

Provided by the author(s) and University College Dublin Library in accordance with publisher policies. Please cite the published version when available.

Title	Where robots and virtual agents meet : a survey of social interaction across Milgram's Reality-Virtuality Continuum
Author(s)	Holz, Thomas; Dragone, Mauro; O'Hare, G. M. P. (Greg M. P.)
Publication Date	2009-01
Publication information	International Journal of Social Robotics, 1 (1): 83-93
Publisher	Springer
Link to publisher's version	http://dx.doi.org/10.1007/s12369-008-0002-2
This item's record/more information	http://hdl.handle.net/10197/1241

Downloaded 2012-10-17T14:19:38Z

Some rights reserved. For more information, please see the item record link above.



Where Robots and Virtual Agents Meet

A Survey of Social Interaction Research across Milgram's Reality-Virtuality Continuum

Thomas Holz · Mauro Dragone · G.M.P. O'Hare

Received: date / Accepted: date

Abstract Traditionally, social interaction research has concentrated on either fully virtually embodied agents (e.g. embodied conversational agents) or fully physically embodied agents (e.g. robots). For some time, however, both areas have started augmenting their agents' capabilities for social interaction using ubiquitous and intelligent environments.

We are placing different agent systems for social interaction along Milgram's Reality-Virtuality Continuum—according to the degree they are embodied in a physical, virtual or mixed reality environment—and show systems that follow the next logical step in this progression, namely social interaction in the middle of Milgram's continuum, that is, agents richly embodied in the physical and virtual world.

This paper surveys the field of social interaction research with embodied agents with a particular view towards their embodiment forms and highlights some of the advantages and issues associated with the very recent field of social interaction with mixed reality agents.

Keywords Social interaction · Mixed reality agents · Human-robot interaction · Human-computer interaction

1 Introduction

In today's digital society, an ever-growing amount of human activities relies on digital technology. Trends

such as inexpensive internet access and the diffusion of wireless computing devices have made ubiquitous or pervasive computing a practical reality that augments the normal physical environment and supports the delivery of services to human users anytime and anywhere. A lot of interfaces for these environments are built on the idea that a social interface, that is, an interface availing of human-like social cues and communication modalities, is the most natural and thus most effective way for humans to interact. This belief is rooted in the assumption that humans react socially to computers [56] and has been the driving force behind a big body of research in the areas of Human-Computer Interaction (HCI) and Human-Robot Interaction (HRI).

But while in the past software agents and robots have usually been seen as distinct artefacts of their respective domains, the modern conception is, in fact, to consider them as particular instances of the same notion of agent—an autonomous entity capable of reactive and pro-active behaviour in the environment it inhabits. From this point of view, the concept of embodiment, intended as the structural coupling between an agent and its environment (e.g. [68]), provides the common ground upon which different strands of agent-related research can be analysed and compared.

Indeed, by placing social interaction research on Milgram's Reality-Virtuality Continuum [43] (see Fig. 1), according to the nature and extent of their embodiment in a physical or virtual environment, we can view robots and virtual agents as embodied at the extremes of said continuum.

These extremes are the areas where most of the traditional social interaction research to date has taken place, and where we first turn our focus, giving a brief overview of the commonalities and differences of traditional HCI and HRI research in Section 2.

T. Holz · M. Dragone · G.M.P. O'Hare
CLARITY Centre for Sensor Web Technologies, School of
Computer Science & Informatics, University College Dublin,
Belfield, Dublin 4, Ireland
Tel.: +353-1-716-2483, Fax: +353-1-269-7262
E-mail: thomas.holz@ucd.ie, mauro.dragone@ucd.ie,
gregory.ohare@ucd.ie




Fig. 1 Milgram’s Reality-Virtuality Continuum (adopted from [43])

However, as Section 3 details, both the virtual and robotic domain have started moving away from the edges for quite some time now by embodying their agents in ubiquitous environments, that is, environments that are “augmented” with digital sensors and devices. This shift has occurred with two main objectives in mind. On the one hand, the ubiquitous environment is supposed to augment the agent’s interaction capabilities, feeding it information about the user’s whereabouts and state to assist social interaction. On the other hand, the agent provides the user with a social interface that acts as a representative of the services the intelligent environment offers.

Following social interaction research from both sides away from the edges and further along Milgram’s continuum naturally leads us to the very centre of the spectrum, where social agents are endowed with mixed reality bodies and placed into mixed reality environments. Only very recently have researchers started to look at this area, as Section 4 will show, but they have already identified some unique possibilities and challenges for social interaction.

2 Social Interaction with Robotic and Virtual Agents—Research at the Extremes of Milgram’s Continuum

When investigating human-agent social interaction (see Fig 2 for examples of real and virtual social agents), both HRI and HCI share many similarities, as they can both draw on insights into human social research.

Fong et al. (2003) use the term “socially interactive robots” to describe robots for which social interaction plays a key role. For Fong et al. these robots are important in domains where robots must exhibit social interaction skills, either because such skills are required for solving specific tasks, e.g. as in scenarios where robots are envisioned working shoulder to shoulder with humans [1], or because the primary function of the robot is to interact socially with people, e.g. as in companionship [16] and educational robots [59].

The specific origin of a robot’s social skills, such as their role within the robot’s cognitive apparatus, remains an open issue which usually also depends on the particular research emphasis, i.e. if in pursuance of a robot-centred or a human-centred perspective [17].

However, there is enough evidence to suggest that these robots need to exhibit a certain degree of social intelligence, for the way they manifest their awareness and react to the presence of humans, in order to be accepted as social peers [47], or simply tolerated within humanly populated environments [53].

Similar issues are being investigated within the 3D and virtual reality communities, as a large body of works now shares an interest in the incorporation of virtual characters into virtual and augmented reality environments. Human-like virtual characters (virtual humans) are being used with success as virtual representatives of human users in virtual conference applications [61], or as fully autonomous agents inhabiting virtual worlds to enhance the user’s experience and ease his interaction with the virtual world. Such characters can make the interaction more engaging and make the user pay more attention but they can also require more effort to interact with the system, e.g. in educational and training applications [11]. A much appreciated feature in the latter type of applications is that virtual humans can provide pedagogical assistance that can be tailored to the needs and preferences of the learner [7].

Studies focusing on how the appearance of virtual characters can affect cooperation, change attitudes, and motivate users [57, 67] indicate that humans treat them as social partners and, in particular, that many of the rules that apply to human-human interaction carry over to human-agent interaction.

The result is that, despite technical and methodological differences between dealing with robotic and virtual domains, today a large number of issues behind the construction of successful social agents cross the boundaries of agent species.

What distinguishes all the research in socially intelligent agents is the emphasis given to the role of the human as a social interaction partner of artificial agents and, subsequently, to the relevance attributed to aspects of human-style social intelligence in informing and shaping such interactions. 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 [15] in human societies is characterised by the ability of recognising each other and developing and managing relationships between individualised agents. Effective human-agent interaction is, by the same token, based on the social relationship and the mutual understanding between user and agent, improving predictability and trust between the two.

One factor in this human-agent relationship is the agent’s identity who makes it possible for a user to dis-




Fig. 2 Examples of real and virtual social agents (from left to right: Kismet, Minerva, Max and Steve)

tinguish his agent from a dozen others like him and to recognise its persona over time. It has been envisaged that users have an expectation that the agent should evolve in terms of the character and interaction style with the relationship developing and emerging whereby the agent becomes empathetic to the user needs [16]. Indeed research indicates that we tend to ascribe anthropomorphic attributes to robotic entities and as such expect them to behave in a socially intelligent and individualised manner to meet our diverse requirements [16].

Another factor that is instrumental in human-human interactions is the ability to understand and anticipate each other’s intentions and emotions [14]. We convey this information through explicit communication (natural language) as well as an array of non-verbal cues, tactile interaction, postures, facial expressions, and gestures [32]. To be effective members of human societies, synthetic agents need to be able to use and understand the same array of non-verbal cues [4].

As a result, many agent architectures in both domains are developed based on deep models of human cognition and social competence. In this type of research, pure rational mechanisms, such as the Belief-Desire-Intention (BDI) paradigm [55], can be extended with anthropomorphic psychological models, e.g. of emotions, attention, creative problem solving, memory retrieval, decision making and learning. In doing so, such architectures build upon the human attitude of employing an intentional stance toward the autonomous agents. Specifically, they improve common understanding between human users and agents, and hence their social interaction capabilities, for the way they fulfil user expectations by adhering to human social norms. The attention mechanism of the robot Kismet, for example, is modelled on that of a human infant [8].

However, research in user modelling is not only important to help humans understand synthetic agents but also to enable agents to reason about the human users with whom they interact. By availing themselves of intentional and/or emotional models of human users, synthetic agents can adapt to, and predict their behaviours and possibly change attitudes, and also ac-

commodate different users, e.g. with varying skills, experience, and knowledge.

To build up these user models, both areas can leverage functionalities intended to recognise social cues in humans, e.g. through vision-based face detection and recognition systems, and techniques intended to classify facial expressions. Whether applied to a video-feed from the robot’s eyes [8, 9] or fixed cameras used by screen-based virtual agents [33], the same techniques are employed to identify human users and also to understand their emotional state.

On a cautionary note, however, it is important to remain aware of the differences between physical and virtual agents, as social interaction with robots and virtual characters is by no means equal. Naturally, one of the most important defining characteristics of robot agents, and their biggest advantage over virtual agents in general, is their physical presence and their ability to interact with the physical world and the human whereas virtual agents are inherently constrained in their interaction capabilities with the physical world—usually restricted to vocal communication with the user. While this might seem an obvious statement, it should nevertheless be kept in mind. Indeed, several studies ([64] and [31] among others) have indicated that humans prefer a real robot to an animated version in one-on-one interactions precisely because their physical nature evokes a higher sense of presence in the user, making them more trustworthy and engaging.

Robots are also free to move within the physical environment. The implications of this simple fact become clearer when one considers, for instance, mobile robots employed in museums. They may act as a guide for physically present visitors [46] or as a tele-presence body for remote visitors [40]. They may also do both, such as in the case of MINERVA [62], a robot that guides visitors through the Smithsonian National Museum of American History and that can also be used to explore the museum remotely via a web interface. Through their physical body, these robots may proactively engage the user and lead him towards a particular exhibit.

Virtual agents in a similar setting, instead, require the user to initiate the action and approach the agent [37], since they are unable to wander freely or to interact with physical objects in the real world. In particular, the divide between the 2D embodiment of screen-based agents and the 3D physical world makes it more difficult for them to attract and engage the user in the 3D physical environment he inhabits, as they are visible only from a limited angle and cannot effectively point (or gaze) in the 3D physical surrounding.

Virtual agents, on the other hand, are not only capable of actions that are simply impossible in the real world, such as mutating their form [41], they can also exhibit a high degree of anthropomorphism [37, 30] with highly expressive interfaces that are easily adjusted and personalised for each user at a fraction of the cost of a robotic interface. While many robots have been built with the explicit goal of investigating human-robot social interaction (see [23] for a good review), their expressional capabilities still remain restricted when compared to virtual agents and thus fail to convey subtler meanings of intention or emotion [5]. Indeed, creating individual robots with a unique physical form and ascribing them with individual visible traits and behaviours is still far from the capabilities of current production lines.

First steps towards building realistic human-like companions with rich visual expressiveness have been taken (e.g. [29, 27]), but they still suffer from limitations and high costs, mainly due to the vast number of actuators needed to animate the robot’s face. Therefore, for the moment, the vision of lifelong robotic companions, for which personalisation has been named a key factor [16], still remains out of reach.

Finally, one obstacle to be wary of in building any new agent, physical or virtual, is that a strong anthropomorphic paradigm—while necessary in works that build upon its evocative power—can overly increase people’s expectations of the system’s performance, and subsequently severely raise the behavioural complexity required for a practical application to succeed. Mori dubbed this problem the *Uncanny Valley* [44]. 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 perfect, and thus appears uncanny and weird, with severe negative effects on human-robot social interaction (see Fig. 3).

3 Moving Closer Together—Social Interaction in Ubiquitous Environments

When leaving the extremes of Milgram’s continuum, we are confronted with agents embodied in ubiquitous environments. These environments extend our normal physical surroundings with embedded computational devices and act as service providers and shared information spaces for the human user.

In such a context, the role of the agents embodied in ubiquitous environments is multi-faceted. From the user’s perspective, the agent is acting as a intelligent user interface to the environment, providing them

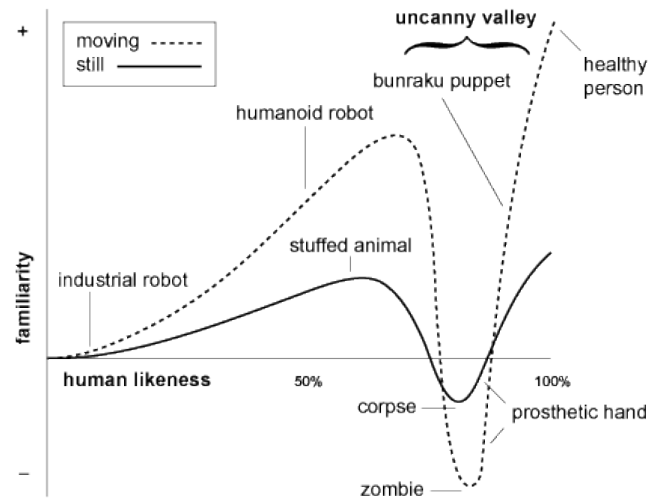


Fig. 3 Mori’s Uncanny Valley (adopted from [39])

with the services they need, anytime and anywhere (see Fig. 4 for examples of ubiquitous agents).

The Ubiquitous Robotic Companion (URC) [26], for example, guards its owner’s home, cleans their rooms and reads to their kids. The URC also provided one of the first mass deployment examples of ubiquitous robots with roughly one thousand URCs distributed to households in Seoul and the vicinity. NEC’s robot PaPeRo™ (Partner-Type Personal Robot) [24] is another example of this interface role, as its network capabilities allow it not only to communicate with other PaPeRos, but also with PCs, PDAs, and with other electrical appliances in the home. Thus, a PaPeRo can check the user’s e-mail, tune the TV onto the user’s favourite channel or, as the URC, access the Internet, e.g. for retrieving stories to narrate to children.

In its child care version, PaPeRo also presents some social attributes thanks to the advanced interaction abilities that are enabled by an array of sensors, including touch sensors, sonar, directional microphones, and cameras. With these sensors, PaPeRo can act as a teacher by locating and identifying the children, taking attendance, imparting lessons and even quizzing them.

The social interaction aspect is also investigated by Philips Research with the iCat robot [10], a research prototype of an emotionally intelligent robot that provides an easy to use and enjoyable interaction style for ubiquitous environments. In particular, iCat can communicate information encoded in coloured light through multi-color LEDs in its feet and ears, speak natural language, and also use facial expression to give emotional feedback to the user.

One of the key research questions investigated within the iCat project is to find out whether these expressions and capabilities can be organised to give the robot



Fig. 4 Examples of ubiquitous agents (from left to right: URC, Ubibot, PaPeRo and iCat)

a personality and which type of personality would be most appropriate during interaction with users. These studies have shown that mechanically-rendered emotions and behaviours can have a significant effect on the way users perceive—and interact with - robots. Moreover, they have also shown that users prefer to interact with a socially-intelligent robot for a range of applications, compared to more conventional interaction means [6].

From the agent’s perspective, being embodied in a ubiquitous environment implies the augmentation of its interaction capabilities with the sensors and services embedded within the ubiquitous environment, for example, to acquire information about the users. A preferred approach to this end is to equip the latter with wearable interfaces.

For instance, wearable RFID tags have been proposed in conjunction with ubiquitous robotic systems [26] to aid the detection and the location of users in the environment, as the RFID tags are sensed by the ubiquitous infrastructure and the relevant information stored in it (user identity and location) is communicated to the robots. Notably, such an approach can substitute conventional (and more difficult to realise) robot-centric perception as even simple robots (e.g. with no camera) can effectively recognise and track the humans in their environment.

Direct wireless communication between robots and user wearable setups can also be used in this sense, for instance, for overcoming the limitations of today’s speech recognition systems. Users of PaPeRo, for example, employ wireless microphones for cancelling the impact of changing background noise.

All these wearable interfaces, as well as the LEDs used in iCat, break social norms—for the way they impose unnatural constraints on the human partner or deviate from natural, human-like capabilities—in order to facilitate and augment human-robot interaction. But while some of these constraints are envisioned to vanish in the future (e.g. as speech recognition software improves), iCat deliberately uses its LEDs to convey information about its state in a way that humans cannot.

Another scenario enabled by ubiquitous environments is to have software agents migrating between different devices and platforms in the environment, such as PDAs or other handheld devices. As virtual agents can easily move between different software environments and adapt to different circumstances, e.g. different computational power or different interface capabilities, they are natural candidates for delivering this type of interface, for instance in the form of personal assistant agents. In this way, they are able to maintain the connection with the user through his daily activities, thereby adopting an advantageous position for interacting with the user in a variety of circumstances, gathering feedback and learning in order to better adapt to the user’s future needs.

Since it is often possible to separate a robot software system from the robot hardware, a ubiquitous robot can also be seen as a software agent in charge of robotic sensors and actuators. Agent-based ubiquitous robots, in particular, build on this fact, enabling applications in which the software can migrate between platforms and thus take control of different physical robotic bodies.

Using migration, software agents in control of robots can then change their body in order to find the most appropriate form for any given situation/task. From this perspective, physical robots simply become another observation point from which software agents can follow the user’s activities and movements. This also increases the number of possible agent-user interaction scenarios, which can then also include a variety of real world situated interaction, well supporting, for example, applications in which the agent needs to learn to adapt to the user.

Notable examples of migration in agent-based ubiquitous robots are investigated within the Ubibot system [34] and the Agent Chameleons project [22, 51, 41]. The Ubibot system is structured as a collective integrating three different types of robots: Sobots (mobile software agents), which can migrate and take control of Embots (embedded devices and sensors) and Mobots (mobile robot). The usability of this concept has been demonstrated with a system developed at the RIT Lab., KAIST, in which a Sobot associated with a 3D synthetic dog character, called Rity, controls an Embot that consists of a face recognition system employing a USB camera. By combining the Sobot reasoning and learning capabilities with the sensing capabilities delivered by potentially different Embots, Rity can follow the user’s movements and can interact in a personalised manner with the user wherever he goes.

The Agent Chameleons project [49, 50, 21] at UCD investigates similar scenarios in the context of personal assistant agents migrating across different user interface

devices and crossing the software/hardware divide. The reasoning and migration capabilities of agent chameleons are based on Agent Factory (AF) [48], an agent programming framework already used in the context of ubiquitous environments [45]. One of the applications considered within the project is the development of personal travel-assistant agents. These can take the form of a virtual character operating on a personal computer that helps the user to find and book a flight to a certain conference. The agent may then migrate to the user's PDA so that it can effectively travel with the user, helping with flight connections and attending tasks such as filtering e-mails or negotiating the transit of the user through the airport (e.g. giving the user's information at check-in gates). Once at the hotel, the agent may then leave its computational environment, by migrating from the user's PDA to a computer in control of a service robot that is gracefully provided by the (futuristic) hotel. By possessing such a physical body, the agent will then be able to continue assisting the user, this time by also using physical capabilities.

Notably, this scenario is substantially different from the one in which the user is simply assigned a service robot on his arrival at the hotel. In the scenario foreseen by the Agent Chameleons project, the robot would effectively act as the user's personal assistant. By having access to his personal data, the robot would, for example, be able to serve the perfect coffee without having to query the user about his preferences. Although such results may be achieved with the personal assistant agent communicating the necessary instructions to the robot, a solution based upon migration may be more efficient and offers the advantage of not having to release more private user details. In addition, an Agent Chameleon goes a step further by making sure that the user knows that all the agents are a mutated form of his personal assistant. For instance, the Agent Chameleon will use the same synthetic voice or will display other characteristic personality traits that are independent from his specific bodily form, thus preserving the advantages of the familiar relationships between the user and its assistant.

4 Human-Agent Interaction in Virtual and Mixed Reality Environments—Where Robots and Virtual Agents Meet

While ubiquitous agents can be seen as a promising first step towards overcoming the physical/digital divide between robots, humans, and virtual agents, they still suffer some of the same limitations as their non-augmented counterparts, e.g. limited expressive capabilities or confinement to a screen. A logical progression



Fig. 5 Examples of social interaction in mixed reality environments (from left to right: Ubiagent, Welbo, Nexus, and PECA)

along Milgram's continuum then is to employ mixed reality technology to enable virtual agents to share the same 3D space as humans. By immersing human users in a mixed reality space that superimposes virtual elements in the same perspective as the real scene, they can interact across the whole spectrum of Milgram's continuum.

In robotics, these techniques have traditionally been used in the context of tele-presence systems to enhance the user's understanding of the remote environment (e.g. [42]) or as a development tool for robot systems, by allowing the user to view sensor data, such as sonar and laser scans, in context with the real world (e.g. [12]). But while robotics research has mostly concerned itself with augmenting the robot's sensing or other technical capabilities, MR interfaces have long helped virtual characters to escape purely virtual environments and to socially engage humans in their physical environments (see Fig. 5 for virtual social agents in mixed reality environments).

The conversational agent Welbo [2], for example, is an interface agent that guides, helps, and serves the user in an *MR Living Room* to visually simulate the location of virtual furniture. Pedagogical Embodied Conversational Agents (PECA) [18] are similar virtual agents that apply proven pedagogical techniques to interact with humans in various learning scenarios, including outdoor student tours of historical buildings. Virtual agents in these types of application extend the user interface paradigm by acting as mediator between the human users and the mixed reality environment for the way they can act on virtual artefacts (the virtual furniture in Welbo or architectural 3D models in PECA).

As user interface agents, they need to adapt to particular users, e.g. to improve and accelerate human learning performance by tailoring their behaviour to the user's personal interests and learning strengths [7]. More often, these agents need to demonstrate context aware intelligence for integrating with pre-existing ubiquitous infrastructures but also for reacting to the user's behaviour and, more generally, for behaving in a believable/realistic manner in the mixed reality environment they share with human users. For this purpose, an in-



Fig. 6 Examples of robotic mixed reality agents (from left to right: MiRA, U-Tsu-Shi-O-Mi, and Jeeves)

creasing number of these systems are employing sophisticated agent control architectures.

In particular, the Nexus [52] and the UbiAgent [3] frameworks, both depicted in Fig. 5, demonstrate how virtual agents, equipped with BDI control systems, can provide the reasoning apparatus for creating believable characters that are responsive to modifications and stimuli in their environment, but are also proactive and goal-oriented. Both systems demonstrate the capabilities of creating applications with multiple virtual agents and multiple users, where virtual agents possess a high degree of autonomy and can watch the mixed reality environment, in particular by sensing each other’s relative positions as well as the movements of other physical objects.

On the other hand, in order to be employable as believable social characters in real environments, virtual characters, in systems such as Nexus and UbiAgent, need to perceive not only the position (for gaze behaviour [63, 54] and the placement of the virtual agent [2]) but also the state of the humans they interact with [35], e.g. their moods, given, for example, by the state of the conversation and their gestures.

Since these agents, in contrast to robots, lack a physical body, they need to cooperate with a ubiquitous infrastructure, e.g. wireless sensor networks including cameras and localisation sensors deployed in the environment where human-agent interaction takes place.

By combining a physical robot with a virtual character, a number of recent applications, depicted in Fig. 6, seek to overcome these limitations while at the same time taking full advantage of their mixed reality nature. By giving a virtual social interface to a physical robot and a physical body to a virtual character, the agents in these systems exhibit tangible physical presence while offering rich expressional capabilities and personalisation features that are complex and expensive to realise with pure hardware-based solutions.

These robots represent a characteristic example of embodiment in ubiquitous environments, in which the agents themselves are composed of virtual and real com-

ponents. As such, they are referred to in the rest of this paper as Mixed Reality Agents (MiRAs).

Dragone et al.’s [19] MiRA system is the first example to combine a MR character with a physical robot platform to provide an expressive social interface. In particular, thanks to its agent-oriented software engineering approach, MiRA offers a versatile implementation of the real/virtual merging, which enables a flexible and adaptive control of both behaviour and appearance of the agent.

The focus of Young et al.’s Jeeves project [66]—which similarly combines a robot (the iRobot Roomba vacuum cleaner) with a cartoon-like character—is instead to investigate 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. Young et al. point out that cartoon art techniques can augment and complement existing robotic interaction metaphors such as speech and gestures, and can capitalise on the physical nature of the robot.

Finally, another recent example in this class of applications is the U-Tsu-Shi-O-Mi Virtual Humanoid [58], which maps a humanoid avatar onto a robot’s anatomically correct, green-cloth surface. The result is a 3D mixed reality humanoid character that the user can touch and interact with.

Together, these applications showcase the advantages of employing mixed reality to combine physical robot platforms with virtual characters. Among the possibilities demonstrated in these systems, the virtual character can be overlaid as a form of 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, or even figuring as the robot’s driver—as in some demonstrations of the MiRA system [28]. In every case, in contrast to robots with virtual characters visualised on a screen placed on top of them, such as GRACE [60] and VALERIE [25], the mixed reality characters are visible from all angles and are not subjected to diminishing visibility at greater distances.

Obviously, compared to the U-Tsu-Shi-O-Mi Virtual Humanoid, MiRA and Jeeves take greater advantage of their virtual components, as they are free from the engineering effort of realising sophisticated mechanical interfaces that mirror the movements of their virtual characters. For example, in MiRA the virtual character has the ability to point and gaze in 3D without having to construct any physical counterpart. In this

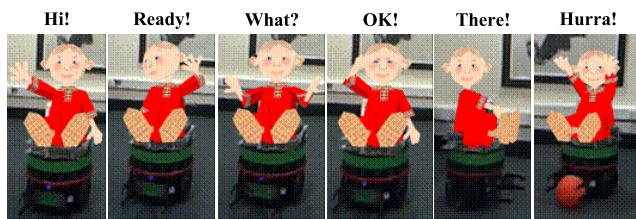


Fig. 7 Gestures and facial expressions in the MiRA system

manner, it can overcome the inherent limitations of screen-based solutions, as well as provide a rich repertoire of gestures and facial expressions, which can be used to advertise its capabilities and communicate its state (see Fig. 7).

Being wearable solutions, these systems are thus also advantageous in applications with a high robot-to-user ratio, as a single wearable interface can augment the interaction capabilities of multiple simple robots (e.g. with no screen, head or arms) and, possibly, even portray different characters to different users at the same time.

Notably, being based on virtual artefacts, behavioural capabilities in MiRA and Jeeves are not just limited to “natural” human-like forms, but can also include more complex effects involving other virtual objects and cartoon-like animations. In MiRA, for example, a flash bulb can be displayed to signal the robot’s perceptions of certain events, whereas Jeeves can leave *thought crumbs* behind, virtual artefacts that can be tagged to real objects.

While these advantages are inherently associated with the wearable and mixed reality nature of these applications, Dragone et al.’s MiRA implementation adds a coordination dimension which is instrumental to maintaining the expressivity of the mixed reality medium without losing the tremendous opportunities for ubiquity and personalisation associated with its virtual component. Specifically, despite having different autonomous software agents in charge of real and virtual components, MiRA creates the impression of one holistic agent to the human by leveraging on the coordination capabilities of the underlying multiagent platform [13]. This enables it to animate the virtual character in line with the behaviour and the state of its robotic counterpart and, correspondingly, to make the robot’s physical behaviour responsive to the state of the associated virtual character.

By processing the user’s state and input, the latter may easily access (or indeed learn) useful personal data about the user that can be consequently used to explicitly negotiate the MiRA’s persona—intended as the aggregation of its form and observable behaviour—to better suit the user.

5 Discussion

As Fig. 8) shows, the research reviewed in this paper spreads across the complete width of the Reality-Virtuality continuum. What type of embodiment one should choose for a particular applications is a very tricky question, however.

As mentioned before, a lot of studies suggest that robots are generally preferred over virtual characters ([64, 31] among others). However, results are far from black and white if one takes into account the nature of the task. Yamato et al. [65], for example, have performed an experiment in which users were confronted with the same task in a physical and in a virtual setting, each time assisted by a virtual character or a robot. They found the virtual character to be more effective in the virtual world and the robot to be more effective in the real world.

To further complicate matters, Lee et al. [38] report findings to indicate that physical embodiment alone raises user expectations. Their results show that a robot can be more negatively perceived than a virtual character even in a physical setting if tactile interaction is restricted.

There are comparably little studies, however, when we leave the ends of the continuum. Wagner et al. [63] found no discernible difference between a virtual agent on a screen and one floating in the user’s physical space through AR visualisation but they also remark that the agent did not take advantage of the spatial abilities of the AR agent. This bears some resemblance to the robot not being able to physically interact and thus should not be taken as proof for the equality of screen and AR characters without further studies.

Moving even further towards the middle of the spectrum, experimental findings become increasingly hard to find. Most studies of ubiquitous robots and agents treat the agent as if it were purely virtually or physically embodied and ignore the ubiquity aspect. In particular, if a social agent is employed as an interface to a ubiquitous environment with dozens of sensors and actuators, what does the user perceive as the agent’s body and what as devices the agent avails itself of?

Mixed reality agents, square in the middle of Milgram’s continuum and the latest addition to the wider area of human-agent interaction, have so far only reported findings of preliminary user studies ([20] and [66] for MiRA and Jeeves, respectively), and no dedicated studies to inform the grounding of mixed reality agents in a theoretical framework—or even just to validate the feasibility of the approach—currently exist.



Fig. 8 Placing agent applications on Milgram's continuum (top row from left to right: Minerva, iCat, MiRA, PECA, Max; bottom row from left to right: Kismet, PaPeRo, Jeeves, Ubibot, Steve)

6 Conclusion

This paper provided a comprehensive review of social interaction with agents embodied as robots, virtual agents and across domains (mixed reality), and illustrated how, when viewed from an embodiment perspective, there currently seems to be a progression towards the centre of Milgram's Reality-Virtuality Continuum.

Applications like the Virtual Humanoid U-Tsu-Shi-O-Mi, Jeeves, and MiRA are the first examples of social agents that populate the centre of this spectrum. Due to their specific combination of robotic and virtual components they can take full advantage of the dual nature of mixed reality embodiment and navigate around a lot of engineering challenges commonly associated with mixed reality environments.

In particular, MiRA explicitly allows for personalisation of the system thanks to its multiagent implementation. By combining a generic robot body with an individual, flexible and personalisable virtual persona, the robot is easily provided with an identity that distinguishes it from other robots with the same physical body. The MiRA infrastructure allows to project the identity of an agent the user is familiar with, using all the expressive power of a virtual character both in appearance and behaviour.

However, besides the lack of studies indicated in the last section, there also obviously exist more practical and down-to-earth obstacles. Most notably, although HMDs and wearable computers are becoming less invasive [36], at the moment these applications still im-

pose bulky and expensive hardware on each user. This clearly encumbers their deployment in scenarios with high user-to-robot ratio, and also complicates their comparison with traditional, non-augmented solutions.

Nonetheless, mixed reality agents constitute a new, exciting field of research, as they exhibit different properties than purely virtual or purely physical agents. As such, this area has numerous potential application scenarios and fascinating implications. Most notably, how do people's perceptions of a mixed reality agent differ from those of a robot or virtual character? How can the virtual agent complement the physical one, and in which areas is the application of mixed reality agents most useful?

The ubiquity and the adaptivity of these systems also open the possibility to explore new interesting scenarios. In particular, thanks to the virtual nature of their social interface, these systems can enable the creation of agents with dynamic embodiment, that is, agent chameleons that are free to mutate their form over time according to the capabilities and the expressions that they wish to utilise.

In addition to deciding their own external appearance (e.g. depending on the identity of the robot and the human user, their respective roles and their past interactions), these agents will also have the potentially unsettling property of being able to project different identities to different users.

One particular advantage this property brings with it, is the potential to project gender-, racial-, and cultural-specific personas to different people at the same time,

thereby opening the possibility to take advantage of the influence these factors have on human users while engaging a diverse group of people

It is the conjecture of this paper that such systems are the best candidate to deliver the intelligence of future ubiquitous environments, empowering them in ways traditional robotics cannot by embracing the robot's dynamic nature as mechanical and digital machine rather than being constrained to human-centred and physical capabilities.

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

References

1. Ambrose RO, Aldridge H, Askew RS, Burrige RR, Bluethmann W, Diftler M, Lovchik C, Magruder D, Rehnmark F (2000) Robonaut: NASA's space humanoid. *IEEE Int Syst* 15(4):57–63
2. Anabuki M, Kakuta H, Yamamoto H, Tamura H (2000) Welbo: an embodied conversational agent living in mixed reality space. In: *CHI '00: Ext. Abstr. on Hum. Factors in Comput. Syst.*, ACM, New York, NY, USA, pp 10–11
3. Barakonyi I, Psik T, Schmalstieg D (2004) Agents that talk and hit back: animated agents in augmented reality. In: *Proc. of the 3rd IEEE and ACM Int. Symp. on Mix. and Augm. Real. (ISMAR '04)*, pp 141–150
4. Bartneck C (2001) How convincing is Mr. Data's smile: Affective expressions of machines. *User Model and User-Adapt Interact* 11(4):279–295
5. Bartneck C (2002) eMuu - an embodied emotional character for the ambient intelligent home. PhD thesis, Eindhoven University of Technology, Eindhoven, Netherlands, unpublished
6. Bartneck C (2003) Interacting with an embodied emotional character. In: *DPPI '03: Proc. of the 2003 Int. Conf. on Design. Pleasurable Prod. and Interfaces*, ACM, New York, NY, USA, pp 55–60
7. Baylor AL, Rosenberg-Kima RB, Plant EA (2006) Interface agents as social models: the impact of appearance on females' attitude toward engineering. In: *CHI '06: Ext. Abstr. on Hum. Factors in Comput. Syst.*, ACM, New York, NY, USA, pp 526–531
8. Breazeal C (2002) *Designing Sociable Robots*. MIT Press, Cambridge, MA, USA
9. Breazeal C (2003) Toward sociable robots. *Robot and Auton Syst* 42(3):167–175
10. van Breemen A, Yan X, Meerbeek B (2005) iCat: an animated user-interface robot with personality. In: *Proc. of the 4th Int. Joint Conf. on Auton. Agents and Multiagent Syst. (AAMAS '05)*, ACM, New York, NY, USA, pp 143–144
11. Chittaro L, Ieronutti L, Rigutti S (2005) Supporting presentation techniques based on virtual humans in educational virtual worlds. In: *CW '05: Proc. of the 2005 Int. Conf. on Cyberworlds*, IEEE Comput. Soc., Washington, DC, USA, pp 245–252
12. Collett T, MacDonald B (2006) Augmented reality visualisation for player. In: *Proc. of the 2006 IEEE Int. Conf. on Robot. and Automa. (ICRA '06)*, pp 3954–3959
13. Collier RW (2001) Agent factory: A framework for the engineering of agent-oriented applications. PhD thesis, University College Dublin, Dublin, Ireland
14. Dautenhahn K (1997) I could be you: The phenomenological dimension of social understanding. *Cybern and Syst* 28:417–453(37)
15. Dautenhahn K (1998) The art of designing socially intelligent agents: Science, fiction, and the human in the loop. *Appl Artif Int* 12:573–617(45)
16. Dautenhahn K (2004) Robots we like to live with?! - a developmental perspective on a personalized, life-long robot companion. *Proc of the 13th IEEE Int Workshop on Robot and Hum Interact Commun (RO-MAN '04)* pp 17–22
17. Dautenhahn K (2007) Socially intelligent robots: Dimensions of human-robot interaction. *Philos Trans of the R Soc B: Biol Sci* 362(1480):679–804
18. Doswell JT (2005) It's virtually pedagogical: pedagogical agents in mixed reality learning environments. In: *ACM SIGGRAPH 2005 Educators Program*, ACM, New York, NY, USA, p 25
19. Dragone M, Holz T, O'Hare GMP (2006) Mixing robotic realities. In: *Proc. of the 11th Int. Conf. on Intell. User Interfaces (IUI '06)*, ACM Press, New York, NY, USA, pp 261–263
20. Dragone M, Holz T, O'Hare GMP (2007) Using mixed reality agents as social interfaces for robots. In: *Proc. of the 16th IEEE Int. Workshop on Robot and Hum. Interact. Commun. (RO-MAN '07)*, IEEE Press, Jeju Island, Korea
21. Duffy BR, O'Hare GMP, O'Donoghue RP, Rooney CF, Collier RW (1999) Reality and virtual reality in mobile robotics. In: *Proc. of the 1st Intern. Workshop on Managing Interact. in Smart Env. (MANSE '99)*, Dec. 13-14, Dublin, Ireland
22. Duffy BR, O'Hare GMP, Martin A, Bradley JF, Schön B (2003) Agent chameleons: agent minds and bodies. In: *Proc. of the 16th Int. Conf. on Comput. Animat. and Social Agents (CASA '03)*, pp 118–125
23. Fong T, Nourbakhsh I, Dautenhahn K (2003) A survey of socially interactive robots. *Robot and Au-*

- ton Syst 42(3-4):143-166
24. Fujita Y (2002) Personal robot PaPeRo. *J of Robot and Mechatron* 14(1):60-63
 25. Gockley R, Bruce A, Forlizzi J, Michalowski M, Mundell A, Rosenthal S, Sellner B, Simmons R, Snipes K, Schultz AC, Wang J (2005) Designing robots for long-term social interaction. In: *Proc. of the 2005 IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS '05)*, pp 1338-1343
 26. Ha YG, Sohn JC, Cho YJ, Yoon H (2005) Towards a ubiquitous robotic companion: Design and implementation of ubiquitous robotic service framework. *ETRI J* 27(6):666-676
 27. Hanson D, Olney A, Prilliman S, Mathews E, Zielke M, Hammons D, Fernandez R, Stephanou H (2005) Upending the uncanny valley. In: *Proc. of the National Conf. on Artif. Int. (AAAI '05)*, Pittsburgh, PA, USA
 28. Holz T, Dragone M, O'Hare GMP, Martin A, Duffy BR (2006) Mixed reality agents as museum guides. In: *ABSHL '06: Agent-Based Syst. for Hum. Learn., AAMAS '06 Workshop*, ACM Press, New York, NY, USA
 29. Ishiguro H (2007) Android science: Toward a new cross-disciplinary framework. In: *Springer Tracts in Adv. Robot., Springer Berlin / Heidelberg*, pp 118-127
 30. Johnson WL, Rickel J (1997) Steve: an animated pedagogical agent for procedural training in virtual environments. *SIGART Bull* 8(1-4):16-21
 31. Kidd CD (2003) Sociable robots: The role of presence and task in human-robot interaction. Master's thesis, MIT Media Lab, Cambridge, MA, USA
 32. Kidd CD, Breazeal C (2005) Sociable robot systems for real-world problems. In: *Proc. of the IEEE Int. Workshop on Robot and Hum. Interact. Commun. (ROMAN '05)*, pp 353-358
 33. Kim BK, Ohara K, Ohba K, Tanikawa T, Hirai S (2005) Design of ubiquitous functions for networked robots in the informative spaces. In: *Proc. of the 2nd Int. Conf. on Ubiquitous Robots and Ambient Int.*, pp 71-76
 34. Kim JH (2005) Ubiquitous robot. In: *Comput. Intell., Theory and Appl.*, vol 33, Springer Berlin / Heidelberg
 35. Klein J, Moon Y, Picard RW (1999) This computer responds to user frustration. In: *CHI '99: Ext. Abstr. on Hum. Factors in Comput. Syst.*, ACM, New York, NY, USA, pp 242-243
 36. Kleweno CP, Seibel EJ, Viirre ES, Kelly JP, Furness III TA (2001) The virtual retinal display as a low-vision computer interface: A pilot study. *J of Rehab Res and Dev* 38(4):431-442
 37. Kopp S, Gesellensetter L, Krämer NC, Wachsmuth I (2005) A conversational agent as museum guide: design and evaluation of a real-world application. In: *Proc. of the 5th Int. Working Conf. on Intell. Virtual Agents (IVA '05)*, Springer-Verlag, London, UK, pp 329-343
 38. Lee KM, Jung Y, Kim J, Kim SR (2006) Are physically embodied social agents better than disembodied social agents?: The effects of physical embodiment, tactile interaction, and people's loneliness in human-robot interaction. *Int J Hum-Comput Stud* 64(10):962-973
 39. MacDorman KF (2005) Androids as an experimental apparatus: Why is there an uncanny valley and can we exploit it? In: *CogSci '05 Workshop: Toward Social Mechanisms of Android Science*, pp 106-118
 40. Maeyama S, Yuta S, Harada A (1998) Remote museum system using a networked robot. *Adv Robot* 13:315-316(2)
 41. Martin A, O'Hare GMP, Duffy BR, Schön B, Bradley JF (2005) Maintaining the identity of dynamically embodied agents. In: *Proc. of the 5th Int. Working Conf. on Intell. Virtual Agents (IVA '05)*, Springer-Verlag, London, UK, pp 454-465
 42. Milgram P, Rastogi A, Grodski JJ (1995) Telerobotic control using augmented reality. In: *Proc. of the 4th IEEE Int. Workshop on Robot and Hum. Commun. (RO-MAN 1995)*, pp 21-29
 43. Milgram P, Takemura H, Utsumi A, Kishino F (1995) Augmented reality: A class of displays on the reality-virtuality continuum. In: *Proc. of the SPIE Conf. on Telemanipulator and Telepresence Technol.*, Boston, MA, USA, *Proc. of SPIE*, vol 2351, pp 282-292
 44. Mori M (1982) *The Buddha in the Robot*. C.E. Tuttle
 45. Muldoon C, O'Hare GMP, O'Grady MJ (2006) Collaborative agent tuning: Performance enhancement on mobile devices. In: *Proc. of the 6th Intern. Workshop on Eng. Soc. in the Agents' World (ESAW '05)*, October 26-28, 2005, Kusadasi, Turkey, Revised Selected and Invited Papers, Springer Berlin / Heidelberg, *Lecture Notes in Computer Science (LNCS)*, vol 3963, pp 241-258
 46. Nourbakhsh IR, Kunz C, Willeke T (2003) The mobot museum robot installations: a five year experiment. In: *Proc. of the 2003 IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS '03)*, vol 4, pp 3636-3641 (3)
 47. Ogden B, Dautenhahn K (2000) Robotic etiquette: Structured interaction in humans and robots. In: *Proc. of the Symp. on Intell. Robot. Syst. (SIRS '00)*, pp 353-361

48. O'Hare GMP (1996) The agent factory: An environment for the fabrication of distributed artificial systems. In: O'Hare GMP, Jennings NR (eds) *Foundations of Distributed Artificial Intelligence*, Sixth Generation Computer Series, Wiley Interscience Publishers, New York, NY, USA, pp 449–484
49. O'Hare GMP, Duffy BR (2002) Agent chameleons: Migration and mutation within and between real and virtual spaces. In: *Proc. of the AISB '02 Symp. on Animat. Expressive Characters for Soc. Interact.*, Apr. 3-5, 2002, Imperial College of Science & Technology, London, UK, pp 63–68
50. O'Hare GMP, Duffy BR, Bradley JF, Martin AN (2003) Moving minds from robots to digital information spaces. In: *Proc. of the 2nd Intern. Symp. on Auton. MiniRobots for Res. and Edutainment*, Feb. 18-21 2003, Brisbane, Australia
51. O'Hare GMP, Duffy BR, Schön B, Martin A, Bradley JF (2003) Agent chameleons: Virtual agents real intelligence. In: *Proc. of the 4th Intern. Working Conf. on Intell. Virtual Agents (IVA '03)*, September 15-17, 2003, Irsee, Germany, *Lecture Notes in Computer Science (LNCS)*, vol 2792, pp 218–225
52. O'Hare GMP, Campbell AG, Stafford JW, Aiken R (2005) NeXuS: Behavioural realism in mixed reality scenarios through virtual sensing. In: *Proc. of the 18th Int. Conf. on Comput. Animat. and Social Agents (CASA '05)*
53. Pacchierotti E, Christensen H, Jensfelt P (2006) Evaluation of passing distance for social robots. *Proc of the 15th IEEE Int Symp on Robot and Hum Interact Commun (RO-MAN '06)* pp 315–320
54. Prendinger H, Dohi H, Wang H, Mayer S, Ishizuka M (2004) Empathic embodied interfaces: Addressing users' affective state. In: *Proc. of the Tutor. and Res. Workshop on Affect. Dialogue Syst.*, Springer Verlag, pp 53–64
55. Rao AS, Georgeff MP (1991) Modeling rational agents within a BDI-architecture. In: Allen J, Fikes R, Sandewall E (eds) *Proc. of the 2nd Int. Conf. on Princ. of Knowl. Represent. and Reason. (KR '91)*, Morgan Kaufmann, San Mateo, CA, USA
56. Reeves B, Nass C (1998) *The media equation: how people treat computers, television, and new media like real people and places*. Cambridge University Press, New York, NY, USA
57. Rickenberg R, Reeves B (2000) The effects of animated characters on anxiety, task performance, and evaluations of user interfaces. In: *CHI '00: Proc. of the SIGCHI Conf. on Hum. Factors in Comput. Syst.*, ACM, New York, NY, USA, pp 49–56
58. Shoji M, Miura K, Konno A (2006) U-Tsu-Shi-O-Mi: the virtual humanoid you can reach. In: *ACM SIGGRAPH 2006 Emerg. Technol.*, ACM, New York, NY, USA, p 34
59. Simmons R, Fernandez J, Goodwin R, Koenig S, O'Sullivan J (2000) Lessons learned from Xavier. *IEEE Robot & Autom Mag* 7(2):33–39
60. Simmons R, Goldberg D, Goode A, Montemerlo M, Roy N, Sellner B, Urmson C, Schultz A, Abramson M, Adams W, Atrash A, Bugajska M, Coblenz M, MacMahon M, Perzanowski D, Horswill I, Zubek R, Kortenkamp D, Wolfe B, Milam T, Maxwell B (2003) GRACE: an autonomous robot for the AAAI robot challenge. *AI Mag* 24(2):51–72
61. Slater M, Pertaub DP, Steed A (1999) Public speaking in virtual reality: facing an audience of avatars. *IEEE Comput Graph and Appl* 19(2):6–9
62. Thrun S, Bennewitz M, Burgard W, Cremers AB, Dellaert F, Fox D, Hähnel D, Rosenberg C, Roy N, Schulte J, Schulz D (1999) MINERVA: a second-generation museum tour-guide robot. In: *Proc. of the 1999 IEEE Int. Conf. on Robot. and Autom.*, vol 3, pp 1999–2005 (3)
63. Wagner D, Billingham M, Schmalstieg D (2006) How real should virtual characters be? In: *ACE '06: Proc. of the 2006 ACM SIGCHI Int. Conf. on Adv. in Comput. Entertain. Technol.*, ACM, New York, NY, USA, p 57
64. Yamato J, Shinozawa K, Naya F, Kogure K (2001) Evaluation of communication with robot and agent: Are robots better social actors than agents? In: *Proc. of the 8th IFIP TC.13 Int. Conf. on Hum.-Comput. Interact. (INTERACT '01)*, Tokyo, Japan, pp 690–691
65. Yamato J, Shinozawa K, Brooks R, Naya F (2003) Human-robot dynamic social interaction. *NTT Tech Rev* 1(6):37–43
66. Young JE, Xin M, Sharlin E (2007) Robot expressionism through cartooning. In: *Proc. of the ACM/IEEE Int. Conf. on Hum.-Robot Interact. (HRI '07)*, ACM, New York, NY, USA, pp 309–316
67. Zambaka CA, Ulinski AC, Goolkasian P, Hodges LF (2007) Social responses to virtual humans: implications for future interface design. In: *CHI '07: Proc. of the SIGCHI Conf. on Hum. Factors in Comput. Syst.*, ACM, New York, NY, USA, pp 1561–1570
68. Ziemke T (2003) What's that thing called embodiment? In: *Proc. of the 25th Ann. Meet. of the Cogn. Sci. Soc.*