FACULTY OF ENGINEERING AND COMPUTER SCIENCE
Department of Computer Science

HAW HAMBURG

Deep Reinforcement Learning for Multi-Agent Systems on the Example of Chess

Deep Reinforcement Learning (DRL) is a significant advancement in for example the manufacturing process. Questions of how to train a robot to perform certain actions without being aware of which actions are possible to perform, or without any understanding of the process itself, are relevant and generalizable to various problems. How can I optimally train a new robot agent to learn without knowing the perfect strategy myself? How can a set of agents learn together to collectively solve problems? In order to investigate these questions a representative setup is built in chess, where multiple agents of varying capabilities cooperatively elaborate a solution.

Methods

Phase 1 – Approximating an action space through a frequency distribution

- Take a dataset of thousands of played games from LichessAPI and clean the dataset
- Convert the algebraic notation of each game into action samples of each individual piece, calculate a frequency distribution and training samples
- Use frequency distribution as a starting place of the action space of each piece

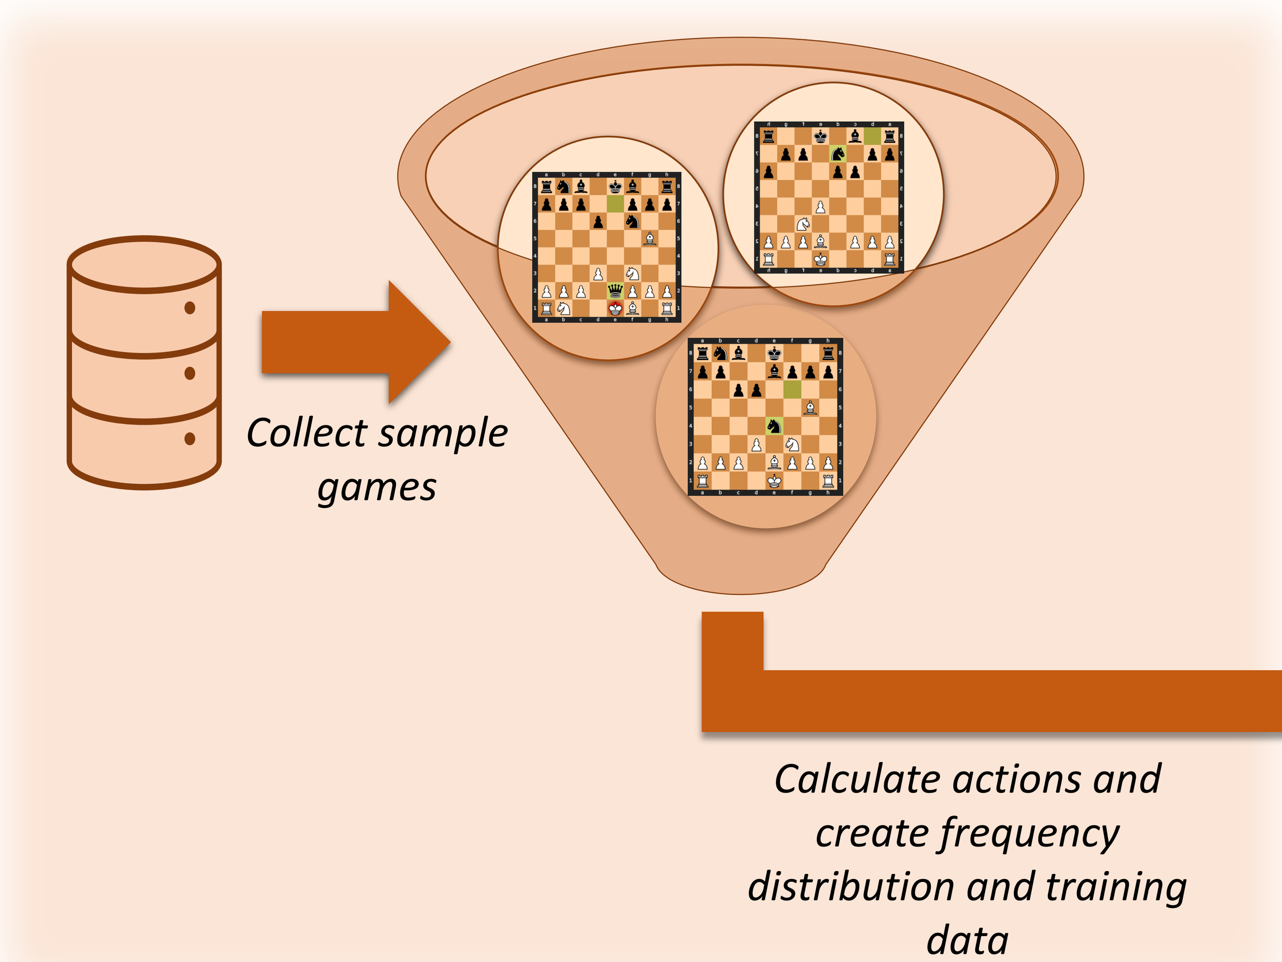

Process of the first phase
Own source: H. Haase

Results

- After performing a small subset of 65 games the frequency distribution can be derived individually per color, piece type, and even for each individual piece
- These distribution yielded a complete action space of each piece, while already incorporating strategic preferences of each piece

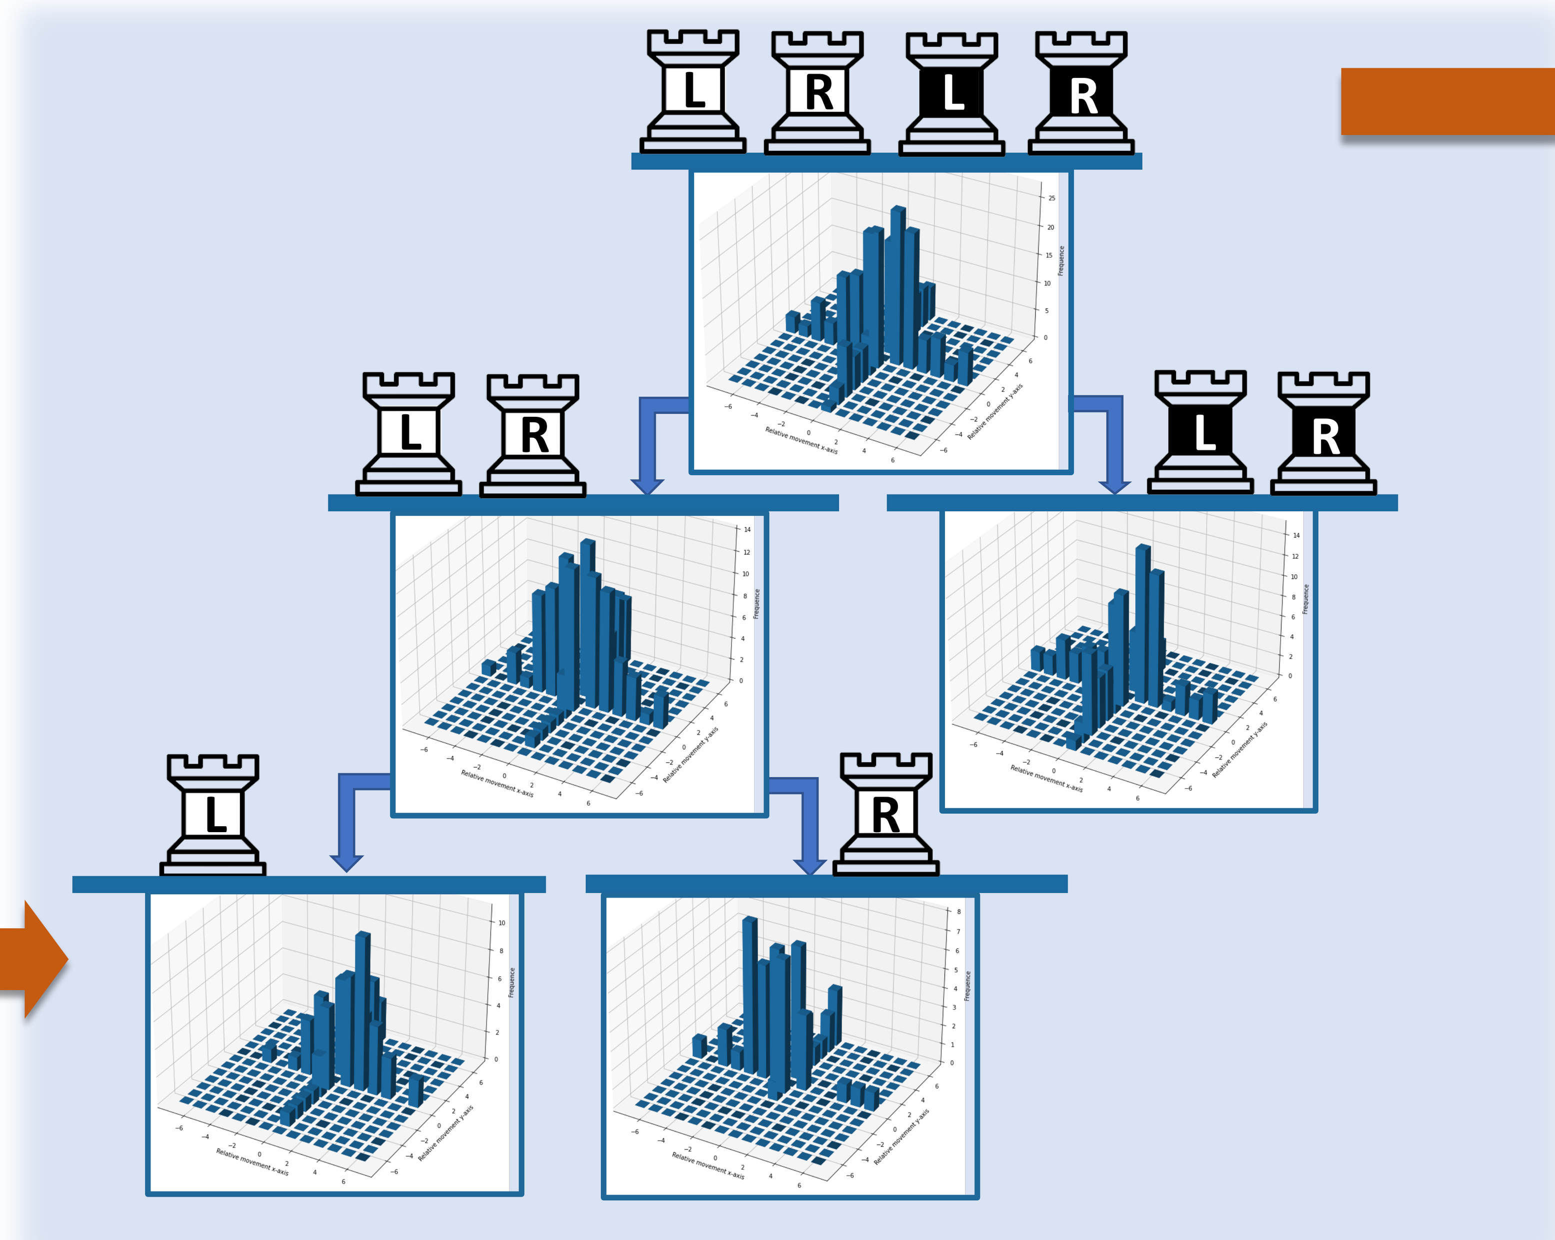

Examplary frequency distributions of rooks
Own source: H. Haase

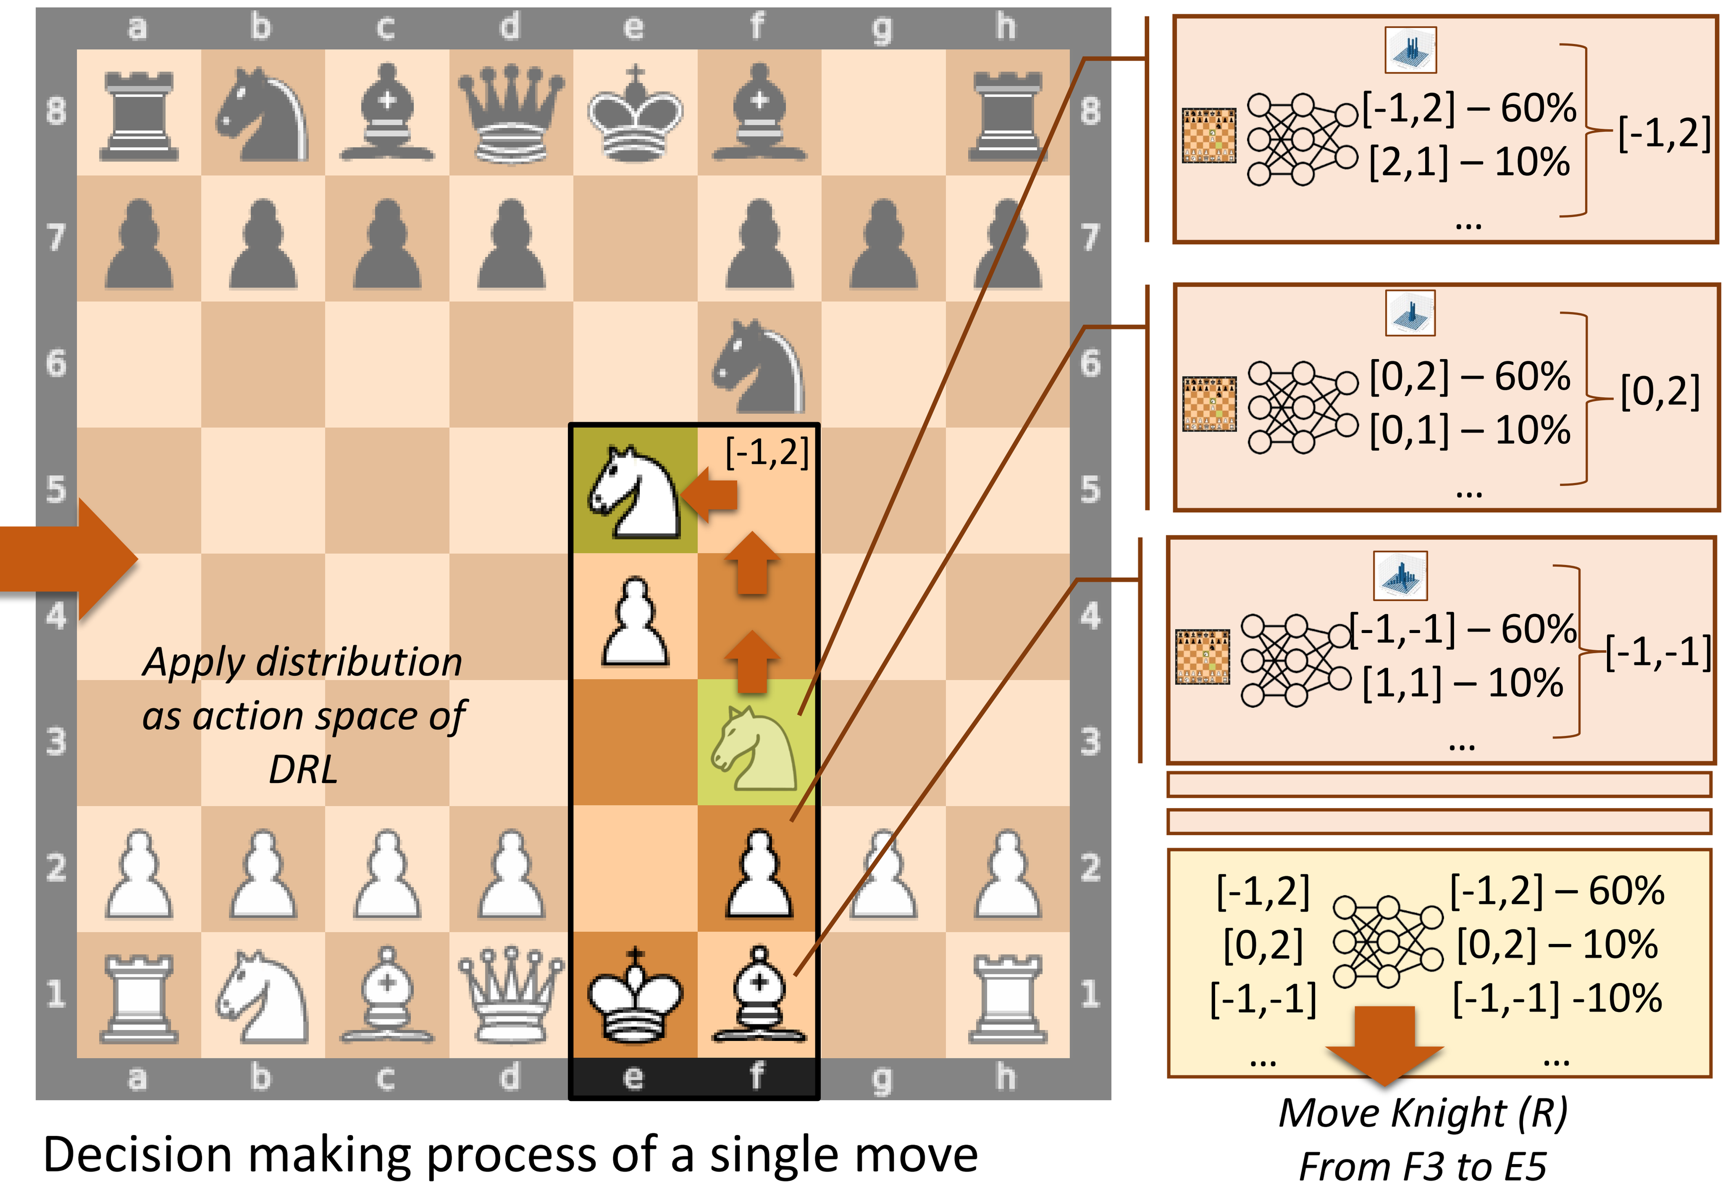

Decision making process of a single move
Own source: H. Haase

Outlook

Phase 2 – Experimenting with the Setup

- Local vs. Global learning and decision making
- Incorporating as little knowledge as possible into the training process
- Knowledge transfer between local pieces
- Neuroevolutive algorithms with dynamic observation and action spaces
- Teaching and Hierarchical Learning on local and global level
- Sparse Rewards and Reward Engineering with no knowledge inclusion
- Competitive and Cooperative acting in MAS

Further Readings

- **[1]** R.S. Sutton and A. G. Barto (2018), Reinforcement learning: An introduction, MIT press,
- **[2]** Zhang, S., Cao, J., Yuan, L., Yu, Y., & Zhan, D. C. (2023). Self-Motivated Multi-Agent Exploration. arXiv preprint arXiv:2301.02083
- **[3]** Mei, Y., Zhou, H., Lan, T., Venkataramani, G., & Wei, P. (2023). MAC-PO: Multi-agent experience replay via collective priority optimization. arXiv preprint arXiv:2302.10418.

MARS GROUP
Multi-Agent Research & Simulation

RESEARCH FOCUS

Reinforcement Learning
Multi-Agent Systems

CONTACT

Helen Haase (B.Sc)
Helen.Haase@haw-hamburg.de

Prof. Dr. Thomas Clemen
Thomas.Clemen@haw-hamburg.de