

Using Real-Time Recognition of Human-Robot Interaction Styles for Creating Adaptive Robot Behaviour in Robot-Assisted Play

Dorothee François, Kerstin Dautenhahn, Daniel Polani

Abstract—This paper presents an application of the Cascaded Information Bottleneck Method for real-time recognition of Human-Robot Interaction styles in robot-assisted play. This method, that we have developed, is implemented here for an adaptive robot that can recognize and adapt to children’s play styles in real time. The robot rewards well-balanced interaction styles and encourages children to engage in the interaction. The potential impact of such an adaptive robot in robot-assisted play for children with autism is evaluated through a study conducted with seven children with autism in a school. A statistical analysis of the results shows the positive impact of such an adaptive robot on the children’s play styles and on their engagement in the interaction with the robot.

I. INTRODUCTION

The work presented in this paper 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 communication, social interaction and imagination and fantasy¹. 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” [4]. Children with autism often play in a repetitive way, which can be linked to the children’s preference for predictable environments. The advantage of enabling children with autism to interact with a robot is that robots enable simple and safe interaction by initially providing a relatively predictable environment for play. Progressively the complexity of the interaction can be increased.

Different possible obstacles have been identified that often prevent children with autism to actualize their potential for play. Among them are impairments in socioemotional intersubjectivity, impairment in joint attention and impairment in Theory of Mind [5]. These impairments negatively influence

Dorothee François, Kerstin Dautenhahn and Daniel Polani are with the Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, Hatfield AL10 9AB, UK (email: {d.francois, k.dautenhahn, d.polani}@herts.ac.uk).

The first author is supported by a research scholarship of the University of Hertfordshire. The work described in this paper 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. The authors would like to thank the teachers at the school, the children who took part in the experiments, and their parents. Thank you to René de Boekhorst for his useful input and comments on the statistical analysis.

¹Autistic Spectrum Disorders can appear at various degrees and refer to different skills and abilities [1; 2]. Communication, social interaction and imagination and fantasy have been identified as the main impairments in autism, [3].

interaction in general and, more specifically, imply a lack of spontaneous and social reciprocity during play. Besides, the difficulty in perceiving the coherence of categories and concepts can be a reason why children with autism perceive an object in its parts and not as a whole, compare the weak central coherence theory [6; 7] for details. However, causes for impaired play are still not very well understood. These causes can vary for different children, depending also on the personality of the child and her past experience of play.

Yet play is an important vehicle for learning. Children can construct some understanding, i.e. active construction of meaning, through play. Besides, 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 a play situation [8]. 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 their quality of life. Play is also an important medium for self-expression [8].

Consequently, here we focus on facilitating play between children with autism and an autonomous robot, and particularly, we investigate the potential of a robot that can detect the children’s play styles and adapt to them accordingly (such a robot is called an ‘adaptive robot’ in contrast to a ‘reactive robot’ which would only respond to current sensory input). Our goal is to encourage the children to engage in play and, when playing, to encourage ‘well balanced’ tactile interaction styles, i.e. neither too forceful nor too weak and within an intermediate frequency of interaction. We therefore address the following research questions:

- Does the adaptive robot, as described above, encourage or discourage the children from engaging in the interaction with the robot? Does their engagement change when interacting with a reactive robot?
- Does a 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 tactile 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?

In order to study these research questions, a hard technical challenge needs to be addressed, namely how to enable the robot to recognize in real time the tactile play styles of a child. This has been achieved by applying the Cascaded Information Bottleneck Method, a method that we developed and that is capable of extracting the temporal information of

a signal such as a time series of sensor data. We introduced it in a previous paper [9]. This method was developed as an extension of the well-known Information Bottleneck Method to the analysis of time series [10]. Section 3 briefly explains the method and provides details on its implementation for the recognition of human-robot interaction styles. We then report on trials conducted in a school with seven children with autism which evaluated the potential impact of such an adaptive robot on the children’s play styles.

II. RELATED WORK

Related work in robot-assisted play for children with autism has shown that when playing with a robot (in contrast to a stuffed animal), children with autism tend to show more behaviours that are typically impaired in autism (e.g. eye contact) [11]. Earlier comparisons between a mobile robot and a toy truck have shown more engaging behaviour towards the autonomous robot [12; 13]. Moreover other studies highlighted the potential role of the robot as a social mediator for children with autism [14; 15; 16]. Most studies were conducted in task oriented settings, e.g. involving imitation [16] or chasing games [15] with reactive (remotely controlled or autonomous) robots. Besides, the role of the experimenter in robot-assisted play has been investigated, firstly by Robins et al. [17] and more recently by François et al. [18].

The current paper focuses on investigating the role of an adaptive robot in robot-assisted play. We investigate whether an adaptive robot, i.e. a robot that could adapt to each child’s play styles in real time would have a positive effect on the children’s play styles and guide them progressively towards more well balanced interaction styles.

III. THE CASCADED INFORMATION BOTTLENECK METHOD

A. Background: The Information Bottleneck Method

The Information Bottleneck Method [10] is a clustering method based on an information theoretic approach whose purpose is to extract the relevant information² 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: $\tilde{X} \leftarrow X \leftarrow Y$ where \tilde{X} is the variable that extracts information about Y through X .

This popular 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 an external distortion measure but directly through a variational principle implementing an information-theoretic formulation of sufficient statistics. 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

²In this context, the relevant information is defined as the information that the (accessible) signal $x \in \mathcal{X}$ provides about another (typically not directly accessible) signal $y \in \mathcal{Y}$.

minimizing the number of bits from the accessible signal X , i.e. maximizing the compression. The optimal assignment can be found by minimizing the functional


$$\mathcal{L}[p(\tilde{x}|x)] = I(\tilde{X}; X) - \beta I(\tilde{X}; Y) \quad (1)$$

$I(X; Y)$ stands for the mutual information between X and Y . For β and the cardinality of \tilde{X} fixed, an expression can be given which specifies implicitly the solution and leads to a fixed-point iteration. For the information bottleneck setting, the Kullback-Leibler divergence $D_{KL}(p(y|x)||p(y|\tilde{x}))$ replaces the distortion function from conventional rate-distortion theory.


The Agglomerative Information Bottleneck algorithm [19] makes the assumption that β tends to ∞ in the Lagrangian equation (Eq. 1). In this specific setting, 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 cardinality 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 for which $p(y|\tilde{x})$ has the smallest $D_{KL}(p(y|x)||p(y|\tilde{x}))$. The hard partition can be softened afterwards, with reverse annealing.

B. The Cascaded Information Bottleneck Method

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



2) *Extrapolation:* The Cascaded Information Bottleneck Method progressively extracts the relevant information from an input sample $X = [X_0, \dots, X_{S-1}]$ by a recall on the

successive components (X_0 for the first step of the cascade, (\tilde{X}_{s-1}, X_s) for step $s > 0$). Each bottleneck (we now discuss only $s > 0$, without loss of generality) is characterized by a probabilistic mapping $p(\tilde{x}_s | (\tilde{x}_{s-1}, x_s))$ which, for the present work using the agglomerative information bottleneck algorithm, is hard, i.e. above probability is 1 for exactly one value \tilde{x}_s of \tilde{X}_s and vanishes otherwise, i.e. it implements a hard mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ (note that the input (\tilde{x}_{s-1}, x_s) corresponds to the input x of the original information bottleneck method).

During the information bottleneck training process, for each step of the cascade successively ($s > 0$), the mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ is built. If, however, at a step s in the cascade a pair (\tilde{x}_{s-1}, x_s) never occurs during the training (we call this an *unseen pair*), the mapping $(\tilde{x}_{s-1}, x_s) \mapsto \tilde{x}_s$ will not be defined for the completed cascade. Upon processing of novel data, however, such a pair may be observed and in this case the cascade has no way to infer the following bottleneck state \tilde{x}_s , since there is no natural a priori correspondence of bottleneck states in successive bottlenecks.

For such cases, we therefore introduce an identification of successive bottleneck states which will provide us with a “default” continuation of a bottleneck state from step $s - 1$ to step s in the case of unseen pairs. Let $\tilde{\mathcal{X}}_{s-1}$ and $\tilde{\mathcal{X}}_s$ be the set of bottleneck states \tilde{x}_{s-1} and \tilde{x}_s , as well as $p(\tilde{x}_{s-1})$ and $p(\tilde{x}_s)$ their empirical probabilities. We consider one-to-one mappings r from $\tilde{\mathcal{X}}_{s-1}$ to $\tilde{\mathcal{X}}_s$ (which, for convenience, we call *permutations*). Each such permutation r provides an identification of successive bottleneck states. We define the informational cost of a permutation as

$$d_{(s-1,s)}(r) = - \sum_{\tilde{x}_{s-1} \in \tilde{\mathcal{X}}_{s-1}} p(\tilde{x}_{s-1}) \log \tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1}) | \tilde{X}_{s-1} = \tilde{x}_{s-1}) \quad (2)$$

Note that $\tilde{p}(\tilde{X}_s = r(\tilde{x}_{s-1}) | \tilde{X}_{s-1} = \tilde{x}_{s-1})$ is, for a given permutation r , the probability that the next state is $r(\tilde{x}_{s-1})$ knowing that the current state is \tilde{x}_{s-1} . The logarithm measures the unpredictability of the next state (i.e. the unpredictability of \tilde{X}_s given \tilde{x}_{s-1}). 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 ∞ .

To define a “default” continuation we now choose a permutation $R(s-1, s)$ that minimizes that unpredictability, weighted by the probability that the state \tilde{x}_{s-1} actually occurs. Note that per construction of the bottleneck cascade, one never has $p(\tilde{x}_{s-1}) = 0$.

$$R(s-1, s) = \arg \min_r d_{(s-1,s)}(r) \quad (3)$$

$R(s-1, s)$ defines now a “default” path between $\tilde{\mathcal{X}}_{s-1}$ and $\tilde{\mathcal{X}}_s$, and thus provides an extrapolation of the succeeding bottleneck state in the case of an unseen pair.

3) *Implementation:* The Cascaded Information Bottleneck Method has been evaluated with two different criteria of interaction, namely the gentleness and the frequency of the interaction in [9]. The criterion gentleness contains two classes, namely ‘gentle’ and ‘strong’ which correspond respectively to non-forceful and forceful tactile interaction. The

frequency of the interaction is categorised into four classes, defined by their typical periodicity of interaction: i) *very low* (S_0): the elapsed time between two tactile interactions is greater than 15 seconds; ii) *middle inferior* (S_1): the elapsed time between two tactile interactions is lower or equal to 15 seconds and greater than 5 seconds; iii) *middle superior* (S_2): the elapsed time between two tactile interactions is lower or equal to 5 seconds and greater than 1 second; iv) *very high* (S_3): the elapsed time between two tactile interactions is lower or equal to 1 second. S_1 and S_2 are considered here as well-balanced frequencies of interaction, while S_0 corresponds to a rare interaction and S_3 to a very intense interaction.

Two different cascades were built independently, one for each criterion of interaction. The gentleness corresponds to a short-term time scale event while the frequency corresponds to a mid-term time scale event (see Fig. 1 which provides the parameters for each cascade). The samples for the training of each 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

Fig. 1. Parameters for each cascade of bottlenecks.

In both cases, a sliding window proceeds on the sensor data time series. 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.

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

The method shows a sound recognition for both short-term and mid-term time scale events and involves only a very short delay for the recognition of short-term time scale events (0.17 seconds on average) [9]. Besides, the training process enables a structure to emerge over the cascade since the conditional entropy between the bottleneck states of two successive bottlenecks is globally decreasing over the cascade (Fig. 2). The Cascaded Information Bottleneck method is transparent and enables control over how much and what new information is taken at which step of the cascade. In particular, the extrapolation process enables to control the degrees of freedom of the system and prevent the cascade from over-learning.

In the next section, we present an application of the

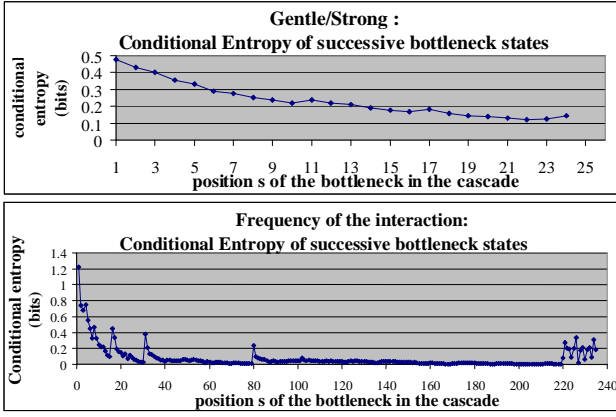


Fig. 2. Conditional entropy $H(\tilde{X}_{s+1}|\tilde{X}_s)$. $H(\tilde{X}_{s+1}|\tilde{X}_s)$ globally decreases over the cascade, pretty quickly, which suggests that a structure is progressively and rapidly emerging over the cascade: at the beginning of the cascade, a lot of new information is needed 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 which suggest that at these steps s , the input data X_s may influence a bit more in the choice of next equivalent state \tilde{X}_{s+1} . Note that the ones at the end of the cascade may reflect the importance of the last steps for distinguishing the classes S_0 and S_1 . 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).

Cascaded Information Bottleneck method for Robot-Assisted Play whereby the method is used to enable a robot to recognize in real time human-robot interaction styles and adapt to them accordingly. In this application the criteria of interaction detected are the gentleness and the frequency of the interaction and the cascades are the ones described in the previous paragraph (and detailed in [9]).

IV. APPLICATION: A REAL-TIME ADAPTIVE ROBOT FOR ROBOT-ASSISTED PLAY

A. The adaptive robot

In the context of this paper, 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 i) which sensor(s) is (are) activated (e.g. head sensor) and ii) the styles of interaction recognised. In contrast, by ‘reactive’ robot, we refer to a robot that can only react differently depending on which sensor is activated (e.g. head sensor or back sensor front), and which will not change its behaviour according to the interaction styles.

1) *Reward of well balanced interaction styles*: The adaptive mode relies on a reward basis for well-balanced interaction styles: the child should get a positive feedback from the robot when he/she plays in an appropriate style of interaction. The idea behind is that 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 e.g. too strong, the robot does not show any reaction. Moreover, the child should be encouraged to engage in the interaction if he/she is not engaged. Therefore, the robot should be both rewarding and engaging.

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 during long-term studies 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.

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 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, then 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 study by the classes S_1 and S_2 . Note, this model is generic and can be applied with different criteria of interactions. Fig. 3 presents the reward schema for the two criteria (gentleness and frequency) considered here.

Frequency			
S_3		Reward for Gentle	No Reward
S_2		Reward for Gentle + Reward for well-balanced Frequency	No Reward
S_1		Reward for Gentle + Reward for well-balanced Frequency	No Reward
S_0		Reward for Gentle	No Reward
No interaction	Engaging, proactive		
	No interaction	GENTLE	STRONG
			Gentleness

Fig. 3. Reward Schema for the two criteria of interaction.

2) *Architecture for Decision-Making Based on Interaction Styles*: The real-time recognition of the interaction styles uses the Cascaded Information Bottleneck Method. The small delay involved in the recognition process 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 pause, 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 exceeds a fixed predefined threshold then the final choice (i.e. the decision) is that the child’s behaviour towards the robot is recognized as a ‘strong interaction style’ and the child will not get a reaction from the robot to his/her stimulation. If below threshold, then the decision ‘gentle interaction style’ is made 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 stimulation³.

As for the evaluation of the child’s disengagement, 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) which is reflected by a null input vector for the frequency of the interaction.

B. Trials

1) *Participants*: Seven children with autism participated in the experiments which took place in a school for moderate learning difficulties in UK. The children had had the chance to play with the robot during several months beforehand and were familiar with both the robot and the experimenter. 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.

2) *Artifact*: The robot was the Aibo ERS-7. It behaved autonomously and operated either in adaptive or reactive mode. 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. 4.

3) *Procedures and Measures*:

Procedures: Each child participated in two sessions and the experiments involved one child at a time. Each session consisted of three successive steps⁴ (also called games or runs), each step being defined by the mode of the robot–

³Note 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.

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

Fig. 4. Mapping between the external tactile sensors of the robot and its behaviours. For child E, the walking has been removed and replaced by the robot’s wagging of the tail as part of the design by immersion.

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, see Fig. 5).

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

Fig. 5. Settings for the different children. Setting 1 corresponds to session 1 and setting 2 corresponds to session 2.

The 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 mode); the 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. The child could see the experimenter putting the sticker on the back 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 (this procedure was considered to help the children cope with different experimental conditions). 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.

During each game, the child could freely interact with the robot. Before the start 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),
- 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 (sensor data were

not collected at this stage), neither did she try to influence the child’s behaviour in any way. The experimenter did not take part in the child-robot interactions in order not to interfere with the purpose of this study which had to focus on dyadic, uninterrupted interactions between the child and the robot, in order to test the potential of an adaptive robot to influence children’s play styles.

Measures: The experiments were video recorded. The sensor 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 offline. For the criterion gentle/strong, we actually looked at the overall proportion of the 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. We then used statistical techniques for non-parametric statistics.

C. Results

1) Statistical analysis of the engagement in the interaction and the gentleness of the strokes:




Fig. 6. 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.




Fig. 7. Graph showing the relative engagement of the children in adaptive and reactive modes.

Engagement in the interaction: In order to study whether the adaptive robot may have a positive impact on the engagement of the children in play we do not consider

the specificity of the strokes, i.e. whether they are gentle or strong. Instead, for each child we are interested in the total number of sensors’ activations that we compare for adaptive and reactive robot modes.

For each child and for each mode, we count the total number of times the sensors were activated (each sensor⁵ activated counts as one activation), namely, $N(Reactive)$, for the reactive mode, and $N(Adaptive)$, for the adaptive mode; for each child, we analyse the relative ratio of each mode⁶, as follows:

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

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

The Wilcoxon test [20] is applied to the data from the seven children for the two following variables (Fig 7): $r(Adaptive)$, representing the adaptive mode, and $r(Reactive)$, representing the reactive mode. The test shows that there is a significant effect of the experimental conditions adaptive versus reactive (for $T = 1.000$, $p < 0.028$, with $N = 7$, Fig. 6). Thus, we can conclude that the children engage significantly more in the interaction during the adaptive mode..

Gentleness of the interaction: Here, we study the nature of the activation in terms of gentleness, i.e. whether an activation is gentle or strong. We therefore consider the percentage of strong strokes (also called strong activations) among the total number of sensor 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 sessions⁷ (Fig. 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

Fig. 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 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

⁵Here we look at the activation of any of the four continuous external sensors, i.e. the head sensor and the three back sensors.

⁶Some 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.

⁷Here we consider 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.

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) 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.

2) *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 the adaptive and reactive modes.

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:

$$\rho_{S_i}(r) = \frac{N_{S_i}(r)}{\sum_{\tilde{r} \in R} N_{S_i}(\tilde{r})} \quad (4)$$

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 $Av_m(\rho)$.

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 $N = 7$. 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). For class S_0 (respectively class S_1) there is no significant difference between the two experimental conditions (adaptive versus reactive), since, for $T = 5.000$ (respectively $T = 4.000$), $p < 0.128$ (respectively $p < 0.173$) with $N = 7$. However, the proportion of cases where $Av_{Adaptive}(\rho_{S_0}) > Av_{Reactive}(\rho_{S_0})$ (respectively $Av_{Adaptive}(\rho_{S_1}) > Av_{Reactive}(\rho_{S_1})$) 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 that the percentage of children for which there are more events related to S_0 (respectively S_1) in the adaptive mode than in the reactive mode deviates significantly from a fifty-fifty ratio. Concerning S_2 (respectively S_3) there is a significant effect of the experimental conditions Adaptive and Reactive since for $T = 1.000$ (respectively $T = 0.000$), $p < 0.028$ (respectively $p < 0.018$) with $N = 7$ (Fig. 9 and Fig. 10). Consequently, in the adaptive mode, there are significantly more events from class S_2 (respectively S_3) than in the reactive mode.

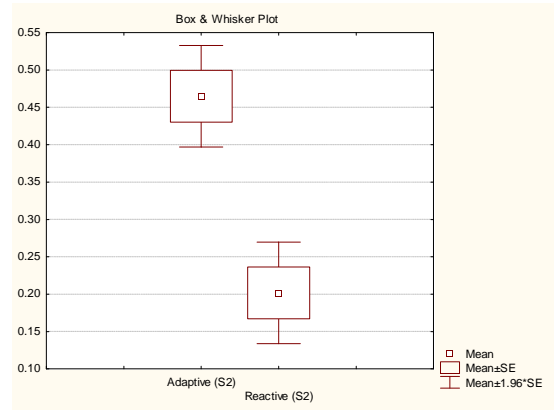


Fig. 9. Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_2 . The two variables are $Av_{Adaptive}(\rho_{S_2})$ and $Av_{Reactive}(\rho_{S_2})$.




Fig. 10. Mean, Standard Error of the Mean (SE) and Confidence Intervals for S_3 . The two variables are $Av_{Adaptive}(\rho_{S_3})$ and $Av_{Reactive}(\rho_{S_3})$.

V. DISCUSSION

This study has shown that a robot that can adapt to child-robot interaction tactile styles can influence positively the children's play styles. Firstly, the children engaged significantly more in the interaction when the robot was adaptive and significantly more children played more gently with the robot in the adaptive mode. Besides, the interactions were significantly richer and higher frequencies including in particular a range of well balanced frequencies were significantly more present in the adaptive mode. The introduction of an adaptive robot in robot-assisted play for children with autism which is able to adapt in real time to children's interaction styles is a novel contribution. This is both a technical and methodological step forward in robot-assisted play. On the one hand, the development of a new computational method that enables robots to recognize in real time human-robot interaction styles is a step forward towards socially adaptive robots. On the other hand, the evaluation of the potential of an adaptive robot in robot-assisted play expands the role that robots could potentially play in the context of autism therapy.

The study conducted here involved only a few children and was short-term. Future work should expand this study with different children with autism. Note, in our study the children were familiar with both the robot and the experimenter and the robot's behaviour mapping had been tailored according to each child's needs and abilities which is very important in the context of autism. Future work should enable the same for the new children involved. Besides, future work should consider possible long-term effects of such an adaptive robot. In particular, future work could expand the model of adaptation by focusing on a larger grid of criteria for the interaction styles: while the child progresses, the robot could increase the range of criteria the child should meet to get a reward. 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 could 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' which has been widely used in developmental robotics [21; 22; 23]. This technique, typically applied for a system learning motor skills, may be transposed to a social system, constituted here by the child and the robot: this social system is freezing some complexity in the interaction to learn more efficiently how to deal with interaction in general.

VI. CONCLUSION

This paper has presented an application of the previously introduced Cascaded Information Bottleneck Method for real-time recognition of Human-Robot interaction styles in the context of robot-assisted play for children with autism. Such an adaptive robot, which can detect the play styles of the children in real time and adapt to them accordingly, has been implemented and tested in the particular context of tactile interaction. The adaptation scheme rewards well balanced interaction styles and encourages the child to engage in the interaction if he/she is disengaged. The potential role of such an adaptive robot (compared to a reactive robot) in robot-assisted play has been evaluated with seven children with autism in a school and a statistical analysis showed that the adaptive mode influenced positively the play styles of the children in the following manner: 1) the children engaged significantly more in the interaction, 2) significantly more children played more gently with the robot, 3) the interactions were significantly richer and 4) the occurrence of higher frequencies, including a range of well balanced frequencies, was significantly increased. Future work should consider a wider study with different children with autism and investigate its long-term term impact. The Cascaded Information Bottleneck Method is generic and could potentially also be used in a variety of other applications in Robotics, Artificial Intelligence and Artificial Life. It is hoped that this work represents a step forward towards socially adaptive robots as well as robot-assisted play for children with autism.

REFERENCES

- [1] S. Powell, *Helping Children with Autism to Learn*. David Fulton Publish, 2000.
- [2] R. Jordan, *Autistic Spectrum Disorders, An Introductory Handbook for Practitioners*. David Fulton publishers Ltd, 1999.
- [3] "National autistic society," 2008, <http://www.nas.org.uk/>.
- [4] A. P. Association, *Diagnostic and Statistical Manual of Mental Disorders DSM-IV, fourth Edition*. American Psychiatric Publishing, 1994.
- [5] S. Baron-Cohen, *Mindblindness, an essay on Autism and Theory of Mind*. The MIT Press, new edition, 1997.
- [6] U. Fritz, *Autism: Explaining the Enigma*. Oxford, Blackwell, 1989.
- [7] K. Dautenhahn and I. Werry, "Towards interactive robots in autism therapy background, motivation and challenges," *Pragmatics and Cognition*, vol. 12(1), pp. 1–35, 2004.
- [8] J. Boucher, "Editorial: interventions with children with autism - methods based on play," *Child language teaching and therapy*, vol. 15, pp. 1–5, 1999.
- [9] D. François, D. Polani, and K. Dautenhahn, "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)*, 2008, pp. 353–359.
- [10] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. of the 37-th annual Allerton Conference on Communication, Control and Computing*, 1999, pp. 368–377.
- [11] C. M. Stanton, P. H. Kahn, R. L. Severson, J. H. Ruckert, and B. T. Gill, "Robotics animals might aid in the social development of children with autism," in *Proceedings of the 3rd International Conference on Human-Robot Interaction*, 2008.
- [12] I. Werry, K. Dautenhahn, and W. Harwin, "Evaluating the response of children with autism to a robot," in *Proc. RESNA 2001*, 2001.
- [13] I. Werry and K. Dautenhahn, *Human-Robot Interaction as a Model for Autism Therapy: An Experimental Study with Children with Autism*. In *Modeling Biology: Structures, Behaviors, Evolution*. Manfred Laubichler and Gerd B. Müller eds., Vienna Series in Theoretical Biology, MIT Press, 2007, pp. 283–299.
- [14] H. Kozima, C. Nakagawa, and Y. Yasuda, "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)*, 2005, pp. 341–346.
- [15] I. Werry, K. Dautenhahn, B. Ogden, and W. Harwin, "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)*, 2001, pp. 57–74.
- [16] B. Robins, K. Dautenhahn, and J. Dubowski, "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*, 2005, pp. 82–88.
- [17] B. Robins and K. Dautenhahn, "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)*, 2006, pp. 646–651.
- [18] D. François, S. Powell, and K. Dautenhahn, "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'*, 2009.
- [19] N. Slonim and N. Tishby, "Agglomerative information bottleneck," in *Proc. of Neural Information Processing Systems (NIPS 99)*, 1999, pp. 617–623.
- [20] S. Siegel and N. J. Castellan, *Nonparametric Statistics for the behavioral sciences*. McGraw-Hill International Editions, 1988.
- [21] L. Berthouze and M. Lungarella, "Motor skills acquisition under environmental perturbations; on the necessity of alternate freezing and freeing of degrees of freedom," *Adaptive Behaviour*, vol. 12(1), pp. 47 – 64, 2004.
- [22] M. Lungarella and L. Berthouze, "Adaptativity via alternate freezing and freeing of degrees of freedom," in *Proc. 9th Int. Conf. on Neural Information Processing*, 2002, pp. 482–487.
- [23] G. Gómez, M. Lungarella, P. E. Hotz, K. Matsushita, and R. Pfeifer, "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*, 2004, pp. 119–122.