A Framework for Building Intelligent Software Assistants for Virtual Worlds

Abraham Guerra
Pace University
awar88@gmail.com

Abstract

This research effort is focused on the use of artificial intelligence in virtual worlds. In particular, the use of artificial intelligence to facilitate the virtual world experience. There are several challenges facing users of virtual worlds and a premise of this research is that artificial intelligence can alleviate some of these challenges. The need for general-purpose bots is another premise of this research. This need is highlighted in this research in view of the fact that an intelligent assistant for virtual worlds needs to be able to handle a broader set of tasks. An intelligent assistant would not be of great value if it only provided assistance with wayfinding or avatar control. An intelligent assistant for virtual worlds needs to offer a wide variety of assistance. Thus, this research provides value for the usability of virtual worlds and general purpose agents.

1. Background

Virtual worlds are challenging for new users. Recognized challenges include "motivation, technical difficulties, controlling one's avatar, interacting with others, and finding compelling activities" [1]. Wayfinding is another recognized challenge of virtual worlds which can be categorized under one of the aforementioned challenges but merits mentioning on its own because "it still remains as one of the main limitations of using 3d worlds in internet" [2]. Wayfinding is the process through which a user navigates a virtual world.

Virtual worlds attempt to address many of these challenges. The most common approach to the wayfinding challenge is to provide maps of the virtual world environment. These maps will show points of interest and will offer teleporting to some of these points of interest. However, these maps may not contain the level of detail desired by users. Figure 1 shows one such map which highlights the lack of detail which may be desired by a user. For example, the map is rendering the Pace campus in Second Life but it has no representation of the cow farm which you can see on the campus (just above the map). This may be a point of interest missing from the map.

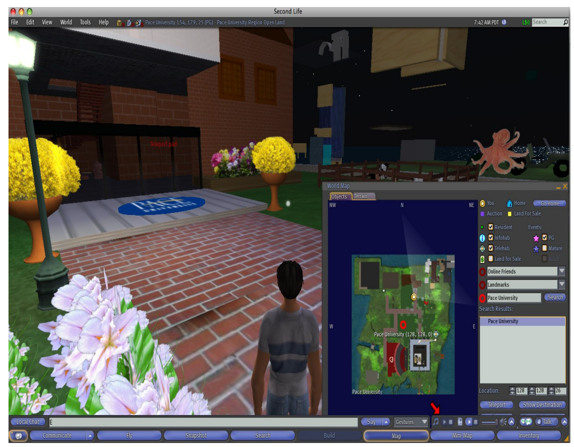

Figure 1. Virtual World Map

Another significant problem with the map-based approach to wayfinding is it may require the user to navigate to the point of interest on their own. Particularly, when users may not be able to teleport to all points found on a map. A user will have to navigate their avatar to the point of interest under these circumstance. This introduces the challenge of reading and referencing the map.

This research helps address these challenges by enhancing existing wayfinding solutions through an intelligent assistant for virtuals worlds. The intelligent assistant is capable of exploring the virtual world and perceiving objects it encounters. The intelligent assistant builds a cognitive map of the objects it perceives which it can then use to compliment map-based approaches to the wayfinding problem. A user can use their intelligent assistant to find object which may be missing from the maps provided by the virtual world engine. Similar approaches have been explored

[2-5] but this approach is unique in that it combines perception, search-ability, and dynamic map-building into a highly interactive system. More importantly, it is unique in that wayfinding is only one of multiple intelligent behaviors contained within an intelligent assistant.

The learning curve associated with the use of virtual worlds is also a significant challenge. Users must learn to control camera angles, avatar appearance, avatar movement, and much more. Solutions to this problem range from creating educational areas in the virtual world (see figure 2) to having real-world users on standby to assist new users. Creating educational areas is an effective way for users to learn these virtual world skills but is susceptible to the challenges of wayfinding. Providing assistance through real-world users is also effective but is difficult to accomplish in both social and corporate virtual worlds because of availability and cost factors, respectively.

Figure 2. Virtual World Training Area

The intelligent assistant proposed in this research can help address the challenges of new user orientation. It can be pre-loaded with this sort of knowledge and provide a more efficient mechanism for disseminating this knowledge to users. More importantly, this highlights the intelligent assistant's ability to handle multiple tasks. The assistant will not only be able to handle wayfinding and new user orientation challenges, it can be extended to contain a multitude of functionalities which can be encapsulated in artificial intelligence modules.

2. Framework Architecture

There are five major components to the architecture of the proposed framework (see figure 3). One of the major components is the artificial intelligence engine (AI Engine in figure 3). It enables the encapsulation of the desired intelligence into a corresponding modules.

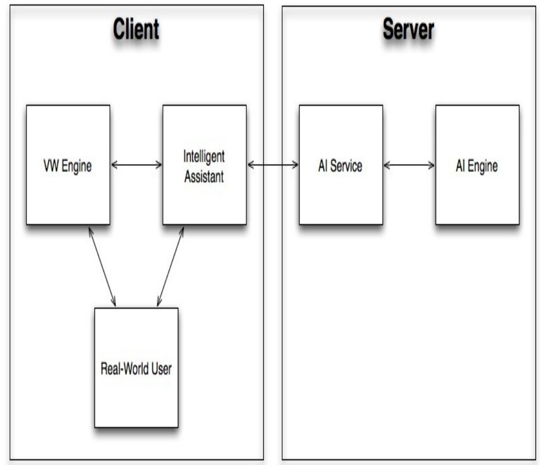

Figure 3. Intelligent Assistant Architecture

The modular approach enables the general-purpose nature of the intelligent assistant.

The artificial intelligence service component (AI Service in figure 3) serves as the nervous system of the intelligent assistant. It orchestrates the varying AI modules within the AI engine. This component offers great flexibility in that it can communicate with any client implementation.

The Intelligent Assistant component provides the interface to the AI components on the server side of the architecture. The user may be a real-world person who is interacting with the intelligent assistant through the virtual world viewer. The user may also be the virtual world engine. This component, in combination with the AI Service, is the key to the flexibility of this framework. The Intelligent Assistant is not limited to any specific virtual world and can thus become a universal intelligent assistant.

The Virtual World Engine component (VW Engine in figure 3) represents the server hosting the virtual world. A virtual world engine may be hosting multiple worlds which, in turn, may be hosting multiple regions. The Intelligent Assistant serves as an abstraction layer which allows for the integration of divergent virtual world engines.

The Real-World User component represents the UI through which a real-world user interacts with the VW Engine and/or the Intelligent Assistant. The Intelligent Assistant component again serves as an abstraction layer which allows for the integration of divergent UI's. It also allows for direct communication between the real-world user and the Intelligent Assistant.

3. Prototype

The ideas behind this research have been validated through a prototype implementation. The implementation is based on Soar and Java on the server side. It is based on OpenSimulator and C# on the client side. A diagram of the prototype is provided below.

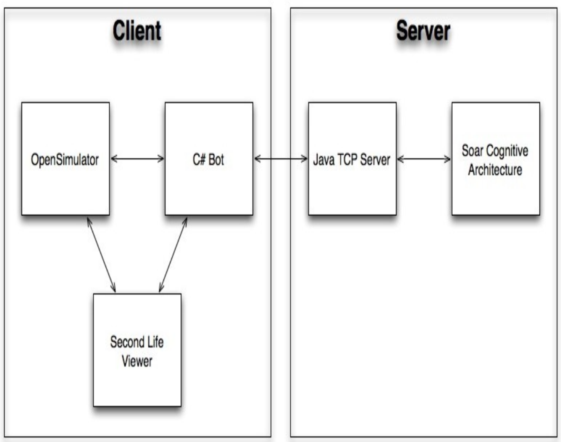

Figure 4. Prototype Implementation

The prototype implementation of the AI Engine component is based on Soar. Soar is a cognitive architecture which simulates cognition [6]. As such, it simulates short-term memory, long-term memory, and learning. It also uses a goal-centric approach to accomplishing tasks. For these reasons, Soar is a great foundation on which to build artificially intelligent agents.

The AI Service component is implemented as a Java-based TCP server. This server manages a TCP connection on which it listens for requests from AI clients and processes them through the AI engine. The server is built with the flexibility to handle multiple AI engines. It will grab requests from the TCP's input stream which will specify AI engines, intelligence modules, and functions within the intelligence modules.

The Intelligent Assistant component is implemented as a C# bot. However, the framework is flexible enough to handle different implementations of this component. This research chose a bot implementation because of the high degree of interactivity offered through a bot implementation. The C# bot, which we call Bert-O, logs into a virtual world and then requests its artificial intelligence through the Java AI Service provider.

The prototype implementation of the VW Engine component is based on OpenSimulator (http://opensimulator.org/wiki/Main_Page). OpenSimulator is an open source server used to run virtual worlds. Aside from being open source, OpenSimulator offers several advantages. OpenSimulator is free, has an active development community, and has rich documentation It is also client compatible with Second Life which helps ensure our prototype has the broadest coverage.

The Real-World User component is implemented through the Second Life Viewer (http://secondlife.com/support/downloads/). The Second Life Viewer is a client browser through which users connect to virtual worlds. The Second Life Viewer was chosen because of its support of OpenSimulator and its popularity (due to Second Life popularity).

3.1. Artificial Intelligence Modules

The Intelligent Assistant presented in this research is enabled through AI modules. The AI modules are loaded into the AI Engine through the AI Service. The prototype implementation of this research uses three AI modules: an orientation module, a sales module, and a wayfinding module.

The orientation module is a valuable task with which the intelligent assistant can help. It is very valuable because of the availability and cost constraints associated with providing this sort of assistance through a real-world user. Particularly in a business environment as expressed in interviews with senior technical staff involved in virtual world development at IBM [7],[8]. The orientation module is used to encapsulate the sort of knowledge disseminated through the virtual training area mentioned above. Providing the knowledge through an intelligent assistant alleviates the challenges mentioned above and in the Introduction of this paper. Figure 5 shows such an interaction.

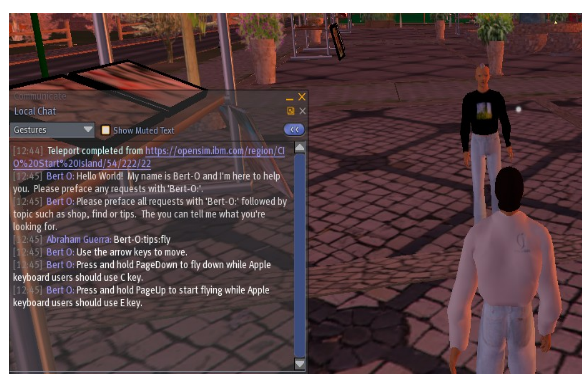

Figure 5. Intelligent Assistant - Orientation Module

The sales module highlights a second task with which the intelligent assistant can help and another virtual world challenge it can help alleviate. It helps alleviate the challenge virtual world users face in finding compelling activities. Similar to the orientation module, it also helps with the availability and cost constraints mentioned above. The sales module contains knowledge regarding stores in the virtual world and can help users with a purchase request. Figure 6 shows such an interaction.

Figure 6. Intelligent Assistant - Sales Module

The wayfinding module highlights the breadth of functionality and intelligence which can be provided through the intelligent assistant. The wayfinding module enables a complete wayfinding solution. The intelligent assistant wanders through the virtual world perceiving objects in the virtual world. It adds the perceived objects to its cognitive map. It can then reference its cognitive map to help users find objects. Figures 7a and 7b show such an interaction.

Figure 7a. Intelligent Assistant - Wayfinding Module

Figure 7b. Intelligent Assistant - Wayfinding Module

4. Future Work

There is much room for the evolution of this research. One area which lends itself to future work is natural language processing. The prototype is built using a strict communication protocol. Integration of a natural language processing component would greatly enhance the interactivity of this prototype.

Another area of future work is integration with a heads-up display (HUD). Initial feedback on the prototype is to implement the intelligent assistant as a HUD rather than a bot [9],[10]. A HUD implementation may prove to be a more effective manner of providing assistance and may be a better mechanism for potentially improving the research through Swarm Intelligence.

A final area of future work worth mentioning at this time is integrating learning components into the orientation and sales modules. Currently, these modules are pre-loaded with information to handle frequently asked questions by new users. This approach is adequate but could be greatly enhanced by adding a learning component so it can learn answers to questions which may not have been known in advance.

5. References

[1] K. Bessière, J.B. Ellis, and W.A. Kellogg, “Acquiring a professional "second life": problems and prospects for the use of virtual worlds in business,” *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, Boston, MA, USA: ACM, 2009, pp. 2883-2898.

[2] M.J. Abásolo and J.M. Della, “Magallanes: 3D navigation for everybody,” *Proceedings of the 5th international conference on Computer graphics and interactive techniques in Australia and Southeast Asia*, Perth, Australia: ACM, 2007, pp. 135-142.

[3] L. Chittaro and S. Burigat, “3D location-pointing as a navigation aid in Virtual Environments,” *Proceedings of the working conference on Advanced visual interfaces - AVI '04*, Gallipoli, Italy: 2004, p. 267.

[4] J.S. Yoon and M.L. Maher, “A swarm algorithm for wayfinding in dynamic virtual worlds,” *Proceedings of the ACM symposium on Virtual reality software and technology - VRST '05*, Monterey, CA, USA: 2005, p. 113.

[5] J. Eno, S. Gauch, and C. Thompson, “Intelligent Crawling in Virtual Worlds,” *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 03*, IEEE Computer Society, 2009, pp. 555-558.

[6] J.E. Laird, A. Newell, and P.S. Rosenbloom, “SOAR: an architecture for general intelligence,” *Artif. Intell.*, vol. 33, 1987, pp. 1-64.

[7] C. Becker, “Value of Artificial Intelligence in Virtual Worlds,” Mar. 2010.

[8] N. Katz, “Value of Artificial Intelligence in Virtual Worlds,” Mar. 2010.

[9] M. Ackerbauer, “Prototype Demonstration,” May. 2010.

[10] R. Alther, “Prototype Demonstration,” May. 2010.