

Towards Adaptive Spoken Dialogue Agents

Diane J. Litman

AT&T Labs - Research
180 Park Avenue
Florham Park, NJ 07932
diane@research.att.com

Abstract

Although it is now possible to build real-time spoken dialogue systems for a wide variety of applications, designing the dialogue strategy of such systems involves a number of nontrivial design choices. These choices can seriously impact system performance, and include choosing appropriate strategies for initiative (whether to accept relatively open-ended versus constrained user utterances) and confirmation (when to confirm a user's previous utterance). Typically, dialogue strategy design has been done in an ad-hoc manner. This paper summarizes how my recent research has addressed the design, implementation, and experimental evaluation of methods for dynamically adapting and optimizing a system's dialogue strategy, in order to improve system performance.

User-Adaptable TOOT

Although users differ with respect to their knowledge of system limitations, and although different dialogue strategies make system limitations more apparent to users, most spoken dialogue systems do not try to improve performance by adapting dialogue behavior to individual users. We have evaluated the utility of allowing users to dynamically adapt (at any point(s) in a dialogue) the initial set of initiative and confirmation dialogue strategies used by TOOT, a spoken dialogue system for retrieving train schedules on the web. (The initiative strategy specifies who controls the dialogue, while the confirmation strategy specifies how TOOT lets the user know what it just understood.) We conducted an experiment in which 20 novice users carried out 4 tasks with adaptable and non-adaptable versions of TOOT, resulting in a corpus of 80 dialogues. Our results showed that adaptable TOOT outperformed non-adaptable TOOT for task success and a subjective measure of user satisfaction, and that the utility of adaptation for other evaluation measures depended on the initial configuration of dialogue strategies (Litman & Pan 1999).

Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Automatically Adaptive TOOT, via Detection of Poor Speech Recognition

We have explored the utility of automatically adapting the behavior of TOOT, in response to inferred poor speech recognition performance. Our evaluations of TOOT and other spoken dialogue systems suggest that speech recognition accuracy is among the most significant predictors of dialogue system performance, e.g. (Walker *et al.* 1997; Kamm, Litman, & Walker 1998; Litman & Pan 1999). First, we demonstrated that it is possible to recognize when dialogues are going poorly with respect to speech recognition. Our training and testing data consisted of 544 dialogues, obtained from TOOT and two other dialogue systems. Each dialogue was assigned a classification of either "good" or "bad", by thresholding on the percentage of user utterances that were misrecognized. Each dialogue was also represented in terms of 23 features that were automatically computed from system logs. The features were based on 5 types of knowledge sources (speech recognition confidence, dialogue efficiency, dialogue quality, lexical information, and experimental conditions). The machine learning program Ripper (Cohen 1996) was used to induce classification rules from a variety of feature representations of our data. The best learned rulesets had cross-validated error rates of approximately 24%, compared to a 48% baseline majority class rate (Litman, Walker, & Kearns 1999).

Next, we used these results to automate the process of dynamically adapting the dialogue behavior of TOOT. We learned a classifier for predicting when TOOT dialogues were problematic with respect to speech recognition, by instantiating (Litman, Walker, & Kearns 1999) for the TOOT corpus. Second, we implemented an adaptive version of TOOT that used this classifier. In each dialogue, adaptive TOOT 1) monitored the features needed for prediction, 2) predicted whether the dialogue was problematic after a pre-specified number of user utterances, and 3) adapted to a more conservative dialogue strategy when a dialogue was classified as problematic. The dialogue status was then reset to non-problematic and the monitoring process repeated. Finally, we conducted an experiment in which

6 novice users carried out 4 tasks with adaptive TOOT. Our results showed that adaptive TOOT outperformed non-adaptive TOOT on task success, and that the performance of adaptive TOOT approached that of user-adaptable TOOT (Litman & Pan 2000).

Finally, we have successfully extended our problematic situation detection methodology (Litman, Walker, & Kearns 1999) to apply to two new prediction tasks (predicting misrecognized utterances in TOOT using prosody (Hirschberg, Litman, & Swerts 1999), and predicting dialogues in which callers were - or should have been - transferred to a human operator in a call-routing dialogue system (Langkilde *et al.* 1999)). However, we have not yet used these learned prediction rules to trigger or evaluate the utility of adaptation.

Adaptation over Time in RLDS, via Reinforcement Learning

While the above research focused on adapting behavior during a single dialogue, we are also interested in adapting behavior over longer periods of time. Recently, it has been proposed that dialogue systems can be treated as Markov decision processes (MDPs), where the goal is to take actions to maximize reward; this makes it possible to apply reinforcement learning (RL) to find a good dialogue strategy (Walker 1993; Biermann & Long 1996; Levin, Pieraccini, & Eckert 1997; Walker, Fromer, & Narayanan 1998). However, the application of RL to dialogue systems faces a number of technical challenges (e.g., representing the dialogue state by the entire dialogue so far is neither feasible nor conceptually useful, and it is not clear what to use as a reward). Our work has further explored the utility of RL for practical dialogue system design.

First, we 1) extended the MDP formalism to handle notions of approximate state, action, and reward, 2) developed an associated software tool to learn an optimal dialogue strategy from corpora, and 3) performed one of the first applications of RL to previously collected human-machine dialogue corpora (from TOOT and the ELVIS email system). Our results demonstrated that sensible dialogue strategies could be learned even given sparse data, and that RL was a useful tool for browsing and understanding correlations in complex, temporally dependent dialogue corpora. Furthermore, our policies were similar even when different rewards (such as subjective user satisfaction versus objective task completion) were used (Singh *et al.* 1999).

Based on the success of this work, we next applied our tool in an online manner to the problem of automated dialogue strategy synthesis, with the goal of obtaining measurable improvements in system performance. We 1) built a dialogue system called RLDS (Reinforcement Learning Dialogue System) for accessing a database of activities in New Jersey, that was deliberately exploratory with respect to initiative and confirmation dialogue strategies, 2) fielded the exploratory system on a set of 54 training subjects, 3) used this

data (311 dialogues) to build an MDP providing a state-based statistical model of user reactions to system actions, then used RL to simultaneously evaluate many dialogue strategies and compute the optimal dialogue strategy of the MDP, and 4) refiled a version of the system that used the optimal dialogue strategy on a set of 21 test subjects (124 dialogues). Our results demonstrated that the use of RL led to statistically significant improvements on objective measures such as task completion (which increased from 52% to 63%), despite violations of the MDP formalism (Litman *et al.* 2000).

References

- Biermann, A. W., and Long, P. M. 1996. The composition of messages in speech-graphics interactive systems. In *Proceedings of the 1996 International Symposium on Spoken Dialogue*, 97-100.
- Cohen, W. 1996. Learning trees and rules with set-valued features. In *14th Conference of the American Association of Artificial Intelligence, AAAI*.
- Hirschberg, J.; Litman, D.; and Swerts, M. 1999. Prosodic cues to recognition errors. In *Proceedings of the Automatic Speech Recognition and Understanding Workshop*.
- Kamm, C.; Litman, D.; and Walker, M. A. 1998. From novice to expert: The effect of tutorials on user expertise with spoken dialogue systems. In *Proceedings of the International Conference on Spoken Language Processing*.
- Langkilde, I.; Walker, M.; Wright, J.; Gorin, A.; and Litman, D. 1999. Automatic prediction of problematic human-computer dialogues in 'how may i help you?'. In *Proceedings of the Automatic Speech Recognition and Understanding Workshop*.
- Levin, E.; Pieraccini, R.; and Eckert, W. 1997. Learning dialogue strategies within the markov decision process framework. In *Proc. IEEE Workshop on Automatic Speech Recognition and Understanding*.
- Litman, D. J., and Pan, S. 1999. Empirically Evaluating an Adaptable Spoken Dialogue System. In *Proceedings of the 7th International Conference on User Modeling*, 55-64.
- Litman, D. J., and Pan, S. 2000. Predicting and adapting to poor speech recognition in a spoken dialogue system. Manuscript submitted for publication.
- Litman, D. J.; Kearns, M. S.; Singh, S.; and Walker, M. A. 2000. Automatic optimization of dialogue management. Manuscript submitted for publication.
- Litman, D. J.; Walker, M. A.; and Kearns, M. J. 1999. Automatic detection of poor speech recognition at the dialogue level. In *Proceedings of the 37th Annual Meeting of the Association of Computational Linguistics*, 309-316.
- Singh, S.; Kearns, M. S.; Litman, D. J.; and Walker, M. A. 1999. Reinforcement learning for spoken dialogue systems. In *Proceedings of NIPS99*.

Walker, M.; Hindle, D.; Fromer, J.; Fabrizio, G. D.; and Mestel, C. 1997. Evaluating competing agent strategies for a voice email agent. In *Proceedings of the European Conference on Speech Communication and Technology*.

Walker, M. A.; Fromer, J. C.; and Narayanan, S. 1998. Learning optimal dialogue strategies: A case study of a spoken dialogue agent for email. In *Proceedings of the 36th Annual Meeting of the Association of Computational Linguistics*, 1345–1352.

Walker, M. A. 1993. *Informational Redundancy and Resource Bounds in Dialogue*. Ph.D. Dissertation, University of Pennsylvania.