Research Repository UCD

Title	Using mixed reality agents as social interfaces for robots
Authors(s)	Dragone, Mauro, Holz, Thomas, O'Hare, G. M. P. (Greg M. P.)
Publication date	2007-08-26
Publication information	Dragone, Mauro, Thomas Holz, and G. M. P. (Greg M. P.) O'Hare. "Using Mixed Reality Agents as Social Interfaces for Robots." IEEE, 2007.
Conference details	The 16th IEEE International Symposium on Robot and Human Interactive Communication 2007 (IEEE RO-MAN 2007), Jeju Island, Korea, August, 2007
Publisher	IEEE
Item record/more information	http://hdl.handle.net/10197/4526
Publisher's version (DOI)	10.1109/ROMAN.2007.4415255

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Using Mixed Reality Agents as Social Interfaces for Robots

Mauro Dragone, Thomas Holz, and G.M.P. O'Hare

School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland,
e-mail:(mauro.dragone, thomas.holz, gregory.ohare)@ucd.ie

Abstract—Endowing robots with a social interface is often costly and difficult. Virtual characters on the other hand are comparatively cheap and well equipped but suffer from other difficulties, most notably their inability to interact with the physical world. This paper details our wearable solution to combining physical robots and virtual characters into a Mixed Reality Agent (MiRA) through mixed reality visualisation.

It describes a pilot study demonstrating our system, and showing how such a technique can offer a viable alternative cost effective approach to enabling a rich social interface for Human-Robot Interaction.

I. INTRODUCTION

Artificial agents, either of a robotic or purely software genre, are gradually populating the human social space. The pervasiveness of robotic platforms, in the form of entertainment robots, household appliances, or assistive technologies, has acted as a catalyst for a large body of research exploring human-robot interaction (HRI). The ultimate goal of social robotics is to develop robots that socially interact with humans, working alongside people and acting as socially competent peers rather than mechanical servants depending on human supervision.

Crucially, the same objectives are addressed in software-only domains, for example, in supporting the development of virtual characters that assist Human-Computer Interaction (HCI), e.g. as game opponents or personal assistants (PAs). Profound technical and methodological differences exist when dealing with robotic and software domains, nevertheless there is an emerging recognition that the construction of truly social agents could benefit from the accommodation of both schools of thought. Specifically, the popularity of social agents has also opened the question of how heterogeneous agent societies, e.g. communities of agents (hardware or software), can be integrated in order to perform useful tasks in human societies.

This paper contributes a novel integration methodology, which combines a physical robotic body and a virtual character displayed through a mixed reality overlay. We believe that such a construction, which we call MiRA (Mixed Reality Agent), offers a viable alternative to merging the offering of physical and virtual agents into one socially competent agent.

After illustrating the principles and the software architecture behind MiRA, we report the lessons learnt from a pilot study which we conducted to gather user feedback, and to demonstrate our approach and our practical realisation.

II. SOCIAL AGENTS

The consensus in social agent research is that effective human-agent interaction greatly leverages the instauration of a human-style social relationship between human and agent. Dautenhahn's model of social intelligence [1] in human societies is characterised by the ability to recognise each other and develop and manage relationships between individual agents. Part of this intelligence is the ability of humans to understand and anticipate each other's intention [2],[3]. Humans convey this information through explicit communication as well as an array of non-verbal cues, such as anticipation (providing cues to an action before undertaking it), tactile interaction and facial expressions. Deictic spatial gestures (e.g. gaze, pointing...) are also a major component of joint visual attention between humans [4]. For these reasons, social agents generally employ key human-centric interaction modalities such as speech, gestures and even models of emotion in order to realise as natural a social interaction as possible

Although these issues are confronted in both hardware and software domains, there are important differences that need to be considered.

A. Robot Agents

Naturally, the biggest advantage over virtual agents is the robot's ability to interact with the physical world. They can not only sense and manipulate physical objects but they can also shake hands with humans and adhere to human social conventions in the way they share the physical space, e.g. joining queues or keeping an adequate distance during interaction with humans.

Due to physical constraints, however, expressional capabilities of most traditional robots remain restricted when compared to humans (see [5] for a review of robotic user interfaces).

Also, creating individual robots with a unique physical form and ascribing them with individual visible traits and behaviours is still a very costly and challenging engineering task. As a consequence, many robots are poorly equipped and simply lack the capabilities necessary for bootstrapping and maintaining adequate social interaction with humans, and for adapting to a human's changing expectations.

In addition, strong anthropomorphic paradigms in robotics overly increase people's expectations of the system's performance, and subsequently severely raise the behavioural

complexity required for a successful robot. Mori dubbed this problem the “The Uncanny Valley” [6]. His thesis is that the more closely a robot resembles a human, the more affection it can engender through familiar human-like communication references. However, there is a region in the design space where the robot becomes too similar but not quite perfect and thus appears uncanny and weird, with negative effects on human-robot social interaction.

B. Virtual Agents

Virtual agents, on the other hand, offer advantages of their own over their robotic cousins. They are capable of exhibiting a high degree of anthropomorphism [3] with highly expressive interfaces that are easily adjusted and personalised for each user at a fraction of the cost of a robotic interface. They are also capable of actions that are simply impossible in the real world, such as mutating their form, for example.

Virtual agents are inherently ubiquitous, as they can easily migrate through the network in order to interact with humans on different computers. They can also be carried around by the user, in handheld computers such as PDAs or other wearable devices. As such, they can be in an advantageous position for interacting with the user in a variety of circumstances, gathering user feedback and learning in order to better adapt to the user’s future needs.

On the other hand, virtual agents incur into problems where they do not meet the user through immersive interfaces. They are obviously limited in their interaction capabilities with the physical world - usually restricted to vocal communication with the user. It is not possible for the agent to leave the screen and wander freely or to interact with physical objects in the real world. In such a context, and with the divide between their 2D screen embodiment and the 3D physical world, it is also more difficult for them to attract and engage the user. For example, they cannot effectively point (or look) in a 3D physical surrounding.

C. Merging the Physical with the Virtual

A logical attempt to overcome these difficulties suggests trying to combine the two worlds, physical and virtual.

For example, in order to deliver context-sensitive behaviour, virtual agents may be integrated in ubiquitous infrastructures composed of a network of sensors and physical actuators such as mobile robots (e.g. [7]). A distinct approach is to form a single agent combining the physical abilities of a real robot and the expressivity of a virtual character. A first incarnation of this principle is the use of a virtual character visualised on a screen placed on top of the robot (e.g. [8]). However, this approach suffers some crucial drawbacks as the screen can only be seen from one angle and the virtual character and its actions are still limited to the screen.

III. MiRA: MIXED REALITY AGENTS

Contrary to traditional approaches, MiRAs can interact with both physical and virtual objects. This interaction transcends the boundaries between the real and the virtual world; the virtual avatar can, for example, point at real things, while

the physical robot can steer around virtual obstacles. As the virtual part is rendered individually on each user’s equipment, it offers the compelling possibility of personalising the content for each user.

Fig. 1 illustrates the MiRA concept by showing the typical setting behind its realisation. It shows a robot and a user wearing a rucksack containing a computer and a video see-through Head-Mounted Display (HMD) with a digital camera and microphone. Users and robots are also wirelessly interconnected. Through the HMD, the user can see the live scene captured by the camera, depicting the real robots augmented by superimposing synthetic imagery. These include the virtual character associated with the robot and other components of the MiRA interface.

Although the objective of MiRA is to interact with humans as one holistic agent, its implementation is the result of the collaboration between several distributed agent components: those controlling the robots (*Robot Agents*), those supervising the user’s HMD (*User Interface Agents*) and those controlling the virtual characters (*Avatar Agents*).

Each robot agent controls a particular robot by attending to its navigation (e.g. obstacle avoidance, autonomous exploration) and other functions such as object tracking. Each user interface agent (one for each user) supervises the interface between the user and other participants (both robots and humans, enabling multi-robot/multi-user scenarios), e.g. by processing user utterances and translating them into ACL directives to be directed to the robots. Finally, since the rendering of the virtual characters is performed on the user’s HMD, an instance of an avatar agent is in charge of a particular virtual character, controlling its behaviour while keeping it in line with the activity and intentions of the associated robot.

Fig. 1: MiRA is a mixed reality, portable, agent-based solution to HRI.

A. System Architecture

We will now briefly describe constituent components of the system architecture which underpins MiRAs.

Agent System: MiRAs are built upon Agent Factory (AF) [9],[10], a Java based, FIPA-compliant, and open-source BDI

Agent toolkit. In addition to the FIPA Agent Communication Language (ACL), MiRAs rely upon an XML-based service for the dissemination of low-level data (e.g. sensor readings, robot pose, tracking information...), which is also exploited for the dynamic discovery of peer observers.

Augmented Reality: AR overlay is realised by tracking the position and orientation of the HMD of the user in the coordinate frame of the observed robots (see Fig. 1). This information is then used to align the image of the virtual characters with the associated real robots. Tracking and rendering functionalities in the current implementation are based upon ARToolkit [11], a software library for the recognition and pose estimation of square markers within a camera image. We arranged five different markers upon the visible faces of a cube (visible in Fig. 1, on top of the Nomad Scout robot and below the snowman avatar) to make the robot traceable from all angles.

Virtual Characters: Virtual characters and behaviour animations (e.g. wave, point) are implemented via OpenVRML (<http://www.openvrml.org>), an open source C++ library for programmatic access to VRML models.

Hardware: The robot in the experiment is a standard differential drive Nomad Scout robot, equipped with a sonar ring, bumper sensors, and extended with a fixed digital camera (50° wide angle) and a simple ball pushing device. A Pentium 1.8 GHz acts as on-board computer. The user's wearable also includes a Pentium 1.8 GHz laptop and an i-Visor (www.personaldisplay.com) HMD, with 800X600 pixels resolution. Robot and user computers are connected via an ad-hoc wireless network (802.11b).

B. Advantages and Disadvantages

As a mixed reality, wearable, and agent-based social interface, MiRA is fundamentally different from screen-based robots or non-augmented social robots in general. As such, it presents a distinctive set of advantages and disadvantages that needs to be properly evaluated when considering any specific applicative scenarios.

Specifically, mixed reality offers tremendous opportunities for the personalisation of both form and associated behaviour. Among the possibilities, the virtual character can be overlaid as a virtual *clothing* that envelops the physical robot and acts as a *visualisation membrane*, de-facto hiding the robot's hardware. Alternatively, the virtual character can be visualised on top of the robot, as a bust protruding from the robot's body (centaur), or even figuring as the robot's driver (see Fig. 2). In every case, in contrast to a screen, the character is visible from all angles and is not subject to diminishing visibility at greater distances. In addition, the interaction capabilities of the robot are augmented by means of virtual limbs and the virtual character's head, giving him the ability to point and gaze in 3D. Also, the virtual character has a rich repertoire of gestures and facial expressions, which can be used to advertise its capabilities and communicate its state (see Fig. 2).

Notably, being virtual artefacts, these behavioural capabilities are not just limited to "natural" human-like forms, but can also include more complex effects, involving cartoon-like

animations, other virtual objects (e.g. a flashbulb, a question mark), or even mutation of the character's form.

Being a *wearable* solution, MiRAs are advantageous in applications with a high robot-to-user ratio (e.g. swarms), as the virtual components (e.g. the virtual character acting as social interface) are associated with the user's HMD. As a consequence, one single MiRA-enabled wearable device can augment the interaction capabilities of multiple simple robots (e.g. with no screen, head or arms). Crucially also, since robots and users are wirelessly interconnected and AR visualisation is based on tracking the robot's pose in respect to the user's HMD camera, this information is relayed to the robot so that the robot can recognise and track the position of the user and the user's gaze (approximated by the head's orientation) with great precision and without the need of other sensors (e.g. a camera).

MiRAs also offer enhanced personalisation capabilities. For example, MiRAs can decide to project different virtual characters to different users, e.g. taking the form of a cartoon character when interacting with a child. Another opportunity is for MiRAs to follow the user through its movements, e.g. by projecting the identity (and personality) of a particular character to different robots. For example, the user may take a flight and find at its destination a service robot combining a local robot platform with the virtual character he is already familiar with.

Instrumental to these degrees of flexibility is the *agent-based* nature of MiRAs. They can easily adapt to different users and different robots as there is no pre-defined coupling between the robot and the appearance and behaviour of its associated virtual character. Instead, thanks to their communication with the robot agent, both user interface agent and avatar agent can take context sensitive decisions in order to deliver a personalised and adaptive HRI interface. Finally, the interface agent can act as intermediate between the user and the robots he interacts with. Specifically, the interface agent can employ its knowledge of the robot's capabilities and of the state of the user to ease human-robot interaction, e.g. by requesting the robot to slow down if it moves too fast for an injured or elderly user. Crucially, this mediation happens in a manner transparent to the user and without disclosing users' private information to the robot.

An obvious disadvantage of employing MiRA's augmentation approach is the cumbersome and expensive hardware imposed on each user, which at the moment clearly hinders the deployment of MiRAs in applications with high user-to-robot ratio. However, this situation is on the verge of change as both HMDs and wearable computers become cheaper and less invasive [12]. The other inherent limitation is that artefacts created with mixed reality cannot be used for tactile interaction.

Finally, our current implementation suffers from problems specific to the particular solution adopted for the tracking necessary for AR visualisation. In particular, the marker based optic tracking, while being a cheap and easy solution, comes with its own shortcomings, most notably, the sensitivity to different lighting conditions and the failure to track

when the cubic marker is partially occluded or too close (e.g. less than 1m) to the user.

IV. OTHER RELATED WORK

Traditionally, AR has been used the context of tele-presence systems to enhance the user's understanding of the remote environment while controlling robotic manipulators [13], or mobile robots [14].

The area closest to our work concerns applications that involve humans working side by side with mobile robots. Contrary to work in social robotics, which emphasises human-like traits, the majority of these systems predominately focus upon the AR ability of visualising geometrical properties such as planned trajectories, sensory and world models. Compared to stationary computer systems, portable solutions enable human operators to gain an understanding and evaluate the state and the intentions of a robot system while moving around in the same environment the robot is operating in.

[15] describes an AR system for enhancing HRI with robotic swarms. While it is highly specialised toward a very interesting but narrow class of applications (search and rescue) it shows how AR can enhance robot interaction capabilities - in this case by allowing each robot in the swarm to point to the direction of the victim without the need of a physical arm.

Jonker and Caarls [16] investigate AR-based human-robot interaction also addressing search and rescue scenarios but by means of a knowledge system shared by robots and humans.

Giesler [17] developed a mobile AR system for the control of an autonomous mobile platform. The self-localisation of the robot is combined with the position and orientation of the user's HMD in the same frame of reference by means of multiple ARToolkit fiducials distributed in the working area. The user also has a pen terminating in a cube of ARToolkit fiducials that make its orientation visible from all angles. It can be used for pointing on the floor (e.g. for way-point navigation), just like moving a mouse on a desktop computer.

Young and Sharlin [18] at Calgary University present a taxonomy for organising and classifying the various interaction techniques offered by what they call an MR Integrated Environment (MRIE) - the combination of physical and MR interaction spaces. They also detail two interaction techniques, thought crumbs and bubblegrams, which are comic-style thought and speech bubbles represented (through AR visualisation) floating next to the robot that generated it.

As a test bed for this kind of interaction, they developed Jeeves [19], to our knowledge, the only other example of AR employed for human-robot social interaction. Jeeves combines a Roomba robot with a portable MR interface. In particular, Jeeves' main focus is investigating the use of cartoon-art, i.e. simplified and exaggerated facial expressions, in support of intuitive social interaction with humans. Such an approach is intended to offer insight into the robot's state while avoiding the uncanny valley represented by more realistic and human-like representations.

In contrast, our architecture [20], which predates Jeeves, offers a more versatile implementation of MiRAs as the use

of agents enables a flexible and adaptive control of both behaviour and appearance of the virtual element associated with each robot. It allows not only for a variety of avatar forms, from very stylised (e.g. as the snowman we initially employed, depicted in Fig. 1) to humanoid (but still cartoon-like, such as the one described in this paper and seen in Fig. 2), but also permits the personalisation of form and function based on user profiles and preferences.

Also relevant to our work are robotic systems involving wearable interfaces, which are used for enhancing HRI capabilities, although without recurring to AR visualisation. For instance, in the area of ubiquitous robots, wearable RFID tags have been proposed for aiding the detection and the localisation of users in the environment [21], as they can communicate relevant information (e.g. users' identity, location or other personal information) to the robots upon detection.

Direct wireless communication between robots and wearable interfaces can also be used in this sense, for instance in overcoming the limitations of today's speech recognition systems (e.g. as in NEC's Papero [22] and [7]). The wearable interface used in [7] also includes hardware to aid the robot in recognising user gestures and pointing.

V. EXPERIMENT

A. Experiment Design

Evaluating interaction with MiRAs is a difficult challenge, especially considering the added novelty factor caused by the use of AR and the limitations of the current equipment.

Our pilot experiment was therefore designed, foremost, to prove the feasibility of our approach, to demonstrate the capability of the system, and to gather users' expectations and reactions in the scope of informing future developments, both in the architecture and in the evaluation methodology. In particular, we were interested in the following questions:

- What capabilities do users ascribe to the agent?
- Can users assess the interaction modes of the agent?
- Are users able to understand the agent's behaviour?

Our experimental design is heavily influenced by prior experimental studies. We were especially interested in what Scholtz et al. define as the Bystander role [23], which assumes no prior knowledge of the user. We were interested in the mental model users would build when observing and interacting with a mixed reality agent.

Subjects were split into two groups. The first group interacted with the mixed reality agent, while the second (control) group interacted with just the physical robot. In both cases, user could only interact via a set of voice commands:

Look at me: The avatar first turns its head toward the user, and then the robot turns its body in the same direction and the avatar waves to greet the user. Now it is waiting for instructions and the avatar strains its ears if none are forthcoming (see Fig. 2a and 2b).

Bring me the ball: The robot turns on the spot until it locates the ball and the avatar points at it. Then the robot

moves toward it, grabs it and brings it back to the user. The avatar's gaze is fixed on the ball until the robot grabs it successfully. Then the avatar cheers before it gazes at the user while the robot brings the ball towards him (see Fig. 2e-g).

Explore: The agent wanders around, exploring its space. Upon meeting an obstacle, it moves off in a random direction.

Stop: The robot ceases any movement and all prior commands are cancelled.

Turn left/right: The avatar turns its head 90° to its left/right before the robot body follows.

The avatar also reacted to user utterances with a nod for understanding and a shrug for not understanding. It also saluted when given an instruction (see Fig. 2c and 2d). To the control group, only the robot actions were visible, of course.

Because voice recognition is inherently difficult without training the system for a particular voice, we wanted to limit distorted results as a consequence of incorrect voice recognition. We, therefore, adopted a Wizard-of-Oz approach, whereby an experimenter remotely controlled the robot via a keyboard according to the user's commands.

B. Participants & Procedure

We recruited an equal number of male and female participants, 20 in total and aged 18-42. All participants were from an academic background (i.e. student, PhD student, post-doc, or lecturer), mainly from within the School of Computer Science & Informatics. About half the participants indicated that they had used a VR system more than once before undertaking the experiment.

The robot was located in a confined 3m square, with a table separating it from the user. An orange ball was placed at the far end of the area.

Subjects entered the room and were first introduced to the mixed reality agent as seen through the HMD or to just the robot. For the control group no HMD was used. Participants were then asked to write down their expectations of the agent's capabilities and interaction modes as well as what specific voice commands or gestures, if any, they thought could be used.

After this pre-experiment questionnaire they were briefed on the actual set of voice commands. A short interaction period with the agent followed in which we asked the participants to try all voice commands in any order they liked and observe the agent's behaviour. These observations were gathered in an open-ended fashion after the interaction.

While users could freely move their head, they were requested to sit for the duration of the trial. Given the limitations of our current implementation, this setup guarantees a high degree of repeatability and reliability and is instrumental in demonstrating the system and assessing user expectations, despite being clearly limited in interaction capabilities offered to the user

VI. RESULTS

Experimental findings demonstrated the reliability of our system, as it always successfully tracked the user's position and never crashed throughout the duration of the trials. Results indicate that AR assists effective user interaction as every subject observed that the avatar was looking at them when requested to do so.

In general, people ascribed more capabilities to the mixed reality agent (see Table 1). While some referred specifically to the avatar's abilities (e.g. references to the avatar's head and arms), they also expected the agent to be able to tell them about itself, its capabilities and its environment. No participant in the control group reported anything similar. Participants in the MiRA group also ascribed more playful capabilities/attributes to the agent, expecting it to be able to play some other kind of game besides playing with the ball.

Other comments seem to confirm that the humanoid avatar creates a mismatch between user's expectation and effective robot's capabilities. For example, some users lamented that the robot was not fetching the ball while it was actively searching for it. The fact is that those users thought the robot should have been able to see the ball in situations in which the ball was out of the visual range of the on-board camera. For this reason we plan to try different avatars in future trials and also to employ a wide angle camera with pan-and-tilt functionalities on-board the robot.

While these heightened expectations might seem an obvious result, they bring with it important implications which ought to be factored into the development of mixed reality agents. The virtual form should match the capabilities of the agent as a whole. Otherwise, results might be skewed by people's disappointed expectations.

Voice was the predominant mode of interaction in all groups. Surprisingly, only 40% in each group made reference to this. Gestures of the head were mentioned by four subjects in the MiRA group, probably due to the existence of the

Fig. 1. (a) Saluting the user (b) Listening and waiting for instructions (c) User's utterance not understood (d) Instruction acknowledged (e) Found ball, pointing at it (f) Looking at the ball while grabbing it (g) Celebrating a successful grabbing manoeuvre.

TABLE 1. FREQUENCIES OF EXPECTED CAPABILITIES

Capabilities	Robot	MiRA
Track user movements	4	3
Tell about itself	0	3
Produce sound	3	3
Gestures	0	2
Play with ball	3	5
Play other games	0	4

HMD. Furthermore, other natural interaction modalities, like hand gestures and touch, were referred to more often in the MiRA group, whereas traditional human-computer interaction modalities (keyboard, joystick, remote), while not exclusive to, were more prominent in the robot group. As the Nomad Scout lacks obvious interaction capabilities, some people failed to see any possibility for interaction at all. A complete overview of the expected interaction modalities is given in Table 2.

No significant results were found to indicate a heightened awareness of the agent's behaviour. We believe that this is in part a fault of our methodology (especially the open-ended questionnaire), and partly, because of the limited capabilities of the agent. The vocal commands were fairly obvious (turn left/right, stop), and even more complex behaviours didn't require much understanding of the actual moment-to-moment behaviour (look at me, bring ball, explore).

Overall, people clearly enjoyed the experience and provided lots of positive feedback. One user even waved back at the agent and used phrases like "come on, good boy" while waiting for the agent to bring the ball.¹

VII. CONCLUSION

This paper has explored the challenge of Human Robot Interaction. It offers an innovative and pioneering approach to HRI entitled MiRA. It describes an architecture that enables the development of MiRAs and the effective overlay of augmented surfaces upon physical robots. Such MiRAs are mobile as are the users with which they interact. The architecture adopts an agent-oriented approach in support of user interaction (potentially personalised and offering adaptivity). The preliminary study presented in this paper offers valuable insights not only into the capabilities and the technical challenges behind our novel solution to HRI but also into understanding how to evaluate such systems.

Future work will focus on a proper and thorough evaluation and explore both the personalisation capabilities and the augmentation of the robot's spatial awareness through the AR tracking of the user and its gaze.

VIII. ACKNOWLEDGMENTS

This research is based upon works supported by the Science Foundation Ireland under Grant No. 03/IN.3/1361.

TABLE 2. FREQUENCIES OF EXPECTED INTERACTION MODALITIES

Interactions	Robot	MiRA
Voice	4	4
Head gestures	0	4
Hand gestures	1	3
Touch	0	1
Keyboard/Joystick/Remote	3	1

REFERENCES

- [1] K. Dautenhahn, "Socially Intelligent Agents", *AI Magazine*, vol. 19, no. 2, 1998, 127.
- [2] C. Breazeal, "Social interactions in HRI: the robot view", *Trans. on Systems, Man, and Cybernetics*, no. 34, 2004, 81-186.
- [3] C. Bartneck, "Affective Expressions of Machines" in *Proc. of the ACM Conference on Human Factors in Computing Systems (CHI '01)*, 2001, 189-190.
- [4] C.D. Kidd et al., "Effects of nonverbal communication in human-robot interaction" in *Proc. Of IEEE/RSJ international Conference on Intelligent Robots and Systems (IROS '05)*, 2005.
- [5] C. Bartneck and M. Okada, "Robotic User Interfaces" in *Proc. of the Human and Computer Conference (HC'01)*, 2001, 130-140.
- [6] Mori, M., *The Buddha in the Robot*. Charles E. Tuttle Co., 1982.
- [7] J. Suomela and A. Halme, "Cognitive Human Machine Interface for a Centauroid Robot" in *Proc. of the ServiceRob Conference*, 2001.
- [8] R. G. Simmons et al. "GRACE: An Autonomous Robot for the AAI Robot Challenge", *AI Magazine*, vol. 24, no. 2, 2003, 51-72.
- [9] G.M.P. O'Hare et al., "Agent Factory: Towards Social Robots", in *Proc. of the 1st International Workshop of Central and Eastern Europe on Multi-Agent Systems (CEEMAS'99)*, 1999.
- [10] B.R. Duffy et al., "Social Robotics: Reality and Virtuality in Agent-Based Robotics", in *Proc. of the Bar-Ilan Symposium on the Foundations of Artificial Intelligence: Bridging Theory and Practice (BISFAI '99)*, 1999.
- [11] H. Kato et al., *ARToolKit*. Technical report, Hiroshima City University, 1999.
- [12] C.P. Kleweno et al., "The Virtual Retinal Display as an Alternative Low Vision Computer Interface: Pilot Study", *Journal of Rehabilitation Research and Development*, no. 38, 2001, 431-442.
- [13] P. Milgram et al., "Applications of augmented reality for human-robot communication" in *Proc. IEEE/RSJ International Conference on Intelligent Robots and System (IROS 93)*, 1993, 3, 1467-1472.
- [14] S.W. Lawson, "Augmented reality for underground pipe inspection and maintenance", in *Proc. of the SPIE Conference on Telemanipulator and Telepresence Technologies*, 1991, vol. 3524, 98-104.
- [15] M. Dailyet al., "World embedded interfaces for human-robot interaction" in *Proc. of the 36th Annual Hawaii International Conference on System Sciences*, 2003, 125-130.
- [16] P.P. Jonker and J. Caarls, "Towards Tele-Presence: Combining Ambient Awareness, Augmented Reality and Autonomous Robots", *ERCIM News*, vol. 53, April, 2003, 29-30.
- [17] B. Giesler et al., "Using augmented reality to interact with an autonomous mobile platform" in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA '04)*, 2004, 1009-1014.
- [18] J. Young and E. Sharlin, "Sharing Spaces with Robots - An Integrated Environment for Human-Robot Interaction", in *Proc. of the International Symposium on Intelligent Environments (ISIE '06)*, 2006.
- [19] J.E. Young et al. "Robot expressionism through cartooning" in *Proc. of the ACM/IEEE international conference on Human-robot interaction*, 2007.
- [20] M. Dragone, T. Holz and G.M.P. O'Hare, "Mixing robotic realities", in *Proc. of the 2006 International Conference on Intelligent User Interfaces (IUI '06)*, 2006.
- [21] B.K. Kim, et al. 2005, "Design of Ubiquitous Functions for Networked Robots in the Informative Spaces," in *Proc. of the 2nd International Conference on Ubiquitous Robots and Ambient Intelligence*, 2005.
- [22] Y. Fujita, "Personal robot PaPeRo," *J. Robot. Mechatron*, vol. 14, no. 1, 2002, 60-63.
- [23] J. Scholtz and S. Bahrami, "Human-robot interaction: development of an evaluation methodology for the bystander role of interaction" in *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, 2003, 4, 3212-3217.

¹ Videos of the user trials are available at <http://sosaa.ucd.ie>