Nina Dethlefs

University of Bremen
Department of Linguistics (FB10)
Bibliothekstraße 1
28334 Bremen, Germany

dethlefs@uni-bremen.de
www-user.uni-bremen.de\~dethlefs

1 Research Interests

My research interests involve **context-sensitive**, or **adaptive, Natural Language Generation** (NLG) for **situated dialogue systems**, especially for **spoken** interaction. Context-sensitive situated dialogue systems are typically required to adapt flexibly to dynamic changes of (a) properties of the situation or the spatial setting, such as visible objects, or the complexity of the environment, (b) properties of the user, such as their prior knowledge, goals, beliefs, and general information need, and (c) the dialogue history. In this context, I am mainly interested in applying **Reinforcement Learning** (RL) with **hierarchical control** and prior knowledge in several contexts of rather **large-scale** systems for complex domains. I have also recently looked into the **joint optimisation** of different system behaviours for interdependent decision making between them.

1.1 Incremental Natural Language Generation

Natural Language Generation systems that work in situated domains and need to generate utterances incrementally are faced with a number of challenges: they need to adapt their decisions to continuously changing linguistic and nonlinguistic contexts, as well as to the user's properties, such as their individual information needs, and verbal or nonverbal responses to each generated utterance. Generation decisions involve the tasks of content selection ('what to say'), utterance planning ('how to organise it') and surface realisation ('how to express it'), which can be in many ways related and interdependent. For the former two tasks, for example, there is a tradeoff between how much information to include in an utterance (to increase task success), and how much a user can actually comprehend online in spoken interaction. With regard to surface realisation, decisions are often made according to a language model of the domain. However, there are other linguistic phenomena, such as alignment, consistency, and variation, which can influence people's assessment of discourse and generated output. Also, in dialogue the most likely sequence may not always be appropriate, because it does not correspond to the user's information need, the user is confused or lost, or the most likely sequence is infelicitous with respect to the dialogue history. In such cases, it is important to optimise surface realisation in a unified fashion with content selection and utterance planning. To this end, we presented (Dethlefs and Cuayáhuitl, 2011) an approach towards optimising content selection and surface realisation in a unified fashion using Hierarchical Reinforcement Learning (HRL) with a reward function specifically suited for surface realisation. We were able to show that optimising these processes jointly achieved better results (in terms of task success and similarity with human surface form data) than optimising them in isolation. I am also currently working on including the behaviour of other system components.

1.2 Reinforcement Learning for Situated Dialogue

The idea of optimising a system's behaviour using reinforcement learning is as follows: given a set of system states, a set of actions, and an objective reward function, an optimal system strategy maximises the objective function by choosing the actions leading to the highest reward for every reached state. One problem with flat RL is that it is affected by 'the curse of dimensionality', which means that the state space grows exponentially in the number of state variables, limiting flat RL agents to small-scale applications. Including prior knowledge of a domain into an agent (Singh et al., 2002) therefore has two advantages: it reduces the search space, and allows us to pre-specify certain system strategies in accordance with expert knowledge of the domain, or empirical findings of studies with human participants. Another approach towards making RL scale to larger domains and more complex problems is to use it with a divide-and-conquer approach as in Hierarchical Reinforcement Learning, which has been applied to dialogue strategy learning by Cuayáhuitl (2009). I presented my first work on transferring HRL to language generation in Dethlefs and Cuayáhuitl (2010). Since my own work deals with language generation in a spoken, situated dialogue system with a relatively large state-action space, in Dethlefs et al. (2011) we presented an approach that combines both hierarchical RL with a hierarchical Information State capturing prior knowledge of the domain.

2 Future of Spoken Dialog Research

I see as one of the most interesting and challenging directions for the future the application of adaptive spoken dialogue systems in real-world scenarios. Such systems will need to adapt quickly and flexibly to individual users and unseen situations during the course of the interaction. This will require that systems observe their environment carefully, such as the situational or spatial setting and the user, to detect dynamic changes and be able to adapt their own behaviour and dialogue strategies accordingly. In particular, we will need to enquire into sophisticated methods for inferring fine-grained user models in real-time. Such user models need to be negotiated, and continuously re-negotiated with the user during an ongoing dialogue. In contrast to current dialogue systems, many of which display such adaptive properties already, an important future step will be to increasingly infer or learn the best system behaviour in real time, and therefore enable systems to behave and react optimally, even in unseen scenarios that may differ from the ones they were trained for or developed in.

A further future challenge is that dialogue agents will need to become more natural in their dialogue behaviour and surface realisation choices in order to be more widely accepted as human-like interloctors in everyday interaction scenarios by a wide range of people.

I am convinced that machine learning methods, in particular reinforcement learning, will play an important role in addressing both of the outlined challenges. On the other hand, I believe that empirical research on human cognition and linguistic behaviour will be essential in informing systems and system behaviours in various ways.

Further, several recent approaches towards optimising different behaviours of a system in a unified fashion (Lemon, 2011; Cuayáhuitl and Dethlefs, to appear; Dethlefs and Cuayáhuitl, 2011) have shown promising results for future application.

3 Suggestions for Discussion

- Joint optimisation of the behaviours of different system components to achieve coherence and interdependent decision making.
- Methods for Situated Spoken Dialogue Systems that scale to complex real-world applications.
- Shared tasks and more standardisation of evaluation procedures for better comparability of research approaches, sharing of resources, and simplified replicability of results.

References

Heriberto Cuayáhuitl. 2009. *Hierarchical Reinforcement Learning for Spoken Dialogue Systems*. PhD Thesis, School of Informatics, University of Edinburgh.

Heriberto Cuayáhuitl and Nina Dethlefs. to appear. *Spatially-aware Dialogue Control Using Hierarchical Reinforcement Learning*. ACM Transactions on Speech and Language Processing (Special Issue on Machine Learning for Robust and Adaptive Spoken Dialogue Systems).

Nina Dethlefs and Heriberto Cuayáhuitl. 2010. *Hierarchical Reinforcement Learning for Adaptive Text Generation*. Proceedings of INLG-2010, Dublin, Ireland.

Nina Dethlefs and Heriberto Cuayáhuitl and Jette Viethen. 2011. *Optimising Natural Language Generation Decision Making for Situated Dialogue*. Proceedings of SIGdial-2011, Portland, OR, USA.

Nina Dethlefs and Heriberto Cuayáhuitl. 2011. *Hierarchical Reinforcement Learning and Hidden Markov Models for Task-Oriented Natural Language Generation*. Proceedings of ACL-HLT 2011, Portland, OR, USA.

Peter Heeman. 2007. *Combining Reinforcement Learning with Information-State update rules*. Proceedings of HLT, Rochester, NY, USA, pages 268-275.

Oliver Lemon. 2011. *Learning what to say and how to say it: joint optimization of spoken dialogue management and Natural Language Generation*. Computer Speech and Language, volume 25, issue 2, pages 210-221.

Satinder Singh and Diane J. Litman and Michael S. Kearns and Marilyn A. Walker . 2002. *Optimizing Dialogue Management with Reinforcement Learning: Experiments with the NJFun System*. Journal of Artifical Intelligence Research, volume 16, pages 105-133.

Biographical Sketch

Nina Dethlefs is currently a PhD student and research assistant in the group for Computational Linguistics at the University of Bremen. She studied in Bremen and Edinburgh (School of Informatics). She is a member of the Transregional Research Centre SFB/TR-8 on Spatial Cognition in Bremen. She has also worked on collaborative projects on generating adaptive spoken route instructions with the University of Melbourne, Australia, and on reference generation in situated dialogue with the Centre for Language Technology at Macquarie University in Sydney. She likes travelling, running, good restaurants, and spending nice evenings with friends.