Machine Learning for Adaptivity in Spoken Dialogue Systems

OLIVER LEMON, Heriot Watt University
OLIVIER PIETQUIN, Supélec / UMI 2958 (GeorgiaTech - CNRS)

ACM Reference Format:
Lemon, O., Pietquin, O., 2010. Machine Learning for Adaptivity in Spoken Dialogue Systems ACM Trans. Speech Lang. Process. 0, 0, Article 0 (2011), 3 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

> "I know I've made some very poor decisions recently, but I can give you my complete assurance that my work will be back to normal. I've still got the greatest enthusiasm and confidence in the mission. And I want to help you."

(HAL-9000 in *2001 A Space Odyssey*)

In the 1960s, AI researchers were predicting that machines capable of spoken dialogue behaviour, somewhat like the example above, would be possible within a few decades. Fifty years later, we are still confronted with very difficult problems in AI, but we have powerful new tools with which to address the spoken dialogue problem.

The research landscape in spoken dialogue systems has undergone significant changes over the past decade. This transformation has been the result of new momentum and fresh insights coming from the investigation of data-driven, statistical machine learning methods in three core areas of dialogue system research: Spoken Language Understanding, Dialogue Management, and Natural Language Generation. These methods hold the promise of mathematically precise approaches to system design, optimisation, and evaluation, based on data collected from real user interactions with dialogue systems. The papers collected together in this special issue represent important themes in these research areas, and illustrate current research directions in the field. As such, we hope that they will form a valuable resource for the international community of dialogue system researchers, both in industry and academia.

Speech and language processing techniques have now achieved such a level of maturity that voice-enabled user interfaces are widely deployed, and have created a billion dollar industry. However, the design and development of these interfaces is far from being a simple and standardised process. Indeed, it is not enough to simply plug together speech recognition and synthesis systems: recognised speech must be *understood* in the context of the application, the overall interaction must be appropriately *managed*, and spoken language must be suitably *generated* to improve overall system efficiency

The authors are grateful for funding for this work, from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 216594 (CLASSiC project, www.classic-project.org) and from the EPSRC, project no. EP/G069840/1.
Author's addresses: O. Lemon, Heriot Watt University (MACS), EH14 4AS Edinburgh (UK), , and O. Pietquin, Supélec, 2 rue Edouard Belin, 57070 Metz (France)
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2011 ACM 1550-4875/2011/-ART0 $10.00
DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

and naturalness. Although high level scripting languages exist [W3C 2008] for the rapid design of new speech interfaces, their use is still limited when dialogue tasks become complex.

In the area of dialogue management, machine learning, and especially Reinforcement Learning (RL) [Sutton and Barto 1998], for optimisation of interaction policies is now part of the state-of-the-art. Dialogue management can be represented as a sequential decision making problem where the dialogue manager decides which information to request, confirm, or generate for the user. This decision depends on the current dialogue context and some notion of the overall system goals, which can be represented as a reward (also called 'objective') function for each dialogue decision. This model has therefore been investigated in the paradigm of Markov Decision Processes (MDPs) [Levin et al. 1997; Singh et al. 1999; Scheffler and Young 2001; Pietquin and Dutoit 2006; Frampton and Lemon 2009; Henderson et al. 2008; Rieser and Lemon 2011]. All systems face the problem of uncertainty about the current dialogue context (for instance, due to imperfect speech recognition or speech understanding). Partially Observable MDP (POMDPs) offer methods for directly representing uncertainty, and have become an important recent focus of investigation [J. Williams and Young 2005; Young 2006; Williams and Young 2007].

Another important area concerns the small amount of data generally available for learning and testing dialogue strategies. These data sets do not in general contain enough information to explore the whole space of dialogue states (and of strategies), and so dialogue simulation is often required to expand the existing dataset – therefore stochastic dialogue modeling and simulation has become a research field in its own right [Singh et al. 1999; Scheffler and Young 2001; Pietquin and Dutoit 2006; Levin et al. 2000; López-Cózar et al. 2003; Pietquin 2005; Georgila et al. 2005; Schatzmann et al. 2006]. User simulations for different types of user, and for training adaptive natural language generation, are a particular new focus of interest.

Other areas related to spoken dialogue have also benefited from machine learning research – for example context-sensitive speech recognition [Gabsdil and Lemon 2004; Jonson 2006; Lemon and Konstas 2009], trainable natural language generation (e.g. the SPOT system, [Walker et al. 2001]), statistical parsing for dialogue (e.g. the HVS parser [He and Young 2003]) and adaptive TTS (emphasis, prosody etc.).

In this special issue dedicated to "*Machine Learning for Adaptivity in Spoken Dialogue Systems*", we focus on these areas of dialogue processing that have benefited from recent advances in machine learning. This issue follows two successful special sessions of Interspeech organised at Antwerp (Belgium) in 2007 [Lemon and Pietquin 2007] and Brighton (UK) in 2009, showing the growing interest in this research. This field is also supported by several currently funded national and international research projects, such as the European FP7 project "Computational Learning in Adaptive Systems for Spoken Conversation" (CLASSiC: www.classic-project.org) within which the editors collaborate.

These approaches are now being adopted to varying degrees in commercial systems, and have had significant impact in international shared challenges such as the Spoken Dialogue Challenge (SDC: www.dialrc.org/sdc/), where half of the 2010 competitors fielded statistical systems trained with machine learning techniques.

Looking ahead, it seems that statistical learning-based approaches to dialogue systems are here to stay, but that significant challenges still remain. Not least of these is included in the title of this special issue: "adaptivity". This concept captures the fluid and natural manner in which humans interact conversationally with their dialogue partners. This involves rapid adaptation to a changing context, including repair of mis-hearings and misunderstandings, and the design of dialogue contributions which are appropriate to that context, including tailoring of generated language to a dynamic model of the user. Real users show vast variations in their behaviour, vocabulary, interaction style, preferences, and changing conceptualisation of the task domain, and the next generation of dialogue systems will need to adapt and learn about these variations during interaction. The past decade has given us important new tools with which to approach these problems in a principled and data-driven manner.

We hope that newcomers to this area will find it as exciting and stimulating as we have, and that the cross-section of papers collected here will inspire new ideas and directions for future research.

Oliver Lemon & Olivier Pietquin
Associate Editors

REFERENCES

FRAMPTON, M. AND LEMON, O. 2009. Recent research advances in reinforcement learning in spoken dialogue systems. *Knowledge Engineering Review 24,* 4, 375–408.

GABSDIL, M. AND LEMON, O. 2004. Combining acoustic and pragmatic features to predict recognition performance in spoken dialogue systems. In *Proceedings of ACL-04*. 344–351.

GEORGILA, K., HENDERSON, J., AND LEMON, O. 2005. Learning user simulations for information state update dialogue systems. In *Proceedings of Interspeech*.

HE, Y. AND YOUNG, S. 2003. Hidden Vector State Model for Hierarchical Semantic Parsing. In *Proceedings of the International Conference on Acoustics Speech and Signal Processing*.

HENDERSON, J., LEMON, O., AND GEORGILA, K. 2008. Hybrid Reinforcement/Supervised Learning of Dialogue Policies from Fixed Datasets. *Computational Linguistics 34,* 4, 487–512.

J. WILLIAMS, P. P. AND YOUNG, S. 2005. Partially Observable Markov Decision Processes with Continuous Observations for Dialogue Management. In *Proceedings of the SigDial Workshop*.

JONSON, R. 2006. Dialogue Context-Based Re-ranking of ASR Hypotheses. In *Proceedings IEEE 2006 Workshop on Spoken Language Technology*.

LEMON, O. AND KONSTAS, I. 2009. User simulations for context-sensitive speech recognition in spoken dialogue systems. In *European Conference of the Association for Computational Linguistics*.

LEMON, O. AND PIETQUIN, O. 2007. Machine Learning for Spoken Dialogue Systems. In *Proceedings of the 10th European Conference on Speech Communication and Technologies (Interspeech'07)*. Anvers (Belgium), 2685–2688.

LEVIN, E., PIERACCINI, R., AND ECKERT, W. 1997. Learning Dialogue Strategies within the Markov Decision Process Framework. In *Proc. ASRU*.

LEVIN, E., PIERACCINI, R., AND ECKERT, W. 2000. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on Speech and Audio Processing 8,* 1, 11–23.

LÓPEZ-CÓZAR, R., DE LA TORRE, A., SEGURA, J., AND RUBIO, A. 2003. Assessment of dialogue systems by means of a new simulation technique. *Speech Communication 40,* 3, 387–407.

PIETQUIN, O. 2005. A probabilistic description of man-machine spoken communication. In *Proc. ICME'05*.

PIETQUIN, O. AND DUTOIT, T. 2006. A probabilistic framework for dialog simulation and optimal strategy learning. *IEEE Transactions on Audio, Speech and Language Processing 14,* 2, 589–599.

RIESER, V. AND LEMON, O. 2011. Learning and evaluation of dialogue strategies for new applications: Empirical methods for optimization from small data sets. *Computational Linguistics 37,* 1.

SCHATZMANN, J., WEILHAMMER, K., STUTTLE, M. N., AND YOUNG, S. J. 2006. A Survey of Statistical User Simulation Techniques for Reinforcement-Learning of Dialogue Management Strategies. *Knowledge Engineering Review 21*, 97–126.

SCHEFFLER, K. AND YOUNG, S. 2001. Corpus-based dialogue simulation for automatic strategy learning and evaluation. In *Proc. NAACL Workshop on Adaptation in Dialogue Systems*.

SINGH, S., KEARNS, M., LITMAN, D., AND WALKER, M. 1999. Reinforcement learning for spoken dialogue systems. In *Proc. NIPS*.

SUTTON, R. AND BARTO, A. 1998. *Reinforcement Learning : An Introduction*. MIT Press.

W3C 2008. *VoiceXML 3.0 Specifications*. W3C. http://www.w3.org/TR/voicexml30/.

WALKER, M., RAMBOW, O., AND ROGATI, M. 2001. Spot: A trainable sentence planner. In *In Proceedings of the North American Meeting of the Association for Computational Linguistics, 2001.*

WILLIAMS, J. AND YOUNG, S. 2007. Partially Observable Markov Decision Processes for Spoken Dialog Systems. *Computer Speech and Language 21,* 2, 231–422.

YOUNG, S. 2006. Using POMDPs for Dialog Management. In *Proceedings of the 1st IEEE/ACL Workshop on Spoken Language Technologies (SLT'06)*.