Neural networks – simulation and application in strategy games

# Abstract

Tobias Schlatter Daniel Meister

May 2007

The aim of our project was to understand more about artificial neural networks and to explore the limits of this technology. Artificial neural networks are self-learning algorithms that are inspired by the way biological nervous systems, such as the brain, process information. They are suitable for solving complex problems in an efficient way; they do not necessarily find the perfect solution. This has to be taken into consideration when problems are chosen. The advantages of artificial neural networks are their adaptability, learning aptitude, robustness and the possibility of generalisation. We tested the limits of artificial neural networks by developing a program that learns to win a board game (*The Settlers of Catan*) involving sophisticated strategies and complex rules.

There are two common approaches to design the learning process for strategy games: one is so-called *supervised learning*, in which the network is trained by an expert using sample game cases. The other is *reinforcement learning*, in which the learning system receives only one reward or penalty; it is left to discover the underlying strategy of the game by itself. We chose the latter approach for our project because with reinforcement learning the potential performance of the neural network is not limited by the quality of the sample cases or by the playing ability of the expert. In our review of the literature, we discovered a program called *TD-Gammon* which is based on reinforcement learning and is now capable of holding its own against the best backgammon players in the world.

After reviewing the literature and becoming familiar with theories of artificial neural networks, we implemented small sample networks for simple problems, such as the Tic-tac-toe game, to gain experience. Because it was important to us that our program would rapidly develop reasonable strategies as well as allow the user to interact with it easily, our objectives included a graphical user interface for *The Settlers of Catan*. Once we completed the implementation of the logic of the game, the neural network, the "neuronal" player and the graphical user interface, we started to train the program. To obtain better results, we wrote a second program that enabled us to train on several computers in parallel.

The results of our work are encouraging. We managed to train several networks to the point that they were able to play well against each other. Playing against human opponents turned out to be very difficult because the behaviour of the "neuronal" player is incomprehensible to humans. Improving performance would require further research into several open issues, which are described in the last chapter of our paper. The final chapter also includes suggestions for future work with the goal of reaching or even exceeding human playing ability. We hope that this is just the beginning of further research and that we will have the opportunity to develop our ideas further when we are at university.