A DATA-DRIVEN APPROACH TO NFT TRADING: Q-LEARNING BASED SIMULATOR

A Thesis

by

Süleyman Kamalak

Submitted to the
Graduate School of Sciences and Engineering
in Partial Fulfillment of the Requirements for
the Degree of

Master of Science

in
Department of Data Science

Özyeğin University
August 2023

A DATA-DRIVEN APPROACH TO NFT TRADING: Q-LEARNING BASED SIMULATOR

Approved by:

Asst. Prof. Erinç Albey, Advisor
Dept. of Industrial Eng.
Özyeğin University

Asst. Prof. Mehmet Önal
Dept. of Industrial Eng.
Özyeğin University

Prof. Mehmet Güray Güler
Dept. of Industrial Eng.
Yıldız Technical University

Date Approved: August 7, 2023

To my precious family and my dear love...

ABSTRACT

Non-fungible tokens (NFTs) have garnered considerable attention in recent years due to their broad range of applications and potential as a lucrative investment opportunity. Given the nascent nature of the technology and the scarcity of comprehensive studies, there is a pressing need for a holistic trading framework that not only addresses the complexities inherent in the trading process but also proposes viable solutions to existing challenges. This study introduces a novel approach for navigating the NFT trading landscape, effectively confronting the various challenges and suggesting practical solutions. The trading environment is modeled using a Markov Decision Process (MDP), with Q-learning employed to simulate the environment and resolve the MDP problem. The study proposes machine learning models to tackle key challenges, including defining the market state, appraising NFT tokens, and addressing the illiquidity issue prevalent in the NFT market. The proposed approach yields an NFT trading strategy that has shown to outperform traditional strategies, generating substantial profits even amidst bearish market conditions. The Bored Ape Yacht Club (BAYC) collection serves as the primary data set, with the agent trained from June 1, 2021, through January 1, 2023. In testing period from January 1 to June 10, 2023, the proposed model outperformed traditional benchmarks, achieving a profit of 21.14% as opposed to a 20.39% loss for the best-performing benchmark. We assert that this forms a robust foundation for future research into NFT trading simulation and backtesting. We also identify potential areas for future enhancements, particularly possible improvements in the trading strategy and Q-learning approach. The insights gleaned significantly enhance the understanding of the importance of AI applications in the rapidly evolving field of NFT trading.

ÖZETÇE

Nitelikli-Fikri Tapular (NFT'ler), geniş uygulama alanları ve kazançlı bir yatırım fırsatı olma potansiyelleri sebebiyle son yıllarda büyük ilgi toplamıştır. NFT teknolojisinin yeni bir teknoloji olması ve kapsamlı çalışmaların yetersizliği göz önüne alındığında, NFT alım-satım sürecindeki zorlukları tanımlayıp bunlara çözüm öneren bütünsel bir alım-satım yönetim şemasına ciddi bir ihtiyaç vardır. Bu çalışma NFT alım-satım sürecindeki zorlukları tanımlayan ve bunlara pratik çözümler öneren bir yaklaşım sunmaktadır. Alım-satım süreci Markov Karar Süreci (MDP) kullanılarak modellenmiş ve Q-öğrenme ile çevresel koşullar simüle edilerek MDP problemi çözülmüştür. Çalışma, NFT marketinin durumunu tanımlamak, NFT token'larını değerlemek ve NFT piyasasında yaygın olan likidite sorununu ele almak gibi temel zorluklara çözüm olarak makine öğrenmesi modelleri geliştirmiştir. Önerilen yaklaşım, geleneksel alım-satım stratejilerinden daha iyi performans gösteren bir NFT alım-satım stratejisi sunmuştur. Önerilen model, kötü piyasa koşullarında bile önemli kazançlar elde etmiştir. 2023 yılının 1 Ocak tarihinden 10 Haziran'a kadar olan test döneminde en iyi performans gösteren geleneksel alım-satım stratejisinin %20.39'luk kaybına karşı %21.14 kar elde etmiştir. Bu çalışma, NFT alım-satım simülasyonu ve stratejileri geriye dönük test etme konularında gelecekte yapılacak araştırmalar için sağlam bir temel oluşturmaktadır. Çalışmada, önerilen yaklaşımda kullanılan alım-satım stratejisi ve Q-öğrenme yaklaşımında yapılabilecek olası iyileştirmelerle ilgili öneriler sunulmuştur. Elde edilen bulgular, NFT alım-satım alanındaki yapay zeka uygulamalarının önemini kavramada ciddi bir rol oynamaktadır.

ACKNOWLEDGEMENTS

First and foremost, I wish to express my deepest gratitude to my supervisor, Erinç Albey, for his invaluable advice, insightful guidance, and continuous encouragement throughout my studies. His unending patience and wisdom have invariably guided me towards the best course of action in both my education and professional life.

I extend my thanks to all the faculty members in the Data Science Department who equipped me with the knowledge and skills necessary to complete this study. I would also like to acknowledge Gökalp Erbeyoğlu for his constant support and consistently warm and welcoming help during my research.

I am sincerely grateful to my family for their continuous support and motivation throughout my academic journey. I owe a debt of gratitude to my mother and my father for their relentless efforts and guidance during this process. I am equally thankful to my siblings Rumeysa, Elif, and Ekrem for their unwavering trust and love.

I am also profoundly grateful to my dear love Fatma, who has steadfastly supported me throughout this challenging and intense journey, always being a beacon of light in my life. It wouldn't have been the same without you.

My close friend Hanifi, who has been a brother to me, deserves my sincerest thanks. His unwavering support and solidarity have been invaluable. Similarly, my dear friend Yunus, whose uplifting conversations have been a source of comfort during challenging times, deserves a special mention.

Lastly, I would like to extend my appreciation to those whose names I have not mentioned but have nonetheless offered their support.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

1.1 Motivation

Blockchain technology represents a revolutionary breakthrough, offering a path into a decentralized world. This shift has become increasingly necessary given the dramatic growth of online interactions over the past few decades. People are now using the internet to shop, socialize, manage everyday tasks, pay bills, listen to music, get news, and much more. This online activity is progressively disconnecting people from the physical world and immersing us deeper into a virtual environment. Such a shift requires decentralized solutions that does not require real-world interaction. Blockchain allows for money and online transactions to be managed in a decentralized way, providing security and proof of ownership without needing a central authority or third parties. This gives us a new and secure way of conducting transactions in this increasingly virtual world.

Bitcoin was designed to become the 'gold' of the internet. Its creators stated, "A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution" [5]. Following the creation of this electronic payment system, further development led to the development of Ethereum, a platform that broadens the usage areas of blockchain applications. They stated, the technology can be used in many other applications, "using on-blockchain digital assets to represent custom currencies and financial instruments (colored coins), the ownership of an underlying physical device (smart property), non-fungible assets such as domain names (Namecoin), as well as more complex applications involving having digital assets being directly controlled

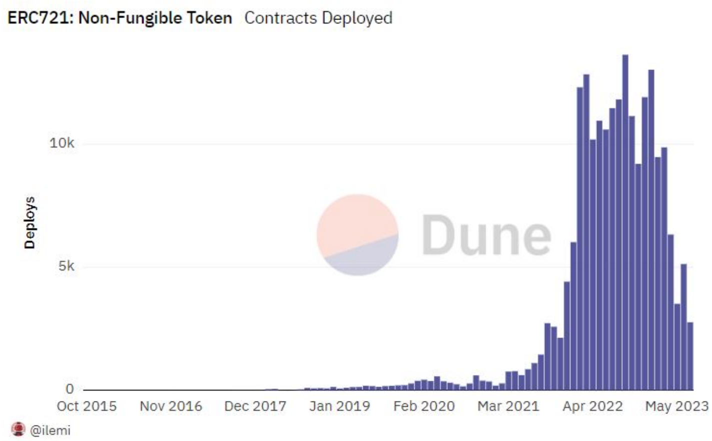

Figure 1: Number of ERC721 contract deployments [1]

by a piece of code implementing arbitrary rules (smart contracts) or even blockchain-based ”decentralized autonomous organizations” (DAOs).” [6].

The advent of the Ethereum blockchain was indeed a second revolution in blockchain technology. This development fostered the creation of more blockchains and thousands of decentralized applications (DApps). Our study focuses on non-fungible tokens (NFTs), which are created using various smart contract standards, with the most common and first being ERC721. The number of deployed ERC721 contracts is demonstrated in Figure 1. There was a significant surge in contract deployments towards the end of 2021, with monthly deployments jumping from thousands to tens of thousands. To date, more than two hundred thousand contracts had been already created only in ETH mainnet [1], highlighting the immense interest in NFTs for a wide range of applications.

Beyond the high number of contract deployments, the transaction volume of NFTs has also dramatically increased. As seen in Figure 2, the market volume grew exponentially from 2018 to 2023. In early 2018, the monthly total market volume in the ETH chain was around USD 10,000. It reached an all-time high of USD 16.6 billion in January 2022 [2]. Meanwhile, other blockchains have also emerged and produced significant transaction volumes.

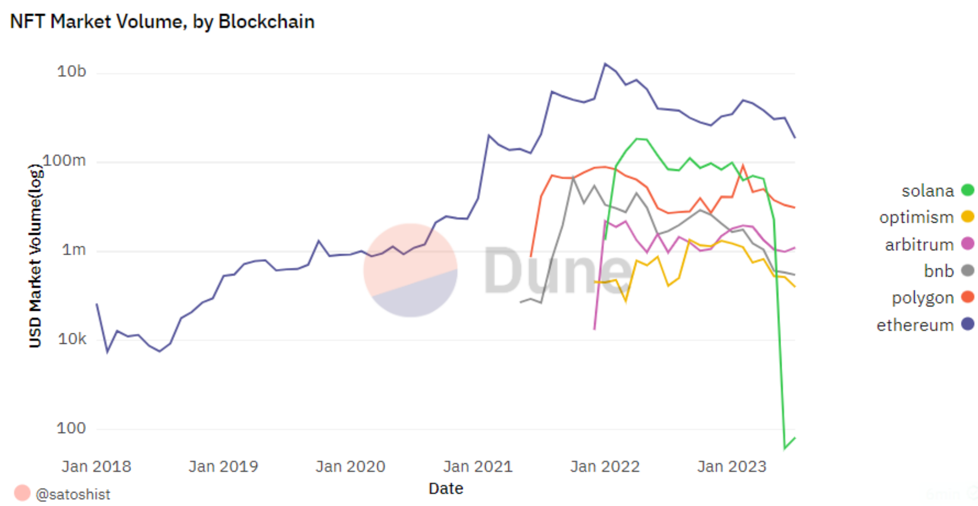

Figure 2: NFT market volume in USD [2]

The technology behind NFTs is promising, with its diverse application areas, significant market volume, and the keen interest from a large number of people. The dynamics of the NFT market are distinct in many ways, such as the uniqueness of items, its illiquidity, and its ever-evolving nature. Given that NFTs are a relatively new concept, there's much to be discovered. So far, the majority of research in the field of NFTs has primarily focused on descriptive analytics, anomaly detection, and efforts to understand the network of NFTs. In addition to their technological benefits, NFTs also offers attractive investment opportunities. Many people strive to understand the NFT market and engage in trading activities, but they encounter many challenges along the way. We designed a comprehensive scheme for NFT trading that addresses some of the main challenges, such as market state analysis, NFT appraisal, NFT liquidation, and the development of a simulation bot and trading strategy.

Our work mainly focuses on a single collection, the Bored Ape Yacht Club (BAYC). The rationale for this particular choice is multifaceted. Above all, BAYC is among the most recognized and popular collections in the NFT market. Characterized by intense market activity and the great interest of influential wallets known as "whales", this collection consistently maintains one of the highest floor prices and market volumes. In addition, the BAYC collection's life span of more than two years provides a comprehensive set of historical data. This comprehensive data repository serves as a crucial resource for our research, allowing for an in-depth and detailed analysis of the dynamics of NFT market.

1.2 Organization of the Thesis

The remainder of this thesis is organized as follows:

- Chapter 2 offers background information pertaining to blockchain, NFTs and network analysis.
- Chapter 3 presents a literature review of relevant topics.
- Chapter 4 provides a detailed explanation of the problem this research seeks to address.
- Chapter 5 describes the data sources and architecture.
- Chapter 6 introduces the methodology and the proposed system for addressing the problem.
- Chapter 7 contains the final results, a discussion of those results, and suggestions for potential future work.
- Finally, Chapter 8 concludes the thesis.

CHAPTER II

BACKGROUND INFORMATION

2.1 Blockchain, NFTs and NFT Marketplaces

Blockchain is a public ledger that stores transactions between peers. This technology uses cryptographic hashing to secure data, and instead of relying on a centralized authority, it is sustained by a distributed network of computers. Prior to the introduction of Bitcoin in 2008 by Satoshi Nakamoto, numerous attempts were made to establish a digital currency. However, all these efforts failed before reaching to the mass, often due to the challenge of achieving decentralized consensus [7]. Bitcoin was the first successful implementation of a peer-to-peer electronic cash system that rewarded miners who solve computationally complex problems to ensure system security, with the primary network coin, BTC.

Following the success of Bitcoin, various other blockchain projects emerged. The technology gained a wide variety of applications, particularly after the introduction of Ethereum in 2014 by Vitalik Buterin. Ethereum's main contribution to blockchain technology was the introduction of smart contracts, which facilitated the automatic execution of digital agreements without the need for an intermediary. Smart contracts enabled the development of decentralized applications (DApps) that operate on the blockchain network. In 2018, the ERC721 standard for non-fungible tokens (NFTs) was introduced by William Entriken, Dieter Shirley, Jacob Evans, and Nastassia Sachs [8]. The key difference between fungible token standard (ERC20) and ERC721 is the fungibility of tokens, with ERC20 tokens being fungible and ERC721 tokens being unique and immutable.

NFTs are digital representations of real-world assets that leverage blockchain technology to ensure asset ownership [9]. NFTs can take the form of an image, video, audio, or any other digital asset. The ERC721 standard allows for the creation of unique tokens, known as minting, and the transfer of these tokens between peers [8]. The first applications of NFTs that drew significant attention were collectible collections and games. CryptoPunks, one of the first NFT projects, was launched by Larva Labs in 2017 and consisted of 10,000 unique collectible characters stored as ERC721 tokens. The main component of an NFT is its metadata, which defines the NFT through its name, description, image, and attributes.

Unlike ERC20 tokens, the price of an NFT within a collection is not fixed, which adds a level of complexity to NFT trading. In many other investment instruments, the price of an asset is determined by supply and demand, but the price of a unique NFT token operates on multiple levels. The first level is the price of the collection among other collections. While there is not a clear definition of the price of a collection, the industry widely accepts the floor price of a collection (the lowest listed price of an NFT in the collection) as an indicator of the collection's value. The second level is the price of an NFT relative to other NFTs in the same collection, generally determined by the rarity of the NFT. Rarity can be defined by the attributes of the NFT and how rare these attributes are within the collection. However, defining the rarity of an NFT is not straightforward, and different calculation methods assign different weights to the NFT's attributes. Hence, an NFT's rarity does not directly correspond to its price, but it does provide valuable insights for valuation. As a result, determining the price of an NFT is not a straightforward task and can be approached with different perspectives.

To facilitate NFT trading, NFT marketplaces have been established. These platforms allow users to explore collection information, NFT metadata, and price related data, as well as trade NFTs. Notable NFT marketplaces include OpenSea, Blur and

LooksRare. While most NFT marketplaces store their data off-chain (not stored on the blockchain), many share their data via REST API or through certain data providers. Marketplaces sustain their operations by charging a commission on each sale made on their platform or by issuing a fungible or non-fungible project. In addition to marketplace commissions, traders must also account for the gas fee, paid to the network miners for the transaction, and the royalty fee, paid to the creator of the NFT collection, when calculating their costs.

2.2 Network Analysis

Network analysis is a field of study that focuses on examining the relationships between nodes within a network. This analytical approach is often used in various domains such as social networks, biological networks, computer networks, and transportation and logistics. In network analysis, nodes represent the individual entities that are interconnected. These connections, or edges, are established as nodes interact with one another. Numerous network metrics have been developed to facilitate this analysis, and the choice of the metrics depends largely on the specific problem being addressed.

In this study, we utilize network analysis to identify influential nodes in the network, which correspond to key traders within the NFT market. This identification process is crucial because, as indicated in [10], the top 10 percent of traders are responsible for 85 percent of total transactions. Moreover, these traders have been involved in 97 percent of all NFT trades made during the study period, which spans from June 23, 2017, to April 27, 2021. This period encompasses 6.1 million trades involving 4.7 million NFTs.

Given this background, we devised a network analysis method to determine the influential nodes within the network using network metrics. Various metrics are developed to understand the market belief of these nodes based on their transactional activity. These metrics are then discretized and included in the state space definition used in the Q-learning algorithm.

CHAPTER III

LITERATURE REVIEW

In the existing literature, no studies have proposed a comprehensive NFT trading simulation similar to our approach. Given that NFTs are a relatively new concept, most existing literature in NFTs primarily focuses on defining NFTs, examining their value proposition, and identifying opportunities and anomalies in the network.

Financially-focused research on NFTs incorporates these digital assets into their portfolio management approaches. For instance, the study by [11] reveals that adding NFTs into an investment portfolio diversifies the portfolio and increases the risk in the portfolio, but could also potentially improve overall portfolio performance. Similarly, other studies like [12] introduce NFTs into their portfolio optimization problems, evaluating their performance over a defined time period. However, these studies do not attempt to create an NFT trading scheme that addresses the specific challenges inherent in NFT trading.

Several studies have offered valuable insights that have informed our approach to different aspects of the challenges in NFT trading, such as wash trade detection, network analysis, appraisal model, and the application of the Q-learning algorithm to the trading problem. In the following sections, we will discuss these relevant studies in their respective context.

3.1 Wash Trades

The rise in popularity of NFTs has led to new ways for illicit activities, and recent research highlights the existence of wash trading and potential money laundering in this emerging market. Wash Trading, a practice in which a trader manipulates perceived asset value by trading between their own addresses, is becoming a notable concern.

A review of this phenomenon was conducted by [13], and 262 users were identified as suspected of such behavior. The study revealed that these activities, despite being undertaken by a small group, resulted in an overall profit of approximately $8.9 million for 110 traders.

Meanwhile, a separate study focusing on wash trading in NFT markets on the Ethereum blockchain from 2018 to 2021 found similar manipulative behavior [4]. By analyzing the 52 largest collections by volume, the authors discovered that about 3.93% of addresses were responsible for about 2.04% of sales transactions; this suggests potential market manipulation that could inflate trading volumes by as much as $149.5 million. They proposed a method to detect wash trading in the NFT market using transaction graphs. These graphs, composed of nodes (addresses) and edges (transactions), are analyzed for closed cycle patterns, which suggest suspicious activity. Traders might avoid these cycles for rapid trades, so transactions executed in short time frames or with minor price deviations are also considered suspicious. This method leverages the unique features of NFTs and the transparency of the Ethereum blockchain to observe individual account behaviors.

Both studies highlight the need for regulation to protect market integrity by highlighting the potential risks associated with NFT markets. While these illegal activities make up a small portion of overall transactions, their existence and potential growth requires serious attention to avoid undermining the integrity of the NFT market.

3.2 Q-Learning

The history of automated trading traces back to the 1970s, marking an era of significant technological advancements in financial markets [14]. Initially, trading algorithms were primarily designed to mimic the actions and strategies of human traders. However, with time and continuous refinement, these automated agents have gained complexity and sophistication. Currently, automated-trading is responsible for a considerable proportion of the trading volume in both traditional and cryptocurrency markets. In 2022, the global algorithmic trading market was valued at an impressive USD 2.03 billion [15]. Recently, NFTs have emerged as an attractive investment instrument, yet there is not any study in the academic literature that proposes an automated NFT trading schema.

Numerous strategies have been advanced in the literature, many of which models the problem with Markov Decision Process (MDP) and leverages the Q-learning algorithm to optimize trading outcomes. [16] introduces a new trading model, called Human-Aligned Trading (HAT), that is designed to align machine trading agents with human behaviors using reinforcement learning, specifically Q-learning. The model uses a novel multi-loss function combining supervised learning, single-step and multi-step Q-learning, and imitation learning to create trading behaviors that mimic those of human traders. However, the study acknowledges the limitation of reinforcement learning when dealing with partially observable Markov Decision Processes, as the model only accesses time-series data from the most recent days.

[17] focuses on the use of Q-learning to establish dynamic trading strategies in stock markets. Two distinct models are proposed for representing the discrete states of the trading environment. The first model uses the unsupervised learning method k-Means to define environment states of Q-learning by grouping historical stock market data into clusters. The second model bases its state formation on candlestick patterns. Tested on Indian and American stock markets, both models demonstrated superior performance over traditional Buy-and-Hold and Decision-Tree trading strategies.

[18] introduces a novel portfolio management framework that leverages Q-learning. This framework is composed of local agents, responsible for understanding individual asset behaviors, and a global agent that determines the overall reward function. Tested on a cryptocurrency portfolio, the study demonstrates that using multiple

deep reinforcement learning approaches offers promising results for dynamic portfolio optimization.

[19] showcases various techniques for optimizing trading strategy parameters using Q-Learning and Deep Reinforcement Learning frameworks with a focus on the Relative Strength Index (RSI) indicator. It establishes that the Deep Reinforcement Learning (DRL) approach with a Double Deep Q-Network (DDQN) setting, and the Bayesian Optimization (BO) approach can both produce positive average returns, with the former offering superior short-term trading results and the latter promising high long-term profitability. Strategies that depend on RSI indicator are widely used in classical trading algorithms, this strategy can be utilized as a benchmark for different trading algorithms. Trading problems requires multi-phase solutions. An example to this approach is [20], they addresses challenges in trading systems using reinforcement learning by introducing three new methods to maximize profits and mirror real market situations. It employs a deep neural network (DNN) regressor with the deep Q-network to predict share quantities, different Q-value action strategies of Q-learning algorithm for volatile markets, and transfer learning to control overfitting due to inadequate financial data. These researches do not focus on NFT trading but developes a similar trading approach to stock market and cryptocurrencies.

3.3 Network Analysis

As in [17], we need to define the states for the Q-learning algorithm for our agent, which in this case should represent the state of the NFT market. Assessing the market state accurately is challenging, requiring an in-depth analysis of market dynamics. A trading bot's primary objective is to maximize profit and minimize risk, both of which can be facilitated by analyzing the historical transactions of other traders.

The concept of "smart money" is described as "the funds invested by individuals or entities with extensive financial experience, knowledge, and a keen eye for lucrative opportunities" by CoinMarketCap [21]. [10] showed that transaction numbers by traders in the NFT market can be modeled with a power law with λ = -1.85. This distribution implies that a small number of traders, specifically the top 10%, are responsible for a large percentage of transactions, approximately 85% [10]. This skewness towards smaller values indicates a heavy-tailed and unbalanced distribution of transactions, where a few traders dominate the market. In light of these analyses, we decided to track the transactions of top traders, "smart money" or NFT whales, and define the market state according to their belief in the market.

Previous studies [10], [22], and [23] employed network analysis to understand the structure of the NFT network. They found that the structure of the NFT network shares similarities to that of general social networks [22], and graph network metrics can be effectively applied to the NFT network. In our study, we used network analysis to identify influential nodes, predicting the market state based on their activity.

3.4 NFT Apprisal

Digital asset valuation is a widely studied concept in literature but the sources in the literature that apply valuation models to the NFT valuation are scarce. In the context of digital asset valuation in NFTs, a recent study has proposed a methodology to assess the valuation of an NFT, aiding investment decisions in NFT software development [24]. The study suggests that an NFT's valuation is largely influenced by community engagement and scarcity, though these factors show weak and negative correlations respectively with total sales volume. The research's future scope includes refining the quantification of the community factor using various social media platforms and enhancing the accuracy of results through big data analysis, thereby supporting informed investment decisions in NFT projects.

[25] reviewed eight studies about NFT valuation, and the insights from these

papers informed the development of two LSTM machine learning models for predicting NFT sale prices. The models were applied to multiple multivariate time series datasets related to the NFT market. The first prototype achieved a Root Mean Squared Error (RMSE) of 0.2975, suggesting a need for improvement. The second prototype demonstrated superior performance with a reduced RMSE of 0.24.

Some other studies also worked on the relation between NFT price and social media interactions, and search trends. Twitter sentiment and related interactions significantly influence the value and growth of NFTs, as found in [26]. [27] and [28] have shown a direct relationship between search trends, public market data, and NFT valuations.

3.5 Liquidation Model

Our research presents the pioneering development of a liquidity model designed to predict the likelihood of selling a Non-Fungible Token (NFT), given its price, and a specified timeframe. Similar challenges have been addressed in other subjects, particularly in relation to real-world assets such as real estate. The study by [29], for instance, focuses on the 'Days on Market' (DOM), a metric that quantifies the duration an asset is available for sale. Their approach primarily aims to forecast the DOM based on an array of factors, including the house profile, community attributes, geographical specifics, and temporal elements. Although their approach is similar to ours, targets of the models diverge, with the DOM study predicting listing duration rather than the probability of a sale within a certain period.

However, our approach could potentially be applied to solve the liquidity problem in real estate given a similar feature space. Supporting our hypothesis, the DOM study identified that the average DOM of recently sold houses emerged as a significant indicator, mirroring our own methodology where we evaluate the ratio of sold tokens to recent listings among among 'neighbor' tokens, defined by their rarity and sale price.

Our approach also finds validation in another study [30] that examined the correlation between the duration a house is on the market, its price distribution, and its liquidity. The researchers found that although the time-on-market is a significant factor affecting liquidity, the unique price distribution of houses also plays a vital and independent role. This conclusion affirms our strategy of examining the liquidation of a group of similar tokens based on their prices and rarities, when predicting NFT sale likelihood.

CHAPTER IV

PROBLEM DEFINITION

The primary problem the thesis addresses revolves around the challenges of trading non-fungible tokens (NFTs) within a rapidly evolving market landscape. While the appeal and novelty of NFTs have attracted significant attention, their unique characteristics and the complexity of the underlying blockchain technology present numerous challenges for traders. The primary issues we manage in this thesis are as follows:

- Valuation of NFTs: Unlike fungible tokens, NFTs are unique, each having distinct attributes that affect their valuation. The lack of standard valuation metrics and the need to consider each token's individual characteristics makes fair and accurate valuation a significant challenge.

- Liquidity: Due to the uniqueness of each NFT, the liquidity in the NFT market operates differently from traditional markets. Sellers must wait for a buyer who values their NFT as they do, leading to an inherently illiquid market.

- Market State Definition and Trading Strategies: The dynamics of the NFT market are complex, influenced by numerous factors. Defining the market state and designing optimal trading strategies based on these market states presents a further problem.

Additionally, a less complex yet equally important challenge is wash trades. These are essentially manipulative trades that create a false impression of market activity, thus making it more difficult to accurately value tokens and collections and comprehend the genuine demand of investors in the market.

In the light of these challenges, the aim of this thesis is to propose a systematic, data-driven approach for NFT trading. We model the trading problem with MDP. For the simulation and finding optimal state-action map, we develop a trading simulator based on the Q-learning algorithm, which also includes machine learning models to estimate an NFT's fair price and liquidity probability. We also removed wash trades from historical sale data using a rule-based approach and provide this clean data to the models.

This thesis is committed to proposing, implementing, and evaluating this solution within the context of a specific NFT collection, Bored Ape Yacht Club (BAYC). Ultimately, we aim to demonstrate that our approach can provide valuable insights and tools for the understanding of NFT market by effectively addressing its inherent challenges and offering viable solutions.

CHAPTER V

DATA ARCHITECTURE

Management of the NFT data is notably challenging when compared to handling data of other financial instruments, such as traditional stocks and cryptocurrencies. This challenge arises from the fact that the nature of NFTs are complex, they have both collection-level and token-level data.

In this chapter, we cover the details of our data sources, the methodologies employed for data collection, our approach to data storage and the architecture of the data pipeline.

5.1 Data Collection

5.1.1 NFT Market and Data

OpenSea [31], being one of the earliest established and most popular NFT marketplaces, played a critical role in shaping the initial landscape of NFT trading. Based on the ERC721 contract, any newly minted token is instantly listed on OpenSea. OpenSea was the market leader until December 2022, Figure 3, before being overtaken by Blur [32].

OpenSea has a long-standing presence in the NFT market and has made substantial strides in improving the user experience over time. However, shifts in the demands and expectations of traders have led to the rise of new platforms, such as Blur. Blur differentiates itself through its simplistic and intuitive user interface, but this is not the sole reason behind its market leadership. Blur launched its platform alongside an airdrop campaign - a marketing strategy aimed at increasing product awareness. This campaign and Blur's decision to refrain from taking commissions from sales on their platform and offering flexible royalty payments to the creators of

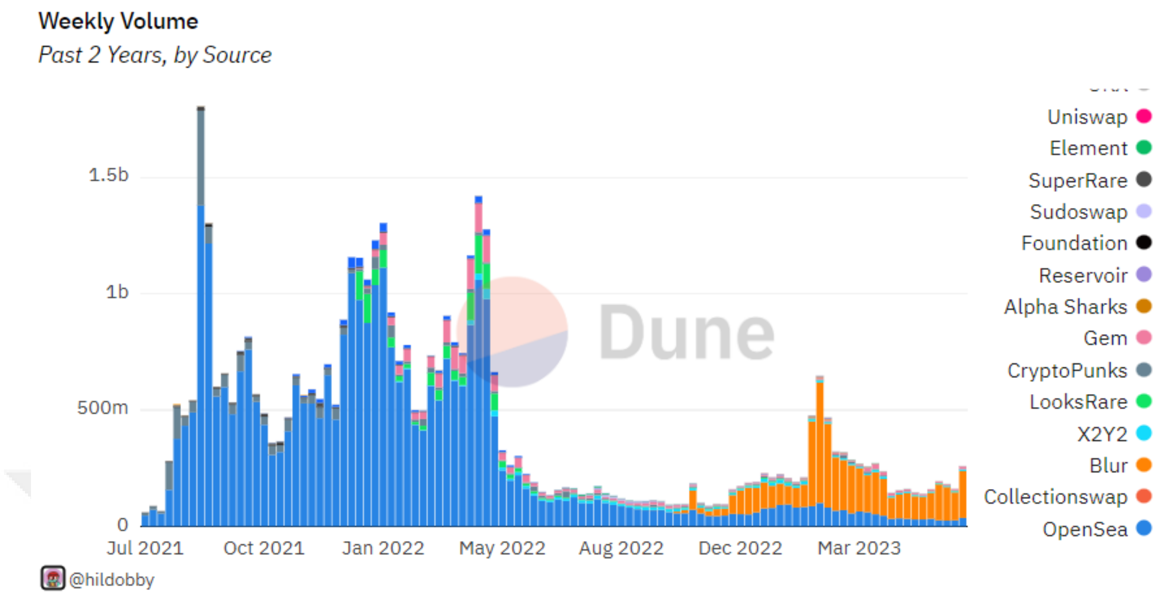

Figure 3: Market share of NFT marketplaces [3]

NFT collections, significantly contributed to Blur's rise to the top.

While OpenSea and Blur are prominent players in the NFT marketplace arena, the landscape is much more diverse and includes numerous other platforms. Some marketplaces cater to specific types of NFTs or support specific blockchains. For instance, Blur currently only supports Ethereum blockchain, while OpenSea is more diverse, supporting up to nine different blockchains, including Ethereum.

Figure 3 illustrates that OpenSea and Blur are not the sole competitors in this space. Other noteworthy platforms include marketplace aggregators such as Gem and NFT marketplaces like X2Y2 and LooksRare. Moreover, certain NFT collections, like CryptoPunks, have launched their own dedicated marketplaces. At the time of this study, the CryptoPunks marketplace ranks as the 5th highest volume marketplace.

Most of these platforms store their data in a centralized database, with NFT listings and bid data constituting a significant portion of this data. NFT trading methods vary, with the most common examples being listing NFTs for a fixed price or auction. Traders can also make a bid for an NFT, and if the owner accepts, the NFT is transferred to the bidder.

To construct an NFT trading simulator, we need to collect blockchain and marketplace data. As the NFT data has a complex nature, there is not any data source that can completely structure the data yet. Some NFT marketplaces provide their data through APIs but most of them do not. Recently some data aggregators emerged, that collects the internal marketplace data and serves through their servers. These aggregators collect data via APIs, scraping, or through private agreements with the marketplaces and serves through Rest APIs and websockets.

Reservoir [33], one of the most prominent data aggregators in the NFT market, covers a significant number of marketplaces, particularly those with high volumes. However, their data coverage starts from April 2022, and our study required older data as the creation of BAYC was in May 2021. Therefore, we supplemented Reservoir's data with data from OpenSea for the period before April 2022.

To conduct a comprehensive study of an NFT collection, we need to assemble both static and market data. Required data can be summarized as:

- Collection metadata: These are the descriptive details about the collection as a whole.
- Token metadata: Information regarding each individual token in the collection.
  - Token attributes: The characteristics that differentiate each token within a collection.
  - Multimedia associated with tokens: Images, videos, or any other media associated with the tokens.
- Market data: This comprises data relevant to the trading of the NFTs.
  - Listing data: Information about the NFTs that have been listed for sale.
  - Sale data: Historical data of NFT transactions, including the prices and dates of sales.

Data extraction from the Reservoir API is a straightforward process due to its well-structured format compared to other sources. We collected all data they offer for our purpose. For the remaining required data, we utilized the OpenSea API.

5.1.1.1 OpenSea

OpenSea provides their data through two technologies: a REST API and a WebSocket. For our data collection, we primarily utilized the REST API. We required collection-level data such as royalty fees, contract address, and collection slug. While the OpenSea API provides an extensive array of data, our study primarily requires token listing-related data along with the aforementioned collection-level data. To accomplish this, we developed a Python script designed to retrieve the data from the OpenSea API and store it in a PostgreSQL database.

5.1.1.2 Reservoir

Reservoir operates as a data aggregator, collecting data from the majority of commonly used marketplaces within the NFT market. To offer their data, they utilize a REST API and WebSocket. They offer both blockchain data and marketplace data. For our study, we accessed token attributes from Reservoir's 'tokens' endpoint, and subsequently stored this information within our PostgreSQL database. Furthermore, we retrieved available listing and sale data from their 'events' and 'sales' endpoints respectively.

5.2 Data Storage

We chose to use PostgreSQL as our primary data storage system due to its ease of use and efficiency in handling relational data. This system stores all our raw data. However, considering the intensive data needs of our simulator, which runs hundreds of episodes, and the requirement to execute multiple simulators in parallel, we realized that a NoSQL structure would be more advantageous for the simulator's operations.

Thus, we store the processed data needed by the simulator locally, on our machine, within the Windows file system.

A graphical representation of our PostgreSQL database structure is shown in Figure 4. It includes several tables:

1. **Collections**: Contains metadata for each NFT collection.
2. **Tokens**: Contains metadata for individual NFTs.
3. **Trait Groups**: These refer to the categories of traits in a collection, such as skin, eyes, mouth, etc.
4. **Traits**: This refers to the specific trait values, like blue eyes, red eyes, white skin, etc.
5. **Tokens-Traits**: This table links each token to its corresponding traits.
6. **Market Events**: Stores all events related to token listings and sales.
7. **Model-fits**: Holds the fitted models for NFT appraisal and the probability of sale.
8. **Model-predictions**: Contains price predictions generated by the appraisal model for each token.

As previously mentioned, while certain data is stored in our PostgreSQL database, some derived data is held locally on our machine for simulator operations.

5.3 Data Pipeline

Figure 5 represents the process of data flow from our data sources to our PostgreSQL database. Initially, we fetch collection metadata and listing events from OpenSea, specifically for data before April 2022. This data undergoes a processing phase where

- history schema
    - market_events
    - model_fits
    - model_predictions
- static data schema
    - collections
    - tokens
    - tokens_traits
    - traits
    - trait_groups

Figure 4: PostgreSQL database structure

we generate the final listing data for the simulator and calculate hourly floor prices for the given period.

Listing data comprises of three types of events: listing, delisting, and sale. Each event is recorded separately, and it is necessary to match listing events with corresponding delisting and sale events, if they exist, to identify the duration for which the listing was valid. Another factor that can invalidate a listing is its expiration, defined by the owner at the time of listing. Consequently, we need to identify which event occurred first, be it expiration, delisting, or sale, to determine the final invalidation date of the listing. Upon labeling the listings with their respective invalidation dates, we proceed to calculate the hourly floor price for each collection. To achieve this, we identify the valid listings at each timestamp, where the creation date is earlier than the timestamp and the invalidation date is later. This allows us to determine a snapshot of the floor price at that timestamp.

We utilize the Reservoir API to fetch token metadata, listing events post-April 2022, and sale events. In addition to basic metadata like name, description, and image of the tokens, we also obtain token attributes. The structure of listing events differs from that provided by the OpenSea API, comprising more event types and information. However, Reservoir provides an order ID for each event, which eliminates

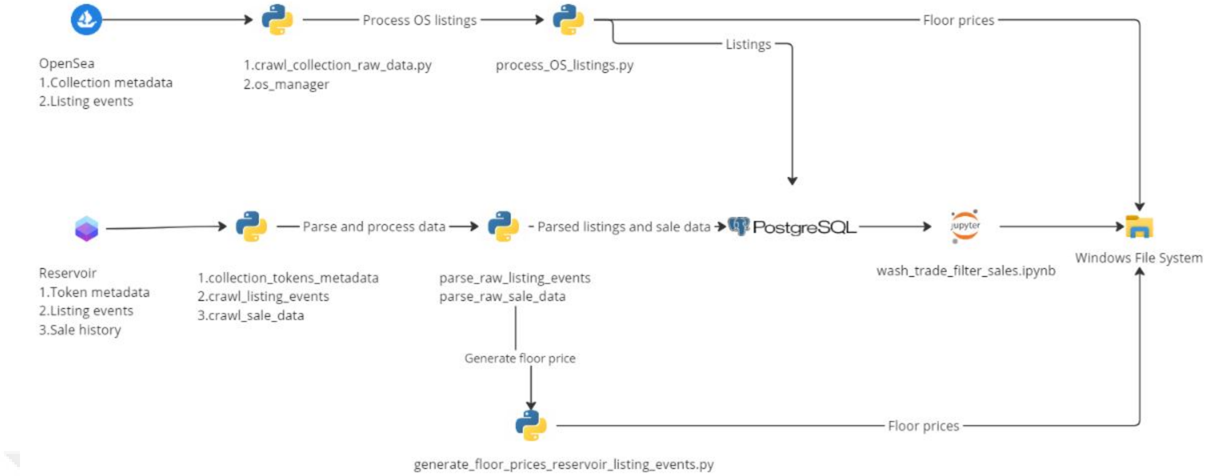

Figure 5: Data pipeline

the need for matching events to identify valid listings. After parsing the data, it is directly used in the simulator. A Python script named 'generate-floor-prices.py' is used to calculate the floor prices for the collections. Similar to the process with OpenSea data, this script identifies the valid listings at each timestamp and calculates the floor price. The floor price data is stored in parquet files.

The sales data is subjected to a wash trade filter, which is a rule-based method designed to remove any fictitious trade. The methodology underlying the wash trade filter is explained in greater detail in the Methodology chapter.

CHAPTER VI

METHODOLOGY

In order to successfully carry out trading in the NFT market, one has to take into consideration the unique characteristics that differentiate it from traditional investment markets. A prevalent assumption in the NFT market is using the floor price as a reference point to value an entire collection. Yet, when developing a trading simulator for NFTs, additional problems must be addressed.

Two prominent issues that arise when trading in the NFT market are the valuation of individual tokens and the overall liquidity of the market. To solve these problems, we employed two key strategies:

- To estimate a fair price for individual tokens at every timestamp, we developed a model leveraging XGBoost, a popular machine learning algorithm known for its efficiency and performance. The algorithm creates an ensemble of decision trees and adjusts the weights of the model iteratively to minimize the errors of predictions.

- The second challenge, illiquidity, refers to the absence of a liquidation price for individual NFT tokens, necessitating the development of an estimation model for backtesting applications. In the context of NFT trading, liquidity plays a crucial role as it can significantly impact the feasibility of sell actions. We addressed this issue by using a model developed by our research group that predicts the probability of selling a token at a given timestamp and selling price. This model was modified to our specific problem and integrated into our system.

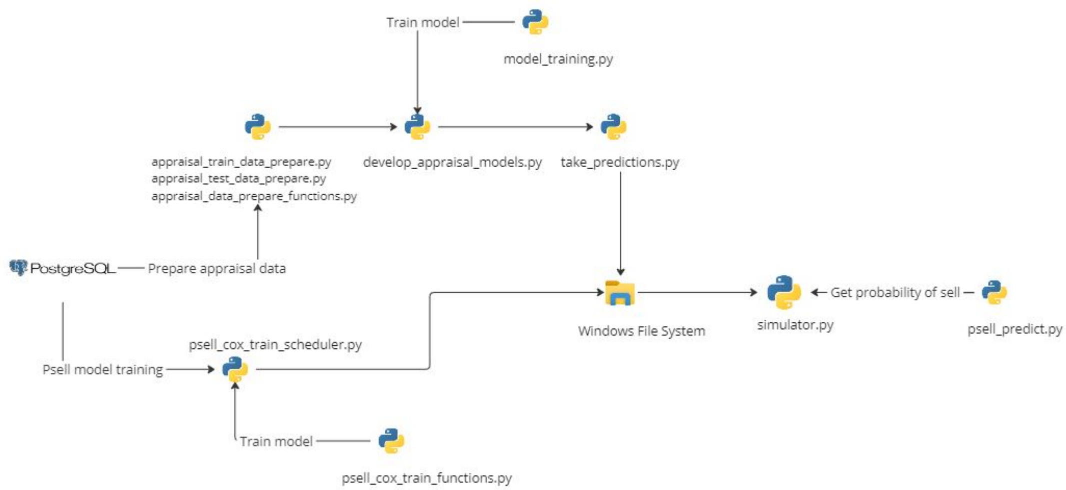

Figure 6: Machine learning pipeline

However, we also require a decision-making mechanism to determine trading actions. To address this, we model the problem as a Markov Decision Process (MDP) and implement a Q-learning algorithm, a form of reinforcement learning, to develop a trading simulator. This simulator, using Q-learning, solves the MDP, effectively determining the optimal sequence of actions—buying, holding, or selling NFT tokens—that would maximize the cumulative rewards over time.

In this section, we will delve deeper into the specifics of our complete approach, which includes applying wash trade filter, defining market state using network analysis, the handling of liquidity and appraisal challenges, modeling the trading problem as a Markov Decision Process (MDP), and applying the Q-learning algorithm in our trading simulator. The flow of machine learning pipeline is visualized in Figure 6.

6.1 Wash Trade Detection

Wash trades are common in NFT trading. Unlike most traditional investment instruments which are regulated and where wash trades are illegal, blockchain and NFTs are yet to be regulated. While there is no clear definition of wash trades in the NFT market, several ways exist to create artificial volume and price trends. This subject has been widely studied in the literature for different contexts.

To understand the impact of wash trades on the market, one can refer to the analysis conducted by Chainalysis [34]. They analyzed the wash trades in the crypto market in 2021, estimating the total cryptocurrency-based wash trades to be around \$8.6 billion. A portion of this volume is from the NFT market. Their analysis indicates an upward trend in wash trading. Figure 7 can be examined to gain insights into recent wash trading volumes in the NFT market. While the calculation methods used in different studies may vary, most of them employ basic and reliable methods to detect wash trades. These methods can provide a basic understanding of how common wash trades are in the market.

As demonstrated in Figure 7, since the beginning of 2022, wash trading has been a consistent presence in the market. There are notable spikes, and on some days, wash trading volume even surpasses normal trading volume. Most of these major spikes align with significant market events. For instance, the significant spike around January 10, 2022, coincides with the launch of LooksRare. They initiated an airdrop campaign based on user trading volumes, causing traders to inflate their trading volumes artificially to gain more from the airdrop. In most cases, traders manipulate the market by buying and selling their own tokens between their wallets. The study by Coffman [35] revealed that a single wallet managed to earn 3.5 million dollars in a day by generating \$44,522,815 in trading volume through the purchase and sale of a single token between the owner's wallets.

There exist numerous instances that entice traders to generate artificial trading volumes and price trends, however, understanding the reasons behind the wash trades is not the primary focus of this study. Our study requires eliminating wash trades to accurately reveal and comprehend the dynamics of the NFT market.

Our wash trade detection approach consists of five rules. Rules 2, 3, and 4 were also employed in an analysis by a data analyst in Dune Analytics [36].

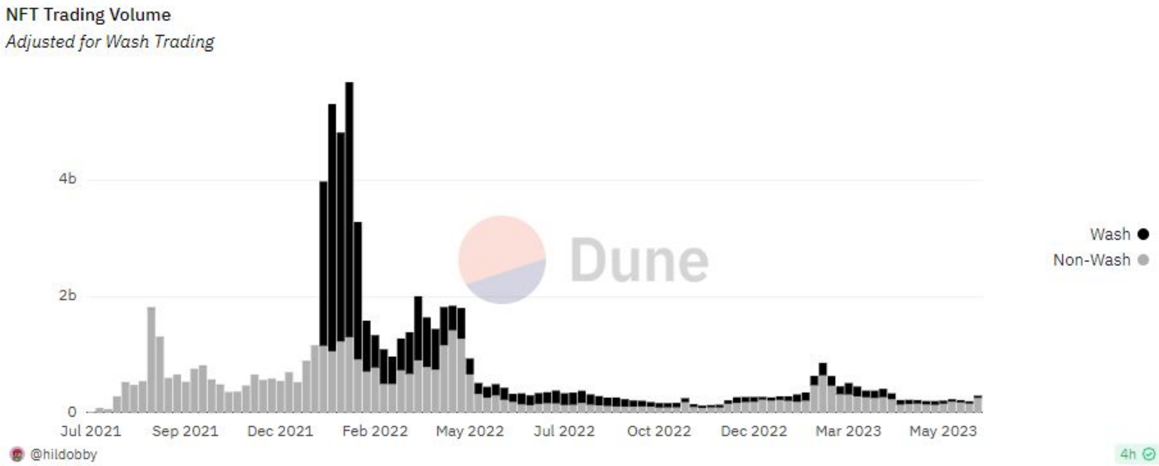

Figure 7: NFT wash trade volume [4]

1. A trade is a wash trade if the same wallet buys the token twice from a different wallet in any time period. For instance, $token_i$ is bought by w_i from w_j, multiple times.

2. A trade is labeled as a wash trade if a token is sold back and forth between two wallets in any time period.

3. If a token is purchased by the same wallet three times or more from any wallet within any time period, it is considered a wash trade.

4. Some traders buy their own tokens from the same wallet, thus the fourth rule checks if the buyer and seller wallets are identical.

5. The final rule states that a trade is a wash trade if the price of the transaction is more than 20 times the median price of the collection's last 50 transactions or if the price is less than one-tenth of the median price of the collection's last 50 transactions. Median price is a fair indicator, as it robust to outliers and reflects traders' perception of the market. The thresholds were chosen to filter out the top and bottom 1 percent of the data.

At the end of our analysis, we flagged 6.58% of transactions as potential wash trades for the selected collection. An analysis by a data analyst in Dune Analytics [36] found that 41.94% of transactions in the NFT market are wash trades. The volume of wash trades varies between collections, as traders may prefer certain collections for various reasons. This could be due to a lower or nonexistent royalty fee in the selected collection, or because the collection is new and traders are attempting to manipulate the price and volume to attract more users.

6.2 Network Analysis

Our purpose to include graph networks in the study is to determine influencial nodes in the network to further analyze their transactions. We leveraged NetworkX library in Python to develop the graph network. The network is a directed weighted graph DG(V,E), where V symbolizes the vertices (or nodes) of the graph network, which in this context are wallets within a transactional network, and E signifies the edges, that is, transactions occurring between the wallets. The mathematical representation of an edge, $e_{i,j}$, is expressed as follows:

$$e_{i,j} = (v_i, v_j) | v_i \in V, v_j \in V \tag{1}$$

In this equation, $e_{i,j}$ signifies transactions (or edge) from wallet v_i to wallet v_j. A maximum of 2 edges, in opposing directions, can be established between any 2 nodes. The number of transactions associated with each edge is considered as the weight, denoted as w, of that edge and is provided as an edge parameter. w is mathematically defined as:

$$w(v_i, v_j) = k | v_i \in V, v_j \in V, k \in \mathbb{N}^+ \tag{2}$$

In this equation, k represents the number of NFT transactions from wallet v_i to wallet v_j. Given that the directionality is vital for understanding which nodes are accumulating NFTs, we preferred directed graph to undirected one, where the edge

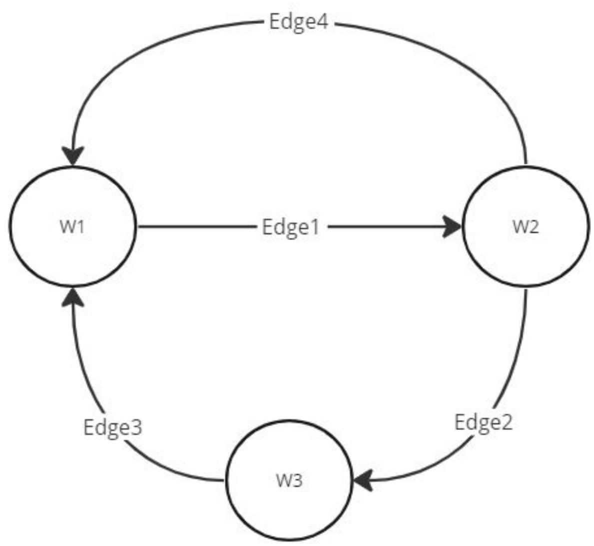

Figure 8: Network structure example

direction is disregarded.

In our graph network, centrality measures are employed to determine the centrality of a node in the network. Essentially, these measures compute a score for each node based on its edges and subsequently create a ranking. Higher ranks indicate a more central node, implying greater influence within the network. We computed five different centrality measures:

- Degree centrality,
- Closeness centrality,
- Betweenness centrality,
- Eigenvector centrality,
- PageRank centrality

In addition, we applied the HITS algorithm to identify hubs and authorities within the network. This algorithm assigns a score to each node based on its incoming or outgoing links to detect authorities and hubs, respectively.

In each episode, the network is constructed using the last 500 transactions. This is a fundamental assumption that excludes wallets that are no longer active traders or holders of the collection. While the number of recent transactions to analyze can be adjusted, we found 500 to be an adequate starting point for understanding the beliefs of the top 10 percent. With the developed network model, the aforementioned metrics are calculated. Nodes are ranked based on these calculated scores, and the average of these ranks is used as the final rank of each node. We chose the top 10 percent of nodes as influential nodes as it has been observed that they account for a significant portion of all transactions within the network [10].

To capture the latest market sentiments of these influential nodes, we calculated additional metrics. These metrics represent the percentage of the number and volume of recent transactions carried out by influential nodes, and the average rank of the tokens they are trading. Alongside the network metrics, the Relative Strength Index (RSI) of the floor price is included in the feature space. These features are then fed into a Decision Tree Classifier (DTC) to categorize the market state into three different labels.

Labelling is performed based on the average floor price of the next three days. If this value is at least 5 percent higher than the current floor price, the market is labelled as 'bullish', and 'bearish' if the value is at least 5 percent lower. Otherwise, the label is 'neutral'. This percentage value can also be tuned. The DTC model has an accuracy of 0.79 on the test data, which, while not perfect, is a notable improvement over the zero rate classifier accuracy of 0.74. Further research could potentially enhance this model, but this falls beyond the main aim of this study.

6.3 Appraisal Model

Appraising an NFT is much like valuing a piece of art, especially for collectible NFTs such as Bored Ape Yacht Club. Unlike traditional assets, there is no fixed liquidation price for an NFT. Instead, its price is determined by the token's owner and the prospective buyer at the time of the transaction. Consequently, the price of an NFT relative to other tokens in the same collection is not constant, but differs at each transaction. Moreover, the pricing pattern of tokens within a collection cannot be interpreted with a linear relationship. However, our goal is not to ascertain the exact price of every token but rather to estimate a fair price for tokens at any given timestamp and market condition.

We created a feature space comprising 80 features which can be categorized under three main groups. The first feature group includes the features that represent the relative rarity of the token in the collection. As the traders mostly cares about the rarest trait of the token, rarity scores of top 3 traits are also included in the features. An overall rarity score in a token-level is calculated. Rarity score of the token i is given as:

$$rarity\ score_i = \sum_{T} \frac{S}{n_{t_j}} | j \in \mathbb{N}^+, i \in 0, 1, 2, ..., S \tag{3}$$

where S is the token supply of the collection, $T(i)$ is the traits of token i and n_{t_j} is the number of tokens that have the trait t_j. The second feature group represents market dynamics, including price and volume-related metrics, as well as some momentum indicators. The last feature group represents the liquidity of the market for the collection and the distribution of holders within the collection. Sale and transfer data are used to calculate the features in this group. These features are calculated for each transaction in the history of the collection.

For each timestamp of the training duration, corresponding feature space is constructed and an XGBoost model is fitted. Two level of hyper-parameter tuning is applied. First, a random search is applied to search the best hyper-parameters within the defined search space. Then, a grid search is applied to find the best hyper-parameters around the ones found in the random search. Appraisal model is developed on a weekly basis and fair prices for other days is adjusted with the change in floor price.

The standard metric employed to measure the performance of NFT appraisal models is the Median Absolute Percentage Error (MdAPE). One of the major advantages of MdApe is that it is more robust against outliers compared to the Mean Absolute Percentage Error (MAPE).

$$\text{MdAPE} = \text{median}\left(\left\{\frac{|y_i - \hat{y}_i|}{|y_i|}\right\}_{i=0}^{n}\right) \times 100$$

In this equation, y represents the actual price of the token, $\hat{y}$ is the predicted price of the token, and n is the number of vectors in the dataset. The model's performance is evaluated using the sales data. For every sale, we have an appraisal of the token. Thus, we can match the sales data with the appraisals at the time of sale. Our model exhibits a MdAPE of 0.106 on sales data from June 6, 2021, to December 31, 2022. The histogram of prediction errors, Figure 9, shows errors concentrated around zero, but we have also relatively high errors. The main reason of high errors is the price of an NFT determined by the owner and the buyer at the time of sale with the knowledge of recent market condition. Hence, the price of an NFT can be changed dramatically in a short time period but our model did not be fitted before every transaction with the latest data because of the computational cost.

Comparing the MdAPE of our model to those developed by industry researchers helps gauge its performance. For instance, Upshot, a leading NFT analytics platform, reported a MdAPE of 0.106 for the Bored Ape Yacht Club collection [37], the same as our model's MdAPE. NFTBank's model, another benchmark, had a slightly better MdAPE of 0.087 [38]. However, neither Upshot nor NFTBank stated the time period of their test data. Without knowing the time period of their test data, a fair comparison cannot be achieved. But, we can say that our model has a comparable performance with the other models in the industry.

The main goal of our model is to calculate an 'opportunity gap' using the estimated fair price and listing price of each token. This helps us identify tokens that have a

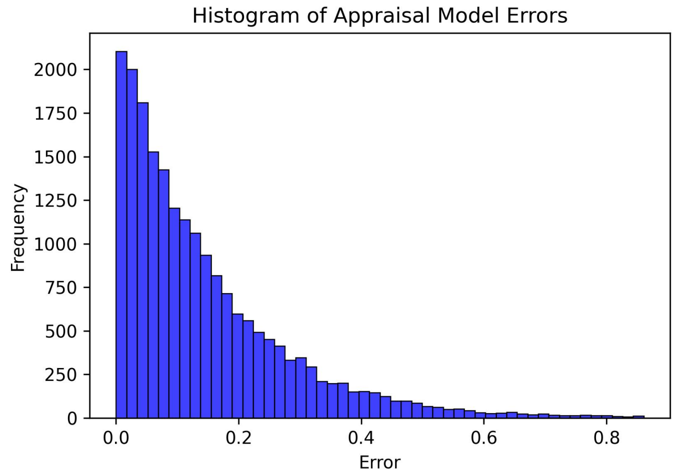

Figure 9: Histogram of appraisal errors

high potential at the specific timestamp. In addition, we use the estimated fair price to set a selling price at the time of listing. Hence, given these requirements, we can confidently state that our model performs well for our purposes.

6.4 Liquidity Model

Liquidity model is trained for each timestamp and stored locally for both training and testing processes, as the pipeline shown in Figure 6. To prepare the data, at each timestamp we matched the last 14 days of listing data before the given timestamp with the sale data from the same period. This allowed us to determine which listings were sold and which were not. Listings that were listed shortly were excluded since they might have sold if they had been listed for a longer period - a minimum threshold of 3 hours was set for this purpose.

For the remaining listings, we apply a neighborhood search for each record. The neighborhood is defined as a region within 1% above and below the rank and price of the token. Within this neighborhood, we compute two parameters: the ratio of the number of sold listings to the total number of listings, and the number of days until a sale was made, if it did indeed happen. The former is used as the exogenous variable, along with the rarity of the token and the price of the listings normalized by both the fair and floor prices, in the Cox model. The latter - the number of days until a sale is made - is used as the 'lifetime' of the subject. The event of a sale (notated as 1 for 'sold' and 0 for 'not sold') is used as the 'death' observation in the model. Model fits are taken for each timestamp with in the duration, and these fits are later used in the simulation for liquidity prediction.

In terms of evaluating the performance of our model, it is worth noting that no existing liquidity models for the NFT market currently exist for comparison. We tested the model's performance using a separate test set, constructed using data from January 3rd to June 10th, 2023. For each day within this period, we estimated the probability of a sale for the active token listings using the fitted models. Our model achieved an average accuracy of 0.83 in this time-frame, with recall and precision values of 0.70 and 0.88, respectively.

Our model's precision greatly exceeds its recall, which implies that false positive samples are considerably lower than false negatives. In other words, the model is more likely to incorrectly label sold listings as negative (false negatives) than it is to falsely label unsold listings as positive (false positives).

6.5 Trading Simulator - Q-learning

The Markov Decision Process (MDP) is a widely used stochastic decision-making model for determining sequential action decisions based on the environment. In MDP it is assumed that all states have Markov property, which means each state only depends on the previous state [17]. An MDP can be formulated by:

- S is the state space that denotes the possible states of the environment,
- A is the action space that defines the possible actions that the agent can undertake,
- P is the transition matrix that represents the transition probability from state

s to s' given a,

$$P(s'|s,a) = \Pr(s_{t+1} = s'|s_t = s, a_t = a) \tag{4}$$

- R is the immediate reward function to calculate the reward given state and action,

$$R_a(s,s') = \mathbb{E}[r_{t+1}|s_t = s, a_t = a, s_{t+1} = s'] \tag{5}$$

In each state the agent takes an action, reaches a new state, and receives an immediate reward. The objective of the MDP is to find an optimal policy π that maps the state space to the action space which maximizes the cumulative reward. But in some problems, like in our problem, future rewards and transition probabilities are not explicitly known. In this case, the Q-learning algorithm, a model-free, off-policy reinforcement learning (RL) method, can be utilized to solve the problem.

is decay rate The Q-learning algorithm holds a Q-table that maps the states of the environment to potential actions. The goal of this algorithm is to learn the optimal Q-table. Depending on its current state, the agent takes an action and computes a reward for this action. The state space is defined based on the market states coming from the network analysis and the fungible holding level in the portfolio. The portfolio fungible level at time t is defined by:

$$fr_t = \frac{fb_t}{pv_t} \tag{6}$$

$$pv_t = fb_t + nfev_t \tag{7}$$

where fr is the fungible ratio, fb is fungible balance that is ETH balance in our case, pv is the portfolio value and $nfev$ is the total estimated value of NFTs. fr values are equally split into 4 ranges between 0 and 1 for state space definition and the DTC that predicts the market state ms depending on the network influencer activity can take 3 different values as explained. So, our model has 12 different states and state at time t can be defined as:

$$s_t = (ms_t, fr_t) \tag{8}$$

Algorithm 1 Q-learning NFT Trading Simulator

Inputs: **initial timestamp (**t_{init}**)**, **last timestamp (**t_{last}**)**, **number of episodes (M)**, ϵ_{min}, ϵ_{init}, **decay rate** (ζ)

Initialize Q table and portfolio, P

$t \leftarrow t_{init}$

for $i = 1$ to M **do**

  while $t \leq t_{last}$ **do**

    Update floor price, fair price and corresponding sales data for time t

    Get market state estimation, ms_t,

    Get fungible state, fr_t,

    Define state tuple as,

$$s_t = (ms_t, fr_t)$$

    if $t \neq t_{init}$ **then**

      Update Q-table:

$$r_t = \frac{(pv_t - pv_{t-1})}{pv_{t-1}} \times 100$$

$$Q(s_{t-1}, a_{t-1}) = Q(s_{t-1}, a_{t-1}) + \alpha * (r_t + \gamma * \max_a Q(s_t, a) - Q(s_{t-1}, a_{t-1}))$$

    end if

    Define *infeasible actions*, and apply action masking,

    $Q_{masked} = mask\_actions(Q, s_t, infeasible\ actions)$

    Select a_t with ϵ-Greedy policy, pick a random number and calculate ϵ, as,

$$\epsilon = max(\epsilon_{min}, \epsilon_{init} * exp(-\zeta * k)),\ k = (t - t_{init}) \text{ in days}$$

$$a_t = \begin{cases} \text{random action from } Q_{masked}(s_t), & \text{if random number} < \epsilon \\ arg\max(Q_{masked}(s_t)), & \text{otherwise} \end{cases}$$

    Take action

    Update P

  end while

  Initialize P, ϵ and t

end for

In each timestamp our agent needs to choose an action at the current state. We first need to define current state. A graph network is constructed with recent data, centrality metrics is calculated for ranking the nodes and influential nodes are determined. We need to examine the transactional activities of the influential nodes, some metrics are designed to reveal market belief of these nodes. These metric space is given to DTC and market state prediction is obtained. Then, fungible state is calculated and the current state tuple is defined. When the agent knows the current state, it can update Q-table for the previous action. It calculates the reward of the previous action as the percentage difference between the portfolio value of the current state and the previous state. Q value of the state-action pair is updated using Temporal Difference method.

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha * (r_t + \gamma * \max_a Q(s_{t+1}, a) - Q(s_t, a_t)) \tag{9}$$

where α is learning rate and γ is discount rate.

The actions of the agent is a tuple of buy and sell decisions,

$$a_t = (buy\ action_t, sell\ action_t) \tag{10}$$

Buy action can be 'buy' or 'do nothing'. Sell actions are 'list bullish', 'list bearish', 'list neural', and 'do nothing'. Some actions are not feasible for some states, to prevent the agent to choose an infeasible action, action masking is applied. We need to examine opportunity gaps of active listings, their listing price and ETH balance in the portfolio for buy actions. We calculate opportunity gap of the active listing l_j at the time t for token i as,

$$opportunity\ gap(l_j) = \frac{fair\ price_{it} - list\ price_{it} - fees}{list\ price_{it}} | j \in \mathbb{N}^+ \tag{11}$$

Fees include the creator fee and marketplace fee. We exclude gas fees in the calculation, because it is negligible compared to the floor price of the selected collection. Rarely gas fees increases dramatically, but lasts a very short time.

Listings with prices that exceed the wallet balance or have an opportunity gap lower than the threshold of 0.05 are removed. The agent then determines which actions are feasible using filtered listings. If not any listing left after constraints or ETH balance is not enough to buy any listing, actions with 'buy' decision are infeasible. Selling actions are not feasible and set to 'do nothing', if the agent does not have any NFT tokens.

Algorithm 2 Action Masking in Q-learning

Require: Q, s_t, $infeasible\_actions$

```
function mask_actions(Q, s_t, infeasible_actions)
    Q_masked ← Q
    for each action in infeasible_actions do
        Q_masked(s_t, action) = −1
    end_for
    return Q_masked
end_function
```

The agent then updates the ϵ value, generates a random number and uses it to decide the next action. If the number is lower than ϵ, it chooses a random action decision from the feasible options; otherwise, it picks the feasible action tuple with the highest Q-value for the current state. The agent takes buy and sell actions consecutively. If the buy action is 'buy', the agent calculates a reward for available listings, by multiplying opportunity gap of the token and its list price. The agent buys up to two tokens that provide the highest reward without exceeding the wallet balance. If it cannot find suitable tokens, it changes the buy action to 'do nothing'.

If the agent holds NFTs and the selling action is not 'do nothing', it will determine selling prices for the tokens. The price depends on the action definition: for 'list_neutral', it's the fair price; for 'list_bullish' or 'list_bearish', it's the fair price increased or decreased by 10%. The agent then estimates the likelihood of sale of the token within a day using the liquidity model. A random number is generated for each token and if the number is lower than the likelihood, the token is sold. After performing the actions, the agent updates the ETH balance and the tokens in the portfolio. This process repeats until it reaches the last timestamp in the training set, then it resets the portfolio and timestamp for the next episode.

The test simulator operates in a similar manner as the training component, albeit with a few minor differences. At the beginning of the simulation, the agent uses the Q-table obtained from the final episode of training. The ϵ value is initiated at 0.3 and decays at the same rate as during training to ensure the learning process continues. However, the agent is biased towards exploitation over exploration by setting the ϵ value below 0.5.

Another difference involves the application of the wash trade filter. In the training phase, the wash trade filter was applied prior to simulation as all data was available at once. In a real-world scenario (and in the backtesting process), only historic data would be available, and thus, the wash trade filter is applied at each step during the test phase.

The test simulator records all transactions and performance metrics, enabling subsequent analysis in the following section.

6.6 Benchmarks

Given the absence of any comprehensive NFT trading simulator to compare our performance, we developed two basic benchmarks that commonly used in traditional stock and cryptocurrency markets. These strategies attempt to understand underlying trends and make buy and sell decisions accordingly.

The Relative Strength Index (RSI) is a popular momentum indicator frequently used in such strategies. It has been the basis for various trading strategies for several decades, primarily because of its simple logic: buy when the market is oversold and sell when the market is overbought [39]. We can apply this principle to the NFT market using the floor prices of collections as our metric.

We implemented two versions of RSI-based strategy, both utilizes the same methodology for RSI as described in [39] and [19], but the portfolio management is different. The first version adopts an aggressive approach, buys tokens with the full ETH balance in the wallet whenever the RSI value returns to 30 from below. It lists for sale all tokens in the portfolio, when the RSI value returns to 70 from above. The second version is a more cautious approach designed to reduce risk. It only uses half of the ETH balance for purchases at each buy decision but sell all tokens in the portfolio at once. Additionally, if the fungible portion of the portfolio falls below 25%, then all the remaining ETH balance is used for purchases. In both strategies, tokens to be bought are selected from listings close to the floor price. A price multiplier is calculated as the ratio of token cost to the floor price at the time of buying. This multiplier is used to determine the sell price by multiplying with the floor price at the time of sell. That is, it is assumed that tokens' price is normalized over floor price.

We also implemented a simple but often robust benchmark strategy: buy and hold. As the analysis in [39], where various strategies including buy and hold, Kelly gambling, and RSI-based approaches were applied to 16 different stocks, the buy and hold strategy proved to be the best performing one for 10 out of the 16 stocks. It generally performs well in the long term. Despite our limited dataset, it can provide a useful performance metric and gives us an idea of how it might perform in the context of NFT trading.

CHAPTER VII

RESULTS AND DISCUSSION

7.1 Simulation Results

In the training set, BAYC's floor price demonstrated a predominantly upward trend. However, this uptrend shifted to a major downtrend after May 2022, as illustrated in Figure 10. The final floor price is 49.5 times higher than the floor price at the beginning of training data. Despite minor fluctuations, the price mostly experienced long periods of either major uptrend or downtrend.

At the last episode of training, our agent effectively grew its portfolio value from an initial 100 ETH to a substantial 3976 ETH. But the increase in floor price is higher than the performance of the agent. Can we say buy and hold strategy performs better than our model in the training period? We are observing a historical price of one of the best assets in the market. While this asset has shown an impressive performance since its minting, the ability to identify such a successful collection with complete certainty of its potential price trajectory is impossible. A 'buy and hold' strategy necessitates the identification of such unusually good-performing collections for it to be successful. However, given the daily release of hundreds of new collections, finding these potential 'unicorns' is not an easy task. Contrary to the 'buy and hold' strategy, our model aims to be profitable in both bullish and bearish market conditions. Theoretically, if we had not imposed the constraint of limiting the purchase action to maximum of two tokens, our model probably could have performed better. The agent could have possibly accumulated more tokens when the floor prices were low, resulting in an increased portfolio value at the end. However, without this buy constraint, the model could under-perform during bearish market conditions.

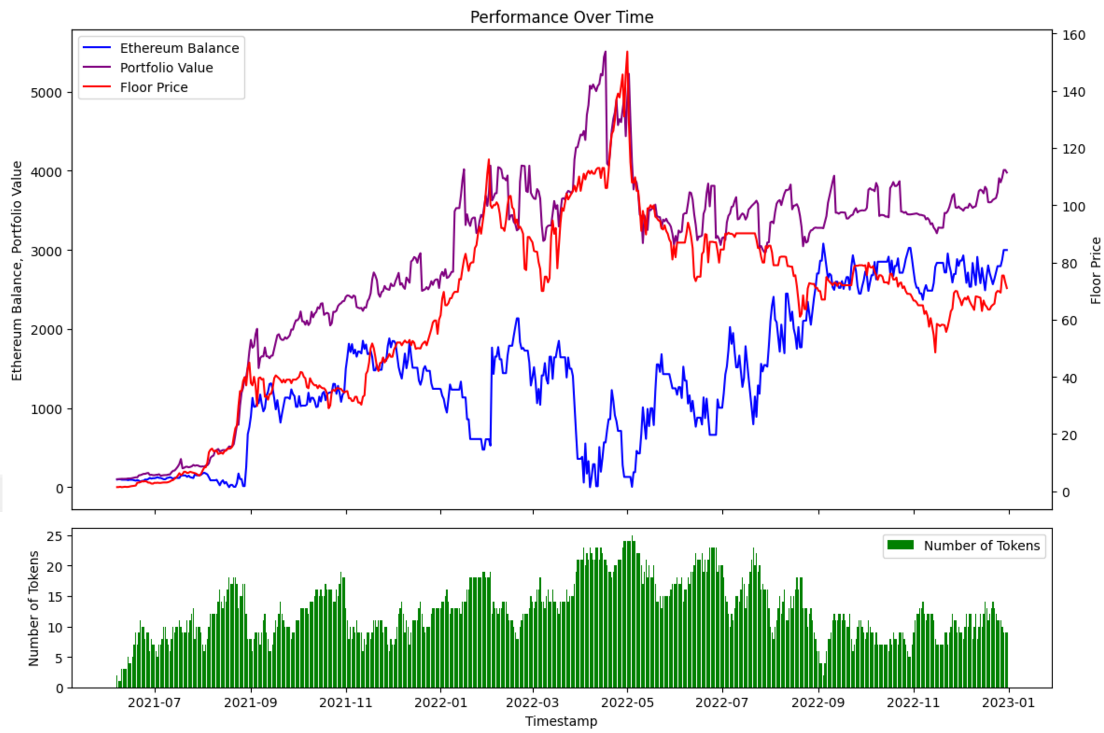

Figure 10: Train simulation performance over time at the last episode

The agent's strategy of purchasing a maximum of two tokens at a time is also designed to reduce the risk of missed opportunities due to insufficient balance in an uptrending market, as well as the risk of sudden declining prices. Therefore, the agent's strategy aims to maintain equilibrium between the ETH balance and the total estimated valuation of the NFTs in the portfolio, prioritizing a balanced approach to risk management over an aggressive accumulation strategy.

After the minting phase, the floor price remains relatively low for a few months. However, at the beginning of August, the floor price experiences a sudden surge, rising from 1 ETH to 10 ETH. This surge accelerates until September 2021. During this period of substantial price increase, the agent increases its NFT holdings. After a point, the price begins to exhibit greater volatility, resulting in corresponding fluctuations in the agent's ETH balance and token holdings. When the all time high is seen at the end of April 2023, the agent reaches the highest number of NFTs in its portfolio. The floor price suddenly drops from 154 ETH to under 100 ETH, and continues to decrease in the upcoming months. During the sudden drop, the agent tactically reduces its NFT holdings and increases its fungible (ETH) holdings, demonstrating a prudent risk management strategy. As a result of these strategic transactions, the portfolio value begins to diverge from the floor price during this period. The opportunity-centered buying strategy contributes to an overall increase in the portfolio value, even in turbulent market conditions.

The agent starts the test period with an optimized Q-table learned during the training phase. For the test simulation, the agent begins with a budget of 1000 ETH, and ϵ initialized at 0.3. Notably, the test data displays a predominant downtrend, unlike the bullish trend observed in the training data.

The performance of the model during the test period is represented in Figure 11. The green line in this figure illustrates the worst-case portfolio value, calculated as the sum of the ETH balance and the total estimated value of the NFTs in the wallet (determined by multiplying the number of tokens by the floor price). Given that the highest liquidity is generally found around the floor price, this line provides an estimate of the portfolio value in case of immediate liquidation. The agent maintains an ETH balance exceeding 500 ETH until the start of May 2023. During this period, the quantity of tokens in the portfolio is in a lower scale. However, after this point, the agent increases its NFT holdings. As the agent holds a high ETH balance during periods of significant price declines, it partially saved the total portfolio value. Towards the end of the test simulation, the floor price experiences the lowest prices after a long time and the agent seizes this opportunity to increase its NFT holdings. During the final 40 days of the test period, 20 days were marked as 'bullish' while only 5 days were 'bearish'. The agent tends to increase NFT holdings during a bull market, as expected.

At the end of test period, total portfolio value is 1218 ETH, reflecting a gain of 21.8%. Performance statistics of the test simulation is summarized in Table 1. In a

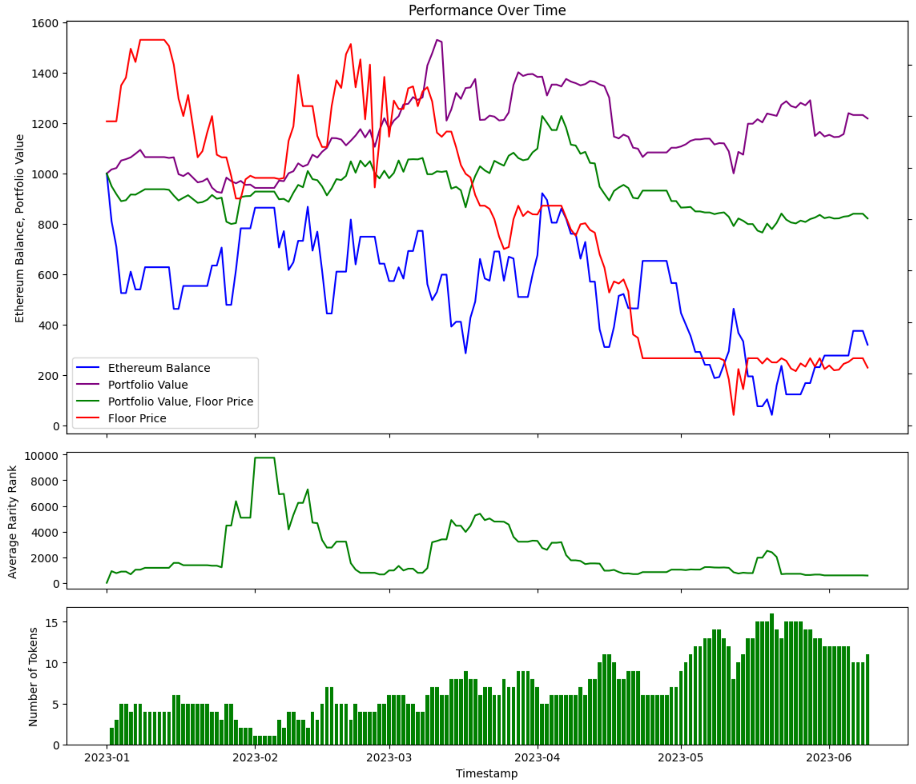

Figure 11: Test simulator performance over time

period of 160 days, the agent bought 96 NFT tokens, 85 of them is sold and 11 token remains in the portfolio at the end. Remarkably, 60% of these sales resulted in profit, with an average gain of 21.14%. The remaining transactions experience an average loss of 18.3%. On average, tokens were sold in 9.51 days after purchase and 11 tokens in the portfolio cannot be sold in 28.91 days. Since the market was predominantly bullish towards the end of the test period in terms of influencer activity, the agent does not list the tokens or lists with a bullish price. That is, the agent want to keep the tokens at hand. During the test period, the floor price decreased by 33.53%, from 69.49 ETH to 45.59 ETH.

The green line ends the test set with a portfolio value of 821.7 ETH, consisting of 320.2 ETH and 11 NFTs. As the floor price essentially serves as a liquidation price for all tokens in the collection, it informs us the worst-case scenario performance, selling all tokens immediately at floor price. In this scenario, the agent does not seem

Table 1: Test simulation results summary

Total number of transactions	181
Total number of tokens bought	96
Total number of tokens sold	85
Number of listed days, sold tokens	9.51
Number of listed days, unsold tokens	28.91
Profitable sale ratio	0.6
Profitable sales average profit (%)	21.14
Non-profitable sales average loss (%)	-18.3
Final portfolio value (ETH)	1218.49
Final profit (ETH)	218.49
Floor price change (%)	-33.53

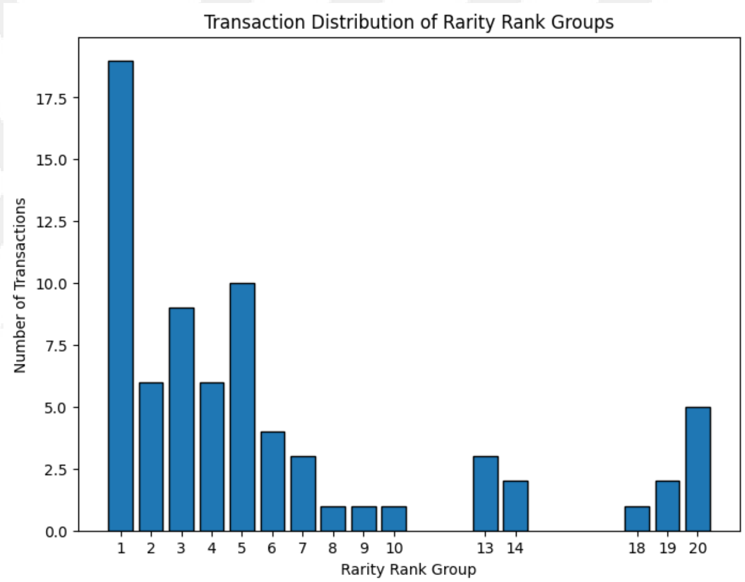

Figure 12: Rarity distribution of transactions in the test period

to have realized a profit, but it's important to remember that selling tokens at the floor price might not be fair, given that the portfolio includes rare items.

As it can be seen in the second subplot of Figure 11, the agent generally purchases items with higher rarity. The distribution of the number of transactions of the agent in the test period across different rarity groups is visualized in Figure 12. Each rarity group includes 500 NFTs, with the first group being the rarest. The most recently traded tokens are in the first rarity group and the rest are accumulated in rare groups. The agent generally prefers buying rare items. However, during periods

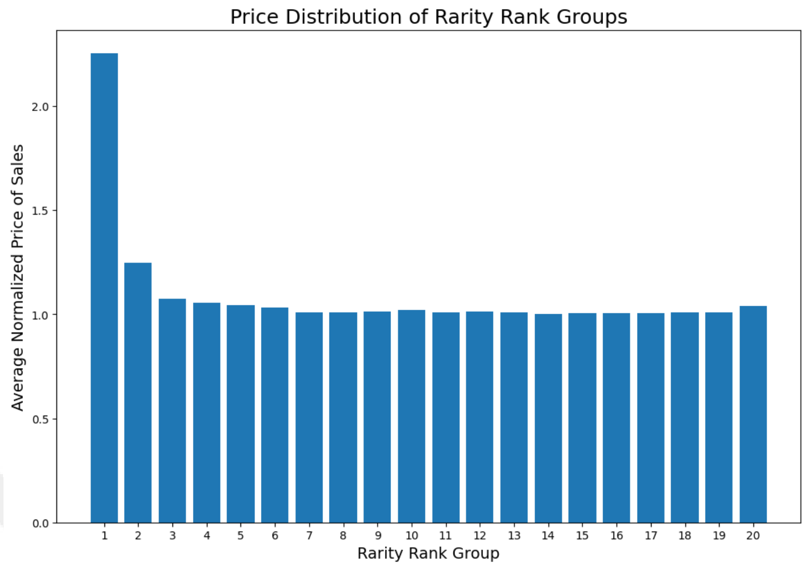

Figure 13: Average normalized prices of rarity rank groups in the historical sales of BAYC

of price decline, it often liquidates these rare assets, which results in an increase in the portfolio's average rarity rank in these periods. Despite this, the average rarity rank of the tokens in the portfolio remains below 15% for approximately 57.5% of the test period. At the end of test period, the average rarity rank of 11 tokens is 5.6%. This demonstrates that the tokens' valuations are likely to be significantly higher than the floor price, as the rarity of a token is a strong indicator of its valuation. In Figure 13, the normalized price distribution of rarity groups of BAYC collection from 1st of January 2023 to 10th of June 2023 is shown. Price is normalized by the floor price of the collection. The average normalized price of the rarest 2 groups are 2.25 and 1.25, respectively. The remaining groups have similar average normalized prices that fluctuate around 1.

In the final portfolio, 6 out of 11 tokens is in the first rarity group (the rarest) and 2 are in the second rarity group. It is obvious that final ETH balance would be substantially higher than the value represented by the green line, even if the tokens were liquidated at their average price. The agent tends to select rare items because it calculates a reward based on the opportunity gap and the price of the token. For rarer tokens with a positive opportunity gap, the reward in ETH is higher. This is not surprising because, as Figure 13 shows, the prices of tokens in common rarity groups tend to fluctuate around the floor price. In contrast, tokens in rarer groups exhibit high volatility, high risk, and potentially high returns.

7.2 Benchmark Results

The first and simplest strategy we examined was the buy-and-hold strategy. As previously mentioned, if a trader buys all tokens at the floor price at the start of the test period and sells all tokens at the floor price at the end, the portfolio value would decrease by 33.53%. If the trader builds a portfolio that includes both common and rare tokens, using the average price at the start and end of the duration can provide an approximation of performance. According to this approach, using a 3-day average price, the performance of the strategy would result in a loss of 31.67%.

The second benchmark is the RSI strategy. Throughout the test period, the first version of daily RSI strategy does make some token purchases, but the sell condition is never met. Consequently, the simulation ends with 16 tokens and 4.3 ETH remaining, for a total portfolio value of 733 ETH. A more cautious version of this strategy also only bought tokens, ending with 14 tokens and 158.4 ETH for a total value of 796.1 ETH. To increase transaction activity, we implemented another benchmark that executes buy and sell actions on an hourly basis. This approach ended with 7 tokens and 27.2 ETH left for a total value of 350.7 ETH. A more cautious version of the hourly strategy had a total value of 214.9 ETH at the end.

While the RSI strategy is widely used across various investment instruments, this performance underlines that the dynamics of the NFT market are distinct. This proves the need for developing strategies tailored specifically for the NFT market to achieve success.

Among the benchmark strategies, the best performance was achieved by cautious daily RSI strategy, which recorded a 20.39% loss. However, our Q-learning simulation, enriched by network analysis and an appraisal model, yielded a profit of 21.14%. This represents a significant improvement over basic strategies and provides a promising start to introducing a comprehensive trading approach to the NFT market.

7.3 Suggestions on Future Work

The point we have arrived is the start of NFT journey, a lot to be improved. No other investment tool shares the structure and the ever-evolving dynamics of NFTs. The initial challenge lies in identifying the market state and selecting one or more collections to trade over time. We leveraged network analysis and transactional analysis to determine the market state. However, our approach could be refined, and alternative strategies could be proposed. These analyses may also be useful in selecting collections for multi-collection trading.

Our appraisal model is quite efficient when compared to the existing models in the literature and the NFT industry. There can be an extention of the model that also accounts the social media interaction of the collection to identify the trends and holder attributes. Furthermore, the state and trend of cryptocurrency markets directly impact the NFT market. Hence, a study on the relationship could enhance our model as pointed out in existing studies.

Liquidity model is one of its kind in the literature for NFT domain. Some similar efforts can be seen in house valuation problem. The liquidation problem in the NFT market might be the toughest to handle. It is as hard as predicting house or car liquidation problems but with greater price volatility. We believe our approach is promising with its performance but more complex models that accounts a wider feature space can outperform our scores.

There is a lot to improve in our simulation and Q-learning approaches. Several issues arise in the simulation process, such as determining trading frequency and buy

& sell strategies. More aggressive trading strategies, like permitting a large number of purchases, could enhance performances in bullish markets. Different price adjustment coefficients could be applied for the selling strategy. Also, increasing trading frequency might help capture as many opportunities as possible.

Additionally, there is scope to refine the Q-learning process. Defining the state and action space is one of the hardest yet crucial issues. Expanding the state and action space could allow the model to distinguish states better but might lead to a sparse Q-table. The number of episodes to run is also a critical factor, given our limited data points high numbers might cause overfitting. The parameters of the ϵ-Greedy algorithm are also crucial. We started each episode with $\epsilon = 1$, in each episode the agent tries to explore for most of the time. ϵ can be dynamically reduced as the number of episodes increases. Defining the reward function is as important as defining the state and action space. Since the model learns according to the rewards of the action taken in a given state, the reward function can entirely shape the agent's behavior. The main challenge in defining a reward function for the problem is the illiquid nature of the market, requires longer-term investments. However, defining the reward function based on the long-term performance of trading actions is not applicable given the complexity of attributing profit or loss in the outcome to individual trading actions. Buy and sell options of multiple tokens create infinite possibilities in the long time horizon.

CHAPTER VIII

CONCLUSION

This study has opened new frontiers in understanding the dynamics of the NFT market, leveraging advanced machine learning techniques and network analysis. Our results indicate that a comprehensive approach, Q-learning model in our case, can indeed outperform traditional investment strategies in this relatively new and volatile market.

Our findings underline the critical importance of tailoring trading strategies to the unique characteristics of the NFT market. A simple buy-and-hold strategy or conventional indicators like RSI, while successful in other markets, may not yield favorable outcomes. The use of our Q-learning model enriched by network analysis and an appraisal model achieved considerable profit, even in a bearish market.

However, it's essential to note that the NFT market is highly complex and continues to evolve rapidly. An effective NFT trading strategy demands a deep understanding of the value proposition of various NFT collections, rarity distribution, token valuation, liquidation challenges, and various other concepts like bids, offers, auctions, price dynamics. These requirements possess a challenging path for most investors, needing a lot of time, effort, and dedication to fully understand and navigate successfully.

While our current simulator subject to further enhancement in many ways, we believe it provides a solid foundation for future exploration into NFT trading simulation and backtesting. Hopefully, our approach will offer a comprehensive overview of the complexities involved in NFT trading and possible solutions to address them.

REFERENCES

[1] ilemi, “Erc & eip starter kit.” https://dune.com/ilemi/erc-and-eip-starter-kit, 2023.

[2] satoshist, “Nft market volume, by blockchain.” https://dune.com/queries/2770431/4612097, 2023.

[3] hildobby, “Nft market overview.” https://dune.com/hildobby/NFTs, 2023.

[4] V. von Wachter, J. R. Jensen, F. Regner, and O. Ross, “NFT wash trading: Quantifying suspicious behaviour in NFT markets,” *CoRR*, vol. abs/2202.03866, 2022.

[5] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system.” https://bitcoin.org/bitcoin.pdf, 2008.

[6] E. Team, “Ethereum whitepaper.” https://ethereum.org/en/whitepaper/, 2023.

[7] J. Wu, J. Liu, Y. Zhao, and Z. Zheng, “Analysis of cryptocurrency transactions from a network perspective: An overview,” *Journal of Network and Computer Applications*, vol. 190, p. 103139, 2021.

[8] J. E. j. N. S. n. William Entriken (@fulldecent), Dieter Shirley ¡dete@axiomzen.co¿, “Erc-721: Non-fungible token standard,” *Ethereum Improvement Proposals*, no. 721, 2018.

[9] J. B. Cho, S. Serneels, and D. S. Matteson, “Non-fungible token transactions: Data and challenges,” *Data Science in Science*, vol. 2, no. 1, p. 2151950, 2023.

[10] A. L. D. G. F. e. a. Nadini, M., “Mapping the nft revolution: market trends, trade networks, and visual features.,” *Scientific Reports*, vol. 11, no. 1, p. 20902, 2021.

[11] H. Ko, B. Son, Y. Lee, H. Jang, and J. Lee, “The economic value of nft: Evidence from a portfolio analysis using mean–variance framework,” *Finance Research Letters*, vol. 47, p. 102784, 2022.

[12] M. Mazur, “Non-fungible tokens (nft). the analysis of risk and return,” 2021.

[13] C. TEAM, “Crime and nfts: Chainalysis detects significant wash trading and some nft money laundering in this emerging asset class.” https://blog.chainalysis.com/reports/2022-crypto-crime-report-preview-nft-wash-trading-money-laundering/, 2022.

[14] X. L. Z. L. Z. Z. Huang B., Huan Y., "Automated trading systems statistical and machine learning methods and hardware implementation: a survey," *Enterprise Information Systems*, vol. 13, no. 1, pp. 132–144, 2019.

[15] F. B. Insights, "The global algorithmic trading market size was valued at $2.03 billion in 2022 is projected to grow from $2.19 billion in 2023 to $3.56 billion by 2030." https://www.fortunebusinessinsights.com/algorithmic-trading-market-107174, Feb 2023.

[16] Z. J. Ye and B. W. Schuller, "Human-aligned trading by imitative multi-loss reinforcement learning," *Expert Systems with Applications*, vol. 234, p. 120939, 2023.

[17] J. B. Chakole, M. S. Kolhe, G. D. Mahapurush, A. Yadav, and M. P. Kurhekar, "A q-learning agent for automated trading in equity stock markets," *Expert Systems with Applications*, vol. 163, p. 113761, 2021.

[18] B. M. Lucarelli, G., "A deep q-learning portfolio management framework for the cryptocurrency market," *Neural Computing and Applications*, vol. 32, pp. 17229–17244, 2020.

[19] M. Tran, D. Pham-Hi, and M. Bui, "Optimizing automated trading systems with deep reinforcement learning," *Algorithms*, vol. 16, no. 1, 2023.

[20] K. H. Jeong G., "Improving financial trading decisions using deep q-learning: Predicting the number of shares, action strategies, and transfer learning," *Expert Systems with Applications*, vol. 117, pp. 125–138, 2019.

[21] CoinMarketCap, "Smart money." `https://coinmarketcap.com/alexandria/glossary/smart-money`, 2023.

[22] S. Casale-Brunet, P. Ribeca, P. Doyle, and M. Mattavelli, "Networks of ethereum non-fungible tokens: A graph-based analysis of the erc-721 ecosystem," in *2021 IEEE International Conference on Blockchain (Blockchain)*, (Los Alamitos, CA, USA), pp. 188–195, IEEE Computer Society, dec 2021.

[23] P. Khati, "Measurement, analysis, and insight of nfts transaction networks," 2022.

[24] A. Cheung and J. Keung, "On the scrutinization of the nft valuation factors," in *2022 29th Asia-Pacific Software Engineering Conference (APSEC)*, pp. 564–565, 2022.

[25] J. Branny, R. Dornberger, and T. Hanne, "Non-fungible token price prediction with multivariate lstm neural networks," in *2022 9th International Conference on Soft Computing Machine Intelligence (ISCMI)*, pp. 56–61, 2022.

[26] A. Kapoor, D. Guhathakurta, M. Mathur, R. Yadav, M. Gupta, and P. Kumaraguru, “Tweetboost: Influence of social media on nft valuation,” WWW ’22, (New York, NY, USA), p. 621–629, Association for Computing Machinery, 2022.

[27] Y. Kaneko, “A time-series analysis of how google trends searches affect cryptocurrency prices for decentralized finance and non-fungible tokens,” in *2021 International Conference on Data Mining Workshops (ICDMW)*, pp. 222–227, 2021.

[28] C. Pinto-Gutiérrez, S. Gaitán, D. Jaramillo, and S. Velasquez, “The NFT Hype: What Draws Attention to Non-Fungible Tokens?,” *Mathematics*, vol. 10, p. 1, January 2022.

[29] H. Zhu, H. Xiong, F. Tang, Q. Liu, Y. Ge, E. Chen, and Y. Fu, “Days on market: Measuring liquidity in real estate markets,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’16, (New York, NY, USA), p. 393–402, Association for Computing Machinery, 2016.

[30] N. Kotova and A. L. Zhang., “Liquidity in residential real estate markets,” tech. rep., American Economic Association, 2020.

[31] “Opensea website.” `https://opensea.io/`. Accessed: 2023-06-10.

[32] “Blur website.” `https://blur.io/`. Accessed: 2023-06-10.

[33] “Reservoir website.” `https://reservoir.tools/`. Accessed: 2023-06-10.

[34] C. TEAM, “2022 biggest year ever for crypto hacking with $3.8 billion stolen, primarily from defi protocols and by north korea-linked attackers.” `https://blog.chainalysis.com/reports/2022-biggest-year-ever-for-crypto-hacking/`, 2023.

[35] C. Coffman, “Wash trading: Who, what, why, and what should we do about it?.” `https://blog.cryptoslam.io/wash-trading-who-what-why-and-what-should-we-do-about-it/`, 2022.

[36] hildobby, “Ethereum nfts wash trading.” `https://dune.com/hildobby/nfts-wash-trading`, 2023.

[37] Upshot, “Upshot api documentation website.” `https://v0.docs.upshot.xyz/api/contract-addresses`, 2022.

[38] NFTBank, “Bored ape yacht club website.” `https://nftbank.ai/collection/ethereum/0xbc4ca0eda7647a8ab7c2061c2e118a18a936f13d`, 2023.

[39] P. Marek and B. Sediva, “Optimization and testing of rsi,” 2017.

VITA

Süleyman Kamalak earned his Bachelor's Degree from Boğaziçi University, specializing in Mechanical Engineering. After finishing his undergraduate studies, he started his journey into the world of Data Science for his Master's degree, under the guidance of Asst. Prof. Erinç Albey. He has a deep interest in the fascinating domain of blockchain technology. Recognizing its potential as an expansive, publicly accessible data source and its promise as a pivotal technology shaping the future of data and the internet. His primary research focus lies in developing descriptive and predictive models that incorporate statistical analysis, machine learning algorithms, and deep learning techniques.