Machine Learning applied to the card game of Differenzler Jass

Degree programme: BSc in Computer Science | Specialisation: Data Engineering
Thesis advisor: Prof. Dr. Erik Graf
Expert: Daniel Voisard
Industrial partner: Swisslos Interkantonale Landeslotterie, Basel

In the past years, advances in Machine Learning have enabled new breakthroughs for computers competing against humans in popular card and board games. This project aims to achieve human-level play by a computer in the «Differenzler» variant of Switzerland's most popular card game «Jass», by applying new methods of Machine and Reinforcement Learning.

While very popular in Switzerland, the game of Differenzler Jass has a geographically very limited player base and thus a narrow range of computer programs playing the game. In most cases the computer's level of play is subpar. Therefore, the goal is to train a competitive computer using new methodologies.

Differenzler Jass
The game's rounds consist of two phases. In the prediction phase, the player has to declare a target score for the round based on the dealt hand and trump suit. In the playing phase, nine tricks are played (following standard Jass rules) in which the player with the highest ranking card collects the trick's points. Each player's final score is computed as the absolute difference between their prediction and scored points during the playing phase. The winner is the player with the lowest final score.

Data Analysis
A dataset of 200'000 human-played rounds, provided by Swisslos intercantonal Lottery from their online platform, were analyzed to gain a deeper understanding of human play and strategies, such as commonly used prediction patterns. The analysis also establishes a benchmark for human-level play.

Prediction
In a first phase, several models were trained and tested on the basis of the real-world data set to improve prediction. Random decision forests and multilayer perceptrons (artificial neural networks) based prediction models have shown to be relatively accurate, improving on the human predictions in over half of the test-cases, despite having no impact on the human playing strategy. Yet, no significant improvement could be measured compared to generally known prediction strategies.

Reinforcement Learning (RL)
A Jass engine was specifically built for the purpose of this project, enabling the simulation of games by different agents and combinations thereof. A Multi-Agent Reinforcement Learning approach was considered but disregarded, as no coordination between the players is allowed and a player's result is independent from other players results (Non-Zero-Sum Game). Since other player's predictions and strategies are unknown, their behaviour can be modeled as random from our perspective.

The game's scoring scheme raises a significant challenge in a RL approach, as the target score can be different in every round. Early tests did not show any significant learning progress when an agent had to play towards a set prediction. Thus, in order to simplify the training goal, the process was flipped and two agents were trained to either minimize or maximize their playing score, completely disregarding any prediction. Especially a (Maskable) Proximal Policy Optimization approach rendered promising results. A matching prediction model was then trained based on a large set of computer simulated games played by each agent to be combined into a complete agent.

Outlook
The resulting agent has shown promising strides in empirical tests. Some large scale test of the agent in a human playing-field has yet to be performed to draw final conclusions. Nevertheless, several aspects - such as training non min/max agents, combining RL with search, optimizing hyperparameters further and putting more resources (time & computing power) into training - could still lead to improvements and some day potentially to Switzerland's AlphaGo moment.

Alexandre Reinhard Moeri
076 320 06 18
alex.moeri@gmail.com