Design and Implementation of a Virtual Salesclerk

Christopher Mumme, Niels Pinkwart, Frank Loll

Clausthal University of Technology
Clausthal-Zellerfeld, Germany

Abstract. This paper describes the design and implementation of a virtual agent that is capable of providing customers in a 3D online shop with advice. Based on a product knowledge base, a conversation model and a model of the shop, the agent communicates with the customer through text based dialogues and leads the customer through the virtual world using gestures.

Keywords: Virtual Agents, 3D Environments, Online Shopping

1 Introduction

Virtual Environments (VEs) are used in a variety of research fields including collaboration systems, marketing [1], medicine [2], computer graphics and psychology [3]. In contrast to traditional 2D representations used in websites or chats, VEs make possible a much higher level of social interaction due to several factors, including the identification of the user with the avatar embodiment, the possibility of expressing emotions and gestures, and the overall representation of the virtual world which may resemble reality much better than is possible in 2D [3]. These factors can also open up new possibilities for eBusiness applications: using 3D representations, products can be presented more realistically and can be arranged next to each other as is done in real shops, making use spatial contexts [1]. For instance, in 2006 the company Adidas made use of this way of presentation by displaying their new sports collection in a virtual shop on Second Life. Also numerous other companies like IBM, ABC.com, Toyota and Deutsche Telekom try to reach additional target groups through this new medium. In addition to saving costs for physical warehouses, virtual shops can offer the customer personalized representations of assortments of goods [4]. For instance, it is possible to show a coach suite in different variants to a customer within the customers' virtual imitated living-room, or to present a shelf of "his preferred products" to a customer within a virtual store. Such a personalization of products is not feasible in real-world shops.

Yet, a large range of products makes it hard for customers to decide. Therefore, many customers like to ask qualified salesclerks for help. This may lead to an increased shop turnover because it simplifies the decisions for the customers. Traditional 2D online shops like Amazon.com work with recommender systems to achieve similar effects. While these can be seen as anonymous agents without a visual representation, there are other approaches which apply the concept of agent in form of chatbots, i.e. graphical representations of agents one can communicate with.

Examples for such agents on websites include the Coca Cola Company which replaced their FAQ by an agent, PayPal which guides their users through the payment process with an agent, or Bol.com which gives information about products via agents. Prominent advantages of using virtual salesclerk agents on websites are, among others, a simplification of navigation (users can ask the agent instead of browsing the website to find the searched information) as well as extended marketing potentials (the agent can actively seek contact with the customer to point his attention to specific products) [5, 6].

While using chatbots as sales agents works fine in 2D environments, a straightforward approach for developing virtual salesclerks for VEs by simply integrate existing chatbots into a virtual online shop is problematic. Typically, chatbots do not support actions like movements and non-verbal communication which are essential interaction techniques in VEs. To fully exploit the potential of 3D environments, more advanced agent models are required. A number of successful agents in differently targeted VEs have been developed. For instance, Kenny et al. [7] have described the use of virtual patients in the training of medical students. Other examples in the field of medicine are a virtual fitness trainer [8] and a virtual therapist [2]. Furthermore, 3D agents have been employed to escort visitors in real museums [9]. For the specific application area of online shops however, there are no research results yet – and sales agents in virtual shops are generally very rare also in environments like Second Life that are used by a number of companies.

Figure 1. Product specification, shopping basket and product demonstration

We developed a virtual video store based on a modified OpenSim server and a Second Life client as a testing framework for a virtual salesclerk agent. To process interaction data (e.g. movement, communication), we made use of the framework described in [10] which also allows for sending feedback directly into the VE. The virtual video shop contains more than 2500 movies and about 500 seasons of series. The shop is designed close to video shops in the real world in order to ensure that the customers behave as they would in real life [11]. The only notable difference is the size of the DVD covers, which is larger than in reality. The reason for this design choice was that in the real world, humans can recognize DVDs covers from some distance. To enable this also in the virtual store with the limited screen resolution, we increased the cover size. To see information about a movie, a customer can open an extra window in the client by clicking on a virtual DVD (Figure 1, right). After that the user puts the movie into his shopping basket which also opens in an extra window (Figure 1, left). The basket is also used to purchase movies.

2 Agent Design

In this section, we discuss the requirements for a virtual salesclerk agent and describe the design of our agent in terms of embodiment, movement, language processing and provision of shopping advice.

2.1 Embodiment & Movement

Embodiment deals with the representation of artificial figures such as avatars and robots, including their appearance and their behaviour. Key research results to consider are that humans transfer social behaviours of real-life situations into virtual environments [11]. The degree of this transfer depends on the level of immersion, which is determined by the design of the virtual world [12]. Often, the question whether users accept a virtual environment is hard to predict and requires empirical studies. As a rule of thumb, a close-to-reality representation (such as offered by a Second Life based environment) is likely to reach a high acceptance [12, 13]. Another factor connected to embodiment is that humans using a virtual world would expect virtual agents to make use of the communication options that this framework offers. This includes moving around and exploiting the available space (in our case, for showing products) and making use of gestures. Figure 1 illustrates how our sales agent points to a DVD, recommending it. The availability of gestures and the sufficiently realistic visual representation were reasons to use OpenSim/SecondLife as a technical base for our implementations.

Concerning movements, our agent is able to go to and point at a DVD in the shop for recommending it. To calculate the shortest paths while avoiding obstacles, we use the Potentialfield/Wavefront algorithm [14]. In addition, in order to provide customers with advice, the agent also needs to select the customers he wants to counsel. It does so as follows (see Figure 2): First of all, the agent addresses customers that are in a logical waiting queue to which customers are added to when they want to get served while the agent already counsels another customer. In this case, the agent asks these clients to wait and then comes back to them. If the queue is empty, the agent updates a list of all users that are in the store. He then looks for the nearest costumer, provided that he hasn't served him within the last three minutes. That way, it is ensured that the agent does not address the same customer consecutively several times. If the agent finds a customer, he offers his assistance. If there are no clients in the shop, the agent moves to a sales counter.

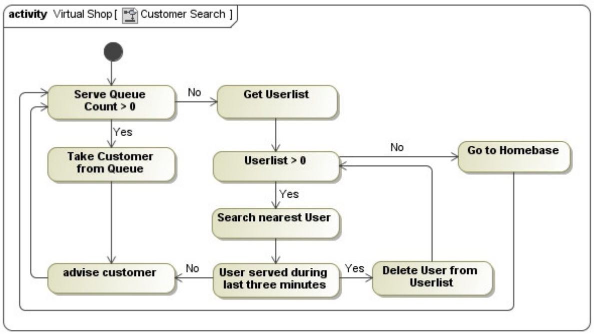

Figure 2. Model for choosing customers

2.2 Language Processing & Provision of Shopping Advice

Evaluating user inputs is the most challenging task of a chatbot engine. Generally, there are two types of talks between a salesclerk and a customer: either the customer looks for a specific product or the customer has general ideas about the type of product he is looking for, but is, however, still undecided. While the former case requires only a comparison with the underlying knowledge base to retrieve the required information, the latter is more difficult, since the agent has to talk actively with the user to collect further information about his needs [6].

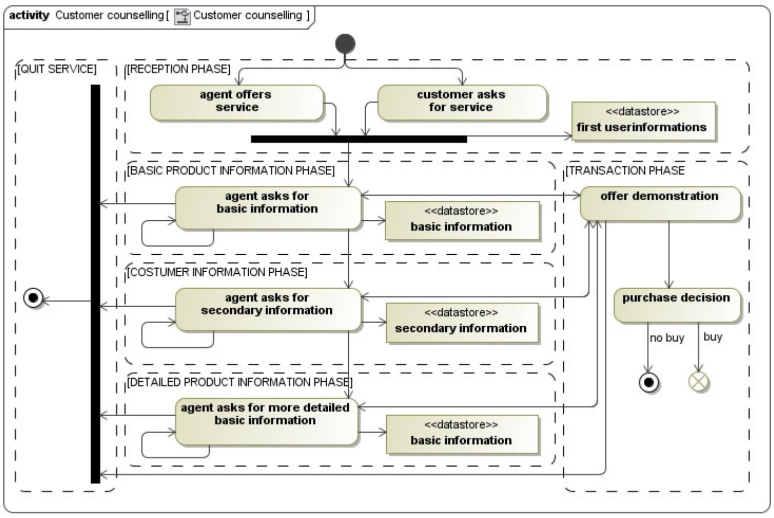

Figure 3. Model of customer counselling

Thus, we developed a sales dialogue model (see Figure 3) consisting of five phases: (1) *Reception*: The welcome phase, initiated either by the customer or the agent. (2) *Basic Product Information*: The agent collects information about the product the customer is looking for and compares it to his database. If there are hits in the database, the agent offers the customer to lead him into an area of the shop or directly to the product, pointing on it before turning back to the customer. (3) *Customer Information:* After having asked specific questions about the product, the agent tries to collect personal data like age or sex about the customer. By comparing this information to statistical data, the agent can infer potentially good recommendations. For instance, kids will not be recommended horror movies, gender-specific typical genre preferences (for instance, men prefer action movies, while women prefer love films) will be considered, and more recent films will be recommended to younger people. (4) *Detailed Product Information*: In this phase, the agent repeatedly asks product specific questions in order to gain further insights into the preferences of the customer. (5) *Transaction*: Finally, the transaction phase introduces the end of the conversation. The agent invites the customer to look around, assuming that they are now at a place in the shop where the desired film is located and thus – due to the shop layout – also other potentially interesting films are. These phases are exemplified in the example interaction sequence in Table 2.

Table 2. Example dialogue between agent and user (including movements and gestures)

1 : [**Reception phase**]	17: Customer: 20 years old.
2: Agent: Hello, can I help you?	18: [**Detailed product information phase**]
3: Customer: Yes.	19: Agent: Can you tell me another genre which you prefer?
4: [**Basic product information phase**]	20: Customer: Yes
5: Agent: What kind of genre do you prefer?	21: Agent: So please tell me.
6: Customer: I´m looking for some action movies.	22: Customer: **Thriler**
7: Agent: Can I show you the action section?	23: Agent: Can I offer you the movie "War of the Worlds"?
8: Customer: Yes, gladly.	24: Customer: Yes
9: Agent: Ok, please follow me.	25: Agent: Ok, please follow me.
10: *… Agent walks, customer follows*	26: *… Agent walks, customer follows.*
11: Agent: Here we are, please take a look around.	27: [**Transaction phase**]
12: *… Agent waits near customer. After 30s:*	28: Agent: Here we are *(pointing at a DVD)*, please take a look around.
13: Agent: Can I still help you?	29: Agent: Can I still help you?
14: Customer: Yes.	30: Customer: No, I think I´ll take this one, thank you!
15: [**Customer information phase**]	31: Agent: You are welcome.
16: Agent: Can you tell me how old the person is for whom you want to buy the movie?	

In each phase of the dialogue, it is possible for the customer to leave the service also without communicating this directly to the agent. Identifying this case when both the agent and the customer move around the shop is not trivial. In our implementation, we make use of timeouts (see Table 3). If a customer does not answer an agent's question within a certain time interval, the agent assumes that the customer does not wish a further consultation. An alternative approach would have been to check if the customer moves away from the agent and thus to use the growing distance between them as indicator whether the counseling should be continued. However, it is imaginable that a customer walks away, wishing the agent to follow. In this case it would be wrong to finish the service.

As a matter of fact, spelling mistakes occur in chats. Thus, the agent's chat engine must be able to deal with it. Therefore, we used the algorithm of Damerau-Levenshtein which delivers a string metric between two words. Depending on word length and distance, the algorithm can determine whether a spelling mistake is present (as in Table 2, line 22). Pilot tests with different parameters have shown that distance values of 0 for word lengths of 0-4 characters, a distance of 1 for 5-8 character words, and a distance of 2 for longer words worked best.

Table 3. Timeouts for quitting counselling

Timeout	Handling
Agent follows customer for more than 30 seconds without customer talking to him	Finish counselling
Agent waits for an answer for more than 30 seconds	Repeat question
Agent waits for an answer for more than 60 seconds	Finish counselling
Agent arrives at a product while counselling and waits for more than 30 seconds for the arrival of the customer	Finish counselling

In the different phases of the conversation, the agent has to give product recommendations to the customer on the basis of the collected information. Generally, recommendation algorithms can be classified into two groups: personalized techniques give advice based on a user profile of the customer (e.g., about his shopping behaviour), while non-personalized techniques give the same advice to any customer (e.g., based on sales volume) [15]. Our agent integrates personalized techniques (cf. "Customer Information phase") with a form of non-personalized recommendations that Schafer et al. [16] called attribute-based recommendations. In our implementation we compare the user input and the user information with the data that is available in the product specifications. The film with the best match is then recommended. For the matching algorithm, in order to consider the temporal structure of sales dialogues where statements made by clients recently are of higher priority than those made half an hour ago, the information collected by the agent during the counseling is evaluated using weights. New information receives a very high weight, while older data loses its weight increasingly (see Figure 4).

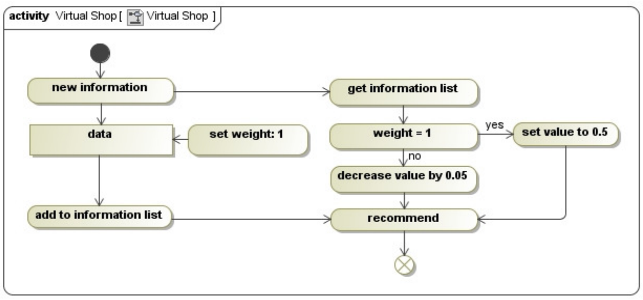

Figure 4. Weighing information over time

3 Conclusion & Outlook

Shops in virtual worlds have a large potential for eBusiness applications. They enable a detailed representation and arrangement of products and a rich interaction in the environment. In this paper, we discussed some differences between classical 2D online shops and modern 3D virtual environments and the design challenges this imposes for implementing sales agents in virtual worlds.

The virtual salesclerk agent for a virtual movie store, presented in this paper, is able to provide customers with advice based on a product and a shop model, information about the user and a process model for the counseling activity. Interacting with the customer, the agent combines text based dialogues with movements and gestures in the virtual world. To evaluate the success of our approach, we conducted a study with 36 participants. While details about this study are not in the focus of this paper, the results showed that the participants generally followed the advice given by the agent and that they liked the interaction with our agent. Many users stated that they can imagine shopping in a virtual shop, receiving advice by a virtual shopping agent. Based on these positive results, we plan to further refine and evaluate the virtual shop and the agent in future studies.

References

1. Herder, J., Jaensch, K., Horst, B., Novotny, T.: Testmärkte in einer Virtuellen Umgebung. Workshop Augmented & Virtual Reality in der Produktentstehung, pp. 97-110, (2004).
2. Pontier, M., Siddiqui, G. F.: A Virtual Therapist That Responds Empathically to Your Answers. In: Proceedings of IVA 2008, pp. 409-416, Tokyo, Japan.
3. Benford, S., Bowers, J., Fahlen, L. E., Greenhalgh C.: Managing Mutual Awareness in Collaborative Virtual Environments. In Proceedings of VRST 1994, ACM Press.
4. Rogoll, T., Piller, F.: Product configuration from the customer's perspective: A comparison of configuration systems in the apparel industry. International Conference on Economic, Technical and Organisational aspects of Product Configuration Systems, Denmark (2004).
5. Mahbubur Rahman, S., Bignall, R. J.: Internet commerce and software agents, IGI Publishing (2001).
6. Braun, A.: Chatbots in der Kundenkommunikation. Springer, Heidelberg (2003).
7. Kenny, P., Parsons, T. D., Gratch, J., Rizzo, A. A.: Evaluation of Justina: A Virtual Patient with PTSD. In: Proceedings of IVA 2008, pp. 394-408, Tokyo, Japan.
8. Ruttkay, Z., Welbergen, H.: Elbows Higher! Performing, Observing and Correcting Exercises by a Virtual Trainer. Proceedings of IVA 2008, 409-416, Tokyo, Japan.
9. Satoh, I.: Context-Aware Agents to Guide Visitors in Museums. Proceedings of IVA 2008, pp. 409-416, Tokyo, Japan.
10. Mumme, C., Olivier, H., Pinkwart, N. (2008). A Framework for Interaction Analysis and Feedback in Collaborative Virtual Worlds. *Proceedings of the 14th International Conference on Concurrent Enterprising,* pp. 143-150.
11. Friedman, D., Steed, A., Slater, M.: Spatial Social Behaviour in Second Life. Proceedings of IVA 2007, pp. 252-263, Paris, France.
12. Schneider E.: Mapping out the Uncanny Valley: a multidisciplinary approach. Proceedings of SIGGRAPH 2008, New York, USA.
13. Benford, S., Bowers, J., Fahlén, L. E., Greenhalgh, C., Snowdon, D.: User embodiment in collaborative virtual environments, In: Proc. of SIGCHI 1995, pp. 242-249, Denver, USA.
14. Barraquand, J.; Langlois, B.; Latombe, J.-C.: Numerical potential field techniques for robot path planning. IEEE Trans. on Systems, Man and Cybernetics, 22(2), pp. 224-241.
15. Runte, M.: Personalisierung im Internet - Individualisierte Angebote mit Collaborative Filtering. Dissertation, Christian-Albrechts-Universität zu Kiel (2000).
16. Schafer, J. B., Konstan, J., Riedl, J.: Recommender Systems in E-Commerce. Proceedings of ACM-EC 1999, Denver, USA.