

# Hey, I'm over here – How can a robot attract people's attention?\*

Markus Finke

*Neuroinformatics and Cognitive Robotics Group  
Faculty of Informatics and Automatization  
Technical University Ilmenau  
P.O.Box 100565, 98684 Ilmenau, Germany  
Markus.Finke@Stud.Tu-Ilmenau.De*

Kheng Lee Koay, Kerstin Dautenhahn,

Chrystopher L. Nehaniv, Michael L. Walters and Joe Saunders  
*Adaptive Systems Research Group  
School of Computer Science, University of Hertfordshire  
College Lane, Hatfield, Herts AL10 9AB, United Kingdom  
K.L.Koay@herts.ac.uk*

**Abstract**— This paper describes how sonar sensors can be used to recognize human movements. The robot distinguishes objects from humans by assuming that only people move by themselves. Two methods using either rules or Hidden Markov Models are described. The robot classifies different movements to provide a basis for judging if a person is interested in an interaction. A comparison of two experiment results is presented. The use of orienting cues by the robot in response to detected human movement for eliciting interaction is also studied.

**Index Terms**— Human-Robot Interaction, Social Robotics.

## I. INTRODUCTION

Service robots are being developed for applications to assist humans which involve dialogue and/or communication with humans. Environments include offices or department stores, where service robots can potentially provide useful information on products (e.g. Boehme et al. [1], Gross et al. [2]). Should a robot directly approach and verbally address a customer? This behaviour might be acceptable if displayed by a human sales assistant, but might be interpreted as “intrusive” or “pushy” behaviour if performed by a robot. Would it be beneficial to have a robot capable of interesting a person in interaction with it in a more “gentle” way (e.g. by leaving it to the customer to approach the robot and initiate the interaction). Could certain movement cues provided by the robot elicit such self-initiated human behaviour? Therefore we investigated the following research questions: 1) How can a robot detect that a human is interested in interacting with it? 2) Can simple orientation movements be used to encourage a person to interact with a robot?

In order to address these research questions, we developed and experimentally evaluated two computational methods for detecting human movements using sonar sensors on a Peoplebot™ robot. Also we studied in experiments the reaction of human subjects towards the robot in conditions involving orientation cues.

The conventional approach for detecting human movement is normally performed using vision systems [3]. We believe

that in a human-robot interaction scenario, the dynamics of multi-modal interaction are often more important than the precise detection of particular features in the environment. It may also be possible to use sensory fusion for detecting human movement in the future.

The commercially available Peoplebot™ robot was used in this experiment and has various sensors, including infrared, sonar, contact sensors, and an onboard camera. We investigated using a sonar-based movement detection system, since non-vision sensors are widely used with success in the field of mobile robotics, especially in the area of obstacle detection. Buchberger et al. [5] uses a combination of laser and sonar sensors which avoid static and dynamic obstacles by recognizing objects in realtime. Salter et al. [4] using arrays of infrared sensors to detect human behaviour.

Sonar sensors can be error-prone as sometimes the data from a sensor can be lost. Sources of error include: ultrasonic waves not being deflected back directly to the sensor, but to other objects, then back to the robot. Detected distances maybe overestimated and the robot may collide with an object. Crosstalk may occur if more than one source emits ultrasonic waves. The received echoes can be sent by another source, so that the sensor detects an object as closer than it really is. Joerg and Berg [6] describe a method that defines an echo for each sensor so that it can distinguish between its own echo and that coming from any other source. They use pseudo-random sequences to get independent ultrasonic waves. Every sensor can also be used to identify echoes from other sensors. This information can be used for triangulation. Finally all sensors can emit ultrasonic waves at the same time so that obstacles can be detected earlier.

The basic research approach is presented in section II. In section III we describe two algorithms for the recognition of human movements. The first one is rule-based and the second one uses Hidden Markov Models (HMMs). A comparison of both methods is shown in section IV. Section V provides an analysis of how human behaviour may be related to the robot's behaviour.

## II. BASIC RESEARCH APPROACH

Sonar sensors cannot distinguish between an object and a person, and can only give two kinds of data: 1) There is an object at a measured distance. 2) There is no other object

\*The work was carried out during an internship of Markus Finke at University of Hertfordshire, hosted by the Adaptive Systems Research Group. The work described in this paper was partially conducted within the EU Integrated Project COGNIRON (“The Cognitive Companion”) funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020.

between that detected and the robot, because ultrasonic waves cannot go through objects.

The change of data over time is important because movements cause variations at every sensor. For the purpose of this paper we assume that only people can move by themselves and that moving objects detected at a height of one meter are usually associated with a moving person.

We also assume that the robot itself is static. Otherwise its movements cause significant changes of the data and the system cannot know if this is caused by a person or by the robot.

In order for the robot to know that a person is interested or wants to interact we take E.T. Hall’s [7], [8] “social distances” into consideration. At a certain proximity, it could be assumed that a person wants to interact with the robot. The spatial distances between a robot and a human are discussed in Walters et al. [9]. The generally recognized personal space zones between humans (e.g. northern Europeans) are well known and are discussed in Lambert [10]. It is also important to note that we cannot classify every person that has entered the robot’s social zone as interested in interacting with the robot, as the person may be just passing through the area. It is also safe to assume that people that are outside the robot’s social zone are probably not interested in interacting with the robot. Therefore, the system should identify different movements before deciding if the person is interested in interacting with the robot.

### III. THE ALGORITHMS

#### A. A rule-based approach

The algorithm (see Fig. 1(a)) concentrates on the following information: 1) The distance between the robot and the human subject, 2) the duration a human subject spends within the detection window of each sonar sensor, and 3) the initial and the final distances between human subject and robot when the human subject entered and left the detection window, respectively.

The collected data shows that sometimes the signal is lost for 3 or 4 timesteps (0.3 - 0.4 seconds). The received values are set to -1 to indicate the sensor error condition which is then ignored.

Sonar sensors readings are never stable, even in a static environment. This limitation does not preclude its usage as human movements usually cause more significant changes in the sensor readings. We used two different threshold values (i.e.  $k1=0.8$  and  $k2=1.35$ ) to assist in identifying human movements. These values were defined based on the ratio of previous and current sensor readings of the distances between the subject and the sensor, and were obtained empirically through trial-and-error. These threshold values are plotted on figure 1(b), where  $k1$  and  $k2$  each represent the border lines that separate regions B and C, and regions C and A respectively. Different regions correspond to a person entering (zone A) or leaving (zone B) the area of detection of the sensor. If the ratio of previous and current sensor readings is in

zone C, this means either no significant change has occurred or the person is still in the area of detection.

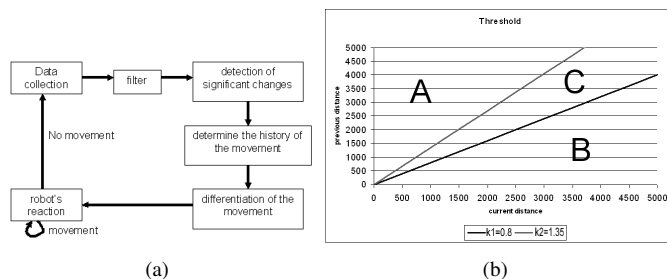


Fig. 1. (a) The main modules of the rule-based algorithm, (b) Thresholds indicating human movements for a single sensor.

Usually only one or two sensors will be involved in detecting a person’s approach behaviour (depending on how far the person is from the sensors and the sensing angle of the sound beams of the sensors). There will be no significant changes in the sensory readings as the person approaches. However by comparing each of the current sensory readings of the involved sensors with the average sensory reading over a period of 10 timesteps (i.e. 1 second), the system can recognise human approach behaviour.

The rows of the matrix in figure 2 each show how a given sensor has been activated over time. If a person approaches the robot, one or two sensors will be activated several times in a short sequence.




Fig. 2. Patterns of movement. The y-axis displays the eight sensors. S1 points to the left, S8 to the right, S4 and S5 point to the front of the robot. The x-axis displays the time. The closer an object is detected the darker the gradient is. Top: Diagonal movement from the far left corner to the right side of the robot. Middle: Movement straight from the right side of the robot to the left. Bottom: Shows the movement of a person approaching the robot.

For detecting other human movement behaviours, the system will have to look at all the sensory readings over a period of 40 or 50 timesteps. The history of the sensory readings is usually stored in a table, where each column represents the sensory readings at each timestep (see figure 2). The system identifies movement by tracking the movement of the darkest gradient – corresponding to the closest object detected — along the sensor axis (i.e. row) over a period of 40 or 50 timesteps (i.e. columns). Movement usually involved the

darkest gradient moving across 4 rows in a sequential manner over a period of 40 or 50 timesteps (i.e. each of four sonar sensors sequentially detect a person over a period of time). By studying these examples of human movement data recorded from the sensors in such matrices, 11 rules were hand-coded to detect and classify the types of motion.

### B. Using Hidden Markov Models

Hidden Markov Models (HMMs) are a technique using finite automata with probabilistic transitions to model the generation of observations corresponding to different patterns of a system's behaviour. The basic ideas are explained in the tutorial by Rabiner [11]. HMMs are widely used for pattern recognition. For example; Billard and Calinon [12] used HMMs to recognize and produce matching behaviours in a skill learning (imitation) context. Westeyn et al. [13] explain a system that detects gestures to control a car radio.

Significant examples are necessary to train the HMMs in order to use them in an application with actual data. Different movements in the environment of a robot cause different patterns in the detected distances over time. Sample data is illustrated in figure 2.

We used the Georgia Tech Gesture Toolkit (GT<sup>2</sup>k) built on top of an HMM system from Cambridge University – the same system used for the above two applications [13]. With a few modifications GT<sup>2</sup>k automatically ran the HMM algorithms. The most extensive task was the preparation of training data for the system.

The HMMs are time-invariant, but cannot easily handle cases where the same movement they were trained to detect occurs either at different distances or in different environments. Therefore, we standardized the data from the eight sensors, but lost the ability to distinguish between movements towards or away from the robot. Also, as movements towards the robot were only recognized badly as a result of using a single HMM approach, a two HMM was used.

Preprocessing the data got rid of sequences where the sonar data was lost, and these were replaced by the distance measured in the following timestep. The system receives the data and continuously averages the last ten timesteps. This is stable if nothing happens in the environment of the robot. Otherwise the difference between the current distance and average will be significant.

We used two eight-state HMMs where transitions cannot go back to a previous state, but can stay at the same one, go to the next one or even skip one state (so-called 'right-to-left models'). The first model is responsible for movements close to the robot, so that people who are interested in the robot or want to interact, are detected. The second model recognizes movements from one side of the robot to the other one. This model classifies people who are not interested, or show only a little interest, but move on. These people might be interested if they notice that the robot watches them and maybe turns towards them.

For detecting a movement, the values of the differences between current distance and average of the last seconds

were saved in a text-file. This was repeated for every sensor independently using the first HMM (i.e. approach detection). Only if no movement towards the robot is detected, will the values of all eight sensors together be saved in a second text-file, which was used as input for the second HMM (i.e. left or right movement detection).

### C. Behaviour of the robot

The robot behaves the same way for both algorithms depending on the recognized movement. If a person approaches the robot closer than one meter, the robot will assume that this person is interested or wants to interact. In this case the robot turns head-on to the person, because people are used to talking face to face during an interaction. The distance of one meter is chosen with regard to E.T. Hall's "social distances". He subdivides the environment of a person into intimate, personal, social and public zones. The personal area is an adequate distance for human-human interaction and we assumed in this work that it also applies to human-robot interaction.<sup>1</sup>

If the robot detects a movement from one side to the other, it will assume that the person is not interested. It is also possible that a person did not realize that the robot was working and watching him. The robot will then turn 45 degrees in the direction the person is moving. This gives feedback to the person that he has been detected. The person might then become interested in the robot and approach. If the person does not come close to the robot but moves on, the robot will turn back to its previous position.

Collection of sonar data was temporarily suspended as soon as the robot started turning. Otherwise, the robot would experience sensory input similar to when a person is moving from left to right, as it turns from right to left, and vice versa. If the robot turns, one sensor will detect distances that its neighboring sensor has detected earlier and there will be a significant variation that would be interpreted as the "detection" of a moving person.<sup>2</sup>

## IV. COMPARISON OF BOTH ALGORITHMS

We carried out two experiments in order to compare both algorithms. The first one took place in the same environment as the training phase. Twelve people, who were not involved in the training, moved around the robot. The second experiment took place in a new environment with people who were neither involved in the training nor in the first experiment. This experiment demonstrated the ability of the algorithms to generalize in a new environment.

<sup>1</sup>But compare also the results of [9], showing strong individual differences between people on whether human-human interaction distances are generalized to their interactions with robots. In particular, some persons appear to treat robots more as objects (to which human-human social distances do not apply).

<sup>2</sup>This limitation of the robotic system is analogous to the fact that humans are blind to changes in visual scenes that occur during saccadic eye movement [14].




Fig. 3. Environment of the first experiment. The movements of the three scenarios are shown by the arrows. Scenario 1 corresponds to the subject moving from “right to left” and “left to right” in front of the robot. Scenario 2.1 and 2.2 have the subject moving past the robot on its right side from front to back and visa versa. Scenario 2.3 has the subject moving along a curved path from the robot’s front to the passage at the robot’s left. Scenario 2.4 is the same but with the subject moving in the opposite direction. Scenario 3 corresponds to the subject moving “forward” in a straight line from anywhere within a semi-circle ahead of the robot, stopping in front of the robot.

### A. Experiment 1

Each person received written instruction for movements subdivided into three scenarios (see fig. 3) before the experiment started. The environment of this experiment is shown in figure 3, where the movements of the three scenarios are indicated by the arrows.

In order to compare both the algorithms using exactly the same human movement data, each algorithm was used for an online test of six of the twelve cases, and the movement data were collected. The data collected during online testing of an algorithm were then later used in offline testing of the other algorithm. Therefore a total of twelve (six online and six offline) human movement cases was tested on each of the algorithm. Note that the movement data recorder stopped recording movement data as soon as an algorithm recognised a movement. Because the HMMs algorithm required more movement data than the rule-based method to classify a movement, it was expected that the HMMs would perform badly during offline testing with data collected from the rule-based method.

The results are shown in table I. The “incorrect” classifications include all examples that were classified incorrectly or were not classified.

Both algorithms were not trained for the movements in scenario 2. Scenarios 2.3 and 2.4 were similar to right→left and left→right training examples, but scenarios 2.1 and 2.2 were completely different from any of the training examples. The result shows that both algorithms could not detect “new” movements very well, especially the HMMs with a correct classification rate of 50% or less.

Table II shows the online and offline test results of right→left, left→right and forward movements. The columns and the rows of the table represents the online and offline test results of the algorithms respectively. This table shows in detail how many of the test examples were detected correctly

by one method, were detected incorrectly by the other method and vice versa.

The results shown in table III indicate that overall the algorithm using HMMs performed better than the rule-based method, but were worse on-line than rule-based offline for the left to right movement.

TABLE I  
ONLINE TEST RESULTS OF THE FIRST EXPERIMENT

Movement	Online Test Results					
	HMMs Algorithm			Rule-Based Algorithm		
	Correct (%)	Incorrect (%)	N	Correct (%)	Incorrect (%)	N
<b>Right to Left</b>	78	22	18	70	30	20
<b>Left to Right</b>	72	28	18	60	40	20
Scenario	2.1	67	33	6	50	6
	2.2	33	67	6	50	6
	2.3	50	50	6	100	6
	2.4	33	67	6	83	6
<b>Forward</b>	70	30	20	90	10	21

TABLE II  
ONLINE VS. OFFLINE TEST RESULTS OF THE FIRST EXPERIMENT

Offline Test Results	Algorithm	Movement	Online Test Results					
			HMMs Algorithm			Rule-Based Algorithm		
			Correct	Incorrect	Total	Correct	Incorrect	Total
Rule-Based	Right to Left	Correct	10	1	11	12	2	14
		Incorrect	4	3	7	2	4	6
		<b>Total</b>	<b>14</b>	<b>4</b>	<b>18</b>	<b>14</b>	<b>6</b>	<b>20</b>
	Left to Right	Correct	12	2	14	12	6	18
		Incorrect	1	3	4	0	2	2
		<b>Total</b>	<b>13</b>	<b>5</b>	<b>18</b>	<b>12</b>	<b>8</b>	<b>20</b>
Forward	Correct	7	5	12	19	1	20	
	Incorrect	7	1	8	0	1	1	
	<b>Total</b>	<b>14</b>	<b>6</b>	<b>20</b>	<b>19</b>	<b>2</b>	<b>21</b>	
HMMs	Right to Left	Correct	10	1	11	12	2	14
		Incorrect	4	3	7	2	4	6
		<b>Total</b>	<b>14</b>	<b>4</b>	<b>18</b>	<b>14</b>	<b>6</b>	<b>20</b>
	Left to Right	Correct	12	2	14	12	6	18
		Incorrect	1	3	4	0	2	2
		<b>Total</b>	<b>13</b>	<b>5</b>	<b>18</b>	<b>12</b>	<b>8</b>	<b>20</b>
Forward	Correct	7	5	12	19	1	20	
	Incorrect	7	1	8	0	1	1	
	<b>Total</b>	<b>14</b>	<b>6</b>	<b>20</b>	<b>19</b>	<b>2</b>	<b>21</b>	

Note: The results of the online and offline tests are shown in the columns and rows of the table respectively.

TABLE III  
EXPERIMENT 1 RESULTS SUMMARY

Movement	Online	Offline	Offline	Online	Overall: Online-Offline	
	HMMs (%)	Rule-Based (%)	HMMs (%)	Rule-Based (%)	HMMs (%)	Rule-Based (%)
<b>Right to Left</b>	78	61	70	70	74	66
<b>Left to right</b>	72	78	90	60	81	69
<b>Forward</b>	70	60	95	90	83	75
<b>Total</b>	73	66	85	73	79	70

### B. Experiment 2

The second experiment took place in a public corridor (see figure 4). The people were passers-by and were not instructed how to move or to behave in front of the robot, nor were they informed that an experiment was in progress. It should show how well the algorithms work in a different environment, with people who were not involved in the training process. We collected data in five different conditions or ‘rounds’. The difference between the conditions was in the behaviour of the robot. As in Experiment 1, the behaviour of the robot depended on the detected movement. In this experiment, the



Fig. 4. Environment of the second experiment. Sample paths that were chosen by the detected people are shown by the arrows.

robot responded differently in different conditions. During the first round the robot did not react. In the second and fourth round it turned 20 or 50 degrees respectively to the direction the person was moving (i.e. follow direction). In the third and fifth round the robot turned 20 or 50 degrees respectively into the direction the person came from (i.e. away direction).

The first condition, when the robot did not turn, lasted ten minutes. The other four conditions lasted five minutes each. Total numbers per round of observed persons varied and overall we tested 152 subjects. Most of the subjects movements were moving from one side to the other (relative to the robot). The robot stood near the entrance of a main corridor. During the trials, we observed that the majority of the subjects slowed down when they noticed the robot, but most of them moved on while looking at the robot. There was only one subject that became very interested in the robot, and approached it.

“Back” means that people came through a door behind the robot, passed its left side and turned left or passed its right side and turned right. In both cases they did not cross in front of the robot. “Forward” means the corresponding movement towards the door behind of the robot.

The robot used the rule-based algorithm during the experiment to trigger movement when a person was detected. The stored data was later tested offline using the HMMs. The results are shown in table IV.

There were eight cases where two subjects moved from two different directions in front of the robot. In six of these cases, the robot managed to detect only one subject’s movement instead of the two movements. For the other two cases, the robot failed to detect these movements. Note, we did not expect the robot to accurately detect simultaneous movements of more than one subject as both algorithms were built for detecting a single subject movement at a time.

The algorithm using HMMs is better overall than the rule-based one. However, for movements from left to right the rule-based method is better than the HMMs. The HMMs need to be trained with different data to improve the recognition of movements from left to right. The comparison of both algorithms’ performance is shown in figure 5.

TABLE IV  
SIMULATION (HMMs) RESULTS OF THE SECOND EXPERIMENT

		Online Test Results of Rule-Based Algorithm										Total			
		Robot in Static		Robot Rotate Follow Direction				Robot Rotate Away Direction							
				20°		50°		20°		50°					
		%	Correct	Incorrect	Correct	Incorrect	Correct	Incorrect	Correct	Incorrect	Correct		Incorrect		
Offline Test Results of HMMs Algorithm	Subject's Movement	Right to Left	Correct	100	0	60	40	50	42	100	0	60	13	93	
			Incorrect	0	0	0	0	8	0	0	0	13	13	7	
		Left to Right	Correct	60	30	43	29	20	20	25	25	25	25	60	
			Incorrect	10	0	14	14	40	20	50	0	50	0	40	
		Forward	Correct	100	0	-	-	-	-	-	-	-	-	-	100
			Incorrect	0	0	-	-	-	-	-	-	-	-	-	0
	Back	Correct	100	0	100	0	33	33	0	100	0	100	93		
		Incorrect	0	0	0	0	0	33	0	0	0	0	7		
	Total		93	7	72	28	51	49	58	42	49	51			




Fig. 5. Experiment 2 Results Summary

## V. ANALYSIS OF THE VIDEO DATA

Finally, we analyzed the video data with regard to the behaviour of people when they noticed the robot. A person’s behaviour is grouped under one of five categories according the robot’s reaction. The total number of each of the five groups differs from the total number of these groups in table IV because sometimes the robot turned the wrong way due to misclassification. The difference in these numbers is also caused by groups of people that were detected by the robot as one person, but not all reacted the same way.

The video data showed the following seven behaviours of people, organized from those that showed the highest interest to those who showed no interest in the robot: **a** - approaches the robot, **b** - stops, speaks (excitedly) then continued along the same path, **c** - stops and watches, then continued along the same path, **d** - watches and slows down, **e** - watches while walking on, **f** - only a short glimpse at the robot while walking on and **g** - ignores robot.

Analysis of human behaviour just by watching it is very difficult because the interpretation is influenced by the personal opinion of the observer. Because of this the video data was interpreted twice by two different people. The level of inter-rater agreement is 84%.

Figure 6 shows the results of this analysis. The surprisingly high ratio of people who ignore the robot was due to some people who moved along the corridor several times, but were not interested in the robot when they came for the second or third time. The ratio was also influenced by some members of our laboratory who already knew the robot and its behaviour.

Table V illustrates subject behaviours in respond to robot’s action. If the robot classified the movement incorrectly, the resulting behaviour of the people was added to the group according to the robot’s action. If the robot did not detect a

movement, the behaviour of the person was classified “doesn’t move” since the robot did not move.

The results with respect to the behaviour of the observed people show that most of them looked at the robot, while walking on. However, the orientation cue of the robot, which was meant to attract the person’s attention (rotation), seemed not to be sufficient for eliciting a response. This may be due to the robot movement detector algorithms taking about 4 or 5 seconds to detect a movement. By the time the robot moves, the person had already passed the robot. Also, the experiments were carried out in a busy university where people often walked past quickly and were not distracted easily. Under these difficult conditions it is probably not surprising that a simple orientation cue did not have any major effect.




Fig. 6. Subjects’ Reaction to Robot’s Behaviours

TABLE V

VIDEO ANALYSIS OF PEOPLE’S REACTION TO ROBOT BEHAVIOUR

Subject's Reaction	Robot's Reaction				
	Static	Follow 20°	Reverse 20°	Follow 50°	Reverse 50°
a	1	0	1	0	0
b	3	0	2	0	0
c	5*	1	0	0	2
d	8	0	4*	0	1
e	26*	4	5	5	10
f	5	0	0	0	0
g	19	7	9	16	18
Total (N)	67	12	21	21	31

Notes: \*2 persons stopped as they wanted to know if they can pass the robot; 2 persons stopped as they talked to each other and then moved to different directions, ♦ once a person appeared shortly after another and talked to him about the robot and ♣ one person looked bemused when the robot turned.

## VI. CONCLUSION

We have created two algorithms to detect human movements in the environment of a robot just by using sonar sensors. One algorithm is rule-based and analyzes the sonar data in order to find significant changes over time. The second one uses Hidden Markov Models to recognize a pattern in the data.

Both algorithms were implemented on a PeopleBot™ and their reliability was compared in two experiments. The second experiment also tested the ability to generalize in different environments. The results of both experiments show that both algorithms work adequately, but the one using Hidden Markov Models works better and detects the movements correctly in approximately 80% of the cases. The reliability of the

algorithms can be improved in the future by incorporating different movements which happen in real scenarios.

The detected movements are used by the robot to interpret the behaviour of a person. We assumed that people who are interested in the robot and want to interact, approach the robot. The reaction of the robot to interact with the subject will depend on this interpretation.

With respect to people’s behaviours, we found that most people looked at the robot when the robot is in sight, while walking on. However, the robot’s orientation cue (rotation) was not enough, perhaps due to the robot’s slow reactions. One of the solutions could be first using a voice system to attract the attention of people that have already moved past the robot, then followed by the orientation cue. Future work needs to investigate other robot cues (e.g. movement, speech, gestures) or more likely, a combination of various robot cues that will be able to attract attention and encourage approach and engagement in an interaction with the robot.

## REFERENCES

- [1] H.-J. Boehme, T. Wilhelm, J. Key, C. Schauer, Ch. Schroeter, H.-M. Gross and T. Hempel, “An Approach to Multi-modal Human-Machine Interaction for Intelligent Service Robots”, *Robotics and Autonomous Systems*, vol. 44, pp. 83-96, Elsevier Science.
- [2] H.-M. Gross, H.-J. Boehme, J. Key and T. Wilhelm, “The PERSES Project - a Vision-based Interactive Mobile Shopping Assistant”, *Künstliche Intelligenz*, vol. 4, pp. 34-36, 2000.
- [3] W. Zajdel, Z. Zivkovic and B. Kröse, “Keeping track of humans: Have I seen this person before?”, Accepted for *ICRA*, 2005.
- [4] T. Salter, R. te Boekhorst and K. Dautenhahn, “Detecting and analysing children’s play styles with autonomous mobile robots: A case study comparing observational data with sensor readings”, *Proc. IAS-8*, 2004, pp. 61-70.
- [5] M. Buchberger, K.-W. Joerg and E. von Puttkamer, “Laserradar and Sonar Based World Modeling and Motion Control for Fast Obstacle Avoidance of the Autonomous Mobile Robot MOBOT-IV”, *IEEE International Conference on Robotics and Automation*, 1993, pp. 534-540.
- [6] K.-W. Joerg and M. Berg, “First Results in Eliminating Crosstalk and Noise by Applying Pseudo-Random Sequences to Mobile Robot Sonar Sensing”, *Proc. EUROBOT '96*, pp. 40-45.
- [7] E. T. Hall, *The Hidden Dimensions: Man’s Use of Space in Public and Private*, The Bodley Head Ltd, London, UK, 1966.
- [8] E. T. Hall, “Proxemics”, *Current Anthropology*, vol. 9, no. 2-3, pp. 83-108, 1968.
- [9] M. L. Walters, K. Dautenhahn, R. te Boekhorst, K. L. Koay, C. Kaouri, S. Woods, C. Nehaniv, D. Lee and I. Werry, “The Influence of Subjects Personality Traits on Personal Spatial Zones in a Human-Robot Interaction Experiment”, *Proc. IEEE Ro-man*, 2005.
- [10] D. Lambert, *Body Language*. HarperCollins, 2004.
- [11] L. R. Rabiner, “A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition”, *Proc. of the IEEE*, vol. 77, no. 2, 1989.
- [12] S. Calinon and A. Billard, “Learning of Gestures by Imitation in a Humanoid Robot”, Dautenhahn, K. and Nehaniv, C.L. (eds.). *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press. 2005. In press.
- [13] T. Westeyn, H. Brashear, A. Atrash and T. Starner, “Georgia Tech Gesture Toolkit: Supporting Experiments in Gesture Recognition”, *Proc. 5th ICMI*, pp. 85-92, 2003.
- [14] J. M. Henderson and A. Hollingworth, “Global transsaccadic blindness during scene perception”, *Psychological Science*, vol.14, pp.493-497, 2003.