

The Hidden Information State Dialogue Manager: A Real-World POMDP-Based System

Steve Young, Jost Schatzmann, Blaise Thomson, Karl Weilhammer, Hui Ye

Cambridge University Engineering Department

Trumpington Street, Cambridge, CB21PZ, United Kingdom

{sjy, js532, brmt2, kw278, hy216}@eng.cam.ac.uk

Abstract

The Hidden Information State (HIS) Dialogue System is the first trainable and scalable implementation of a spoken dialog system based on the Partially-Observable Markov-Decision-Process (POMDP) model of dialogue. The system responds to n-best output from the speech recogniser, maintains multiple concurrent dialogue state hypotheses, and provides a visual display showing how competing hypotheses are ranked. The demo is a prototype application for the Tourist Information Domain and achieved a task completion rate of over 90% in a recent user study.

1 Partially Observable Markov Decision Processes for Dialogue Systems

Recent work on statistical models for spoken dialogue systems has argued that Partially Observable Markov Decision Processes (POMDPs) provide a principled mathematical framework for modeling the uncertainty inherent in human-machine dialogue (Williams, 2006; Young, 2006; Williams and Young, 2007). Briefly speaking, POMDPs extend the traditional fully-observable Markov Decision Process (MDP) framework by maintaining a *belief state*, ie. a probability distribution over dialogue states. This enables the dialogue manager to avoid and recover from recognition errors by sharing and shifting probability mass between multiple hypotheses of the current dialogue state. The framework also naturally

incorporates n-best lists of multiple recognition hypotheses coming from the speech recogniser.

Due to the vast number of possible dialogue states and policies, the use of POMDPs in practical dialogue systems is far from straightforward. The size of the belief state scales linearly with the number of dialogue states and belief state updates at every turn during a dialogue require all state probabilities to be recomputed. This is too computationally intensive to be practical with current technology. Worse than that, the complexity involved in policy optimisation grows exponentially with the number of states and system actions and neither exact nor approximate algorithms exist that provide a tractable solution for systems with thousands of states.

2 The Hidden Information State (HIS) Dialogue Manager

The Hidden Information State (HIS) dialogue manager presented in this demonstration is the first trainable and scalable dialogue system based on the POMDP model. As described in (Young, 2006; Young et al., 2007) it partitions the state space using a tree-based representation of user goals so that only a small set of partition beliefs needs to be updated at every turn. In order to make policy optimisation tractable, a much reduced summary space is maintained in addition to the master state space. Policies are optimised in summary space and the selected summary actions are then mapped back to master space to form system actions. Apart from some very simple ontology definitions, the dialog manager has no application dependent heuristics.

The system uses a grid-based discretisation of the



Figure 1: The HIS Demo System is a Tourist Information application for a fictitious town

state space and online ϵ -greedy policy optimisation. While this offers the potential for online adaptation with real users at a later stage, a simulated user is needed to bootstrap the training process. A novel agenda-based simulation technique was used for this step, as described in (Schatzmann et al., 2007).

3 The HIS Demo System

The HIS demo system is a prototype application for the Tourist Information domain. Users are assumed to be visiting a fictitious town called “Jasonville” (see Fig. 1) and need to find a suitable hotel, bar or restaurant subject to certain constraints. Examples of task scenarios are “*finding a cheap Chinese restaurant near the post office in the centre of town*” or “*a wine bar with Jazz music on the riverside*”. Once a venue is found, users may request further information such as the phone number or the address.

At run-time, the system provides a visual display (see Fig. 2) which shows how competing dialogue state hypotheses are being ranked. This allows developers to gain a better understanding of the internal operation of the system.

4 Demo System Performance

In a recent user study the demo system was evaluated by 40 human subjects. In total, 160 dialogues were recorded with an average Word-Error-Rate of 29.8%. The performance of the system was measured based on the recommendation of a correct venue and achieved a task completion rate of 90.6% with an average number of 5.59 dialogue turns to completion (Thomson et al., 2007).

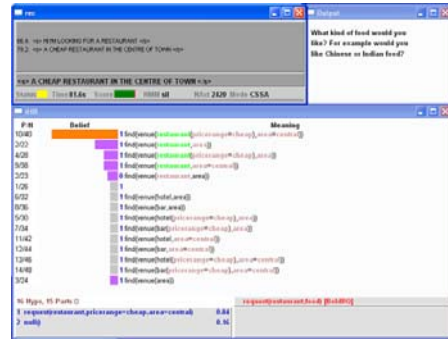


Figure 2: A system screenshot showing the ranking of competing dialogue state hypotheses

The results demonstrate that POMDPs facilitate design and implementation of spoken dialogue systems, and that the implementation used in the HIS dialogue manager can be scaled to handle real world tasks. The user study results also show that a simulated user can be successfully used to train a POMDP dialogue policy that performs well in experiments with real users.

5 Accompanying materials

The demo system and related materials are accessible online at our website

<http://mi.eng.cam.ac.uk/research/dialogue/>

References

- J. Schatzmann, B. Thomson, K. Weilhammer, H. Ye, and S. Young. 2007. Agenda-Based User Simulation for Bootstrapping a POMDP Dialogue System. In *Proceedings of HLT/NAACL, Rochester, NY*.
- B. Thomson, J. Schatzmann, K. Weilhammer, H. Ye, and S. Young. 2007. Training a real-world POMDP-based Dialogue System. In *Proceedings of Bridging the Gap: Academic and Industrial Research in Dialog Technology, Workshop at HLT/NAACL, Rochester, NY*.
- J. D. Williams and S. Young. 2007. Partially Observable Markov Decision Processes for Spoken Dialog Systems. *Computer Speech and Language*, 21(2):231–422.
- J. D. Williams. 2006. *Partially Observable Markov Decision Processes for Spoken Dialogue Management*. Ph.D. thesis, University of Cambridge.
- S. Young, J. Schatzmann, K. Weilhammer, and H. Ye. 2007. The Hidden Information State Approach to Dialog Management. In *Proc. of ICASSP (forthcoming)*, Honolulu, Hawaii.
- S. Young. 2006. Using POMDPs for Dialog Management. In *Proc. of IEEE/ACL SLT, Palm Beach, Aruba*.