Taylor & Francis
Taylor & Francis Group

FROM THE EDITORS OF THE SPECIAL ISSUE

**Zsófia Ruttkay[1], Michael Kipp[2], Anton Nijholt[3],
and Hannes Högni Vilhjálmsson[4]**

[1]*Creative Technology Lab, Moholy-Nagy University of Art and Design, Budapest, Hungary*
[2]*DFKI, Embodied Agents Research Group, Saarbruecken, Germany*
[3]*Human-Media Interaction, University of Twente, Enschede, The Netherlands*
[4]*School of Computer Science, Reykjavík University, Reykjavík, Iceland*

Welcome to the special issue on intelligent virtual agents. Intelligent virtual agents (IVAs) are interactive characters that, in spite of being merely 2D or 3D computer graphics models, exhibit human-like qualities and communicate with humans or with each other using natural human modalities such as speech and gesture. They are capable of real-time perception, cognition, and action, allowing them to participate in a dynamic physical and social environment.

Such full-fledged IVAs are still a long-term research vision. Even if we restrict ourselves to the nonverbal modalities used in communication, current IVAs would not pass the Turing test. However, our research community can be proud of many achievements. First, IVAs are already around in our daily lives: You have probably come across some IVAs in the guise of information providers, sales people, educators or assistants, or as inhabitants in virtual worlds or nonplayer characters in games. Second, we have come far considering how young the field is. It was established less than two decades ago, stimulated by the pioneering work of a few researchers on the synthesis of human motion, hand gestures, and emotional facial expressions. And finally and most importantly, the challenges are enormous: to design and implement virtual characters that, even if they cannot be taken for real people, are engaging enough to "suspend the disbelief" of the human viewer or interlocutor. This is the precondition for the successful application of IVAs that make us study more effectively, do fitness exercises properly, follow their advices, and get emotional. The realization of engaging IVAs is a huge and multidisciplinary task, so reusable modules and tools are of great value. The development of IVAs requires expertise not only in multimodal interaction, but also in several artificial intelligence (AI) fields such as

cognitive modeling, planning, vision, and natural language processing. Computational models are typically based on experimental studies and theories of human–human interaction; conversely, IVA technology may provide interesting lessons on how real people communicate.

Today, we have an established and still growing research community, with dozens of research groups focusing on everything from specific aspects of the cognitive and communicative capabilities to complete IVA system frameworks and supporting tools. There exist interdisciplinary forums, EU projects, and Networks of Excellence to establish common ground between different disciplines, such as psychology and computer graphics. The goal is to gather and disseminate the small pieces of the jigsaw puzzle we face when attempting to reproduce the characteristics of human behavior.

This special issue gives a glimpse of the field through eight revised articles selected from mature work presented at the major event of the field, the 9th International Conference on Virtual Agents, held September 14–16, 2009 in Amsterdam, The Netherlands.[1] The traditionally 2.5-day, single-track conference attracted some 150 participants, featured invited speakers from linguistics and game industry, and one speaker demonstrating the value of combining the artistic and scientific approach. The special application theme of IVA 2009 was games. To date, the creation of game characters and their social behavior has largely relied on careful hand crafting rather than automation. The conference gave an opportunity to reveal, tackle, and discuss the issues that relate to using IVAs in games. Long and short articles and demos, as well as the entertaining and illuminating entries mostly by students to the GALA (Gathering of Life-Like Animated Agents, which can be seen online[2]) gave the impression of a vigorous and nicely cooperative research community.

Before you start browsing through the articles of this volume, it is useful to see how they contribute to the "big picture." One can look at an IVA from different points of view:

1. Body and mind

When endowing a single IVA with behavior capabilities, one has to provide mechanisms at several different layers. On the lowest level, there must be some mechanisms for creating and deforming a body model and synthesizing sounds for human speech. On the next level, these mechanisms can be used to create single gestures, facial expressions, gaze movements, and synthesize words in some language. At a yet higher level these elements can be combined, in sequence and in parallel, to form smooth, well-timed multimodal utterances-enough to create, for example, a virtual news-reader. But if we want to place our IVA in a conversation or in another similar dynamic environment, we also need perception capabilities. And when we expect the IVA to demonstrate its awareness

of the surrounding people and environment, cognitive mechanisms are necessary to generate visible reactive behavior. On the highest level, when the IVA is an autonomous agent, it has to be endowed with common-sense knowledge about the world (including other, virtual, or real humans they encounter), update beliefs about its changing state, maintain goals, and also respond to the experiences emotionally. At this highest level, the bodily behavior ranging from deliberate communicative actions to subtle details of unconscious bodily behavior such as respiration or blinking rate should all be treated as display of the cognitive, emotional and physical state of the character, including the case when the IVA "finds it more appropriate" to hide or mask the "felt" emotions or real physical state to comply to social rules or in order to achieve its goal.

2. Identity and individuality

Another dimension of complexity is the identity of the character, which must be addressed if we deal with several characters or a crowd. But even when a single IVA acts as for example a salesperson, it is useful if the IVA can accommodate the (social, ethnical, age) identity of the (real) customer. IVAs must express their identity in appearance and in behavior. Applying the same motion patterns (e.g., for walking or for facial expressions) to all characters immediately hits the human eye as machine-like. Moreover, characters must suggest identity: On the one hand, each character is recognized as member of some groups (racial, age, culture) or, lacking such features, as a stereotype, a sterile and machine-like virtual human. On the other hand, characters must differ, as people differ in their physical outfit and motion characteristics, personality, and personal history. The characters' (re)actions–both what they do and say and the way they do it-should also reflect their identity. This dimension implies that a new inter-character behavior layer is to be established, as soon as multiple characters are to act in a social and physical space (e.g., taking care that the virtual characters do not collide with each other and that they greet each other in a proper way, according to accepted social strategies).

3. Life span

Virtual agents still count as technological novelty, which may imply a positive bias when evaluating how people perceive them in lab experiments. It is a completely different situation if people get exposed to an IVA as a permanent staff in a frequented shop or as one's daily training companion for weeks. Such real-life situations can be considered real test cases of IVAs' acceptance. How effective and likable IVAs turn out to be often depends on secondary aspects, such as the quality of bodily behavior and the match between the bodily features (look) and nonverbal signals of the real and virtual human. Moreover, if an IVA is to live long, this raises yet further design challenges such as maintaining memory (it is not good to tell the same joke every day!), building long-term relationship with

its (real or virtual) partners, aging, or reflection on the cultural (marked days) and the meteorological (daily weather) local aspects of the calendar.

The articles in this journal represent broad ranges in the above dimensions. The first three articles present work on endowing IVAs with nonverbal bodily behavior. From the first article by de Melo et al. you learn that it is worth the trouble to endow IVAs with such subtle autonomic signals as wrinkles, blushing, sweating, tears, and respiration, because they have a significant effect on the perception of several emotions and cognitive states displayed by the IVA. The article is a good example of the different layers of bodily expressivity, as it discusses the basic mechanisms of producing the signals and their role in expressing the cognitive and emotional state of the IVA. The work also exemplifies the typical iterative process and major topics of the development of expressive IVAs: Engineering is needed to accomplish real-time synthesis of communicative signals, but the signal patterns themselves need to be generated according to findings from psychological studies on signal patterns in humans, and their actual perception needs to be tested experimentally. In the next article, Heloir and Kipp discuss EMBR, a state-of-the-art character animation engine that supports the definition of multimodal gestures (in the broadest sense, since other modalities than the hands are covered) and a rich control of their motion qualities. Besides expressivity and real-time generation for the nonverbal behavior, the other merit of the engine is that it provides elegant mapping from a slim functional control level (provided by an external "mind" module) to the bodily behavior. In the third article, Bergmann and Kopp address the question of what gestures an IVA should use, both concerning the formation (what does the gesture look like, by which hand is it performed, etc.) and the usage of certain gestures, in their case iconic gestures used for indicating characteristics of objects. Based on video recordings of humans gesturing and hand annotations, the authors use and evaluate different machine learning techniques to create models for automatic gesture formation, where these models reflect individual style. They lay out a new path for using machine learning as a research paradigm to find out about gesture usage in humans as well as suggest a novel knowledge representation to be used in the process of generating gesturing in IVAs.

The next few articles explore the second dimension, identity and individuality. In their article, Mascarenhas et al. use a fundamental, value-base, cultural model as a basis for a computational module that drives the actions and influences the verbal and nonverbal behavior of IVAs. They developed a prototype using a dinner party scenario to demonstrate the difference in behavior between an individualistic and a collectivistic culture. The article by Pedica and Vilhjàlmsson addresses the generation of basic-level reactive behavior when groups of people come together.

They offer a computational model of human territorial behaviors based on social norms, which generates visible social dynamics when virtual agents interact. These two works illustrate the bridging of a cognitive layer and a bodily behavior layer along the body–mind dimension. The next article by Rodrigues et al. tackles the problem of steering groups of agents from a completely different angle, in the flavor of game developers: They adopt a space colonization algorithm, more typically used for generating leaf venation patterns. They use the parameterized algorithm to create responsive, realistic-looking group behaviors such as the formation of groups, positional alignment, and collision avoidance.

The last two articles talk about IVAs in actual applications. In the first article Bogdanovych et al. provide a full-fledged environment for building virtual worlds where virtual inhabitants reenact daily life in different historic cultures, providing an example of the application of culture models. The scope and effectiveness of the general AI environment is illustrated with a prototype developed for the City of Uruk, 3000 B.C., that can be experienced on the online Second Life platform. Finally, the selection closes with the article awarded the "Best article" of IVA2009 by Bickmore et al., addressing the topic of long-term user engagement with an interface agent. Such studies are not frequent, because it is much harder to monitor the effectiveness of an IVA over a time span of months or even years than hours or days. The authors have developed their own "virtual laboratory" system to run longitudinal studies on interaction between a human and a virtual agent. They have shown that variability, manifested in subtle changes in language and visual representation, and the "back story" of their fitness assisting virtual agent have a positive effect on engagement and liking.

However, you will not finish this special issue on a blissful note, because there was no positive effect on actual exercise behavior in Bickmore's study. But for an expert in the field, this is not a shocking result. Several other experiments have concluded that there is no correlation between a likable IVA and an effective one. Let's welcome such puzzling open questions and appreciate how they indicate the complexity of human behavior. IVAs are getting smarter in mind, smoother in behavior, and more enjoyable by the year, yet there is still much work ahead of us to understand how we, the *real* people, operate in our everyday life and how we can then develop IVAs that perform comparably well in their human-like jobs.

April 2010

NOTES

1. http://iva09.dfki.de/
2. http://hmi.ewi.utwente.nl/GALA