Context and Attention in Reinforcement Learning

Christian Balkenius
christian.balkenius@lucs.lu.se

Lund University Cognitive Science
Kungshuset, Lundagård
S-222 22 Lund, Sweden

In many situations, it makes sense to divide the input to a reinforcement learning system into two parts[8]: one that codes for the stimulus and one that codes for the context. Current thinking in animal learning theory suggests that the stimulus and the context do not play symmetric roles in learning[7]. Initial learning appears to be insensitive to the context, while relearning makes behavior increasingly context sensitive. By using an asymmetric learning rule of this kind, a reinforcement learning system can be designed that initially generalizes maximally between contexts and later restricts the selection of actions to contexts where they are successful[2]. If the stimulus and context are selected and represented[9] in an appropriate way for the task, this scheme can lead to very fast learning.

One way to build the stimulus and context codes is to use an attentional mechanism[4,5,7]. The current focus of attention acts as the stimulus while a sequence of attentional states make up the context. This implies that the context can be controlled by choosing how attention is allocated. The natural way to do this is to view attentional shifts as any other action[8]. This also allows for the learning of *epistemic actions* that updates the context with relevant aspects of the environmental state. The context can also be changed when prediction errors occur based on fixed orientation reactions[7]. The world is represented only indirectly through the actions that are possible–impossible or desirable–undesirable in each situation, and possibly also their expected outcome[10].

This general architecture has been applied to a wide range of cognitive problem domains including working memory tasks and contextual categorization[2], motor set and task-switching[2], modeling of developmental disorders[3], contextual cueing[4], learning in visual attention[5], perception of dynamical scenes[1] and the acquisition of symbols [6].

1. Balkenius, C., Eriksson, A. P., and Åström, K. (2004). *in preparation.*
2. Balkenius, C., and Winberg, S. (2004). Cognitive modeling with context sensitive reinforcement learning, In Malec, J. et al. (Eds.), *Proceedings of AILS 2004, in press.*
3. Balkenius, C. and Björne, P. (2004). First steps toward a computational theory of autism, *submitted to EpiRob '04.*
4. Balkenius, C. (2003). Cognitive processes in contextual cueing. In Schmalhofer, F., Young, R. M., and Katz, G. (Eds.), *Proceedings of the European Cognitive Science Conference 2003* (pp. 43-47). Mahwah, NJ: Lawrence Erlbaum Associates.
5. Balkenius, C. (2000). Attention, habituation and conditioning: Toward a computational model, *Cognitive Science Quarterly*, 1, 2, 171-214.
6. Balkenius, C., Gärdenfors, P., and Hall, L. (2000). The origin of symbols in the brain. In *Proceedings of The Evolution of Language '00.*
7. Balkenius, C., and Morén, J. (2000). A computational model of context processing. In Meyer, J-A., Berthoz, A., Floreano, D., Roitblat, H. L., Wilson, S. W. (Eds.), *From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behaviour*. Cambridge, MA: MIT Press.
8. Balkenius, C. and Hulth, N. (1999). Attention as selection-for-action: a scheme for active perception. In Schweitzer, G., Burgard, W., Nehmzow, U., and Vestli, S. J. (Eds.), *Proceedings of EUROBOT '99* (pp. 113-119). IEEE Press.
9. Balkenius, C. (1996). Generalization in instrumental learning. In Maes, P., Mataric, M., Meyer, J.-A. , Pollack, J., and Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press.*
10. Balkenius, C. (1995). *Natural intelligence in artificial creatures.* Lund University Cognitive Studies, 37.