

A Mixed-Initiative Conversational Dialogue System for Healthcare

Fabrizio Morbini and Eric Forbell and David DeVault and Kenji Sagae and David R. Traum and Albert A. Rizzo

Institute for Creative Technologies
University of Southern California
Los Angeles, CA 90094, USA

{morbini, forbell, devault, sagae, traum, rizzo}@ict.usc.edu

Abstract

We present a mixed initiative conversational dialogue system designed to address primarily mental health care concerns related to military deployment. It is supported by a new information-state based dialogue manager, FLoReS (Forward-Looking, Reward Seeking dialogue manager), that allows both advanced, flexible, mixed initiative interaction, and efficient policy creation by domain experts. To easily reach its target population this dialogue system is accessible as a web application.

1 Introduction

The SimCoach project is motivated by the challenge of empowering troops and their significant others in regard to their healthcare, especially with respect to issues related to the psychological toll of military deployment. SimCoach virtual humans are not designed to act as therapists, but rather to encourage users to explore available options and seek treatment when needed by fostering comfort and confidence in a safe and anonymous environment where users can express their concerns to an artificial conversational partner without fear of judgment or possible repercussions.

SimCoach presents a rich test case for all components of a dialogue system. The interaction with the virtual human is delivered via the web for easy access. As a trade-off between performance and quality, the virtual human has access to a limited set of pre-rendered animations.

The Natural Language Understanding (NLU) module needs to cope with both chat and military

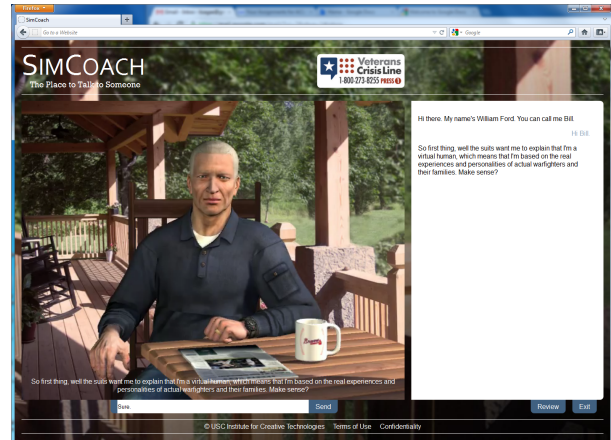


Figure 1: Bill Ford, a SimCoach character. SimCoach virtual humans are accessible through a web browser. The user enters natural language input in the text field on the bottom of the screen. The simcoach responds with text, speech and character animation. The text area to the right shows a transcript of the dialogue.

slang and a broad conversational domain. The dialogue policy authoring module needs to support non-dialogue experts given that important parts of the dialogue policy are contributed by experts in psychometrics and mental health issues in the military, and others with familiarity with the military domain.

The dialogue manager (DM) must be able to take initiative when building rapport or collecting the information it needs, but also respond appropriately when the user takes initiative.

2 Supporting Mixed Initiative Dialogues

There is often a tension between system initiative and performance of the system's decision-making for understanding and actions. A strong system-initiative policy reduces the action state space since

user actions are only allowed at certain points in the dialogue. System initiative also usually makes it easier for a domain expert to design a dialogue policy that will behave as desired.¹ Such systems can work well if the limited options available to the user are what the user wants to do, but can be problematic otherwise, especially if the user has a choice of whether or not to use the system. In particular, this approach may not be well suited to an application like SimCoach. At the other extreme, some systems allow the user to say anything at any time, but have fairly flat dialogue policies, e.g., (Leuski et al., 2006). These systems can work well when the user is naturally in charge, such as in interviewing a character, but may not be suitable for situations in which a character is asking the user questions, or mixed initiative is desired.

True mixed initiative is notoriously difficult for a manually constructed call-flow graph, in which the system might want to take different actions in response to similar stimuli, depending on local utilities. Reinforcement learning approaches (Williams and Young, 2007; English and Heeman, 2005) can be very useful at learning local policy optimizations, but they require large amounts of training data and a well-defined global reward structure, are difficult to apply to a large state-space and remove some of the control, which can be undesirable (Paek and Pieraccini, 2008).

Our approach to this problem is a forward-looking reward seeking agent, similar to that described in (Liu and Schubert, 2010), though with support for complex dialogue interaction and its authoring. Authoring involves design of local subdialogue networks with pre-conditions and effects, and also qualitative reward categories (goals), which can be instantiated with specific reward values. The dialogue manager, called FLoReS, can locally optimize policy decisions, by calculating the highest overall expected reward for the best sequence of subdialogues from a given point. Within a subdialogue, authors can craft the specific structure of interaction.

Briefly, the main modules that form FLoReS are:

- The **information state**, a propositional knowl-

¹Simple structures, such as a call flow graph (Pieraccini and Huerta, 2005) and branching narrative for interactive games (Tavinor, 2009) will suffice for authoring.

edge base that keeps track of the current state of the conversation. The information state supports missing or unknown information by allowing atomic formulas to have 3 possible values: true, false and null.

- A set of **inference rules** that allows the system to add new knowledge to its information state, based on logical reasoning. Forward inference facilitates policy authoring by providing a mechanism to specify information state updates that are independent of the specific dialogue context.²
- An **event handling** system, that allows the information state to be updated based on user input, system action, or other classes of author-defined events (such as system timeouts).
- A set of **operators**. Operators represent local dialogue structure (trees), and can also be thought of as reusable subdialogues. Each state within the subdialogue can include a *reward* for reaching that state. Rewards are functions of the goals of the system, and are the main method used to decide what to do when there is more than one applicable operator. Operators have preconditions and effects. Effects specify changes to the information state. The preconditions define when an operator can be activated.

3 Sample Dialogue

In this demo, the user will interact with the SimCoach character Bill Ford, using a standard web browser and typing text. The virtual human, driven by FLoReS, will respond using pre-rendered animations encoded as H.264 video, delivered via a standard web server. Table 1 shows an excerpt from a sample conversation with Bill Ford that illustrates some of the features of this dialogue manager.

The excerpt starts from a rapport building smalltalk sub-dialogue on the topic of barbecuing which is interrupted by a user question about confidentiality. The system responds to the user interruption and then re-starts the interrupted smalltalk because it is still the most valuable conversation continuation available at that moment.

²For example: every time the user says that s/he has nightmares we want to update the information state to include that s/he also has sleeping problems.

Dialogue transcript	Notes
<p>BBQ Smalltalk</p> <p>Ask anybody about me, and they'll tell you that I love to BBQ</p> <p><i>Is this conversation secret?</i></p> <p>Confidentiality QA</p> <p>We don't share your info with anyone who can personally identify you. The techs can see what we say, but just to tell that the site is working. But they have no idea who said it, just what was said</p> <p>Did that help you?</p> <p><i>Yes it did.</i></p> <p>Great.</p>	<p>The character is equipped with a few operators for smalltalk about a few topics. BBQ is one of them.</p> <p>Here the system is interrupted by a user question and it decides that answering it is the best course of action.</p>
<p>BBQ Smalltalk</p> <p>Like I was saying, I love to BBQ</p> <p><i>What is PTSD?</i></p> <p>What is PTSD QA</p> <p>PTSD, or post-traumatic stress disorder is an anxiety condition associated with serious traumatic events. It can come with survivor guilt, reliving the trauma in dreams, numbness, and lack of involvement with reality.</p>	<p>After answering the question, the best course of action is to awaken the paused operator about the BBQ smalltalk.</p> <p>Again the BBQ smalltalk is interrupted by another question from the user.</p>
<p>PTSD Topic Interest QA</p> <p>So, is PTSD something you're worried about. I only ask, because you've been asking about it. . . .</p>	<p>After answering the second question the system decides to ignore the paused operator and load a follow-up operator related to the important topic raised by the user's question. The selection is based on the expected reward that talking about PTSD can bring to the system.</p>

Table 1: An excerpt of a conversation with Bill Ford that shows opportunistic mixed initiative behavior.

Next, the user asks a question about the important topic of post-traumatic stress disorder (PTSD). That allows operators related to the PTSD topic to become available and at the next chance the most

rewarding operator is no longer the smalltalk sub-dialogue but one that stays on the PTSD topic.

4 Conclusion

We described the SimCoach dialogue system which is designed to facilitate access to difficult health concerns faced by military personnel and their families. To easily reach its target population, the system is available on the web. The dialogue is driven by FLoReS, a new information-state and plan-based DM with opportunistic action selection based on expected rewards that supports non-expert authoring.

Acknowledgments

The effort described here has been sponsored by the U.S. Army. Any opinions, content or information presented does not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

- M.S. English and P.A. Heeman. 2005. Learning mixed initiative dialogue strategies by using reinforcement learning on both conversants. In *HLT-EMNLP*.
- Anton Leuski, Ronakkumar Patel, David Traum, and Brandon Kennedy. 2006. Building effective question answering characters. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, pages 18–27.
- Daphne Liu and Lenhart K. Schubert. 2010. Combining self-motivation with logical planning and inference in a reward-seeking agent. In Joaquim Filipe, Ana L. N. Fred, and Bernadette Sharp, editors, *ICAART (2)*, pages 257–263. INSTICC Press.
- Tim Paek and Roberto Pieraccini. 2008. Automating spoken dialogue management design using machine learning: An industry perspective. *Speech Communication*, 50(89):716 – 729. Evaluating new methods and models for advanced speech-based interactive systems.
- Roberto Pieraccini and Juan Huerta. 2005. Where do we go from here? Research and commercial spoken dialog systems. In *Proceedings of the 6th SIGdial Workshop on Discourse and Dialogue*, Lisbon, Portugal, September.
- Grant Tavinor. 2009. *The art of videogames*. New Directions in Aesthetics. Wiley-Blackwell, Oxford.
- J.D. Williams and S. Young. 2007. Scaling POMDPs for spoken dialog management. *IEEE Trans. on Audio, Speech, and Language Processing*, 15(7):2116–2129.