Knowledge-Based Systems 22 (2009) 540–544

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

Deploying embodied AI into virtual worlds

David J.H. Burden *

Daden Limited, 103 Oxford Rd., Moseley, Birmingham, B13 9SG, United Kingdom

ARTICLE INFO

Article history:
Available online 9 January 2009

Keywords:
Artificial intelligence
Virtual worlds
Second life
Turing test
Avatar

ABSTRACT

Unlike computer games where Non-Player-Character avatars are common, in most virtual worlds they are the exception. Deploying an embodied AI into a virtual world offers a unique opportunity to evaluate embodied AIs, and to develop them within an environment where human and computer are on almost equal terms. This paper presents an architecture being used for the deployment of chatbot driven avatars within the Second Life virtual world, looks at the challenges of deploying an AI within such a virtual world, the possible implications for the Turing Test, and identifies research directions for the future.

1. Introduction

In the half-century since the Turing Imitation Game [35] was created technology has changed dramatically. The paper, card and text interface was replaced by the interactive terminal, and then by the desktop PC with a windows, icon, mouse and pointer (WIMP) interface. Technologies such as Flash and text-to-speech have enabled us to create avatars on the desktop or on web pages – human like characters which can attempt to converse with users when controlled by a chatbot [9] programme. The most well known consumer chatbot is probably Anna on the Ikea site [17]. However chatbots have never really caught on, possibly partly due the immaturity of the chatbot engines, but also due to the way that the conversation model breaks the WIMP metaphor of the rest of the interface.

The last couple of years though have seen the emergence of a new interaction model – the virtual world. Here, the computer creates a complete 3D environment, and the user, represented by their own avatar, can move around the 3D space, meet and interact with avatars controlled by other users, and change and build new environments and new devices. Linden Lab's Second Life [29] is probably the best current example of an open virtual world – one in which users have almost as much freedom of action as they do in the real world. Second Life grew from 100,000 registrations in April 2006 to over 13m registrations by April 2008. More significantly organisations ranging from the BBC and IBM to Save The Children Fund and the British Council have started using virtual worlds both within and external to their organisations.

2. Non-Player Characters

Virtual worlds themselves partly grew out of Multi-User-Dungeons (MUDs), Massively Multiplayer On-Line Role-Playing Games (MMORPGS) and computer (and even paper) role-playing games. In all of these the "Non-Player Character" (NPC) has always been present. The NPC is a character (or avatar) controlled by the game (computer) and which is used either to impart knowledge or things, act as friend or foe, or just provide local colour. Whilst their scope to have an open ended conversation has been limited (and usually not even present), the fact is they were designed to blend in with the human avatars and the computer generated environment.

There has also been significant work (e.g. Hubal [16] on using NPCs as tutors within eLearning systems). Their aim is typically to provide help, motivation and feedback to the student during a learning scenario. However as with web based chatbots the model is typically one student-one tutor, the tutor is explicitly an NPC, and the NPC has privileged access to information.

In virtual worlds such as Second Life the NPC has been more or less completely absent. Partly this was due to an initial resistance to such characters from Linden Lab (this was after all meant to be a shared virtual world where human to human interaction and socialisation were paramount), and partly since technical limitations (lack of avatar control, lack of web interface or powerful programming language) made it hard to create even a basic NPC.

Now, however, these technical limitations have gone, and Linden Lab is taking a more open view. It is also worth noting that some competitor virtual world platforms (such as Active Worlds [15] and Kaneva [18]) offer NPCs "out the box". The result is that we are now able to fully experiment with chatbot controlled NPCs

* Tel.: +44 121 698 8520; fax: +44 121 698 8600.
E-mail address: david.burden@daden.co.uk.

0950-7051/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.knosys.2008.10.001

within a virtual world. The roles that such an NPC can play within a virtual world range from virtual receptionists, greeters and guides, to personal assistants, mentors and tutors, and ultimately perhaps as personal agents – controlling a user's avatar in their absence, or even death.

3. From chatbot to robotar

Chatbot development is reasonably well studied ever since the TIG was first proposed. ELIZA [12] was the first famous chatbot, and ALICE [36] was another milestone. The Loebner Prize [20] and The Chatterbox Challenge [33] are both annual competitions which have their roots in the TIG.

However, these are typically text-only experiments – although some limited visual components are often added. The focus is on whether through the text exchange alone we can replicate human "behaviour". However, with virtual worlds we have the ability to embody the chatbot. The new challenge is:

> "*Are we able to create an NPC within a virtual world which is indistinguishable in its complete behaviour from a player character/avatar*".

And if we can do so, will we have passed the Turing Test?

We try to keep away from using the term Artificial Intelligence since there appears to be no commonly agreed definition of what Artificial Intelligence [2] is. The term "AI" brings with it the grand visions of science fiction of powerful artificial intelligences, and more founded concepts such as the Singularity [34], and in academic parlance the term AI is now being replaced by Artificial General Intelligence [14] (no doubt partly for the two reasons above). The defining characteristics of AGIs appear to be around problem solving [24], learning [24] and coping within insufficient information [37].

We are under no illusion as to what we are trying to make or study – we simply aim to create computer programmes which *mimic* human behaviour. As such we prefer to refer to our creations as either robotars – an avatar in a virtual world which is controlled by a computer, rather than a person. We make no special claim to their "intelligence", they could be very simple, or ultimately very advanced.

The Turing Test (or a virtual world version) is however still a good milestone for chatbot development. Indeed, for perhaps the first time within the context of the Turing Test, virtual worlds place the human and the computer on an equal footing [10]. Both are operating at a level of abstraction beyond their normal environment. Both are operating as avatars in a virtual world (and until recently both were constrained to text chat). As such the computer is finally presented with a level playing field in which to take the Turing Test or a virtual equivalent (although such a window may be closing as voice gets introduced to virtual worlds).

A significant aspect of a virtual world such as Second Life (but interestingly not in a MMOPRG like World of Warcraft [38]) is that the working assumption that a human player has when they encounter another avatar is that the other avatar is also being controlled by a human. In a conventional Turing Test the challenge that a computer has is to prove that it is a human. Within a virtual world the challenge is subtly different – it must not give-away the fact that it isn't a human.

Whilst developing chatbots as a purely academic exercise has its attractions, we are a commercial organisation. This means that our developments are focused on creating chatbots with commercial uses. For instance we are already involved in deploying chatbots as:

- web-based guides to web sites (e.g. Yoma [39]),
- web-based virtual analysts and advisors (e.g. Atlas Intelligence's AIRA [3]),
- avatar based NPCs within virtual worlds such as Second Life as virtual receptionists, sales staff, and advisors.

To us the most immediate test of a chatbot's salience is the satisfaction of the customers using it. This may or may not have a correlation with a bot's Turing performance.

By limiting the scope and context of our robotars to such commercial applications we are to an extent making the task of creating a convincing AI easier. We certainly constrain the domain of knowledge and range of capabilities that the robotar should require. The users may also come with constrained expectations of what the avatar (human or robotar) might be able to do. However, in our experience users always have high expectations and expect the bot to be able to do many "common sense" things even if the bot is within a constrained role (e.g. our receptionist bot is often asked to do maths). An interesting exercise might be to try the same range of questions out on a real receptionist – or perhaps the fact that such questions are being asked shows that the bot has "failed" to mimic a human.

4. Example embodied chatbots

Since late 2007 we have deployed two robotars within the Second Life virtual world. Abi is our virtual receptionist. She staffs our office 24/7, greeting visitors and answering basic questions about Daden and what we do. She also knows her way around the office area. In fact Abi has two simultaneous existences, one as an embodied avatar in Second Life, and one as a head-and-shoulders Flash based avatar on our web site. Both share the same chatbot engine. Halo is our robotar test-bed. She has been running on the web for over 4 years, and now has a Second Life existence. We are building her her own "home" in Second Life, starting with a garden, and she is being given a far more open set of goals in terms of what she can do – ultimately to be driven by her own motivation model.

In the last 3 months Abi has had 1260 conversations, of which about 140 were in Second Life (the rest being on the web), and Halo has had 840 conversations, of which 32 were in Second Life.

We aim to do a formal assessment of the effectiveness of these bots as part of our research on emotions in robotars with the University of Wolverhampton, but we are already seeing positive coverage in Second Life related media [5].

5. Technical architecture

Within our Chatbot system we have taken a very pragmatic approach drawing on technologies such as:

- Artificial Intelligence Markup Language [8] (AIML), increasingly just for marshalling the conversation, rather than actually storing pattern/response pairs.
- Resource Descriptor Framework (RDF) [27] and Topic Map [32] approaches for storing information and memories, together with related ontologies.
- Web services, in order to access existing information on the web (such as Google, Wikipedia, Amazon, and RSS feeds etc.).
- Existing repositories of chatbot data such as WordNet [1], Alice Superbot [31] and OpenCyc [22].

Our chatbot engine into which we are incorporating these features is called Discourse [7].

Although we have experimented [6] with mimic based automated learning systems (such as that used by the George chatbot