

Robot Interactive Learning through Human Assistance

G.Ferrer, A. Garrell, M. Villamizar, I.Huerta, A. Sanfeliu

1 Introduction

Humans live interacting with other people and do everyday tasks in individual and collective ways. Robotics people are interested in building robots that can interact with people in the same way that humans do. In order to reach this goal robots must learn from the interaction with humans and learn the human skills used in everyday life to acquire robot social behaviors that can then be used in a wide range of real-world scenarios: domestic tasks, shopping, assistance, guidance, entertainment, surveillance, rescue or industrial shop-floor.

There are many examples where this interaction occurs, but some of them are very basic and people do not realize the extreme difficulty that entails executing such tasks for a robot. For example, the navigation in crowded environments, such as crossing streets or shopping malls, or the social engagement for initiate a conversation, are simple examples where this interaction occurs. In the last years important academic and private research efforts have been carried out in this field. Examples can be seen in automatic exploration sites [29], evacuation of people in emergency situations [4], crafting robots that operate as team members [26], therapists [7], robotics services [22] or robot guiding [15, 14]

In this chapter we will present some examples where the robots learn from the interaction with humans using the general multimodal interaction framework. We will show how the general multimodal system is used in three specific tasks namely: interactive motion learning for robot companion; robot's proactively seeking interaction; and online face learning using robot vision.

There has been already described the general multimodal interactive framework shown in Fig. 1. As it can be seen in the framework, the model can be learned off-line or on-line, and the human (the oracle) uses the information coming from inputs

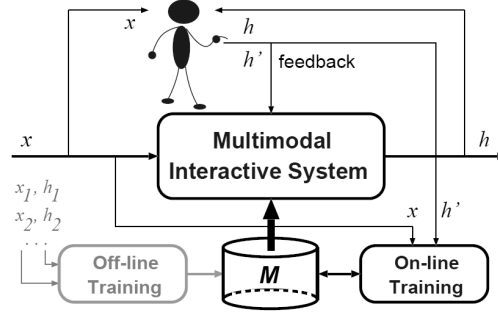


Fig. 1 General multimodal interactive framework.

and the outputs to train again the system in order to improve the model. We will see in the three examples how this framework is used.

We have developed three prototypes where the interaction occurs and it is used to improve the systems. The first prototype is "interactive motion learning for robot companion". The objective is to learn how a robot has to approach to a pedestrian who is going to a destination, minimizing the disturbances to the expected person's path. In this prototype, the robot has to detect the person's path, predict where the person is going to move and approach to the target while taking into account the person intentionality.

In the second prototype, "robot's proactively seeking interaction", the objective is to invite a person to approach the robot to initiate a dialogue. The robot has to take into account the person behavior (reactions) to convince the person to approach the robot. The robot uses a perception system to know the person position and orientation and uses a dialogue and robot motions to invite the person to approach.

The third prototype, "online face learning using robot vision", uses the second prototype to approach the person to the robot and the objective is to learn the face of the person that is invited for a dialogue. The system learns the face of the person by means of a sequence of images that the robot vision system captures while the person is in front of the robot. The robot asks the person when the captured face image is very different with respect to the learned face model. If the person agrees with the new face image, the robot uses this image as a positive image to improve the face classifier. In case that the person rejects that face image, the robot uses the image as a negative image to also improve the face classifier. The on-line face learning is done in real-time and is robust to varying environment conditions such as lighting changes. Moreover, it is robust to different people independently on the aspect and gender.

Throughout the three prototypes, the multimodal interactive system improves the accuracy and robustness of the prototypes thanks to the use of a human in the loop. The human plays the role of a teacher with the robots, that is, it evaluates and corrects the results of the robots' tasks in changing environment conditions and human behaviors. The system has been tested in real-life situations and the tests

shown the improvements of using this framework with respect using classical non-interactive approaches in several robot tasks.

The remainder of the chapter is organized as follows. In section 2, the interactive motion learning for robot companion approach towards humans is explained. How the robot performs his active behavior is presented in the section 3.1. Section 3 describes the online face learning using robot vision to detect and identify the people. Finally, the last section briefly reviews the topics discussed in the different sections of this chapter and establishes the final concluding remarks of this work.

2 Interactive Motion Learning for Robot Companion

Navigation in crowded urban environments, such as crossing streets or shopping malls, is an easy task for humans. However, it is extremely difficult for a robot due to the high environment uncertainties and the variability of the human behavior.

The uncertainties associated to the problem can be partially overcome using the multimodal interaction (MI) framework, where the human can teach specific issues of the robot companion approach. The best robot navigation behaviors can be taught by the human using the feedback of the system, through discarding the worst behaviors, modifying the present behaviors or adding new behaviors.

The aim of this prototype is to show how a robot can learn to navigate safely and naturally in urban settings, minimizing the disturbances to the expected person's paths in two different situations: when crossing the path of a person and when approaching a person to guide him/her to a destination. In order to implement the system we have to take into account several requirements which are:

- The human detector has to track the person path, while handle occlusions and crossings.
- The human motion prediction must infer the person motion intentionality (goal), forecasting the path required to get there.
- The robot motion in a guiding task must take into account the person motion intentionality.

The prototype scheme is depicted in Fig. 2. We can realize that it shares some issues of the general multimodal interaction framework shown in Fig. 1. The input to the system is the robot motion and the pedestrian paths, which are obtained through the robot motion features and the robot laser/vision tracker. The output of the system is the robot motion approaching or guiding the person. The human in the loop can modify the robot motion behavior in different ways. We have used in this prototype the on-line feedback of the person by using a subjective measure of comfortableness of the target being approached or guided. This measure allows to learn some parameters of the robot motion.

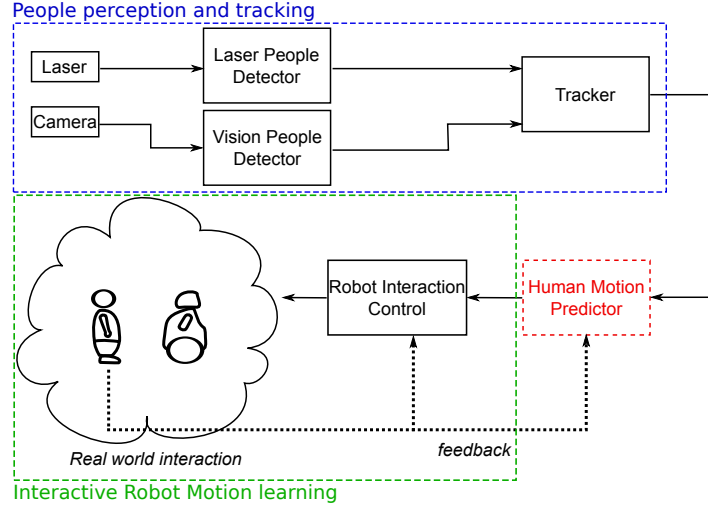


Fig. 2 Schematic prototype of the interactive motion learning for robot companion.

2.1 People Detection and Tracking

People detection is needed to extract the learning parameters for comfortable robot navigation in urban sites. Our implementation of the laser detector is based on [2], and the vision descriptor used for the person classifier is the Histogram of Oriented Gradient [6]. The people tracking follows a straightforward implementation of the work of [23] and [1] with some variations, for example instead of using a Kalman filter, we use a particle filter.

We have used a parallel design for the two detectors in order to obtain a robust detection of the people, fusing the information of the laser and vision detectors. The output of this fusion is used as the tracker input. The integration of these two sensors have entailed the calibration of both elements, camera and laser, being the camera at the moving robot head and the laser is at the bottom part of the robot.

2.2 Human Motion Prediction

We require a model capable of forecasting the set of trajectories any person might describe at any time, specifically in urban settings. The literature on this issue is extensive. Approaches like the works of Bennewitz [3], or Zielbart [35] learn basic patterns of human motion trajectories. Another approach to predict human movements makes use of the geometry and a final position which a person aims to, like the work of Foka [11], or Ferrer [10]. However, these models do not consider any kind of interaction, which is essential in the study of human-robot interaction. The

social force model proposed by Helbing [18] takes into account both destinations and interactions by defining a summation of existing forces determining people trajectories. Zanolungo *et al.* [34] propose a variation of the social force model taking into account the time of collision.

Previously to consider the motion prediction problem, it is greatly useful to obtain an intentionality predictor as well. We have proposed a geometrical approach in which a Bayesian predictor calculates the person posteriori probabilities to reach all destinations in the scene. The best candidate destination is used for our purposes.

In order to implement the human motion predictor module, we used the outcomes of the works of Helbing [18] and Zanolungo [34]. The Helbing's approach treats each pedestrian as a particle abiding the laws of Newtonian mechanics. We have incorporated the social work to the multimodal interaction approach, where the robot is initially commanded by an expert controller, while simultaneously, a single person tries to perform a determined trajectory aiming to a given destination.

The human-robot interaction is provided under the shape of the person's response to the motion stimuli generated by the robot. Under these settings, the expert controller perturbs the expected target person trajectory, understood as the path generated under no-external force conditions. Thus, we decouple the experiment settings into a two step optimization: first we optimize the parameters of the model forces describing the expected human trajectories under no external constraints and second, we optimize the parameters of the force interaction model under the presence of a moving robot, taken into account that these are the only external force altering the outcome of the described trajectory. All optimizations are carried out using genetic optimization algorithms.

To this end, we have recorded two different databases: one force-free and the other under the influence of robot force. The second one is used to determine the value of the existing parameters describing the human motion prediction models.

The proposed multimodal interaction model helps to enlighten the nature of the model, in addition to generate controlled interaction forces that otherwise would be extremely complicated to generate. We are using an external person to command the robot movements, while the inputs provided by the expert agent would help the human prediction module as well as the robot interaction module in their respective learning parameters. In order the human enter the information to the system a video-game controller is used, consisting of a direction pad and a stop key, although controlling the robot is a hard task. Thus, the velocities that the robot can achieve while in tele-operated mode are quite limited, prioritizing safety over all.

It will be shown in the following section the direct applicability of the motion prediction model, constituting the base for the robot-approach model.

2.3 Interactive Robot Motion Learning

The system described below only can work if the previous described modules are working properly. The robot motion is based in the social forces commented in

the previous sections, and the robot autonomous move to the destination goal, first looking for the person and then accompany him/her to the destination goal. In the robot motion the interaction is done continuously, through the social forces and also using the human feedback of comfortableness, to learn different robot approach behaviors, which adds extra difficulty to the problem.

There are few articles at the present time in this topic and the ones that have been developed take into account different hypotheses and constrains. Examples of them are the work of Fox [12] or more recently the work of Fraichard [13] which analyzes dynamical obstacle avoidance strategies for robot navigation; the work of Kanda *et al.* [21] that uses prediction strategies in social robots in a train station; or the works of Chung *et al.* [5] or Henry [19] in robot control design.

The on-line feedback comes from the target person to whom the robot tries to approach, as part of the multimodal interaction provided by a human agent by using a video-game controller, the same used in the model learning phase. However, instead of a direct position control, we expect to receive a feedback measure of the subjective comfortableness of the target being approached.

This feedback is a subjective measure, which varies some parameters of the system by weighting the contribution of all the active forces:

- Force to the target destination: we infer the target destination by using the intentionality prediction described at section 2.2 and thus the robot aims to the most expectable target destination.
- Force aiming to the target: either the current target position as well the expected motion prediction are known. Taking into account this information, we consequently choose the target destination.
- Force of interaction: that is a repulsive force due to the relative position and velocity between the robot and the target.

The combination of these three forces determines the behavior of the robot while the robot is approaching the person. Although we want to obtain a general approaching rule, it highly varies from person to person in addition to the highly noisy environment in which we are working. Accordingly, more sophisticated estimation methods are used to obtain the approaching behavior, such as the fuzzy logic methods. While iteratively repeating the robot physical approach, the provided feedback refines the weights of the force parameters and we can infer a basic interactive behavior where the person feels comfortable under the presence of the robot.

As can be seen in Fig. 3, we have reproduced the experiment under controlled conditions. The left figure shows the robot motion and after a few approaches to the target, the robot captures the behavior of the person, by heading towards the most expectable destination of the target. The attractive force to the target destination is plotted as the blue arrow, and the force approaching the person is plotted as the lilac arrow. The interaction force represent the repulsion generated by the target towards the robot. This force is important to reach the state where the robot does not approach too close to the target, as this behavior will most likely produce repulsion. The result of all the weighted forces is represented as the red arrow. The quantification of the approach is done qualitatively. Once the target does not use any




Fig. 3 Illustration of the experiment.

additional feedback from the command to correct the weighted forces, we assume the system is reach the good behavior.

3 Autonomous Mobile Robot Seeking Interaction for Human-Assisted Learning

In the last years, great efforts have been carried out by researchers around the world with the aim of creating robots capable of initiate and keep dynamic and coherent conversations with humans [?]. If robots are able to start a conversation, they create an active engagement with people which can be used to seek assistance from them. This engagement is particular convenient to improve some robot skills. For example, a human can act as a teacher to guide and correct the robot's behavior or its response. This active interaction leads to improve the robot capabilities using the human knowledge.

In this section, we present a multi-modal framework where robot and human interact actively to compute an on-line and discriminative face detector. To achieve this objective, the proposed framework consists of two main components or steps. The first one corresponds to create the engagement between the robot and a human, whereas the second step refers to the computation of the on-line face detector once the engagement and the dialogue are established.

More specifically, during the first step, the robot seeks and approaches to a human in order to initiate the conversation or interaction. This is done using its sensors and approaching algorithms. Once the conversation is initialized, a coherent dialogue is conducted during the second step to compute and refine the face detector using the human assistance. This results in a robust and discriminative face detector that is computed on the fly and is assisted in difficult circumstances.

The proposed framework is described in the following. Sec. 3.1 shows the proactively seeking interaction between the robot and humans (first step), and Sec. 3.2

describes the on-line face detector and the procedure used to assist the classifier using human-robot interactions (second step).

3.1 Robot's Proactively Seeking Interaction

Recently, social robots have begun to move from laboratories to real environments to perform daily life activities [27, 28, 32]. To this end, the robots must be able to interact with people in a natural way. Recent studies have shown robots which are able to encourage people to begin interaction [8, 17], but using a strategy based on people approaching to the robot in order to establish the interaction and dialogue. Contrary, we present, in this section, a method where the robot is proactive and approaches to people to initiate the interaction and establish the engagement. This is exemplified in Fig. 4.

This proactive way of creating engagements between people and robots enables numerous applications such as guiding robots, tourism robots, or robots focused in approaching people for providing information about a specific urban area. On the other hand, this engagement can be also useful to assist the robot and improve its skills. For example, using the human help, the robot can improve its vision skills. Therefore, it can detect objects and faces in a more robust and discriminative manner. The human can assist the robot to validate or correct the robot responses when it has uncertainty about its predictions. In this way, the robot capabilities are improved along with the number of human interventions. This is particular application is addressed in Sec. 3.2.



Fig. 4 Robot approaching. The TIBI robot approaches to a human to start the interaction.

To seek the interaction with humans, the robot has a people detector that allows to localize and identify humans in its neighbourhood. Once the person is localized, the robot approaches and invites the human to initiate and participate in the interaction. The robot is also able to respond according to human reactions. For instance, if the robot invites a person to approach, and he ignores it, the robot will return to insist. However, if human does not approach, the robot will search for another volunteer. Furthermore, if a person shows interest in the robot, it will start the interaction process with this person.

The active robot's behavior is performed developing a finite state machine. This state machine allows robot to react depending on people's behavior. The robot is able to decide if humans are interested in starting the interaction by tracking people positions only.

Hall [16] presented a conceptual framework known as "proxemics" which studied human perception and the use of the space. This work proposed a basic classification of distances between individuals:

- Intimate distance: the presence of other person is unmistakable, close friends or lovers (0-45cm).
- Personal distance: comfortable spacing, friends (45cm-1.22m).
- Social distance: limited involvement, non-friends interaction (1.22m-3m).
- Public distance: outside circle of involvement, public speaking (>3m).

Based on these proxemics, Michalowski et al. [24] classified the space around a robot to distinguish human's levels of engagement while interacting or moving around a robot. In the present work, our robot tries to maintain a social distance through voice messages and movements.

In Table 1 some sample phrases uttered by the robot are presented. Allowing the robot to acquire the proactive behavior, the number of interactions between the robot and people increases, so, as it will be explained in section 3.2, humans are able to assist the robot in the the computation of an on-line method for face recognition.

Invitation to create an engagement	Hey, how are you? I'm Tibi. I'm trying to learn to detect faces, will you help me?
	Hi, I am Tibi, I'd like to learn how to recognize different objects, can you be my teacher?
Invitation to continue the interaction	I only want to talk to you, can you stay in front of me?
	Please, don't go. It will take just two
	Let me explain you the purpose of the experiment, and then, you can decide if you want to stay.
Invitation to start the engagement	Thanks for your patience. Let's start the demonstration.
	Now we are ready to start. I'm so happy you'll help me.

Table 1 Robot's utterances. Some utterances used during the human-robot interaction to keep an active and coherent conversation.

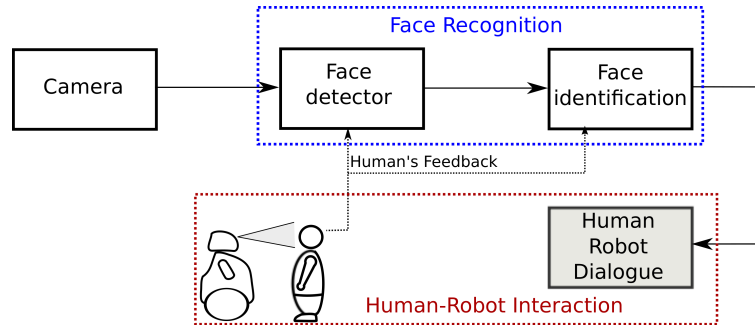


Fig. 5 On-line face learning. The proposed approach consists, mainly, of a face recognition module and a human-robot interaction module. The first module is in charge of detecting and identifying faces, whereas the second one establishes a dialog with a human. The synergically combination of both modules allows to compute a robust and efficient classifier for recognizing faces using a mobile robot.

3.2 On-line Face Learning Approach

In order to detect and identify faces in images, we use an on-line and discriminative classifier. Particularly, this classifier is based on on-line random ferns [30, 20], which can be progressively learned using its own hypotheses as new training samples. To avoid feeding the classifier with false positive samples, the robot will ask for the human assistance when dealing with uncertain hypotheses. This particular combination of human and robot skills allows to compute a discriminative and robust face classifier that outperforms a completely off-line random ferns [25], both in terms of recognition rate and number of false positives.

Following, the main components of the proposed approach are described in detail. Fig. 5 sketches these constituents and the overview scheme. The synergically combination of a face recognition system with a human-robot interaction module gives the proposed approach: *on-line face learning*.

Human-Robot Interaction. The on-line classifier is learned and assisted using the mobile robot and its interaction with a human. To this end, the robot is equipped with devices such as a keyboard and a screen that enable a dynamic and efficient interaction with the human. The interaction is carried out by formulating a set of concise questions (Fig. 6(Left)), that expect for a ‘yes’ or ‘not’ answer. In addition, the robot has been programmed with behaviors that avoid having large latency times, specially when the human does not know exactly how to proceed. Strategies for approaching the person in a safe and social manner, or attracting people’s attention have been designed for this purpose [9, 33].

On-line Face Classifier. The on-line classifier consists of a random ferns classifier [25] that, in contrast to its original formulation, is learned, updated and improved on the fly [30]. This yields a robust and discriminative classifier which is continu-

Greeting	Nice to meet you Can you teach me to detect faces/objects?
Assistance	Is your face inside the rectangle? I'm not sure if I see you, am I?
No detection	I can't see you, move a little bit. Can you stand in front of me?
Farewell	Thank you for your help, nice to meet you I hope I see you soon.




Fig. 6 Human-Robot Interaction. **Left:** Sample phrases uttered by the robot to allow the human assistance. **Right:** The interaction is carried out using diverse devices such as keyboard or touchscreen.

ously adapted to changing scene conditions and copes with different face gestures and appearance.

Random Ferns (RFs) are random and simple binary features computed from pixel intensities [25]. More formally, each Fern F_t is a set of m binary features $\{f_1^t, f_2^t, \dots, f_m^t\}$, whose outputs are Boolean values comparing two pixel intensities over an image I . Each feature can be expressed as:

$$f(x) = \begin{cases} 1 & I(\mathbf{x}_a) > I(\mathbf{x}_b) \\ 0 & I(\mathbf{x}_a) \leq I(\mathbf{x}_b) \end{cases}, \quad (1)$$

where \mathbf{x}_a and \mathbf{x}_b are the pixel coordinates. These coordinates are defined at random during the learning stage. The Fern output is represented by the combination of their Boolean feature outputs. For instance, the output z_t of a Fern F_t made of $m = 3$ features, with outputs $\{0, 1, 0\}$, is $(010)_2 = 2$.

On-line Random Ferns (ORFs) are Random Ferns which are continuously updated and refined using their own detection hypotheses or predictions. Initially, the parameters of this classifier are set using the first frame. To this end, the opencv face detector is used to find a face candidate with which to start the on-line learning procedure. Subsequently, several random affine deformations are applied to this training face sample in order to enlarge the initial training set, and initialize the RFs. In addition, the classifier is computed sharing a small set of RFs with the aim of increasing its efficiency, both for the training and detection stages [31].

As shown in Fig. 7(Left), during the on-line training, the number of positive p_z and negative n_z samples falling within each output of each Fern is accumulated. Then, given a sample bounding box centered at x and a Fern F_t , the probability that x belongs to the positive class is approximated by $P(F_t = z|x) = p_z / (p_z + n_z)$, where z is the Fern output [20, 30]. The average of all Fern probabilities gives the response of the on-line classifier:

$$H(x) = \frac{1}{k} \sum_{t=1}^k P(F_t|x), \quad (2)$$




Fig. 7 On-line Random Ferns. Left: Ferns probabilities. **Right:** Human-assistance criterion.

where $\frac{1}{k}$ is a normalization factor. If the classifier confidence $H(x)$ is above 0.5, the sample x will be assigned to the positive (face) class. Otherwise, it will be assigned to the negative (background) class.

The classifier is updated every frame using its own hypotheses or predictions. In particular, the classifier selects the hypothesis (bounding box) with the highest confidence as the new face location. Using this hypothesis as reference, nearby hypotheses are considered as new positive samples, while hypotheses which are far away are considered as new false positive samples. These positive and false positive samples are then evaluated for all the Ferns to update the aforementioned p_z and n_z parameters, see Fig. 7(Left).

Human Assistance. ORFs are continuously updated using their own detection predictions. However, in difficult situations in which the classifier is not confident about its response, the human assistance will be required. The degree of confidence is determined by the response $H(x)$. Ideally, if $H(x) > 0.5$ the sample should be classified as a positive. Yet, as shown in Fig. 7(Right), a range of values θ (centered on $H(x) = 0.5$) is defined for which the system is not truly confident about the classifier response. Note that the width of θ represents a trade off between the frequency of required human interventions, and the recognition rates. A concise evaluation of this parameter is performed in the experimental section.

3.3 Experiments

The on-line face learning method is evaluated on a face dataset acquired using a mobile robot. This face dataset has 12 sequences of 6 different persons (2 sequences per person). Each face classifier is learned using an image sequence and tested in the second one. The dataset is quite challenging as faces appear under partial occlusions, 3D rotations and at different scales. Also, fast motions and face gestures disturb the learning method [30].




Fig. 8 Face Recognition Rates. **Left:** Precision-Recall curves for different detection approaches. **Right:** Recognition rates in terms of human assistance.

More precisely, the learning/recognition method is evaluated using three different strategies for building the classifier. First, an offline Random Ferns approach (RFs) is considered. This classifier is learned using just the first frame of the training sequence and is not updated anymore. The second approach considers an ORFs methodology without human intervention. Finally, the proposed human-assisted approach which is denoted by A-ORFs. Remind that the human assistance is only required during the learning stage. During the test, all classifiers remain constant, with no further updating or assistance.

Fig. 8(Left) shows the Precision-Recall curves of the three methodologies, and Fig. 9(Left) depicts the Equal Error Rates (EER). Both graphs show that the A-ORFs consistently outperform the other two approaches. This was in fact expected, as the A-ORFs significantly reduce the risk of drifting, for which both the RFs and ORFs are very sensitive, especially when dealing with large variations of the learning sequence.

What is remarkable about the proposed approach is that its higher performance can be achieved with very little human effort. This is shown both in the last 4 rows of the table in Fig. 9(Left) and in Fig. 8(Right), where it is seen how the amount of human assistance influences the detection rates. Observe that with just assisting in a 4% of the training frames, the detection rate with respect to ORFs increases a 2%. This improvement grows to an 8% when the human assists on a 25% of the frames.

Finally, Fig. 9(Right) shows a few sample frames of the detection results, once the classifier learning is saturated (i.e., when no further human intervention is required). The on-line face classifier is able to handle large occlusions, scalings and rotations, at about 5 fps.

Method	θ	PR-EER	Human Assistance
RFs	—	55.81	—
ORFs	—	74.79	—
A-ORFs	0.05	76.31	$4.66\% \pm 0.46$
A-ORFs	0.1	76.51	$9.54\% \pm 0.87$
A-ORFs	0.2	79.44	$16.25\% \pm 1.09$
A-ORFs	0.3	82.06	$25.72\% \pm 1.65$




Fig. 9 Recognition Results. **Left:** Face recognition rates for different learning approaches: off-line Random Ferns (RFs), On-line Random Ferns (ORFs) and On-line Human-Assisted Random Ferns (A-ORFs). **Right:** Face detection examples given by the proposed human-assisted method.

4 Conclusions

In this chapter we have presented three different ways of robot learning using the interaction with humans and we have presented three different prototypes: interactive motion learning for robot companion; robot's proactively seeking interaction; and online face learning using robot vision.

We have presented a complete interactive motion learning for robot companion, the "interactive motion learning for robot companion" prototype, in three stages. The first initial design, the perception module, has been implemented and tested extensively in indoor environments. The implementation of the second design, where an external agent moves the robot, was a key step in order to obtain a human intentionality predictor and a motion predictor. A database has been collected of the robot approach to a walking human and the data was used to calculate the model parameters of the intrinsic forces and the interaction forces. For the final stage, we have implemented a multimodal feedback system, where a behavior inference of the weighting parameters of the contributing forces is implemented on-line. All this stages went through intensive real experimentation in outdoor scenarios, by far more challenging scenarios. A qualitative measure of the results gives information regarding the success of the system, in which the subjective comfortableness of the target and up being achieved.

In the "robot's proactively seeking interaction" prototype the human-robot interaction is performed in a very dynamic and efficient manner. Robot's proactive behavior has advantages in comparison with passive conducts. Firstly, invitation service, a robot offers information and invites people to interact with it. And, secondly, this behavior increases the number of interactions, and therefore, people can assist the robot to improve its skills continuously.

In the "online face learning using robot vision" prototype we have realized that using the interactive multimodal framework, we are able to handle large occlusions, scaling and rotations in different environment and with diverse number of people.

Acknowledgements *This research was conducted at the Institut de Robòtica i Informàtica Industrial (CSIC- UPC). It was partially supported by the CICYT project RobTaskCoop (DPI2010-17112) and the MIPRCV Ingenio Consolider 2010 (CSD2007-018).

References

1. K.O. Arras, S. Grzonka, M. Luber, and W. Burgard. Efficient people tracking in laser range data using a multi-hypothesis leg-tracker with adaptive occlusion probabilities. *IEEE International Conference on Robotics and Automation*, pages 1710–1715, May 2008.
2. K.O. Arras, O.M. Mozos, and W. Burgard. Using boosted features for the detection of people in 2d range data. In *IEEE International Conference on Robotics and Automation*, number April, pages 3402–3407. IEEE, 2007.
3. M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research*, 24(1):31, 2005.
4. J. Casper and R.R. Murphy. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE transactions on systems, man, and cybernetics, Part B.*, 33(3):367–385, 2003.
5. S.Y. Chung and H. Huang. A Mobile Robot That Understands Pedestrian Spatial Behaviors. *Learning*, pages 5861–5866, 2010.
6. N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE CVPR'05*, June 2005.
7. K. Dautenhahn. Robots as social actors: Aurora and the case of autism. In *The Third International Cognitive Technology Conference*, pages 359–374, 1999.
8. K. Dautenhahn, M. Walters, S. Woods, K.L. Koay, CL Nehaniv, A. Sisbot, R. Alami, and T. Siméon. How may i serve you?: a robot companion approaching a seated person in a helping context. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 172–179. ACM, 2006.
9. D. Feil-Seifer and M.J. Mataric. Defining socially assistive robotics. In *Proc. of the International Conference on Robotics and Automation*, pages 465–468, 2005.
10. G. Ferrer and A. Sanfeliu. Comparative analysis of human motion trajectory prediction using minimum variance curvature. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 135–136, Lausanne, Switzerland, 2011.
11. A.F. Foka and P.E. Trahanias. Probabilistic Autonomous Robot Navigation in Dynamic Environments with Human Motion Prediction. *International Journal of Social Robotics*, 2(1):79–94, 2010.
12. D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics; Automation Magazine*, 4(1):23–33, March 1997.
13. T. Fraichard and J.J. Kuffner. Guaranteeing motion safety for robots. *Autonomous Robots*, (January), February 2012.
14. A. Garrell and A. Sanfeliu. Local optimization of cooperative robot movements for guiding and regrouping people in a guiding mission. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010.
15. A. Garrell, A. Sanfeliu, and F. Moreno-Noguer. Discrete time motion model for guiding people in urban areas using multiple robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems.*, pages 486–491. IEEE, 2009.
16. E. T. Hall. The hidden dimension, man's use of space in public and private. *Great Britain, London: The Bodley Head Ltd*, 1966.
17. K. Hayashi, D. Sakamoto, T. Kanda, M. Shiomi, S. Koizumi, H. Ishiguro, T. Ogasawara, and N. Hagita. Humanoid robots as a passive-social medium: a field experiment at a train station. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction*, pages 137–144, 2007.

18. D. Helbing and P. Molnár. Social force model for pedestrian dynamics. In *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics*, volume 51, pages 4282–4286. May 1995.
19. P. Henry, C. Vollmer, and B. Ferris. Learning to navigate through crowded environments. *Robotics and Automation*, 2010.
20. Z. Kalal, J. Matas, and K. Mikolajczyk. P-n learning: Bootstrapping binary classifiers by structural constraints. In *Computer Vision and Pattern Recognition*, 2010.
21. T. Kanda, D.F. Glas, M. Shiomi, H. Ishiguro, and N. Hagita. Who will be the customer?: a social robot that anticipates people’s behavior from their trajectories. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 380–389. ACM, 2008.
22. K. Kawamura, R.T. Pack, M. Bishay, and M. Iskarous. Design philosophy for service robots. *Robotics and Autonomous Systems*, 18(1–2):109–116, 1996.
23. M. Luber, G. Diego Tipaldi, and K.O. Arras. Place-dependent people tracking. *The International Journal of Robotics Research*, 30(3):280, January 2011.
24. M.P. Michalowski, S. Sabanovic, and R. Simmons. A spatial model of engagement for a social robot. pages 762–767, 2006.
25. M. Ozuysal, M. Calonder, V. Lepetit, and P. Fua. Fast keypoint recognition using random ferns. In *IEEE Transactions Pattern Analysis and Machine Intelligence*, pages 448–461, 2010.
26. J.C. Scholtz. Human-robot interactions: Creating synergistic cyber froces. In *Multi-robot systems: from swarms to intelligent automata: proceedings from the NRL workshop on multi-robot systems*, page 177, 2002.
27. R. Siegwart, K.O. Arras, S. Bouabdallah, D. Burnier, G. Froidevaux, X. Greppin, B. Jensen, A. Lorotte, L. Mayor, and M. Meisser. Robox at expo. 02: A large-scale installation of personal robots. *Robotics and Autonomous Systems*, 42(3):203–222, 2003.
28. T. Tasaki, S. Matsumoto, H. Ohba, M. Toda, K. Komatani, T. Ogata, and H.G Okuno. Dynamic communication of humanoid robot with multiple people based on interaction distance. In *RO-MAN, 13th IEEE International Workshop on Robot and Human Interactive Communication*, pages 71–76. IEEE, 2004.
29. C. Trevai, Y. Fukazawa, J. Ota, H. Yuasa, T. Arai, and H. Asama. Cooperative exploration of mobile robots using reaction-diffusion equation on a graph. *ICRA*, 2003.
30. M. Villamizar, A. Garrell, A. Sanfeliu, and F. Moreno-Noguer. Online human-assisted learning using random ferns. In *International Conference on Pattern Recognition, Tsukuba, Japan*, 2012.
31. M. Villamizar, F. Moreno-Noguer, J. Andrade-Cetto, and A. Sanfeliu. Shared random ferns for efficient detection of multiple categories. In *International Conference on Pattern Recognition*, 2010.
32. K. Wada, T. Shibata, T. Saito, and K. Tanie. Analysis of factors that bring mental effects to elderly people in robot assisted activity. volume 2, pages 1710–1715, 2002.
33. D.M. Wilkes, R.T. Pack, A. Alford, and K. Kawamura. Hudl, a design philosophy for socially intelligent service robots. In *American Association for Artificial Intelligence Conference*, 1997.
34. F. Zanlungo, T. Ikeda, and T. Kanda. Social force model with explicit collision prediction. *EPL (Europhysics Letters)*, 93(6):68005, March 2011.
35. B.D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. Andrew, M. Bagnell, A. Hebert, K. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3931–3936, October 2009.