

Presenting through Performing: On the Use of Multiple Lifelike Characters in Knowledge- Based Presentation Systems

Elisabeth André, Thomas Rist

German Research Center for Artificial Intelligence GmbH
Stuhlsatzenhausweg 3
D-66123 Saarbrücken, Germany
{andre,rist}@dfki.de

ABSTRACT

In this paper, we investigate a new style for presenting information. We introduce the notion of presentation teams which – rather than addressing the user directly – convey information in the style of performances to be observed by him or her. The paper presents an approach to the automated generation of performances which has been tested in two different application scenarios, car sales dialogues and soccer commentary.

Keywords

Presentation teams, animated characters, conversational embodied agents, believable dialogues

INTRODUCTION

Trying to imitate the skills of human presenters, some R&D projects have begun to deploy animated characters (or agents) in knowledge-based presentation systems. A particular strength of animated characters is their ability to express emotions in a believable way by combining facial expressions with body gestures and affective speech. Furthermore, they provide effective means of conveying conversational signals, such as taking turns or awaiting feedback, which are difficult to communicate in traditional media, such as text or graphics alone. Last but not least, results of empirical studies show that animated characters may have a strong motivational impact – many users experience presentations given by animated characters as being more lively and engaging [13,18].

Frequently, systems that use presentation agents rely on settings in which the agent addresses the user directly as it were a face-to-face conversation between human beings. Such a setting seems quite appropriate for a number of applications that draw on a distinguished agent-user relationship. For example, an agent may serve as a personal guide or assistant in information spaces like the world-wide

web (as the AiA Persona [1]), or it can be a user's personal consultant or tutor (as Herman the Bug [13] and Steve [22]), or it may come as a real estate sales person that tries to convince an individual customer (as the REA agent [5]). However, there are also situations in which the emulation of a direct agent-to-user communication is not necessarily the most effective way to present information. First of all, there is no "standard user". Rather, the members of a user population can differ widely with regard to personality and individual preferences for a certain style of acquiring new information. In fact, some people feel less comfortable when being approached directly by an agent.

In this paper, we investigate a new style for presenting information. We introduce the notion of presentation teams which – rather than addressing the user directly – convey information in the style of performances to be observed by him or her. So-called infotainment and edutainment transmissions on TV and also some well-designed advertisement clips are examples that demonstrate how information can be conveyed in an appealing manner by multiple presenters with complementary characters and role castings.

To avoid any misunderstanding, we emphasize that our work is not intended to argue for degrading the user to the role of a passive viewer with the only difference that this time it is not a TV, but a computer screen. Rather, we argue that performances by presentation teams are useful additions to the repertoire of presentation techniques for intelligent presentation systems. In fact, presentation teams can contribute to the success of a presentation with regard to the following aspects:

- Presentation teams can convey certain relationships among information units in a more canonical way. Among other things, this is of benefit in decision support systems where the user has to be informed about different and incompatible points of views, pairs of arguments and counter arguments, or alternative conclusions and suggestions. For solving such presentation tasks, it seems quite natural to structure presentations according to argumentative and rhetorical strategies common in real dialogues with two or more

conversational partners. For instance, a debate between two characters representing contrary opinions is an effective means of providing an audience with the pros and cons of an issue.

- The single members of a presentation team can serve as indices which help the user to classify the conveyed information. For example, in a sales presentation we may present financial and technical aspects strictly separated from each other by a sales assistant and a technician. Or we may indicate the occurrence of important events through a unique notification agent across different applications. Furthermore, different agents can also be used to convey meta-information, such as the origin of an information unit and the reliability of its source. We are not aware of any empirical studies on whether there is a labeling effect on the viewer's recall of information. However, a character has an audio visual appearance. Thus, if a character is recognized as an index, we might expect effects similar to those found in studies supporting the dual encoding theory [21].
- Presentation teams can serve as rhetorical devices that allow for a continuous reinforcement of beliefs. For example, they enable us to repeat one and the same information in a less monotonous manner simply by employing different agents to convey it. Furthermore, arguments may be reinforced by presenting them through agents that are more likely to be accepted by the audience. For example, it is quite common in TV ads to make use of (pseudo-)experts to increase the credibility of the conveyed information.

With regards to recent attempts in the area of collaborative browsing, the use of multiple presenters would also allow for performances that account to a certain extent for the different interest profiles of a diverse audience.

The rest of the paper is organized as follows. The next section discusses related work. After that, we describe the basic steps of our approach to the automated generation of performances with multiple characters. This approach has been applied to two different scenarios: sales dialogues and soccer commentary. Finally, we provide a conclusion and an outlook on future research.

RELATED WORK ON ANIMATED PRESENTERS

A number of research projects has discovered lifelike agents as a new means of computer-based presentation.

Noma and Badler [20] created a virtual human-like weather reporter. Thalmann and Kalra [24] produced some animation sequences for a virtual character acting as a television presenter. The PPP and AiA Personas [1] developed at DFKI operate as desktop assistants or web chauffeurs. However, all these systems employ just one agent for presenting information.

The Agneta & Frida system [8] incorporates narratives into a web environment by placing two characters on the user's

desktop. These characters watch the user during the browsing process and make comments on the visited web pages. In contrast to the approach presented here, the system relies on pre-authored scripts and no generative mechanism is employed. Consequently, the system operates on predefined web pages only.

Cassell and colleagues [4] automatically generate and animate dialogues between a bank teller and a bank employee with appropriate synchronized speech, intonation, facial expressions and hand gestures. However, they do not aim at conveying information from different points of view, but restrict themselves to a question-answering dialogue between the two animated agents.

Mr. Bengo [19] is a system for the resolution of disputes which employs three agents: a judge, a prosecutor and an attorney which is controlled by the user. The prosecutor and the attorney discuss the interpretation of legal rules. Finally, the judge decides on the winner. The virtual agents are able to exhibit some basic emotions, such as anger, sadness and surprise, by means of facial expressions. However, they do not rely on any other means, such as linguistic style, to convey personality or emotions.

Hayes-Roth and colleagues [7] have implemented several scenarios following the metaphor of a virtual theatre. Their characters are not directly associated with a specific personality. Instead, they are assigned a role and have to express a personality which is in agreement with this role. A key concept of their approach is improvisation. That is characters spontaneously and cooperatively work out the details of a story at performance time taking into account the constraints of directions either coming from the system or a human user. Even though the main focus of the work by Hayes-Roth and colleagues was not the communication of information by means of performances, the metaphor of a virtual theatre can be employed in presentation scenarios as well.

DESIGNING PRESENTATION DIALOGUES: BASIC STEPS

Our approach is based on the observation that vivid and believable dialogues are in fact a means to present information to an audience. Given a certain discourse purpose and a set of information units to be presented, we have to determine an appropriate dialogue type, define roles for the characters to be involved, recruit concrete characters with personality profiles that go together with the assigned roles, and finally, work out the details of the single dialogue turns and have them performed by the characters.

Dialogue Types and Character Roles

The structure of a performance is predetermined by the choice of the dialogue type which depends on the overall presentation goal. In this paper, we restrict ourselves to sales dialogues and chats about jointly watched events. Once a certain dialogue type has been chosen, we need to define the roles to be occupied by the characters. Most dialogue types induce certain constraints on the required

roles. For instance, in a debate on a certain subject matter, there is at least a proponent and an opponent role to be filled. In a sales scenario, we need at least a seller and a customer.

The next step is the occupation of the designated roles with appropriate characters. To generate effective performances, we cannot simply multiply an existing character. Rather, characters have to be realized as distinguishable individuals with their own areas of expertise, interest profiles, personalities and audio/visual appearance taking into account their specific task in a given context.

An agent's personality is represented by a vector of discrete values along a number of psychological traits, such as extraversion, openness or agreeableness, that uniquely characterize an individual. Personality traits are not influenced by current events, but remain stable over a longer period of time. Closely related to personality is the concept of emotion. In contrast to personality, emotions are short-lived and are influenced by the character's current situation [17]. The intensity of an emotion strongly depends on the character's personality. For instance, a hot-tempered soccer fan is likely to get more angry if its team misses a goal chance than a more balanced character. A further important component of a character's profile is its audio/visual appearance. We start from a given set of characters and basic gestures which are either freely available or have been designed by a professional artist with a specific presentation task in mind. In our example applications, we offer the user the possibility to select a team of presenters from this set and assign roles and personalities to them. Unlike roles and personalities, emotions are automatically computed considering the character's momentary situation at runtime.

Generation of Dialogue Contributions

After a team of presenters has been recruited by the user, our system automatically generates a performance. Following a speech-act theoretic view, we represent simulated dialogues as a sequence of communicative acts to achieve certain goals. To automatically generate such dialogues, we are investigating the following two approaches:

- *Actors with Scripted Behaviors*

In this approach, the system appears in the role of a producer which generates a script for the actors of a play. The script specifies the dialogue acts to be carried out as well as their temporal coordination. The approach facilitates the generation of coherent dialogues since the script writer completely controls the structure of the dialogue. However, it requires that the knowledge to be communicated is a priori known. From a technical point of view, this approach may be realized by a central planning component which decomposes a complex presentation goal into elementary dialogue acts which are then allocated to the single agents. Knowledge concerning the

decomposition process is then realized by operators of the planning component.

- *Autonomous Actors*

In this approach, the single agents will be assigned a set of communicative goals which they try to achieve. That is both the determination and assignment of dialogue contributions is handled by the agents themselves. To accomplish this task, each agent has a repertoire of dialogue strategies at its disposal. However, since the agents have only limited knowledge concerning what other agents may do or say next, this approach puts much higher demands on the agents' reactive capabilities. Furthermore, it is much more difficult to ensure the coherence of the dialogue. Think of two people giving a talk together without clarifying in advance who is going to explain what. From a technical point of view, this approach may be realized by assigning each agent its own reactive planner. The agents' dialogue strategies are then realized as operators of the single planners.

Depending on their role and personality, characters may pursue completely different goals. For instance, a customer in a sales situation usually tries to get information on a certain product in order to make a decision while the seller aims at presenting this product in a positive light. To generate believable dialogues, we have to ensure that the assigned dialogue contributions do not conflict with the character's goal. Characters differ not only with respect to their communicative goals, but also with respect to their communicative behavior. Depending on their personality and emotions, they may apply completely different dialogue strategies. For instance, a shy agent will less likely take the initiative in a dialogue and exhibit a more passive behavior. Finally, what an agent is able to say depends on its area of expertise. Both planning approaches allow us to consider the characters' profile by treating it as an additional constraint during the selection and instantiation of dialogue strategies.

Even if the agents have to strictly follow a script as in the script-based approach, there is still enough room for improvisation at performance time. In particular, a script leaves open how to render the dialogue contributions to make. Agents with a different personality should not only differ in their high-level dialogue behaviors, but also perform elementary dialogue acts in a character-specific way. Furthermore, the rendering of dialogue acts depends on an agent's emotional state. Important means of conveying an agent's personality and emotions are verbal and acoustic realization, facial expressions and body gestures (see [6] for an overview of empirical studies on emotive expression). To consider such parameters, the planner(s) enhance the input of the animation module and the speech synthesizer with additional instructions, e.g. in an XML-based mark-up language.

INHABITED MARKET PLACE

As a first example, we address the generation of animated sales dialogues. For the graphical realization of this scenario, we use the Microsoft Agent™ package [16] that includes a programmable interface to four predefined characters: *Genie*, *Robby*, *Peedy* and *Merlin*.

Fig. 1 shows a dialogue between Merlin as a car seller and Genie and Robby as buyers. Genie has uttered some concerns about the high running costs which Merlin tries to play down. From the point of view of the system, the presentation goal is to provide the observer – who is assumed to be the real customer - with facts about a certain car. However, the presentation is not just a mere enumeration of the plain facts about the car. Rather, the facts are presented along with an evaluation under consideration of the observer's interest profile. This scenario was inspired by Jameson and colleagues [11] who developed a dialogue system which models non-cooperative dialogues between a car seller and a buyer. However, while the objective of Jameson and colleagues is the generation of dialogue contributions which meet the goals of the system which may either take on the role of the seller or the buyer, our focus is on the development of animated agents that convey information by giving performances.



Fig. 1: Screenshot of the Inhabited Market Place

Character Profiles

To support experiments with different character settings, the user has the possibility of choosing three out of the four characters and assigning roles to them. For instance, he or she may have Merlin appear in the role of a seller or buyer. Furthermore, he or she may ascribe to each character certain preferences and interests (see Fig. 2). Personality traits may be set by the user as well. We have decided to model the following two personality factors:

- *Extraversion* with the possible values: *extravert*, *normal* or *introvert*
- *Agreeableness* with the possible values: *agreeable*, *neutral* or *disagreeable*

In the first version of the sales scenario, we decided just to model one dimension of emotional response: *valence* with the possible values *positive*, *neutral* and *negative*. Emotions are triggered by the state of goal achievement. For instance, an agent that wants to present a product in a positive light, will be satisfied if it is asked a question on a attribute with a favorable extension. Our characters do not lie in the sense that they exhibit emotions which they do not have (even though this might be quite common in sales scenarios).



Fig. 2: Role Casting Interface for the Car Sales Scenario

Source, Structure and Representation of the Information to be Communicated

Part of the domain knowledge is an ordinary product database, e.g., organized in the form of an n-dimensional attribute vector per product. In our current scenario, the products are cars with attributes, such as model type, maximum speed, horsepower, fuel consumption, price, air conditioning, electric window lifters, airbag type etc. Thus, to a large extent, the contents of the database determines what an agent can say about a product. However, products and their attributes are described in a technical language with which the user may not be familiar with. Therefore, it seems much more appropriate to maintain a further description of the products - one that reflects the impact of the product attributes on the value dimensions of potential customers. Such an approach can be modeled in the framework of multi-attribute utility theory (e.g. see [27]), and has already been used for the identification of customer profiles in an electronic bourse for used cars [15]. In this project, the car database was provided from a large German/American car producer and retailer, whereas the value dimensions for the product "car" have been adopted from a study of the German car market [23] that suggests that safety, economy, comfort, sportiness, prestige, family and environmental friendliness are the most relevant. In addition, it was represented how difficult it is to infer such implications. The work presented here follows this approach even though we employ a simplified model. For instance, we use the expressions:

```
FACT value "ccar1" 8;
FACT polarity "ccar1" "environment" "neg";
FACT difficulty "ccar1" "environment" "low";
```

to represent that a certain car consumes 8 liters, that this fact has a negative impact on the dimension "environment" and this implication is not difficult to infer.

Design of Product Information Dialogues

To automatically generate product information dialogues, we use a central planning component which decomposes a complex goal into more elementary goals. The result of this process is a dialogue script that represents the elementary dialogue acts to be executed by the single agents as well as their temporal order. Dialogue acts include not only the propositional contents of an utterance, but also its communicative function, such as taking turns or responding to a question. This is in line with [5] who regard conversational behaviors as fulfilling propositional and interactional conversational functions.

Knowledge concerning the generation of scripts is represented by means of plan operators. In the sales scenario, plan operators that implement argumentative strategies play a central role. There has been a great deal of work on argumentation ranging from formal models of argument structure, such as Toulmin's classical work [25], to generative approaches, such as [14] and [28]. Our work differs from these approaches in that it does not just generate arguments for a single agent, but allocates the parts of an argumentative discourse to a team of presenters. Consequently, our plan operators do not only handle the specification of dialogue acts, but also the distribution of these acts onto the individual agents. An example of a plan operator is listed in Fig. 3.

```

NAME: "DiscussValue1"
GOAL: PERFORM DiscussValue $attribute;
PRECONDITION:
FACT polarity $attribute $dimension "neg";
FACT difficulty $attribute $dimension "low";
FACT Buyer $buyer;
FACT Disagreeable $buyer;
FACT Seller $seller;
BODY:
PERFORM NegativeResp $buyer $dimension;
PERFORM RespNegResp $seller $attribute $dimension;

```

Fig. 3: Example of a plan operator for discussion an attribute value

The operator represents a scenario where two agents discuss a feature of an object. It only applies if the feature has a negative impact on any dimension and if this relationship can be easily inferred. According to the operator, any disagreeable buyer produces a negative comment referring to this dimension (*NegativeResp*). The negative comment is followed by a response from the seller (*RespNegResp*).

When defining dialogue strategies for the sales scenario, we implicitly started from the assumption that the single agents collaborate with each other in order to achieve a common goal, namely to provide information on a certain product. Nevertheless, the applied methodology is general enough to allow for the synthesis of non-cooperative dialogues in which one agent e.g. refuses to provide an answer to a question.

The implementation of the planning approach is based on the Java-based JAM Agents architecture framework [10].

The outcome of the planning process is an HTML file that includes control sequences for the Microsoft Agents. The performances can be played in the Microsoft Internet Explorer.

What about this car? Two Generation Examples

In the following, we present a short dialogue fragment to illustrate how the agents' personality and interest profiles influence the contents and the structure of the sales dialogue. We use extreme parameter settings for the agents' personality traits and interest profiles in order to demonstrate the differences in the agents' behavior.

Agent	Role	Personality factors	Interests
Robby	seller	extravert, agreeable	sportiness
Peedy	buyer	introvert, disagreeable	environment
Merlin	buyer	extravert, agreeable	safety

Robby: *Hello, I'm Robby. What can I do for you?*

;;; starts the conversation because it is extravert

Merlin: *We are interested in this car.*

;;; responds to the question because it is extravert

Robby: *This is a very sporty car. It can drive 100 miles per hour.*

;;; emphasizes a dimension which is important to him and

;;; mentions an attribute which has a positive impact on this

;;; dimension

Merlin: *Does it have airbags?*

;;; starts asking questions because it is extravert

;;; wants to know whether the car got airbags, since this has

;;; an impact on safety which is important to him

Robby: *Sure.*

;;; retrieves the value of the attribute airbags from the data

;;; database

Merlin: *Excellent.*

;;; positive evaluation because it is agreeable, powerful

;;; language because it is extravert

...

Peedy: *How much gas does it consume?*

;;; gas consumption has an impact on the environment

Robby: *It consumes 8 liters per 100 km.*

Peedy: *Isn't that bad for the environment?*

;;; negative comment because it is disagreeable, less

;;; direct speech because it is introvert

Robby: *Bad for the environment? It has a catalytic converter. It is made of recyclable material.*

;;; questions the negative impact and provides counter

;;; arguments

...

The dialogues are based on just a few dialogue strategies. Essentially, each agent asks after the values of features which might have any impact – positive or negative – on a dimension it is interested in. After that, the value of this attribute is discussed. The dialogue terminates after all relevant attributes of the car under consideration have been discussed.

GERD & MATZE COMMENTATING ROBOCUP SOCCER GAMES

The second application for our work on multiple presentation agents is Rocco II, an automated live report system for the simulator league of RoboCup, the Robot World-Cup Soccer. Fig. 4 shows a screenshot of the system which was taken during a typical session. In the upper window, a previously recorded game is played back while being commented by two soccer fans: Gerd and Matze sitting on a sofa. Unlike the agents of our sales scenario, Gerd and Matze have been specifically designed for soccer commentary. Furthermore, this application is based on our own Java-based Persona Engine [1].



Fig. 4: Commentator Team Gerd & Matze

Character Profiles

Apart from being smokers and beer drinkers, Gerd and Matze are characterized by their sympathy for a certain team, their level of extraversion (extravert, neutral, or introvert) and openness (open, neutral, not open). As in the previous application, these values may be interactively changed. We decided to focus on two emotional dispositions which are characteristic of the soccer domain: *Arousal* with the values *calm*, *neutral* and *excited* and *Valence* with values *positive*, *neutral* and *negative*. Emotions are influenced by the current state of the game. For instance, both agents get excited if the ball approaches one of the goals and calm down in phases of little activity. An agent gets enthusiastic if the team it supports performs a successful action and disappointed if it fails.

Source, Structure and Representation of the Information to be Communicated

Rocco II concentrates on the RoboCup simulator league, which involves software agents only (as opposed to the real robot leagues). Thus, the soccer games to be commented are not observed visually. Rather, the system obtains its basic input from the Soccer Server [12] which delivers: player location and orientation (for all players), ball location and game score and play modes (such as throw-ins, goal kicks, etc.). Based on these data, Rocco's incremental event recognition component performs a higher level analysis of the scene in order to recognize conceptual units at a higher level of abstractions, such as spatial relations or typical motion patterns. The interpretation results of the time-varying scene together with the original input data provide the required basic material for Gerd's and Matze's commentary [2].

Generation of Live Reports for Commentator Teams

Unlike the agents in the car sales scenario, Gerd and Matze have been realized as (semi-) autonomous agents. That is each agent is triggered by events occurring in the scene or by dialogue contributions of the other agent.

For the generation of natural-language, we rely on a parameterized template-based generator. To obtain a rich repertoire of templates, 13.5 hours of TV soccer reports in English have been transcribed from which we manually extracted about 300 basic templates. Each template was annotated with the following linguistic features: *Verbosity* referring to the length of a pattern, *Specificity* referring to the degree of detail the template provides, *Force* with the values: powerful, normal and hesitant, *Floridity* with the values: dry, normal and flowery, *Formality* with the values: formal, colloquial and slang and *Bias* with the values: negative, neutral and positive. The choice of the features has been inspired by Hovy [9] who presents one of the first approaches to natural language generation that also considers social factors, such as the relationship between the speaker and the hearer, when producing an utterance. To select among several applicable templates, we apply a four-phase filtering process. Only the best templates of each filtering phase will be considered for the next evaluation step. The first filtering phase tries to accommodate for the specific needs of a real-time live report. If time pressure is high, only short templates will pass this filtering phase where more specific templates will be given preference over less specific ones. In the second phase, templates which have been used only recently will be eliminated in order to avoid monotonous repetitions. The third phase serves to communicate the speaker's attitude. If the speaker is strongly in favor of a certain team, templates with a positive bias will be preferred for describing the activities of this team. The fourth phase finally considers the speakers' personality. For instance, forceful language is used for extravert commentators, flowery language for open commentators which are characterized as being creative and imaginative.

Another important means of conveying personality is acoustic realization. We have not yet addressed this issue, but simply designed two voices which may be easily distinguished by the user. Acoustic realization is, however, used for the expression of emotions. Drawing upon Cahn's pioneering work [3], we have been examining how we can generate affective speech by parameterizing the TrueTalk™ speech synthesizer. Currently, we mainly vary pitch accent, pitch range and speed. For instance, excitement is expressed by a higher talking speed and pitch range. Unfortunately, the TrueTalk™ speech synthesizer only allows for setting very few parameters. Consequently, we cannot only simulate a small subset of the effects investigated by Cahn.

Kasuga against Andhill Commented by Gerd & Matze

In the car sales example, personality is essentially conveyed by the choice of dialogue acts. Gerd & Matze portray their personality and emotions essentially by body gestures and linguistic style which refers to the semantic content, the syntactic form and the acoustic realization of an utterance [26]. In the first version of Rocco II, each commentator concentrates on the activities of a certain team. That is there is an implicit agreement between the characters concerning the distribution of dialogue contributions. Responses to the dialogue contributions of the other commentator are possible provided that the speed of the game allows for it. Furthermore, the commentators may provide background information on the game and the involved teams. This information is simply retrieved from a database. We present a protocol of a system run with the following parameter settings:

Agent	Attitude	Personality factors
Gerd	in favor of team Kasuga	extravert, open
Matze	neutral	introvert, not open

Gerd: *Kasuga kicks off*

;;; recognized event: kick off

Matze: *Andhill 5*

;;; recognized event: ball possession, time pressure

Gerd: *We're live from an exciting game, team Andhill in red versus Kasuga in yellow*

;;; time for background information

Matze: *Now Andhill 9*

;;; recognized event: ball possession

Gerd: *Super interception by yellow 4*

;;; recognized event: loss of ball, attitude: pro Kasuga,

;;; forceful language because it is extravert

still number 4

;;; recognized event: ball possession, number 4 is

;;; topicalized

Matze: *Andhill 9 is arriving*

;;; recognized event: approach

Gerd: *ball hacked away by Kasuga 4*

;;; recognized event: shot, flowery language since it is

;;; creative

...

CONCLUSION

In this paper, we proposed performances given by a team of characters as a new form of presentation. The basic idea is to communicate information by means of simulated dialogues that are observed by an audience. We have investigated these issues in two different application scenarios and implemented demonstrator systems for each of them. In the first application, a sales scenario, the dialogue contributions of the involved characters are pre-determined by a script. Since the knowledge to be communicated was a priori stored in a knowledge base, this approach seemed adequate. In contrast, the characters in the soccer scenario have to respond immediately to a rapidly changing environment. Therefore, we decided to realize them as (semi-)autonomous agents. A main feature of our presentations is that the characters do not only communicate the plain facts about a certain subject matter, but present them from a point of view that reflects their specific personality traits and interest profiles. Consequently, our presentations do not only depend on the knowledge that is to be communicated, but also on who presents it.

The purpose of our demonstration systems was not to implement a more or less complete model of personality for characters, such as a seller, a customer or a soccer fan. Rather, the systems have been designed as test beds that allow for experiments with various personalities and roles. First informal system tests were encouraging. Even though it was not our intention to make use of humor as the authors of the Agneta & Frida system, people found both scenarios entertaining and amusing. Furthermore, people were very eager about to test various role castings in order to find out which effect this would have on the generated presentations. These observations suggest that people possibly learn more about a subject matter because they are willing to spend more time with a system. In the future, we will concentrate on more formal evaluations in order to shed light on questions, such as: What is the optimal number of roles and what should an optimal casting look like? Furthermore, we would like to investigate how to actively involve humans in a presentation – either as co-presenters that are assisted by an animated presenter or as part of the audience that is allowed to provide feedback during a performance.

ACKNOWLEDGMENTS

We would like to thank Martin Klesen for reformulating the presentation strategies of the car sales scenario as plans in the JAM framework, and for his implementation of the interface. We are also grateful to Marc Huber for his support with the JAM framework. We thank Stephan Baldes for his engagement in the development of the Rocco II system. Many thanks also to Sabine Bein and Peter Rist for providing us with the Gerd & Matze cartoons. Finally, we would like to thank Justine Cassell and the five reviewers for their valuable comments.

REFERENCES

1. André, E., Rist, T., and Müller, J. Employing AI Methods to Control the Behavior of Animated Interface Agents. *Applied Artificial Intelligence* 13:415-448, 1999.
2. André, E., Binsted, K., Tanaka-Ishii, K., Luke, S., Herzog, G., and Rist, T. Three RoboCup Simulation League Commentator Systems. *AI Magazine*, to appear.
3. Cahn, J. The Generation of Affect in Synthesized Speech. *Journal of the American Voice I/O Society*, 8:1-19, 1990.
4. Cassell, J., Pelachaud, C., Badler, N.I., Steedman, M., Achorn, B., Becket, T., Douville, B., Prevost, S., and Stone, M. Animated Conversation: Rule-Based Generation of Facial Expression, Gesture and Spoken Intonation for Multiple Conversational Agents, *Computer Graphics*, 28(4):413-420, 1994.
5. Cassell, J., Bickmore, T., Camphell, L., Vilhjalmsson, H., and Yan, H. The Human Conversation as a System Framework: Designing Embodied Conversational Agents. in Cassell, J., Sullivan, J., Prevost, S., and Churchill, E. (eds.): *Embodied Conversational Agents*, MIT Press: Cambridge, MA, to appear.
6. Collier, G. *Emotional Expression*. Hillsdale, N.J.: Lawrence Erlbaum, 1985.
7. Hayes-Roth, B., van Gent, R., and Huber, D. Acting in Character. in Trapp, R., and Petta, P. (eds.): *Creating Personalities for Synthetic Actors*, 92-112, Springer, Berlin, Heidelberg, New York, 1997.
8. Höök, K., Sjölander, M., Ereback, A.-L., and Persson, P. Dealing with the Lurking Lutheran View on Interfaces: Evaluation of the Agneta & Frida System. in *Proc. of the i3 Spring Days Workshop on Behavior Planning for Lifelike Characters and Avatars*, 125-136, Sitges, Spain, 1999.
9. Hovy, E. Some pragmatic decision criteria in generation. in Kempen, G. (ed.): *Natural Language Generation*, 3-17, Martinus Nijhoff Publishers: Dordrecht, Bosten, Lancaster, 1987.
10. Huber, M. JAM: A BDI-theoretic Mobile Agent Architecture. in: *Proc. of Autonomous Agents '99*, 236-243, 1999.
11. Jameson, A., Schäfer, R., Simons, J., and Weis, T. Adaptive Provision of Evaluation-Oriented Information: Tasks and Techniques. in *Proc. of IJCAI '95*, 1886-1893, Morgan Kaufmann: San Mateo, CA, 1995.
12. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., and Matsubara, H. RoboCup: A Challenging Problem for AI. *AI Magazine* 18(1):73-85, 1997.
13. Lester, J., Converse, S.A., Stone, B.A., and Kahler, S.E. Animated Pedagogical Agents and Problem-Solving Effectiveness: A Large-Scale Empirical Evaluation. *Artificial Intelligence in Education*, IOS Press: Amsterdam, 23-30, 1999.
14. Maybury, M. Communicative Acts for Generating Natural Language Arguments. in *Proc. of AAAI-93*, 357-364, Washington, DC, 1993.
15. Mehlmann, O., Landvogt, J., Jameson, A., Rist, T., and Schäfer, R. Einsatz Bayes'scher Netze zur Identifikation von Kundenwünschen im Internet, *Künstliche Intelligenz* 3(98):43-48, 1998.
16. *Microsoft Agent: Software Development Kit*, Microsoft Press, Redmond Washington, 1999.
17. Moffat, D. Personality Parameters and Programs. in Trapp, R., and Petta, P. (eds.): *Creating Personalities for Synthetic Actors*, 120-165, Springer: Berlin, 1997.
18. van Mulken, S., André, E., and Müller, J. The Persona Effect: How Substantial is it? in *Proc. of HCI'98*, Sheffield, England, 53-66, 1998.
19. Nitta, K., Hasegawa, O., Akiba, T., Kamishima, T., Kurita, T., Hayamizu, S., Itoh, K., Ishizuka, M., Dohi, H., and Okamura, M. An Experimental Multimodal Disputation System. in *Proc. of IJCAI '97 Workshop on Intelligent Multimodal Systems*, 1997.
20. Noma, T., and Badler, N. A Virtual Human Presenter. in: *Proc. of IJCAI'97 Workshop on Animated Interface Agents: Making them Intelligent*, 45-51, 1997.
21. Paivio, A. The empirical case for dual coding. in Yuille, J.C. (ed.): *Imagery, memory and cognition*, 307-332, Erlbaum: Hillsdale, 1983.
22. Rickel, J. and Johnson, L. Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control. *Applied Artificial Intelligence* 13:343-382, 1999.
23. Spiegel-Verlag, *SPIEGEL-Dokumentation: Auto, Verkehr und Umwelt*, Augstein: Hamburg, 1993.
24. Thalmann, N.M. and Kalra, P. The Simulation of a Virtual TV Presenter. *Computer Graphics and Applications*, 9-21. World Scientific Press: Singapore, 1995.
25. Toulmin, S. *The Uses of Argument*. Cambridge University Press: Cambridge, England, 1958.
26. Walker, M., Cahn, J., and Whittaker, S.J. Improving Linguistic Style: Social and Affective Bases for Agent Personality. in *Proc. of Autonomous Agents '97*, Marina del Rey, 96-105, 1997.
27. Winterfeldt, D., and Edwards, W. *Decision Analysis and Behavioral Research*. Cambridge University Press: Cambridge, England, 1986.
28. Zukerman, I., McConachy, R., and K.B. Korb, K.B. Attention During Argument Generation and Presentation. in *Proc. of the 9th International Workshop on Natural Language Generation*, 148-157, Niagara-on-the-Lake, Canada, 1998.