ARTIFICIAL INTELLIGENCE AND ENTERTAINMENT

MARCO GORI
Dipartimento di Ingegneria dell'Informazione - Università degli Studi di Siena

OLIVIERO STOCK
ITC-irst - Trento

1 Introduction

Computers are now seen as companions that are expected to help not only with efficiency and effectiveness, but also in supporting our basic need for entertainment, pleasure, and aesthetics. Research in the field of entertainment has also been addressing the development of perceptual, cognitive and social capabilities. Reasoning and learning will be at the core, and so modeling of the user, discourse, context, domain, media.

Models of emotion/entertainment (e.g., fear, joy, surprise) will become important both for virtual characters and for understanding human user input (e.g., user satisfaction, and frustration). There is a long tradition in artificial intelligence in this direction, starting with popular games such as chess, checkers, and other puzzles, in which computers are exhibiting increasing performance in what is often regarded as one of the most evident manifestations of intelligence. Recent advances in artificial intelligence and human computer interaction offer unprecedented and seemingly endless opportunities for enhancing traditional forms of entertainment and supporting the creation of new ones.

At a recently held INTETAIN conference [14], the first in Europe to embrace the whole field, intelligent technologies for interactive entertainment were presented. These include adaptive media presentations, recommendation systems in media scalable crossmedia, affective user interfaces, intelligent speech interfaces, tele-presence in entertainment, collaborative user models and group behavior, collaborative and virtual environments, cross domain user models, animation and virtual characters, holographic interfaces, augmented, virtual and mixed reality, computer graphics and multimedia, pervasive multimedia, creative language environments, computational humor, and the list goes on. An important role is proposed for novel underlying Interactive Device Technologies (for example mobile devices, home entertainment centers, haptic devices, wall screen displays, holographic displays, fog screens, distributed smart sensors, immersive screens and wearable devices). Interactive applications for entertainment were to include, among others: intelligent interactive games, intelligent music systems, interactive cinema, edutainment, interactive art, interactive museum guides, city and tourism explorer assistants, shopping assistants, interactive real TV, interactive social networks, interactive storytelling, personal diaries, websites and blogs, comprehensive assisting environments for special segments of populations (e.g. impaired, children, elderly).

In this paper we shall not cover all aspects of intelligent technologies for interactive entertainment. Instead, we shall focus mainly on some of the areas that have seen a wide range of original contributions in Italy.

2 Games

As we cannot talk about entertainment without discussing games, let us first mention that a tendency toward automatic learning prevails also in games. In the 2005 edition of the AAAI international conference, the American Association for Artificial Intelligence introduced a general game playing competition that would "test the abilities of general game playing systems by comparing their performance on a variety of games". Players are not told anything about the games to be played before the competition and the rules of all games are electronically transmitted to the players at the beginning of each game using a Game Description Language (games.stanford.edu/language.html). Systems compete against one another for a $10,000$ USD prize.

Concerning the strategic role of games in research, Minsky pointed out that

> It is not that the games and mathematical problems (puzzles) are chosen because they are clear and simple; rather it is that they give us, for the smallest initial structures, the greatest complexity.

This provides significant motivation for the study of games especially for their impact in different fields. But let us go back to the best known challenge of AI in the domain of games.

2.1 Computer chess

The ambition of being able to compete with the best chess players has been one of the original challenges that have characterized artificial intelligence. At the time of the Dartmouth Conference, in 1956, the belief existed that within some ten years it would have been possible to realize a program that would prevail over the best human players. The truth is that at the beginning of the 1990's, this goal was far from being achieved–the best chess programs were not at the level of the hundredth best human chess player. Later IBM invested enormously in a special project that led to the development of Deep Thought, first, and then of Deep Blue. Deep Blue was based on special hardware and substantially exploited brute force, with its capability of exploring 200 million moves per second. The system, after a first failed challenge, succeeded in 1997 in overcoming Garry Kasparov, generally considered the best human chess player of all time. Kasparov left the game site furious, claiming that the operators of Deep Blue had cheated, for instance by contravening the agreements about operating the computer between games. In any case, he never accepted the result. Nonetheless, IBM declared victory, and Deep Blue never played again. The results of Deep Blue formed the basis of a number of other practical projects that IBM sustained in the following years.

Even though Deep Blue was not known to have incorporated aspects that are typical of human intelligence, the AI community took advantage of the result and declared that one of the best known goals of AI, put forward by the founders of the discipline, had been attained. 1997 marked a milestone, but it was not the end of the story. The following years saw the development of a new generation of programs, running on traditional hardware - powerful PCs, but based on strategies that come closer to the ones used by humans. Deep Junior for instance, is a program developed by two Israelis, Shay Bushinsky and Amir Ban, that has been world champion three times in the 2000s. The system has played also in Italy against the best Italian player, Michele Godena and against one of the best female world players, Almira Skripchenko. Deep Junior includes pattern matching, machine learning, opponent modeling, reasoning and planning. In general the new programs, like Deep Junior (world champion in 2006), Shredder, or Zappa, are specifically built for competing at the world championship of computer programs. Yet in 2003, Deep Junior challenged Garry Kasparov, with a tie as the final result. However this time Kasparov declared himself satisfied with a tie result, and clearly played the final two games on the defensive, impressed by Junior's capability to change strategies in the course of the game. Kasparov seemed to recognize that the future of "chess intelligence" had been declared. As for the match against Michele Godena, held at IRST, Trento on September 27, 2004, Deep Junior appeared to give less importance to material advantage than to strategic positioning. The program also displayed moves that the expert chess commentators described as indicators of authentic creativity. In the end Godena expressed his admiration. A similar result was obtained in the match against Skripchenko. Bushinsky declared that there are many open problems and research challenges in computer chess; but the one role he sees for chess programs in the future is to propose creative solutions: just the opposite of what folk wisdom attributes to computers.

In May 2006, Italy was at center stage, hosting the World Championship in Turin for the first time in conjunction with the human chess Olympiads. Deep Junior, with some recent improvements won the title again, even though with much less computing power than other competitors: for instance, Zappa was using 512 processors on the NCSA supercomputer.

2.2 Cracking crosswords

Crosswords are one of the most popular linguistic puzzles worldwide and represent a very interesting challenge for machines in both their production and solution. Problems of this category have been informally defined as AI-complete [13], meaning that there is no closed-world assumption and they require human-level knowledge to be attacked.

Interestingly, for the first time since AI's kick-off, there is a first nucleus of technology, such as search engines, information retrieval and machine learning techniques, that enable computers to capture real-life concepts with semantics. Research towards both the generation and the solution of crosswords has been carried out at the University of Siena in the last three years. The main approach is to attack crosswords making massive use of the Web as a sort of self-updating repository of human knowledge. The research activity in crossword generation is still in its infancy. The main challenge is not allocating words of a dictionary under the crossword's constraints, but rather producing meaningful and creative cues. In principle, one could simply access definitions from a database of crosswords' clues. Not only is this not interesting, but it would also be useless since it would infringe on the copyright on published crosswords used to construct the database. While this is a very controversial issue and many crosswords contain definitions that have already been published elsewhere, the problem would surely arise because of the systematic copy of clues. A very stimulating challenge is to make them by appropriate processing of web pages, aimed at discovering the relevant definition of the term.

Unlike generation, significant advances have been obtained in crossword puzzle solving under the WebCrow project (http://webcrow.dii.unisi.it) that was announced the first time during the afore-mentioned chess competition at IRST in September 2004. A few weeks later, the system was reviewed in the news of Nature [16] and, later on, a brief description appeared on the AAAI Web site concerning games and puzzles (http://www.aaai.org/AITopics/html/crosswd.html), and in [6]. To the best of our knowledge, the only significant attempt reported in the literature to tackle this problem is the Proverb system, which has reached human-like performances [19].

Unlike Proverb, WebCrow neither possesses any knowledge-specific expert module nor a crossword database of great dimensions. In order to stress the generality of its knowledge and language-independency, WebCrow has a module at its core performing a special form of question answering on the Web, referred to as call clue-answering. The second goal of the system, filling the crossword grid with the best set of candidate answers, has been tackled as a Probabilistic-CSP. WebCrow has a striking advantage with respect to Proverb due to its web-based architecture that makes it possible to attack crosswords in different languages without substantial changes. Most of its modules are in fact based on text mining techniques that invoke machine learning and that can be reused in a straightforward manner when changing the language. Interestingly, also the language-dependent parts of the systems have been designed in such a way to minimize the porting to a different language.

WebCrow works currently for Italian and English, though it has been massively experimented for Italian crosswords only. WebCrow has been recently involved in a competition that took place at the University of Siena against 181 students, distributed in five labs. WebCrow reached the 55-th position, exhibiting better performance on hard more than on easy crosswords (see the details at http://webcrow.dii.unisi.it). Concerning English, WebCrow is expected to play against volunteers at the ECAI-06 conference. Moreover, a new game for bilingual people will be proposed in which crosswords will be composed of both English and Italian clues. Interestingly, while the competition resembles popular chess events, there is also a significant difference: there is no widely recognized tradition of "crossword masters" like there is in chess and this is why the idea of involving as many people as possible is reasonable to assess actual developments.

The WebCrow project has been supported by Google Inc. under the Research Program Awards.

2.3 Web games

The Web has been offering a natural environment for games that are relevant not only for entertainment but also for the evolution of technical issues, like labeling. For instance, the ESP game (available at www.espgame.org) is played by two people on the Web who cannot communicate with each other [23]. They don't know anything about each other: the server pairs them and chooses an image from the web, shows it to both partners simultaneously, and asks them to label it. Each player's goal is to guess what the other is typing, without any communication. As soon as they both type the same phrase, the system offers a new image and they play again. Meanwhile, behind the scenes, that phrase is recorded as metadata for the image. This is a simple yet effective means of producing reliable metadata.

Tagsocratic, a system for coordinating labeling in the blogosphere and for facilitating the social interaction among bloggers, originated from this simple idea and is currently being developed at IRST, Trento [8]. TagSocratic offers a new type of collaborative entertainment in which emergent topics of shared interest are discussed and developed online. The system is based on an agent-based architecture where agents interact to learn their respective topic competence.

Other commendable initiatives concern the impact in education. A Web game has been developed at the University of Genoa to support studies for discovering European heritage. The project, called ChiKho is supported by the European Union [2] ChiKho is essentially a web-distributed educational game which allows players to share and improve their knowledge about the heritage of European countries. The game structure is similar to a hurdle-path game. The system handles a huge amount of multimedia contents that are made available to several network configurations. The massive experimentation of the system has demonstrated that participants have been given the opportunity to acquire interesting information and to discover intriguing links between European cities.

2.4 Puzzles and problem solving

The 15-puzzle conceived by Sam Loyd, the Rubik's cube and related problems are popular games for both humans and machines. They represent the ideal prototype problems to illustrate the aim of problem solving: given an initial state and a set of operators (admissible moves), find a goal state. For both mentioned puzzles, the problem is reasonably easy to solve if we are satisfied with just a "path-solution," regardless of whether or not it is the optimal one. These problems become significantly harder if the purpose is to discover the optimal parth from the initial to the goal state. Interestingly, these are typical problems in which machines already significantly overcome human skills. On the other hand, the discovery of optimal solutions is very expensive and, as a matter of fact, can be achieved only for small-scale instances of the puzzles.

For instance, while traditional heuristics for attacking the Sam Loyd Game, like the Manhattan distance, are quite effective for facing the eight- and most frequent instances of the fifteen-puzzle, the problem soon becomes intractable for higher dimensions. A crucial ingredient for the success of searching in the space of states is discovery of smart heuristics and proper use of the solution of sub-problems that gave rise to a method referred to as disjoint pattern database [18].

A more recent approach, followed at the University of Siena, is based on special heuristics that are learned from examples. In particular, likely-admissible heuristics have been introduced whose admissibility requirement is relaxed in a probabilistic sense. Instead of providing an upper-bound to the cost, it is guaranteed to end up with optimal solutions with a given probability. Interestingly, likely-admissible heuristics can be obtained naturally by statistical learning techniques such as artificial neural networks, which can learn the expected value of the cost to reach the target from examples. The experiments have shown that neural heuristics are the first online source capable to return a certain amount 29% of optimal solutions of the 15-puzzle with time and space costs lower than Disjoint Pattern Databases (non-reflected). When coupling two neural networks for the computation of the heuristics, the optimality degree reached approximated 50% with non growing search complexity. Compared to Manhattan Distance this search performed over 500 times faster by using 1/13000 of the memory resources [5].

3 Entertaining communication

Communication can be highly entertaining. Computer communication can be expected to have the same basic properties of human communication when artificial agents will be part of our ecological scene. These properties include: humor, engaging storytelling, or, novel human communication capabilities, like producing good real life pictures or video clips. Human computer interaction is one of the traditional areas of artificial intelligence. Aspects of an intelligent interface include multimodality, tailored interaction, interactivity, cooperation, and mixed initiative. Multimodal interfaces support multimedia (e.g., language (speech or text), gesture, gaze, graphics) and multimodal (e.g., audio, visual, haptic) input analysis and output generation as well as multimodal discourse processing (e.g., supporting within and across modal reference and discourse). Of course, entertaining communication can have a practical role in many situations: for instance one can think of the role that humor plays in advertisements, due to its capacity of getting the attention of the audience, and facilitating memorization of the message. Or simply consider the essential role that entertaining communication plays within children's education environments. In another context, if a story teller is able to engage the listener, it can provide a good companion to an elderly person and convince him/her of the importance of taking medicine and doing exercise. Or for instance it can provide attractive sport comments. Video clips can be a very effective way of transmitting knowledge when used in a documentary. Entertaining communication is beginning to deliver some initial prototypes. Even though humor is a very complex capability to reproduce, and has been considered AI-complete, it is realistic to model some types of humor production and to aim at implementing this capability in computational systems. Some prototypes are able to produce expressions limited in humor typology but meant to work in unrestricted domains [3]. Besides humor production, some initial works on humor recognition have been reported. Work has been conducted also on automatic emotional story telling, where the artificial character evokes empathy in the audience [12]. Automatic video production is a topic of research that is producing some results, especially in the simplified context that does not involve shooting in the real world, but starting from 2D existing images [4]. This is just the beginning. Robot-based systems for producing good quality pictures of events also constitute a promising avenue. All these themes open substantial challenges for the future. They will also require more insight on emotion, cognition and computation. There are a number of intriguing open questions that might be addressed in the near future such as, what is the nature of a surprise in communication? How to keep attention high?

3.1 Generating Humor Expressions Automatically

Humor is one of the most interesting and puzzling aspects of human behavior. Despite the attention it has received in fields such as philosophy, linguistics, and psychology, there have been few attempts to create computational models for humor recognition or production. Yet computational humor has the potential to change computers into extraordinarily creative and motivational tools. Computer-human interaction needs to evolve beyond usability and productivity. There is a wide perception in the field that the future is in themes such as entertainment, fun, emotions, aesthetic pleasure, motivation, attention, engagement and so on. Humor is an essential element in communication: it is strictly related to the themes mentioned above, and probably humans could not survive without it. While it is generally considered merely a way to induce amusement, humor provides an important way to influence the mental state of people to improve their activity. Even though humor is a very complex capability to reproduce, it is realistic to model some types of humor production and to aim at implementing this capability in computational systems. While humor is relatively well studied in scientific fields such as linguistics [1] and psychology [7], to date there is only a limited number of research contributions made toward the construction of computational humor prototypes. Almost all the approaches try to deal with incongruity theory at various levels of refinement [17, 1]. Incongruity theory focuses on the element of surprise. It states that humor is created out of a conflict between what is expected and what actually occurs in the joke. This accounts for the most obvious features of a large part of humor phenomena: ambiguity or double meaning. One of the first attempts is perhaps the work described in [3], where a formal model of semantic and syntactic regularities was devised, underlying some types of puns (punning riddles). The model was then exploited in a system called JAPE that was able to automatically generate amusing puns. Another humour-generation project was the HAHAcronym project [21], whose goal was to develop a system able to automatically generates humorous versions of existing acronyms, or to produce a new amusing acronym constrained to be a valid vocabulary word, starting with concepts provided by the user. The comic effect was achieved mainly by exploiting incongruity theories (e.g. finding a variation concerning religion for a technical acronym).

Humor recognition has received less attention. It is worth mentioning the work of [15] that investigated the application of text categorization techniques to humor recognition. In particular they showed that classification techniques are a viable approach for distinguishing between humorous and non-humorous text, through experiments performed on very large data sets.

3.2 Interactive TV

The recent impressive growth of TV content, digital TV networks and broadband requires smarter methods for making the interaction of the user more friendly. A research activity in this direction is carried out at Dipartimento di Informatica, University of Turin. The project was in cooperation with Telecom Italia Lab (formerly CSELT). Its main goal was the introduction of personalization techniques for the customization of future television services and the exploitation of such techniques within a prototype system for the generation of personalized Electronic Program Guides (EPGs). The customization of the EPG concerns the personalized selection of the TV programs to be advertized, on the basis of the user's interests. The proposed system is based on a multi-agent architecture, where specialized agents collect data about the available TV programs, monitor the user's behavior to acquire his interests and select the events to be advertised in the personalized EPG, depending on the user's preferences at the time of day he wants to watch TV. The system exploits multi-agent technologies for supporting the agent communication and it runs locally within the user's Set Top Box, where the tasks for the management of the EPGs are executed. The system is remarkably flexible and can easily be extended with agents providing new functionalities. In addition, the user's data can be handled locally to the Set Top Box, thus making it possible to pay special attention to privacy issues.

4 Educational entertainment

Already a number of interesting intelligent interactive elements are being incorporated into so called "edutainment" Reading, writing, and drawing assistants exist to help novices learn basic skills. Some incorporate, for example, speech processing (including accent models) so that young readers receive immediate and individualized performance feedback as they learn their native or a foreign language. Incorporation of interactive games seems to enhance user motivation. In the future, incorporating question answering systems could help to provide more tailored mentoring. Animated pedagogical agents in interactive learning environments enable effective situated learning [10]. For example in Cavazza and Simo [9] a qualitative simulation of the cardiac cycle in a virtual patient provides a safe and effective 3D virtual emergency room training environment for student physicians. Let us now focus on a kind of application that is of particular importance for Italy.

4.1 Museum guides

Various projects have introduced technology for mobile presentations triggered by the position of the visitor in the physical space. Such technology typically takes advantage of a localization system (for instance based on devices that generate an infrared signal from fixed positions, or based on triangulation through emitters/receivers of wireless digital signals, or on very sensitive GPS systems, which recently have achieved enough sensitivity to work inside buildings). The visitor typically carries a small portable device (for example a PDA or a simple headset), and receives information relevant to items at their particular location. From a methodological point of view, most of these projects concentrate mainly on the design of one technology and present a prototype and possibly a limited user study.

In the PEACH project, a complex and elaborated view was proposed [20]. Mobile and stationary components are integrated into the system seamlessly so that one can stop at a large screen and obtain a presentation appropriate for that device before moving on to presentations on a mobile device that take into account previous interaction. The system is initialized with a user profile and then, in the course of the visit, it adapts to the behavior of the visitor, proposing personalized, context-dependent presentations. Presentations themselves are multimodal, and in particular we have developed technology that combines language presentations and small, tailored, visual documentaries, meant to provide a coherent network of support for the visitor, for instance when guiding the visitor's gaze toward particular details of an exhibit. Another component of the system helps in identifying the focus of attention of the visitor when in front of a large bi-dimensional space, like a painting. The artificial vision-based technology pinpoints what region the visitor is currently looking at so that relevant presentations about the details of the area in the focus of the visitor's attention can be provided. In the course of a visit, the system capitalizes on each individual visitor's feedback whenever possible to guarantee appropriate presentations for their interests and taste. The system observes the visitor's behavior, and the elements comprising that behavior are interpreted as implicit input. Explicit input, on the other hand, is gathered in a very simple and "affective" way: feedback on the part of the user is limited to simple communication of liking or disliking the current piece of the presentation, with consequences for subsequent presentations. At the end of the visit, an overall personalized report that summarizes the key aspects of the visit experience is automatically produced for the visitor to take home or to receive as an electronic diary and an entry point for their subsequent cultural experiences.

4.2 Virtual reality for educational and medical applications

Augmented reality technology permits the concurrent interaction of computer-generated virtual objects with the real environment, thus making it an an interesting technology for developing educational applications for manipulation and visualization. This work has been carryed out at Fondazione GraphiTech in collaboration with Università della Calabria [22]. The interaction with complex objects that are deformable is an hard task. Studies on this subject are being carried out at the University of Siena where dynamic models of visuo-haptic interaction are studied to make the interaction as natural as possible. These studies have turned out to be important also for the development of Fetus-Touch, a system that allows one to haptically interact with 3D reconstructions of fetuses obtained by medical ultrasound imaging [11].

5 Conclusions

The research at the crossroad of artificial intelligence and entertainment has been growing fast for at least two independent reasons. First, as pointed out by Minsky, the studies on games give fundamental insights on the solution of many different problems that we believe might be very important in real world applications. For instance, while attacking puzzles like the Rubik's cube, one devises methods for planning the subsequent selection of operators that can be important in practical problems of planning. Likewise, the special clue answering scheme on the Web adopted for discovering good candidates for crossword solving might be of interest in itself and there are plenty of techniques studied for chess solvers that are of practical interest.

On the other hand, most of the mentioned research activities are of interest for their actual entertainment role. Amongst others, the exhibition of humor and story telling ability are changing the face of machines in society, with significant consequences in education and in many other aspects of life.

REFERENCES

[1] S Attardo. *Linguistic Theory of Humor*. Mouton de Gruyter, 2002.

[2] F. Bellotti, E. Ferretti, and A. De Gloria. Discovering the european heritage through the chikho educational web game. In *Proceedings of the INTETAIN*, pages 13–21, December 2005.

[3] K. Binsted and G. Ritchie. Computational rules for punning riddles. *Humor*, 10(1), 1997.

[4] C. Callaway, T. Kuflik, E. Not, A. Novello, O. Stock, and M. Zancanaro. Personal reporting of a museum visit as an entrypoint to future cultural experience. In *Proceedings of IUI-2005, International Conference of Intelligent User Interfaces*, San Diego, 2005.

[5] M. Ernandes, G. Angelini, and M. Gori. Likely-admissible and sub-symbolic heuristics. In *Proceedings of the 16th Eureopean Conference on ArtiÞcial Intelligence, ECAI2004*, 2004.

[6] M. Ernandes, G. Angelini, and M. Gori. Webcrow: a web-based system for crossword solving. In *Proceedings of the sixteenth conference of the American Association for ArtiÞcial Intelligence AAAI2005*, July 2005.

[7] S. Freud. *Jokes and their Relation to the Unconscious*. Norton. Frijda, Nico H., 1905.

[8] D. Goren-Bar, O. Levi, C. Hayes, and P. Avesani. Tagsocratic: Learning shared concepts on the blogosphere. In *Proceedings of the INTETAIN*, pages 324–327, December 2005.

[9] S. Hartley, M. Cavazza, L. Bec, J.-L. Lugrin, and S. Crooks. Simulating artificial organisms with qualitative physiology. In *Advances in ArtiÞcial Life: 8th European Conference, ECAL 2005, Canterbury, UK*, September 5-9, 2005.

[10] W. L. Johnson, J. W. Rickel, and J. C. Lester. Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of ArtiÞcial Intelligence in Education*, 11:47–78, 2000.

[11] B. La Torre, D. Prattichizzo, F. Barbagli, and A. Vicino. The fetouch project. In *Proceedings. ICRA Õ03. IEEE International Conference on Robotics and Automation ICRA Õ03*, pages 1259–1263, 2004.

[12] M. Y. Lim, R. Aylett, and C. M. Jones. Affective guide with attitude. In *First International Conference on Affective Computing and Intelligent Interaction*, Beijing 2005.

[13] M.L. Littman. Review: computer language games. *Computer and Games*, 134:396–404, 2000.

[14] W. Wahlster (Eds.) M. Maybury, O. Stock. *Intelligent Technologies for Interactive Entertainment*. LNAI 3814, Springer, 2005.

[15] R. Mihalcea and C. Strapparava. Computational laughing: Automatic recognition of humorous one-liners. In *Proc. of the 27th Annual Conference of the Cognitive Science Society*, 2005.

[16] News-Nature. Program crosses web to fill in puzzling words. *Nature*, 431:620, 10 2004.

[17] V. Raskin. *Semantic Mechanisms of Humor*. Dordrecht/ Boston/Lancaster, 1985.

[18] R.E. Korf and A. Felner. Disjoint pattern database heuristics. *ArtiÞcial Intelligence*, 132:9–22, 2002.

[19] N.M. Shazeer, G.A. Keim, and M.L. Littman. Solving crosswords as probabilistic contraint satisfaction. In *Proceeding of AAAI Õ99*, pages 156–152, 1999.

[20] O. Stock and M. Zancanaro (Eds). Peach: Intelligent interfaces for museum visits. In *Springer, Cognitive Technologies Series*, Berlin, 2006.

[21] O. Stock and C. Strapparava. Getting serious with the development of computational humor. In *Proceedings of IJCAI-03, the Eighteenth International Joint Conference on ArtiÞcial Intelligence*, 2003.

[22] G. Uccelli, G. Conti, R. De Amicis, and R. Servidio. Learning using augmented reality technology: Multiple means of interaction for teaching children theory of colours. In *Proceedings of the INTETAIN*, pages 193–202, December 2005.

[23] L. von Ahn and L. Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326, 2004.