Escaping from Children's Abuse of Social Robots

Dražen Brščić[1] Hiroyuki Kidokoro[1,2] Yoshitaka Suehiro[1,2] Takayuki Kanda[1,2]

[1] ATR Intelligent Robotics and Communication Laboratory, 2-2-2 Hikaridai, Seika-cho Keihanna Science City, Kyoto, Japan

[2] Osaka University, Graduate school of Information Science and Technology, 2-1 Yamadaoka, Suita, Osaka, Japan

{drazen, kidokoro, suehiro, kanda}@atr.jp

ABSTRACT
Social robots working in public space often stimulate children's curiosity. However, sometimes children also show abusive behavior toward robots. In our case studies, we observed in many cases that children persistently obstruct the robot's activity. Some actually abused the robot by saying bad things, and at times even kicking or punching the robot. We developed a statistical model of occurrence of children's abuse. Using this model together with a simulator of pedestrian behavior, we enabled the robot to predict the possibility of an abuse situation and escape before it happens. We demonstrated that with the model the robot successfully lowered the occurrence of abuse in a real shopping mall.

Categories and Subject Descriptors
H.5.2 [**Information Interfaces and Presentation**]: User Interfaces - *Interaction styles*; I.2.9 [**Artificial Intelligence**]: Robotics

General Terms
Design, Experimentation, Human Factors.

Keywords
human-robot interaction; children; robot abuse.

1. INTRODUCTION
Many robots are being developed for the use in open public environments. For instance, previous studies revealed robots can be successfully used in a museum [1, 2], supermarket [3], transit station [4], or urban sidewalk [5]. All these works reported that people exhibit great curiosity and actively interact with robots.

Similarly, in our experiments in public spaces we too observed that many people gathered around our robot. This includes children, who were usually with their parents and behaved well. However, we noticed that the situation often changed when there were not many people close to the robot and children were left alone to play with it. This sometimes lead to children showing anti-social behavior toward the robot, such as blocking its way, calling it names or even acting violently toward it.

Figure 1 shows two scenes where our robot was in trouble with children. They frequently stood in its way and stopped it from moving (Figure 1 left). Sometimes their behavior escalated further. They said bad words, punched, kicked, and/or pushed the robot (Figure 1 right). Even if the robot asked them to stop, its request was often ignored. They typically did not stop such misbehavior until they got bored or their parents stopped them.

Figure 1. Children's abusive behavior toward the robot.

We believe that occurrence of children's abusive behavior will be a real problem for social robots. Due to such behavior robots' execution of tasks would be hindered. Furthermore, children's abusive behaviors could make people uncomfortable [6] and perhaps also be problematic for children's healthy development [7, 8]. Our aim in this work was to study the children's abuse of robots and find a way to prevent or mitigate the problem.

2. RELATED WORKS
2.1 Robot Abuse
There are a few pioneering researches on the concept of abuse of robots. Bartneck and his colleagues first used the term "robot abuse" [9] and revealed that people show less concern about the abuse of robots than abuse of humans, yet they hesitate to destroy a robot when they perceive it intelligent [10]. There is a report that people felt bad and perceived empathy toward the robot when they saw a video in which a robot was tortured [6].

Salvini and his colleagues reported abusive behavior, or bullying behavior, toward their robots [11]. In their open-public demonstration, they observed that people approached their robot out of curiosity, but quite often escalated into aggressive behaviors, such as kicking, punching or slapping the robots. Such behavior only occurred when the robot was not attended by human operators.

Overall, none of previous studies have analyzed the pattern of occurrence of robot abuse, and did not propose a method to address such situation. The novelty of our study is that it is the first to attempt to find a method to moderate the abuse of robots.

2.2 Abuse in Other Interactions
Apart from robots, children can show abusive behavior in other social interactions, too. Perhaps the most prominent example is the bullying among children, which [11] defines as deliberate,

HRI '15, March 2-5, 2015, Portland, OR, USA

http://dx.doi.org/10.1145/2696454.2696468

repeated or long-term exposure to negative acts performed by a person or group of persons regarded of higher status or greater strength than the victim. Although similarities do exist, one important difference to the abuse of robots is that bullying among children is typically studied in long-term interactions, e.g. in schools, where the bully and victim know each other well.

Abusive behavior has been reported also in human-animal interaction research. Children sometimes engage in animal abuse, defined as non-accidental, socially unacceptable behavior that causes pain, suffering or distress and occasionally also death of an animal [13]. One view is to see animal abuse as an impulsive act without instrumental benefit, and thus consider it as an early symptom of conduct-disorder in children. It was reported that 25% of conduct-disorder children exhibited animal cruelty [8]. Research also shows that children who do animal abuse often also engage in bullying [7]. On the other hand, Arluke [14] argued that animal abuse has various other reasons, such as play, a form of knowledge creation, or a group activity with playmates. In any case, a concern remains that there is a link between animal abuse and violence toward humans, and that it is therefore better to prevent it.

In this work we start from the premise that robot abuse should also be avoided, and consider the modeling and robot planning to achieve this.

2.3 Simulation-based Planning in HRI

Human behavior is complex and it is hard to accurately predict what people would do in the future. This is a primary source of difficulty for using planning in robot interactions with people.

Recent studies started to use simulation in planning in human-robot interaction. For instance, Hoffman and Breazeal [15] conducted a study revealing that anticipatory computation enables fluent interaction, and further considered that such similarity in perception would work as a perceptual simulation. In a different work [16], a simulation of two people's side-by-side walking was made and the model for simulated agent was used for anticipatory computation.

Regarding the interaction with pedestrians, previous studies started to investigate the way to predict or anticipate people's behavior, and some of them used pedestrian simulation for planning. For instance, Bennewitz and her colleagues developed a model learnt from people's trajectories and used it to plan to avoid collisions with people [17]. Henry et al., trained a collision avoiding algorithm using a pedestrian simulator [18]. Garrell and Sanfeliu used a pedestrian simulation to test a robot's capability for handling groups of people [19], and used a pedestrian simulation to compute better position for a robot to avoid collisions [20]. In [21] we used a pedestrian simulator to predict the occurrence of crowding around a robot and used a planner to choose a path to minimize dissatisfaction to people passing by.

This study follows the paradigm of simulation-based planning. Nevertheless, the originality of the study is in that it successfully simulates occurrences of children's abuse of the robot, which are complex emergent phenomena. Further, this simulation is used for planning and avoidance of the robot-abuse problem.

3. DATA COLLECTION

3.1 Environment and Infrastructure

We conducted our study in a part of the ATC shopping mall in Osaka, Japan (Figure 2), where we set up a sensing infrastructure, which allowed us to track the people that are inside the area. The

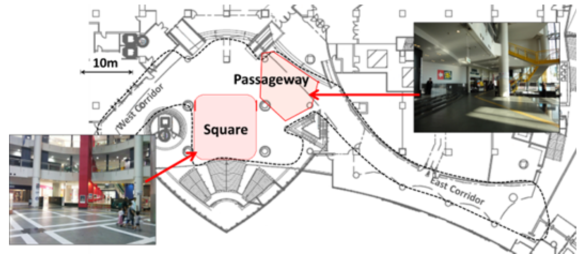

Figure 2. Environment where we conducted the field study.

observed area consisted of a central square, a long corridor leading toward east and a shorter corridor connecting to the west. The main flow of people went through the east corridor, where the density was largest. On the other hand, the square was typically less crowded.

The tracking system consisted of in total 49 3-D range sensors of different types, arranged such that they cover the area of interest. The total covered area was around 900 m2 (inside the dashed line in the figure). We used the tracking algorithm proposed in [22], which allowed us to track the position of all the people in the area at 20 Hz with accuracy of around 30 cm.

We used the human-like robot Robovie2 (Figure 1). The robot had a microphone to record the people's speech, and two laser range sensors which were used for safety. During autonomous operation, the range sensors were also used for localizing the robot using a particle filter on an occupancy grid map [23].

3.2 Setting

We let the robot patrol in two areas where different flow of people was observed: the square, and the area between the east corridor and square, which we refer to as passageway (see Figure 2). The patrol task consisted in repeated moving in straight lines between several fixed way points. In both areas we performed the data collection on a weekend for 2 hours. During the data collection all operations of the robot were controlled remotely by an operator.

When a person stopped in front of the robot, the operator stopped the robot and started the robot's conversation. The robot asked the person to step aside by saying: "I am Robovie. I am now patrolling, please let me go through." If the person stayed for 3 seconds, the robot said: "I wish to go through, could you please open the way?" After 3 seconds if the person still stood in the way, the robot turned toward the previous waypoint, and tried to move in the opposite direction. In case its way was blocked again, it started the above process from the beginning. When the person moved aside, it said "Thank you", and reassumed the patrolling task.

We collected videos and people's trajectories using the tracking system, and both of them were used to label all the instances where a person obstructed the robot.

3.3 Definition of Robot Abuse

In order to do the labelling we needed a definition of robot abuse, so we made a preliminary analysis of the collected data. When the robot asked to open the way, people typically followed the request; however, there were also a number of cases when they ignored or refused. Most people only did so once or twice, maybe to see the robot's reaction and test its capability. However some children continued this behavior even after the robot requested multiple times, intentionally obstructing the robot. In some cases their behavior gradually escalated. We observed three types of abuse, examples of which we illustrate next.

Type 1: Persistent obstruction. In many instances children persistently obstructed the robot. Figure 3 shows an example scene. At the beginning, the child would step aside after the robot's request, but then would quickly come back in front of the robot. After a while, she started ignoring the requests and just kept standing in front of the robot. Eventually, she also started to verbally express her intention of obstruction by saying “I won’t” or “No-no” when requested to move. At times other children also joined her. The whole scene lasted 19 minutes, until the girl’s mother came and made her leave.

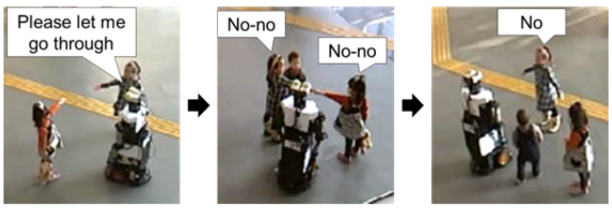

Figure 3. Persistent obstruction.

Type 2: Offensive utterances. In one example, four children surrounded the robot and ignored its repeated requests. After the fifth request, one child approached the robot and said a series of bad words to the robot. He said “You idiot” 8 times, “I won’t” 2 times, and “Go away” once. In addition, he agitated other children to be aggressive toward the robot saying “Let’s surround it.”

Type 3: Violence. Similar to the previous cases, at the beginning usually one child refused to let the robot go even after several requests, and soon other children gathered and after a while violent behavior occurred. One boy bent the robot’s neck (Figure 4a), and another boy hit its head with a plastic bottle (Figure 4b). We observed other cases of violence outside of the data collection, too. For example, one boy first hit the robot’s head with his hand saying “Go away”, and then threw a soccer ball onto the robot’s head (Figure 4c). In another case, 3 boys started hitting the robot with plastic bottles. This gradually escalated until they started hitting as strong as they could and throwing the bottles on the robot (Figure 4d). There was no hardware damage in any of these situations which could suggest that the children’s violent actions were not meant for seriously breaking the robot.

Although there is diversity in the seriousness of abuse, common to all above cases is that children intentionally continued to obstruct the robot’s task beyond curiosity. Our research goal was to prevent abuse which hinders the robot from performing its tasks, so all of above cases are certainly beyond the threshold to be judged as abuse. Although we introduced three cases, in this study we do not distinguish between them and treat all of them as abusive interaction. Informed by these observations, we proposed the following definition of abuse towards the robot:

> ***Robot abuse****:* Persistent offensive action, either verbal or non-verbal, or physical violence that violates the robot’s role or its human-like (or animal-like) nature.

(a) Bending the neck

(b) Hitting with plastic bottle

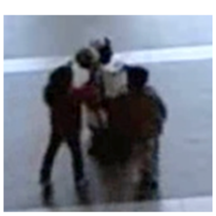

(c) Hitting with ball

(d) Throwing a plastic bottle

Figure 4. Children’s violence toward the robot.

We also noted that it was always children who abused the robot. Adults were rarely present and not at all involved in the abuse, so in most cases the robot had to deal directly with children. Moreover, the abuse occurred more on the square where fewer people are present, rather than in the passageway where people are continuously passing.

3.4 Data Coding

For each instance of interaction between people and the robot, a human coder examined the interaction from the start until the end and decided whether there was any abuse toward the robot.

Due to the human-likeness of the robot and its behavior, people have a tendency to anthropomorphize the robot. The coder who watched the videos of interaction judged whether a child’s action was as a case of “robot abuse” or not, based on the definition given above. For difficult to decide situations the coders were told to judge as abuse situations the cases which they themselves would dislike, or be irritated or fed up with the child’s behavior. For instance, the above type 1 example was judged as abuse because the robot was not able to function in its patrolling role due to children’s persistent obstructive action. In the type 2 it was due to the anthropomorphizing of the robot and the coder’s belief that a person would dislike such a situation.

For validating the results, a second coder performed the same task for randomly-selected 10% of the interaction instances extracted by the first coder. A good matching between the results of the coders was obtained (Cohen’s kappa .714). For all cases of abusive interaction we also noted the number of children involved and whether their parents were present. In total 185 interactions were labelled, out of which 19 were judged as abuse cases.

From the coded dataset we obtained the information whether there was an abuse, how long the interaction lasted, the number of involved children, presence of parents, and location. A summary of the data is given in Table 1. In addition, we analyzed the statistics of the cases where abuse did and did not occur, shown in Table 2.

Table 1. Statistics of interaction (SD values in brackets).

	probability of abuse	interaction time [s]	pedestrian density [per./min]	other children [child/15s]	parents present [%]
Passage-way	0.085 (±0.0104)	45.2 (±10.70)	40.6 (±1.17)	0.78 (±0.074)	54.1 (±5.4)
Square	0.139 (±0.0110)	72.6 (±10.60)	12.6 (±0.28)	1.52 (±0.122)	50.0 (±5.0)

Table 2. Comparison of abuse and no abuse cases (SD values in brackets).

	interaction time [s]	number of other children	parents present [%]
Abuse	238.3 (±49.0)	4.94 (±0.89)	21.1 (±9.6)
No abuse	39.6 (±4.2)	2.12 (±0.12)	55.4 (±3.9)

4. MODEL OF CHILDREN'S BEHAVIOR

4.1 Overview

We developed a statistical model, which provides an estimate of the probability that a child who interacts with the robot will cause abuse. Figure 5 illustrates the model. It consists of two sub-models. For simulations (section 4.5) we needed to decide how long a child will stay with the robot, and therefore the modeling of the interaction time is done separately. The model of interaction time describes how long the child will interact with the robot. The model of occurrence of abuse computes the probability that the child will engage in abusive behavior. These sub-models are explained in the following subsections.

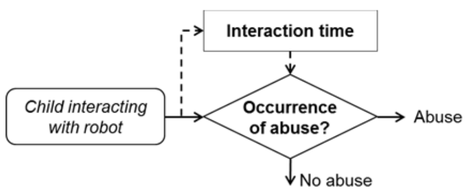

Figure 5. Estimation of occurrence of abuse toward robot.

4.2 Model of Interaction Time

We analyzed the statistics of interaction times, obtained from the data collection reported in section 3. We found out that the shape of the distribution of the interaction times is close to an exponential function. Therefore we modeled it using an exponential distribution:

$$p_{int}(t,c) = \lambda(c) \cdot e^{-\lambda(c)t}, \quad (1)$$

where $p_{int}(t,c)$ is the probability child c will stop interacting after time t, and λ is the parameter of the distribution. Figure 6 shows the histogram of interaction times and a fitted exponential function with parameter $\lambda = 0.032$.

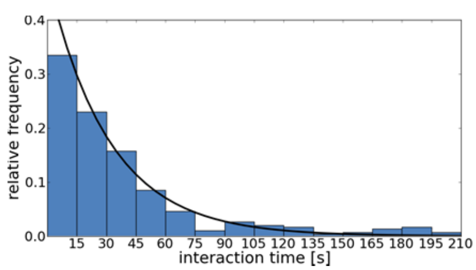

Figure 6. Histogram of interaction times.

Results in Table 1 show that interaction times differ across locations, suggesting that there is an influence of factors such as density on the interaction time; we modeled such influence into the λ. We started with all possibly relevant variables:

a) ***Pedestrian density*** *(density)* – the density of people inside the area where interaction happened. This was defined as the expected number of people that come within 10 meters from the robot during 1 minute, and the values were obtained from 3 hours of pedestrian trajectories collected when there was no robot in the environment.

b) ***Number of other children*** *(child)* – the average number of other children who simultaneously interacted.

c) ***Presence of parents*** *(parent)* – 1 if the child's parents are in the vicinity (set to be within 3m of the child), 0 otherwise.

After fitting to the data we found that the number of children had little influence. Finally, we obtained the following model for λ:

$$\lambda(c) = a_{density} \cdot density + a_{parent} \cdot parent + b_1' \quad (2)$$

Fitting the data gave the coefficients: $a_{density} = 0.000219$, $a_{parent} = 0.0116$, $b_1' = 0.0725$. These values tell us that the interaction times are shorter in areas with higher person density, and that the influence of parents is large, namely that the child is more likely to stop the interaction when a parent is present.

4.3 Model of Occurrence of Abuse

All variables noted in Table 2 appear to be relevant in describing the possibility of abuse, with abuse more likely occurring in cases when children interact longer, with more children around, and without having their parents nearby.

We wished not only to have an estimate if abuse occurs or not, but also to have a value for the probability of occurrence. Therefore we applied a logistic regression model to the result of the above analysis. In addition to the three variables used in the subsection 4.2, we included the following variable:

d) ***Interaction time*** *(interact)* – representing the time from the start of the interaction, i.e. after the child approached the robot and stayed within a defined interaction distance (experimentally set to 0.89 m) from it.

We started with all variables and fit the data. We found that $parent$ and $density$ had very little influence, so we removed them from the model. An interpretation of this is that their primary influence on the occurrence of abuse is indirect through the interaction time (i.e. if with a parent or in a space with more people around, the child stays shorter around the robot). Finally, we obtained the following model:

$$p_{abuse}(t,c) = \frac{1}{1+\exp(-(a_{child} \cdot child + a_{interact} \cdot interact + b_2))} \quad (3)$$

where $a_{child} = 0.5935$, $a_{interact} = 0.003218$, $b_2 = -3.084$.

In summary, the modeling shows that probability of abuse increases with the number of children around the robot and the interaction time. For example, short 30 s interactions with 2 children around the robot (i.e. $child = 1$) give $p_{abuse} = 0.048$. With same interaction length and the total number of children equal to 4 the probability becomes 0.23, whereas with 2 children and a longer interaction time of 200 seconds, the probability increases to 0.136.

4.4 Model of the Effect of Abuse on the Robot

The performance of the robot can be affected by the abuse. For planning purposes (section 5), we wish to model the net effect of possible abuse on the robot's task. We considered the influence on the navigation – the effect of abuse on the robot's average speed.

This was modeled in the following way. At time t, the robot's speed $v_r(t)$ was defined to be proportional to the sum of the abuse probabilities for all children:

$$v_r(t) = a_{abuse} \cdot \sum_{s \in N} p_{abuse}(t, c) + v_0 \quad (4)$$

We fit the model to the data. By minimizing least square mean error, we obtained the coefficient a_{abuse} = -115 and the intercept v_0 = 297 mm/s. In other words, the speed of the robot decreases with the increased probability of abuse.

4.5 Simulation

We developed a simulation in which the movement of pedestrians is reproduced, using the framework reported in [21]. It simulates the movement of pedestrians in the environment, where some of them will approach the robot, and a number of them will also abuse it.

The simulated pedestrians are generated based on the *model of pedestrian flow*, which describes the statistical nature of the movement of pedestrians, Figure 7. The model defines the flow of the simulation: the creation of new simulated pedestrians, generation of the trajectories each pedestrian will traverse, and finally the type of interaction they will have with the robot. This follows the framework in [21], with a few additions explained below. All parameters of the model are calibrated using data from the tracking system described in 3.1.

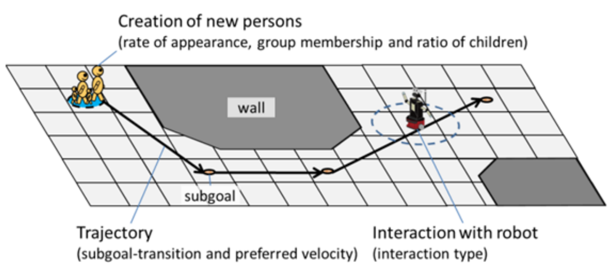

Figure 7. Basic elements of model of pedestrian flow.

Creation of new people: Apart from the statistics describing the time and place of creation of new simulated pedestrians and the distribution of groups, we also added the *ratio of children in groups* – for each group size, the distribution of the number of children is recorded and afterward used in simulation.

Trajectories: After being created, the simulated pedestrians move along the path defined by the "subgoal-transition model" (based on the technique reported in [24]) and according to the preferred speed statistic. All people in the same group have the same subgoals and preferred velocity.

Interaction with robot: Interaction around the robot is simulated with the following models.

- *Interaction type*: In a real environment some of pedestrians approach the robot, while others do not. We adopted the model reported in [25] to reproduce this behavior. If the *interaction type* is "stop to interact", the pedestrian changes his/her course to approach the robot when the robot is visible, stays for a while, and then continues moving toward his/her original subgoals. Ratio of the "stop to interact" behavior in the collected data was 35.48% for children and 4.21% for adults.
- *Parent-child relationship*: We extended the original interaction type model [25] to include the parent-child relationship, since it is important in this study. When a child approaches the robot to interact, the parent will typically watch from the side until the child to end the interaction (Figure 8 left). We model this in the following way. The parent and child are assigned interaction types which can be different. When the child's interaction type is "stop to interact", after the simulated child enters the interaction distance (0.89 m) from the robot, the simulated parent will stop walking when the child becomes "out of sight" (the child is more than 90 degrees from the parent's direction of movement), and wait for the child to finish the interaction.

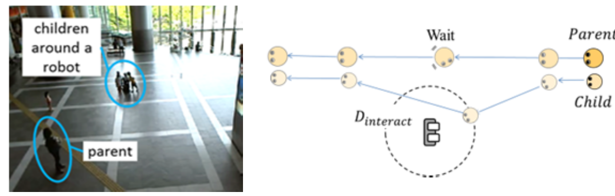

Figure 8. Simulation of parent's behavior when child stops to interact.

Overall, in the simulation, each pedestrian's position is periodically updated with above models. When simulated children interact with the robot, the models in section 4 are used to determine the probability they will abuse the robot, and the net effect on the simulated robot's speed.

4.6 Evaluation of the Simulation

To quantitatively test the validity of the model, we used the simulation to reproduce the state of the environment during the data collection. We ran 20 2-hour simulations for both the passageway and square area. The model of pedestrian flow was generated for each trajectory data obtained from data collection. The simulated robot moved between the same fixed locations as it did in the data collection.

Table 3 shows the result of the simulation. The statistics are obtained as average from all simulation runs. When compared with the statistics from real data from Table 1 we see that for all values the difference is within one standard deviation, showing that the simulation adequately reproduces the real world behavior.

Table 3. Simulation results.

	probability of abuse	interaction time [s]	flow density [per./min]	other children [child/15s]	parents present [%]
Passage-way	0.094	44.6.	41.6	0.78	55.6
Square	0.134	69.6	12.6	1.43	51.5

5. BEHAVIOR PLANNING SYSTEM

5.1 Architecture

We developed a prototype system, in which the developed simulation of children's abuse behavior is used for robot behavior planning. The robot's task was to navigate around the environment and travel as long as possible within limited amount of time, i.e. be as efficient as possible. It is a simple task, but can be easily extended to various future tasks, such as patrolling, searching (e.g. to find a child who got lost), guiding users to destinations, and carrying. All of these services require efficient navigations.

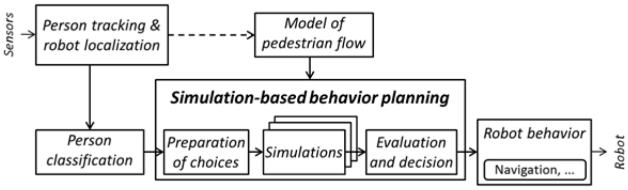

Figure 9. Architecture of the behavior planning.

Figure 9 illustrates the architecture of the developed system. The central component is the *simulation-based behavior planning* module, in which pedestrian simulations (section 4.5) were used to predict the near future behavior of pedestrians.

5.2 Person Classification

The people's trajectories were observed using the tracking system described in section 3.1. Based on these data the following attributes were extracted for the use in simulation:

- *Group membership*: Following the algorithm in [26], people who stay within 3 m for 60 % of time together were estimated as members of the same group.
- *Children/adult*: This was estimated from the observed person's height. We used a simple threshold (1.38 m) for distinction. When an adult and a child belonged to the same group, the adult was marked as the parent of the child.
- *Interaction type*: This attribute was initially probabilistically assigned according to the observed distribution, however once the person passed by the robot, it was updated based on his/her real behavior. If the person stopped close to the robot, the type was estimated as "stop to interact".

5.3 Simulation-based Behavior Planning

First a number of possible robot's behavior choices were prepared, and then multiple simulations were run. Based on the outcome of simulation, the most appropriate robot's behavior was selected.

Preparation of choices: The system prepared a series of possible destinations $d_1, d_2, \cdots, d_{|D|}$ where the robot can move to next. Three types of choices were prepared as follows: (1) the current destination; (2) 10 randomly chosen locations; (3) location nearby parents of the children who are currently around the robot. These choices were used because, according to the analysis in section 4, both destinations with higher person densities (part of type 2) and close to parents (type 3) can be effective for escaping from abuse.

Simulation: For each choice of destination d_i the robot ran multiple pedestrian simulations to predict the future development of the situation. The simulations were initialized on the current situation (locations and attributes of currently observed pedestrians). Finally, the average speeds of the robot during the simulations $(v_1^{d_i}, v_2^{d_i}, \cdots, v_{|Q|}^{d_i})$ were saved. The duration of each simulation was empirically set to 60 seconds, and the number of simulations for each destination choice to $|Q| = 50$.

Evaluation and decision: For each destination d_i the robot's efficiency (the expected travel speed V^{d_i}) was computed as:

$$V^{d_i} = 1/|Q| \cdot \sum_{q=1}^{|Q|} v_q^{d_i} \quad (5)$$

Finally, the destination that yields the highest V^{d_i} was chosen as the target destination. This planning was done every second. To prevent too frequent switching between the destinations, which causes loss of time by continuously rotating towards different destinations, the current goal's efficiency value was multiplied by a weight factor, empirically set to 1.25.

5.4 Robot Behavior

The *robot behavior* module received the destination from the *simulation-based behavior planning* module, and navigated the robot toward the given destination. Unless its frontal direction is blocked, the robot moved straight toward the destination, and it stopped its locomotion if a person stops in front of the robot.

When the robot stopped, its behavior followed the one explained in the section 3.2, except that in this case the behavior was automated. That is, the detection was based on the location of person provided by the person tracking: when there was a person in front of it within 50 cm from the center and within 90 degrees from the forward direction, the robot stopped and asked the person to move aside. The same wording and sequence of actions as in the data collection (section 3.2) was used.

6. FIELD EVALUATION

6.1 Hypothesis and Predictions

If the simulation-based planning really enables the robot to anticipate the future situation, it should be able to lower the chance of occurrence of abuse. Since abuse results in the robot often having to wait and rotate in space which can last for a long duration of time, the navigation should also be more efficient. We made the following predictions:

Prediction 1: There will be fewer occurrences of abuse and when the robot uses the proposed simulation-based planning method than when it simply plans to maximize its efficiency without taking abuse into account.

Prediction 2: The robot navigation will also be more efficient, i.e., the resulting average speed of the robot will be higher.

6.2 Conditions

We compared the following two conditions:

Simulation-based planning: The robot choses the destination as described in section 5.

Maximum efficiency planning: The alternative strategy ignores the influence of abuse and only maximizes the travel efficiency. The robot uses a simple planning algorithm, where after reaching a goal it chooses the next destination as the one which will result in maximum travel velocity, assuming that there are no people interacting with the robot.

6.3 Procedure

We conducted the evaluation in the shopping mall environment described in section 3.1. To have a fair evaluation of both conditions, we prepared a number of paired time-slots to which the conditions were randomly assigned. This is done to make the environmental conditions approximately equivalent across conditions. We conducted the evaluation experiment on a weekend, for in total 40 minutes for each condition. In both conditions the robot moved fully autonomously.

6.4 Measurements

The evaluation criteria were defined as follows:

Occurrence of abuse: the number of occurrences of abuse (judged based on the criteria reported in section 3.4)

Travel efficiency: the average travel velocity of the robot during the experiments.

6.5 Evaluation Results

6.5.1 Overall trend

In the condition of maximum efficiency planning, the robot had a tendency to choose long straight paths, as they result in higher average speeds. Because of that it often passed through areas with fewer people where, similar to what we described in the section 3, several cases of abuse happened. In one example situation, while the robot was traversing the square a child stood in front of it with arms open wide (Figure 10a left). The child's actions escalated and just when the child's mother was coming back they turned into physical abuse (Figure 10a right). The mother then scolded the child and moved it away from the robot.

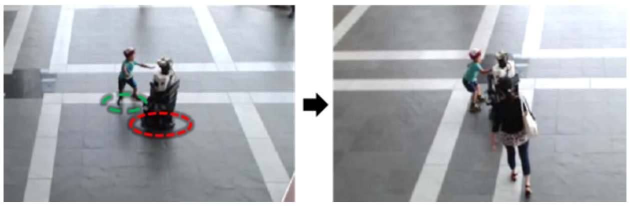

(a) Maximum efficiency planning condition: occurrence of abuse.

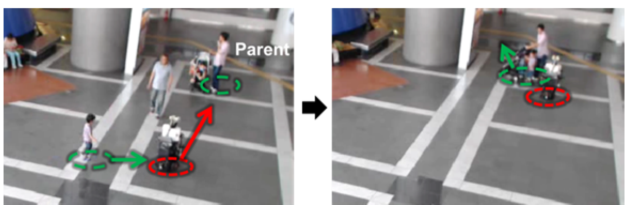

(b) Simulation-based planning condition: escaping from abuse.

Figure 10. Example situations during evaluation.

In the simulation-based planning condition, when there were no children in the environment the robot crossed the square from side to side, similar to the maximum efficiency planning case. When a child appeared and tried to approach the robot, the robot would move either toward an area with high person flow or toward the parent of the child. We observed that in that case the parent would take the child away, or in the case of the high flow area, the time of interaction tended to be shorter. For example, Figure 10b shows a situation where after detecting an approaching child, the robot changed its course toward a parent standing on the side of the square, after which the parent left with the child.

6.5.2 Qualitative evaluation

During the evaluation experiment 4064 people were detected in the tracking area (2147 in the maximum efficiency planning and 1917 in the simulation-based planning condition). Out of them, 57 children in total approached the robot (31 and 26, for the maximum efficiency and simulation-based planning, respectively). Two coders labeled the interaction cases for abuse of robot. Good matching between coders was obtained (Cohen's kappa .818).

The result of the evaluation is shown in Table 4. A Chi-squared test on the result which showed there is a statistically significant difference between the ratios of abuse cases in the two conditions ($\chi^2 = 4.11$, $df = 1$, $p < .05$). This shows that *Prediction 1* is supported and that the simulation-based planning is effective in restricting and avoiding the occurrence of abuse toward the robot.

We also evaluated the robot's travel velocity. The average velocity during the experiments in the maximum efficiency planning condition was 260.7±4.96 mm/s, and in the simulation-based planning condition 321.3±5.25 mm/s. There is a statistically significant difference between the results ($t = 8.402$, $df = 4353.4$, $p < .001$), which shows that using simulation-based planning results in higher speeds. We conclude that *Prediction 2* is supported.

Table 4. Occurrence of abuse.

	Abuse	No abuse
Maximum efficiency planning	7	24
Simulation-based planning	1	25

7. DISCUSSION

7.1 Can the Robot "Overcome" the Abuse?

In this work the robot's strategy to prevent abuse was to "escape", i.e. move to a location where it is less likely abuse will occur. One could ask why the robot cannot overcome the abuse. In our preliminary trials we have tried several approaches, but we found that it is very difficult for the robot to persuade children not to abuse it. For example, we changed the robot's wordings in many ways, using strong words, emotional or polite expressions, but none of them were successful. One partially successful strategy was the robot 'physically' pushing children. When its way was blocked, it would just try to keep going and behave as if it will collide into children and force its way through (under careful monitoring from a human operator). We observed that children at first accepted the robot's requests and obeyed them; but, very soon they learned that they are stronger than the robot so they can win if they push, and also that they can stop it by pushing the bumper switch (attached on the robot for safety). After realizing that, they just continued with the abusive behavior. Obviously having a stronger robot would present a problem for safety and social acceptance so dealing with such abusive situations remains difficult.

7.2 Robot Ethics

An important ethical question implicit in this work is: what type of behavior toward the robot is appropriate? The definition of robot abuse that we use is based on an anthropomorphized view of the robot, and we showed that if we use such a definition, abuse of robots is certainly a real problem in public spaces. Whether this definition is justified, and if not how should we treat robots – these ethical questions remain open.

7.3 Generalizability and Limitations

We only modeled the situation where the robot is just navigating around. For the future use of social robots, it will also be necessary to consider situations where the robot engages in other tasks, such as conversation. To do so, some model parameters would need to be adjusted. Depending on factors such as the robot's speed, its reactions, and type of tasks, the expected time people spend with the robot would differ significantly. Moreover, in our system we inferred the occurrence of abuse from observed people's positions and their interaction time; in other cases it might be necessary to directly detect abusive behavior.

In this study, we let the robot escape from the occurrence of abuse only by moving to a different position. Of course, there could be other solutions, like calling a guard if abuse happens.

8. CONCLUSION

We found that our robot is affected by children's abusive behavior. We analyzed how the abuse occurs, and modelled the statistical nature of the occurrences. Only children caused abuse. If their parents were not close and if there were fewer pedestrians around, children tended to stay longer around the robot. When they stayed longer, and especially if more children were present, they had a greater tendency to cause abuse to the robot. This emergent nature of abuse was simulated in a pedestrian simulator. We found that the simulation can reproduce reasonably well the situations that happen in the actual environment. Based on the simulation, we developed a planning technique for avoiding the occurrence of abuse. Our field trial demonstrated the efficacy of this approach.

9. ACKNOWLEDGMENTS

This work was supported by JST, CREST. We thank the staff of the ATC shopping mall for their support.

10. REFERENCES

[1] Thrun, S., Bennewitz, M., Burgard, W., Cremers, A.B., Dellaert, F., Fox, D., Hahnel, D., Rosenberg, C., Roy, N., Schulte, J., Schulz, D., "MINERVA: A Second-Generation Museum Tour-Guide Robot", *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp.1999-2005, 1999

[2] Siegwart, R., Arras, K.O., Bouabdallah, S., Burnier, D., Froidevaux, G., Greppin, X., Jensen, B., Lorotte, A., Mayor, L., Meisser, M., Philippsen, R., Piguet, R., Ramel, G., Terrien, G., Tomatis, N., "Robox at Expo.02: A Large Scale Installation of Personal Robots", *Robotics and Autonomous Systems*, vol. 42, pp.203-222, 2003

[3] Gross, H.-M., Boehme, H.-J., Schroeter, C., Mueller, S., Koenig, A., Martin, Ch., Merten, M., Bley, A., "ShopBot: progress in developing an interactive mobile shopping assistant for everyday use", *Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC)*, pp.3471-3478, 2008

[4] Svenstrup, M., Bak, T., Maler, O., Andersen, H.J., Jensen, O.B., "Pilot Study of Person Robot Interaction in a Public Transit Space", *Research and Education in Robotics - EUROBOT 2008*, vol.33, pp.96-106, 2009

[5] Weiss, A., Bernhaupt, R., Tscheligi, M., Wollherr, D., Kühnlenz, K., Buss, M., "A Methodological Variation for Acceptance Evaluation of Human-Robot Interaction in Public Places", *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication (RO-MAN)*, pp.713-718, 2008

[6] Rosenthal-von der Pütten, A.M., Krämer, N.C. Hoffmann, L., Sobieraj, S., Eimler, S.C., "An Experimental Study on Emotional Reactions Towards a Robot", *Int. J. of Social Robotics*, vol.5, pp.17-34, 2013

[7] Gullone, E. and Robertson, N., "The relationship between bullying and animal abuse behaviors in adolescents: The importance of witnessing animal abuse", *J. of Applied Developmental Psychology*, vol.29, pp.371-379, 2008

[8] Miller, C. "Childhood animal cruelty and interpersonal violence", *Clinical Psych. Review*, vol.21, pp.735-749, 2001

[9] Bartneck, C., Rosalia, C., Menges, R. and Deckers, I., "Robot abuse–a limitation of the media equation", *Proc. Interact 2005 Workshop on Agent Abuse,* 2005

[10] Bartneck, C. and Hu, J., "Exploring the Abuse of Robots", *Interaction Studies - Social Behaviour and Communication in Biological and Artificial Systems*, vol.9, pp.415-433, 2008

[11] Salvini, P., Ciaravella, G., Yu, W., Ferri, G., Manzi, A., Mazzolai, B., Laschi, C., Oh, S.-R., Dario, P., "How safe are service robots in urban environments? Bullying a Robot", *Proc. IEEE Int. Symp. on Robot and Human Interactive Communication (RO-MAN)*, pp.1-7, 2010

[12] Olweus, D., *Bullying at school: What we know and what we can do*, Oxford: Blackwell Publishers, 1993

[13] Ascione, F. R. and Shapiro, K., "People and Animals, Kindness and Cruelty: Research Directions and Policy Implications", *J. of Social Issues*, vol.65, pp.569-587, 2009

[14] Arluke, A., "Animal Abuse as Dirty Play", *Symbolic Interaction*, vol.25, pp.405-430, 2002Hoffman, G. and Breazeal, C., "Effects of anticipatory perceptual simulation on practiced human-robot tasks", *Autonomous Robots*, vol.28, pp.403-423, 2010

[15] Morales Saiki, L.Y., Satake, S., Huq, R., Glas, D., Kanda, T., Hagita, N., "How Do People Walk Side-By-Side? -Using A Computational Model Of Human Behavior For A Social Robot", *Proc. ACM/IEEE Int. Conf. on Human Robot Interaction (HRI)*, pp.301-308, 2012

[16] Bennewitz, M., Burgard, W., Cielniak, G. and Thrun, S., "Learning Motion Patterns of People for Compliant Robot Motion", *Int. J. of Robotics Research*, vol.24, pp.31-48, 2005

[17] Henry, P., Vollmer, C., Ferris, B. and Fox, D., "Learning to Navigate Through Crowded Environments", *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp.981-986, 2010

[18] Garrell, A. and Sanfeliu, A., "Model validation: robot behavior in people guidance mission using DTM model and estimation of human motion behavior", *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp.5836-5841, 2010

[19] Ferrer, G., Garrell, A. and Sanfeliu, A., "Robot Companion: A Social-Force based approach with Human Awareness-Navigation in Crowded Environments", *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp.1688-1694, 2013

[20] Kidokoro, H., Kanda, T., Brščić, D. and Shiomi, M., "Will I bother here? -A robot anticipating its influence on pedestrian walking comfort". *Proc. ACM/IEEE Int. Conf. on Human-Robot Interaction (HRI)*, pp.259-266, 2013

[21] Brščić, D., Kanda, T., Ikeda, T. and Miyashita, T., "Person tracking in large public spaces using 3D range sensors", *IEEE Trans. on Human-Machine Systems*, 43, 6, pp.522-534, 2013

[22] Thrun, S., Burgard, W. and Fox, D., *Probabilistic Robotics*, The MIT Press, 2005

[23] Ikeda, T., Chigodo, Y., Rea, D., Zanlungo, F., Shiomi, M., Kanda, T., "Modeling and Prediction of Pedestrian Behavior based on the Sub-goal Concept", *Robotics: Science and Systems*, 2012

[24] Shiomi, M., Zanlungo, F., Hayashi, K. and Kanda, T., "A Framework with a Pedestrian Simulator for Deploying Robots into a Real Environment", *Proc. Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)*, 2012

[25] Kanda, T, Shiomi, M., Perrin, L., Nomura, T., Ishiguro, H., Hagita, N., "Analysis of people trajectories with ubiquitous sensors in a science museum", *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp.4846-4853, 2007