

# Domain-Oriented Conversation with H. C. Andersen

Niels Ole Bernsen and Laila Dybkjær

NISLab, University of Southern Denmark, Odense  
nob@nis.sdu.dk, laila@nis.sdu.dk

**Abstract.** This paper describes the first running prototype of a system for domain-oriented spoken conversation with life-like animated fairy tale author Hans Christian Andersen. Following a brief description of the system architecture, we present our approach to the highly interrelated issues of making Andersen life-like, capable of domain-oriented conversation, and affective. The paper concludes with a brief report on the recently completed first user test.

## 1 Introduction

In recent years, animated interface agents have become a sub-specialty among developers of multimodal dialogue systems. The community is growing fast, as witnessed by, e.g., the large attendance at the 2003 Intelligent Virtual Agents workshop [9]. Basically, animated interface agents are characterised by the on-screen display of a more or less human-like animated face. Some animated interface agents are embodied as well, and distinction may be made between cartoon (face or embodied) and life-like (face or embodied) human interface agents, depending on the level of rendering realism. To most users, human animated interface agents are not very expressive without the use of output speech. The talking face of news reader Ananova illustrates a basic animated interface agent who speaks [1]. When today's animated interface agents are interactive, this is primarily accomplished through spoken input and output, essentially turning the agent into a multimodal spoken dialogue system-cum-animated human output graphics. These agents are often called conversational interface agents despite the fact that the large majority of them are still task-oriented [5] and hence, as argued below, *not* conversational in an important sense of this term. Increasingly, future animated interface agents will be able to interpret both verbal and non-verbal input communication, they will gradually become more life-like, and they will go beyond the level of task-oriented systems.

This paper describes the first prototype of a non-task-oriented, life-like animated agent system which also understands 2D gesture input. The prototype has been developed in the NICE (Natural Interactive Conversation for Edutainment) project [8]. Work in NICE aims to demonstrate English domain-oriented conversation with fairy tale author Hans Christian Andersen (HCA) in his study and Swedish spoken computer game-style interaction with some of his fairy tale characters in the adjacent fairy tale world. The present paper addresses the former goal. We propose the term *domain-oriented conversation* to designate a half-way post between task-oriented spoken dialogue [4, 6], and Turing-test compliant conversation [10]. In domain-

oriented conversation, the system is able to conduct unconstrained conversation about topics within its knowledge domain(s). The target users of the HCA system are 10-18 years old kids. The primary use setting is in museums and other public locations where interactions are expected to have an average duration of 10-15 minutes.

In the following, Section 2 discusses an influential empirical generalisation on user perception of agent life-likeness. Section 3 outlines the general architecture of the HCA system and of the HCA character module. Section 4 describes our key conversation strategies for making HCA a new kind of believable virtual semi-human. Section 5 focuses on one particular such strategy, i.e. emotional HCA. Section 6 briefly reports on the first user test of the HCA system. Section 7 concludes the paper.

## **2 The need for new classes of agents**

Although we are still in the dark in most areas, empirical generalisations are beginning to emerge from evaluations of interactive animated interface agents. One finding is that there seems to exist a “user-system togetherness” problem zone separating two generic classes of agents. Due to the primitive nature of their interactive behaviour, some agents are so different from their human interlocutors that they are (almost) invariably perceived as systems rather than humans. This class includes, among others, simple task-oriented unimodal spoken dialogue systems speaking with a “computer voice” [11], primitive cartoon-style agents and other not very life-like agents [5]. However, as graphical life-likeness, conversational abilities, and/or persona expressiveness improve, users appear to start forming unconscious expectations to the effect that they are facing a system with human-like capabilities. If these expectations are thwarted, as they mostly are with today’s interactive agents, frustration results. The user actually believed to be together with another human but wasn’t. The message for interactive animated interface agent research seems to be to find ways to safely pass beyond the problem zone by building interactive agents which no longer frustrate their users but, rather, constitute entirely new kinds of believable virtual semi-humans. Some of the means towards this goal are: to endow interactive agents not only with life-like graphical quality and domain-oriented conversation but also with non-stereotypical personalities, personal agendas and consistent emotional behaviour. Our aim is for HCA to become such a character, or agent.

## **3 The NICE HCA system**

Two important goals in developing the HCA system are to investigate (i) how to successfully integrate spoken interaction with gesture input and non-verbal animated character output, and (ii) the use of spoken conversation for education and entertainment. The key goal, however, is to (iii) investigate non-task-oriented spoken conversation in a potentially realistic application. Arguably, the achievement of those goals requires a new kind of “self-reliant” animated conversational agents which no longer cause user frustration (Section 2).

We have developed a first prototype of the HCA system which was tested with target group users in January 2004. The prototype is running with simulated recognition. The recogniser still needs to be trained on large amounts of language data from non-native English speaking children and will be included in the second prototype.

Figure 3.1 shows the architecture of the HCA prototype which is described in more detail in [3]. NISLab is responsible for HCA's natural language understanding, character modelling and response generation functionalities. The other components in Figure 3.1 are being developed by other NICE project partners or are (based on) freeware (gesture recognition, message broker and speech synthesis). The project partners are TeliaSonera, Sweden, Liquid Media, Sweden, Scansoft, Belgium, and LIMSI, France.

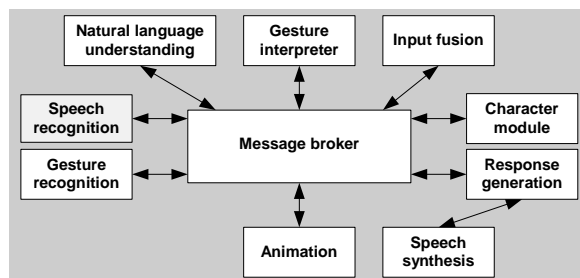


Figure 3.1. General NICE HCA system architecture.

The focus in this paper is on the HCA character module which is responsible for conversation management. Figure 3.2 shows the architecture of this module.

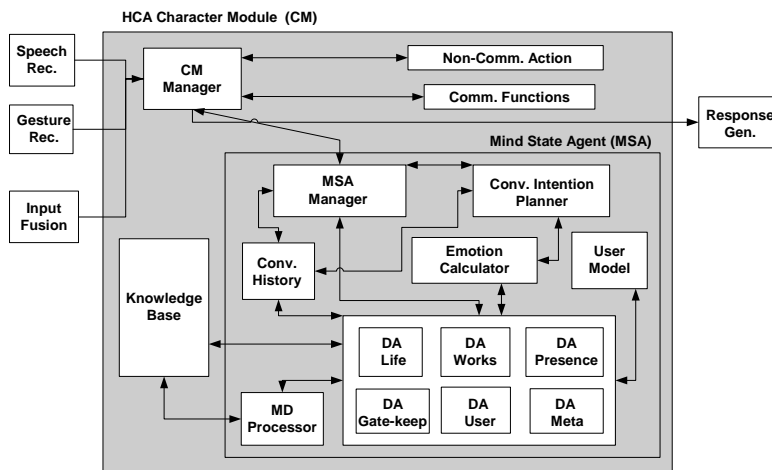


Figure 3.2. HCA character module. DA is domain agent. MD is mini-dialogue.

The character module is always in one of three output states, producing either *non-communicative action* output when HCA is alone in his study, *communicative function* output when HCA is listening, or paying attention, to a visitor's contribution to the conversation, or *communicative action* when HCA produces a conversational contribution, cf. the mind state agent in Figure 3.2 and Section 4.4.

The mind-state agent generates HCA's conversational contributions, managing HCA's conversational agenda, interpreting the user's spoken and/or gesture input in context, deciding on conversation initiative, and planning HCA's verbal and non-verbal output. The conversational intention planner applies HCA's conversational agenda to the user's current input and keeps track of agenda achievement (see Section 4.5). Six domain agents (DAs), one per knowledge domain, take care of domain-specific reasoning, including meta-communication, and user model maintenance (Section 4.3). The emotion calculator updates HCA's emotional state (Section 5).

Mind-state agent processing is supported by three additional modules. The conversation history stores a representation of the emerging discourse context for consultation by other mind-state agent modules. The knowledge base maintains the system's ontology, including references to HCA output. Finally, the finite-state machine mini-dialogue (MD) processor processes all user-HCA mini-dialogues, i.e. predefined small dialogues of the kind familiar from task-oriented systems. The output references retrieved from the knowledge base are sent to response generation via the mind state agent manager and the character module manager.

## **4 Life-likeness and conversation**

In view of the discussion in Section 2, virtual HCA should not, on the one hand, pose as the real HCA, nor, on the other, should the character be trapped in the "togetherness" problem zone in which interactive agents frustrate their users. To address this challenge, the first HCA prototype uses strategies such as the following: (i) a cover story, (ii) life-like output graphics, (iii) life-like domains of discourse, (iv) life-like in-and-out-of-conversation behaviour, (v) a conversation agenda, (vi) conversational principles, (vii) error handling, and (viii) emotional behaviour. (i) through (vii) are discussed in the present section, (viii) forms the subject of Section 5.

### **4.1 Cover story**

The cover story for HCA's limited knowledge about his domains of conversation is that HCA is coming back! However, he still has to re-learn much of what he once knew. If the user would do him the favour of visiting him later, he is convinced that he will have become much more of what he once was. In addition to the very true information provided by this cover story, the story may help convince users that HCA is not (yet) a full virtual person. It may be added that HCA does not tell the cover story up front to new users. Rather, users are likely to come across the cover story if they either explicitly ask what HCA knows about, or can do, or if they show too much interest in things he does not know about (yet).

### **4.2 Life-like output graphics**

The HCA computer graphics has been developed by Swedish computer games company Liquid Media. Figure 4.1 shows 55-year old HCA surrounded by artefacts

in his study. Users can use gesture and speech to indicate an artefact which HCA might want to tell a story about. The study is a rendering of HCA's study on display in Copenhagen, modified so that he can walk around freely and so that a pair of doors lead into the fairy tale games world (cf. Section 1). Also, pictures clearly relating to HCA's knowledge domains have been hung on the walls.



**Figure 4.1.** HCA in his study.

### 4.3 Life-like domains of discourse

Development of domain-oriented conversation requires selection of one or several knowledge domains for the character. In the first NICE prototype, HCA's knowledge domains are: his fairy tales (works), his childhood in Odense (life), his physical presence in his study (presence), getting information about the user (user), his role as "gatekeeper" for the fairy tale games world (gatekeeper), and the "meta" domain of resolving problems of miscommunication (meta). These domains are probably those which most users would expect anyway.

Since we want to explore the domain development challenges "breadth-first" in order to investigate, among other things, how to handle potential cross-domain, multiple-domain, and super-domain issues, none of those domains have been developed to their full depth in the first prototype. For instance, HCA only has in-depth knowledge of three of his most famous fairy tales, the Little Mermaid, the Ugly Duckling, and the Princess and the Pea. If a user asks about some other fairy tale, the user is told some version of HCA's cover story.

HCA has two mechanisms for in-depth conversation. The fairy tales are stored in template-style fashion in the knowledge base, enabling HCA to tell stories about, e.g., the main character in some fairy tale or the morale of a particular fairy tale. Mini-dialogues are used for structured, in-depth conversation about some topic, such as game-playing. HCA will show interest in games played by kids and adolescents today and he understands terms for games he is interested in, such as 'computer games' and 'football'. HCA also conducts a mini-dialogue in order to gather knowledge about the present user. The knowledge HCA collects about the user is stored in the user model for use during conversation (Figure 3.2).

#### 4.4 Life-like in-and-out-of-conversation behaviour

As shown in Figure 3.2, HCA behaves as a human dedicated to fairy tale authoring both when he is alone in his study, when paying attention to the user's speech and/or gesture input, and when producing natural interactive output. In the non-communicative action output state, HCA goes about his work in his study as displayed through a quasi-endless loop of micro-varied behaviours. We cannot have him walk around on his own yet, however, because he may walk through the walls and crash the system due the not-fully-debugged graphics rendering. In the communicative functions output state, HCA pays attention to the user's speech and/or gesture input through conversational recipient behaviours, such as looking at the user, nodding, etc. For this to happen in real time, the character module will soon have fast-track connections to the speech recogniser and the gesture recogniser to be able to act as soon as one or both of them receive input. In the communicative action output state, HCA responds to input through verbal and non-verbal communicative action [2].

#### 4.5 Conversation agenda

HCA follows his own agenda during conversation. The agenda reflects his personal interests, e.g. his interest in collecting knowledge about the user and in having a good long conversation with motivated users. The agenda ensures some amount of conversational continuity on HCA's part, making sure that a domain is pretty thoroughly addressed before moving to another, unless the user changes domain and is allowed to do so by HCA. HCA makes sure by keeping track of what has been addressed in each domain so far, which also helps him avoid repeating himself. Also, since many users are likely to leave HCA's study when learning that the double doors lead to the fairy tale world, HCA is reluctant to embark on the "gatekeeper" domain until the other domains have been addressed to quite some extent. If a user embarks on "gatekeeper" too early, HCA changes the domain of conversation.

#### 4.6 Conversational principles

Conversation, properly so-called, is very different from task-oriented dialogue. In addressing conversation, the seasoned spoken dialogue designer discovers the absence of the comforting and richly constraining limitations imposed by the inherent logic and combinatorics of dialogue about some particular task. Instead, the developer finds a different, and often contrary or even contradictory, richness which is that of spoken conversation. HCA follows a set of principles for *successful, prototypical human-human conversation* which we have developed for the purpose in the apparent absence of an authoritative account in the literature which could lend itself to easy adaptation for our purposes. The principles are:

1. initially, in a polite and friendly way, the interlocutors *search for common ground*, such as basic personal information, shared interests, shared knowledge, and similarity of character and personality, to be pursued in the conversation;

2. the conversation is successful to the extent that the interlocutors *find enough common ground* to want to continue the conversation;
3. the interlocutors provide, by and large, *symmetrical contributions* to the conversation, for instance by taking turns in acting as experts in different domains of common interest, so that one partner does not end up in the role of passive hearer/spectator, like, e.g., the novice who is being educated by the other(s);
4. to a significant extent, the conversation is characterised by the participants taking turns in *telling stories*, such as anecdotes, descriptions of items within their domains of expertise, jokes, etc.;
5. conversation is *rhapsodic*, i.e. highly tolerant to digression, the introduction of new topics before the current topic has been exhausted, etc.; and
6. conversation, when successful, leaves the partners with a sense that it has been *worthwhile*.

The reader may have noted that the above list does not mention entertainment at all, despite the fact that the HCA system has an *edutainment* goal. This is partly because we assume that successful conversation is itself entertaining and partly because we want to focus on computer gaming-style entertainment in the second HCA prototype.

The ways in which HCA pursues the principles listed above are the following. He assumes, of course, that the user is interested in his life and fairy tales (1,2). However, he is aware that common ground not only has an HCA aspect but also a user aspect. He therefore tries to elicit user opinions on his fairy tales, on his visible self and on his study. However, he also tries to make the user the expert (3) by asking about games played by children and adolescents today, demonstrating interest in football, computers, and the like. During Wizard of Oz collection of 30 hours and approx. 500 spoken conversations with young users in the summer of 2003, we found that the users had strong interest in telling HCA about contemporary game-playing and also about technical inventions made after HCA's times. HCA himself, in turn, does not just answer questions, or ask them, but tells stories – about his life, about his fairy tales, about wall pictures in his room, etc. (3,4).

HCA's main problem seems to be that he cannot always pursue in depth a topic launched by his interlocutor because, at this stage of development, at least, his knowledge and conversational skills are still somewhat limited, and we do not have sufficient information about the key interest zones of his target audience. This is where the rhapsodic nature of conversation (5) may come to his rescue to some extent. When, during conversation, and despite his following an agenda in conversation, HCA is lost and repeatedly does not understand what the user is saying, he changes topic or even domain in order to recover conversational control.

Analysis of data from the user test of the system will, we hope, provide substantial information on the extent to which our implementation of the conversational strategies described above promise to achieve domain-oriented conversation, including evidence on whether the conversation is considered worthwhile by the users (6).

#### **4.7 Error handling**

Error handling meta-communication is still rather primitive in the first HCA prototype. We have considered four types of user-initiated meta-communication, i.e. clari-

fication, correction, repetition, and insult, and four types of system-initiated meta-communication, i.e. clarification, repetition, “kukkasse”, and start of conversation.

*User clarification* is not handled in PT1. This is generally hard to do and we don’t really know which kinds of clarification may occur. Thus, we have decided to wait for PT1 evaluation data before taking action. *User correction* is not treated as meta-communication in PT1 but is handled as new input to which HCA will reply if he can. The user can ask for *repetition* and get the latest output repeated. The user may also *insult* HCA. In this case, HCA will react emotionally and provide rather rude verbal output. Repetition and insult are handled by the meta domain agent.

*HCA clarification* is only handled to a limited extent in some of the mini-dialogues in PT1. PT2 is expected to allow direct clarification questions, e.g., concerning which picture the user pointed to. When HCA does not understand what the user said (low confidence score), he will ask for *repetition* or otherwise indicate that the input was not understood. HCA has various ways of expressing this, depending on how many times in succession he has not been able to understand what the user said. He also has a rhapsodic escape option from this situation, which is to jump to something completely different. To this end, he has a so-called “*kukkasse*” which is a collection of phrases that, quite obviously, are out of context, e.g. “In China, as you know, the emperor is Chinese” or “Do you think that my nose is too big?”. The hypothesis is that such rhapsodic phrases will make the user change topic instead of trying to re-express something which HCA cannot understand. If no conversation is going on but HCA receives spoken input with a low confidence score, he will address the potential user to find out if a new *conversation is starting* by saying, e.g., “Would you like a chat with me?”. Asking for repetition, “*kukkasse*”, and figuring out if a conversation is starting are all handled by the meta domain agent.

Finally, it may be mentioned that, if the system receives low confidence score gesture input, HCA does not react. This is to avoid inappropriate system behaviour in cases when a user is fiddling with the gesture input device (mouse or touchscreen).


## 5 HCA’s emotional life

Endowing HCA with emulated emotions serves two purposes. The first purpose is to add to his fundamental human-like features, the second, to make conversation with him more entertaining, due, for instance, to the occasional eruption of extreme emotional behaviour. This section describes HCA’s current emotional life.

### 5.1 Modelling emotions

HCA has the simple emotional state space model shown in Figure 5.1. His default emotional state is friendly, which is how he welcomes a new user. During conversation, his emotional state may shift towards happiness, sadness, anger, or a mixture of anger and sadness. At any time, his current emotional state is represented as ES: [h: , s: , a: ]. Each attribute has a value between 0 and 10. If h (happiness) is non-zero, s (sadness) and a (anger) are zero. If s and/or a are non-zero, h is zero. The default friendly state is ES: [h: 0, s: 0, a: 0].





**Figure 5.1.** HCA's emotional state space.

## 5.2 Eliciting emotional change

HCA's emotional state changes as a function of the user's input, for instance if the user insults HCA, wants to know his age, or shows a keen interest in the Ugly Duckling. Emotional changes caused by input semantics are identified in the knowledge base by domain agents. Emotional changes are called *emotion increments* and are represented as EI: [h: , s: , a: ]. Increment values range from 1 to 10 and only a single emotional attribute is incremented per emotion increment. Each time an emotion increment is identified, it is sent to the emotion calculator (Figure 3.2) which updates and returns HCA's emotional state. As in humans, the strength of HCA's non-default emotions decrease over time. Thus, for each user input which does not elicit any emotion increments, and as long as HCA's emotional state is different from the default ES: [h: 0, s: 0, a: 0], the state converges towards the default by (1h) or (1s+1a).

## 5.3 Expressing emotion

HCA expresses his emotional state verbally and non-verbally. A threshold function is applied for selecting knowledge base output according to HCA's current emotional state. In the friendly core (+/-6) area of happiness, sadness, and anger values, he expresses himself in a friendly manner. Beyond those values, and so far to a limited extent-only, he expresses himself in a pronounced happy, sad, or angry manner.

## 5.4 Challenges for the second HCA prototype

As described above, HCA's first-prototype emotional characteristics include: a four-emotion state space, ability to react emotionally to input, and emotional state-dependent verbal and non-verbal output. Obviously, we need to evaluate those emotional characteristics as part of the user evaluation of PT1, before making strong design decisions concerning emotion in the second prototype. However, we have identified several potential improvements in emotional behaviour which are candidates for PT2 implementation. These are described in the following.

First, HCA *may* need a more articulate emotional state space in PT2. However, compared to the points below, and despite the fact that more complex sets of emotions abound in the literature [9], this is not a top priority. As long as HCA's mechanisms for emotion *expression* are strongly limited, there does not seem to be sufficient reason for endowing him with a richer internal emotional state space.

Secondly, we would like to experiment with ways of systematically modifying verbalisation as a function of emotional state, for instance by using emotion tags for modifying HCA's verbal conversational contributions on-line.

Thirdly, we hope to be able to fine-tune HCA's non-verbal expression of emotion to a far greater extent than in the first prototype. One way of doing this is to use his current emotional state to modulate the non-verbal behaviour parameters *amplitude* and *speed*. Thus, HCA would, e.g., smile more broadly or gesture more widely the more happy he is, gesture faster the more angry he is, and act and communicate more slowly the more sad he is. A second approach, compatible with the one just mentioned, is to use rules for adding or deleting emotion tags in the response generator as a function of the current emotional state. This approach may also involve a layered model of non-verbal behaviour, so that basic posture is modified as a function of emotional state prior to modifying all posture-based non-verbal expressions.

Fourthly, it is a well-known fact that humans sometimes change at least some of their emotions by expressing them. For instance, when a person expresses anger, the anger sometimes diminishes as a result. In such cases, although the expression of anger *is* a function of the user's input, the reduced anger is *not* a function of the input but, rather, a function of the actual expression of the anger. We would like to have HCA do the same, cf. [7].

## 6 User test

The first HCA prototype was tested at NISLab in January 2004 with 18 users, nine boys and nine girls, from the target user group of 10-18 years old children and teenagers. The users' spoken input was fast-typed whereupon the system did the rest. This resulted in approximately 11 hours of audio, video, and logfile-recorded interaction and 18 sets of structured interview notes. Each user session had a duration of 60-75 minutes. A test session included conversation with HCA in two different conditions followed by a post-test interview. In the first condition, the users only received basic instructions on how to operate the system, i.e. speak using the headset, control HCA's movement, control the four camera angles, and gesture using mouse or touchscreen. After 15 minutes the session was interrupted and the user received a set of thirteen typed scenario problems to be solved through speech or gesture input in the second session, such as "Find out if HCA has a preferred fairy tale and what it is" and "Tell HCA about games you like or know". The problems might be addressed in any order and the user was not necessarily expected to carry out all of them. The purpose was to ensure a significant amount of user initiative to explore how the system would respond under the resulting pressure.

Figure 6.1 shows three excerpts from the user test conversations, illustrating, among other things, HCA's story-telling strategy, "Kukkasse" evasions (\*), user gesture input (\*\*), and what may be an overly temperamental HCA response (\*\*\*)).

**HCA:** If you know it what do you think of the character the little mermaid?  
**U:** she is very beautiful  
**HCA:** Can you tell me about the fairy tale The Little Mermaid?  
**U:** there was a little mermaid who could not wait till she turned fifteen so she could get up of the water  
**HCA:** What do you think of the prince the character from the fairy tale The Little Mermaid?  
**U:** she is very very curious  
**HCA:** Can you tell me about the prince the character from the fairy tale The Little Mermaid?  
**U:** he is a young prince who sees the little mermaid and falls in love with her  
**HCA:** I hear that the present crown prince Frederik has now found his princess. What do you think about it?  
**U:** I think it is very good for him and I think mary will be a good queen

---

**U:** what is your favourite fairy tale  
**HCA:** I like several of my fairy tales very much. One of my favourites is the ugly duckling. The ugly duckling is not accepted by the other ducks and animals because it looks different. It has a hard time of suffering and troubles. Eventually it wishes only to die till it discovers that it is in fact a beautiful swan. The morale of the fairy tale is that if you are different then you will get bullied.  
**U:** how old are you  
**HCA:** I am fifty-five years old. And I am in a good shape occasionally I ride on horses and go for long walks and when I travel I even sometimes hike.  
**U:** what is your interests  
**\*HCA:** Did I already tell you that only the bible has been translated into more languages than my fairy tales.  
**U:** really

---

**\*\*U:** picturejenny lind  
**HCA:** This is a picture showing the Swedish nightingale Jenny Lind. She sang beautifully and is one of the most beautiful women I have ever met I fell very much in love with her. Unfortunately she didn't fall in love with me but she became a dear friend of mine a dear friend. You should try to point at the things in my study. I like to tell my guests about some of them.  
**U:** stop it  
**\*\*\*HCA:** Don't piss me off!  
**U:** do you have anything to tell me about these two  
**HCA:** Sometimes I write fairy tales while I am travelling but I always finish the writing here at home in my study.

**Figure 6.1.** Three conversation snippets from the user test.

The data collected in the first HCA system user test is being analysed at the time of writing. Based on the analysis made so far it seems fair to conclude that, overall, the prototype was received remarkably well by the target users. As predicted, the tested system version performed significantly better in the first condition in which HCA faced less aggressive and totally novice users 15 of whom never spoke to a system before. In the second condition in which all users had a tall scenario-based agenda of their own, the system's limitations in handling sudden and frequent user-initiated domain and topic shifts and out-of-domain input became more evident, providing invaluable data for improving HCA's conversational abilities.

## 7 Conclusion and future work

We have described the recently completed first prototype of a life-like embodied domain-oriented conversational agent system. Following brief presentation of the system architecture, we focused on the strategies employed for making fairy tale author H. C. Andersen represent a new breed of believable embodied conversational characters who, like their taciturn professional computer game counterparts but more so, help define the notion of a virtual person rather than emulate real persons. We briefly described the recent target user test of the prototype. Ongoing work includes in-depth analysis of the collected data, addition of speech recognition to the system, and, based on these steps, requirements and design specification of the second HCA system prototype with particular emphasis on conversational smoothness and flexibility.

## Acknowledgement

The work presented is being supported by the EU Human Language Technologies programme under contract IST-2001-35293. We gratefully acknowledge the support.

## References

1. Ananova: <http://www.ananova.com/>
2. Bernsen, N. O.: When H. C. Andersen is not talking back In Rist, T., Aylet, R., Ballin, D. and Rickel, J. (Eds.): *Proceedings of the Fourth International Working Conference on Intelligent Virtual Agents (IVA'2003)*. Berlin: Springer Verlag 2003, 27-30.
3. Bernsen, N. O., Charfuelán, M., Corradini, A., Dybkjær, L., Hansen, T., Kiilerich, S., Kolodnytsky, M., Kupkin, D., and Mehta, M.: Conversational H.C. Andersen. First Prototype Description. In this volume of *Proceedings of the Workshop on Affective Dialogue Systems*, Kloster Irsee, Germany, 2004.
4. Bernsen, N. O., Dybkjær, H. and Dybkjær, L.: *Designing Interactive Speech Systems. From First Ideas to User Testing*. Springer Verlag 1998.
5. Cassell, J., Sullivan, J., Prevost, S., and Churchill, E. (Eds.): *Embodied conversational agents*. Cambridge, MS: MIT Press 2000.
6. Charfuelán, C. and Bernsen, N. O.: A task and dialogue model independent dialogue manager. In Angelova, G., Bontcheva, K., Mitkov, R., Nicolov, N., and Nikolov, N. (Eds.): *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP-2003)*, Borovets, Bulgaria, 2003. INCOMA, Shoumen, Bulgaria 2003, 91-97.
7. Marsella, S. C.: Interactive pedagogical drama: Carmen's bright IDEAS assessed. In Rist et al. 2003, 1-4.
8. NICE: <http://www.niceproject.com/>
9. Rist, T., Aylet, R., Ballin, D. and Rickel, J. (Eds.): *Proceedings of the Fourth International Working Conference on Intelligent Virtual Agents (IVA'2003)*, Kloster Irsee, Germany, Berlin: Springer Verlag 2003.
10. Turing, A.: Computing machinery and intelligence. *Mind* 59, 1950, 433-60.
11. Zoltan-Ford, E.: How to get people to say and type what computers can understand. *International Journal of Man-Machine Studies* 34, 1991, 527-547.