

This is a reprint from a paper published in the Proceedings of the IADIS International Conferences
IADIS, <http://www.iadis.org>

IMPLICIT TRAINING OF VIRTUAL SHOPPING ASSISTANTS IN 3D ELECTRONIC INSTITUTIONS

Anton Bogdanovych, Simeon Simoff
*Faculty of Information Technology
University of Technology, Sydney, NSW, Australia*

Carles Sierra
*Artificial Intelligence Research Institute (IIIA-CSIC)
Barcelona, Catalonia, Spain*

Helmut Berger
*Electronic Commerce Competence Center – EC3,
Donau-City-Strasse 1, A-1220 Wien, Austria,*

ABSTRACT

The growing demand for shopping assistants in E-Commerce was identified by many researchers. Some retailers try to address this need by introducing totally autonomous agents; others make extensive use of human resources, shifting operators from telephone lines to chat-based interactions with online customers. The recently developed 3D Electronic Institutions methodology provides facilities to conveniently combine these two approaches. Initially, an autonomous agent tries to deal with customer's requests. When the limitations of its intelligence are reached a human operator takes over and satisfies the "out-of-scope" inquiry. At the same time, the agent observes the human operator and learns how to handle similar inquiries in the future. We argue that this learning aspect can be realized by means of 3D Electronic Institutions and believe that this "agent training" will be feasible and far more successful than it is possible in nowadays form-based E-Commerce solutions.

KEYWORDS

Sales Assistants, 3D Electronic Institutions, Learning.

1. INTRODUCTION

Millions of people around the world spend an average of 20 hours per week in Virtual Worlds based computer games. Usually those users have to pay a monthly fee of 10\$ just for being able to participate (Hunter & Lastowka, 2003). Apparently, outside of pornography and eBay, 3D Virtual Worlds are one of the few online businesses that are making money on the Web. Norrath, the online world created by Sony, has more residents than Miami and higher per capita GNP than Bulgaria. In Virtual Worlds people are not only paying real money for the experience of being there, but also trade virtual goods. The most expensive virtual item (a virtual island) was sold for 26500 USD in the Project Entropia multi-user game (Hunter & Lastowka, 2003). If selling virtual items in Virtual Worlds is that profitable, why not selling real items there?

The use of 3D technology for selling real items was a hot topic in late nineties. Many people tried to enhance their sites with 3D models to offer customers better product presentation facilities. Unfortunately, the majority of the investments in 3D shopping didn't pay off. The major reason for the lack of success of 3D technology in E-Commerce is coupled with issues like costs and performance (Hurst, 2000). Nevertheless, recent developments prove that in the near future 3D applications may be faster and cheaper to create than quality photographs (Früh et al, 2005). The broad availability of fast Internet access also supports this trend.

Proving the usefulness of 3D visualization Daugherty et al. (2002) claim that a virtual experience (3D product presentation) has the potential to be richer than both direct (manipulation with a real product) and indirect (2D form-based product presentation) experience because it can be simulated, framed, annotated and contextualized. However, the conducted study showed that when tactile affordances are the most relevant for the product (e.g. touching the fabric), a virtual experience may have the same effect as indirect experience.

The possibility of scanning physical interaction behaviors of 3D objects proposed by Pai et al. (2001) draws even more optimistic picture. Researchers have managed to produce a fully automatic device for measuring the characteristics of different materials and attach this information to the 3D models of the scanned objects. These characteristics include deformation response, contact textures and contact sounds. Using this technology in combination with haptic devices that allow users to feel virtual objects (Whitton, 2003) may have a consequence that the feeling of touching the fabric becomes transmittable through the web.

In addition to product presentation 3D Virtual Worlds can provide several other benefits to E-Commerce environments. Social interactions play an important role in real world commerce and will definitely be important factors in the future of E-Commerce (Preece & Maloney-Krichmar, 2003). Satisfaction of impulsive decisions, similarity of the user interface with the real world, more ways to convey information, collaborative shopping possibility are just a few more examples (Bogdanovych et al, 2004) of potential benefits of the technology. Moreover, 3D E-Commerce environments built as 3D Electronic Institutions (Bogdanovych et al, 2005) provide facilities for implicit training to improve the believability of shopping assistants. Under "implicit training" we understand training of the autonomous agents by observing humans to help those agents to act like humans, without any explicit training efforts from the humans.

Many researchers, whose work is focused on virtual characters (and shopping assistants in particular), face the problem of making these characters believable. The believability has a lot of different characteristics, e.g. personality, social role awareness etc. New aspects of believability are constantly discovered and introduced to the research community (e.g. Prendinger & Ishizuka, 2001; Magnenat-Thalmann, 2005). The implementation complexity grows but passing the Turing test is still far from being possible.

In our work, instead of trying to discover and implement different characteristics of believability, we focus on the simulation theory. The main hypothesis of this theory can be best summarized by the cliché "to know a man is to walk a mile in his shoes". We strongly believe that simulation and imitation are the key technologies for achieving believability. To increase the believability of shopping assistants we propose that autonomous agents observe the behavior of human operators in E-Commerce environments and imitate it.

The better observation means the agents have the better can they imitate their principals. An efficient way to achieve full observation is to have a human fully immersed into an E-Commerce environment based on our 3D Electronic Institution technology. This technology provides context, background knowledge and similar embodiment for all participants, including humans and the autonomous agents who imitate the humans. This information valuably facilitates the training of the autonomous agents.

The remainder of the paper is structured as follows. Section 2 presents the state of the art for virtual shopping assistants in E-Commerce. Section 3 describes the 3D Electronic Institutions technology. In Section 4 we present the extension of the 3D Electronic Institutions technology to be used for implicit training. Concluding remarks and details of the future work are given in Section 5.

2. STATE OF THE ART

"Shopping assistant" is one of the most persuasive sales tools in traditional commerce. Shopping assistants offer help in a store, provide additional information on products and simplify decision making process helping to find a good that satisfies customer's requirements and various constrains. One of the major drawbacks that E-Commerce is facing today is the lack of such sales clerks. There is strong evidence that in brick and mortar stores customers find interaction with a sales person very beneficial. People value and are willing to pay for the reduction of perceived risk, the optimal configuration of the transaction for their specific usage context, and the enhancement of the in-use experience, which shopping assistants can provide them with. Therefore, shopping assistants are able to cause dramatic increase in sales. A study conducted by "Celio" showed that contact with a shopping assistant resulted in 18% more purchases (Chowdhury, 2004).

The demand for customer service and assistance in E-Commerce is going to be growing. Modahl (2000) suggests that we are now in the ten year transition in the way consumers shop and save. The majority of

internet purchases are still done by technology optimists, which are ready to deal with the lack of support for the sake of new and convenient technology, while mainstream customers are just starting to shop online. When the majority of mainstream shoppers arrive, either sites are going to have to offer more customer service and sales support on their sites, or they are going to see skyrocketing customer service costs.

If there is a problem to be clarified, online shoppers of today have to use expensive telephone service or accept long waiting times from e-mail communication. In case of the telephone service, either consumers who have to pay high per minute rates or the vendors that provide free customer service numbers have to deal with significant telephone service costs. Moreover, the majority of customers don't feel enthusiastic about switching from Web browsing to the telephone service. The situation with e-mail as the way of customer support is not much better. Many e-mail inquires take multiple correspondences to be satisfied, e-mail isn't real time while customers expect the answers back right away, quality assurance is difficult. According to Gutzman (2000) the only customer service option that seems to fit the speed of the Internet is real-time chat.

Commercial success of the chat technology was documented by Earthlink (2005). Tracking customer actions at Earthlink's site for ISP services showed that 70% of the customers were leaving without purchasing even when they saw broadband was available in their zip code for a "best offer" price. Following this observation the shopping team decided to integrate a chat feature into the website to provide confused customers with real-time online support, and conducted a case study for better understanding of the attitude of the shoppers towards chat service. The results showed that 15% of people who chatted with sales assistants converted and 80% of consumers gave the chat a good or excellent rating. Moreover, 61% of chat users preferred chat to other options of customer support.

One of the great advantages of chat from the site owners' point of view is that it is also much less expensive than the telephone service. A phone call costs in average around \$6.00, but the average chat is \$2.40 (Maguire, 2002). But even further cost reduction can potentially be achieved with the help of automated chat tools where customers' questions are answered by a computer program (autonomous agent) rather than a human being. These automations combine natural language processing algorithms with complex databases to give consumers answers to their queries and can even connect a customer with a human if the autonomous agent is unable to answer a question after several attempts (Prince, 2005). The drawback of this approach is that the intelligence of such autonomous agents is usually very poor. In order for them to be able to act believable and provide customers with reasonable feedback without switching to human-operators straight away enormous resources have to be used for explicit training. This training usually means that human operators have to type in huge list of possible customers' questions and for each of the questions provide an enormous variety of answers, which is very inefficient and resource demanding.

Surprisingly enough, none of the present solutions consider implicit training. The most important characteristic of implicit training is that human-operators train software agents by doing their routine job (answering customers' questions) in contrast to doing training as a routine job. We propose a scenario where initially only human operators answer customers' requests while corresponding autonomous agents observe their behavior and learn from them. Eventually, after testing and validation of the behaviors learned by the agents, they will start communicating with the customers without human help. In case a customer puts an inquiry that an agent is not trained to deal with, the agent requests help from an available human operator. The human replaces the software agent (with or without explicitly notifying the customer about this) and answers the non-standard inquiry. Unlike explicit training, the autonomous agent meanwhile observes the human operator and learns how to deal with similar enquires in the future (extends his intelligence). We believe that such human-computer cooperative approach offers the most feasible way to achieve reasonable efficiency in providing shoppers with trustful expertise.

The idea of training of autonomous agent by humans is not new, but is not as popular within the community as the other types of learning. One of the early examples of this approach is "programming by demonstration" (Bauer et al, 2000), where agents are trained by humans to recognize relevant text on a web page. Similar technique is also used in robotics, where humans (or other robots) are used to demonstrate some actions which a robot tries to imitate (Ekvall & Kragic 2005). In (Aleotti et al, 2003) robots acquire their behaviors from the provided demonstration of how to solve a certain task, given the necessary initial knowledge to the system. The trainee (robot) then automatically interprets what is to be done from the observed task, thus eliminating the need for tedious textual or iconic programming. Alissandrakis et al, 2001 proposed a successful approach for training robots by humans despite dissimilarities in the embodiments.

There are also some achievements in the area of training chatter bots using natural language (Rule, 2003). Transplantable Artificial Neurological Unit (TANU) chatter bots demonstrate the ability to participate in a

conversation with a human in a very reasonable way, after some training by a human is conducted. Authors claim that an average person only goes through about 70,000 important states in a 5 year span. So creating 70,000 states properly interconnected with transitions will make a very smart chatbot. The TANU network runs language aware transitions, so a transition that supports an event “Can you teach?” will be associated automatically with a transition from the source state when it receives “Can you educate?”, “Can you tutor?”, “Can you lecture?”, “Can you instruct?” and “Can you edify?”. Unfortunately, adding new states and transitions to the network doesn’t happen automatically but have to be done manually by human operators.

As we see, the major drawback of all the aforementioned approaches is the explicit nature of training. Explicit training is a very resource demanding technique, which becomes too expensive with complex scenarios. For these scenarios implicit training is much more appropriate. We anticipate that implicit training is feasible only when the teacher is fully observable. In case of E-Commerce, it is possible when a human operator is present “in” the World Wide Web, rather than “on” the World Wide Web. We argue that to achieve this humans and autonomous agents have to be embodied as avatars in 3D Electronic Institutions.

3. 3D ELECTRONIC INSTITUTIONS

3D Electronic Institutions is a concept that appeared from the combination of the Electronic Institutions and 3D Virtual Worlds technologies. This combination resulted in a working methodology, supported by a number of tools, for designing highly secure and reliable immersive 3D E-Commerce solutions. Applying 3D Electronic Institutions methodology requires 3 important steps to be accomplished:

- Specification of an Electronic Institution using ISLANDER (Esteva et al, 2002).
- Annotation of the Electronic Institution specification with components of the 3D Virtual World.
- Automatic generation of the corresponding 3D environment.

The specification phase introduces a dramatic difference to the development of 3D shopping environments compared to the majority of present agent-based solutions. Instead of focusing on the implementation details of each particular agent (agent-oriented approach), a system-oriented view is taken. We assume that participating agents may be heterogeneous and self-interested, and we cannot rely on their correct behavior. Therefore, the institution is designed as a set of limitations which every participant have to comply with. This assumption permits that agents behave autonomously and make their decisions freely up to the limits imposed by the institution. The limitations include a common ontology, a set of permitted activities (scenes) that the agents can be involved in, participants’ role flow within the institution (connections between scenes), synchronization points (transitions), institutional norms (obligations) and the dialogues that govern the enactment of different scenes. The specification is produced with ISLANDER tool, a UML-like editor that also performs several verifications (integrity, protocol correctness, and norm correctness).

While the specification strictly defines the limitations, it also helps to understand what participants need in order to operate in the institution. Some elements of the specification have conceptual similarities with building blocks in 3D Virtual Worlds, which makes it possible to automatically generate a 3D representation of the specification. The scenes and transitions, for example, are transformed into 3D rooms, connections correspond to doors, and the number of participants allowed in a scene determines the size of a room (see Bogdanovych et al, 2005 for details).

The created 3D Electronic Institution is ready to be executed. The generated 3D Virtual World can be visualized and the 3D Electronic Institutions infrastructure will take care of the validity of interactions between participants, verify the permissions of participants to access different scenes and will make sure that all the institutional norms and obligations are imposed. To make sure that the institution is fully protected from deviant behavior a special type of agent (governor) is used as a mediator between the institution and a participant. The only way for a participant to talk and act in the institution is through a governor.

Conceptually, 3D Electronic Institutions have two different levels of execution: *institutional level* and *social level*. The institutional level makes sure that the institutional rules are not violated. On this level all a participant can do is to send a text message to the corresponding governor for verification, requesting to perform an action in the institution. If under the current circumstances the action is allowed to be performed without violating the rules, the agent receives back a response and the action is performed (visualized). The way these actions are visualized can not be changed, as well as their execution can not be terminated. For human participants sending of text messages is transparent as it happens as a result of the actions in the

Virtual World (e.g. when an avatar representing a human collides with the exit door of the room, the corresponding “exitScene()” message will be automatically generated and sent to the governor).

Actions that are not controlled by the institution are performed at the social level. These actions are executed directly by the participant without prior verification by the governor. For example, there is no need for the institution to specify how the participants should walk from one room to another.

Figure 1 presents an example of a simple institution and shows the actions performed on both social and institutional level. The institutional level actions include: (enterScene, exitScene, enterTransition, exitTransition and login). On the social level these are: moving, clicking, colliding, rotating etc. The black arrows on the picture show the trajectory of the participant’s movement through registrationRoom, meetingRoom to TradeRoom. The black figure represents the participant, other figures correspond to internal agents (employees of the institution) Receptionist and Auctioneer. The Receptionist welcomes the participant in the RegistrationRoom, verifies the login and password and unlocks the doors to other scenes if the identity of the participant is proven. The Auctioneer sales different goods in the TradeRoom. It announces the current product to be auctioned, waits for incoming bids and sells it to the winner of the auction. The MeetingRoom is used for social interaction between buyers.

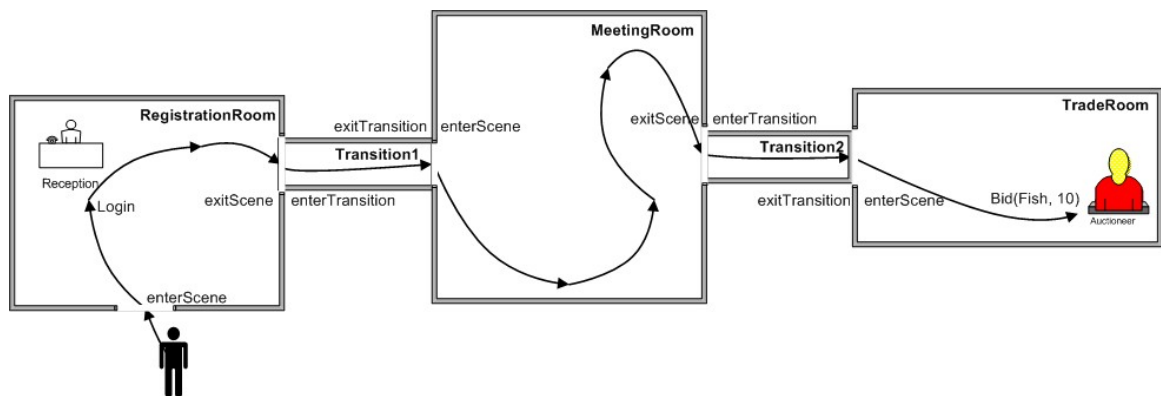


Figure 1. Two Levels of Execution in 3D Electronic Institutions

4. IMPLICIT TRAINING IN 3D ELECTRONIC INSTITUTIONS

In 3D Electronic Institutions the couple agent/principal is represented by an avatar. The autonomous agent is always active, and when the human is driving the avatar the autonomous part is observing the behavior of the human part, learning from his/her behavior patterns. The vision that inspires our proposal is graphically expressed in Figure 2. The avatars are fully controlled by either humans or autonomous agents through an *interface* (the interface is a sort of *glove puppet* that translates all decisions of its *puppeteer* into terms of the institution machine understandable language). Our view is that the agent and the human co-operate in the solution of the tasks the human has to deal with. We want to permit that either the human takes full control over the avatar or that the autonomous agent is in full charge of the decision making process. Moreover, we want to allow other types of interaction among them, such as the human giving guidelines to the agent, or the agent suggesting potential solutions to the human (via the interface), in a sort of “expanded intelligence” mechanism similar to the “expanded reality” that nowadays virtual reality tools offer.

Such duality (agent/principal) is a general feature of 3D Electronic Institutions and every participant (either a buyer or a shopping assistant) is integrated into the system via such architecture. This particularly means that each autonomous agent in the institution is trained by humans in the same way, which permits a decentralized approach to user modeling. Each autonomous agent only observes its principal and dynamically updates the user model of the principal. In case an agent (e.g. shopping assistant) needs to obtain the information about another agent (e.g. buyer), instead of trying to observe the behavior of other’s avatar and use very sophisticated modeling techniques, the agent simply sends a direct request to the autonomous agent responsible for the avatar. If the agent on the other end agrees to share the information it will reply with the relevant part of its current user profile. Such solution significantly reduces the amount of computations and

the size of stored data. It also permits easier control over privacy (e.g. if a participant doesn't want to be observed he/she just prohibits the agent to share personal information with others).




Figure 2. Implicit Training and User Modeling

As it was described in the previous section, the 3D Electronic Institutions comprise two different levels of execution. For effective building of user profiles and for making the implicit training of participants possible the actions from both levels have to be observed by autonomous agents. In the model we propose the institutional level gives context for learning and the actions of the social level are very important for teaching human characteristics to autonomous agents.

The actions of the institutional level, on the one hand, help the autonomous agent to understand when to start and stop recording the actions of the social level and which context to assign to the recorded sequences. On the other hand, analyzing the sequence of institutional level actions helps, in a long run, to understand how to reach different scenes and states in those scenes.

One of the valuable sources of information that can be obtained in the social level is the trajectory of the user. According to (Bauer & Deru, 2005) the trajectory of human's movements is tightly connected with the cognitive load. Figure 2 outlines the case of a human controlling the avatar marked as buyer and an autonomous agent controls the avatar of a shopping assistant. The simple scenario starts by the shopping assistant representative agent that notices a buyer approaching its avatar. Before getting involved into a conversation the shopping assistant requests information on human's "shopping mood" from the buyer representative agent. The agent analyzes the trajectory of the principal (from the moment when the principal entered the institution) and realizes that the human is in a hurry. This information is sent back to the shopping assistant representative agent. Before the human starts a conversation, the shopping assistants already knows that its responses have to be very short and precise.

Technologically such scenario can be supported by shallow reasoning. We record a classification list of trajectories generated by avatars driven by humans, the time the avatar spent on the trajectory and the number of utterances in which the avatar was engaged. Depending on these values, there is a class label associated with each trajectory (e.g. "in a hurry", "exploring", "waiting"). The final classification is refined by a human expert. At classification stage we use only the information about the initial (approaching) trajectory. In order to classify the trajectory of the avatar we compare it with every trajectory in a classification list and identify it as the most similar one from the list. The classification list is filled manually by system designers and each trajectory there is annotated with a text label (e.g. "in a hurry"). Technically, the trajectories are specified as arrays of landmarks. Each of the landmarks corresponds to the position of the avatar in a given moment in time. The position is permanently updated by the system every 10 Msec, so the information about the avatar's velocity is easily reconstructed from the distance between two neighboring landmarks. This simple representation allows efficient trajectory comparison and classification. To increase the performance of the classification on the first step of the algorithm the irrelevant landmarks (noise) are removed using the

approach presented at (Perng et al, 2000). After this, a combination of *Levenshtein Distance* and *Euclidean Distance* algorithms is applied to compare the analyzed trajectory with each trajectory stored in the classification list. As the result of the comparison, the trajectory from the classification list with the lowest distance value is selected and the corresponding text value is extracted to be used as a behavior label.

The actions of the social level are used not only to predict the “shopping mood” of the human. The autonomous agents record all those actions to imitate human behavior in a believable way. Regarding this point we base our research on the Bayesian method proposed by Le Hy et al, (2004). The application of this approach involves the observation of the set of environment variables in combination with the actions made by a human to produce a probabilistic table for predicting the next action to be executed by the agent.

To better explain how the training on each level is happening we refer back to the example on Figure 1. Imagine that an agent tries to learn how to represent a human in this simple institution. Initially, when a 3D representation is generated from the specification, the visualization of behavior of the autonomous agents in the institutional level is usually preprogrammed in a very simple way (e.g. entering a scene is visualized as making the avatar disappear in the previous room and showing it in the middle of the entered room). The actions of the social level can not be visualized on this stage. Now let’s observe what happens if the human takes control over the avatar to complete the scenario from Figure 1. The autonomous agent starts recording the actions of the social level when it receives the message “enterScene” as the result of human entering the RegistrationRoom. When the login message is sent by the human through the interface the agent realizes that the context in the social level has changed (which means that all new actions have to be observed in a new context). In the similar way the agent records all other actions and assigns context to them when human moves through scenes and transitions to the TradingRoom. In the institutional level the agent records the whole sequence of the institutional messages sent by the human (enterScene, exitScene, enterTransition, exitTransition, login, bid). Now let’s assume that after participating in the fish auction the human has bought a box of fish for the price of 10\$ and left the institution. The next time the human enters, he/she expects that the corresponding autonomous agent is already smart enough, and uses a special command “Do:bid(fish, 10)” to instruct the agent to buy the fish. The first thing what the agent does is searching the prerecorded sequence of institutional actions for the appearances of bid() function. Once the function is found, the agent knows which sequence of the institutional level actions will lead to achieving the goal. On the social level the agent knows which actions it has to execute for believable imitation of the human performance. The aforementioned reasoning will result in the following behavior: the avatar enters through the door to RegistrationRoom, the avatar reproduces the trajectory of the human and approaches the reception desk, the request for login information is received and the agent sends the login details. In the similar way the agent continues its movement to the TradingRoom, where it offers 10\$ for the box of fish. If the agent wins the lot – the scenario is finished; if the price this time is higher – the agent will request the human intervention.

5. CONCLUSION AND FUTURE WORK

We presented the concept of implicit training that is used for teaching human behavioral characteristics to autonomous agents in 3D Electronic Institutions. These environments and the proposed approach are beneficial for E-Commerce as an embedded training technology for increasing the acceptance and believability of virtual shopping assistants. Though we undertook a very detailed exploratory research of the implicit training concept the validation of it was not yet performed due to the insufficient implementation. The goal of the paper was to present our theoretical findings and position our research and described technological solution in the context of a new generation E-Commerce technology. Future work will include the implementation of the aforementioned algorithm for implicit training and extension of the already existing path recognition software with the possibility of automated generation of the path classification list.

ACKNOWLEDGEMENT

This research is partially supported by an ARC Discovery Grant DP0451692, Institute for Information and Communication Technologies (IICT) and the UTS e-Markets Research Program. Carles Sierra acknowledges partial support by the Spanish project WEBI2, TIC2003-08763-C02-00.

REFERENCES

- Aleotti, J. et al, 2003. Toward Programming of Assembly Tasks by Demonstration in Virtual Environments, *12th IEEE Workshop Robot and Human Interactive Communication*. October 31 - November 2, 2003.
- Alissandrakis, A. et al, 2001. Through the Looking-Glass with ALICE — Trying to Imitate using Correspondences. *Proceedings of the First International Workshop on Epigenetic Robotics*.
- Bauer, M. et al, 2000. Programming by example: programming by demonstration for information agents. *Communications of the ACM*. Volume 43, Issue 3. pp. 98-103.
- Bauer, M. and Deru, M. 2005. Motion Based Adaptation of Information Services for Mobile Users. *Working Notes of the UM05 Workshop on Machine Learning for User Modeling*. Edinburgh, Scotland.
- Bogdanovych, A. et al, 2004. E-Commerce Environments as 3D Electronic Institutions. In *Proceedings of the E-Commerce 2004 Conference (EC2004)*, Lisbon, Portugal, pp. 237-242.
- Bogdanovych, A. et al, 2005. Narrowing the Gap between Humans and Agents in E-Commerce: 3D Electronic Institutions. *Proc. of the 6-th International Conference on Electronic Commerce and Web Technologies*, pp. 128-137.
- Chowdhury, S. 2004. *Next Generation Business Handbook. New Strategies from Tomorrow's Thought Leaders*. John Wiley & Sons, pp. 494-495.
- Daugherty, T. et al, 2005. Experiential E-Commerce: A Summary of Research Investigating the Impact of Virtual Experience on Consumer Learning. *Online Consumer Psychology: Understanding and Influencing Consumer Behavior in the Virtual World*. Lawrence Erlbaum Associates, pp. 457-489.
- Eartlink. 2005. How to Convert More Online Shoppers with Live Chat Invites. *Practical Know-How & Case Studies*. Sherpa marketing.
- Edwards, S. & Harshavardhan, G. 2001. The Novelty of 3D Product Presentations Online. *Journal of Interactive Advertising*, 2(1).
- Ekvall, S. and Kragic D. 2005. Grasp Recognition for Programming by Demonstration, In *IEEE/RSJ International Conference on Robotics and Automation*.
- Esteva, M. et al, 2002. ISLANDER: An Electronic Institutions Editor. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent systems (AAMAS '02)*, pp. 1045-1052.
- Früh, C. et al, 2005. Data Processing Algorithms for Generating Textured 3D Building Facade Meshes from Laser Scans and Camera Images. *International Journal of Computer Vision*, 61(2), pp. 159-184.
- Gutzman, A. 2000. Online Customer Service: Do or Die the Law of Averages, *CRM on the Web CRM Library*.
- Hunter, D. & Lastowka, F. G. 2003. To Kill an Avatar. *Legal Affairs Magazine*. Issue July-August 2003, pp. 20-24.
- Hurst, M. 2000. Why 3D shopping makes no sense. [online] cited at <http://www.goodexperience.com>.
- Klein, S. & Loebbecke, C. 1999. Signaling and Segmentation on Electronic Markets: Innovative Pricing Strategies for Improved Resource Allocation. *Proceedings of the 6th Symposium on Emerging Electronic Markets*, pp. 127-142.
- Le Hy, R. et al, 2004. Teaching Bayesian Behaviours to Video Game Characters. *Robotics and Autonomous Systems*. Vol. 47. pp. 117-185.
- Magnenat-Thalmann, N. et al, 2005. Believability and Interaction in Virtual Worlds. *Proceedings of the 11th International Multimedia Modelling Conference, (MMM 2005)*. pp. 2-9
- Maguire, J. 2002. Case Study: LivePerson Searching for the Personal in E-Commerce. *Ecommerce-guide.com*, 23.02. 02.
- Modahl, M. 2000. *Now or never: how companies must change today to win the battle for internet consumers*, New York, HarperBusiness. ISBN: 0066620120.
- Pai, D. K. et al, 2001. Scanning Physical Interaction Behavior of 3D Objects, in *Computer Graphics (ACM SIGGRAPH 2001) Conference Proceedings*.
- Perng, C. et al, 2000. Landmarks: a New Model for Similaritybased Pattern Querying in Time Series Databases. *Proceedings of the 16th International Conference on Data Engineering*, pp. 33-42.
- Preece, J. & Maloney-Krichmar, D. 2003, *The Human-Computer Interaction Handbook*, Lawrence Erlbaum Associates Inc., chapter Online Communities: Sociability and Usability, Mahwah, USA, pp. 596-620.
- Prendinger, H., Ishizuka, M. 2001. Let's talk! Socially intelligent agents for language conversation training. *IEEE Transactions on Systems, Man, and Cybernetics*, Part A 31(5): pp. 465-471
- Prince, M. 2005. E-Commerce: Help May Be on the Way, *The Wall Street Journal*, March 21, 2005.
- Rafaeli, S. & Noy, A. 2005. Social Presence: Influence on Bidders in Internet Auctions. *EM-Electronic Markets*. Volume 15(2). pp. 158-175.
- Rule, D. S. 2003. Artificial Intelligence (AI) Chatbot http://www.codetools.com/useritems/AI_Chatbot.asp. 12.2003.
- Whitton, M.C. 2003. Making Virtual Environments Compelling. *Communications of the ACM*. Volume 46, Issue 7 (July 2003), SPECIAL ISSUE: "A game experience in every application", pp. 40-47.