Goal-Driven Approach To Open-Ended Dialogue Management using BDI Agents

(Extended Abstract)

Wilson Wong, Lawrence Cavedon, John Thangarajah and Lin Padgham
School of Computer Science and IT, RMIT University, Melbourne, Australia
{wilson.wong,lawrence.cavedon,john.thangarajah,lin.padgham}@rmit.edu.au

ABSTRACT

We describe a BDI (Belief, Desire, Intention) approach and architecture for a conversational virtual companion embodied as a child's Toy. Our aim is to support both structured conversation-based activities (e.g., story-telling, collaborative games) as well as more free-flowing, engaging dialogue with variation and some unpredictability. We argue that a goal-oriented approach to the agent's conversational capabilities provides these competing capabilities.

Keywords

BDI architecture, Dialogue management, Conversational agent

1. OVERVIEW

We propose a BDI architecture as shown in Figure 1 for a conversational agent that supports both task-oriented dialogue as well as "chatty" conversations. The BDI agent model has been used successfully in a range of applications requiring a mix of reactive behaviour and goal-directed reasoning, and its design model supports different means for achieving a goal depending on context and other factors [3]. The BDI framework thus allows the conversational agent to select different strategies for satisfying a conversational goal where a conversational goal may involve playing a collaborative game such as a role-play, satisfying a request from the user such as an information request, or simply conversing with the user about a pertinent topic. BDI agent-based approaches to dialogue management have been previously proposed [2]. However, these have typically been for task-oriented conversations where the outcome was to support the user in performing a given task (e.g., accessing email). A significant novelty of our use of the BDI approach is to provide multiple plans to satisfy a given goal (e.g., chat, engage in a shared role-play), and to support variation in the way that goal is then achieved including enforcing variation in the agent's contributions to the conversation. When interacting with the child, the Toy suggests possible *Conversational Activities* such as a cooking game, a story, a quiz, etc. These activities are represented as goal-plan structures, which are a set of plan templates in the Toy's `Plan Library`. These plans are used to guide the different aspects of the activity and the selection of fragments for the Toy to utter in pursuit of that activity. More importantly, the specific utterances are **not** part of the activity structure. The plans can provide contextual information which is then used by the `Conversation Manager` to select the appropriate outputs from the `Fragment Library`. The goal-plan tree which is induced by the `Plan Library` gives a structure that is essentially an AND/OR tree. This provides a large number of possible executions within a relatively compact structure [3]. It is this which we exploit to achieve the desired variability, while retaining a coherent, goal-oriented dialogue. In Figure 2, we show a partial goal-plan tree for a particular activity in the Toy, namely, a cooking role-play activity. The top-level goal has a single plan which guides the struc-

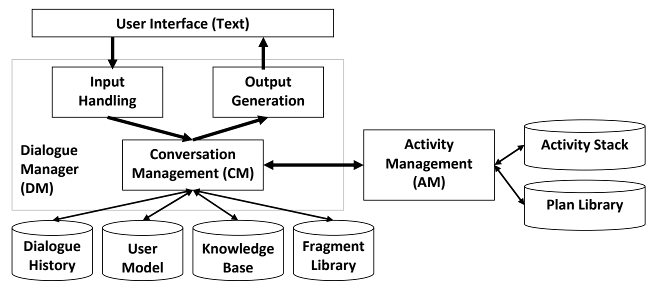

Figure 1: System architecture

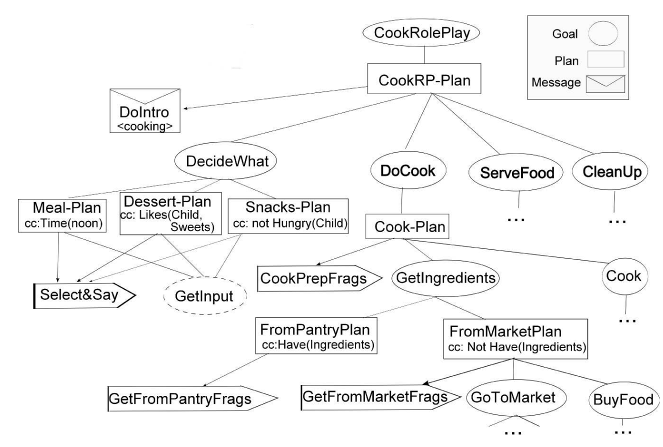

Figure 2: Example activity: Cooking role play

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

ture of the activity. It is possible to have different plans to choose from at the top level to provide more variety. This plan has a sequential set of subgoals, each of which has a set of plans to choose from, and so on. We see that the first goal is `DoIntro` which is a goal which carries information about the current activity (Cooking) and triggers a plan in the `Conversation Manager` to select a suitable introductory fragment for this activity, and prepend it to the next system output (i.e., it is not a fragment with any expected response). Following this is a choice of plans, one of which will be selected in any given conversation. It will then do a number of things such as decide how many interactions to have in this subgoal. Importantly, it will provide some additional keyphrases to be added to the ongoing collection from the dialogue history to assist in fragment selection. It initiates an `Interact` goal which results in the `Conversation Manager` determining an output fragment and analysing the user response, which is then provided back to the plan in the form of keyphrases and a response category. Assuming the response is accepted, when the plan has completed its interactions, it decides (based on the keywords collected) what food it believes is going to be prepared and the activity progresses onto the next subgoal `CookProcess`, which is managed in a similar way. Activity suspension, resumption and abortion, which are the responsibility of the `Activity Manager`, are not discussed here due to space constraints.

2. EXAMPLE AND DISCUSSIONS

Figure 3 shows an example interaction with the Toy that demonstrates a number of the features discussed in the previous sections. We highlight the features of this interaction here. Items 2 to 12 are mainly about the dynamic building of context for fragment selection. Item 13 shows an example of activity selection based on relatedness. Items 13 to 15 illustrates several interactions guided by the goal-plan activity structure. Item 16 is a specific user request which results in the suspension of the current activity. Items 23 to 24 shows the resumption of the suspended activity together with a transition fragment. A brief suspension and resumption is shown in Items 27 to 29. Finally, an activity finishes and a new activity is selected in Item 35.

We performed an initial evaluation of the Toy involving nine human users with a total of 229 system utterances (25 on average) and 162 user inputs (18 on average). The users are asked to rate each of the Toy's responses on a Likert scale of 0 (i.e., *very odd response*) to 4 (i.e., *very sensible response*). The ratings of the majority of the utterances tend towards level 3 (i.e., *sensible*) with a standard deviation between 0.71 to 1.12. At the end of each 10-minute interaction, the users were asked to rate on two 5-point Likert scales, one on overall coherence (*very scattered* versus *very coherent*), and the other regarding variability (*too repetitive* to *good variation*). 7 out of the 9 ratings were on level 3 for *reasonable variation*, with one rating each at levels 2 and 4. However, 4 out of the 9 participants selected level 2 (i.e. *scattered*) for overall coherence. The results were promising even though the coverage of the `Knowledge Base` is currently limited. Our analysis of the evaluation data suggests we could expect increase in overall user satisfaction by improving input processing, extending the amount of content, and improving the coherence measure for new-topic selection.

Our future work includes extending the capabilities listed above together with evaluating the system in the context of noisy automated speech recognition. We have been mining social forums to automatically extend our `Knowledge Base` [4]. More relevant to the goal-oriented approach, we are developing a model of *proactive engagement*, whereby the conversational agent actively monitors user engagement [1] and applies conversational strategies to increase engagement within the current or a new activity.

```
01: USER: hi.
02: SYSM: hello. what have you been doing?
03: USER: i just had chocolate.
04: SYSM: do you like chocolate?
05: USER: i love chocolate.
06: SYSM: i love it too. i wish i could eat it
          for breakfast...
07: USER: chocolate for breakfast? yucks!
08: SYSM: what do you eat for breakfast instead
          of chocolate?
09: USER: toast.
10: SYSM: what do you like to eat on your toast.
11: USER: ice cream.
12: SYSM: yum yum! delicious.
13: SYSM: do you want to make some ice cream?
14: USER: yeah.
15: SYSM: do you have the recipe with you?
16: USER: can you please tell me a story?
17: SYSM: sure. susie the ice cream moose was...
...
23: SYSM: that's the end of the story.
24: SYSM: let's continue with our cooking. so, do
          you have the recipe with you?
25: USER: no.
26: SYSM: you'll need milk and chocolate.
27: USER: where does chocolate come from?
28: SYSM: chocolate comes from cacao trees.
29: SYSM: let's continue with our cooking, do you
          need to get out a saucepan or something?
...
34: USER: it's done!
35: SYSM: yum yum! your ice cream is ready.
36: SYSM: now, let's talk about zoo, ok?
37: USER: i like visiting the zoo.
38: SYSM: do you go to the zoo often?
...
```

Figure 3: An example execution trace.

Acknowledgment

This work is partially supported by the Australian Research Council and Real Thing Entertainment Pty. Ltd. under Linkage grant number LP0882013.

3. REFERENCES

[1] G. Castellano, A. Pereira, I. Leit, A. Paiva, and P. McOwan. Detecting user engagement with a robot companion using task and social interaction-based features. In *Proceedings of the 11th ICMI*, Cambridge, USA, 2009.

[2] A. Nguyen and W. Wobcke. An agent-based approach to dialogue management in personal assistants. In *Proceedings of the IUI Conference*, San Diego, California, 2005.

[3] L. Padgham and M. Winikoff. *Developing Intelligent Agent Systems: A practical guide*. Wiley Series in Agent Technology. John Wiley and Sons, 2004.

[4] W. Wong, J. Thangarajah, and L. Padgham. Health conversational system based on contextual matching of community-driven question-answer pairs. In *Proceedings of the 20th ACM CIKM (Demo)*, pages 2577–2580, Glasgow, UK, 2011.