robot developers. Researchers are now beginning to recognize the importance of natural touch as a means of communication with a robot, especially people working with children or in therapy; e.g. [15,16,6]. According to Ito "*in the case of very young children, non-verbal communication is more important than verbal communication. Especially, tactile communication is very important*" [16]. Most touch sensors used in robotics are based on microswitches, which are very different from biological touch or surface sensors, e.g. the skin. A human being's skin can facilitate a variety of more natural touch interaction including stroking, gentle touching etc. A move away from the typical use of microswitches for detecting touch is realized in the seal robot Paro, which was developed for robot assisted activity in hospitals or homes for the elderly. In this system, physical contact with the robot, for example touching, is recognized by a system based on balloons [6]. Our work further addresses the question of how robot–human interfaces can be improved using inexpensive, reliable and robust IR sensors, which can recognize a very natural form of touch without the need to exert force or to press down in a specific area.

1.2. Detecting individuality

Our specific research interest is the use of IR sensors to ultimately detect individual differences in the play and interaction styles of children [15,17]. Individuality is a distinct feature of humans that also needs to be acknowledged in robot–human interfaces. It has been argued previously that in the context of human–agent interface design "*Humans are individuals and they want to be treated as such*" [18, p. 609]. A related view has been expressed by [19, p. 146], and again in [13]. Detecting differences among individual users is an important requirement in therapeutic contexts where we find a large spectrum of individual people's needs and abilities. Generally, detecting and responding to differences in how individual people interact with robots is likely to benefit the development of individualized robot companions in a variety of application areas [20].

1.3. Adaptation

In the AuRoRA project [1] we require robots that can adapt to individuals *easily*. Our work with autistic children does not allow for any lengthy process of decision making or learning onboard the robot. Children simply want to play with our robots and they want them to respond in real time. The method presented here will in the future allow our robots to adapt to the 'type' of child they are interacting with in real time. This will be the first step in adapting to individuality. Adaptation will be useful in the AuRoRA project as the autistic children that interact with our robots can be extremely different in their general 'personality types'. For example, some of the children with autism can be anxious and might be cautious or wary of the robot, whereas others will be totally unafraid and very physically aggressive with the robot. It is a concern that we do not want the robot's behaviour to be such that it will overwhelm a child, and, on the other hand, we do not want the robot to bore a child. Therefore if the robot detects the general type of child it is interacting with it can adapt its behaviour to suit that particular child. For example, for a more cautious child the robot's behaviour could perhaps be slow and have an unthreatening manner, for a confident child perhaps the behaviour could be faster and more exciting. This will be of added benefit as children with autism are often very rigid in their behaviour, finding it hard to progress, change or learn new things. A robot that can detect an individual could adapt its behaviour to suit that particular child and progress onto more complex behaviours at a time that suits the child.

1.4. Recording movement and behaviour

Detecting human interaction through touch can lead to the development of a method that will not only assist in developing human–robot communication and interaction methods but also provide a method to help quantify and study human behavioural characteristics. The idea of using sensors to record human activity is not a novel idea outside the field of robotics. There are five general methods for assessing human activity: behavioural observations, questionnaires, activity diaries, physiological markers (e.g. heart rate), calorimetry, and mechanical and electronic motion sensors [21]. It is through the use of sensor data that we are now beginning to be able to quantify human behaviour. Haigh et al. [10] have loosely termed this 'behaviour recognition'. One of the first projects that used infrared sensors to monitor people's behaviour or activity in their own homes was that of Yamaguchi et al. [22]. Monitoring of people's behaviour through the use of sensors is now a standard approach in so-called "Smart Homes" as can be seen in [23,10]. The approach is to analyse people's behaviour by using stationary sensors located in the environment of elderly people's houses. These sensors are there to investigate patterns of human behaviour that could be useful in monitoring the normal routine of the occupants [10,24]. MacCulloch [25] attempted to use electronic data to study the play patterns of children, using a sensitive floor to record the movements of children. In our work, the infrared sensors used to record movement and behaviour are located on a simple, autonomously moving, robot.

2. Working hypothesis

We argued above that if we are to develop useful, believable and socially intelligent autonomous robots that can interact with people in a more 'natural' way, then an essential property will be the ability to adapt to differences in an individual's interaction style. In this work we hope to develop a control architecture that can be used in robots. The control architecture will give robot developers a method of detecting human interaction. This will enable the robot to adapt to the highly variable behaviours of people. The main aim of the AuRoRA project [1] is to develop scenarios, methods and control architectures for robots that assist autistic children in a therapeutic and educational context. Robots are already being studied in a variety of therapeutic or educational contexts; e.g. [26–29].

2.1. Different patterns of behaviour and interaction

To design a control architecture for a robot that can adapt to individuals we first need to know *what characteristics of human behaviour the robot should adapt to*. To approach these matters, we are investigating how typically developing children of different types play with the same robot. We are analysing how typically developing children interact with a robot so we can then later compare these patterns with those of autistic children. Our working hypothesis for the current work is that characteristic patterns in the behaviour of the different children will be revealed as 'fingerprints' of the child–robot interaction in the registration of the IR sensor data. It is crucial to future development that we learn about different types of interaction pattern. Once we can identify certain patterns of interaction coming from corresponding groups of children then this information can be used to adapt the behaviour of the robot. We will compare the readings from sensor data to established observational methodologies currently used, for example by psychologists.

3. The robot and behaviour exhibited

We use a commercially available medium sized mobile robot called Pekee (see Fig. 1) [30]. The robot comes equipped with a ring of 15 infrared sensors located around its rim; six of these are located to the front, three to each side and three to the rear. Measurements taken from these sensors are used to record the interaction from the children. It has two front wheels, each with its own motor, and a castor wheel at the back. Pekee is 400 mm long, 255 mm wide and 210 mm high and weighs approximately 2.9 kg. It has a maximum speed of 6 km/h, although during the experiments the speed did not exceed 3 km/h. The mobility of the robot gives the children a variety of ways and positions for interacting with it. The robot executes a wandering and simple obstacle avoidance behaviour, treating the child as an 'obstacle' regardless of how the child responds to the robot. This allows us to study the different characteristics of the individual children's styles of interaction with the robot without any direct effect caused by the robot's behaviour. Therefore the experimental set-up and behaviour of the robot were exactly the same for each child. The rationale of this approach was to achieve an (as far as possible) unbiased record of the different types of interaction exhibited by the children that will be utilized at a later date to allow the robot to recognize that type of child. The robot does not have a traditional goal such as navigation. The sole purpose of the robot is to engage the children and encourage interaction. The environment from the robot's point of view is highly dynamic and unpredictable: 'anything can happen' while children interact with the robot (children often treat the robot quite roughly). As such, the control architecture cannot rely on any precise measurements due to the pushing, pulling and picking up, etc. This also makes it preferable for the robot to exhibit a simple behaviour so that the robot continues to function under these extreme conditions. The robot measures distances to objects in the environment by means of its infrared sensors and changes the wheel direction and speed accordingly so as to clearly avoid them. This allows the robot free movement in uncluttered environments.[3] At the same time, close contact or interaction (less than 2 cm) is recorded.[4]

Fig. 1. Pictures showing the arena and different styles of play.

4. The experimental set-up

In therapy situations, the environment quite often consists of an individual child with the therapy aid and care worker(s). In such situations the environment is usually highly structured in order to focus attention, e.g. on a few relevant stimuli. Thus in these rooms the level of clutter (other objects in the environment) is usually low. We have carried out two different trials. In the first trial we invited six children to the University of Hertfordshire and tested interaction with the mobile robot under laboratory conditions. This environment might be similar to the type of environment used in a therapy context. However the environment in this experiment is unknown to the children and unfriendly (hard surfaces, planks of wood etc.) whereas it would normally be at the child's school and thus friendly and familiar to them. The children were taken to the experimental room one at a time, with only the experimenter present (Fig. 1). This room contains an arena, the 'pen' of the robot, of approximately 2 m^2 and is enclosed by four shallow wooden walls. For the experiments, the robot was connected to the network with a long flexible lead.[5] The children were asked to step inside the arena, to play with Pekee, or to do whatever they liked. The experiment lasted between 1 and $1\frac{1}{2}$ min. Some

[3] Experiments investigating how different levels of clutter within an environment can influence the interaction levels detected by the robot are reported in [31].

[4] Note, in this paper we refer to 'interaction' as any measurement that falls into this category, based on the robot's IR readings, and regardless of whether these interactions were caused by encounters with animate or inanimate entities (e.g. children or the walls surrounding the pen).

[5] The lead did not seem to interfere with the children's play, but the onboard computer has since been updated and now does not require this connection.

initial trials were carried out for longer periods but it was found that shy children grew uncomfortable in the unfriendly environment of the robotics laboratory and wanted to leave; for the experiments proper we accordingly shortened the time to around one minute. During the experiment the children had complete freedom to do as they pleased, playing with the robot or watching as it moved about; no child completely ignored the robot or pursued an activity that had no regard to it. In the second set of trials we repeated this experiment but with a larger sample size of typically developing children in a real school environment. In this trial each child was exposed to the robot several times and for a longer period of time (5 min). The children were not hand-picked as in this experiment (see the next section). Details of the experiments, that are in line with the findings in this paper, can be found in [15].

5. The children

In this study all the children used were typically developing boys between 5 and 7 years of age. Their general character or personality types, according to their parents, ranged from boisterous to shy. It was the first time the children had been exposed to a robot. The children were chosen as they clearly fitted into one of three broad psychological groupings that we defined: Boisterous/Active (Type A), Average (Type B) and Passive/Shy (Type C). A child was categorized as type 'A' if (i) he was considered naughty at home and school, and (ii) he seemed confident, unafraid and active. Children were placed in the type 'C' category if (i) they usually did as they were told both at school and at home, (ii) they seemed to require the security of being with a familiar adult, and (iii) they did not readily explore on their own. A child not falling into either of these categories was classified as type 'B'. The information about the children came from discussion with the children's parents, acquiring information from them about the child's behaviour at school and at home. The reason for using clearly defined children was to learn what sensor readings for these clearly defined types of child looked like. The six children used were classified as follows: two as type A, two as type B and two as type C.

6. Data and data analysis

The data were obtained from two sources: (1) sensor readings, and (2) direct observation from video footage taken of the children. Statistical analysis techniques were used in order to compare which data source was more successful for identifying people's general personality type or behavioural characteristics: sensor readings or observational techniques.

6.1. Sensor readings

Sensor readings are written to a text file onboard the robot. Eight times per second any touch or interaction (i.e. close proximity less than 2 cm distance) for each sensor is recorded. These data 'profiles' characterize the interaction pattern of that particular child with the robot for that particular run, leaving behind a 'fingerprint' of the child for analysis. To check the validity of the sensor readings prior experiments were conducted by leaving the robot moving (carrying out obstacle avoidance) on its own in the pen. No interaction was detected. Other tests were carried out: for example, where a person entered the pen and touched a single zone of the robot, e.g. the tail, to confirm that correct data were being recorded. Sometimes events during the trials were misclassified as interaction by the sensors, e.g. when the children forced the robot against the wall (cf. [31]). However, even in these cases (which could be detected by matching the video data with the sensor data) the information is still valuable. Prima facie it shows high levels of interaction, and possibly a child that behaved more aggressively towards the robot. We investigated how the children of different types played with the robot. We were interested to see whether there were any patterns in sensor activation; e.g. did the type A children play with the robot in a similar way to each other and was this different from how the type B and type C children played with the robot?

Three different techniques for analysing the robot's sensor data have been considered. **(1)** *Clustering the different children's profiles (overall sensor data)*—We performed a cluster analysis [32] on a matrix containing the children's profiles (overall interaction levels from the 15 IR sensors). This shows us the degree of similarity among these profiles. As we wanted to investigate how the children played with the robot regardless of their overall activity frequency (certain children were much more active than others), we logarithmically transformed the data prior to the analysis. The children (rows in matrix) were clustered over the sensors (columns). This shows which children were similar in terms of sensor activation patterns. See the results in Fig. 2(ii). **(2)** *Interaction levels from sensor data*—We investigated whether there were any significant differences in the levels of sensor activation, e.g. whether active children produced higher activity levels. See the results in Table 1. **(3)** *Time series data from the different sessions*—Graphs were produced that show the dynamics of the sessions. Patterns of activity can be seen from these graphs. These graphs show us the activation of the sensors and allow us to visually see the 'fingerprint' left behind by the child. See the examples in Fig. 3.

6.2. Behavioural observations

The behavioural observation data analysis technique was based on an established methodology, used in [33,26]. The tapes were viewed and 21 behaviours were identified. These fell into three general categories: General Movement Behaviours such as 'Lie on Floor' and 'Kneel'; Robot Centred Behaviours such as 'Touch Robot', 'Approach Robot' and 'Talk to Robot'; and Person Centred Behaviours such as 'Talk to Experimenter' and 'Look at Experimenter'. While viewing the videotapes, the 21 behaviours were recorded on a protocol sheet in the sequence in which they occurred.

Two different techniques were used to analyse the observational data. **(1)** Clustering of behavioural data—The same technique as applied to the sensor data. See the results in

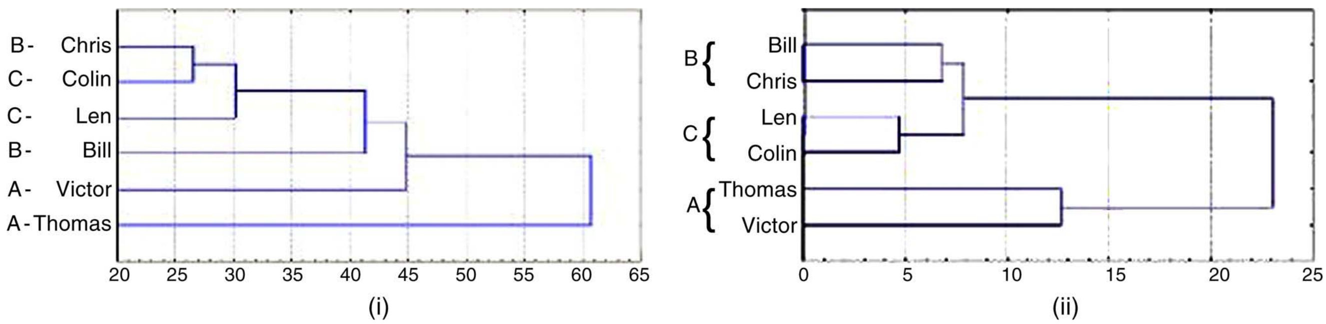

Fig. 2. Dendrogram showing (i) boys clustered with observational data and (ii) boys clustered with sensor data.

Table 1
Overall interaction levels for all the children

Name	Colin (C)	Len (C)	Chris (B)	Bill (B)	Thomas (A)	Victor (A)
Level	9	30	115	124	725	4330

Classified type is shown next to the name.

Table 2
Table showing overall levels of behavioural activity

Name	Colin (C)	Len (C)	Chris (B)	Bill (B)	Thomas (A)	Victor (A)
Level	33	34	28	35	56	35

Classified type is shown next to the name.

Fig. 2(i). Analysis was performed on the accumulative count of the activities (i.e. on how many times each child performed each of the 21 defined behaviours). **(2)** Activity levels—The levels of the 21 behaviours performed were added together to give an overall behavioural activity level. See the results in Table 2.

7. Results

Each of the children played with Pekee in his own way and all appeared to enjoy the experience. Behaviour ranged from cautious and wary, to very confident or physically aggressive. The children did play with the robot according to their classified type, that is to say, type A children (boisterous, active) played with the robot in a more physically aggressive and active manner than the other two groupings. Type C children (shy, cautious) played with the robot in a cautious manner. This could be clearly seen from the video footage and sensor data but not from statistical analysis of observational data. This therefore confirmed our initial hypothesis: the sensor readings did detect and classify the children into their prior psychological groupings. It was not so easy to classify the children by means of traditional behavioural observation techniques. Below we see a breakdown of the results from both the sensor readings and behavioural observations.

7.1. Sensor readings

Data produced from the robot's sensors clearly indicate differences in the way the children played with the robot. **(1)** The cluster analysis of the sensor data showed interesting results. There were indeed patterns in the activation of the sensors; the children were grouped into clusters that corresponded to the prior psychological classification (Fig. 2(ii)). This indicates that children can be classified into personality types on the basis of the sensor activation.

(2) A simple indication as to the type of child interacting with the robot comes from the overall interaction level from the sensor readings; see Table 1. Active children have very high interaction levels, passive children have low interaction levels and type B (average) children have interaction levels in a middle section. Extremely high interaction levels were not expected, as the trial lasted only between 1 and $1\frac{1}{2}$ min[6] and it was also the first time the children had been exposed to a robot. However the most active child had a very high interaction level because he interacted with the robot straight away. Results like this might in the future enable a very simple technique for adaptation in robots.

(3) Time series of the sessions were produced that show the different patterns of play over time. These diagrams show the dynamics of the sessions and a dramatic difference between the most active child (Victor) and the most passive child (Colin); see Fig. 3 for examples. Each time a black line is seen in the diagram it shows a sensor being touched or interacted with. Sensors are shown along the vertical axis and time along the horizontal axis. Sensor number 0 appears at the top of the axis and number 14 appears at the bottom of the axis. If all the sensors were activated simultaneously we would see a straight black line starting at the top of the graph and proceeding straight down to the bottom of the graph. For active children we see lots of black lines, i.e. lots of interaction or touching of sensors. With passive children we see a small amount of black lines indicating little interaction or touching of sensors from that particular child. Hence the more black seen in a graph the more activation or touching took place. In the three examples shown we see a type A child (this child was very confident with the robot) touching the robot, jumping over the robot and pushing the robot whilst it continued with its simple wandering and obstacle avoidance behaviour. This can be seen from the large amount of black lines on the graph. The type B child was curious about the robot and followed the robot without making a lot of close contact. He did however once investigate the robot at close proximity. This can be seen from the black lines in

[6] The most recent trial carried out was a longitudinal study where each session lasted 5 min (see [15]).

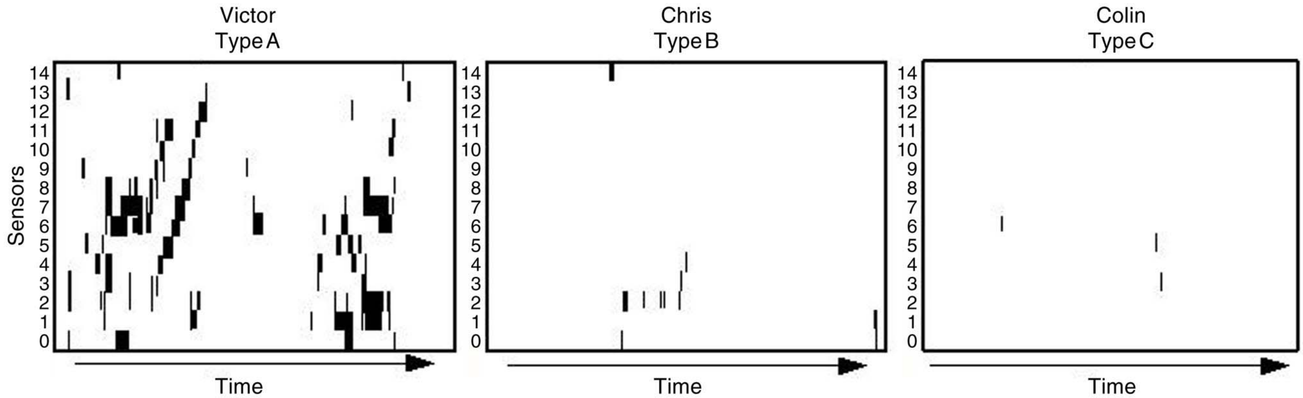

Fig. 3. Time Series plots for the three types of child. Black lines indicate the activation or touching of a sensor. Sensors are shown on the vertical axis and time along the horizontal axis.

the central part of the graph. The type C child was cautious of the robot; he spoke to the experimenter but stayed away from the robot. Once he approached the robot, touched it and then quickly backed away. This can be seen from the small amount of black lines shortly after the central part of the graph.

7.2. Behavioural observations

It was much harder to recognize clearly distinguished groups within the observational data, i.e. the children did not seem to cluster into clear groups. However, closer analysis does reveal one similarity between the previous psychological classification made of the children and the observations taken from the video. As can be seen from Fig. 2 (i), type A boys (active, boisterous) are typically split off from the rest of the children. Again the overall behavioural activity levels do not show any clear patterns with regard to the type of child; see Table 2.

8. Summary of results and discussion

Our initial working hypothesis was confirmed: we identified, using the robot's sensor data, that the children had different styles of playing with robot. This allowed categorization of the children through information from the robot's sensor readings. It was not so easy to classify the children from observational techniques. Specifically, our findings are:

- The analysis based on the robot's sensor readings supports the prior psychological classification of the children. The analysis of observational data gives less clear results.
- Irrespective of widely differing interaction levels, children of different types exhibited patterns of behaviour which showed similarities within the type group and differences between the groups.
- A very simple indication of the type of child interacting with the robot comes from the overall interaction levels.

In this paper we argued that infrared sensors located on a mobile robot can be used to detect and distinguish human interaction. We have also indicated that this detection might then be used to adapt the robot's behaviour to an individual child or children of different types. Our findings represent a first step towards reaching this long-term goal.[7] The importance of this is threefold:

(1) It shows that personality type is reflected in a child's behaviour while playing with a robot.
(2) It is a prerequisite for building robots that adapt to human behaviour.
(3) Within the educational, observational or therapeutic fields, it potentially gives a researcher, teacher, parent, or therapist an extra tool for quantifying and assessing children's behaviour.

These findings suggest that touch could have an important role to play in developing natural human–robot interfaces that might allow a robot to adapt to the individual that it is interacting with. Since writing this paper we have conducted a new set of trials, with larger sample sizes, and repeated exposure of children to the robot. The new trials were held at the children's schools in order to provide a more familiar environment; see [15] for our detailed findings. As these subsequent trials were over a longer period, we have found that the results also confirm findings from Kanda et al. [13]. Kanda et al. found that typically developing children get bored by robots, i.e. the novelty effect wears off over time. This therefore supports our argument for developing robots that can adapt to different interaction levels, based on personality styles, and also based on long-term effects. For example, a robot interacting with a shy child could react by slowing down, stopping at times and thus being non-threatening; similarly it could speed up, spin and beep so as to hold the interest of the confident child, or a child that did have a high interaction and has now lost interest in the robot. In this paper we established a technique that could be a stepping stone towards robots that can adapt to human behaviour. This is interesting when compared to the behavioural observation data results, as it is perhaps an indicator that when considering human–robot interfaces developers should not only consider vision and speech but also touch. Our results indicate that touch really does have an important role to play if we are to

[7] We should emphasize however that the quantitative approach we pursue will be able to point out only certain aspects of a child's behaviour, while other methods might be useful for a more detailed assessment of children; cf. [34].

develop robots that can interact and communicate with humans naturally.

Acknowledgements

This article is based on [17]. We would like to thank anonymous reviewers for many constructive comments that helped us to improve an earlier version of this paper.

References

[1] AuRoRA Project, http://www.aurora-project.com, 2003 (last accessed 23/4/03).
[2] C. Martens, O. Radchenko, A. Pape, H. She, I. Volosyak, A. Gräser, Autonomous and robust beverage serving-task with the rehabilitation robotic system friend, in: 7th European Conference for the Advancement of Assistive Technology (AAATE), Dublin, Ireland, 2003.
[3] M. Hans, B. Graf, R.D. Schraft, Robotic home assistant care-o-bot: Past–present–future, in: IEEE Ro-man 2002, 11th International Workshop on Robot and Human Interactive Communication, Berlin, Germany, 2002, pp. 380–385.
[4] H. Hüttenrauch, L. Oestreicher, K. Severinson-Eklundh, A. Green, Fetch-and-carry with cero: Observations from a long-term user study with a service robot, in: IEEE Ro-man 2002, 11th International Workshop on Robot and Human Interactive Communication, Berlin, Germany, 2002, pp. 158–163.
[5] M. Hillman, Rehabilitation robotics from past to present—a historical perspective, in: IRCORR 2003, The Eighth International Conference on Rehabilitation Robotics, KAIST, Daejeon, Korea, 2003.
[6] D. Wada, T. Shibata, T. Saito, K. Tanie, Robot assisted activity for elderly people and nurses at a day service center, in: IEEE International Conference on Robotics and Automation, Washington, DC, 2002, pp. 1416–1421.
[7] Sony, http://www.aibo-europe.com, 2004 (lasted accessed 06/10/04).
[8] B. Graf, O. Barth, Entertainment robotics: Examples, key technologies and perspectives, in: IROS-Workshop Robots in Exhibitions, 2002.
[9] Honda, http://world.honda.com/ASIMO/, 2004 (last accessed 06/10/04).
[10] K. Haigh, L. Kiff, J. Myers, V. Guralnik, K. Krichbaum, J. Phelps, T. Plocher, D. Toms, The independent lifestyle assistant (i.l.s.a.): Lessons Learned, Tech. Rep., Honeywell Laboratories, 3660 Technology Drive, Minneapolis, 2003.
[11] B. Jensen, G. Froidevaux, X. Greppin, A. Lorotte, L. Mayor, M. Meisser, G. Ramel, R. Siegwart, The interactive autonomous mobile system robox, in: IROS 2002, IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE Press, 2002, pp. 1221–1227.
[12] T. Watanabe, R. Danbara, M. Okubo, Interactor: Speech-driven embodied interactive actor, in: IEEE Ro-man 2002, 11th International Workshop on Robot and Human Interactive Communication, IEEE Press, Berlin, Germany, 2002, pp. 430–435.
[13] T. Kanda, T. Hirano, D. Eaton, H. Ishiguro, Interactive robots as social partners and peer tutors for children: A field trial, Human–Computer Interaction 19 (2004) 61–84.
[14] A. Billard, Robota: Clever toy and educational tool, Robotics and Autonomous Systems 42 (2003) 259–269.
[15] T. Salter, K. Dautenhahn, R. te Boekhorst, Robots moving out of the laboratory—detecting interaction levels and human contact in noisy school environments, in: IEEE Ro-man 2004, 13th IEEE International Workshop on Robot and Human Interactive Communication, Kurashiki, Okayama, Japan, 2004, pp. 563–568.
[16] T. Ito, NEC Personal Robot Center, How children perceive robots. http://www.incx.nec.co.jp/robot/english/univ/05/univ_e05.html, 2003 (last accessed 06/10/04).
[17] T. Salter, R. te Boekhorst, K. Dautenhahn, Detecting and analysing children's play styles with autonomous mobile robots: A case study comparing observational data with sensor readings, in: IAS-8, 8th Conference on Intelligent Autonomous Systems, Amsterdam, NL, 2004, pp. 61–70.
[18] K. Dautenhahn, The art of designing socially intelligent agents—science, fiction, and the human in the loop, Applied Artificial Intelligence Journal 12 (1998) 573–617.
[19] T. Fong, I. Nourbakhsh, K. Dautenhahn, A survey of socially interactive robots, Robotics and Autonomous Systems (2003) 143–166.
[20] K. Dautenhahn, Robots we like to live with?!—a developmental perspective on a personalized, life-long robot companion, in: IEEE Ro-man 2004, 13th IEEE International Workshop on Robot and Human Interactive Communication, Kurashiki, Okayama, Japan, 2004, pp. 17–22.
[21] C.V.C. Bouten, Assessment of daily physical activity by registration of body movement, Ph.D. Thesis, Universiteit Eindhoven, 1995.
[22] A. Yamaguchi, M. Ogawa, T. Tamura, T. Togawa, Monitoring behaviour in the home using positioning sensors, in: 20th Annual Intern Conf IEEE/EMBS, 1998, pp. 1977–1979.
[23] V. Guralnik, K. Haigh, Learning models of human behavior with sequential patterns, in: AAAI-02 Workshop on Automation as Caregiver, 2002, pp. 24–30.
[24] R. Orpwood, C. Gibbs, T. Adlam, R. Faulkner, D. Meegahawatte, The Gloucester smart house for people with dementia—user-interface aspects, in: Designing a More Inclusive World, Springer-Verlag, 2004.
[25] M. MacCulloch, Objective methods and implications of recording children's behaviour, in: Movement and Child Development, Spastics International Medical Publication, London, 1975, pp. 132–144.
[26] K. Dautenhahn, I. Werry, A quantitative technique for analysing robot–human interactions, in: IROS2002, IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE Press, Lausanne, 2002, pp. 1132–1138.
[27] F. Michaud, C. Theberge-Turmel, Mobile robotic toys and autism, in: K. Dautenhahn, A. Bond, L. Cañamero, B. Edmonds (Eds.), Socially Intelligent Agents—Creating Relationships with Computers and Robots, Kluwer Academic Publishers, 2002, pp. 125–132.
[28] B. Robins, K. Dautenhahn, R. te Boekhorst, A. Billard, Effects of repeated exposure of a humanoid robot on children with autism, Universal Access and Assistive Technology (CWUAAT) (2004) 225–236.
[29] B. Prazak, A. Hochgatterer, G. Kronreif, M. Furst, Robot supported play—new possibilities for physically handicapped children?!, in: AAATE 2003, Dublin, Ireland, 2003.
[30] Wany Robotics, http://www.wanyrobotics.com, 2003 (lasted accessed 21-09-03).
[31] T. Salter, K. Dautenhahn, Guidelines for robot–human environments in therapy, in: IEEE Ro-man 2004, 13th IEEE International Workshop on Robot and Human Interactive Communication, Kurashiki, Okayama, Japan, 2004, pp. 41–46.
[32] http://www.statsoftinc.com, 2004 (last accessed 10/05/04).
[33] P. Smith, K. Connolly, The Ecology of Preschool Behaviour, Cambridge University Press, 1980.
[34] K. Dautenhahn, I. Werry, J. Rae, P. Dickerson, P. Stribling, B. Ogden, Robotic playmates: Analysing interactive competencies of children with autism playing with a mobile robot, in: K. Dautenhahn, A. Bond, L. Cañamero, B. Edmonds (Eds.), Socially Intelligent Agents—Creating Relationships with Computers and Robots, 2002, pp. 117–124.

Tamie Salter holds a degree in Computer Science. She is a PhD student in the Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire. Her research interests lie in investigating robots that could be useful in the therapy for and education of children with autism. Specifically she is interested in how analysis of onboard sensor readings can enable a robot to adapt to the individual user. She is currently working on the techniques listed in this paper at the Laborius Research Laboratory on Mobile Robotics and Intelligent Systems, Université de Sherbrooke, Quebec, Canada.

Kerstin Dautenhahn is Professor of Artificial Intelligence and coordinator of the Adaptive Systems Research Group at the University of Hertfordshire in UK. She has initiated and led research projects on socially intelligent agents and social robotics, including the Aurora project which studies the potential use of robotic toys in therapy and education of children with autism. She is currently involved in several European projects on social agents and social robots (Victec, Elvis, Cogniron, Robot-Cub), and she directs the Robotics and Interactive Systems Laboratory at the University of Hertfordshire, Hatfield, UK. Prof. Dautenhahn is an Editor-in-Chief of the journal Interaction Studies—Social Behaviour and Communication in Biological and Artificial Systems.

René te Boekhorst obtained his PhD at the University of Utrecht. His dissertation was on the social ethology and behavioural ecology of great apes (chimpanzees, orang-utans and gorillas) based on field data and individual oriented models. He moved to Switzerland, where he worked at the A.I. Lab of Prof. R. Pfeifer on collective robotics. Currently he is Senior Lecturer at the Computer Science Department at the University of Hertfordshire were he is involved in the application of statistical analysis to behavioural robotics. His interest is in non-linear time series analysis of complex (artificial, biological) systems and in dynamical systems modelling of motivation.