

Can We Talk? Prospects for Automatically Training Spoken Dialogue Systems

Marilyn A. Walker

University of Sheffield
Regent Court, 20 Portobello St, Sheffield, England S1 4DP
m.a.walker@sheffield.ac.uk

Abstract

There is a strong relationship between evaluation and methods for automatically training language processing systems, where generally the same resource and metrics are used both to train system components and to evaluate them. To date, in dialogue systems research, this general methodology is not typically applied to the dialogue manager and spoken language generator. I will argue that any metric that can be used to evaluate system performance should also be usable as a feedback function for automatically training the system. My argument is motivated with examples of the application of reinforcement learning to dialogue manager optimization, and the use of boosting to train the spoken language generator.

Introduction

It is obvious that there is a strong relationship between evaluating a system and automatically training it: the relationship arises from the fact that it is generally possible to use the same data to train a system as to evaluate it. Examples from language technology include: (1) speech recognition, where the training data consists of a corpus of the speech signal and the corresponding human transcriptions, and evaluation compares the recognizer output with the transcription; and (2) part of speech tagging, where the training data consists of a corpus of word strings labeled with part-of-speech tags, and evaluation compares the labeled part-of-speech tags with the output of an automatic tagger. There are many other examples.

However, it makes no sense to evaluate a dialogue system's outputs by comparing them to previously stored outputs collected in another dialogue interaction. This is primarily due to the fact that, unlike the other language processing problems mentioned above, there are many possible correct responses at any point in a dialogue. The most that can be said is that some responses are better than others, for particular users, given particular dialogue contexts, or that some dialogue interactions go better than others, for particular users, given particular tasks. The good news is that this means that it is still possible to use the same data to train a dialogue system as one uses to evaluate it.

An algorithm for automatically training a system requires a metric which the training algorithm attempts to optimize; this metric is called an *objective function*. Dialogue systems are typically evaluated using several types of scalar metrics; (1) Task completion or transaction success measures; (2) Efficiency measures such as time to completion; (3) Subjective measures derived from direct user feedback, such as user satisfaction measures based on a user-survey, or user ratings of system utterances in context. Any of these measures, or a combination of them, could in principle, be used as an objective function for training the dialogue manager and response generation components.

Our research has experimented with (1) automatically training the dialogue manager, by applying reinforcement learning (Sutton and Barto, 1990); and (2) automatically

training the spoken language generator using the Rankboost algorithm, a form of boosting (Schapire, 1999). Our overall goal has been to make it possible to produce an end-to-end trainable dialogue system, in a similar manner to that described in (Young, 2002).

The reinforcement learning experiments employ user satisfaction as the objective function in one prototype system (Walker, 2000), and task completion in another (Litman, Singh, Kearns and Walker 2002). We demonstrate statistically significant performance improvements in dialogue strategy selection. The boosting experiments employ user ratings of the generator output in particular dialogue contexts as the objective function, in one case for information gathering utterances (Walker, Rambow and Rogati, 2002) and in the other case for information presentation (Walker, Stent and Prasad, 2003). In both cases, we demonstrate that an automatically trained response generator is statistically better than several baselines, and statistically equivalent to a hand-crafted template-based generator for the same task. The talk will describe these experiments and results in more detail.

References

- Litman, D., S. Singh, M. Kearns and M. Walker (2000). Automatic Optimization of Dialogue Management. In Proc. of COLING 2000.
- Schapire, R.E. (1999) Theoretical views of Boosting. In Computational Learning Theory: Proc. of the 4th European Conference, EuroCOLT'99.
- Sutton, R.S. and Barto A.G., (1998). Reinforcement Learning: An Introduction. MIT Press.
- Walker, M.A. (2000) An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. Journal of Artificial Intelligence Research.
- Walker, M.A., O. Rambow and M. Rogati (2002). Training a Sentence Planner for Spoken Dialogue Using Boosting. Computer Speech and Language.
- Walker, M.A., A. Stent and R. Prasad (2003). A Trainable Generator for Recommendations in Multimodal Dialogue. In Proc. of EUROSPEECH.
- Young, S.J. (2002). Talking to Machines: Statistically Speaking. In Proc. of ICSLP 2002.