

Domestic Interaction on a Segway Base

W. Bradley Knox, Juhyun Lee, and Peter Stone

Department of Computer Sciences,
University of Texas at Austin
{bradknox, impjdi, pstone}@cs.utexas.edu
<http://www.cs.utexas.edu/~AustinVilla>

Abstract. To be useful in a home environment, an assistive robot needs to be capable of a broad range of interactive activities such as locating objects, following specific people, and distinguishing among different people. This paper presents a Segway-based robot that successfully performed all of these tasks en route to a second place finish in the RoboCup@Home 2007 competition. The main contribution is a complete description and analysis of the robot system and its implemented algorithms that enabled the robot’s successful human-robot interaction in this broad and challenging forum. We describe in detail a novel person recognition algorithm, a key component of our overall success, that included two co-trained classifiers, each focusing on different aspects of the person (face and shirt color).

1 Introduction

The population distribution of the world’s industrialized countries is becoming denser in the age range that may require assisted living. The U.S. Census Bureau estimates that from 2005 to 2030, the percentage of the country’s population that is above 85 years of age will increase by 50% [1]. The cost of providing such human care will be massive. The mean cost of assisted living in the U.S., according to a recent assessment, is \$21,600 per year¹. To give adequate care, we must find ways to decrease the economic cost per retired person. One solution for the not-so-distant future may be replacing the human home assistants with robots.

An assistive robot needs to be capable of a broad range of interactions. Among other capabilities, it must be able to locate and identify common objects; it must be able to follow people or guide people to places of interest; and it must be able to distinguish the set of people with whom it commonly interacts, while also successfully identifying strangers.

RoboCup@Home is an international competition designed to foster research on such interactive robots, with a particular focus on domestic environments. In 2007, its second year of existence, RoboCup@Home attracted eleven custom-built robots from ten different countries and five different continents.

This paper presents the UT Austin Villa RoboCup@Home 2007 entry, a Segway-based robot and the second-place finisher in the competition. The robot demonstrated its ability to complete versions of all three of the tasks mentioned above. The main contribution of this paper is a complete description of the robot

¹ www.eldercare.gov/eldercare/Public/resources/fact_sheets/assisted_living.asp

system and its implemented algorithms which enabled the robot’s successful human-robot interaction in this broad, challenging, and relevant event. A key component of our overall success was a novel person recognition algorithm that included two, co-trained classifiers, each focusing on different aspects of the person (face and shirt color).

The remainder of the paper is organized as follows. Section 2 briefly describes the RoboCup@Home competition. Section 3 introduces the UT Austin Villa robot. Section 5 describes our specific solutions and performance for each task. The final competition results are described in Section 6. Section 7 discusses related work and Section 8 concludes the paper.

2 RoboCup@Home

RoboCup@Home is an international research initiative that aims “to foster the development of useful robotic applications that can assist humans in everyday life”². In The 2007 Competition, robots in a living room and kitchen environment (see Figure 1) had to complete up to four of six specified tasks. These tasks can be considered fundamental building blocks toward the complex behavior and capabilities that would be required of a fully functional home assistant robot. The specific tasks are described in Table 1.



Fig. 1: The RoboCup@Home domestic setting.

Task	Description
Navigate	navigate to a commanded location
Manipulate	manipulate one of three chosen objects
Follow and Guide a Human	follow a human around the room
Lost and Found	search for and locate previously seen objects
Who Is Who	differentiate previously seen and unseen humans
Copycat	copy a human’s block movement in a game-like setting

Table 1: 2007 RoboCup@Home tasks

Within each task, there were two levels of difficulty. The easier level, called the first phase, existed as a proof of concept and often abstracted away part of the problem (e.g. object recognition or mapping and navigation). The second, more difficult phase of each task was structured similarly to how the task would need to be performed in a real domestic setting. During each phase, there was a ten minute time limit to complete the task objectives.

After the specific tasks, all teams performed a free-form demonstration in what was called the *Open Challenge*, during which they showed off their most impressive technical achievements to a panel of other team leaders. Each event was scored and five teams advanced to the Finals. In the Finals, the five finalists performed demonstrations for trustees of the RoboCup organization, who determined the final standings.

² www.ai.rug.nl/robocupathome/documents/rulebook.pdf

Our robot attempted three of the six possible @Home tasks. These tasks were *Lost and Found*, *Follow and Guide a Human*, and *Who Is Who*. Each task is described in the following subsections. Our specific approaches to the three tasks are detailed in Section 5.

2.1 Lost and Found

This task tested a robot’s ability to find an object that had been “lost” in the home environment. We competed in only the first phase of the *Lost and Found* task. In that phase, a team would hide a chosen object somewhere in the living environment at least five meters from their robot and out of its view. Then the task began. The task ended successfully when the robot had moved within 50 cm of the item and had announced that it found it.

2.2 Follow and Guide a Human

In *Follow and Guide a Human*, a robot followed a designated human as he or she walked throughout the home and then, optionally, returned to the starting position (thus “guiding” the human).

First Phase In the first phase, a team member led his or her robot across a path determined by the competition referees. The leader was permitted to wear any clothing or markers he chose.

Second Phase The rules were the same except that the human leader was a volunteer chosen from the audience. Therefore the algorithm needed to robustly identify a person without markers or pre-planned clothing.

2.3 Who Is Who

The *Who Is Who* task tested person-recognition capabilities on a mobile robot. Both phases of the task involved the robot learning to recognize four people, the referees rearranging the people and adding one new person (a “stranger”), and the robot subsequently identifying the four known people and the stranger accurately.

First Phase In the first phase of the *Who Is Who* task, the four people lined up side-to-side while a robot moved among them and learned their appearances and names. Once the robot finished training, the four people and a stranger were arranged into a new order by the referees. Then, the robot again moved among the people, announcing their names as each was identified. One mistake was allowed.

Second Phase The second phase was much like the first, but after the robot finishes training, the four known people and the stranger were placed by the referees in various locations around the entire living room and kitchen environment. The robot then had to search them out *and* correctly identify them.

3 The Segway Robot

This section introduces the hardware and software systems of the UT Austin Villa RoboCup@Home 2007 entry, shown in Figure 2.



Fig. 2: Our Segway home assistant robot.

The robot consists of a Segway Robotic Mobility Platform (RMP) 100³, supporting an on-board computer and various sensors. The Segway provides controlled power in a relatively small package. This suits a domestic environment well, for it is small enough to maneuver a living environment built for humans and powerful enough to reliably traverse varying indoor terrain including rugs, power cords, tile, and other uneven surfaces. The large wheels easily navigate small bumps that challenged other indoor robots during the competition. The two-wheeled, self-balancing robot reaches speeds up to six mph, exerts two horsepower, and has a zero turning radius. The Segway moves with two degrees of freedom, receiving motion commands in the form of forward velocity (m/sec) and angular velocity (radians/sec). It provides proprioceptive feedback in the form of measurements of odometry and pitch. With a payload capacity of 100–150 lbs., the Segway could easily carry several times the weight of its current load.

A 1GHz Fujitsu tablet PC sits atop the Segway platform, performing all sensory processing, behavior generation, and the generation of motor commands on-board. It interfaces with the Segway via USB at 20 Hz.

Two cameras and one laser range finder are available to sense the robot's environment. The Videre Design STOC camera⁴ provides depth information, but is not used for the tasks and experiments described in this paper. Higher picture quality is obtained by the second camera, an inexpensive Intel webcam which sends 30 frames per second. The attached Hokuyo URG-04LX⁵ is a short range, high resolution laser range finder that is well-suited for indoor environments. It collects 769 readings across 270 degrees at 10 Hz. Also, a Logitech microphone and USB speakers are attached.

The Segway RMP 100 is based on the p-Series Segway line for human transport. Despite its power, the robot is quite safe, featuring safety mechanisms such as automatic shut-off, an emergency kill rope, and speed caps at both the hardware and software levels.

A multi-threaded program, written from scratch, operates the robot. The program's structure can be divided into six modules: the camera input processing, the laser range finder input processing, the motion input/output, speech output, the high-level behavior unit, and the GUI.

4 Dual-Classifer Person Recognition

We use a dual-classifier system for person recognition. Face recognition is used as a starting point, but it is augmented by tracking the more frequently visible, but perhaps less uniquely identifying, characteristic of shirt color. Primary,

³ www.segway.com

⁴ www.videredesign.com/vision/stereo_products.htm

⁵ www.hokuyo-aut.jp/02sensor/07scanner/urg.html

uniquely identifying facial characteristics are dynamically associated with secondary, more ambiguous, possibly transient, but more easily computable characteristics (shirt colors). When primary characteristics are identifiable, they can be used to provide additional training data for the secondary characteristics visible on the person. The secondary characteristics can then be used to track the person, even when the primary characteristics are undetectable.

In this section, we summarize the full version of our novel person-recognition algorithm that includes two co-trained classifiers. Since the main contribution of this paper is the description of our overall, integrated system, we refer the reader to [14, 13] for the full details and related experiments. We used this algorithm in the finals of the competition (Section 5.4) and the second phase of *Follow and Guide a Human* (Section 5.2). A modification of the algorithm was used for the more constrained, earlier task of *Who Is Who* (Section 5.3). We note the changes in that section as appropriate.

The algorithm described in this section learns a single person’s characteristics and differentiates that person from unknown people. In Section 5.3, we extend it to differentiate among multiple known people.

4.1 Primary Classifier (Face)

We use the face as the input for our primary classifier. Challenging in its own right, face-recognition becomes even more difficult when performed by a mobile robot. Perspectives change as both robots and humans move. Illumination changes as humans move through different lighting conditions. When the mobile robot is in motion the image quality from its camera(s) is decreased. Also, computational limitations are acute. Faces must be recognized in real time with what computational resources are connected to the robot. A number of successful face-recognition algorithms exist (e.g. [2, 12]), but we found none that fit the needs of a home assistant robot.

Our algorithm is aided by two previously published algorithms. We detect the person’s face in the camera image using Viola and Jones’ real-time face-detection algorithm [19], and facial features are extracted as scale-invariant feature transforms (SIFT) features [16], shown in Figure 4), .

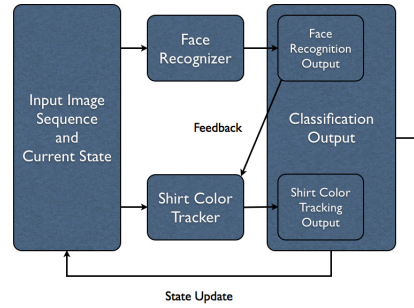



Fig. 3: A conceptual diagram of our dual-classifier person-recognition algorithm.



(a) same person (b) different people

Fig. 4: Matched SIFT features are connected with a line.

Training Using the face detection algorithm, we collect 50 image samples of the person being learned. SIFT features are extracted from these samples. The collected features are stored for use during testing.

Testing After training, the task of the face classifier is to determine whether the known person is visible in the camera image. To do this, it first identifies a face in the picture using Viola and Jones' face-detection algorithm. If one is found, it then extracts SIFT features from the face image. It then attempts to find matches between these SIFT features and the known identity's bag of SIFT features. If the number of SIFT feature matches are above a certain threshold, the face image is positively classified as that of the known person. Otherwise, the face image is deemed to be that of a stranger.

4.2 Secondary Classifier (Shirt)

A person's shirt is more often visible than his or her face, so we support face recognition by using the shirt as a secondary classifier.

Training To sample the person's shirt during training, first the algorithm scans incoming video images for a face, again using Viola and Jones' face detection algorithm. If it detects a face, the bounding box of the face (as given by face detection) is used to choose three other bounding boxes that provide positive and negative shirt samples. One box is drawn directly below the face to provide the positive sample pixels from the shirt. For negative samples, two other bounding boxes are drawn to the left and to the right of the face. Figure 5 illustrates the training process. Training data consists of 50 samples drawn from face-containing frames, taken simultaneously with the face classifier samples.




Fig. 5: Training and testing of the shirt-classifying algorithm. Top-left: face-detection; top-right: positive and negative sampling of the shirt; bottom-left: Boolean mapping of shirt colors onto image; bottom-right: shirt detection

At this point the collected samples are analyzed. Both the positive and negative sample pixels are inserted into respective histograms, each of which is a $16 \times 16 \times 16$ RGB cube. Each histogram is normalized and a final histogram is created by subtracting the negative histogram from the positive histogram. Values in this final histogram that are above a threshold are considered positive shirt colors. From this final histogram, an RGB color cube was created in which each RGB value contained a Boolean value indicating whether or not the RGB value was associated with the shirt color.

Testing Once the training is over, our robot is ready to track the person. The classifier maps a 320×240 webcam image to a same-sized image with Boolean values replacing RGB values at each pixel location. After the mapping, the classifier looks for blobs of Boolean true pixels in the mapped image. Among many

possible shirt candidate blobs, the blob representing the shirt is chosen by its height and its proximity to recent shirt blobs.

4.3 Inter-classifier Feedback

In this system, the secondary classifier acts as a backup to the primary classifier. When the robot does not detect a face, the secondary classifier is acted upon. If a face is detected, the primary classifier's output determines the output of the person classifier. Additionally, the shirt pixels below the detected face were used to check the accuracy of the shirt classifier. If the two classifiers disagreed (e.g. the face is classified as negative and the shirt is positive), the secondary classifier was retrained using newly taken samples. A conceptual diagram of the algorithm is shown in Figure 3.

5 UT Austin Villa Approach and Performance

This section describes the strategies and algorithms the Segway used in the tasks described in Section 2. All tasks were performed in the same home environment, shown differently in Figure 1 and Figure 6.




Fig. 6: An approximate recreation of the RoboCup@Home floor plan.

5.1 Lost and Found

In *Lost And Found*, a robot searched for a known object that had been placed in an unknown location in the home environment. The task setup is described in Section 2.1. Our robot competed in the first phase of *Lost and Found*.

First Phase We used an ARTag marker as the target object[8]. ARTag is a system of 2D fiducial markers and vision-based detection. The markers are robustly detected from impressive distances (more than 5 meters at 320×240 resolution in our lab with a $20\text{cm} \times 20\text{cm}$ marker) with varying light and even partial occlusion. Each marker is mapped to an integer by the provided software library. We did not observe any false positives from our ARTag system.

For the *Lost and Found* task, our robot searched the environment using a reflexive, model-free algorithm that relied on a fusion of range data and camera input. The Segway moved forward until its laser range finder detected an obstacle in its path. It would then look for free space, defined as an unoccupied rectangular section of the laser plane 0.75 m deep and a few centimeters wider than the Segway, to the left and right and turned until facing the free space. If both sides were free, the robot randomly chose a direction. If neither side was free, it turned to the right until it found free space. Algorithmically, free space was determined by a robustly tuned set of pie pieces in the laser data which overlapped to approximate a rectangle (see Figure 7).




Fig. 7: The laser range finder data is checked for occupancy at three different ranges and angles to approximate a rectangle that is a bit wider and deeper than the Segway.

We placed the object on a table at the opposite end from where the Segway began. A straight line between the two would have passed through a television, shelves, and a kitchen table. The robot had neither prior knowledge of the object’s location nor any model of the environment. The Segway successfully completed its search with more than three minutes to spare. Of the six teams that attempted *Lost and Found*, only three teams, including our team, completed it.

5.2 Follow and Guide a Human

In this task, a robot followed behind a human as he or she walked within the home environment, winding around the furniture. Its setup is described in Section 2.2.

First Phase We attempted only the following (not guiding) portion of this phase. We did not attempt the extension because time constraints and technical difficulties left the Segway without functional mapping software. (No team finished the extension of returning back to the starting point.) Again, we used an ARTag marker on the shirt of the leading human. The robot flawlessly followed the human leader, without touching furniture or the human. Six of eight teams that attempted the first phase of *Follow and Guide a Human* completed this portion of the task.

Second Phase Without the ARTags of the first phase, the robot instead trained and used a shirt classifier as described in Section 4.2. Since we anticipated following a human with his back turned, the face recognition component of our person recognition algorithm was not used.

In the competition, the referees chose a volunteer wearing a white shirt. Much of the background was also white, so the negative samples collected during training were not distinguishable by our algorithm from the samples of his shirt. Instead of tracking the volunteer’s shirt as intended, the robot classified a large portion of the wall as the person and was unable to follow the volunteer. The choice of volunteer revealed a weakness in our algorithm, but in later rounds (described in Section 5.4) we showed that, given a shirt color that is distinguishable from the background colors, the robot can follow a person for whom it has no a priori data.

5.3 Who Is Who

The *Who Is Who* task tested a mobile robot’s ability to meet and later recognize humans. We used a modification of the Section 4 dual-classifier person-recognition algorithm for both phases of the *Who Is Who* task, so we explain it here before describing the specific phases.

To learn the faces of multiple people, we train a face classifier for each person as described in Section 4.1. For *Who Is Who*, the output of the multiple-face classifier is the set of identities which had a number of SIFT feature matches above an empirically determined threshold. If the output set is empty, then the threshold is lowered and the classifier is rerun.

Given the set of candidate identities, a shirt classifier takes over. This classifier gathers samples as described in Section 4.2, but otherwise the shirt classifier is different, having been modified to eliminate blob selection. Since the face is

easily detectable in this task, the shirt pixels are simply taken from below the face as in training. For each candidate identity, the Euclidean distance between the average R, G, and B values of the pixels on the persons shirt (a 3-tuple) and the average R, G, and B values of the specific identity’s shirt samples is calculated. If at least one candidate’s shirt distance is above a shirt threshold, then the candidate with the shortest distance is chosen as the identity of the person. If none are above the shirt threshold, the person is announced as a stranger.

First Phase In the first phase, we chose the four people and their shirts. We gave them strongly distinguishable shirt colors – red, green, blue, and yellow. Our robot correctly identified four of the five people. The stranger was misidentified as one of the known people.

We believe this error occurred specifically on the stranger for two reasons. First, the volunteer’s SIFT features matched many features of at least one of the known people. Second, the volunteer’s shirt was colored similarly to the person whose SIFT features were similar. With both the primary characteristic (the face) and the secondary characteristic (the shirt) testing as false positives, the person tracker did not correctly classify the stranger.

Of seven teams that attempted this task, some of which used commercial software packages, only one other received points by identifying at least four of the five people.

Second Phase The training of the second phase is the same as in the first, except the persons were chosen randomly by the committee. The testing is especially more challenging in the second phase. The five people (four known and one stranger) are not standing in a line anymore, but are instead randomly distributed throughout the home environment.

As in the *Lost and Found* task, we used a stochastic search to look for candidate people as recognized by positive identification from the face detection module. During the allotted time, the robot found one of the people and correctly identified him. No other team identified a single person during the second phase.

5.4 Open Challenge and Finals

Once all teams had attempted their specific tasks, each competed in what was called the *Open Challenge*. This consisted of a presentation and free-form demonstration. Going into this event, after receiving scores from the specific tasks, UT Austin Villa ranked third of eleven. A jury of the other team’s leaders ranked us second for the *Open Challenge*. The robot’s demonstration was a simplified version of the one performed in the Finals, so it will not be described.

Finals The top five teams competed in the Finals. Having ranked third in the specific tasks and second in the open challenge, UT Austin Villa advanced, along with Pumas from UNAM in Mexico, AllemaniACs from RWTH Aachen in Germany, RH2-Y from iAi in Austria, and Robot Cognition Lab from NCRM in France. The Finals were judged by a panel of trustees of the RoboCup organization, all well-known robotics researchers.

Before describing the demonstration itself, we begin with some motivation for the scenario we focused on. Accurate person-recognition will be essential

in any fully functional home assistant robot. Rigidly learning a person’s exact appearance at one moment will likely not be sufficient to identify him or her after a significant change in appearance (e.g. haircut, aging, etc.). A home assistant robot will need to be flexible, adapting to such changes.

Our scenario was designed to display our algorithm’s robustness and adaptability. Specifically, it shows person identification using shirt color as a secondary classifier in the absence of the primary classifier, the face. It also mimics the frequent occurrence of a human changing clothes, showing the robot adapt to this change in the secondary classifier’s input. Lastly, it shows that the Segway robot can effectively follow a recently learned person without markers, as we unfortunately were unable to show during the second phase of the *Follow and Guide a Human* task. The difference was that we used shirt colors which stood out from the colors of the background (as a white shirt did not).

Before the demonstration, we again presented a short talk about the robot and our algorithms. A video of the presentation and demonstration can be found at our team webpage, <http://www.cs.utexas.edu/AustinVilla/?p=athome>.

The demonstration involved two people, one with whom the robot intended to interact and another who was unrelated to the robot’s primary task (stranger). The person-recognition component of the demonstration algorithm is described in Section 4. At the beginning, the robot trains classifiers for the intended person’s face and shirt. It then follows the learned person based on only shirt color when his face is not visible, first with a green shirt and later with a red shirt. The Segway twice gets “passed” to a stranger, whose back is turned (i.e. face invisible) and is wearing the same shirt color. Each time, it follows the stranger until it can see his face. At that point, the face classifier returns a negative classification and supersedes the shirt classifier, and the robot announces that it has lost the learned person and turns away to look for him. Upon finding the original person based on a positive facial classification, it retraining the person’s shirt, subsequently stating whether the shirt color has changed.

The demonstration went smoothly, with only one noticeable flaw. At one point, the robot turned away from the red-shirted human and towards a mahogany set of shelves, again revealing the same limitation in our shirt classifier that hurt us in the second phase of *Follow and Guide a Human*.

6 @Home Final Results

The panel of judges scored the presentations and demonstrations of each finalist, determining each team’s final standing in RoboCup@Home 2007. We finished in second place (full results in Table 8). Of the top three teams, we had a couple of unique characteristics. Our at-RoboCup@Home team size of three people was half that of the next smallest team. We were also the only team in the top three that was competing for the first time. We were very successful as well in the specific tasks in which we competed. We received more points than any other team in the person-recognition task of *Who Is Who* and

Team	Final Score
AllemaniACs	256
UT - Austin	238
Pumas	217
RH2-Y	199
Robot Cognition Lab	190

Fig. 8: 2007 RoboCup@Home final results

accomplished all tasks that we attempted in the first phases of *Lost and Found* and *Follow and Guide a Human*.

7 Related Work

A variety of home assistant robots have been created in the past decade. Many exhibited impressive specific capabilities. Care-O-bot II [10] brought items to a human user and took them away in a domestic setting. It also functioned as a walking aid, with handles and an interactive motion system that could be controlled directly or given a destination. Earlier systems include HERMES [4] and MOVAID [5].

Various studies have been conducted on face recognition [19, 12, 2]. In contrast with these recognition methods, which rely on careful face alignment, we extract SIFT features [16] from faces, similar to work proposed in [3], and classify faces by counting the matching SIFT features in near real-time. To track faces through changing perspectives and inconsistent lighting, we augment a face classifier with another classifier (the shirt classifier). Previous work on integrating multiple classifiers has shown that integrating multiple weak learners (ensemble methods) can improve classification accuracy [17]. In [15], multiple visual detectors are co-trained to improve classification performance. This method merges classifiers that attempt to classify the same target function, possibly using different input features. In contrast, the classifiers we merge are trained on different concepts (faces and shirts) and integrated primarily by associating their target classes with one another in order to provide redundant recognition and dynamically revised training labels to one another. Tracking faces and shirts is a known technique [20], but previous work did not utilize inter-classifier feedback.

Person-following specifically has received much attention from researchers. A recent laser-based person-tracking method was developed by Gockley et al. [9]. Their robot Grace combined effective following with social interaction. Asoh et al. [6] couple face recognition with a sound detection module, using the sound's direction to indicate where the face is likely to be. A vision-based approach similar to our own was created by Schlegel et al. [18]. In their system, the robot also tracked shirts using color blobs, but the shirts had to be manually labeled in the training images. Some more recent approaches have used stereo vision and color-based methods to track humans [7, 11].

8 Conclusion and Future Work

RoboCup@Home will continue to move towards higher robotic autonomy and tasks that demand more effective navigation, object and person-recognition, and object manipulation. In order to keep pace with these rising challenges, our system will need to move towards general object recognition instead of relying on the ARTag system, and it will need to include mapping capabilities. Meanwhile, although person recognition is the current strength of our system, it can be improved by strengthening the face classifier's robustness to low-contrast conditions, improving its accuracy at identifying unknown people, and adding the ability to learn patterns of different colors on shirts.

The main contribution of this paper was the complete description of our Segway-based platform that performed successfully in the 2007 RoboCup@Home competition. Leveraging our main technical innovation of using co-training classifiers for different characteristics of a person (face and shirt), it was able to follow a person, distinguish different people, identify them by name, and ultimately combine these abilities into a single robust behavior, adapting to a person changing his or her clothes.

9 Acknowledgments

This research is supported in part by NSF CAREER award IIS-0237699 and ONR YIP award N00014-04-1-0545.

References

1. Interim state population projections. US Administration on Aging within the Population Division, U.S. Census Bureau, 2005.
2. P. Belhumeur, J. Hesphana, and D. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *PAMI*, 1997.
3. M. Bicego, A. Logorio, E. Grosso, and M. Tistarelli. On the use of sift features for face authentication. In *CVPR Workshop*, 2006.
4. R. Bischoff. Design concept and realization of the humanoid service robot hermes, 1998.
5. P. Dario, E. Guglielmelli, C. Laschi, and G. Teti. Movaid: a mobile robotic system for residential care to disabled and elderly people. *MobiNet Symposium*, 1997.
6. H. A. et al. Jijo-2: An office robot that communicates and learns. *IEEE Intelligent Systems*, 16(5):46–55, 2001.
7. V. E. et al. Active stereo vision-based mobile robot navigation for person tracking. *Integr. Comput.-Aided Eng.*, 2006.
8. M. Fiala. Artag, a fiducial marker system using digital techniques. *CVPR*, 2005.
9. R. Gockley, J. Forlizzi, and R. Simmons. Natural person-following behavior for social robots. In *Proceedings of Human-Robot Interaction*, pages 17–24, 2007.
10. B. Graf, M. Hans, and R. D. Schraft. Care-o-bot ii – development of a next generation robotic home assistant. *Autonomous Robots*, 2004.
11. H. Kwon, Y. Yoon, J. B. Park, and A. Kak. Human-following mobile robot in a distributed intelligent sensor network. *ICRA*, 2005.
12. D. Lee and H. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.
13. J. Lee, W. B. Knox, and P. Stone. Inter-classifier feedback for human-robot interaction in a domestic setting. *Journal of Physical Agents*, To appear.
14. J. Lee and P. Stone. Person tracking on a mobile robot with heterogeneous inter-characteristic feedback. *ICRA*, 2008.
15. A. Levin, P. Viola, and Y. Freund. Unsupervised improvement of visual detectors using co-training. In *ICCV*, 2003.
16. D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004.
17. R. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictors. *Machine Learning*, 1999.
18. C. Schlegel, J. Illmann, H. Jaberg, M. Schuster, and R. Worz. Vision based person tracking with a mobile robot. *BMVC*, 1998.
19. P. Viola and M. Jones. Robust real-time face detection. *IJCV*, 2004.
20. S. Waldherr, R. Romero, and S. Thrun. A gesture based interface for human-robot interaction. *Autonomous Robots*, 9(2):151–173, 2000.