

Mastering the Game of Skat using Decision Transformers.

Attention: You will need knowledge in the game of Skat to work on this topic.

Background

In the past years, AI-based engines achieved super human performance in many classical games such as chess, shogi, and Go [5]. To this day, these techniques struggle in imperfect information games, e.g. poker, where the player does not know the entire state of the game at hand and superhuman performance has only been reached years after the success of AI in perfect information games [1]. One recently published method that could prove as a promising approach for such imperfect information games is the decision transformer (DT) [2]. In natural language processing, transformer models such as BERT and GPT-3 have led to a rapid improvement of state of the art performance in translation and reasoning on text by leveraging the concept of attention while analysing long sequences. DTs are based on the same attention mechanisms and are designed to learn the relation between states in long sequences of moves. One of the games, where this is of special importance is the game of Skat, since in every move, the player must consider all previously played moves. We therefore propose to investigate the application of DTs in mastering the game of Skat.

Skat is a traditional German card game that is played by all generations and has been included in the Federal Register of Intangible Cultural Heritage in 2016. The game is played by three players, where one player plays against the other two. Skat has a very large variety of game modes that can be played based on the cards on the hand. The players decide the teams in a special bidding process, which is a crucial part of the game itself. The manifoldness of possible games makes the game of Skat complicated for traditional AI methods. To this day, few works have been published that achieved close to human performance [3, 4] in some game modes but no method surpassed the human so far.

Description

The scope of this thesis is to investigate the potential of DTs in Skat. For this, you will setup a fitting state representation for a game of Skat that can be fed into a DT. You will receive a large dataset of played Skat games on a master level with the kind support of skat-spielen.de. You will train a DT on this dataset and test the performance against human play. To evaluate your policy, you need a Skat environment, where the computer can play against itself and against humans. This can also be taken from previous work, such as [3]. After offline training, the algorithm can be extended to learn online by playing against itself.




figure 1: A DT playing a round of Skat.



Technische Universität München



Fakultät für Informatik

Lehrstuhl für Echtzeitsysteme und Robotik

Supervisor:

Prof. Dr.-Ing. Matthias Althoff

Advisor:

Jakob Thumm and Jonathan Külz

Research project:

Type:

Semester thesis or internship

Research area:

Offline reinforcement learning

Programming language:

Python

Required skills:

Python, Experience in machine learning, Knowledge of Skat.

Language:

English, German

Date of submission:

2. Mai 2022

For more information please contact us:

Phone:

E-Mail: jakob.thumm@tum.de

Internet:

<https://www.in.tum.de/i06>

Tasks

- Perform a literature review on imperfect information games, offline RL, decision transformers, and Skat AI's
- Setup the data set and develop a good state representation for skat games with DTs.
- Train DTs on a large data set of human games.
- Evaluate your policies.
- Train an online DT to play against itself.

Key learnings

- State-of-the-art offline RL algorithms.
- Challenges in learning complex imperfect information games.

We offer the opportunity to publish relevant results and methodology in a conference or journal.

References

- [1] Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- [2] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. In *Advances in Neural Information Processing Systems*, volume 34, pages 15084–15097.
- [3] Stefan Edelkamp. Challenging human supremacy in skat. In *Proc. of the 12th Int. Symp. on Combinatorial Search (SoCS)*, pages 52–60.
- [4] Douglas Rebstock, Christopher Solinas, and Michael Buro. Learning policies from human data for skat. In *IEEE Conference on Games (CoG)*, pages 1–8.
- [5] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. 362(6419):1140–1144.



Technische Universität München



Fakultät für Informatik

Lehrstuhl für Echtzeitsysteme und Robotik