Exposé:

A Deep Reinforcement Learning Environment for the Limit Orderbook Intraday Electricity Market: Development and efficient Implementation

Philipp Petermeier
Registration number: 5250994
Supervisors:
Prof. Dr. Nicolas Stein, THM Section MND
Dr. Christoph Scholz, Fraunhofer IEE,
Section Energy Informatics

October 2, 2024

Abstract

The Thesis discusses the challenge of applying Deep Reinforcement Learning to intraday markets due to a poor fit of model paradigm and training data. It presents the resulting requirements of a possible solution and proposes a theoretical solution along with the development process. The extent to which the requirements are fulfilled is shown through experiments conducted with the developed Python package on a high-performance cluster. The main requirement, runtime performance of the implementation, is further analyzed with the collected data. The discussion of usage aims to determine the extent of successful implementation. Positive results have been found, with the ability to simulate an average of 35,000 discrete events in trading history within an average runtime of 10.7 seconds, based on historical data. Synthetic data is generated in the process and used to train a Deep Neural Network on. The distribution of runtime usage within training processes and scaling with parallelisation is presented as well.

1 Exposé

1.1 Overview

From a technical standpoint, artificial intelligence (AI) necessitates the presence of hardware, software implementing specific algorithms, and data to train those algorithms on. The successful implementation of a solution hinges on a congruence between the aforementioned elements and the actual problem delineated in the use case. In the majority of cases, the availability and structure of data pertinent to the problem are already known, and a fit between the use case, data, software and hardware is achieved by selecting the most suitable machine learning (ML) algorithm and implementation, while fixing the use case and data as given.

1.2 Topic

This thesis presents a use case of automated trading in the intraday electricity market in Germany and a class of ML algorithms, namely Reinforcement Learning (RL), which when combined result in a poor compatibility of available data with the used algorithm.
The document specifies the requirements for a Python Package that is necessary for successful scientific practice concerning the above mentioned combination. It aligns the ML-Algorithm and Use Case by transforming the provided data to act as part of an interface between the Package and the Use Case. The necessary steps are presented on a theoretical level. An architecture of the Package is developed from the requirements, and the resulting implementation is discussed. The configuration of the auxiliary MongoDB and the programming of the necessary data pipelines in Python are also presented. The primary objective of this research is to ascertain the runtime efficiency of the final product in different configurations on the HPC-Cluster specified in the requirements. The results of these experiments is analysed.
In conclusion, the research provides a concrete example the quantity and nature of work, architectural design, resources and skills required to establish a foundation for the successful development, deployment, and iterative improvement of data science, machine learning, and artificial intelligence.

1.3 Related work & contribution

A limit order book (LOB) is a standard used to facilitate trading in financial markets across the globe. The resulting data architecture is highly idiosyncratic and has been the subject of study with the goal developing automated trading application in various ways as documented in the "Related work"-subsection of the thesis' introduction. The specific contributions consist of applying a generalised solution, RL, to a general problem, the processes documented by the historical data of the LOB, in the very specific context of energy trading on the intraday market. Customizing this general problem set to the specific use case is the first contribution. Additionally the presented solution implements the generation of plausible synthetically generated data used for training based on historical data. Finally the modular software architecture enables future customisation of specifically trained market agents, with each agent paired with their own version of a concrete renewable electricity production facility or even a virtual power plant (VPP).

2 Methodology

Out of the presented specific elements of the task, the electricity market design in germany and RL, software package requirements (technical and non-technical) for future research tasks at Fraunhofer IEE are developed. A formal model that meets the technical specifications has been developed and its implementation process is documented to varying degrees. A particular focus has been placed on custom written ETL-Pipelines and the runtime performance of the final product which were identified as the most critical non-technical requirements at the outset of the project. In order to evaluate the degree to which this specific requirement has been fulfilled, experiments have been designed and executed on the specific runtime environment of the use case, namely one of the on-premise HPC-Clusters.

3 Results

The data gathered in the experiments is analysed and execution time attributed to specific elements of the given configurations of different experiments in order to determine the load distribution of each and identify potential further optimisation options. These options include additional code optimisation as well as improved utilisation of available resources on the HPC-Cluster, either by better administration of the used orchestration package (Ray) or data concerning the upward scaling of when using it. Finally, the success with which the initial requirements are fulfilled are is presented, discussed and an overview of potential avenues for further work is provided.