The Intermediary Agent's Brain: Supporting Learning to Collaborate at the Inter-Personal Level

(Short Paper)

Juan Martínez-Miranda, Bernhard Jung, Sabine Payr and Paolo Petta
Austrian Research Institute for Artificial Intelligence, OFAI
Freyung 6/6, A 1010 Vienna, Austria (EU)
{juan.martinez|bernhard.jung|sabine.payr|paolo.petta}@ofai.at

ABSTRACT

We discuss the design of the Intermediary Agent's brain, the control module of an embodied conversational virtual peer in a simulation game aimed at providing learning experiences regarding the dynamics of collaboration at the inter-personal (IP) level. We derive the overall aims of the game from theoretical foundations in collaboration theory and pedagogical theory and related requirements for the virtual peer; present the overall modular design of the system; and then detail the design perspectives and the interplay of the related operationalised concepts leading to the control architecture of the Intermediary Agent, that is realised as a simple cognitive appraisal process driven by direct and indirect effects of the mission-oriented and social interactions of players and agent on the agent's level of trust in its human peers. We conclude with coverage of related work and insights from first deployment experiences.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents*; K.3.1 [**Computers and Education**]: Computer Uses in Education—*Collaborative learning*

General Terms

Design, Experimentation, Human Factors, Theory

Keywords

Virtual Characters, Models of Personality, Serious Games

1. INTRODUCTION

The subject of Collaboration has attracted attention in research areas including management [7], organisational dynamics [12] and education [16], mainly because effective collaboration dynamics are fundamental to learning, knowledge exchange, and development/innovation processes in a wide variety of contexts. Simulation and games-based learning experiences built on dynamic models of human behaviour in organisational contexts have emerged prominently, providing learners with the experience of achieving realistic missions that require them to come in touch with and influence the behaviour of simulated characters displaying different types of attitudes [1, 2, 4, 6].

Cite as: The Intermediary Agent's Brain: Supporting Learning to Collaborate at the Inter-Personal Level (Short Paper), Juan Martínez-Miranda et al., *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May,12-16.,2008,Estoril,Portugal,pp.1277-1280.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

This work is part of an international effort aimed at improving the understanding of factors inhibiting effective collaboration dynamics and leading to the failure of collaboration initiatives, and the interventions required to reduce these risks. The adopted method relies importantly on the development of simulation games and their deployment in workshop scenarios, where they provide a shared reference of experience for subsequent facilitated debriefing sessions. The game presented here is aimed at providing learning experiences regarding the dynamics of collaboration at the inter-personal (IP) level. It focuses on factors that determine both motivation and capability to collaborate at the individual level, their manifestations in inter-personal conversational exchanges, and the possibilities to influence them through one-on-one interactions. Players face a scenario where mission accomplishment requires them to collaborate successfully with a simulated peer, the Intermediary Agent (IA). In this paper, we focus on design aspects of the IA's control architecture: Section 2 introduces theoretical foundations and pedagogical aims; section 3 presents the overall design of the simulation game; section 4 discusses the design elements of the control architecture of the modelled virtual peer; section 5 covers some related work, and we conclude with first findings from empirical evaluations.

2. THEORETICAL FOUNDATIONS

Collaboration can take place at the inter-personal (IP) level, often operating within one or more groups or teams, in an organisational or inter-organisational context. Particular to the IP level, specific dynamics determine success or failure in the collaboration relationship, such as involving trust; power; autonomy; and the impact of individual differences in personalities and motivational or cognitive abilities [8, 9]. Our IP level simulation game addresses these collaboration dynamics, providing teams of players/learners with experiences of how difficult inter-personal collaboration with an individual (virtual) peer can be. The game is a main component of a learning experience (a workshop of up to one day) designed for facilitated groups of participants interested in extending their understanding of the collaboration dynamics in inter-personal contexts. The game supports learning about important IP collaboration dynamics and breakdowns through instrumental mission-oriented and social interactions with a virtual peer, and about communication skills in challenging mission settings; it provides intense experiences, analysed in a debriefing.

Out of the range of concepts identified as influential for the dynamics of collaboration relationships mentioned above, we picked the *trust building cycle* model [17] as first reference: a framework of nurturing activity to establish and maintain a certain level of trust in collaboration. At present, we explicitly model the dynamics of

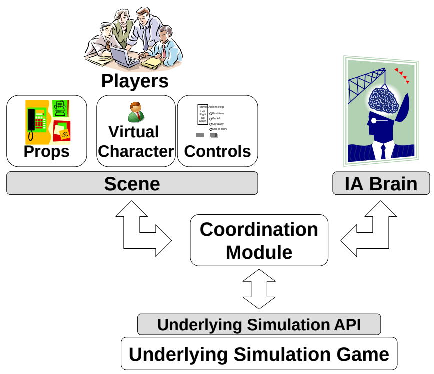

Figure 1: General architecture of the IP level game

the relational attitudes of the IA towards the player (the IA's trust in the player), based on a simple cognitive appraisal process driven by direct and indirect effects of interactions on the IA's trust level. A dynamic choice of dialogue moves (utterances) allows players to probe and try to influence the IA's state, and to portray themselves. The assumptions about trust adopted are: trust increases collaborativeness; (un)successful initiatives increase (decrease) trust; consistently poor performance decreases trust; a small wins strategy thus is more likely to be successful in building trust; openness increases trust; social (as opposed to mission-oriented) interaction can increase trust, if pursued in appropriate contexts.

3. OVERALL DESIGN

Fig. 1 provides an overview of the modules of the high-level architecture of the game. The key concept of the Intermediary Agent (IA) is realised by two components: the Brain, which encapsulates operationalisations derived from theoretical models of collaboration dynamics (section 2), and the IA's appearance as a Character in the Scene of the user interface. Another key element is the background environment for the players' mission requiring collaboration with the IA: the *underlying* simulation game.

The Underlying Simulation Game. This component (interfaced via an API providing some degree of independence from the specific game instance) provides the goal-oriented mission context for the collaboration setting. Progress with the mission in the underlying simulation is influenced by the (un)collaborative behaviour of the IA, which in turn is influenced by its interactions with the player. We use a remotely hosted instance of EIS [3] (an organisational change management mission, wherein the IA is integrated as the defined organisational contact), taking advantage of its following features: EIS is *turn-based*, and consumes *simulated time*: all changes within the simulation that are due to some external input occur one after another and take a specific amount of simulated time. It is *initiative-based*: each *initiative* is a mission-oriented action[1] and can have pre-conditions or other requirements; for each initiative chosen, a textual characterisation of its outcome is provided (in addition to causing changes to the internal simulation state). *Accessibility of state*: internal state needs to be accessible to the IA to some degree.

The Scene. The scene provides the means for the players to interact with the IA and directly with the underlying simulation game. It is run locally, and currently contains a fixed menu of the EIS managerial initiatives (out of which the next is chosen, to be issued by the IA or to be issued directly, bypassing the IA and with related collaborative consequences), and a dynamically adapted short selection of utterances to be directed to the IA; information displays reporting on the current status in the underlying simulation game; the virtual character[2], which communicates via synthesised speech and speech bubbles and portrays an expressive visual rendering of the IA's embodiment; and a conversation history, allowing the players to study the evolution of the interaction with the IA.

The Intermediary Agent Brain. The Brain module runs locally and encapsulates the mechanisms producing the IA's (un)collaborative behaviours. As conceptual entity, the Brain mediates between player and underlying game, issuing requests for initiatives to be implemented and informing about related outcomes.

The Coordination Module. This remotely hosted coordination medium [14] is realised as Web-based message-board, where other components can add and retrieve XML-encoded messages.

4. THE BRAIN DISSECTED

The Brain manages the IA's interaction with the player via *branching dialogues*, building on the modelled concepts of *trust*; *trust change tendency*; *Dynamic Attitude towards Collaboration (DAC-levels)*; and *responsibility* for initiative choice to generate believable behaviour. A *personality profile* shapes the IA's overall pattern of behaviour and informs its *emotional reactions* to events. Finally, the IA has some bounded knowledge about the effectiveness of initiatives[3] and personal preferences (friends and foes among the top managers), that become effective under conditions of limited/no collaborativeness.

4.1 The Conversation Cycle

The behavioural repertoire of the IA comprises two main classes: *talking* to the player and *idling/waiting* for the player's next dialogue move; the two categories are connected by consistent expressiveness of the verbal and non-verbal behaviours. Conceptually, the player–IA interaction develops over *conversation cycles*, defined in terms of initiatives issued in sequence in the underlying simulation game. A conversation cycle is structured into the following stages:

Introduction. This is either some initial greeting of the peers (e.g., by the IA: *"Hello! I am Julie. I will be working with you in this mission."*), or some statement signalling that the previous cycle (analysis of the results of the initiative most recently issued) is concluded, and the next cycle has begun.

Asking for suggestion. This optional stage is entered when players select an utterance asking the IA for initiatives to implement next. Depending on its current level of collaborativeness, the IA can make some proposal, or resist providing a suggestion.

Proposing an initiative. The players propose the initiative to implement: this need not match any suggestions by the IA. Again, the IA can accept or resist issuing the initiative in the underlying simulation. Such resistance allows to instantiate specific sources of collaboration breakdowns (e.g., unavailability of the IA: *"I need to go out now, please issue your request again later on."*; lack of commitment: *"Sorry, I have some other urgent work to finish first!"*; or inter-cultural differences: *"I cannot do that, local holidays are coming and I have no time for any request."*). All of these are modelled in the Brain in terms of specific thematic episodes following such hallmark resistance moves.

Implementing an initiative. Usually, the IA will eventually decide to implement an initiative, issuing it in the underlying simula-

[1]EIS models 18 managerial initiatives, e.g., to seek advice of members of the top management team (2 simulated days), or producing an article for the organisation's internal magazine (3 days).

[2]Realised using the LivingActor technology by Cantoche.

[3]Static broad categorisations in terms of riskiness or absolute "No-No"s (e.g., issuing directives).

tion. This can occur in full compliance with the player's request, can comply only partially (with some or all parameters changed), or can even be an altogether different initiative.

Coping with results. After an initiative has been implemented, the IA reports (with varying degree of detail) feedback obtained about the effect achieved. Players can react with utterances reflecting their evaluation of the result (e.g., assigning credits/blame: "*Come on! How could this happen? That is really bad!*"); they may ask for more information, at which the IA may disclose information initially kept back; or issue utterances related to purely social themes ("*Say, have you seen the latest news on TV?*").

4.2 Brain Dynamics

The following variables modulate the IA's behaviour:

Responsibility represents who takes the final decision in issuing an initiative: *Shared*, when the players adopt a suggestion by the IA (and the IA implements it unchanged); *IA*, when the players' choice is altered by the IA; *Players*, when the IA is requested to issue an initiative it was not consulted about previously.

Trust change tendency represents the agent's tendency to alter its trust towards the player. It is influenced by the IA's evaluation of the player's choice of utterance as positive (e.g., being asked in a friendly manner), negative (offensive authoritative style), or neutral. It increases when players comply with a suggestion, and decreases when players disregard it or disagree.

Trust level represents the how much the IA currently trusts the player. It is updated after an initiative was implemented in the underlying simulation (based on the IA's rating of the outcome, and the assigned responsibility); after players' reactions in the *Coping with results* stage; and (only negatively) when the players exceed a threshold for repeated moves (e.g., insisting on the IA to provide suggestions, or to disclose further information).

Dynamic Attitude towards Collaboration (DAC) represents the different classes of collaboration of the IA: for a value of *collaboration breakdown*, the IA does not implement any requested initiative and implements and suggests only initiatives with bad consequences (if any); when *non-collaborative*, the IA does neither offer suggestions nor implement requested initiatives most of the time; at *limitedly collaborative* level, the IA may equally provide suggestions; implement initiatives requested; and provide a complete report of the outcomes, or not; a *collaborative* IA makes (supposedly) good suggestions and implements initiatives as requested most of the time; a *super-collaborative* IA always provides the best suggestions it can, prevents the player from choosing initiatives with bad consequences, and gives complete reports of initiatives' effects. Changes between DAC levels depend on the trust level, the modelled personality, and the simulation time elapsed. It takes a personality-dependent amount of increase or decrease over the current Trust level for the DAC to change; thresholds for changes between DAC levels increase with simulated time, so that e.g. it takes exceptionally good progress towards reaching the mission's goal to manage to change an uncollaborative IA for the better.

4.3 Personalities and Emotions

Without personality, the behaviour of the IA would vary only with the level of trust, reflected in its demonstrated degree of collaborativeness. To improve believability [13, 15], we employ personality profiles. Agreeable and disagreeable profiles and related utterances and interaction themes were defined, given the high impact of trait agreeableness on collaborative topics such as *cooperation* and *social harmony* (Fig. 2). Agreeableness also relates to bipolar facets such as empathy; friendliness; and helpfulness; of relevance in the social relationship of IA and player, enabling scenarios with an unfriendly but highly collaborative, or a friendly but limitedly collaborative IA.

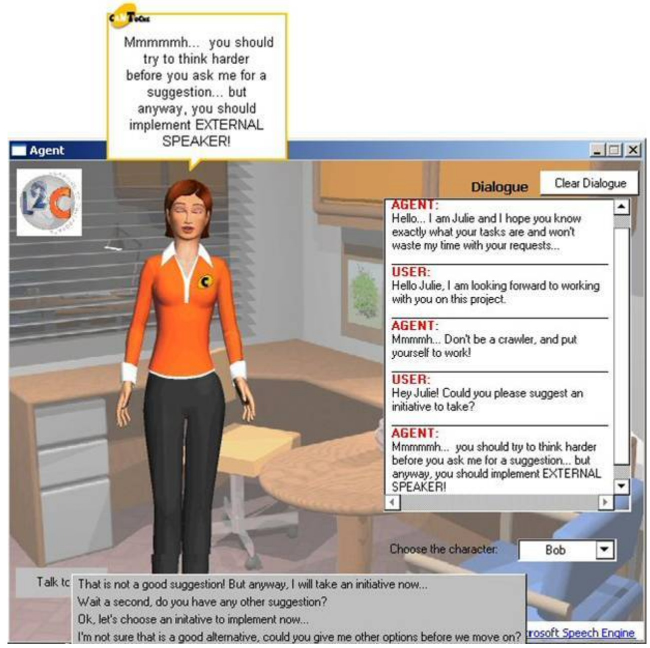

Figure 2: Welcoming stage with a *disagreeable* IA.

Using the internal parameters, the reactions of the IA to salient internal and external events are modelled as emotional appraisals. Resulting action tendencies are mapped to behaviour parameters for expressive animation, allowing to model longer-term "mood" (defined by the current DAC level) with superimposed immediate but shorter valenced reactions in a principled fashion.

5. RELATED WORK

[5] covers a complex model of trust for a reversed scenario, in which it is the agent's task to establish social relationships. Recent efforts aimed on training include the Tactical Language Training (TLT) System [10] combining an ITS with a 3D game to practise language skills in simulated social situations. This application employs sophisticated language technologies and aims at directly supporting real-time skills; in contrast, we aim for broad deployability, focus on more reflected skills, and adopt a workshop scenario, where human facilitators and learning technologies are used in a complementary fashion. A similar argument holds vs. the SASO-ST project and related efforts [6]. This whole list of "heavy" systems could be seen at the opposite end from ours, but certainly related in terms of individual component elements identified. Our overall guiding motif is the development of versatile simulation technologies and quickly customisable applications; the recently published game-based ELECT BiLAT system [11] interestingly includes a number of elements very similar to those employed in our application: trust as key variable governing the actions of the virtual tutor; classifying activities in terms of required, usual, and avoids (cf. stages and themes in the conversation cycle); and different kinds of coaching feedback messages (we exploit semantically annotated feedback from EIS). Another distinguishing aspect of all applications developed in context of our umbrella project, is the social set-up, where *groups* of players discuss turns to take: In the group setting, people reflect on their motivations for moves and their internal models of the IA: reflection (and some learning) happens *during* the game.

6. FIRST RESULTS AND CONCLUSIONS

A first evaluation is based on sustained sessions (completing the mission or exhausting the available simulated mission time) of two individuals and three workshops (two single-team and one two-team; teams of four players: MBA students and post-docs). Overall, the learners appreciated the concrete scenario and the interactive experience, described as engaging and interesting. The game was seen to lead to unanticipated situations and capable of securing the players' attention over the one hour allotted for completion of the mission (but for the "break-down group", see below). The simulation based workshop was assessed more involving than traditional teaching methods.

Beyond this general indication of adequacy of the basic design (corroborating prior WOz studies), the expected issue of a constrained conversational repertoire was raised. The IA is currently limited to just over two hundred utterances (for each of the agreeable and unagreeable personalities). These suffice to model a few variants of the stages and themes of conversation cycles at different levels of collaborativeness, but need to be augmented to fully meet the requirements of a 60 minutes playing time at an average of 2–3 utterances per minute. Such improvement of conversational competence in perceived sustained variability also requires extending the explicit history of interaction: The IA currently relies on the implicit representation in the current DAC and Trust levels (plus Trust Change Tendency within single conversation cycles). Still, the utterance-based branching dialogue appears to be scalable enough for the bounded universe of interaction.

An individual tester got locked in at a level of limited collaboration, consistently failing to identify the last one or two moves in a series to be "rewarded" a qualitative DAC level increase. To prevent such degenerate situations, the original design of the DAC included a momentum term favouring consistent change in either direction; however, this element was discarded because of early insight that the clearer discrimination of the competence of players induced was overly unforgiving, especially given that these games are meant to be played only a very limited number of times at most. Even so, one learner/player group went through the frustrating experience of falling into a *collaboration breakdown*, as a result of "gaming" attitude characterised by trial and error; shifts in strategy; and sampling of different interaction approaches, from politeness to aggression.

These experiences help elucidate the performance criteria of the present system (and also differences over the related work): given the purpose of its deployment, catching and resolving degenerate livelock situations as the ones described falls into the responsibility of the workshop facilitator, who then has to explain how the problems encountered were in fact due to the scarcity of the players' choices, rather than of the IA, thereby assisting them to realise an actual learning experience. At the same time, alongside the evident need to expand the range of utterances, the comments of players faring better in their collaboration experience admonished to improve the coverage in terms of interaction themes, including a wider variety of resistance episodes, but also larger flexibility in turn-taking. This is where we come full circle and elements of the related work become highly relevant. One challenge ahead is how to best exploit the capabilities of models as employed there, without relinquishing the deliberately shallow approach pursued. In addition, in spite of the encouraging early findings, reconciling the gap between wall-clock "real-world" interaction with the IA (including the real-time gaps caused by discussions among players) and the underlying simulation game, and how to best support suspension of disbelief/preservation of immersion in this kind of scenarios remain a fascinating research issues.

7. ACKNOWLEDGMENTS

This paper reflects only the author's views. The European Community is not liable for any use that may be made of the information contained herein. This research is carried out within the EU FP6 Project L2C Learning to Collaborate [IST-2004-2.4.10 027288]. OFAI is supported by the Austrian Federal Ministries for Science and Research and for Transport, Innovation and Technology.

8. REFERENCES

[1] C. Aldrich. Learning by Doing: A Comprehensive Guide to Simulations, Computer Games, and Pedagogy in E-Learning and Other Educational Experiences, Jossey-Bass, 2005.

[2] A.A. Angehrn. Designing Innovation Games for Community-based Learning and Knowledge Exchange, *Int'l Journal of Knowledge and Learning*, 1(3):210–228, 2005.

[3] A.A. Angehrn. Learning to manage innovation and change through organizational and people dynamics simulations, in Proc. of the Int'l Simulation & Gaming Association Conference (ISAGA 2005), 2005.

[4] A.A. Angehrn. Designing SmallWorld Simulations: Experiences and Developments, in *Proc. ICALT 2006*, IEEE Computer Society Press, 413–414, 2006.

[5] J. Cassell, T. Bickmore: Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents. *UMUAI*, 13(1–2):89–132, 2003.

[6] M. Core, D. Traum, H.C. Lane, W. Swartout, J. Gratch, M. van Lent, S. Marsella: Teaching Negotiation Skills through Practice and Reflection with Virtual Humans, *Simulation* 82(11):685–701, 2006.

[7] M.T. Hansen, N. Nohria. How to build collaborative advantage. *MIT Sloan Mgmt. Review*. 46(11):22–30, 2004.

[8] R.A. Hattori, T. Lapidus. Collaboration, trust and innovative change. *Journal of Change Management*, June 2004, 98.

[9] C. Huxham, S. Vangen. Managing to Collaborate: The Theory and Practice of Collaborative Advantage. London: Routledge, 2005.

[10] W.L. Johnson W.L., C. Beal, A. Fowles-Winkler, et al. Tactical Language Training System: An Interim Report, in Lester J.C. et al. (eds.) ITS 2004, Springer LNCS 3220, 336–345, 2004.

[11] H.C. Lane, M.G. Core, D. Gomboc, A. Karnavat, M. Rosenberg: Intelligent tutoring for interpersonal and intercultural skills, I/ITSEC Conference 2007, Paper 7417.

[12] R.C. Mayer, J.H. Davis, F.D. Schoorman. An integrative model of organisational trust. *Academy of Management Review*, 20(3):709–734, 1995.

[13] D. Moffat: Personality Parameters and Programs, in R. Trappl, P. Petta (eds.), Creating Personalities for Synthetic Actors, Springer LNAI 1195, 120–165, 1997.

[14] A. Omicini, F. Zambonelli, M. Klusch, R. Tolksdorf: Coordination of Internet Agents, Springer Berlin, 2001.

[15] A. Ortony: On Making Believable Emotional Agents Believable, in Trappl R. et al. (eds.), Emotions in Humans and Artifacts, MIT Press, 189–212, 2003.

[16] R.D. Pea. Seeing What We Build Together: Distributed Multimedia Learning Environments for Transformative Communications. *Journal of the Learning Sciences*, 3(3):285–299, 1994.

[17] S. Vangen, C. Huxham. Nurturing collaborative relations: building trust in interorganisational collaboration, *Journal of Applied Behavioural Science*, 39(1):5–31, 2003.

A 'Companion' ECA with Planning and Activity Modelling (Short Paper)

Marc Cavazza, Cameron Smith, Daniel Charlton, Li Zhang
University of Teesside
School of Computing
Middlesbrough, United Kingdom
+44 1642 218121
{m.o.cavazza, c.g.smith, d.charlton, l.zhang}@tees.ac.uk

Markku Turunen, Jaakko Hakulinen
University of Tampere
Department of Computer Sciences
Tampere, Finland
+358 3 3551 8559
{Markku.Turunen, Jaakko.Hakulinen}@cs.uta.fi

ABSTRACT
In this paper, we describe the development of an Embodied Conversational Agent (ECA) implementing the concept of a companion, i.e. an agent supporting the persistent representation of user activities and dialogue-based communication with the user. This first experiment implements a Health and Fitness companion aimed at promoting a healthier lifestyle. The system operates by generating an 'ideal' plan of daily activities from background knowledge and dialogue interaction with the user. This plan then becomes an activity model, which will later be instantiated by reports from the user and analysed by the agent from the perspective of initial objectives. At various stages of the day, the plan can still be adapted through further dialogue. The agent is embodied using a wireless rabbit (Nabaztag™) device situated in the user's home. After describing the planning component, based on Hierarchical Task Networks (HTN) and the spoken dialogue system, we present a working example from the system illustrating its behaviour through various phases of user activity generation, updating and re-planning.

Categories and Subject Descriptors
H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems.

General Terms
Algorithms, Human Factors.

Keywords
Embodied Conversational Agents, Planning, Human-Computer Dialogue, Assistive Systems.

1. INTRODUCTION
The successful development of ECA opens the way for many new applications. Alongside training, education and entertainment applications, virtual advisors [4] and personal assistants [2] of all kinds have attracted considerable interest in recent years. A new paradigm for virtual assistants has emerged in the form of companions [21], defined by Forbus and Hinrichs [7] as being able to interact with users over sustained periods of time, while also possessing robust reasoning abilities.

Cite as: A 'Companion' ECA with Planning and Activity Modelling (Short Paper), Cavazza, M., Smith, C., Charlton, D., Zhang, L., Turunen, M., Hakulinen, J., *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1281-1284.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Figure 1: The Nabaztag™ device

In this paper, we describe the development of a physically embodied Health and Fitness Companion (HFC), which aims at promoting healthier lifestyle for a typical user as office worker. It is a central feature of this application to operate in an anytime, persistent fashion both in terms of knowledge use and in terms of dialogue sessions. The user can decide to interact with the HFC to request specific advice but, in the long term, its main mode of operation should be to embed such advice inside more open conversation whose topics will be dictated by the context in which they take place (time of the day, user expected or intended activities).

The HFC is embodied using the Nabaztag™ device (Figure 1), a commercial wireless rabbit character [18] already recognised as one of the most successful ubiquitous computing devices in terms of consumer adoption and potential for applications.

2. RELATION TO PREVIOUS WORK
This work relates to previous research in several ways, both in terms of similar applications and through its underlying technical choices in planning and dialogue.

Several groups have described assistive systems for daily life or office work, although not all of them as ECA. The *Autominder* system [12] [14] is an autonomous mobile robot that can 'live' in the home of an older individual, and provide him or her with reminders about daily plans". The CALO project aims at developing a personal assistant helping an office worker to deal with information and task overload [13] [2]. The POLLy system

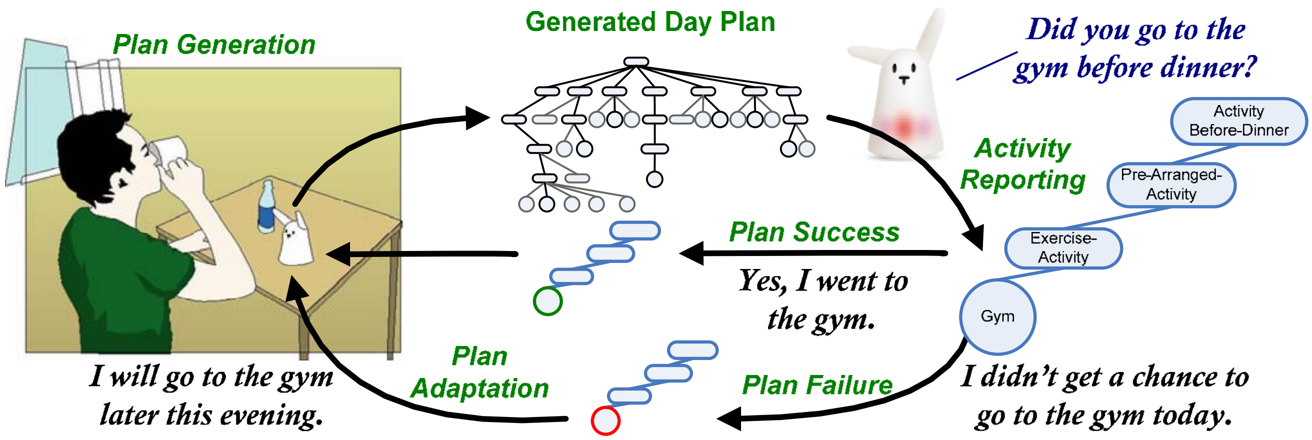

Figure 2: The various phases of interaction with the HFC: Plan generation, activity reporting, plan adaptation or replanning.

[10] has been developed to research politeness in the context of task-based interactions (more specifically, cooking). In this system, dialogue would however take place over plan execution.

Several dialogue systems have used plans as underlying knowledge models, in particular for the representation of joint user-system tasks since the original TRAINS project [5]. Similar approaches to decompositional planning as task representations or baseline plans have been described, for instance, in TRIPS' "straw plans" [6] or WITAS' "recipes" [9].

3. SYSTEM OVERVIEW

For this first prototype, we have devised an interaction scenario which assumes that the HFC is located at the user's home, and consequently the user will only interact with it during specific phases of his working day: in the morning before leaving for work and in the evening just after returning from work but before any further leisure activities. This in turn determines various phases for relating dialogue to planning (Figure 2), and for the nature of dialogue itself:

- plan generation: to plan the day's activities ahead (e.g. in the morning, sometimes leaving certain options open for later in the day).

- activity reporting: to report on activities which took place during the day to instantiate *a posteriori* the task model. This type of dialogue depends mostly on the user but has to be primed by relevant questions from the Nabaztag™.

- plan adaptation: to adapt a portion of the plan or re-plan an entire phase of the day before it takes place, depending on changing user conditions rather than on the outcome of previous phases

In line with the philosophy of a companion agent, we want to depart from task-related dialogue sessions during which the user would be systematically asked for required parameters, with the system leading dialogue and acknowledging all user input. More natural and asynchronous communication can be based on the fact that the agent possesses background knowledge on the user's preferences and her activities. For instance, when elaborating a plan for the user's daily activities, the system will only enquire about specific situations (e.g. the weather conditions or the user's mood) or the user's preferences.

4. SYSTEM ARCHITECTURE

The first HFC prototype, as presented in Figure 2, is implemented with a generic agent-based architecture designed for adaptive spoken dialogue systems [19]. It has been used in several spoken dialogue systems, including a multilingual spoken dialogue system [20]. In the HFC, this architecture is extended to support interaction with virtual and physical Companions.

Our architecture is based on distributed but coordinated components, shared system knowledge and a general system-level adaptation mechanism. The system architecture is distributed so that different managers and agents can run on different computers and platforms. It is similar to certain central components found in other speech architectures, such as the HUB in the Communicator architecture [16], and the Facilitator in the Open Agent Architecture [11].

4.1 Speech Input and Output

The Communication Manager handles all input and output management. It includes devices and engines that provide interfaces to technology components. Most importantly, in the HFC it includes components to control Loquendo™ ASR and TTS components and the physical agent interface. The system uses recognition grammars in "Speech Recognition Grammar Specification" (W3C) format that are dynamically selected by the Modality Manager according to the current dialogue state. Dynamic grammar generation also takes place in certain situations.

In the first prototype natural language understanding is based on the concept-spotting approach, using heavily "Semantic Interpretation for Speech Recognition (SISR) Version 1.0" (W3C) format information. Semantic information provided by the SISR tags is combined with the dialogue state to construct predicates compatible with the planning domain.

Natural language generation is implemented with a concept-based approach, mostly using templates. The main starting point is predicate-form task descriptions formed by the cognitive model. Further details and contextual information are retrieved from dialogue history, the user model, and potentially other sources. Finally, SSML (Speech Synthesis Markup Language) 1.0 tags are used for controlling the Loquendo™ synthesizer.

4.2 The Physical Agent Interface

For a physical agent interface, the jNabServer software was created to handle communication with the Nabaztag™. The Nabaztag™ device can handle various forms of interaction, from voice to touch (button press), and from RFID 'sniffing' to ear movements. It can respond by moving its ears, by displaying or changing the color of its four LED lights. It can also play sounds which can be music, synthesised speech or other voices.

4.3 Dialogue Management

The Dialogue Manager takes care of conversational strategies and communicates with the planner that generates the user activity model. Together, they use hierarchical task decomposition and a dialogue stack similar to CMU Agenda [15] and RavenClaw [3] systems. The dialogue manager maintains a dialogue history tree and communicates facts and user preferences to the planner at the various stages of plan elaboration and task instantiation (Figure 2). The planner (implemented in Allegro Common Lisp) is connected to the software architecture via the Cognitive Model Manager. The integration between the Planner and the dialogue system is based on a mapping between the dialogue lexicon semantics and the Planning domain, as presented in detail in Section 5.

5. ACTIVITY MODEL PLANNING

In order to fulfil his role as an assistant, the system generates a global plan corresponding to an ideal course of action for the user's daily activities. The system uses planning techniques to generate a reasoned top-down decomposition of user activities, implicitly ordered to follow the rhythm of a normal day itself. The central idea of our approach is that this plan in turn becomes a task model representing potential user activities which will be instantiated by user reports.

Establishing the plan consists in generating user activities in a way which maximizes energy expenditure and minimises food intake, within the boundaries of normal activities. There is an implicit agreement that the user will actually follow the plan for the 'standard' part of her activities, and for those actions explicitly discussed with the Nabaztag™.

5.1 Plan Generation

We use Hierarchical Task Network Planning with a total-order forward decomposition algorithm [8], which has been specifically extended to incorporate semantic knowledge in the decomposition process. That is to say, when there are multiple applicable methods, selection of the most appropriate method is based on a heuristic approach that uses semantic categorisation.

To illustrate this, we look at the high level task of travelling to work from home ('Medium-Distance-Travel Home Work', which is part of the Plan-Day domain). The task can be decomposed into eight different options depending on how the user will travel to work. In terms of the AND/OR tree, this involves a root node holding the 'Medium-Distance Travel Home Work' task with an OR-branch node holding five task nodes (see Figure 3).

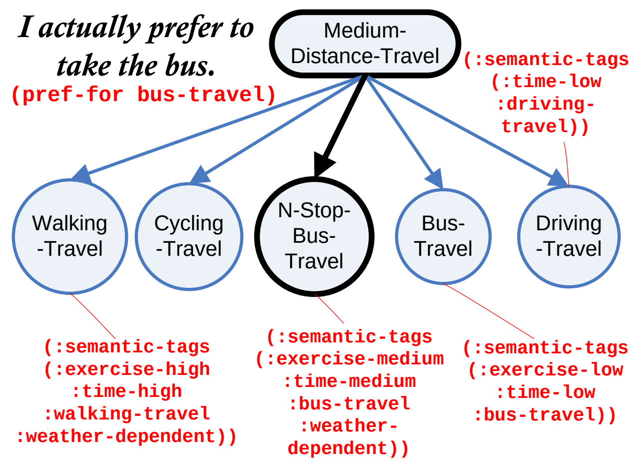

Figure 3: Plan Generation with Semantic Knowledge

The various options outline various ways of getting to work including walking or taking the bus. (The 'n Stop' option for bus indicates getting off a couple of stops early, so as to get more exercise.) The definition of each in the domain includes semantic tags which are used to contrast the differing properties of each option. These semantic descriptions correspond to domain knowledge which should be activated from dialogue.

The initial state of the planner (in Figure 3) contains a recent preference for bus travel generated from dialogue with the user. This preference ensures that Medium-Distance-Travel-n-Stop-Bus scores higher than the other Medium-Distance-Travel options and thus this task is selected to decompose further.

5.2 Activity Reporting

Once a plan has been generated it becomes a task model for user activities and rests on the assumption that the user will generally follow the plan, with however potential for departing from it. It is thus necessary to update the task from the user herself at different stages, for instance when the user returns from work, following the cycle of interaction described on Figure 2. This is done by traversing the AND/OR graph defining the plan and marking task nodes as completed or failed based on the information available (although strictly speaking, this is not a case of the plan "failing" as it is only used as a resource).

5.3 Plan Adaptation

Plan adaptation consists in surface modifications to the planned activities [17]: from a task decomposition perspective, adaptation can be formalised as only involving the lower levels of task decomposition. After the plan has been generated the user may wish to change some aspect of it without generating a whole new plan. This is accomplished by the user rejecting a current task which results in the planner being re-activated and backtracking to the nearest overarching OR branch and generating a new sub-plan from the remaining nodes.

6. CONCLUSIONS AND FURTHER WORK

We have described a first implementation of a 'Companion' ECA generating and analysing user activities so as to influence his/her behaviour. We have adapted the level of plan elaboration to several factors, amongst which the constraints of interacting only when the user is at home as well as a desire to allow more flexible interaction and to avoid the type of complex negotiation and acknowledgement seen in related dialogue systems.

However, a natural extension of the system is to support some phases of real-time dialogue-based Mixed-Initiative Planning [1], in which the user would take a greater interest in the details of his daily activities. There is probably a balance to be found between user control and the burden of interaction and negotiation.

7. ACKNOWLEDGMENTS

This work was funded in part by the Companions project (www.companions-project.org) sponsored by the European Commission as part of the Information Society Technologies (IST) programme under EC grant number IST-FP6-034434.

Nabaztag™ is a trademark of Violet™, who is thanked for authorizing the development by some of the authors of the local web server “jNabServer” used in these experiments.

8. REFERENCES

[1] Allen, J., Chambers, N., Ferguson, G., Galescu, L., Jung, H., Swift, M., and Taysom, W., 2007. PLOW: A Collaborative Task Learning Agent. Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07). Vancouver, Canada, pp. 22-26.

[2] Berry, P.M., Albright, C., Bowring, E., Conley, K., Nitz, K., Pearce, J.P, Peintner, B., Saadati, S., Tambe, M., Uribe, T.E., Yorke-Smith, N., 2006, Conflict negotiation among personal calendar agents. AAMAS 2006: 1467-1468

[3] Bohus, D., Rudnicky, A, 2003. RavenClaw:. Dialog Management Using Hierarchical Task Decomposition and an Expectation Agenda. In Proceedings of Eurospeech 2003: 597-600.

[4] Cavalluzzi, A., Carofiglio, V., de Rosis, F., 2004. Affective Advice Giving Dialogs. In: Elisabeth André, Laila Dybkjær, Wolfgang Minker, Paul Heisterkamp (Eds.): Affective Dialogue Systems, Tutorial and Research Workshop, ADS 2004, Kloster Irsee, Germany, June 14-16, 2004, Lecture Notes in Computer Science n. 3068, Springer, pp. 77-88.

[5] Ferguson, G., Allen, J. and Miller, B., 1996. TRAINS-95: Towards a mixed-initiative planning assistant, in Brian Drabble, ed., Proc. Third Conference on Artificial Intelligence Planning Systems (AIPS-96), pages 70–77, Edinburgh, Scotland, 1996.

[6] Ferguson, G., and. Allen, J. F., 1998. TRIPS: An Integrated Intelligent Problem-Solving Assistant, Proc. of the National Conference on Artificial Intelligence (AAAI-98), Madison, WI, pages 567–573.

[7] Forbus, K. and Hinrichs, T., 2004. Companion Cognitive Systems: A step towards human-level AI. To appear in Proceedings of the 2004 AAAI Fall Symposium on Achieving Human-level Intelligence through Integrated Systems and Research.

[8] Ghallab, M.; Nau, D.; and Traverso, P., 2004. Automated Planning: Theory and Practice. Morgan Kaufmann.

[9] Gruenstein, A., and Cavedon, L., 2004. Using an activity model to address issues in task-oriented dialogue interaction over extended periods, AAAI Spring Symposium on Interaction between Humans and Autonomous Systems over Extended Operation.

[10] Gupta, S., Walker, M.A. and Romano, D.M., 2007. How Rude are You?: Evaluating Politeness and Affect in Interaction. Second International Conference on Affective Computing and Intelligent Interaction, ACII 2007, Lisbon, Portugal, pp. 203-217.

[11] Martin, D. L., Cheyer, A. J., Moran, D. B. The Open Agent Architecture: A framework for building distributed software systems. Applied Artificial Intelligence: An International Journal. Volume 13, Number 1-2, January-March 1999: 91-128, 1999.

[12] McCarthy, C.E. and Pollack, M. A, 2002. Plan-Based Personalized Cognitive Orthotic. Proceedings of the AIPS 2002 Conference.

[13] Myers, K. L. and Yorke-Smith, N., 2007. Proactivity in an Intentionally Helpful Personal Assistive Agent, in Proceedings of the AAAI Spring Symposium on Intentions in Intelligent Systems, AAAI Press, March 2007.

[14] Pollack, M. E., 2002. Planning technology for intelligent cognitive orthotics. In Proceedings of the 6th International Conference on Automated Planning and Scheduling.

[15] Rudnicky, A., and Xu,W, 1999. An agenda-based dialog management architecture for spoken language systems. In Proceedings of IEEE Automatic Speech Recognition and Understanding Workshop, I–337.

[16] Seneff, S., Hurley, E., Lau, R., Pao, C., Schmid, P., Zue, V. Galaxy-II: a Reference Architecture for Conversational System Development. In Proceedings of ICSLP98, 931-934, 1998.

[17] Tate, A., Levine, J., Jarvis, P., and Dalton, J., 2000, Using AI Planning Technology for Army Small Unit Operations, Proceedings of the Fifth International Confer-ence on Artificial Intelligence Planning Systems (AIPS 2000), May 2000.

[18] Tomitsch, M., Kappel, K., Lehner, A. & Grechenig, T., 2007. Towards A Taxonomy For Ambient Information Systems. Pervasive 2007 Workshop on the Issues of Designing and Evaluating Ambient Information Systems.

[19] Turunen, M., Hakulinen, J., Räihä, K.-J., Salonen, E.-P., Kainulainen, A., Prusi, P, 2005. An architecture and applications for speech-based accessibility systems. IBM Systems Journal, Vol. 44, No. 3 (pp. 485-504).

[20] Turunen, M., Salonen, E.-P., Hartikainen, M., Hakulinen, J, 2004. Robust and Adaptive Architecture for Multilingual Spoken Dialogue Systems. In Proceedings of ICSLP 2004.

[21] Wilks, Y., 2006. Artificial Companions as a new kind of interface to the future Internet, Oxford Internet Institute, Research Report No. 13.

Emotional Reading of Medical Texts Using Conversational Agents (Short Paper)

Gersende Georg
Centre des Cordeliers UMRS 872, Eq20
15 rue de l'Ecole de Médecine, F-75006

HAS, 2 Avenue du Stade de France
Saint-Denis La Plaine Cedex, F-93218
gersende.georg@spim.jussieu.fr

Catherine Pelachaud
University of Paris 8, INRIA
INRIA Rocquencourt, Mirages,
78153 Le Chesnay Cedex, France
catherine.pelachaud@inria.fr

Marc Cavazza
School of Computing
University of Teesside
TS1 3BA Middlesbrough
United Kingdom
m.o.cavazza@tees.ac.uk

ABSTRACT
In this paper, we present a prototype that helps visualizing the relative importance of sentences extracted from medical texts using Embodied Conversational Agents (ECA). We propose to map rhetorical structures automatically recognized in the documents onto a set of communicative acts controlling the expression of an ECA. As a consequence, the ECA will dramatize a sentence to reflect its perceived importance and rhetorical strength (advice, requirement, open proposal, etc). This prototype is constituted of three sub-systems: i) G-DEE, a text analysis module ii) a mapping module which converts rhetorical structures produced by the text analysis module into communicative functions driving the ECA animation and iii) an ECA system. By bringing the text to life, this system could help their authors (in our application, expert physicians) to reflect on the potential impact of the writing style they have adopted. The use of ECA re-introduces an affective element which cannot easily be captured by other methods for analyzing document's style.

Categories and Subject Descriptors
H.5.1 [**Multimedia Information Systems**] *Animations;* J.3 [**Life and Medical Sciences**]: Medical information systems; I.2.11 [**Document and Text Processing**]: Document Preparation - *Markup languages - Hypertext/hypermedia*.

General Terms
Algorithms, Human Factors.

Keywords
Embodied Conversational Agents, Document Engineering, Markup languages.

1. INTRODUCTION
The conversion of text to other modalities has been proposed initially as a means to facilitate access to its informational content. In recent years, the use of ECA to read aloud documents using Text-To-Speech (TTS) has gained increased popularity, due to progress in animation and speech synthesis. However, more sophisticated applications can be envisioned if one realises the potential of an ECA to reflect more than just the informational content of the text [2, 7, 13]. ECAs have been demonstrated to bring added value (such as disambiguating text, adding communicative and affective information) to many applications for which a more human-like presentation [8] is beneficial, including assistance, help and guidance [1,2].

In this paper, we introduce a first prototype developed to visualize the importance of specific sentences within medical documents using an ECA. Clinical guidelines are normative texts, aimed at physicians, produced by various Health authorities, which promote best practice in Medicine, based on the concept of *evidence-based medicine*. They are complex documents which require significant amounts of specialized knowledge for their production. Clinical guidelines are based on the notion of *recommendation*, which are syntactic constructs associated to a strong rhetorical value. For instance, "The administration of low doses of aspirin (75 mg/day) is recommended for hypertensive patients with type 2 diabetes in primary care." One main challenge associated to the clinical guidelines' production is to be able to anticipate the impact of the specific recommendations they contain as a function of the style used. This is why we propose the automatic visualization of recommendations, as animating a recommendation through an ECA to restore the link between document content and the original committee discussion which decided on its formulation.

2. SYSTEM OVERVIEW AND ARCHITECTURE
The system presents itself as an ECA interface "reading aloud" specific recommendations extracted from a clinical guideline. It is actually constituted of three sub-systems: i) a document engineering environment, G-DEE [6] (Guidelines Document Engineering Environment) which automatically identifies the most relevant sentences of a guideline (the recommendations), ii) a mapping module which converts those recommendations into the communicative act format used by the ECA, a mark-up language known as APML [5] and iii) an ECA system called Greta [10]. The system operates as follows. Firstly, G-DEE is run offline to analyse the clinical guideline as a whole. It produces a document in which all recommendations are identified through a set of specific mark-ups for their operators and the contents they apply to (referred to as the *scopes* of the operator). An example of scopes marking-up is:

"**<Front-scope>** The administration of low doses of aspirin (75 mg/day) **</Front-scope>** **<Op_Reco>** is recommended **</Op_Reco>** **<Back-scope>** for hypertensive patients with diabetes type 2 in primary care **</Back-scope>**."

Cite as: Emotional Reading of Medical Texts Using Conversational Agents (Short Paper), Georg G, Pelachaud C, Cavazza M. *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1285-1288.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

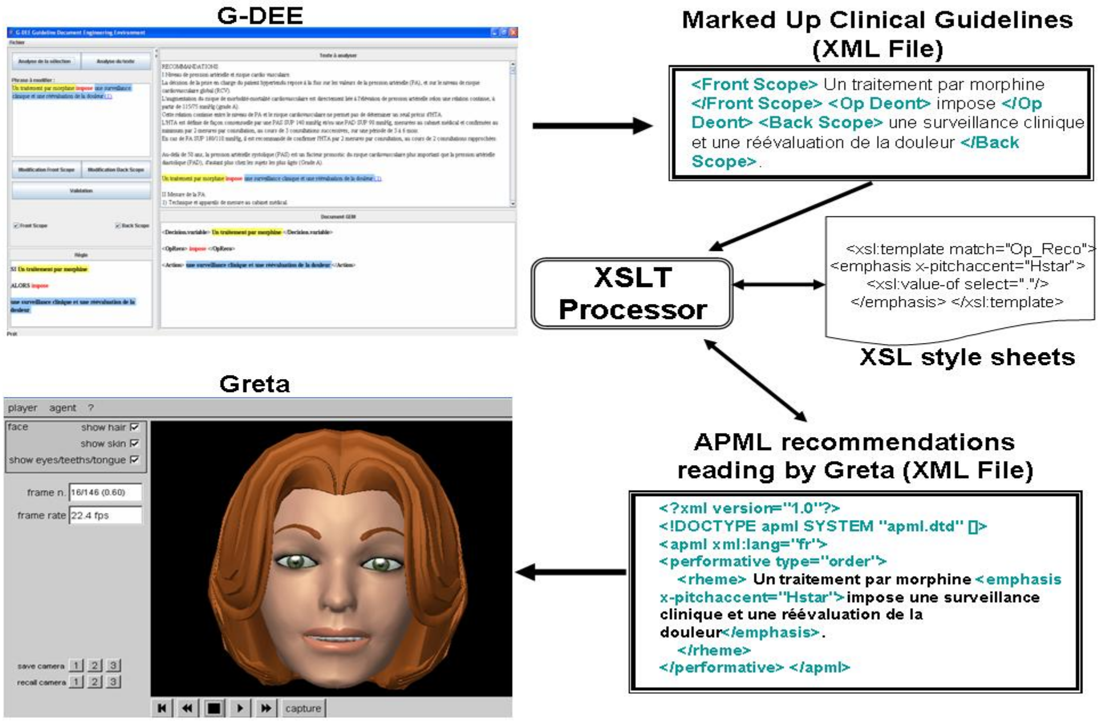

Figure 1. Overview of the architecture.

A marked-up recommendation appears as highlighted text in G-DEE (Figure 1). This text fragment can be selected interactively, which then triggers the generation of an APML file animating Greta on that sentence (the generation uses an XSLT conversion module). During this process, tags on communicative acts linked to recommendation strength are added to the text, as the result of the automatic mapping of rhetorical structures. Finally, Greta processes this APML file and utters the corresponding recommendation, displaying appropriate nonverbal behaviour, which reflects the importance of the recommendation and places emphasis on relevant scopes. In this way, the actual strength of the recommendation and its potential impact can be visualized.

3. AUTOMATIC EXTRACTION OF RECOMMENDATIONS FROM TEXTS

3.1 The Document Engineering Environment

We are interested in clinical guidelines that belong to the generic category of normative texts, to which much research has been dedicated. G-DEE [6] supports multiple document processing functions including the automatic recognition of recommendations using shallow NLP techniques recognizing *deontic operators* in medical texts such as "authorize", "forbid", "ought to". Let us now consider the different aspects that determine the strength and emphasis of a recommendation. Firstly, deontic operators fall within the broad categories of permission, obligation or interdiction. Within these broad categories, specific deontic operators can be classified according to their "strength". Strength is not just an issue of vocabulary, but relates also to syntactic constructs (which have been uncovered as part of the process of deontic operator extraction). In other words, that a specific drug "should not be used" is stronger than it being "not recommended". It can also be noted that this concept bears some similarity with the illocutionary strength of communicative acts (which constituted our initial inspiration for this project).

3.2 The Greta Platform

The Greta agent [10] used in these experiments is a platform developed for research in non-verbal behavior, including an animation system with facial parameters supporting detailed expressive animations synchronized to a TTS system. Greta's animations are controlled using instructions in the APML language [4]. Communicative acts are gathered in classes depending on the information they convey [11]. In particular, previous study [12] have been conducted to elaborate the links between performatives (communicative acts), such as: suggest, propose, refuse, etc., and facial expressions. Three main classes of performatives have been considered: *request*, *inform* and *question*. Performatives of the class *request* have been characterized along three dimensions: i) to whom is the action requested, ii) how certain one is of the information being provided and iii) the power relationship between interactants [12]. Based on the representation of performatives along these three dimensions, we have proposed a mapping between each of these dimensions and the ECA's facial expressions. That is, the facial expression associated to a given performative is obtained by combining the expressions arising from each dimension. Being certain or uncertain can be shown on the eyebrow region: one frowns when being very much certain of what one says, but raises eyebrows if uncertain. Head orientations (such as head kept straight up or tilted aside) can be a sign of a power relation: submissiveness is often shown by displaying our neck [4] while dominance is characterized by a straight up head). Performatives also contain an intrinsic emotional factor. A frown marks the performative 'order' as one can get angry if the interlocutor does not comply with the requested action.

4. RELATED WORK

Our work focuses on conversational agents for visualising rhetorical structures extracted from medical texts. It is related to storytelling agents [3], or emotionally expressive agents [1], with two important differences. The first one is that the 'emotional' content to be visualised is actually related to the importance and authority of a text fragment, rather than to its dramatic qualities. The second is naturally the application area, and the practical use of such a system to estimate the impact and readability of a given document style. It should be also emphasized that these documents have no less emotional impact because they're directed at a physician's audience: issues of importance, authority and responsibility generate powerful emotional responses as well. No work has yet reported the use of ECA to explore medical text perception by physicians. Clinical guidelines are based on the notion of *recommendation* which is a rhetoric structure advising or forbidding a specific course of action (from a pragmatic perspective this corresponds to a *deontic operator*). These recommendations have a significant emotional content which is linked to notions of authority and responsibility.

5. IDENTIFYING THE RHETORICAL STRENGTH OF RECOMMENDATIONS

Physicians do not always identify the most important information when they read clinical guidelines because of the variable quality of their formulation, and phenomena of ambiguity, imprecision, and vagueness [9]. The physician's background has also been shown to play a role in their interpretation of guidelines [6]. In order to formalize the concept of rhetorical strength of a recommendation, we conducted a study involving 14 medical experts from INSERM (French National Institute for Health) and the French National Authority for Health (HAS). These experts rated the strength of 37 recommendations extracted from recent clinical guidelines published by the HAS. They ranked the strength of each recommendation according to a predefined six-point scale defined as follows:

CAT1- well-identified best practice, which is compulsory
CAT2- a practice well adapted to the clinical situation that presents demonstrable benefits
CAT3- accepted practice which can be advised, or to be considered
CAT4- a possible practice left to the discretion of the physician
CAT5- a statement explaining a given clinical practice
CAT6- a useful information item

Figure 2. Categories for evaluating the strength of recommendations.

For each deontic verb used in recommendations, we are able to associate a numerical score quantifying its rhetoric strength. This analysis will serve as a starting point to map the rhetorical strength of deontic expressions onto the emotional categories of Greta.

6. MAPPING RHETORICAL STRUCTURES ONTO MULTIMODAL COMMUNICATIVE ACTS

The process by which the rhetorical strength of textual recommendations will be visualized rests on a mapping from deontic operators onto multimodal communicative acts. These can be described as the dynamic expression of traditional speech acts (order, advice, propose …), using speech parameters and dynamic animation of non-verbal behavior, in particular facial expressions. The rationale for such a mapping derives from the pre-existing commonality between certain deontic operators, used in the description of recommendations, and the set of primitive speech acts embedded in the APML control language (which already contains speech acts such as *advise*), although the two were developed independently by different authors. This mapping attempts to generalize these commonalities by relating deontic operators to communicative acts, but also their perceived strength to the *rheme* part [14] of APML expressions, corresponding to the intentional structure that contains the new information. We have elaborated the mapping between the six categories of the strength scale and the performatives by considering the common values for these 3 dimensions (Figure 3).

CAT1 (to impose / APML: 'order') - only the frown is kept, as the other behaviours are also power signs. To highlight the importance, emphasis is added through head nods.
CAT2 (to recommend / APML: 'recommend') - represented by a less intense frown.
CAT3 (to propose / APML: 'advice') - displayed using the eyebrow shape (slight rising of the eyebrows).
CAT4 (may / APML: 'suggest') - characterized by raised eyebrows and tilted head.
CAT5 (rarely indicate / APML: 'inform + emphasis') - translated by looking at one's addressee and performing a head nod on the emphasised word.
CAT6 (should be suspected / APML: 'inform') - displayed through gaze behaviour, namely looking at the addressee.

Figure 3. Mapping between strength and performative type.

The following example corresponds to Category 2. The dedicated style sheet enables to transform a marked-up recommendation to an APML format (Figure 4) that supports the mapping of the "*il est recommandé*" ("it is recommended") deontic verb to the *recommend* performative type.

<**apml**> <**performative type="recommend"**> <**rheme**> <**emphasis x-pitchaccent="Hstar"**>*Il est recommandé*</**emphasis**> *de réaliser un écho-Doppler veineux lors de la prise en charge de tous les patients présentant un ulcère des membres inférieurs.* </**rheme**></**performative**> </**apml**>

Figure 4. The resulting APML file corresponding to a *recommend* performative type.

The corresponding expression for Greta (Figure 5- left) consists of a recommendation with an emphasis on the deontic verb "*il est recommandé"* (it is recommended) and a raised eyebrow, while the *suggest* communicative act (Figure 5- right) is associated to a slight raising of the eyebrows and a head nod.

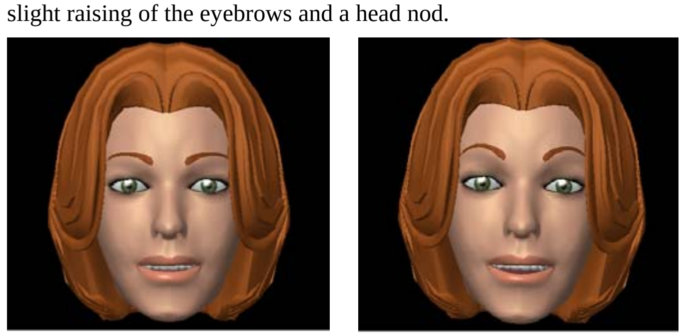

Figure 5. Expressions for *recommend* (left) and *suggest* (right).

7. PRELIMINARY USER EVALUATION

We conducted a preliminary evaluation of the system with 6 medical experts drawn from the group of the 14 experts that participated in the definition of recommendations' strengths. For this evaluation, 9 recommendations, automatically extracted by G-DEE, were visualized by Greta according to their rhetorical strength. The main objective of this evaluation consists of determining whether Greta improves the perception of the recommendations' strength, for instance by generating a stronger consensus or helping to disambiguate between categories. For this evaluation, we produced 9 videos representing Greta reading the 9 sentences with their corresponding communicative acts. These videos were presented to each of the 6 medical experts to rate the recommendation strength they perceived when Greta read the different recommendations. The average strength as well as the standard deviation were calculated for each recommendation, with and without Greta (Figure 6).

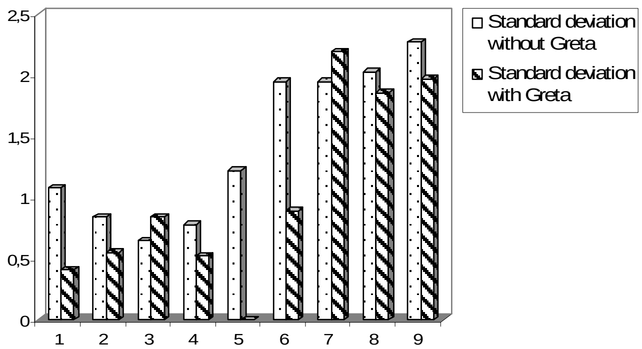

Figure 6. Impact of Greta on the standard deviation of experts' judgments of recommendations' strength.

8. DISCUSSION AND CONCLUSION

Finding the best formulation for a recommendation is a complex process, which often involves multiple cycles of discussion and negotiation within expert working groups. However, these revisions often take place after the initial document has been assembled. They are then disconnected from the consensus group discussions in which social and nonverbal behaviour plays an important part in highlighting the importance of specific recommendations. To a large extent, the system presented here can restore the link between the wording of a recommendation and its intended impact on the reader. It should help selecting the appropriate level of emphasis required, as well as balancing the importance of recommendations across the document as a whole. Our preliminary results suggest that Greta has an impact of the perception of recommendations strength. The significance of the overall distribution was tested by one-way ANOVA which showed this result to be statistically significant ($P < 0.0474$). Most importantly, we observed a significant effect of Greta on the standard deviation of perceived recommendations' strength, and that effect is more pronounced for intermediate categories, such as CAT3 and CAT5. We can argue that the diminution of the standard deviation with Greta corresponds to a better consensus between medical experts. These first results are encouraging and future work will consist of evaluating this approach with a larger test set of recommendations, also using more sophisticated expressive mechanisms such as gestures.

9. ACKNOWLEDGMENTS

Gersende Georg is partly funded through a post-doctoral fellowship from "Region Ile-de-France". We thank medical experts from the French National Health Authority (HAS) and INSERM (French National Institute of Health) for their participation in data collection and in preliminary evaluation experiments.

10. REFERENCES

1. Allbeck J. and Badler N. Toward Representing Agent Behaviors Modified by Personality and Emotion. In *Workshop Embodied conversational agents - let's specify and evaluate them! AAMAS*, (Bologna, 2002).
2. André E, Rist T and Müller J. Guiding the user through dynamically generated hypermedia presentations with a life-like character. In *Proceedings of the 3rd international conference on Intelligent user interfaces*, (San Francisco, California, United States, 1998), 21-28.
3. Cavazza M, Charles F and Mead SJ. Character-Based Interactive Storytelling. *IEEE Intelligent Systems* (2002), 17(4): 17-24.
4. Darwin CR. *The expression of emotions in man and animals*. Murray, London, 1872.
5. De Carolis B, Pelachaud C, Poggi I and Steedman M. APML, a Markup Language for Believable Behavior Generation. In H Prendinger, M Ishizuka (eds). *Life-like Characters. Tools, Affective Functions and Applications*, Springer, 2003, 65-86.
6. Georg G and Jaulent M-C. A Document Engineering Environment for Clinical Guidelines. In *Proceedings of the 2007 ACM Symposium on Document Engineering*, (Winnipeg, Manitoba, Canada, 2007), ACM Press, New York NY, USA, 69-78.
7. Gratch J, Rickel J, André E, Cassell J, Petajan E and Badler N. Creating Interactive Virtual Humans: Some Assembly Required. *IEEE Intelligent Systems* (2002), 54-63.
8. Hoorn J and Konijn E. Personification: Crossover between Metaphor and Fictional Character in Computer Mediated Communication. In *The annual meeting of the International Communication Association.* (San Diego, CA, 2003).
9. Patel V, Arocha J, Diermeier M, How J and Mottur-Pilson C. Cognitive psychological studies of representation and use of clinical practice guidelines. *Int J Med Inf.* (2001), *63* (3);147-167.
10. Pelachaud C. Multimodal expressive embodied conversational agent. In *ACM Multimedia, Brave New Topics session*, (Singapore, 2005), 683-689.
11. Poggi I. Mind Markers. In M Rector, I Poggi, N Trigo eds. *Gestures, Meaning and use*, University Fernando Pessoa Press, Oporto, Portugal, 2003, 203-207.
12. Poggi I and Pelachaud C. Performative faces. *Speech Communication* (1998), *26*, 5-21.
13. Rist T, André E, Baldes S, Gebhard P, Klesen M, Kipp M, Rist P and Schmitt M. A Review of the Development of Embodied Presentation Agents and Their Application Fields. In H Prendinger, M Ishizuka (eds). *Life-Like Characters: Tools, Affective Functions, and Applications*, Springer, 2003, 377-404.
14. Steedman M. *The syntactic process*. MIT Press, Cambridge, MA, 2000.

Individual Differences in Expressive Response: A Challenge for ECA Design (Short Paper)

Ning Wang
Institute for Creative Technologies
University of Southern California
13574 Fiji Way
Marina del Rey, CA 90292 USA
+1(310) 574-5700
nwang@ict.usc.edu

Stacy Marsella
Information Science Institute
University of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292 USA
+1(310) 822-1511
marsella@isi.edu

Tim Hawkins
Institute for Creative Technologies
University of Southern California
13574 Fiji Way
Marina del Rey, CA 90292 USA
+1(310) 574-5700
timh@ict.usc.edu

ABSTRACT
To create realistic and expressive virtual humans, we need to develop better models of the processes and dynamics of human emotions and expressions. A first step in this effort is to develop means to systematically induce and capture realistic expressions in real humans. We conducted a series of studies on human emotions and facial expression using the Emotion Evoking Game (EVG) and a high-speed video camera. In this paper, we discuss a detailed analysis of facial expressions in response to a surprise situation. We provide details on the rich dynamics of facial expressions, along with data useful for animation of virtual human. The analysis of the data also revealed considerable individual differences in whether surprise was evoked and how it was expressed.

Categories and Subject Descriptors
I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agents
I.3.7 [**Three-Dimensional Graphics and Realism**]: Animation

General Terms
Measurement, design, experimentation, human factors, theory.

Keywords
Facial expression, emotions, virtual human expressiveness.

1. INTRODUCTION
The expression of emotion promises to be the elixir that can make an embodied agent come to life. It is not surprising that as work on embodied agents has progressed, there has been an increasing interest in creating agents with human-like emotions and expressive facial expressions. Significant progress has been made in this area, but the promise has not been fully realized.

What's wrong with embodied agent's facial expression and how can we improve it? One approach is to draw on research on human emotions and emotional expression. Existing research in psychology often has not looked at human emotions and facial expression at the level of detail needed to inform agent design. For example, questions concerning the dynamics of emotional expression have largely not been addressed. In our work we have undertaken to closely study human emotions and emotional expression to develop improved ways of modeling emotions and their expression. The methodology we employ requires first a systematic method for emotion evocation and second a method to record in detail the facial expression.

Traditionally, researchers have employed a wide range of stimuli to evoke emotions. These include displaying images or videos with emotional impact (Lang et al., 1999), recall emotional events (Frijda et al., 1989), interacting with a human confederate (Stemmler et al., 2001), and etc. In this study, we used a computer game called Emotion Evoking Game (Wang, Marsella 2006). EVG allows researchers to systematically explore factors that elicit emotion. The use of computer video games promises several benefits over the traditional approaches such as inducing task-related emotions and social emotions. Previous study found that EVG can reliably induce emotions and facial expressions (Wang, Marsella 2006).

Given EVG to systematically evoke emotions, we still need a way to record in detail the resulting facial expression. Earlier work on EVG clearly identified the weakness of using standard video cameras to record facial expressions. Much of the fine detail in the dynamics was lost at standard frame rates. This is not too surprising. Some facial expressions can be fleeting. Ekman (1985) argues that micro-expressions can be on the order of 40 ms. We also know that they can be subtle (Ekman 1985), with dynamic properties that can impact human interpretation (Parkinson et al., 2005). To study human facial expression closely, we need a high speed camera to capture the richness and subtlety of facial expression at a fine grain level.

Armed with EVG and a high-speed camera, we have begun to study facial expressions in earnest. In this paper, we discuss further evaluation of EVG's ability to evoke emotions systematically. We investigate what are the dynamics of human facial expression and what do those dynamics tell us about modeling embodied agents. The study reported here reveals the highly dynamic nature of facial expression, providing detailed timing information that can guide animation design.

2. Related Work
There is a large body of research that addresses questions concerning the relation of facial expressions to underlying emotions, and the impact of facial expressions as a communicative function that mediates social interaction. Studies by Ekman, et al. (1982) indicate that facial expressions can provide accurate information about emotion. Fridlund (1994)

Cite as: Individual Differences in Expressive Response: A Challenge for ECA Design (Short Paper), Ning Wang, Stacy Marsella, Tim Hawkins, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1289-1292.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

argues that expressions do not correlate to underlying emotions and rather has evolved to elicit behaviors from others. He contends that expressions are inherently social.

Research by Ekman (1982) shows that facial expression is a pattern of activities across the face. Darwin (1872) suggested that surprise is a biologically determined facial display consisting of three components: eyebrow raise, widening of the eyes, and opening of the mouth/jaw drop. Other research argues that facial expressions of emotion are more often partial than complete (Carroll, Russell 1997; Reisenzein 2000). Studies by Reisenzein (2006) find that surprise doesn't correspond to the three component display model.

EVG (Wang, Marsella 2006) is built on the ideas first realized in the GAME (Kaiser, Wehrle 1996). As a platform for conducting facial expression experiments, EVG provides us with the opportunity to study these different theories and explore the significance for embodied agents design.

3. EVG: The Emotion Evoking Game

EVG is adapted from a game called Egoboo (2000). It is implemented as a role-playing dungeon adventure game. The current setup includes events targeted to evoke five emotions: boredom, surprise, joy, anger and disappointment, in order. The story in the current study is that the player, accompanied by a teammate (a non-player character), starts out in an underground palace to collect 2000 units of gold. In the end, the player defeats the enemies and successfully collects 2000 units of gold. Then the teammate betrays the player by killing him and stealing the gold. There are five main emotion evoking phrases of this setup. This paper focuses on the stage called "Shock-and-Awe", during which the player faced sudden appearance of powerful enemies for the first time. Detailed descriptions of the other four stages can be found in Wang, Marsella (2006).

4. EVG Study

The focus of the study is emotions and expressions of player during Shock-n-Awe. We had the following hypotheses.

H1: Shock-n-awe event will induce self-reported surprise.

H2a: Subject would display raised eye-brow in response to Shock-n-awe event.

H2b: Subject would display mouth open / jaw drop in response to Shock-n-awe event.

H2c: Subject would display widened eyes in response to Shock-n-awe event.

H3: There is a correlation of self-report of surprise and display of surprise facial expression in response to Shock-n-Awe event.

4.1 Method

Participants: Thirty-five people (40% women, 60% men) participated in this study. They were recruited from craigslist.com and were compensated $20.

Procedure: Subject first read and signed the consent form and then filled out the pre-questionnaire packet. Next, the subject sat in front of the experiment computer and read the following message shown on the welcome screen of EVG:

"*Collect gold in the underground palace. Your goal is to collect 2000 gold. Your name is Louis. Alexis is your team member. Alexis can help you heal. Alexis has the key to the last chamber*."

The subject then went through a training level to get familiar with the game controller. Next the subject started to play EVG. After that, the subject filled out the post-questionnaire packet.

Apparatus: A Vision Research Phantom v10 camera was used to capture facial expression at 240 fps. To produce enough light for the camera, the computer room was lit by 15 floor lamps with 3 100-Watt equivalent florescent light bulbs on each lamp.

Measures: Self-report of appraisal and emotion is measured using five copies of a questionnaire modified from Geneva Appraisal Questionnaire (GAO). Subjects were asked to report five events or moments that he/she felt emotions during the game. Two minutes of subject's facial expression (last two minutes before the game ends) was captured. A certified FACS coder viewed the video and marked appearance of raised eyebrows (AU1 and AU2), widened eyes (AU5) and mouth open/jaw drop (AU25, 26 and 27) after Shock-n-Awe event.

4.2 Result

4.2.1 Testing of Hypothesis

Data from 6 subjects are excluded due to technical difficulties. As a result, data from 29 subjects are reported.

In the post-questionnaire, 65.5% of the subjects reported feeling surprise at Shock-n-awe. Table 1 compares the display of different components of surprise facial expression between all the subjects and those subjects who reported feeling surprise. Out of the three components of the surprise facial expression, mouth open and jaw drop was displayed most often. But only 5 subjects showed widened eyes with low intensity. Interestingly, even though over half of the subjects displayed at least one of the components of the surprise facial expression, no subject showed all three components described by Darwin (1872). In addition, 47.4% of the subjects who reported surprise didn't show any of the three components.

Table 1: Percentage of the subjects that displayed different components of surprise facial expression

	Overall	Reported Surprise
Raised Eyebrow	20.7%	21.1%
Widened Eyes	17.2%	10.5%
Mouth Open/Jaw Drop	41.4%	52.6%
Raised Eyebrow + Widened Eyes	3.4%	0
Raised Eyebrow + Jaw Drop	17.2%	21.1%
Widened Eyes + Jaw Drop	6.9%	10.5%
Any one component	51.7%	52.6%
Any two components	24.1%	31.6%
All three components	0%	0%
None of the three components	48.3%	47.4%

4.2.2 A Closer look at surprise facial expression

In our data, we noticed great richness and dynamics of expressions across all subjects. Figure 1 shows one subject's response to the Shock-n-Awe event. The subject started with slightly parted lips and tightening of the eyebrows as he first walked into the last chamber (frame 0). We noticed a very high percentage of the subjects displayed tightening of the eyebrows at this stage. This could probably due to confusion or the lighting in the room. As the subject in Figure 1 saw the enemy appear, his eyes started to widen (frame 25), followed by raising eyebrows, further tightening of eyebrows and opening his mouth (frame 55). Then, the subject appeared to realize that he is under attack by more powerful enemies. We start to see funneling of the lips and further tightening of eyebrows (frame 110). Next, the subject looked down on his game controller to search for the attack button (frame 215), probably because he's still not very familiar with the controller. After finding the attack button, the subject's inner eyebrows were more relaxed and lips were less funneled (frame 265). As he getting ready to fight the enemy, subject's eyebrows started to raise (frame 295), lips started to tighten (frame 340) then funneled again (frame 370). Gradually, subject returned to a face similar to when he started (frame 425 to 505). All these happened within 506 frames, slightly over 2 seconds.

To further analyze the timing of different components of the facial expression, we annotated the start, apex, sustain and end of each facial expression. Onset is the time between start and apex. Offset is the time between end of sustain and the end. In

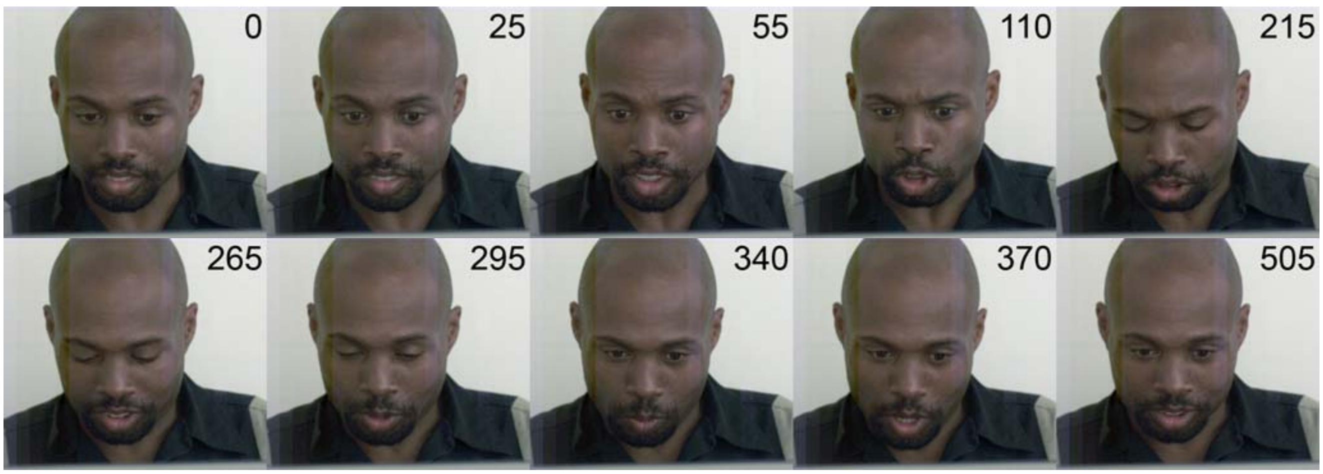

Figure 1. Richness dynamics of facial expression change in response to Shock-n-Awe event

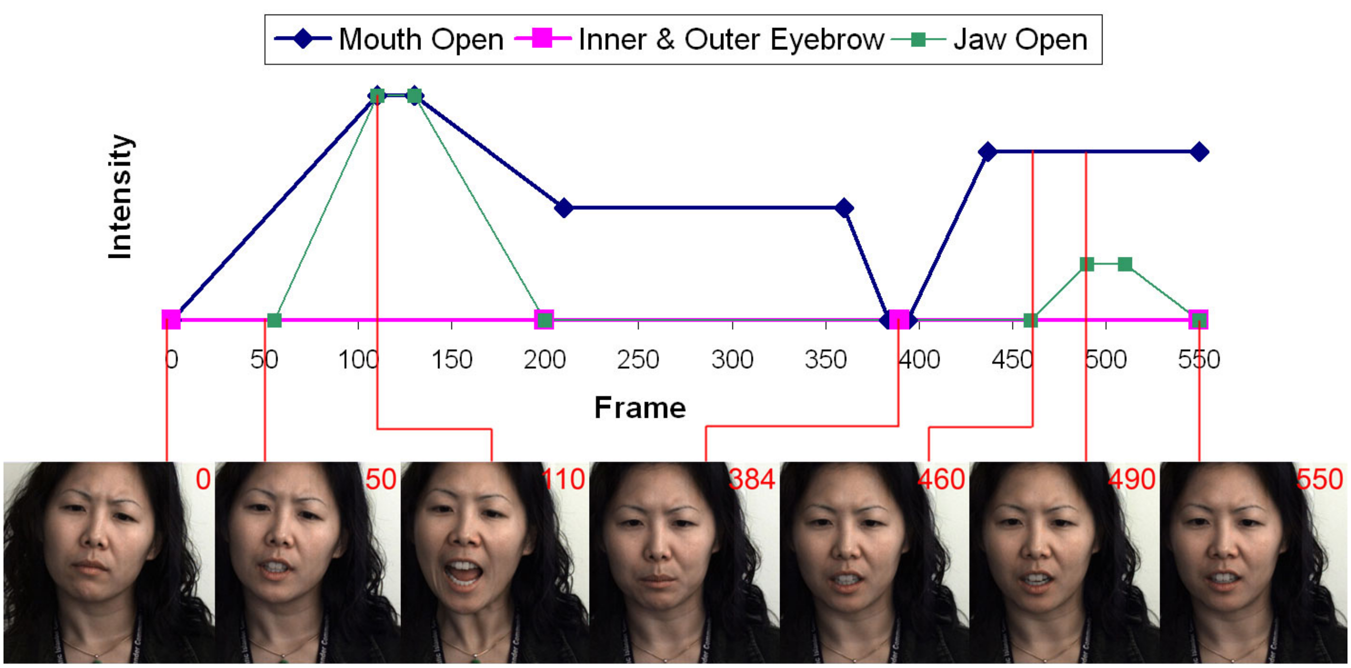

Figure 2. Timing of different components of surprise facial expression in reaction to Shock-n-Awe event

our sample, the average onset of mouth open / jaw drop is .49 seconds. The average onset of eyebrow raise is also about .49 seconds. Both onsets range from 1/10 of a second to just over a second. Even though the average onsets of mouth opening and eyebrow raise are the same, in most cases, the onsets of these two components are different. We didn't have enough data to compute the average onset of eyes widen or the difference between the start of different facial components.

In our data, some subjects "completed" the surprise facial expression. Their face returned to what it was before the surprise expression started. However, there is great diversity in the offset of the surprise facial expression. For the subject in Figure 2, after the expression reaches the apex, the intensity of the components that involved in the expression gradually decrease but don't return to the intensity before they started. They either drop to a lower intensity, stay that intensity for a long time, or other facial components start to take action and change the facial expression. For example, instead of closing the mouth, the open mouth would morph into a smile.

5. Conclusion

Overall, the results reveal considerable differences across subjects in terms of the emotions evoked and expressed. For example, in subjects who reported surprise, 47.4% showed none of Darwin's three components of surprise. And no one showed all three components. Only 17.2% of subjects showed widened eyes. In some cases, instead of showing widen eyes (AU5), subjects showed a decrease in intensity of eye closer (AU43). This means that the eye lids changed from a relaxed more closed state to a relaxed more open state instead of tightened more open state. Sometimes the subject's head is much higher than the computer monitor. This probably made the subject to look down by dropping the eye lids.

Careful study of these high-speed captures reveals remarkable dynamics and variability of the facial expression over time. This suggests that a repository of such capture will be a valuable asset for animating and/or computer modeling of facial expression.

In conclusion, we evaluated EVG and explored the relation between emotion and emotional expression. We found EVG was very successful in evoking emotions and a wide range of facial expression. We found considerable variability between surprise and facial display of surprise. Our data also reveals the highly dynamic nature of facial expression and that emotions are expressed differently from one individual to another. These results suggest that virtual humans can't be one-size-fits-all. We need to design embodied characters in more flexible ways that accommodate and convey individual differences. More importantly, these results argue that more emphasis on the dynamics of facial expression is required and we may need to evolve beyond fixed, canned animations.

6. Acknowledgement

We would like to thank Paul Debevec, Tomas Pereira and Jonathan Gratch from the USC Institute for Creative Technologies. We would also like to thank Intelligent System Division of USC Information Science Institute. This work was sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM), and the content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

7. REFERENCES

[1] Carroll, J. M., Russell, J. A.: Facial expressions in Hollywood's portrayal of emotion. Journal of Personality and Social Psychology, 72, 164-176. (1997)

[2] Cohn, J.F., Schmidt, K., Gross, R., Ekman, P: Individual differences in facial expression: stability over time, relation to self-reported emotion, and ability to inform person identification. Fourth IEEE International Conference on Multimodal Interfaces, 491-496. (2002).

[3] Darwin, C. The expression of the emotions in man and animals. London: Murray. (1872)

[4] Ekman, Paul. Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage. New York: Norton. (1985)

[5] Ekman, P., Rosenberg, E. What the Face Reveals. New York: Oxford University Press. (1997)

[6] Ekman, P.: Emotion in the human face. New York: Cambridge University Press. (1982)

[7] Egoboo. http://zippy-egoboo.sourceforge.net/ (2000)

[8] Fridlund, A. Human Facial Expression: An Evolutionary View. San Diego, CA: Academic Press. (1994)

[9] Frijda, N.H., Kuipers, P., ter Schure, E.: Relations among emotion, appraisal, and emotional action readiness. Journal of Personality and Social Psychology, 57, 212-228. (1989)

[10] IPIP. International Personality Item Pool. http://ipip.ori.org/ipip/ (2007)

[11] Kaiser, S., Wehrle, T.: Situated emotional problem solving in interactive computer games. In Frijda, N.H., (ed.), Proceedings of the VIXth Conference of the International Society for Research on Emotions, 276--280. ISRE Publications (1996)

[12] Lang, P. J., Bradley, M. M., Cuthbert, B. N.: International Affective Picture System (IAPS): Technical manual and affective ratings. Gainsville: Center for Research in Psychophysiology, University of Florida. (1999)

[13] Parkinson, B., Fischer, A. H., & Manstead, A. S. R. Emotion in social relations: Cultural, group, and interpersonal processes. New York: Psychology Press. (2005)

[14] Reisenzein, R., Bördgen, S., Holtbernd, T., Matz, D.: Evidence for strong dissociation between emotion and facial displays: The case of surprise. Journal of Personality and Social Psychology, 91, 295-315. (2006)

[15] Reisenzein, R.: Exploring the strength of association between the components of emotion syndromes: The case of surprise. Cognition and Emotion, 14, 1-38. (2000)

[16] Stemmler, G., Heldmann, M., Pauls, C. A., Scherer, T.: Constraints for emotion specificity in fear and anger: The context counts. Psychophysiology, 38, 275-291. (2001)

[17] Wang, N., Marsella, S. Introducing EVG: An Emotion Evoking Game, Intelligent Virtual Agent (2006)

Another Look at Search-Based Drama Management (Short Paper)

Mark J. Nelson
College of Computing
Georgia Institute of Technology
mnelson@cc.gatech.edu

Michael Mateas
Computer Science Department
University of California, Santa Cruz
michaelm@soe.ucsc.edu

ABSTRACT

A drama manager (DM) is a system that monitors an interactive experience, such as a computer game, and intervenes to keep the global experience in line with the author's goals without decreasing a player's interactive agency. In declarative optimization-based drama management (DODM), an author declaratively specifies desired properties of the experience; the DM intervenes in a way that optimizes the specified metric. The initial DODM approach used online search to optimize an experience-quality function. Later work questioned both online search as a technical approach and the experience-quality optimization framework. Recent work on targeted trajectory distribution Markov decision processes (TTD-MDPs) replaced the experience-quality metric with a metric and associated algorithm based on targeting experience distributions. We show that, though apparently quite different on the surface, the original optimization formulation and TTD-MDPs are actually variants of the same underlying search algorithm, and that offline cached search, as is done by the TTD-MDP algorithm, allows the original search-based systems to achieve similar results to TTD-MDPs. Furthermore, we argue that the original idea of optimizing an experience-quality function does not destroy interactive agency, as had previously been argued, and that in fact it can capture that goal directly.

1. INTRODUCTION

Interactive drama is an interactive experience in which a player interacts with a rich story world in a way that gives a feeling of strong interactive agency while creating, as a result of those interactions, a narrative experience that is dramatic, interesting, and coherent. Putting the player in a story world populated by believable agents does not necessarily create interactive drama: An interactive drama must be designed such that the series of agent-player interactions results in a globally coherent and interesting narrative. A drama manager (DM) is a central coordinator that directs and adapts the agents and other contents of a story world as an experience unfolds to maintain global narrative goals, without removing the player's interactive agency.

One approach is declarative optimization-based drama management (DODM). In DODM, the author specifies a list of the narratively important events that could occur in the experience, called *plot points*; a set of *DM actions* that the DM can take to intervene in the experience; and an *evaluation function* that rates the quality of complete experiences.

Plot points include things such as the player engaging in a particular conversation with an agent in the story world or acquiring an object. They hvae ordering constraints that capture the physical possibilities of the story world. For example, a player cannot interact with a genie in a lamp without having first found the lamp. Plot points are also annotated with information that may be useful to the evaluation function, such as where it happens. DM actions can *cause* a plot point to happen, *hint* to make it more likely that it will happen, *deny* it so it cannot happen, or *undeny* a previously denied plot point. For example, the DM might tell a non-player character to go up to the player and reveal some information, causing the plot point in which the player gains the information. The set of plot points and DM actions, when combined with a player model, provides an abstract, high-level view of the unfolding experience. The evaluation function takes this view of a completed experience and assigns it a rating. The drama manager's job is then to optimize its use of DM actions so as to maximize this evaluation.

The original DODM system, proposed as search-based drama management (SBDM), used a search algorithm to maximize this experience evaluation function [1, 9]. Recent work has questioned both the technical feasibility of search as the optimization method [5], and the conceptual usefulness of having a DM maximize an experience-quality function [8]. In particular, Targeted Trajectory Distribution Markov Decision Processes (TTD-MDPs) have proposed a new goal, with associated algorithms, of targeting an author-specified distribution of experiences [8, 2].

We revisit these criticisms. We show that, although they appear quite different as originally described, the SBDM and TTD-MDP algorithms are actually variants of the same underlying search algorithm. Furthermore, when the original search algorithm is enhanced by caching, as the TTD-MDP one is, it performs at the same level. As a conceptual matter, we argue that the original idea of optimizing an experience-quality function rather than targeting an experience distribution does not destroy player agency, and that to the contrary an experience-evaluation function can directly include interactive agency as a goal, whereas simply adding nondeterminism via TTD-MDPs does not.

Cite as: Another Look at Search-Based Drama Management (Short Paper), Mark J. Nelson and Michael Mateas, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1293-1296.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

```
Build a large tree of possible experience trajectories
for all nodes n in a post-order (leaf-first) traversal do
  if n is terminal then
    n.value ← terminalValue(n)
  else
    n.policy ← optimize(n.actions, n.children)
    n.value ← backup(n.children, n.policy)
  end if
end for
```

Figure 1: Pseudocode for generic cached search. The TTD-MDP and SBDM algorithms share this structure, but differ in how they define terminal values, carry out the optimization, and perform backups.

2. SBDM AND TTD-MDP

DODM was proposed and developed by Bates [1] and Weyhrauch [9] as search-based drama management (SBDM). They proposed a game-tree-search analogy: the player makes "user moves" (*plot points*) through their interaction with the game world, and the DM responds with its own "system moves" (*DM actions*). The DM chooses its "moves" using an author-supplied experience quality function that rates completed experiences, and expectimax search. The expectimax search alternates between maximizing over the available DM actions, and averaging over the possible plot points that could follow, weighted according to a model of likely player behavior. A fairly simple player model is used: the player is assumed to be equally likely to make each of the next possible plot points happen, except for those which have been hinted at, which are considered more likely by a multiplier that the author specifies in an annotation to the hint. To keep things tractable, a sampling search, called SAS+, is used past a certain depth.

Roberts *et al.* [8] proposed a change to the basic formulation. They argued that when the goal is to maximize an evaluation function, the only source of gameplay variation will be unpredictability on the part of the player—and that given sufficiently powerful DM actions, the DM could force an "optimal" story on the player, destroying the truly interactive aspects of the experience. They therefore proposed to start with a desired distribution of experiences (trajectories through the story space), and aim to use the DM actions in a way that would make the actual distribution come as close to that target as possible. Algorithmically, the TTD-MDP system builds a large tree sampled from the space of all possible trajectories; each node in the tree then solves an optimization problem to find a distribution over its available actions that will, according to the player model, cause a resulting distribution over successor plot points that is as close as possible to the distribution specified by the author.

3. OPTIMIZATION BY CACHED SEARCH

SBDM uses an online expectimax search that, to remain computationally tractable, past a cutoff search depth limit switches to sampling trajectories and averaging their evaluations instead of performing full search. The TTD-MDP algorithm [8] operates offline, sampling many possible trajectories through the story world and building them into a tree, and then solving an optimization problem at each node. When a trajectory is seen that wasn't among those sampled in the tree, it falls back to online search. These two algorithms are quite similar when SBDM also uses a tree of cached results. Both build a cached tree, perform an optimization at each node starting from the leaves and working upwards, and back results up the tree, as shown in the generic pseudocode in Figure 1. The main differences are that they choose actions at each node using a different objective function, and assign and back up values to nodes based on different evaluation criteria.

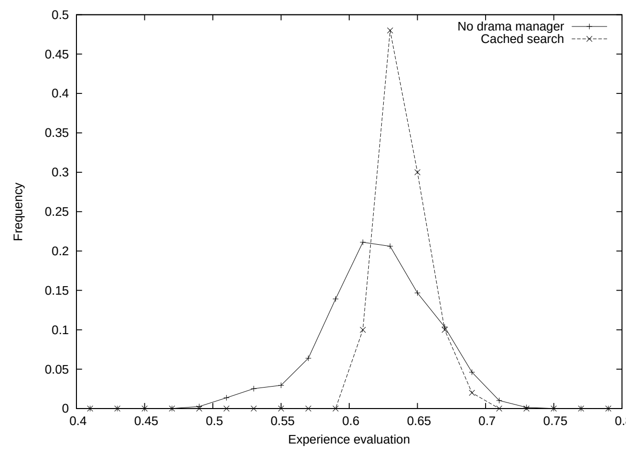

Figure 2: Frequency with which experiences of different qualities (as measured by the evaluation function) occur for a simulated user with a DM guided by iterative-deepening search assuming a minute between plot points, versus a baseline of no DM.

In SBDM, the terminal nodes have their values given by the experience-evaluation function. The policy at each node is to take the DM action that maximizes expected evaluation value when averaged over its children nodes according to the player model. The node's own value is then set to the expected value of this action. In the TTD-MDP algorithm, the terminal nodes have target probabilities as their values. The policy at each node is the distribution over actions that minimizes expected divergence from the target distribution specified by the node's children, with the expectation computed according to the player model. The node's own value is then set to the sum of its children's target probabilities.

Both algorithms can be made to adaptively fill their cached trees during gameplay, using background processor cycles between the occurence of plot points [2]. In fact, once we note the connection with search, we can consider well-known space versus time tradeoffs to avoid literally maintaining a large cache in memory at all. The tree in memory in which we fill in nodes at the frontiers is essentially breadth-first search, which has nice execution-time properties but exponentially large memory requirements. A common alternative is iterative deepening search, which performs a series of fixed-depth depth-first searches with increasing depths, stopping and returning the result of the deepest completed search when the next decision is needed.

Figure 2 shows histograms, both with and without a DM, of the frequency with which experiences of varying quality appear over a number of runs with a simulated player (the same acting-randomly-except-for-hints simulated player used by all previous work), as measured by the authorspecificied evaluation function. The DM in this setup uses a "synthetic" set of DM actions consisting of a causer, denier, and reenabler for every possible plot point; this was the hypothetical maximally powerful setup in which Nelson & Mateas [5] found that search still could not work well. One curve shows the results without a DM, and the other with the iterative-deepening DM. As can be seen by the fact that the curves move towards the right—indicating more frequent highly-rated stories and less frequent low-rated stories—cached search, as a technical matter, functions well in this story, contrasting with the previous results.

4. WHAT TO OPTIMIZE

Since the two algorithms operate similarly, the main question in deciding between SBDM and TTD-MDPs is sorting out what it is a DM should be optimizing: what constitutes a good interactive drama? The main goal is a narratively interesting and coherent experience with strong player agency.

4.1 Maximizing experience quality

DODM envisions an evaluation function that, given a completed experience (a sequence of plot points and DM actions), will rate it based on various features that the author thinks the experience should have had. This function rates the quality of *interactive experiences*, not the quality of plot-point sequences considered as if they were non-interactive stories. That is, DODM does not create interactive drama by taking a set of desiderata for *non-interactive* drama and trying to bring it about in the face of interactivity. Rather, it takes a set of desiderata for the interactive dramatic experience itself, and tries to maintain those. Some DM systems do frame the drama-management problem as one of mediating between authorial narrative goals and player freedom [10, 4, 6]. In DODM, however, the DM starts with a more general notion of what constitutes a narratively interesting experience, and intervenes when necessary to make sure the player has one.

Looking in particular at Weyhrauch's evaluation function, it specifies a number of weighted features that capture his notion of a good experience in his *Tea for Three* story world.

One group of features mainly encourages narrative coherence: *thought flow* prefers stories where subsequent actions relate to each other; *activity flow* prefers stories that have some spatial locality of action; and *momentum* prefers certain pairs of plot points that build on each other well. Separately, the *motivation* feature prefers stories in which at least some plot points are motivated by previous plot points. Note that these are preferences for the interactive experience, and would not necessarily be the same if evaluating a linear story. Weyhrauch doesn't argue that it's necessarily bad for narratives to have the action move around frequently between different locations; rather, he argues that if each plot point happens in a different location from the last in an interactive experience, it was probably the case that the experience contained a lot of uninteresting wandering around the world.

Given only these features, however, there is a danger that the system could identify certain plot-point progressions as ideal and force the player into them, defeating the goal of interactive agency. To avoid this outcome, two versions of an additional evaluation feature, one proposed by Weyhrauch and one by Nelson & Mateas, aim at encoding interactive agency, though from different perspectives.

Weyhrauch's *options* feature identifies twelve meaningful goals a player could have at various points in *Tea for Three*. For example, the goal "talk to George about the new will" is considered to be active between the time the player finds a note mentioning a new will and the time that the player either talks to George about it or is prevented from doing so by other events. The number of goals active at any given time is a rough measure of the degree of interactive agency available. The *options* feature encodes a preference for many such meaningful options to be available towards the beginning of the game, decreasing to fewer towards the end.

Nelson & Mateas's *choices* feature captures a more local notion of agency, measuring how many plot points could have followed any given point in the story, given the ordering constraints in the world and the effects of causers and deniers. If at some point only one plot point could possibly come next, then the same bit of story would play out next regardless of what the player did. If on the other hand many plot points could come next, then the player could locally influence the story to a much greater extent. The *choices* feature has the advantage that it can be computed automatically for any story, but the *options* feature has the advantage that it captures a higher-level notion of meaningful interactive agency.

Finally, a *manipulativity* feature penalizes uses of DM actions that are likely to be particularly noticeable, like clumsy hints or moving objects that the player can see. This is a meta-feature of sorts encoding a preference for the DM's operation to be invisible. Although we use agents in service of a narrative rather than merely simulating them as believable agents in their own right, we do still want them to avoid doing things that would break believability.

4.2 Targeting an experience distribution

Roberts *et al.* [8] criticize the idea of maximizing a story-quality function, arguing that an effective DM can simply bring about the same highly-rated story each time, destroying interactive agency and replayability. They propose instead that the goal of the DM should be to target a distribution of experiences, specified either by some mapping from an evaluation function (*e.g.* bad experiences should never happen, and good ones should happen in proportion to their quality), or by having the author specify a few prototype experiences and then targeting a distribution over experiences similar but not identical to the prototypes [7].

Since the goal of DODM is to maximize experience quality rather than story quality, though, an evaluation function should measure not only the quality of the story that a series of interactions produces, but also the quality of the interaction itself, including elements such as interactive agency; hence the *options* and *choices* features. Moreover, targeting a distribution of experiences does not necessarily coerce the player less than even targeting a single maximum-quality story would. With enough causers and deniers, an TTD-MDP system can directly cause its desired distribution of experiences to come about, by randomly selecting (according to the desired distribution) which DM actions to take in each play-through. Although that would vary *which* story the player is forced into each time, it still uses the DM actions to produce a specific story with no input from the player—randomly selecting a different story to force the user into each time does not create interactive agency.

Indeed we find similar levels of coerciveness if we look at the DM actions performed by the TTD-MDP based system and the SBDM system on the version of *Anchorhead* with a "synthetic" set of DM actions that Roberts *et al.* use as a point of comparison. The "synthetic" set of actions consists of a causer, denier, and reenabler for every possible plot point in the story, thus creating a hypothetical situation where the DM has a maximally powerful set of actions available. The TTD-MDP system claimed better replayability in this case, since it produced a wider variety of stories. However, both the TTD-MDP system and the SBDM system acted almost maximally coercively: they each performed an average of around 15 DM actions per experience, in an experience only 16 plot-points long. The TTD-MDP system varied which specific coercion it performed from run to run, but that again does not constitute interactive agency, which requires that the player, rather than system nondeterminism, be able to meaningfully influence the outcome.

That both systems are quite coercive, however, does point to a failure in the particular experience-quality evaluation function that both used. We can correct this by simply putting a greater weight on the *choices* feature, which emphasizes that giving the player many choices in what to do really is an important part of an interactive experience. When we increase *choices* from being 15% of the total evaluation weight to 50%, both systems drop to using an average of around 5 DM actions per experience.

How to best write evaluation functions does remain an issue that would benefit from additional experimentation in the context of specific real interactive dramas. It is worth noting that all the recent systems have focused on the "synthetic" model of *Anchorhead* that has only causers, deniers, and reenablers, and lacks the hint actions that a DM could use to provide more narrative guidance to the player without unduly removing interactive agency; by contrast, a real application would likely use hints frequently.

Whether the TTD-MDP formulation does still improve matters in a different way depends on the particular way in which the target distribution is defined, and on what we consider to be the goals of interactive drama. In the case where the target distribution is generated by a mapping from an experience-quality function, the results will be fairly similar to the results from an evaluation-function-maximizing approach, since both systems will be trying to avoid low-rated experiences and increase the probability of highly-rated ones according to the same function. The TTD-MDP approach will add some more nondeterminism in doing so; how much depends on how the mapping is constructed. Alternate ways of specifying a target distribution of experiences for TTD-MDPs, however, such as specifying several prototype experiences and inducing a distribution over experiences similar to those prototypes [7], suffer from a greater loss of interactive agency. If the player is being forced into one of several prototype experiences or minor variants thereof, the fact that the specific experience they're forced into is chosen nondeterministically does not preserve interactive agency.

4.3 Non-dramatic interactive experiences

We focus on authoring interactive drama. Similar techniques can be used for other kinds of interactive experiences, which may have different considerations. For example, we argue that in interactive drama, a DM shouldn't be seen as balancing externally imposed constraints with a player's freedom of action, but rather as a system that helps to ensure that there is enough narrative for the player to have a coherent and interesting experience.

Other experiences, however, may have genuinely external constraints that must be imposed in a way that could conflict with the user's freedom and goals. For example, a TTD-MDP based system was proposed for guiding museum tours [3]. In that domain, the goal of reducing congestion really is an external goal imposed on the visitors, and is reasonably expressed by targeting a specific distribution of experiences so as to keep visitors nicely spread out.

5. CONCLUSIONS

By separating the issues of what to optimize and how to carry out the optimization, we showed that the algorithms used by targeted trajectory Markov decision processes (TTD-MDPs) and by search-based drama management (SBDM) are versions of a generic search-based algorithm to which caching or offline computation may be added separately from the consideration of what to optimize.

On the conceptual issue, we defended the original formulation of drama management that sought to maximize an experience-quality function. We pointed out that experience-quality functions are not equivalent to story-quality functions that rate experiences as if they were non-interactive narratives, but are rather functions that explicitly take into account elements of a good interactive experience, such as the notion of interactive agency. We showed that TTD-MDPs, by contrast, primarily serve to add nondeterminism to their actions, which is a separate concern from interactive agency and does not necessarily produce agency.

6. REFERENCES

[1] J. Bates. Virtual reality, art, and entertainment. *Presence*, 1(1):133–138, 1992.

[2] S. Bhat, D. Roberts, M. J. Nelson, C. L. Isbell, and M. Mateas. A globally optimal algorithm for TTD-MDPs. In *Proceedings of AAMAS*, 2007.

[3] A. Cantino, D. L. Roberts, and C. L. Isbell. Autonomous nondeterministic tour guides: Improving quality of experience with TTD-MDPs. In *Proceedings of AAMAS*, 2007.

[4] B. Magerko. Story representation and interactive drama. In *Proceedings of AIIDE*, 2005.

[5] M. J. Nelson and M. Mateas. Search-based drama management in the interactive fiction Anchorhead. In *Proceedings of AIIDE*, 2005.

[6] M. O. Riedl, C. J. Saretto, and R. M. Young. Managing interaction between users and agents in a multi-agent storytelling environment. In *Proceedings of AAMAS*, 2003.

[7] D. L. Roberts, S. Bhat, K. St. Clair, and C. L. Isbell. Authorial idioms for target distributions in TTD-MDPs. In *Proceedings of AAAI*, 2007.

[8] D. L. Roberts, M. J. Nelson, C. L. Isbell, M. Mateas, and M. L. Littman. Targeting specific distributions of trajectories in MDPs. In *Proceedings of AAAI*, 2006.

[9] P. Weyhrauch. *Guiding Interactive Drama*. PhD thesis, Carnegie Mellon University, 1997.

[10] R. M. Young, M. O. Riedl, M. Branly, A. Jhala, R. J. Martin, and C. J. Saretto. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development*, 1(1), 2004.