Statistical User Simulation for Spoken Dialogue Systems: What for, Which Data, Which Future? *

Olivier Pietquin
SUPELEC - UMI 2958 (GeorgiaTech - CNRS)
2 rue Edouard Belin
57070 Metz - France
olivier.pietquin@supelec.fr

Abstract

There has been a lot of interest for user simulation in the field of spoken dialogue systems during the last decades. User simulation was first proposed to assess the performance of SDS before a public release. Since the late 90's, user simulation is also used for dialogue management optimisation. In this position paper, we focus on statistical methods for user simulation, their main advantages and drawbacks. We initiate a reflection about the utility of such methods and give some insights of what their future should be.

1 Introduction

User simulation for Spoken Dialogue Systems (SDS) aims at generating artificial interactions supposed to be representative of what would be an actual dialogue between a human user and a given dialogue system. User simulation is thus different from user modeling which is often included into the systems to infer user goals from observable clues (user's utterances, intonations etc.) (Zukerman and Albrecht, 2001). In this paper we focus on statistical methods for user simulation, that is methods purely based on data and statistical models and not cognitive models. Also, we only address user simulations working at the intention level, that is generating dialog acts and not speech or natural language (Schatzmann et al., 2006). User modeling, used to infer user intentions in dialogue systems is not addressed.

The aim of user simulation was initially to assess the performance of a SDS before a public release (Eckert et al., 1997). Given a performance metric and a simulation method, the natural idea of automatically optimizing SDS (using reinforcement learning RL) appeared in the literature in the late 90's (Levin et al., 2000).

2 Is user simulation useful?

Initially, SDS optimisation required a lot of data because of inefficiency of RL algorithms, justifying the use of simulation. In recent years, sample efficient RL methods were applied to SDS optimization. This allows learning optimal dialogue strategies directly from batches of data collected between suboptimal systems and actual users (Li et al., 2009; Pietquin et al., 2011b) but also from online interactions (Pietquin et al., 2011a; Gasic et al., 2011). Do we have to conclude that user simulation is useless?

3 Do we need to train models?

It is commonly admitted that learning parameters of user simulation models is hard because most of variables are hidden (user goal, mental states etc.) and tricky to annotate. This is why current user simulators are trainable but rarely trained (Pietquin, 2006; Schatzmann et al., 2007). Do we really need to train user simulation models? If so, which data and annotation schemes do we need?

4 Does simulation reach the target?

User simulation aims at reproducing plausible interactions but in contexts that were not seen in the data

*This work as been partially funded by the INTERREG IVa project ALLEGRO and the Région Lorraine

NAACL-HLT 2012 Workshop on Future directions and needs in the Spoken Dialog Community: Tools and Data, pages 9–10,
Montréal, Canada, June 7, 2012.

collected to train the model. It is generally hard to assess the quality of such models. Especially, it is hard to find a single metric to assess user simulation performances (Pietquin and Hastie, 2011). Also, it has been shown that user simulation affects a lot the result of SDS strategy optimisation (Schatzmann et al., 2005). What should be assessed? Statistical consistency, ability to generalize, ability to generate sequences of interactions similar to real dialogues, ability to produce optimal strategies by RL? If one wants to learn an optimal simulation model, there is a need for a single optimality criterion.

5 What's the future of user simulation for SDS?

Whatever the use one wants to make of user simulation (learning or assessment for SDS), the future of this research field relies probably on a redefinition of the role of user simulation. So far, user simulation is seen as a generative systems, generating dialog acts according to the context. Current user simulation models are therefore based on a large amount of conditional probabilities which are hard to learn, and the training (if there is one) requires a lot of prior knowledge, the introduction of smoothing parameters etc.

We believe that user simulation should be redefined as a sequential decision making problem in which a user tries to reach a goal in a natural and efficient way, helped by an artificial agent (the SDS). One major difference between this vision and the common probabilistic one is that it takes into account the fact that human users adapt their behavior to the performances and the strategy of the SDS. This can be called “co-adaptation” between human users and artificial systems and justifies that user simulation should still be studied.

Recently, user simulation models based on inverse reinforcement learning have been proposed (Chandramohan et al., 2011). In this framework, a user is modeled as optimizing it's behavior according to some unknown reward which is inferred from recorded data. This might be an answer to the co-adaptation problem. Yet, is user simulation still useful in this framework? Knowing the reward of the user, do we still need simulation or is it possible to compute directly an optimal dialogue strategy?

References

S. Chandramohan, M. Geist, F. Lefèvre, and O. Pietquin. 2011. User Simulation in Dialogue Systems using Inverse Reinforcement Learning. In *Proc. of Interspeech 2011*, Florence (Italy).

W. Eckert, E. Levin, and R. Pieraccini. 1997. User Modeling for Spoken Dialogue System Evaluation. In *Proc. of ASRU'97*, Santa Barbara (USA).

M. Gasic, F. Jurcicek, B. Thomson, K. Yu, and S. Young. 2011. On-line policy optimisation of spoken dialogue systems via live interaction with human subjects”. In *Proc. of ASRU 2011*, Hawaii (USA).

E. Levin, R. Pieraccini, and W. Eckert. 2000. A Stochastic Model of Human-Machine Interaction for learning dialog Strategies. *IEEE Transactions on Speech and Audio Processing*, 8:11–23.

L. Li, S. Balakrishnan, and J. Williams. 2009. Reinforcement Learning for Dialog Management using Least-Squares Policy Iteration and Fast Feature Selection. In *Proc. of InterSpeech'09*, Brighton (UK).

O. Pietquin and H. Hastie. 2011. A survey on metrics for the evaluation of user simulations. *Knowledge Engineering Review*.

O. Pietquin, M. Geist, and S. Chandramohan. 2011a. Sample Efficient On-line Learning of Optimal Dialogue Policies with Kalman Temporal Differences. In *Proc. of IJCAI 2011*, Barcelona, Spain.

O. Pietquin, M. Geist, S. Chandramohan, and H. Frezza-Buet. 2011b. Sample-Efficient Batch Reinforcement Learning for Dialogue Management Optimization. *ACM Transactions on Speech and Language Processing*, 7(3):7:1–7:21, May.

O. Pietquin. 2006. Consistent goal-directed user model for realistic man-machine task-oriented spoken dialogue simulation. In *ICME'06*, Toronto (Canada).

J. Schatzmann, M.Stuttle, K. Weilhammer, and S. Young. 2005. Effects of the user model on simulation-based learning of dialogue strategies. In *Proc. of ASRU'05*.

J. Schatzmann, K. Weilhammer, M. Stuttle, and S. Young. 2006. A survey of statistical user simulation techniques for reinforcement-learning of dialogue management strategies. *Knowledge Engineering Review, vol. 21(2), pp. 97–126*.

J. Schatzmann, B. Thomson, K. Weilhammer, H. Ye, and S. Young. 2007. Agenda-based User Simulation for Bootstrapping a POMDP Dialogue System. In *Proc. of HLT NAACL*.

I. Zukerman and D. Albrecht. 2001. Predictive statistical models for user modeling. *User Modeling and User-Adapted Interaction*, 11(1-2):5–18. invited paper.