Learning to Predict the Effects of Complex Utterances

Michael Klein (kleinml@ims.uni-stuttgart.de)[1,2] and Kenji Doya[2,3]
[1] IMS, Stuttgart Univ.; [2] ATR Comp. Neurosc. Lab; [3] Crest Jap. Sc. Tech. Ag.

Our paper deals with the question, how the human brain acquires, represents, and processes the meaning of natural language expressions. The main focus lies on the goal-directedness of language, i.e. on the fact, that communication normally serves some purpose. Regarding the brain as a goal-directed system (with the strength of such systems being the selection of actions with respect to goals), the goal-directedness of language can be treated as a special case of goal-directed actions and language and meaning are treated as one means of achieving its goals. Thus, learning of language becomes an optimization process, where the performance of the language learner in handling the environment increases with a higher level of linguistic capacity.

To accomplish complex learning tasks, such as language, the brain uses subsystems, which differ especially with respect to their learning strategies (Doya, 1999), but interact to achieve the global goals of the system. In particular, *supervised learning* can be used to train forward models (Jordan and Rummelhart, 1992) for prediction, and *reinforcement learning* can be used to train a value function (Sutton and Barto, 1998).

In our aproach, the forward model is used to predict context-dependent utterance effects. It is trained by observing the changes in the environment that utterances of other agents bring about. In contrast to the simple forward model described in Klein et al. (2004), in this study, we use a more complex model, which enables the prediction of the effects of compositional utterances. In a first step, a feature representation (color, form, movement) is extracted from the utterance, which is used to identify and locate the relevant object in space. Attention is directed on this particular object and a change of this object in the context is predicted. With this forward model, agents are able to understand and produce sentences which they have never heard before, i.e. they derive the meaning of the complex utterance from the meaning of its parts. A special feature of our approach is the use of the forward model in language understanding to compute the communicative intention of the speaker. A value function, which determines how desirable states of the world are, is used along with the forward model for action or utterance selection. Verbal and non-verbal actions can be selected, depending on whether speaking or manipulating the environment *directly* is more likely to bring about the desired change. We present the results of simulation experiments in multi-agent language games.

References

Kenji Doya. What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks*, 12:961–974, 1999.

Michael Jordan and David E. Rummelhart. Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16:307–354, 1992.

Michael Klein, Hans Kamp, Guenther Palm, and Kenji Doya. Expressing and understanding desires in language games. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9) Boston, Massachusetts September 12-15th 2004*, page (in print), Boston, Massachusetts, 2004. MIT press.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - An Introduction*. MIT Press, 1998.