Research Repository UCD

Title	MiRA - mixed reality agents
Authors(s)	Holz, Thomas, Campbell, Abraham G., O'Hare, G. M. P. (Greg M. P.), Stafford, John W., Dragone, Mauro, Martin, A. (Alan)
Publication date	2011-04
Publication information	Holz, Thomas, Abraham G. Campbell, G. M. P. (Greg M. P.) O'Hare, John W. Stafford, Mauro Dragone, and A. (Alan) Martin. "MiRA - Mixed Reality Agents" 69, no. 4 (April, 2011).
Publisher	Elsevier
Item record/more information	http://hdl.handle.net/10197/3195
Publisher's statement	This is the author's version of a work that was accepted for publication in International Journal of Human-Computer Studies. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in International Journal of Human-Computer Studies (IJHCS), 69 (4): 251-268 DOI: 10.1016/j.ijhcs.2010.10.001
Publisher's version (DOI)	10.1016/j.ijhcs.2010.10.001

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

MiRA - Mixed Reality Agents

Thomas Holz^a, Abraham G. Campbell^a, Gregory M.P. O'Hare^a, John W. Stafford^b,
Alan Martin^b, Mauro Dragone^a

^a CLARITY: Centre for Sensor Web Technologies, School of Computer Science & Informatics, University
College Dublin, Belfield, Dublin 4, Ireland

^b School of Computer Science and Informatics, University College Dublin, Belfield, Dublin 4, Ireland

---

Abstract

In recent years, an increasing number of Mixed Reality (MR) applications have been developed using agent technology — both for the underlying software and as an interface metaphor. However, no unifying field or theory currently exists that can act as a common frame of reference for these varied works. As a result, much duplication of research is evidenced in the literature. This paper seeks to fill this important gap by outlining for the first time a formal field of research that has hitherto gone unacknowledged, namely the field of *Mixed Reality Agents* (MiRAs), which are defined as agents embodied in a Mixed Reality environment.

Based on this definition, a taxonomy is offered that classifies MiRAs along three axes: *agency*, based on the weak and strong notions outlined by Wooldridge and Jennings (1995); *corporeal presence*, which describes the degree of virtual or physical representation (body) of a MiRA; and *interactive capacity*, which characterises its ability to sense and act on the virtual and physical environment.

Furthermore, this paper offers the first comprehensive survey of the state of the art of MiRA research and places each project within the proposed taxonomy. Finally, common trends and future directions for MiRA research are discussed.

By defining Mixed Reality Agents as a formal field, establishing a common taxonomy, and retrospectively placing existing MiRA projects within it, future researchers can effectively position their research within this landscape, thereby avoiding duplication and fostering reuse and interoperability.

Key words: mixed reality agents, mixed reality, interaction metaphors

---

1. Introduction

As Mixed Reality (MR) research and development has progressed and intensified over the past decade, an increasing amount of focus has shifted towards the nature of the environments that MR applications occupy. An MR environment is an environment within which the domains of the virtual and the physical are fused in a spatially coherent manner. The development of MR environments of growing sophistication has brought about a desire to imbue said environments with greater intelligence. Consequently, a number of projects have realised the potential of agent technology to provide

such intelligence, leading to agents that are defined by the MR environments they exist in.

The concept of agents and humans coexisting in a shared MR environment is one that has been present in popular culture for some time. Examples include the BBC documentary *Hyperland* (Adams, 1990), wherein author Douglas Adams discusses his vision of a ‘software agent’, and the Japanese animated series *Dennō Coil* (Iso, 2007), a series that follows a group of children exploring the half physical, half virtual fictional city of Daikoku. While present technology has not quite reached the levels envisioned by science-fiction writers, MR environments are growing increasingly rich and complex. This technological maturation has brought about a scenario whereby previously fantastical notions of shared physical-virtual experiences have moved closer to becoming a tangible reality.

Notwithstanding such advancements, agent-based MR applications have yet to reach the requisite levels of complexity and stability whereby truly intuitive interaction with a user in a shared environment is possible. The motivation behind this paper is to help such technology realise its potential by defining a formal field of research under which the concepts behind Mixed Reality Agent projects can be unified. This formalisation will provide the necessary scaffolding upon which research in this space can be positioned and categorised, thereby preventing the repetition of work, helping frame existing work, promoting interoperability and encouraging the development of standards.

In summary, the contributions and motivations of this paper are:

- To introduce the novel concept of a *Mixed Reality Agent (MiRA)*, in order to define a distinct field of research and provide a common frame of reference for existing and future researchers;
- To introduce the notions of *corporeal presence* and *interactive capacity* to distinguish between a MiRA’s visible form and its ability to sense and act on the environment;
- To develop a taxonomy of Mixed Reality Agents based on these notions, in order to allow discussion and comparison of different research projects;
- To provide a comprehensive overview of the state of the art of MiRA research, locating each existing project within the proposed taxonomy;
- To identify current and future trends and practices of MiRA research.

2. Mixed Reality Agents

The space that MiRA research occupies in relation to Mixed Reality and agent research can be visually represented in a Venn diagram, as depicted in Figure 1. To achieve the goal of establishing a common reference point for agent-based MR developers, one must start by first establishing a broad definition of a Mixed Reality Agent.

2.1. Definition

A *Mixed Reality Agent* (MiRA) is an agent embodied in a Mixed Reality environment.

While this might sound like a trivial definition, in and of itself, each of the key concepts, *Agent*, *Embodiment* and *Mixed Reality Environment*, has historically been used for, and understood to mean, different things. Therefore, in order to clarify the above definition, this section further examines the notions of agency, embodiment and Mixed Reality environments and defines them as they are used within the context of this research.

Figure 1: A Venn Diagram representing MiRA research

2.2. Agents

The term agent can be somewhat ambiguous and may convey different meanings depending upon the research in question. The principle reason for this ambiguity is that the word ‘agent’ is used extensively outside of the software community, while within the community it is “an umbrella term for a heterogeneous body of research and development” (Nwana, 1996). It is therefore necessary to define the term as it is used within the context of this paper. Based upon the notions of agency introduced by Wooldridge and Jennings (1994), an *agent* is defined as a hardware- or software-based entity characterised by the following attributes:

Autonomy: Agents can operate without the direct intervention of humans or others, and have control over their actions and internal state.

Social Ability: Agents can interact with other agents as well as with their users.

Reactivity: Agents can perceive their environment and respond in a timely fashion to changes within it.

Pro-Activity: In addition to responding to their environments, agents can take the initiative and exhibit goal directed behaviour.

Implied in the above definition is *situatedness*; the agent operates within an environment, sensing events in that environment and acting upon it.

Wooldridge and Jennings further distinguish between a *weak agent*, that possesses the above base criteria, and a *strong agent* (first outlined by Shoham (1993)), that is assigned mentalistic attitudes such as knowledge, belief, intention and obligation. The latter agent's mental state is formed from the assignment of these and other such attitudes. These mental attitudes should form components of the agent's reasoning, driven by a formal theory with clear semantics, and should correspond to the common-sense use of the terms. A well-known model in this realm is the Belief-Desire-Intention (BDI) model (Rao and Georgeff, 1995), which defines an agent in terms of its beliefs (information about the world), desires (long-term goals and motivations) and intentions (short-term actions).

Multi-agent systems are today regarded as a general purpose paradigm for the engineering of complex computational systems. Agent-based computing promotes the design and development of applications using multiple autonomous agents. Such an approach is ideally tailored to tackling complex, scalable problems due to the characteristic suitability of agents in dynamic and unpredictable scenarios. Agent-based techniques are, therefore, natural candidates for the implementation of intelligent, distributed, and adaptive systems.

In the area of Human-Computer Interaction (HCI), autonomous agents are usually referred to as interface agents or user agents (Laurel, 1990). Here, by virtue of their autonomy, software agents can be instructed using high-level descriptive terms in place of traditional methods. Ideally, users may simply describe the goal they want the agent to achieve rather than directly supervise their operation. This enables the user to carry out more complex operations by focusing upon high-level supervision while the agent performs all the necessary lower-level computation. In this sense, these agents act as assistants on behalf of the user rather than being mere tools.

2.3. Embodiment

Embodiment is an important concept, yet one for which a varying array of definitions exist. It is therefore necessary (as with the term 'agent') to establish a definition of embodiment as it is used within this research. Based upon the analysis carried out by Ziemke (2001), it is possible to identify a number of distinct definitions of embodiment.

These include: *structural coupling*, in which a system is embodied if it is able to sense, affect and be affected by its environment, as part of an agent-environment dynamic (Quick et al., 1999); *historical embodiment*, in which a system is only seen as embodied if it has a history of structural coupling (Riegler, 2002); *physical embodiment*, in which an agent is embodied if it has a body through which it interacts with the environment (Pfeifer, 2002) and *biological embodiment*, in which only a living system can be seen as embodied (Sharkey and Ziemke, 2000).

Another important consideration regarding embodiment is the social environment and the idea of social embodiment. As outlined by Dourish (2001), social embodiment can be seen as the agent's relationship with its social environment, in a similar way to how physical embodiment can be seen as the agent's relationship with its physical environment. The importance of social embodiment has also been argued by Dauten-

hahn (1998, 1999) and Duffy (2000), who claim that the social environment includes the agent's interactions with other agents or human users.

Within the context of this research, *embodiment* is seen as the strong provision of environmental context (structural coupling) with a social element included. An agent is embodied if it is situated in a particular environment, has a *body*, and senses and interacts with that environment, and any other individuals located therein.

This definition of embodiment coincides with that of Dourish (2001), who emphasised the importance of an embodied approach to Human-Computer Interaction, in light of new developments in Ubiquitous and Social Computing and proposed a number of design guidelines for the development of Embodied Interaction.

Early Artificial Intelligence (AI) research focused upon reasoning based upon search of abstract symbol structures (Newell and Simon, 1976). But this unembodied approach, sometimes referred to as 'Good Old Fashioned AI' (Haugeland, 1985), had a number of flaws. As noted by Steels (2000) and Dautenhahn (1999), humans have a tendency to 'animate' the world and are unlikely to see an unembodied agent as intelligent. Therefore embodiment has, in recent years, come to be seen as an important requirement in the development of an intelligent system (Duffy et al., 2005).

The move away from the unembodied approach was triggered by a series of papers by Brooks (1991a,b), who emphasised the situatedness and embodiment of an agent. Brooks' popularisation of the reactive approach served as a catalyst for the creation of a more embodied approach to AI, where an agent must be structurally coupled with its environment if it is to be seen as intelligent.

While robot agents are embodied in a physical form, using sensors and actuators to perceive and act upon the physical world, virtual agents can also be considered embodied in their simulated environment, at least to the extent to which the simulation manages to create a structural coupling between the agent and the simulated environment. Both strands are motivated by the desire to create agents that are capable of behaving and interacting in an intelligent manner with other agents. Crucially, both robotic and virtual agents can be considered synthetic characters, although differently embodied, i.e. in physical or digital form.

2.4. Mixed Reality Environments

Figure 2: Milgram's Reality-Virtuality Continuum (Milgram and Kishino, 1994)

Key to the definition of Mixed Reality Agents, as outlined within this paper, is the idea of a Mixed Reality environment. Milgram and Kishino (1994) define MR in terms of their Reality-Virtuality Continuum (Figure 2) whereby Mixed Reality is the space between a purely physical (or 'real', as they describe it) environment and a purely virtual environment. Each MR environment can be seen, to a greater or lesser

degree, as consisting of two sub-spaces: the real, or physical, sub-space and the virtual sub-space, the latter consisting of all the virtual artefacts in the MR Environment.

The motivation behind developing MR systems in general is to not only to enhance a user's perception of the real world by superimposing generated digital information, as in Augmented Reality (AR) systems, but also to leverage existing visual and spatial skills to increase the user's interaction capabilities. One of the earliest, and best, examples of such a system is the MagicBook (Billinghurst et al., 2001) project. In utilising the interface metaphor of a physical book, it combines virtual imagery with the tangible nature of a physical object, allowing multiple users to view a shared Mixed Reality object.

The MagicBook also vividly illustrates the continuous nature of Milgram and Kishino's definition of MR. On the real end of the continuum, the book can be handled and read like any normal, real-world book. Using an AR display, however, the real book is enhanced with virtual imagery, creating a Mixed Reality experience that the authors liken to an enhanced pop-up book. Finally, users can fully immerse themselves in the virtual scene that is playing out on the real book pages, thereby reaching the other end of the continuum. All these experiences are internally consistent, with both immersed and AR users being represented, the former as characters in the scene, the latter as virtual heads floating in the sky. As such, the MagicBook can be seen to span the whole of the Reality-Virtuality Continuum.

Other notable examples of the combination of digital content with physical objects through MR technology include applications that aid the design and construction of buildings, seamlessly blending virtual and real buildings to assist planning and early evaluation of urban development projects (Wang and Jeong Kim, 2009), and training and learning tools such as medical applications that assist surgeons during operations by providing images of interior body structures overlaid on patients' bodies (Bornik et al., 2003).

The most sophisticated examples are human-in-the-loop simulation systems, such as the birth simulator developed at the TU München (Sielhorst et al., 2004) that superimposes an X-ray view of the womb on a scale model of a woman's torso. It is also equipped with a set of birth pliers that provide haptic feedback for emulating the extraction of the baby.

In general, the construction of all these MR applications offers characteristic engineering challenges. For example, their effectiveness requires a geometric correspondence between the two sub-spaces constituting the MR environment. Registration between sub-spaces is usually supported by tracking the users' viewpoint using techniques such as magnetic and other head tracking devices. The task becomes considerably more challenging for simulation systems, especially when user interaction is considered. In these cases the virtual artefacts should also align, with reasonable accuracy, to the physical and biological laws that apply to their real equivalents.

3. The MiRA Cube Taxonomy

All Mixed Reality Agents are necessarily embodied within the MR environment by virtue of having a body and being able to sense and act upon the environment (i.e.

being structurally coupled to the environment). However, while for completely virtual or completely physical agents, the ability to sense and act is usually associated with the agent's body, this is not necessarily the case for MiRAs.

For example, a robot that can sense and manipulate virtual objects in its environment does not necessarily possess a virtual representation. Neither is a virtual agent considered to have a physical body just because it uses sensors in the physical world to, for example, track and identify physical objects and human users. In fact, MiRAs that are equally present in both the virtual and the physical domain, in the sense of having a physical and a virtual representation, are currently in the minority.

To truly define the MiRA field, it is therefore necessary to distinguish between an agent's ability to sense and act on the virtual or physical environment, hereafter referred to as its *interactive capacity*, and its degree of virtual or physical representation, hereafter referred to as the *corporeal presence* of a MiRA. Notably, corporeal presence is distinctly different from interactive capacity, as it is an attribute of any artefact in the environment and is not merely associated to the concept of agent-hood. For instance, a real teapot has physical corporeal presence while a virtual teapot has virtual corporeal presence, without any of them being an agent.

However, corporeal presence is more than simply visual representation. As has been noted previously, geometric correspondence between the virtual and physical subspace is necessary for the effectiveness of MR environments and, therefore, a basic property of corporeal presence. Related characteristics, such as occlusion, collision detection, or adherence to other physical laws, increase corporeal presence. For example, if a virtual teapot is moved off a physical table, falls and shatters on the physical floor into virtual pieces, it has a higher physical corporeal presence than if it was just a virtual teapot floating in the air. By the same token, a physical robot that avoids virtual objects has a higher virtual corporeal presence than one that moves through them.

For physical agents, physical corporeal presence is therefore inevitably high — a robot cannot escape the laws of physics. For example, if there is an object between the robot and an observer, the robot will get occluded. Virtual corporeal presence, on the other hand, is often engineered at an environment level since the same features are usually inherent in all virtual objects.

Taking all this into account, one can differentiate between MiRA applications based on their varying degrees of corporeal presence and interactive capacity. By adding the level of agency as a third axis, a vector space is created, the *MiRA Cube* (see Figure 3), which may act as a scaffolding within which projects can now be situated.

Quantifying corporeal presence and interactive capacity in an absolute frame of reference would, at best, be arbitrary and might be mistaken for a form of ranking where no such thing is implied. Instead, in keeping with the spirit of Milgram's continuum, these axes, spanning from the physical extreme to the virtual, are considered relative. As such, a MiRA is located within the taxonomy by contrasting its individual degrees of virtual and physical interactive capacity and corporeal presence and by identifying its level of agency. Because of the relative nature of the axes, a number of categories is proposed in order to enable comparison between different MiRAs.

For corporeal presence, MiRAs can be distinguished by their degree of representation in the virtual and physical domain, leading to three distinct categories:

Figure 3: A diagram of the MiRA Cube Taxonomy spanned by the three axes agency, corporeal presence, and interactive capacity. Employing Wooldridge and Jennings's notions of weak and strong agency (Wooldridge and Jennings, 1994) produces two separate 'planes' of MiRAs.

- Agents that have a stronger corporeal presence in the virtual domain than in the physical domain of their MR environment.
- Agents that have a stronger corporeal presence in the physical domain than in the virtual domain of their MR environment.
- Agents that have an equally strong corporeal presence in both the virtual and physical domains of their MR environment.

Figure 4 illustrates where the three different categories of corporeal presence exist in relation to a modified version of Milgram's Reality-Virtuality Continuum (Milgram and Kishino, 1994). The most notable modification is the change in terminology at the extremes of the continuum. These changes are made to clarify that a MiRA's reality is an environment that mixes the domains of the virtual and the physical and that both are equally 'real' from the MiRA's perspective.

Figure 4: An adaptation of Milgram's Reality-Virtuality Continuum indicating where different categories of MiRAs exist in Mixed Reality space

A similar distinction can be made for interactive capacity, where MiRAs can be distinguished by their sensing and acting capabilities in the virtual and physical domain. Most agents that have a stronger corporeal presence in one of the domains are able to

both sense and act in that particular domain while, most of the time, being only able to sense the other domain. For example, a lot of MiRA research projects employ a virtual character that is only able to sense the physical domain but cannot act on it. MiRAs with equal corporeal presence, on the other hand, can usually interact faithfully with both physical and virtual domains on equal terms. As will be discussed later, however, there are a number of projects who differ from this schema.

For agency, finally, Wooldridge and Jennings's notion of weak and strong agency (Wooldridge and Jennings, 1994) provides a convenient and well-acknowledged differentiation that leads to two distinct 'planes' of MiRAs (see Figure 3). As such, agents that possess the base capabilities of autonomy, social ability, reactivity, and pro-activity are located in the weak agency plane, while agents that also incorporate mentalistic notions, such as beliefs, emotions and desires, are located in the strong agency plane. However, even with such a coarse distinction, accurate categorisation is quite a challenge, as many of the reviewed projects do not describe their agent architecture in sufficient detail.

A taxonomy based on the above three axes facilitates the ordering of agents in terms of their degree of embodiment in the shared MR environment. For instance, a virtual character that avoids real obstacles in its path (such as chairs and other real artefacts) is still less embodied than a physical robot, which has the same capabilities but may also affect these objects thanks to the capabilities associated with its physical corporeal presence.

Table 1: Overview of the reviewed MiRA systems and their categorisation according to the introduced taxonomy.

Project	Agency	CP	IC	Domain	Agent System	Virtual CP	Physical CP	Virtual IC	Physical IC
SHEEP (Sheep & Wolves)	Weak	VCP	VIC	Ent	BE	Virtual body	N/A	Interaction with a complete virtual world	Object recognition
Virtual Autonomy Assistant	Weak	VCP	VIC	Ent	BE	Virtual body	N/A	Sense virtual objects	Location awareness
Robbie	Weak	VCP	VIC	Edu	FSM	Virtual body	Obstacle avoidance, occlusion	Interaction with a complete virtual world	User tracking, spatial awareness
MARA	Weak	VCP	VIC	Ent	FSM	Virtual body	N/A	Sense virtual models of physical buildings	User tracking, speech synthesis, user tracking, gaze tracking, object recognition, location awareness
PECA	N/A	VCP	VIC	Edu	Unknown	Virtual body	N/A	Sense/Act on virtual objects	Gesture recognition, NL interaction, object recognition, user tracking
Welbo	N/A	VCP	VIC	Ent	Unknown	Virtual body	N/A	Sense/Act on virtual objects	Voice recognition, speech synthesis, object recognition, user tracking
Virtual Raft / EcoRaft	Weak	VCP	VIC	Edu	BE	Multiple avatars	Reacts to inertia, gravity	Interaction with a complete virtual world	User tracking, gesture recognition, audio feedback
ALIVE	Strong	VCP	VIC	Ent	Motivational model	Virtual body	Interaction of user with virtual objects	Interaction with a complete virtual world	User tracking, gesture recognition, NL interaction
Virtual Gunslinger	Strong	VCP	VIC	Ent	STE	Multiple avatars	N/A	Interaction with a complete virtual world	User tracking, gesture recognition, NL interaction
GHEST	Strong	VCP	VIC	Edu	STE	Multiple avatars	N/A	Sense/Act on virtual objects	NL interaction, user tracking, gaze tracking, object recognition, location awareness
Invisible Person	Strong	VCP	VIC	Ent	Emotional model	Virtual body	Occlusion	Sense & place virtual markers	User tracking, gesture recognition, body posture
MRIS	Strong	VCP	VIC	Ent	STE	Multiple avatars	Obstacle avoidance & occlusion for inserted user image, interaction of user with virtual objects	Interaction with a complete virtual world	User tracking, gesture recognition, NL interaction
UnAgent Nexus	Strong	VCP	VIC	Ent	BDI	Virtual body	Obstacle avoidance, occlusion	Sense/Act on virtual objects	Marker recognition, user tracking
Virtual Brownies	Strong	VCP	VIC	Oh	BDI MAS	Multiple avatars	N/A	Sense/Act on virtual objects	Voice recognition, speech synthesis, marker recognition, user tracking
Virtual Room Inhabitant	Weak	VCP	EIC	Ent	?	Multiple avatars	N/A	Interaction other avatars	Tracking & manipulation of physical objects
Max	Strong	VCP	EIC	Ent	FSM	Virtual body	N/A	N/A	User tracking, object tracking
MACK (Sheep & Wolves)	Strong	VCP	EIC	Ent	Hybrid BDI	Virtual body	N/A	N/A	Multiple user tracking, object tracking, physical world map
Giesler et al.	Weak	PCP	PIC	Ent	ECA	Virtual body	Mobile quadruped robot	Virtual world map	User tracking, physical world map
Stilman et al.	Weak	PCP	PIC	Ent	BE	N/A	Mobile robot	Sense virtual markers	Mobile robots, NL interaction (text)
Virtual Phantomes	N/A	PCP	PIC	Dev	Unknown	N/A	Mobile robot	Sense virtual markers	Mobile robot, speech synthesis, voice recognition
Virtual Synergy	Weak	ECP	ECP	S&R	BE	N/A	Mobile biped robot	Sense & place virtual objects	Mobile robot, object tracking
Virtual Humanoid	Weak	ECP	EIC	S&R	Unknown	N/A	Mobile robot	Virtual world map, sense & place virtual markers	Mobile robots with basic sensors and IR transmitters
Joyces	N/A	ECP	PIC	Ent	Unknown	Virtual model of the physical body	Humanoid robot	Virtual world map, sense & place virtual markers	Mobile robots of various kinds
MRA	Strong	ECP	EIC	Ent	Hybrid BDI MAS	Virtual face	Humanoid robot	Facial expressions	Humanoid robot
				Oh	Unknown	Virtual clothing	Mobile robot	Sense/Act on virtual objects	Mobile robot, user tracking, gaze tracking, voice recognition, speech synthesis

Legend: CP (Cognitive Presence), IC (Interactive Capacity), ECP (Equal Cognitive Presence), VCP (Greater Virtual Cognitive Presence), PCP (Greater Physical Cognitive Presence), EIC (Equal Interactive Capacity), VIC (Greater Virtual Interactive Capacity), PIC (Greater Physical Interactive Capacity), BD (Belief-Desire-Intention Architecture), MAS (Multi-Agent System), STE (Story-Telling Engine), ECA (Embodied Conversational Agent), BE (Behaviour Engine), FSM (Finite State Machine), Ent (Entertainment), Edu (Education), Int (Interface), S&R (Search & Rescue), Dev (Development Tool), Oh (Other), NL (Natural Language).

4. Populating the MiRA Cube

Using the definition of a MiRA outlined above, the resulting taxonomy can now be exercised to classify a compelling number of research projects as Mixed Reality Agents. The authors have identified over twenty-five existing projects that fall under the proposed taxonomy and that can now be located within the resulting space. Table 1 gives an overview of each project’s individual features and its categorisation within the taxonomy. The following section describes each project in detail and offers the first comprehensive survey of the state-of-the-art of Mixed Reality Agents.

Figure 5: A top-down view of the MiRA Cube depicting how existing projects form different clusters, especially when only taking corporeal presence and interactive capacity into account. Bold project names denote MiRAs that are considered strong agents (mostly employing a BDI system or a story-telling engine), while italic names denote weak agents (which are predominantly driven by finite state machines or behaviour engines). NB: This diagram assumes, arguably, that projects which do not publish details on their agent systems fall under the heading of weak agency.

When projecting both agency planes on top of each other and only charting embodiment as a function of corporeal presence and interactive capacity, a number of striking clusters and, conversely, empty spaces appear within the resulting square (see Figure 5). MiRAs with stronger virtual embodiment possess corporeal presence and interactive capacity that are both greater in the virtual domain. Analogous to this, MiRAs with stronger interactive capacity and corporeal presence in the physical domain are considered to be predominantly physically embodied. MiRAs with Mixed Reality

embodiment, on the other hand, possess full corporeal presence and interactive capacity in both domains. Occupying the remaining corners of the diagram are MiRAs with greater virtual corporeal presence that predominantly interact with the physical world and MiRAs with greater physical corporeal presence that predominantly interact with the virtual world. Each distinct MiRA cluster will now be considered and characterised.

4.1. MiRAs with Stronger Virtual Embodiment

Figure 6: MiRAs with stronger virtual embodiment that employ a head-mounted display. Clockwise from top left: GEIST, Herding Sheep, Ritchie, MARA, Welbo, NeXuS, and UbiAgent.

Perhaps not surprisingly, most current MiRA research concentrates on agents that are predominantly embodied in the virtual domain, i.e. agents who possess both greater virtual corporeal presence and greater virtual interactive capacity. These MiRAs are often motivated by the desire to integrate virtual characters into the physical world in order to allow them to share the same space as the human user.

This ability is especially pronounced in outdoor systems, which allow for location-based and context-aware interaction with the user wherever he goes. A notable example of this type of system is the mobile city guide and personal assistant agent MARA (Schmeil and Broll, 2007), which is built on a two-tier hierarchical state machine. In the physical domain, MARA is limited to sensing the user's position and orientation, while in the virtual, which is populated with spatially aligned models of buildings and other virtual objects, interaction is much more varied and versatile. The user interacts with MARA via predefined spoken commands and by spatially referencing objects and places via gaze.

Still being confined to a purely virtual body, however, the physical interactive capacity of agents such as MARA is greatly lowered, as they can only sense the physical world but cannot affect it. But interaction with the physical space is seldom of prime importance for agents in these types of application, as their main function usually is to act as mediator between the human users and the virtual artefacts in the MR space.

This scenario is fittingly illustrated by Welbo (Anabuki et al., 2000), a virtual agent floating in midair that helps users with interior decoration planning in an “MR Living Room”. Instead of directly manipulating the placement of virtual tables and chairs in the real room, the user gives verbal instructions to Welbo who then places the furniture for him. The authors do not describe Welbo’s level of agency — a common occurrence, unfortunately, in many of the reviewed works whose focus lies primarily with MiRAs as an interaction metaphor and less with the application of the agent paradigm to Mixed Reality environments. Most notably about Welbo, however, are the authors’ reported results from an informal study that evaluated the impact of spatial factors on Welbo’s acceptance

“Through the experiments, we understood that people preferred a size such that they can see Welbo’s whole body in their field of view. Similarly, people like it to stay some distance away from them. As people feel uncomfortable when others look down on them, Welbo gives an unfavorable impression when it floats over them.” (Anabuki et al., 2000)

These findings illustrate the importance of an agent’s embodiment with regard to effective social interaction with humans, something that developers should be even more mindful of when MiRAs are equipped with the ability to move between different environments, as their corporeal presence and interactive capacity might change with regard to the environment they are currently embodied in.

As a result, some MiRAs that do move between environments possess knowledge about their embodiment and have the ability to adjust their corporeal presence and interactive capacity to any given environment. Agents in the PECA (Pedagogical Embodied Conversational Agents) (Doswell, 2005) system, for example, can not only move between different environments but can indeed can intelligently adapt to and choose the environment best suited to a given task, with the goal of improving motivation, accelerating human learning and providing a more enhanced and natural e-learning platform,

The conversational agent Ritchie (Dorfmueller-Ulhaas and André, 2005) can move between a purely virtual and a Mixed Reality environment. Specifically, Ritchie and the user can jointly explore a virtual model of the German city of Augsburg, either from a first-person perspective in an immersive virtual environment, or from a bird’s eye view onto an AR table top model. By placing a special marker onto the table top model, a user can trigger the transition, arriving in the virtual environment where the marker was placed.

An extension of Ritchie, the Virtual Anatomy Assistant (VAA) (Wiendl et al., 2007), teaches visitors the locations of organs in the human body by allowing users to place virtual organs within a physical skeleton using a magic wand, with the agent commenting on the users’ performance. While the agent does not move between environments, the project supports occlusion and even throws realistic shadows onto physical objects, which greatly increases its physical corporeal presence.

A completely different interpretation of agents moving between environments is evidenced by the sister projects Virtual Raft (Tomlinson et al., 2005) and EcoRaft (Tomlinson et al., 2006). In the first, a user can transport virtual characters via a tablet PC (the raft) between several islands (desktop PCs). The characters are sensitive to the motion of the raft, having to balance themselves to not fall off. This gives them a basic level of physical corporeal presence and qualifies the agents, which are controlled by a behaviour engine, as MiRAs. Once deposited on an island, the virtual characters proceed to interact with the other inhabitants. The EcoRaft project applies the same system to an educational context, teaching children about restoration ecology. Users employ the same raft metaphor to move virtual plants and animals to an island in order to restore its ecosphere.

Admittedly, MiRAs that are able to migrate between environments are not necessarily aware of the environment within which they currently reside. The SHEEP project (MacWilliams et al., 2003), for example, simulates virtual sheep grazing on a pasture that is visualised on a table top similar to Ritchie. The sheep employ simple behaviours (based on a simulation algorithm for a flock of birds) to stay close to each other but also flock to a physical sheep figurine which the user can place onto the landscape so as to direct the herd. Individual sheep can be created and deleted using a ‘magic wand’ (i.e. a wand whose position and orientation are tracked by the application). Using a special marker, sheep can be picked up and transferred to a PDA and back down on the pasture with the agents being none the wiser.

The aforementioned MiRAs are primarily concerned with agents as an interface metaphor rather than a system engineering technique. As such, they mostly employ finite state machines or behaviour engines to drive the agent’s actions and are consequently located in the weak plane of agency.

Strong agent architectures, on the other hand, explicitly model an agent’s internal state in terms of folk psychological terms, such as desires, intentions, or emotions. By allowing agents to control these internal states and their reactions to it as well as processing their rich sensory MR environment, developers hope to create more behaviourally realistic MR applications (O’Hare et al., 2005b).

Figure 7: Strong Mixed reality agents with stronger virtual embodiment. From left to right: The Invisible Person, ALIVE, GEIST and MRIS.

In the ALIVE system (Maes et al., 1995), for example, the user’s environment is augmented with a virtual dog character that can be directed using hand gestures. The agent architecture models a set of needs and motivations for the dog character and chooses its behaviours accordingly. ALIVE is also remarkable for being the very first

Mixed Reality Agent recorded in scientific literature and was years ahead of similar efforts.

The Invisible Person (Psik et al., 2003) is based on the ALIVE system and employs a humanoid virtual character in an effort to engage visitors at the Vienna Museum of Technology in a game of Tic Tac Toe. The game board is digitally added onto the floor, and both user and character control the game via body postures and hand gestures. The character's internal state is based on an emotional system that directs its actions, facial expressions and manner of interaction with the user. Feedback from visitors of the exhibition commend the lifelikeness of the agent.

Storytelling engines also explicitly model individual agents' goals and motivations in order to create dynamic narratives from the interplay of these goals and the users actions. One of the first projects to apply digital narrative to Mixed Reality is the cultural heritage application GEIST (Kretschmer et al., 2001). GEIST¹ immerses the user in a thrilling adventure involving events from the Thirty Years' War. As the user roams the old town of Heidelberg, he can enter certain 'hotspots' in which ghosts from the past appear in the form of virtual characters. They plead for the user's help in solving the mystery surrounding their death, creating a quest around the city in which the user has to learn about the history of places and events in order to succeed.

In the physical domain, GEIST agents are limited to sensing the user's position and orientation, while in the virtual, which is populated with spatially aligned models of buildings and other virtual objects, interaction is much more varied and versatile. However, due to the spirit nature of the GEIST agents, corporeal presence of the ghosts is inhibited, as they appear translucent and float in midair.

Another prominent example of MR storytelling is the Mixed-Reality Interactive Storytelling (MRIS) project (Charles et al., 2004), which allows a user to immerse himself into a spy thriller story in the role of the villain. It does so by capturing the user's image in real time, extracting it from the background, and inserting it into a virtual world populated by autonomous synthetic actors with which the user then interacts using natural language and gestures. The resulting image is projected onto a large screen facing the user, who sees his own image embedded in the virtual stage alongside the synthetic actors. Notably, when viewed from the perspective of Milgram's continuum of MR displays (see Figure 2), MRIS is a rare example exemplifying the concept of Augmented Virtuality, i.e. a virtual world with added 'real' components.

Finally, Virtual Gunslinger (Hartholt et al., 2009), is a more recent example of a Mixed Reality storytelling experience. In it, the user plays the character of a cowboy in a Wild West saloon who gets challenged to a duel. The user is placed in an environment featuring a real bar counter and a virtual bartender and outlaw, both of which are projected onto screens placed in the room. The user can interact with the agents using natural language dialogues and gestures, e.g. moving his arm as if to pull a gun when duelling with the outlaw.

Common to all these strong agents is that the agent architecture facilitates the development of agents that exhibit realistic and lifelike behaviour. But strong agent system are also often used to realise highly complex and distributed systems that deal with dy-

---

¹German for ghost

namic and uncertain environments. NeXus (O’Hare et al., 2005a, 2004) and UbiAgent (Barakonyi, 2006), for example, are two fully autonomous and goal-driven BDI agent projects that favour a strong agent model for exactly these reasons.

Barakonyi illustrates the UbiAgent concept through a number of applications. In the AR Lego application, for example, a virtual agent instructs the user in situ on how to assemble a LEGO® Mindstorms robot by standing next to it and overlaying a virtual model of the next brick onto the physical robot. Corporeal presence is heightened, as the instructor gets occluded by the physical LEGO® model as he walks around it.

In contrast to UbiAgent, where each virtual character is encapsulated as a single agent, NeXus is an inherently multi-agent based framework that uses AgentFactory (O’Hare et al., 1998) agents to control all aspects of the environment. This allows developers to create exceedingly complex MR scenarios and highly immersive environments imbued with context-sensitive intelligence. In particular, the NeXus approach allows an agent ‘ghost’ to take hold of a ‘shell’ (i.e. a corporeal presence and its associated interactive capacity) in order to take on a visible embodiment. While this approach can be seen as a separation of deliberation and embodiment, it is by no means a reprise of ‘Good Old Fashioned AI’ (Haugeland, 1985). On the contrary, NeXus not only strongly advocates the notion of embodiment but actually extends it by allowing agents to reason about their corporeal presence and interactive capacity.

4.2. *Virtual Agents Interacting with the Physical World & Physical Agents Interacting with the Virtual World*

Figure 8: MiRAs with stronger virtual corporeal presence that predominantly interact with the physical world. From left to right: The Virtual Room Inhabitant, Max, and MACK.

While the MiRAs mentioned in the last section are primarily represented in, and interact with, the virtual world, examples also exist for MiRAs that, despite possessing only a virtual representation, predominantly interact with the physical world (see Figure 8). Endowed with spatial awareness and knowledge about the physical world, MiRAs in this type of scenario usually serve as interface agents to the user’s physical environment, providing context-sensitive and spatially referenced information. As there is no virtual world as such to sense, these agents only act on the virtual domain by means of animating their bodies. Similarly, they only sense the physical environment but cannot affect it.

MACK (Cassell et al., 2002), the Media Lab Autonomous Conversation Kiosk, for example, combines natural language output with directional gestures to direct visitors at the MIT Media Lab. MACK uses himself (“behind me”), the user (“to your left”), and landmarks (“next to the printer”) as spatial reference points when giving directions. The conversational agent Max spatially references museum exhibits in a similar manner in the Heinz Nixdorf MuseumsForum in Paderborn, Germany. Displayed life-sized on a screen within the museum, Max engages visitors in small talk and provides them with information about the exhibits. Max is based on a sophisticated reactive-deliberative BDI architecture and can track and distinguish between multiple users.

Another example of such an interface agent is the Virtual Room Inhabitant (VRI) (Kruppa et al., 2005). Contrary to Max and MACK, however, the VRI is not displayed on a fixed screen but projected onto the walls of a shop. By tracking the position of users and goods (via infrared and RFID sensors) the VRI act as a smart shopping assistant, projecting itself next to products of interest and providing information about them while facing the user. While this approach allows a virtual character to share the user’s physical world beyond the confines of a screen without the aid of a HMD, it does necessitate a virtual model of the store in order to compensate for the distortion that results from non-normal projection.

As previously pointed out, all of the above interface agents do not act on the physical environment but merely observe it. One, and so far the only, MiRA with a stronger virtual corporeal presence that can physically affect the real world is realised in the Virtual Brownies (Aoki et al., 2005) project. In it, small virtual creatures, called Kobito, can move around physical objects such as a tea caddy. While no details of the agent system are given, the agents seem to exhibit autonomous, reactive and pro-active behaviour and interact with each other and (indirectly) with the user, making them, at the very least, weak agents. This project also very neatly illustrates the interplay of corporeal presence and interactive capacity.

The agents’ virtual bodies define their virtual corporeal presence, while their movement with regard to the physical table and other objects endows them with a certain degree of physical corporeal presence. The ability to sense the physical objects on the table and to move the physical tea caddy defines their physical interactive capacity. In the current incarnation, their virtual interactive capacity only seems to consist of interacting with each other, as there are no purely virtual objects in the environment.

Contrary to virtual agents that interact with the physical world, where a number of possible applications come to mind, one might seem hard pressed to come up with a conceivable scenario for a MiRA of the other extreme, i.e. a physical agent that predominantly interacts with the virtual world. Such a MiRA would essentially instantiate itself as a robot that senses and affects virtual objects but does not sense physical ones (as it arguably has to act on the physical domain in order to constitute a robot). But even though no such MiRA might exist at present, this is not to suggest that research in this area should not be undertaken.

4.3. MiRAs with Stronger Physical Embodiment

MiRAs with stronger physical corporeal presence (see Figure 9), i.e. MiRAs with a robotic body, can be seen, from a robotics point-of-view, as the extension of model-based robot control architectures, whereby robots employ a virtual representation of the

Figure 9: MiRAs with stronger physical embodiment. Clockwise from top left: Sheep & Wolves, Virtual Pheromones, Virtual Synergy, Giesler et al.’s AR robot control system, and Stilman et al.’s AR path-planning visualisation

environment they inhabit. By leveraging the steady progresses achieved in the areas of model-building and model-matching techniques, these virtual representations are now commonly used in robotics to augment the robots’ immediate sensing capabilities and to provide a memory of past observations and events (e.g. associated with specific locations). They can also allow robots to benefit from a global or overhead perspective of the environment.

Many of these robotic systems can be considered examples of MiRAs due to the way their environment models do not only mirror the real world (ideally faithfully), but also augment it with purely virtual elements to improve the robot’s interaction capabilities. Payton et al. (2001), for instance, developed a robot swarm for search-and-rescue scenarios that can lead human users to a victim without the need for a physical pointing component, such as an arm. The swarm uses a bio-inspired virtual pheromone infrastructure that is associated with a shared map of the environment. Each robot releases virtual pheromone trails into this infrastructure that appear as virtual arrows, superimposed on each robot in the swarm, which point in the direction of the nearest victim.

Notably, the effectiveness of this system depends upon the ability of each robot to localise itself within a shared global frame of reference. Only in this manner can position updates and pheromones refer to the shared model and be used to coordinate the robots in the real world, thereby enabling coordination by stigmergy.

As a consequence of the registration between a robot’s position and its environmen-

tal model, human operatives can also be aware of the robot's activities by monitoring their own version of the map, often displayed as a 3D model. Collaborative Virtual Environments (CVEs) can then be used to support human interaction with robot collectives. However, rather than in the real world, CVE-enabled interaction happens in a virtual environment, which can either be fully simulated or mirror a real (and usually remote) environment. For example, in the Virtual Synergy search-and-rescue system (Tejada et al., 2004), real robots and human users are represented in the CVE by avatars mirroring the actions of their real counterparts. As they navigate their avatars through this shared virtual environment, users can give commands to robots, such as "come here", which instructs a robot to approach the physical position corresponding to the virtual location of the user's avatar. They may also leave 'hotspots' in the virtual environment as suggestions of where the robots should explore next.

While CVEs can act as interaction medium, allowing humans to supervise and direct robots from remote, stationary workstations, there are also examples of systems addressing co-located human-robot interaction. These systems usually rely on an MR visualisation, especially for its ability to visualise geometrical properties such as planned trajectories, sensor data and world models. This enables human operators to evaluate the state, and gain an understanding, of a robot system whilst occupying the same environment. To do this, these applications demand tracking the position of the user's HMD in respect to the robots so that users can see both the real world and the synthetic imagery (e.g. user interface components, robot sensor data) registered with and superimposed over the real elements in the scene.

A notable example of such systems is given by Stilman et al. (2005), who equipped a laboratory space with a motion capture system, enabling large area coverage for accurate and fast object localisation and tracking. Their application, addressing development and experimentation with humanoid robots, tracks both robots and obstacles, such as chairs and tables, using retro-reflective markers. It combines these with virtual artefacts, such as simulated robots and simulated objects, creating a hybrid real/virtual environment. By equipping researchers with HMDs and tracking them like any other object, the hybrid scene can be visualised using an immersive AR interface. In this manner, for example, researchers can view real or simulated robots during walking tests among both real and virtual obstacles. In addition, they can also avail of virtual artefacts, such as the visualisation of the robot's intended footsteps on the floor, to aid their judgement of the robot's sensing and path planning capabilities.

Rather than just enhance human observation capabilities, these systems can also support more sophisticated co-located Human-Robot Interaction (HRI). For example, Giesler et al. (2004) developed a mobile AR system for the control of an autonomous mobile platform. Like the above robotic MiRAs, this system combines robot self-localisation with the localisation of the HMD-wearing user in the same frame of reference. The position and orientation of the user's HMD is found by means of multiple ARToolkit (Kato and Billinghurst, 1999) fiducials distributed in the working area. The user is also equipped with a magic wand, in form of a pen terminating in a cube formed by ARToolkit fiducials. The orientation of the cube is visible from every angle, and the user can point with it to imaginary points on the floor in a manner similar to a mouse. Giesler's particular application allows the user to rapidly prototype and manipulate topological maps interactively defined in the working environment, and then

Figure 10: MiRAs with Mixed Reality embodiment. From left to right: U-Tsu-Shi-O-Mi (Virtual Humanoid), Jeeves, and Dragone et al.’s MiRA.

subsequently instruct the robot.

Another HRI example, although for a different application, is constructed in the Sheep & Wolves project (Xin and Sharlin, 2006), whereby the context of a board game is employed to study human-robot interaction, using four Aibo robots (the wolves) and a virtual sheep. Sheep & Wolves is of particular note for being the only project that explicitly mixes different types of MiRAs, i.e. agents with stronger physical embodiment (the wolves) and agents with stronger virtual embodiment (the sheep), in a single MR environment. Three of the robots and the sheep are autonomous agents whereas the fourth robot is controlled by the user. Xin and Sharlin use the scenario as a testbed to examine people’s attitudes when confronted with different robot personalities, expressed during every round of play when the wolf agents and the human user discuss the next move amongst themselves.

4.4. *MiRAs with Mixed Reality Embodiment*

A notable precursor of systems with a strong physical and virtual corporeal presence is the Agent Chameleon project (O’Hare et al., 2003b; Duffy et al., 2003). The specific scenario explored by this project was to develop agents capable of migration and mutation within and between both real and virtual spaces, thereby expanding the traditional concept of an agent’s environment and its constraints. Agent Chameleons are capable of migration between embodiments within virtual environments (e.g. a virtual avatar), physical environments (e.g. a robot) and software environments (e.g. PCs, PDAs). This allows the Agent Chameleon to sense and act within each of these environments. For example, the Agent Chameleon may be located within a virtual environment, interacting with its user, when it is tasked with performing an action within the real world. It can then migrate to a robot to perform this action before returning to the virtual environment to report the results. When instantiated within each of these environments, the agent has knowledge of that particular environment and of its capabilities therein.

Of particular note is the Agent Chameleon agent’s ability to mutate its embodiment form, which further distinguishes it from previously mentioned MiRAs that move between environments, such as Ritchie, PECA, and the Herding Sheep project. This is particularly true when located within virtual environments, where the agent is free of constraints that may exist in the real world. An Agent Chameleon is capable of modifying its embodiment instantiation in response to environmental and task specific events (O’Hare et al., 2003a). For instance, the agent can adapt to a change in its environment, such as it starting to rain, by mutating its form, e.g. by sprouting an umbrella.

Agent Chameleons can be seen as a step towards MiRAs. While they are capable of sensing and acting within both real and virtual worlds, the environment within which they operate is not a true MR environment, but rather two distinct worlds — one physical, one virtual — with an agent capable of moving between the two.

On the other hand, a number of more recent applications truly combine physical and virtual embodiment by equipping a robot with a virtual character through MR visualisation, seeking to overcome the limitations of a purely real or virtual embodiment while taking advantage of both. The robot agents in these systems exhibit tangible physical presence whilst the virtual persona offers rich expressional capabilities and personalisation features that are complex and expensive to realise with purely hardware-based solutions.

Indeed, contrary to agents in the Virtual Synergy project, whose virtual and physical corporeal presence are essentially identical, both from the agent’s and the user’s point of view, the very identity of the following Mixed Reality Agents is constituted by both virtual and real components, making them a characteristic example of embodiment in the MR environment. This means that, although the agents are composed of distinct virtual and real elements, they exhibit cohesion and behavioural consistency to the observer. Such consistency has been termed “behavioural realism” (O’Hare et al., 2005b) and stands in stark contrast to the traditional driver of visual realism.

Dragone et al.’s MiRA system (Dragone et al., 2006) constitutes the first example of this class of application. It seeks to seamlessly integrate physical robots and virtual characters visualised through augmented overlays. Instrumental to the fulfilment of these potentials is the coordination between the real and virtual parts enabled by the agent-based software engineering methodology employed in the implementation of Dragone et al.’s MiRA system. Specifically, the functionalities of a MiRA emerge from the collaboration of a network of distributed agent components: those in control of the robotic platforms, those managing the user interfaces, and those in control of the virtual characters. Such an organisation enables the agent to adapt to different users and different robots as there is no predefined coupling between the robot and the appearance or 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.

As such, Dragone et al.’s MiRA is an ideal test-bed for the Agent Chameleons scenario. It opens the possibility to investigate robotic Agent Chameleons that can mutate their form in both the virtual and Mixed Reality environments according to the capabilities and expressions that they wish to utilise, or displaying the same characters to the user in both, preserving all the advantages of the familiar relationships between the user and its assistant.

The focus of the Jeeves project (Young et al., 2007), on the other hand, which similarly combines a robot (the iRobot Roomba vacuum cleaner) with a cartoon-like character in order 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 insights into the robot’s state while, at the same time, avoiding to overly increase people’s expectations as more realistic and human-like representations could. Young et al. point out that cartoon art techniques can augment and compliment 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, Shoji et al. (2006)), 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. In stark contrast to Dragone et al.’s MiRA system and Jeeves, the Virtual Humanoid does not sense the virtual domain, making it the only MiRA with equal virtual and physical corporeal presence but greater virtual interactive capacity.

Collectively, these applications showcase the advantages of employing Mixed Reality to combine physical robot platforms with virtual characters. 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. In every case, and in contrast to robots with virtual characters visualised on a screen attached to them, such as GRACE (Simmons et al., 2003) and VALERIE (Gockley et al., 2005), the Mixed Reality characters are visible from all angles and are not subjected to diminishing visibility at greater distances.

Furthermore, the virtual characters are in a unique position to augment the HRI capabilities of the robot. For example, the virtual character can act as the robot’s social interface, helping the user to infer the robot’s intentions and actions via gaze, facial expressions, deictic gestures, and through choosing a particular form of corporeal presence, as in some demonstrations of the MiRA system (Holz et al., 2006).

Due to their wearable nature, these systems are advantageous in applications with a high robot-to-user ratio, as they can substitute cheap robotic hardware for expensive humanoid robots. A single wearable interface can augment the interaction capabilities of multiple simple robots (e.g. without screens, heads or arms) and, possibly, even portray different characters to different users simultaneously.

5. The Efficacy of the MiRA Cube

The exemplar MiRA systems detailed in the previous sections have been carried out in diverse domains, such as robotics or embodied conversational agents. As a result of this diversity, developers are often unaware of the full range of Mixed Reality Agent research and are therefore prone to replicate prior efforts.

This paper surveys and draws together diverse strands of research into a single unified taxonomy and, in so doing, provides a common portal for the investigation

of Mixed Reality Agents. The establishment of a common field of research gives an instant advantage to anybody interested in developing similar systems. Researchers become aware of related work, and the unified taxonomy acts as a catalyst for the sharing of methodologies, examination of interoperability, and development of standard frameworks for the creation of MiRAs.

This paper has introduced the MiRA Cube as a mechanism for achieving the christening and categorisation of Mixed Reality Agents. The proposed principal axes of agency, corporeal presence, and interactive capacity have been further subdivided into categories, namely those of weak and strong agency; stronger virtual, stronger physical, and equal virtual and physical corporeal presence and interactive capacity, respectively. These categories possess somewhat inexact and malleable boundaries, as it is rather arbitrary to try to quantify the amount of, say, virtual and physical sensing and acting capabilities.

In spite — or indeed because — of this, the MiRA Cube has proven to be a highly effective classification aid. It offers a three-dimensional scaffolding upon which pre-existing MiRAs can be placed. Experience has shown that situating existing research within the three-dimensional cube is both intuitive and informative, as the MiRA Cube has facilitated the identification of coherent and cohesive clusters which further assist in characterising the space.

The taxonomy also enables the user to move through the three-dimensional space, observing closely related and indeed disparate projects. The artefact thus provides a valuable visualisation metaphor for MiRA researchers. Indeed, one of the few MiRAs that does not conveniently locate itself within the taxonomy, is the Virtual Synergy project (Tejada et al., 2004). Corporeal presence for these agents is equally strong in the physical domain as it is in the virtual. However, while both the robotic body and its virtual representation are inextricably linked, the parts do not contribute to a single, unified corporeal presence as, for example, Dragone et al.'s MiRA does. Instead, one is a copy of the other and, from the agent's and user's point of view, both are essentially one and the same thing.

Another unusual case, which does not fall under any easy heading, is the Virtual Humanoid (Shoji et al., 2006). Although the virtual representation faithfully mirrors the shape of the physical robot, it contributes something very unique to the Virtual Humanoid's appearance. However, the complete occlusion of its physical body and the lack of interaction with any virtual elements stops the Virtual Humanoid from realising the full potential of its Mixed Reality embodiment.

Finally, Sheep & Wolves (Xin and Sharlin, 2006) deserves special mention, being the only project in which different types of MiRAs interact, offering a glimpse into the potential of combining different types of MiRAs into a complete system.

6. Future Research

In defining a taxonomy for MiRAs and populating it with existing work, a number of potential future avenues of research become apparent. In visualising MiRA projects (Figure 5), it can be clearly observed that the spread of projects across the field is not uniform. As mentioned earlier, projects involving physical agents interacting with

the virtual world are currently non-existent. Further research could tell if this is due to a problem in engineering such agents or if there is little possibility of meaningful applications.

A second research possibility is suggested by the comparative lack of virtual agents truly interacting with the physical world. While the Virtual Brownies (Aoki et al., 2005) project shows one possible solution to realising such agents, their approach, which uses magnets placed underneath the table, is cumbersome and difficult to deploy in a larger fashion. A different example of a MiRA with strong virtual corporeal presence affecting the physical world that suggests itself, would be a spatially aware virtual character that serves as an interface to a ubiquitous home. If such an agent were to monitor and control physical appliances such as the TV, window blinds, and the heating, it would indeed possess stronger physical interactive capacity while having only a virtual representation.

However, only detailed interaction studies with human users can tell which MiRAs are most effective. Indeed, throughout the reviewed literature exists a noticeable gap in user studies. Only a handful of the systems were formally evaluated, mostly in terms of acceptance and likeability. This not necessary the fault of the researchers involved, as these project are technically challenging to construct in their own right. However, user studies are necessary for the field to mature to establish MiRAs as a human-computer interaction tool. Studies examining users' experiences of MiRAs could lead to best practice outlines for which level of virtual and physical corporal presence is most suited to any given task. The level of interactive capacity chosen should also be examined, as users may not be comfortable allowing a previous virtual entity to interact with objects in the physical world.

The common language developed to define the taxonomy is aimed at stopping the cycle of reinvention that has beset the projects listed in the previous section. But a taxonomy should not constitute an absolute frame of reference. Instead, future efforts might well cause the definitions to develop over time, changing what would and would not be considered a Mixed Reality Agent. For example, a more nuanced definition of agency would result in more planes than just the two presented. With a continuous definition of agency, the taxonomy would truly form a three-dimensional cube. In general, focusing or, indeed, widening the definitions of agency and embodiment would, consequently, reduce or increase the sphere of projects that could be classified MiRAs.

7. Conclusions

This paper considers the compelling interaction metaphor of Mixed Reality. Specifically it has defined and characterised Mixed Reality Agents (MiRA), an emerging field of research endeavour. It has postulated a unified MiRA taxonomy, a common framework which facilitates awareness, comparison, interoperability, reuse, and the emergence of standards.

MiRAs constitute a new and exciting field of research with numerous potential application scenarios and fascinating implications. MiRAs have the potential to become a truly *disruptive* technology, offering a more natural interface that facilitates user access to, and interaction with, ever more compelling and complex environments. The combination of factors like the advent of cheap and relatively powerful mobile phones

and the plummeting costs of HMDs, coupled with ongoing research into natural tracking (Wagner et al., 2008), suggest that mass market applications of Mixed Reality are a real near-term possibility.

The underlying architecture needed to achieve this necessitates both scalability and extensibility, permitting the use of disparate local resources and heterogeneous devices. Multi-agent systems suggest themselves as a natural choice for delivering the requisite adaptivity and intelligence. Previous research in Agent Chameleons (O'Hare et al., 2003b) has shown how agents can change body forms from virtual to physical and back again, in adapting to any given task. MiRAs are in essence an extension of this concept and thus could prove ideal in any mobile based environment where resources are constrained.

8. Acknowledgements

This work is supported by Science Foundation Ireland under grant 07/CE/I1147.

References

- Adams, D., 1990. Hyperland. BBC.
- Anabuki, M., Kakuta, H., Yamamoto, H., Tamura, H., April 2000. Welbo: An embodied conversational agent living in mixed reality space. In: Proceedings of the Conference on Human Factors in Computing Systems - CHI '00. The Hague, The Netherlands, pp. 10–11.
- Aoki, T., Matsushita, T., Iio, Y., Mitake, H., Toyama, T., Hasegawa, S., Ayukawa, R., Ichikawa, H., Sato, M., Kuriyama, T., Asano, K., Kawase, T., Matumura, I., 2005. Kobito: Virtual bbrownies. In: SIGGRAPH '05: ACM SIGGRAPH 2005 Emerging technologies. ACM, New York, NY, USA, p. 11.
- Barakonyi, I., October 2006. Ubiquitous animated agents for augmented reality. Ph.D. thesis, Technische Universität Wien.
- Billinghurst, M., Kato, H., Poupyrev, I., 2001. The magicbook: a transitional ar interface. *Computers & Graphics* 25 (5), 745–753.
- Bornik, A., Beichel, R., Reitingner, B., Gotschuli, G., Sorantin, E., Leberl, F. W., Sonka, M., 2003. Computer-aided liver surgery planning: An augmented reality approach. In: Proceedings of SPIE Medical Imaging.
- Brooks, R. A., 1991a. Intelligence without reason. In: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence - IJCAI 91. Sydney, New South Wales, Australia, pp. 569–595.
- Brooks, R. A., January 1991b. Intelligence without representation. *Artificial Intelligence* 47 (1-3), 139–159.

- Cassell, J., Stocky, T., Bickmore, T., Gao, Y., Nakano, Y., Ryokai, K., Tversky, D., Vaucelle, C., Vilhjálmsson, H., 2002. MACK: Media lab Autonomous Conversational Kiosk. In: Proceedings of IMAGINA'02.
- Charles, F., Cavazza, M., Mead, S. J., Martin, O., Nandi, A., Marichal, X., June 2004. Compelling experiences in mixed reality interactive storytelling. In: Proceedings of the 2004 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology - ACE 04. Singapore, pp. 32–40.
- Dautenhahn, K., October 1998. The art of designing socially intelligent agents - science, fiction and the human in the loop. *Applied Artificial Intelligence* 12 (8-9), 573–617.
- Dautenhahn, K., 1999. Embodiment and intelligence in socially intelligent life-like agents. In: Nehavim, C. L. (Ed.), *Computation for Metaphors, Analogy, and Agents*. Vol. 1562 of *Lecture Notes in Artificial Intelligence*. Springer.
- Dorfmueller-Ulhaas, K., André, E., 2005. The synthetic character Ritchie: First steps towards a virtual companion for mixed reality. In: *Fourth IEEE International Symposium on Mixed and Augmented Reality - ISMAR '05*. IEEE Computer Society, Washington, DC, USA, pp. 178–179.
- Doswell, J. T., 2005. It's virtually pedagogical: Pedagogical agents in mixed reality learning environments. In: *Proceedings of the Thirtysecond International Conference on Computer Graphics and Interactive Techniques - SIGGRAPH 2005 (Educators Program)*. Los Angeles, California, p. 25.
- Dourish, P., 2001. *Where The Action Is: The Foundations of Embodied Interaction*. MIT Press.
- Dragone, M., Holz, T., O'Hare, G. M. P., January 2006. Mixing robotic realities. In: *Proceedings of the 2006 International Conference on Intelligent User Interfaces - IUI 2006*. Sydney, New South Wales, Australia.
- Duffy, B., November 2000. *The social robot*. Ph.D. thesis, University College Dublin.
- Duffy, B. R., O'Hare, G. M. P., Bradley, J. F., Martin, A., Schön, B., 2005. Future reasoning machines: Mind & body. *Kybernetes* 34 (9/10), 1404–1420.
- Duffy, B. R., O'Hare, G. M. P., Martin, A., Bradley, J. F., Schön, B., May 2003. Agent chameleons: Agent minds and bodies. In: *Proceedings of the Sixteenth International Conference on Computer Animation and Social Agents - CASA 2003*. New Jersey, USA.
- Giesler, B., Salb, T., Steinhaus, P., Dillmann, R., April 2004. Using augmented reality to interact with an autonomous mobile platform. In: *Proceedings of the IEEE International Conference on Robotics and Automation - ICRA 2004*. New Orleans, Louisiana, USA.

- Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A. C., Wang, J., August 2005. Designing robots for long-term social interaction. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems - IROS 2005.
- Hartholt, A., Gratch, J., Weiss, L., The Gunslinger Team, 2009. At the virtual frontier: Introducing gunslinger, a multi-character, mixed-reality, story-driven experience. In: Intelligent Virtual Agents. Vol. 5773/2009 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 500–501.
- Haugeland, J., 1985. Artificial Intelligence: The Very Idea. Massachusetts Institute of Technology, Cambridge, MA, USA.
- Holz, T., Dragone, M., O'Hare, G. M. P., Martin, A., Duffy, B. R., 2006. Mixed reality agents as museum guides. In: ABSHL '06: Agent-Based Systems for Human Learning, AAMAS 2006 Workshop. ACM Press, New York, NY, USA.
- Iso, M., 2007. Dennō coil. NHK Educational.
- Kato, H., Billinghurst, M., October 1999. Marker tracking and HMD calibration for a video-based augmented reality conferencing system. In: Proceedings of the Second International IEEE Workshop on Augmented Reality - IWAR 99. San Francisco, California, USA.
- Kretschmer, U., Coors, V., Spierling, U., Grasbon, D., Schneider, K., Rojas, I., Malaka, R., November 2001. Meeting the spirit of history. In: Proceedings of the 2001 Conference on Virtual Reality, Archeology and Cultural Heritage - VAST 01. Glyfada, Greece, pp. 141–152.
- Kruppa, M., Spassova, L., Schmitz, M., 2005. The Virtual Room Inhabitant – Intuitive Interaction with Intelligent Environments. Vol. 3809 of Lecture Notes in Computer Science. Springer-Verlag, Berlin/Heidelberg, Ch. 25, pp. 225–234.
- Laurel, B., 1990. Interface agents: Metaphors with character. In: Laurel, B. (Ed.), The Art of Human-Computer Interface Design. Addison-Wesley.
- MacWilliams, A., Sandor, C., Wagner, M., Bauer, M., Klinker, G., Bruegge, B., 2003. Herding sheep: Live system development for distributed augmented reality. In: Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality - ISMAR '03. IEEE Computer Society, Washington, DC, USA, pp. 123–133.
- Maes, P., Darrell, T., Blumberg, B., Pentland, A., April 1995. The ALIVE system: Full-body interaction with autonomous agents. In: Proceedings of Computer Animation '95. Geneva, Switzerland.
- Milgram, P., Kishino, F., December 1994. A taxonomy of mixed reality visual displays E77-D (12).
- Newell, A., Simon, H. A., March 1976. Computer science as empirical inquiry: Symbols and search. Communications of the ACM 19 (3), 113–126.

- Nwana, H. S., October/November 1996. Software agents: An overview. *Knowledge Engineering Review* 11 (3), 205–244.
- O'Hare, G. M., Campbell, A. G., Stafford, J. W., Aiken, R., October 2005a. NeXuS: Behavioural realism in mixed reality scenarios through virtual sensing. In: *Proceedings of the Eighteenth International Conference on Computer Animation and Social Agents - CASA 2005*. Hong Kong.
- O'Hare, G. M. P., Campbell, A. G., Stafford, J. W., May 2005b. NeXuS: delivering behavioural realism through intentional agents. In: *Proceedings of the 3rd International Conference on Active Media Technology - AMT 2005*. IEEE Press.
- O'Hare, G. M. P., Collier, R., Conlon, J., Abbas, S., August 1998. Agent factory: An environment for constructing and visualising agent communities. In: *Proceedings of the Ninth Irish Conference on Artificial Intelligence and Cognitive Science - AICS '98*. Dublin, Ireland, pp. 249–261.
- O'Hare, G. M. P., Duffy, B. R., Bradley, J. F., Martin, A. N., February 2003a. Agent chameleons: Moving minds from robots to digital information spaces. In: *Autonomous Minirobots for Research and Edutainment - AMiRE 2003*.
- O'Hare, G. M. P., Duffy, B. R., Campbell, A. G., July 2004. NeXuS: Mixed reality experiments with embodied intentional agents. In: *Proceedings of the Seventeenth International Conference on Computer Animation and Social Agents - CASA 2004*. Geneva, Switzerland.
- O'Hare, G. M. P., Duffy, B. R., Schön, B., Martin, A., Bradley, J. F., September 2003b. Agent chameleons: Virtual agents real intelligence. In: *Proceedings of the Fourth International Working Conference on Intelligent Virtual Agents - IVA 2003*. Vol. 2792 of *Lecture Notes in Computer Science*. Irsee, Germany, pp. 218–225.
- Payton, D., Daily, M., Estowski, R., Howard, M., Lee, C., November 2001. Pheromone robotics. *Autonomous Robots* 11 (3), 319–324.
- Pfeifer, R., August 2002. On the role of embodiment in the emergence of cognition: Grey Walter's turtles and beyond. In: *Proceedings of the International Workshop on Biologically-Inspired Robotics: The Legacy of W. Grey Walter*. Bristol, England.
- Psik, T., Matković, K., Sainitzer, R., Petta, P., Szalavari, Z., 2003. The invisible person: advanced interaction using an embedded interface. In: *Proceedings of the Workshop on Virtual Environments - EGVE '03*. ACM Press, New York, NY, USA, pp. 29–37.
- Quick, T., Dautenhahn, K., Nehaniv, C. L., Roberts, G., September 1999. On bots and bacteria: Ontology independent embodiment. In: *Proceedings of the European Conference on Artificial Life - ECAL'99*. Lausanne, Switzerland, pp. 339–343.
- Rao, A. S., Georgeff, M. P., June 1995. BDI agents: From theory to practice. In: *Proceedings of the First International Conference on Multi-Agent Systems - ICMAS 95*. San Francisco, California, USA, pp. 312–319.

- Riegler, A., September 2002. When is a cognitive system embodied? *Cognitive Systems Research* 3 (3), 339–348.
- Schmeil, A., Broll, W., March 2007. MARA: A mobile augmented reality-based virtual assistant. In: *Proceedings of the IEEE Virtual Reality Conference - VR 2007*. Charlotte, North Carolina, USA, pp. 267–270.
- Sharkey, N., Ziemke, T., March 2000. Life, mind and robots: The ins and outs of embodied cognition. In: Wermter, S., Sun, R. (Eds.), *Hybrid Neural Systems*. Vol. 1778 of *Lecture Notes in Artificial Intelligence*. Springer.
- Shoham, Y., March 1993. Agent-oriented programming. *Artificial Intelligence* 60 (1), 51–92.
- Shoji, M., Miura, K., Konno, A., August 2006. U-Tsu-Shi-O-Mi: The virtual humanoid you can reach. In: *Proceedings of the Thirtythird International Conference on Computer Graphics and Interactive Techniques - SIGGRAPH 2006*. Boston, Massachusetts, USA.
- Sielhorst, T., Obst, T., Burgkart, R., Riener, R., Navab, N., 2004. An augmented reality delivery simulator for medical training. In: *International Workshop on Augmented Environments for Medical Imaging - MICCAI Satellite Workshop*.
- Simmons, R., Goldberg, D., Goode, A., Monetmerlo, 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., Milman, T., Maxwell, B., 2003. Grace: An autonomous robot for the AAAI robot challenge. *AI Magazine* 24, 51–72.
- Steels, L., 2000. Digital angels. In: *Die Gegenwart der Zukunft*. Vol. 367 of *Wagenbach Taschenbuch*. Verlag Klaus Wagenbach, translated from original German by the author.
- Stilman, M., Chestnutt, J., Michel, P., Nishiwaki, K., Kagami, S., Kuffner, J., 2005. Augmented reality for robot development and experimentation. Tech. Rep. CMU-RI-TR-05-55, Robotics Institute, Carnegie Mellon University.
- Tejada, S., Tarapore, S., Cristina, A., Goodwyne, P., O'Hara, R., June 2004. Mixed-initiative interface for human, robot, agent collaboration in urban search and rescue teams. In: *Proceedings of the World Automation Congress*. Seville, Spain, pp. 467–472.
- Tomlinson, B., Yau, M. L., Baumer, E., Goetz, S., Carpenter, L., Pratt, R., Young, K., May-Tobin, C., 2006. The ecraft project: a multi-device interactive graphical exhibit for learning about restoration ecology. In: *CHI '06 extended abstracts on Human factors in computing systems*. ACM, New York, NY, USA, pp. 1445–1450.
- Tomlinson, B., Yau, M. L., O'Connell, J., Williams, K., Yamaoka, S., 2005. The virtual raft project: A mobile interface for interacting with communities of autonomous characters. In: *CHI '05 extended abstracts on Human factors in computing systems*. ACM, New York, NY, USA, pp. 1150–1151.

- Wagner, D., Reitmayr, G., Mulloni, A., Drummond, T., Schmalstieg, D., 2008. Pose tracking from natural features on mobile phones. In: Seventh IEEE/ACM International Symposium on Mixed and Augmented Reality - ISMAR '08. pp. 125–134.
- Wang, X., Jeong Kim, M., 2009. Exploring presence and performance in mixed reality-based design space. In: Wang, Xiangyu; Schnabel, M. A. (Ed.), *Mixed Reality In Architecture, Design, And Construction*. Springer.
- Wiendl, V., Dorfmueller-Ulhaas, K., Schulz, N., André, E., 2007. Integrating a virtual agent into the real world: The virtual anatomy assistant Ritchie. In: *Proceedings of the 7th International Conference on Intelligent Virtual Agents - IVA '07*. Springer-Verlag, pp. 211–224.
- Wooldridge, M., Jennings, N. R., 1994. Intelligent agents: Theory and practice. *Knowledge Engineering Review* 10 (2), 115–152.
- Wooldridge, M., Jennings, N. R., 1995. Intelligent agents: Theory and practice. *Knowledge Engineering Review* 10 (2), 115–152.
- Xin, M., Sharlin, E., April 2006. Sheep and wolves: Test bed for human-robot interaction. In: *Proceedings of the Conference on Human Factors in Computing Systems - CHI 2006*. Montral, Québec, Canada.
- Young, J. E., Xin, M., Sharlin, E., March 2007. Robot expressionism through cartooning. In: *Proceedings of the 2007 ACM/IEEE International Conference on Human-Robot Interaction*. Arlington, Virginia, USA.
- Ziemke, T., September 2001. Are robots embodied? In: *Proceedings of the First International Workshop on Epigenetic Robots: Modeling Cognitive Developments in Robotic Systems*. Vol. 85 of Lund University Cognitive Studies. Lund, Sweden, pp. 75–93.

A. Visualising the MiRA Cube

Mixed Reality research has taught us that three dimensions are better than two. Consequently, a special *MiRA Marker* is provided as an appendix to this paper, which, together with a small program on the authors' homepage (<http://mira.ucd.ie>), enables visualisation of the MiRA Cube, populated with the projects reviewed in this paper. The MiRA Cube is intended as an additional tool to help understand the differences between MiRA projects.

Figure 11: A screenshot of the MiRA Cube AR application.

The program was created using a modified version of the ARToolkit 2.69 VRML Viewer and assembled as an installer using “Inno Setup 5”. As ARToolkit is GPL-based, the source code for the MiRA Cube `simpleVRML.c` has been included in the setup. Using a web cam and a printed MiRA Marker, a user can interact with a live 3D representation of the MiRA Cube. Each project has an image displayed on the bottom

left corner to help the user identify the project.

A.1. Download

The MiRA Cube set-up program can be downloaded from <http://mira.ucd.ie/downloads/MiRACubesetup.exe>. The marker needed to view the cube is provided at the end of this paper or can be printed from [pattMiRA.jpg](#), which is located in the root directory of the MiRA Cube program. Alternatively, it can be downloaded from <http://mira.ucd.ie/downloads/pattMiRA.jpg>.

B. Instructions

The program allows a user to focus on an individual project by using the ‘left’ and ‘right’ cursor keys. Using the ‘up’ and ‘down’ cursor keys allows examination of different slices of the cube. For instance, by selecting only the top of the cube, a user can focus on projects that use strong agents.

The program also contains two different labelling aids which can be selected by using the ‘page up’ and ‘page down’ keys.

The background image can be turned off by pressing the ‘z’ key and made transparent by pressing the ‘o’ key. Also if a user wants to examine the cube without projects, they can press the ‘v’ key.

Please print this marker and visit <http://mira.ucd.ie> where you can download a program to view the location of all reviewed systems within the MiRA Taxonomy Cube in 3D.

The logo consists of the word "MIRA" in a bold, black, sans-serif font. The text is centered within a white square, which is itself centered within a larger black square.

MIRA