Article | Received 15 August 2024; Accepted 19 December 2024; Published 26 December 2024
https://doi.org/10.55092/blockchain20240010

ARTEMIX: A community-boosting-based framework for airdrop hunter detection in the Web3 community

Yuyang Qin[1, 2, 3], Tengfei Ma[1, 2, 3], Hongzhou Chen[3, 4] and Haihan Duan[1, 2, *]

[1] Artificial Intelligence Research Institute, Shenzhen MSU-BIT University, Shenzhen, Guangdong, China
[2] Guangdong-Hong Kong-Macao Joint Laboratory for Emotion Intelligence and Pervasive Computing, Shenzhen MSU-BIT University, Shenzhen, Guangdong, China
[3] The Chinese University of Hong Kong, Shenzhen, Guangdong, China
[4] CKB Eco Fund, Singapore

* Correspondence author; E-mail: duanhaihan@smbu.edu.cn.

Abstract: Airdrops represent a pivotal strategic instrument for Web3 projects, serving to distribute free tokens and motivate early adoption. However, the popularity of these tokens has fueled the emergence of airdrop hunters—individuals who exploit multiple transactions to acquire disproportionate amounts of tokens unfairly. This phenomenon threatens the integrity and fairness of the Web3 community. Current detection methods struggle with high false-positive rates, harming legitimate users, and require significant computational resources for training. Furthermore, these methods face challenges in adapting to the evolving tactics of airdrop hunters, leading to diminished detection accuracy and efficiency. We introduce ARTEMIX, a community-boosting-based framework that integrates custom-engineered features and community detection techniques to identify airdrop hunters in NFT transactions. Using data from the Blur NFT market, ARTEMIX demonstrates superior accuracy and efficiency, outperforming existing graph-based inference models, achieving an F1 score of 0.898. This approach provides a scalable and effective solution to anomaly detection in the Web3 ecosystem, promoting a more secure and equitable environment for token distributions.

Keywords: airdrop; airdrop hunter; sybil detection; Web3; boosting

1. Introduction

In recent years, airdrops have emerged as one of the most prominent topics in the Web3 ecosystem. In the context of blockchain technology and cryptocurrencies, an airdrop refers to the free distribution of tokens or coins to a large number of wallet addresses [1]. Airdrops are typically used as a marketing strategy to increase awareness and adoption of a specific blockchain project or token[1, 2]. Normally, airdrops can incentivize user engagement through actions such as following the project on social media, sharing updates, or contributing to the project's development[1]. A critical aspect of airdrops is their role in decentralization; more centralized projects tend to rely on a few large token holders, a scenario viewed unfavorably by the Web3 community[3]. Hence, airdrops can foster a robust user community, essential for the long-term success of blockchain projects[4].

To qualify for an airdrop, users often need to hold a specific amount of another cryptocurrency or complete particular tasks, such as engaging in on-chain activities, participating in ecosystem projects, or purchasing related NFTs[1]. However, if a project's tokenomics are not well-designed, strategies such as airdrop can have detrimental effects[5]. While they may create short-term excitement and increase token prices, the subsequent rise in circulating tokens can lead to long-term price declines[6]. Additionally, projects may be targeted by fraudsters who create hundreds of bots (fake accounts on social networks) and participate in dozens of airdrops daily to increase their chances of winning tokens[7]. These individuals, known as "airdrop hunters," collect wallet addresses and interact with contracts to obtain lucrative token giveaways. By exploiting blockchain anonymity to register multiple accounts and interact with DApps for maximum profit, these individuals undermine long-term development by rapidly selling tokens and driving down asset prices, thus harming enthusiasts who wish to hold tokens long-term and actively engage in community activities[4, 8]. These actors undermine decentralization by concentrating tokens in a few hands, disrupt financial inclusion by excluding genuine users, and harm system efficiency by introducing fraudulent transaction patterns. Therefore, airdrop teams must seek an efficient and effective method to identify airdrop hunters hidden among a large number of addresses.

Despite the emergence of airdrops and their corresponding hunters as a new business model and community, there is still limited research on the topic. The study by Fan *et al.* [8] demonstrates identifiable and observable patterns in the activities of airdrop hunter addresses. A typical example is “wash trading,” where airdrop hunters often repeatedly transfer assets between two or more addresses[9]. This mimics normal on-chain behavior, creating deceptive interaction data. As the scale increases, such behavior develops more complex strategies, eventually forming large clusters of airdrop hunters composed of hundreds of addresses[4]. Zhou *et al.* propose the ARTEMIS, a systematic airdrop hunter detection based on graph learning. The study by Victor[10] and Liu [7] introduces heuristic algorithms based on searching typical trading patterns to identify wash trading and suspicious sybil addresses. These works provide valuable insights and inspiration for identifying airdrop hunters and reveal many regular patterns in token transfers by airdrop hunters.

However, their performance in practical detection tasks may not be satisfactory for the following reasons:1)In real airdrop events, the data volume may be tens or even hundreds of times larger than in simulated experiments. Similar graph learning algorithms require significant computational resources and training time. Such costs increase significantly with data growth, and graph-based search algorithms’ detection time grows exponentially when dealing with complex and extensive real interaction data. This is particularly unfriendly for airdrop issuers in web3, who typically need to adjust strategies based on real-time community feedback; 2)Existing detection methods perform poorly in avoiding false positives of normal community users. From the perspective of incentivizing users through airdrops, we should focus more on not harming beneficial community participants[11];3)The interpretability of neural networks remains a notable challenge. Airdrop issuers require clear, understandable reasons why certain users are flagged, to make informed decisions and adjust their strategies accordingly. The black-box nature of neural networks makes it difficult to provide such transparency, leading to potential mistrust and inefficiency in their application[12].

Our work introduces ARTEMIX: **A**i**R**drop hun**TE**r detection via a boosting-based **M**erging, **I**ntegration, and e**X**traction framework. A novel model for identifying airdrop hunters, focusing on suspicious patterns typically found only in airdrop hunter accounts. In simple terms, we first built three key classifiers from different perspectives based on the revealed patterns of various airdrop hunters: 1)Nodes participating in typical hunter trading patterns. 2)Nodes with typical trading time patterns. 3)Nodes with typical trading characteristics. We also extracted important features from the constructed complete transaction graph: community detection and classification results based on 11 transaction node basic features using the Louvain algorithm [13]. Finally, we constructed a community-boosting-based framework, integrating the extracted node features and classifier detection results. This system focuses on representative trading features and greatly reduces the model’s training and inference costs while maintaining excellent scalability within the model framework. Experimental results show that our model outperforms existing graph inference models across various metrics in identifying airdrop hunters, providing a new and effective framework for addressing airdrop issues in the community. The codes are available at https://github.com/qCanoe/ARTEMIX-2024.git.

In summary, the contributions of this work are threefold:

- We introduce ARTEMIX, a novel community-boosting-based framework that integrates multiple detection modules and community detection techniques to identify airdrop hunters. This framework enhances detection accuracy and efficiency, significantly outperforming existing graph neural network methods.
- We developed three custom-designed feature extraction components, focusing on trading patterns, trading time patterns, and trading characteristics, which enables effective utilization of critical features for detecting airdrop hunters, emphasizing model interpretability.
- Our experimental results demonstrate that our model achieves state-of-the-art performance, outperforming existing graph inference models across various metrics in identifying airdrop hunters, providing a new and effective framework for addressing airdrop issues and other on-chain anomaly detection tasks in the Web3 community.

2. Background and related work

2.1. Web3 and decentralized applications

In recent years, Web3 (also known as Web 3.0) technology has demonstrated tremendous potential across various fields. Web3 aims to use decentralized and serverless architecture to create a user-centric internet[14]. Unlike the current internet (Web 2.0), which is primarily controlled by centralized entities such as large tech companies, Web3 leverages blockchain technology and decentralized protocols to empower users and enhance transparency, security, and privacy[15]. Web3 has become an all-encompassing term representing a vision for a new and improved internet. At its core, Web3 seeks to return power to users in the form of ownership through blockchain, cryptocurrencies, and non-fungible tokens (NFTs)[16]. It has shown potential applications in various aspects of society, including finance, fostering the emergence of many new design mechanisms[17, 18, 19, 20, 21]. For instance, the development of Decentralized Autonomous Organizations (DAOs) showcases how these organizations can manage and enforce rules through smart contracts, eliminating intermediaries in traditional organizational structures and enabling more democratic and transparent governance[22].

The emergence of Ethereum aims to address some of the limitations and challenges of Bitcoin. It provides developers with a tightly integrated end-to-end system for building software in mainstream computing paradigms: a trusted object messaging computing framework with smart contracts, which are scripts that can run synchronously across multiple nodes on a distributed ledger without external trusted institutions[23]. On Ethereum, we can classify transactions into two types: external transactions and internal transactions. External transactions are initiated by user addresses and can be direct transactions between user addresses or function calls to smart contracts. Internal transactions are initiated by smart contracts, which can transfer among themselves or send tokens to users, such as in the case of airdrop contracts.

Ethereum and smart contracts provide a powerful platform for the popularity of decentralized applications (DApps). Currently, Ethereum is the distributed ledger technology (DLT) with the largest DApp market[24]. Decentralized applications (DApps) are applications that can operate autonomously and typically execute on a decentralized computing blockchain system using smart contracts[18]. Like traditional applications, DApps offer some functions or utilities to their users. However, unlike traditional applications, DApps execute without human intervention and do not belong to any single entity. Instead, DApps distribute tokens that represent ownership. These tokens are allocated to system users according to algorithmic rules, diluting the ownership and control of the DApp[25]. In the absence of any single entity controlling the system, the application becomes decentralized. Tokens play multiple roles within the DApps ecosystem, covering aspects such as functionality, governance, liquidity, user participation, value transfer, and smart contract execution[26]. For instance, many DApps use governance tokens to achieve decentralized governance, allowing token holders to vote on the future direction of the project or use tokens for payments and providing liquidity.

2.2. Airdrop and airdrop hunter

In the design of tokenomics, DApp teams must consider various key aspects that significantly influence a project's success. Tokenomics refers to the entire economic structure of a project's token, encompassing its supply, distribution, incentive mechanisms, and utility. A meticulously crafted tokenomics strategy not only facilitates a seamless project launch and operation but also stimulates community engagement and long-term development[27]. A well-defined tokenomics structure should align the interests of the community with the project's goals, balancing token supply to avoid inflation while ensuring sufficient liquidity.

Airdrops are a popular incentive mechanism used by DApp teams to promote their projects and attract users. By distributing a certain amount of tokens for free to specific groups or the broader public, projects can enhance their visibility and user engagement[28]. Airdrops effectively onboard new users, reward early adopters and active community members and foster a robust ecosystem around the DApp[4]. The mechanism of airdrops can vary, ranging from requiring users to complete specific tasks such as social media promotions, joining community channels, participating in beta testing, or simply holding a certain amount of cryptocurrency[29]. These tasks not only help spread awareness but also ensure that recipients are somewhat invested in the project[28].

Airdrops, intended to distribute tokens widely, are susceptible to Sybil attacks. In such attacks, malicious actors create multiple wallets to receive more tokens than they are due. We refer to these actors as *airdrop hunters*, who fabricate fake accounts and manipulate activities to unfairly gain more airdropped tokens[8]. In our paper, we use the term "Airdrop Hunter" to denote what is commonly known in the Web3 community as an "Airdrop Sybil." For example, if a protocol requires specific transactions or interactions with a smart contract for airdrop eligibility, airdrop hunters can perform these actions from multiple wallets to claim excess tokens[30]. These hunters often employ scripts or bots to generate numerous fake accounts on the target platform. These scripts can automate the creation of random usernames and emails, fill out registration forms, and even use specialized services for captcha verification[4].

The presence of airdrop hunters within the Web3 ecosystem introduces several detrimental effects that jeopardize the foundational principles of blockchain technology. Decentralization, a cornerstone of blockchain systems, is compromised when airdrop hunters disproportionately acquire tokens. Zhang *et al.* [31] highlighted how Sybil attacks, such as those orchestrated by these hunters, result in token concentration among a small number of malicious actors, thereby undermining the decentralization goals inherent in tokenomics strategies. Airdrop hunters also distort transaction volumes by engaging in repetitive and meaningless activities, such as wash trading, back-and-forth trading, and cyclic transactions. These practices impose significant computational and economic burdens on blockchain networks. For instance, Liu *et al.* [32] demonstrated that inefficient transaction patterns inflate gas fees and prolong transaction times, creating barriers for legitimate users. Moreover, financial inclusion, a key tenet of blockchain systems aimed at democratizing access to resources and opportunities, is adversely affected by the activities of airdrop hunters. Exploiting their technical expertise and resources, these actors monopolize rewards, effectively excluding smaller participants who lack the capacity to compete. Cong *et al.* [33] argued that such exclusionary practices hinder the democratization of blockchain ecosystems and contribute to systemic inequities.

Many significant airdrops have exposed gaps in anti-Sybil measures. For instance, Aptos lacked effective anti-Sybil rules during its airdrop, resulting in hunters acquiring many $APT tokens, which they later sold in large quantities, causing market disruptions. Researchers found that Sybil addresses accounted for 40% of the tokens deposited into exchanges[34]. Similarly, in the case of Blur, one of the largest NFT marketplaces, analysts revealed that 50% of Blur's NFT trading volume came from fewer than 300 wallets, while 1% of "whales" held 84% of the total value locked in Blur's bid pools [8].

2.3. Airdrop Hunter Detection

To combat the prevalence of airdrop hunting, some cryptocurrency projects have developed anti-Sybil technologies. Notably, the Arbitrum Foundation implemented rules to determine which addresses were eligible for ARB token airdrops[35]. These rules included:

- Limited Operations: Addresses with few operations within a 48-hour period.
- Balance Threshold: Addresses with a balance of less than 0.005 ETH at the time of the snapshot.
- Prior Identification: Addresses identified as Sybil in the Hop Protocol airdrop.

It is speculated that users verifying multiple wallets from a single IP address on the airdrop's official website were disqualified, although this has not been officially confirmed[36]. Many believe that implementing KYC (Know Your Customer) systems could help solve this problem. However, permissionless and anonymous participation are core values of Web3[36]. While identity verification can prevent the creation of Sybil accounts, it also increases user friction and compromises privacy. Additionally, KYC technology is susceptible to forgery, identity theft, and phishing scams. Preventing KYC fraud often requires new efforts and strategies[37].

Recent research highlights various advanced techniques for detecting fraudulent activities in cryptocurrency networks, particularly focusing on identifying complex transaction patterns and wash trading, which contribute to detecting airdrop hunters. Victor and Weintraud [10] concentrate on detecting and quantifying wash trading on decentralized exchanges. They achieve this by creating token transaction graphs and employing trade volume matching along with strongly connected components (SCCs) heuristic methods to identify and quantify wash trading activities. Similarly, Victor [38]proposes heuristic methods for clustering Ethereum addresses, utilizing deposit address reuse, multiple airdrop participations, and self-authorization to group addresses on the Ethereum blockchain, thereby identifying entities controlling multiple addresses. In the realm of transaction network detection, Wu *et al.* [39]explore the use of hybrid pattern detection for mixing services in Bitcoin transactions. They provide a feature-based network analysis framework to identify the statistical properties of mixing services from three levels—network level, account level, and transaction level—and propose the concept of attributed temporal heterogeneous (ATH) motifs to better characterize the transaction patterns of different types of addresses.

In the task of detecting airdrop hunters, Liu and Zhu [7] proposed an innovative mechanism to combat Sybil attacks during airdrop events. Their approach combines comprehensive address behavior analysis and pattern recognition to filter out malicious actors, effectively distinguishing genuine participants from fraudulent ones. Zhou *et al.* [8]introduced ARTEMIS, a graph neural network system designed to identify airdrop hunters in the NFT market. ARTEMIS analyzes NFT transaction patterns and employs multimodal deep learning techniques to extract insights from NFT metadata, significantly outperforming existing methods in detecting airdrop hunters. Meanwhile, with the rapid advancement of large language models (LLMs), their application in blockchain technology and airdrop hunter detection has garnered increasing attention. ARTEMIS integrates LLMs to analyze metadata and unstructured textual information—such as NFT descriptions and user comments—enabling semantic-level identification of potential malicious behaviors. Furthermore, a research project from Berkeley explores the integration of LLMs with blockchain systems to enhance transaction analysis, anomaly detection, and smart contract auditing [40]. This project highlights the potential of LLMs in capturing complex semantic relationships and transaction patterns, laying a foundation for more accurate detection of airdrop hunters and transaction analysis.

In the industry, Trusta Labs' anti-Sybil detection mechanism is the mainstream approach for identifying airdrop hunters[41]. Their two-stage AI-ML framework uses clustering algorithms to identify Sybil communities. The first stage employs community detection algorithms like Louvain and K-Core to analyze the asset transfer graph (ATG) and detect tightly connected, suspicious Sybil groups. The second stage involves calculating user profiles and activities for each address, refining clusters with K-means to reduce false positives from the first stage. This method is consistent with the primary approach used across different airdrop projects. Overall, most transaction graph-driven algorithms face two major drawbacks: high computational costs and suboptimal accuracy. Therefore, we introduce a community-enhanced fusion framework to integrate different detection modules, improving the overall performance of the task.

3. Dataset

This section introduces the dataset utilized in the study, derived from Blur's NFT transaction records spanning October 2022 to April 2023. The dataset comprises 2,453,280 NFT transactions involving 203,370 unique addresses, with 4,808 labeled as airdrop hunters. Typical behaviors such as self-trading, back-and-forth trading, and short transaction cycles are analyzed to identify patterns indicative of airdrop hunters. These patterns form the basis for feature extraction and model training, enabling effective differentiation between hunters and legitimate users.

3.1. Airdrop hunters in blur

After defining the detection of airdrop hunters as our initial objective, we used transaction records collected from the Blur (https://blur.io/) project as our data source for the experiment. Blur is a decentralized NFT marketplace and aggregator platform. With the explosive growth of the NFT market, it has attracted a diverse range of participants. However, existing platforms often overlook the needs of professional traders. Blur emerged in early 2022, aiming to address this gap by offering a feature-rich, zero-fee marketplace specifically designed for experienced traders[42]. It provides real-time marketplace aggregation, allowing users to sweep and list across NFT marketplaces, snipe reveals, and manage their portfolios efficiently. Blur is also the first marketplace to introduce incentivized royalties [42] (Blur's airdrop system rewards users who list NFTs with royalties, encouraging them to support creators within the ecosystem).

In the NFT domain, OpenSea has long dominated the field due to its first-mover advantage. Blur needs to gain recognition and acquisition in an already competitive market, with rivals such as OpenSea, GEM, and LooksRare. To achieve this goal, the team leveraged point-based incentives in one of the most diverse ways ever seen in the Web3 space. Blur strategically distributed its native token $BLUR through a series of airdrops, which not only attracted users but also fostered a sense of ownership and community. During Blur's second token airdrop on February 15, 2023, over 300 million tokens (more than 10% of the total supply) were distributed, attracting 115,834 users, surpassing OpenSea[43].

On-chain analysis revealed that the sudden increase in Blur's NFT trading volume was primarily driven by whales (traders holding large amounts of specific assets) who continuously bought and sold NFTs through the market's bid pools to "mine" the next airdrop's token rewards[44]. The airdrop incentive mechanism significantly facilitated the spread of wash trading. Post-airdrop on-chain data analysis showed that 50% of Blur's NFT trading volume came from fewer than 300 wallets, and 1% of "whales" held 84% of the total value locked in Blur's bid pools[45]. This phenomenon negatively impacts the community's growth, potentially causing severe token price manipulation and harming the project's overall health.

3.2. Data description

The dataset used in this study comes from the Blur airdrop hunter dataset constructed by Zhou *et al.*[8]. This dataset was compiled using the Etherscan API (https://etherscan.io/) to collect all NFT transaction data and airdrop records related to Blur from October 19, 2022, to April 1, 2023. For traded NFTs, comprehensive metadata was collected, including NFT images, descriptions, and attributes. Clustering techniques were used to process transaction information, and subsequent labeling was employed to meticulously compare airdrop records to identify airdrop hunters. The dataset includes 2,453,280 NFT transactions involving 203,370 unique user addresses. The official Blur address (0xf2d15c0a89428c9251d71a0e29b39ff1e86bce25) conducted a total of 123,815 airdrops, of which 4,808 (approximately 4%) were labeled as airdrop hunters, with the rest being regular traders. Each transaction's timestamp, type (buy or sell), value (based on ETH tokens), sending/receiving addresses, NFT collection, and relevant NFT ID were logged. Historical transaction and smart contract interaction records for each wallet were also compiled. For each transaction, the dataset captures critical details such as:

- Timestamps: The exact time each transaction occurred.
- Transaction Types: Whether the transaction involved buying or selling.
- Transaction Values: Measured in ETH tokens.
- Addresses: Including sender and receiver wallet addresses.
- NFT Information: Metadata including NFT collection and unique identifiers.

Although the Blur dataset has been very supportive regarding the findings in this study, there are a number of limitations that affect generalizing these findings. More specifically, this dataset is based mainly on Blur's airdrop campaigns; hence, it does not reflect all the strategies of airdrop hunters on other platforms. For instance, the variations in the airdrop rules of the platforms, such as interaction requirements or user screening mechanisms, may evoke differences in behavioral pattern variation. Also, the Blur dataset primarily targets the NFT market, and thus it cannot feature the critical variance of the wider blockchain domain, including DeFi or token-oriented airdrops.

Thus, our model places an emphasis on the identification of specific anomalous traits that are innate in the transactional patterns between addresses for improving the applicability and generalizability of the findings. This kind of method is particularly aptly positioned for challenges brought about by differences in platforms and ecosystems. In this regard, it does the extraction of attributes that are independent of platforms from transactions, including transaction frequency, temporal pattern, and structural features of transaction networks. Our model only tries to capture universal behavioral patterns of airdrop hunters rather than relying on the rules or environment of a particular platform. This methodology enhances the performance of this model on the Blur dataset and provides a theoretical basis for extending its applicability to other platforms.

3.3. Typical hunter example

To uncover the behavioral patterns of airdrop hunters, we conducted a comprehensive analysis of the Blur dataset by constructing transaction graphs and employing advanced graph analysis techniques. Specifically, we modeled transaction records as directed graphs, where nodes represent user addresses, and edges denote transactions between them. From the dataset, we manually identified three representative interaction patterns of airdrop hunters, which are illustrated in Figure 1. All addresses and transaction records used in the figure are sourced directly from our Blur dataset.

The Figure 1 illustrates three typical transaction patterns observed in our dataset to identify potential airdrop hunters. The first pattern shows star trading pattern (0xe5e...580) where a central node connects to multiple peripheral nodes with numerous inbound and outbound transactions, indicating a single entity interacting with many others to simulate

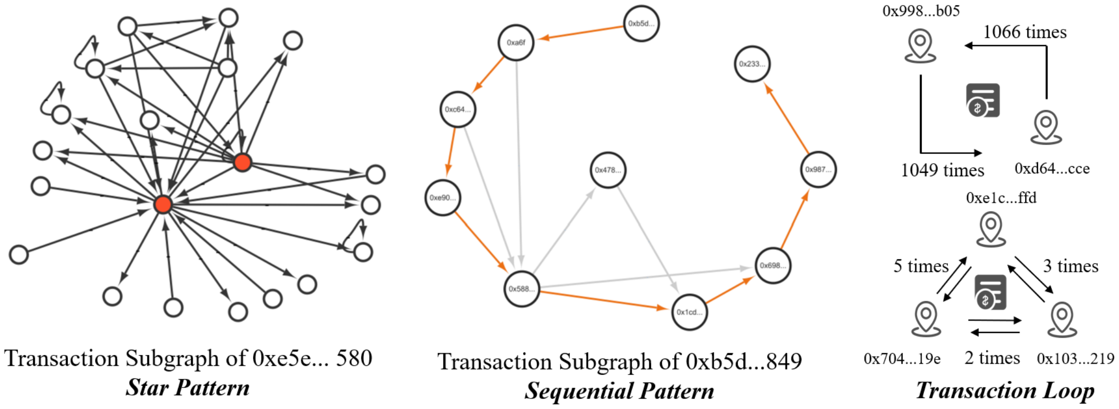

Transaction Subgraph of 0xe5e... 580
Star Pattern

Transaction Subgraph of 0xb5d...849
Sequential Pattern

Transaction Loop

Figure 1. Airdrop hunter transaction patterns: insights extracted from the blur dataset.

high activity. The second pattern is a sequential pattern subgraph (0xb5d...849), depicting a circular chain of transactions among several nodes, forming a loop that suggests coordinated activity to create the appearance of genuine transactions. The third pattern is a transaction loop example showing repetitive transactions between specific addresses, highlighting frequent exchanges that indicate manipulative behavior, such as the loop involving nodes 0x998...b05 and 0xd64...cce with over 1000 transactions in both directions, and another loop among nodes 0x704...19e, 0x103...219, and 0xe1c...ffd, with numerous transactions exchanged among them. These patterns are indicative of behaviors often associated with airdrop hunters, who manipulate transactions to appear more active in the network and thus qualify for airdrop rewards. These patterns are characteristic of behaviors commonly linked to airdrop hunters, who strategically manipulate transactions to appear more active within the network and qualify for airdrop rewards. Accurately identifying the key features of these intricate transaction patterns is essential for our model.

4. Methodology

4.1. Method overview

To achieve superior performance in detecting airdrop hunters, we have introduced ARTEMIX. The methodology of ARTEMIX is primarily divided into two phases as shown in Figure 2. The first phase involves constructing three main classifiers based on the different patterns of airdrop hunters, including nodes participating in typical hunter trading patterns, nodes exhibiting typical trading time patterns, and nodes with characteristic trading behaviors. Key node classification features were extracted from the transaction graph. This process involved embedding edge weights using cosine similarity based on 11 fundamental node features and applying the Louvain algorithm to obtain community detection results. In the second phase, these features and classifier results are integrated into a community learning model utilizing a boosting algorithm for combined training, resulting in the final inference model. This approach effectively highlights essential features, reduces model training and inference costs, and maintains scalability. In this section, we will elucidate our design rationale and introduce the various modules of ARTEMIX.

4.2. Typical hunter trading patterns

Airdrop hunters frequently employ customized trading strategies to curate accounts that fulfill more extensive criteria for airdrops in the Web3 ecosystem. Within the vast network of trading addresses and complex trading records, several distinctive patterns can be identified. These

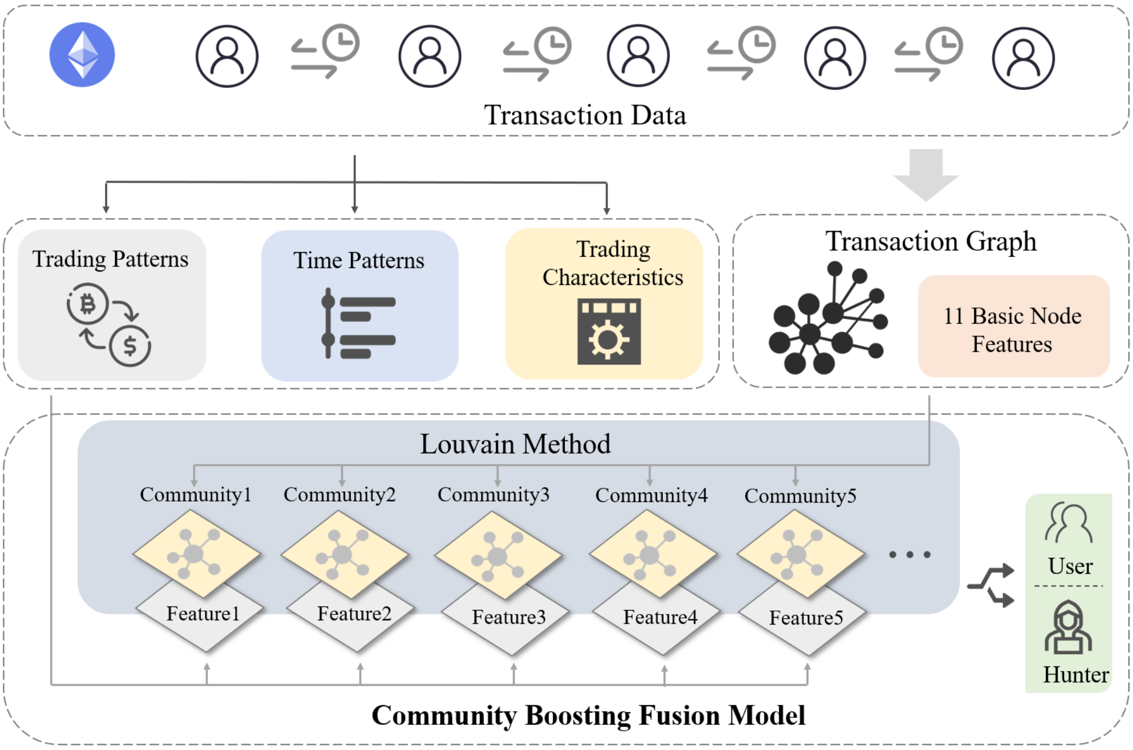

Figure 2. The framework of ARTEMIX.

patterns are effective in flagging suspicious airdrop hunter addresses in the trading ecosystem. These behavioral traits, embedded in the complex transaction graph, are often subtle yet discernible. Despite their relative simplicity, these patterns tend to emerge consistently among airdrop hunters, whether consciously or unconsciously, over extended periods of interaction with multiple accounts. Subsequent analyses reveal that these filters are notably effective in identifying genuine airdrop hunters.

4.2.1. Self-trading

Self-trading involves an account engaging in transactions with itself. This tactic is often employed to create artificial trading volume, manipulate market perceptions, or simulate activity to meet airdrop eligibility criteria. Although self-trading is generally not the primary strategy for airdrop hunters, it is prevalent in a significant number of transactions among certain addresses. Due to the signature mechanism of blockchain transfers, transactions between addresses undergo multiple confirmations by entities, ensuring that each transaction has a clear and deliberate purpose. Self-trading occurs on most marketplaces as a method of wash trading, which is commonly associated with airdrop hunters. Wash trading involves buying and selling the same asset to create misleading market activity and can be used to inflate the perceived demand for a token or meet specific trading volume thresholds necessary for airdrop eligibility. By examining transaction histories, we can identify such nodes where the buying and selling addresses are identical. This identification process is crucial for distinguishing legitimate trading activities from manipulative practices.

In our dataset of 203,101 addresses, we detected 1303 addresses that have engaged in self-trading transactions. By comparing these addresses with the airdrop hunter labels in our dataset, we found that 694 out of these 1303 addresses are identified as airdrop hunters. This indicates that using this single filter allows us to successfully identify approximately 14.4% of airdrop hunters from the 4808 airdrop hunters present in our dataset. This identification indicates that self-trading transactions can serve as a significant indicator for detecting airdrop hunters, but it also is not the sole indicator.

4.2.2. Back-and-forth trading

This behavior is characterized by the repetitive trading of the same or different NFTs between two addresses. This pattern directly reflects the traits of airdrop hunters: increasing community engagement and on-chain transaction volumes across multiple accounts through numerous meaningless transactions. The primary objective is to incur only the gas fees for these transactions in order to receive airdrop rewards. Although this method is not the main interaction tool for hunters, we can still identify many addresses that fit this pattern in on-chain records. However, this has to account for the fact that some legitimate users of the community may use multiple addresses for the purpose of securing their assets or facilitating their transactions. Therefore, we tweak our filters to work effectively by setting a reasonable threshold. This is in the context of the minimum number of transactions that should be between two identified accounts. Derived from the ROC curve in Figure 3, we can identify the optimized threshold of 4 through choosing the peak Youden's index. Evaluation results show this approach has ultimately filtered out 3374 addresses, and among them, it can correctly identify 1977 airdrop hunter addresses, taking up to 58.6% of all airdrop hunter addresses. This indicates that this behavioral characteristic provides strong evidence for identifying airdrop hunters.

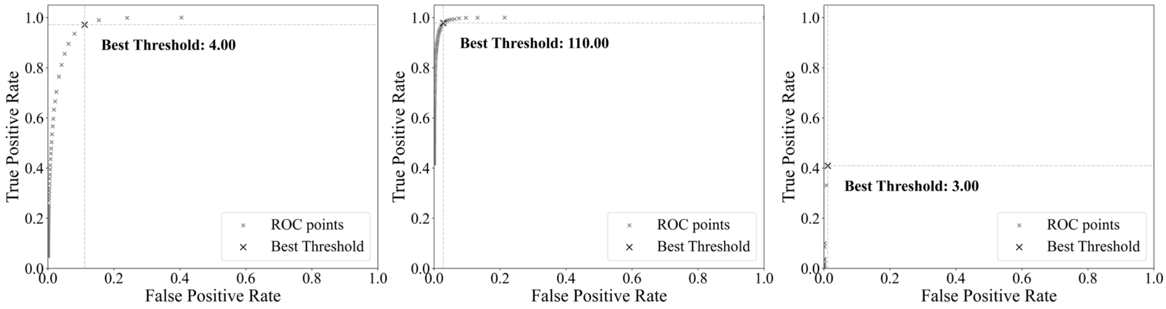

a. Back-and-Forth Trading (ROC) Curve b. Short Cycles Trading (ROC) Curve c. NFT Cycle Trading (ROC) Curve

Figure 3. ROC curve for three typical trading patterns.

4.2.3. Short cycles

Detecting suspicious addresses involved in short cycles among a few nodes is crucial for uncovering potential airdrop hunting activities. These short cycles involve a small group of addresses repeatedly transacting among themselves, creating the illusion of increased activity and engagement on the blockchain. This behavior is often used to exploit airdrop distribution mechanisms that reward higher levels of interaction and transaction volume.

To detect such patterns, we construct a directed graph where nodes represent addresses and edges represent transactions between them. We then implement a cycle detection algorithm focusing on cycles of lengths 2 to 3, as these are indicative of suspicious activity aimed at mimicking legitimate interactions. By setting a threshold for the number of cycles an address must participate in to be flagged as suspicious (Table 3), we can filter out legitimate users who might occasionally participate in short cycles for valid reasons.

During the evaluation, this filter demonstrated remarkable performance. Out of a total of 3439 addresses filtered, 92.56% were marked as airdrop hunters in the dataset. This provides strong evidence for identifying the trading habits of some airdrop hunter manipulators. This method effectively uncovers airdrop hunters engaging in repetitive transactions among a few addresses to exploit the system. By applying the cycle detection algorithm to the transaction graph and filtering addresses based on the established threshold, we identify those that participate in an unusually high number of short cycles. Evaluating and validating the results against known airdrop hunter addresses ensures the approach minimizes false positives and accurately identifies malicious behavior.

4.2.4. Holding the same NFT multiple times

This trait is marked by an account repeatedly acquiring and selling the same NFT. Such behavior may be used to inflate the perceived value of the NFT or to simulate trading activity to meet airdrop eligibility criteria that require holding specific assets. We can spot these patterns by tracking the ownership history of NFTs and identifying accounts that repeatedly interact with the same tokens.

An NFT's Unique Identifier (UI) is a distinctive code or label used to distinguish it from all other NFTs. This identifier is a unique string of numbers or characters that uniquely identifies and differentiates each NFT. On the blockchain, each NFT has a UI, ensuring that even if NFTs share the same name, description, or appearance, they remain distinct and independent. When the exact same NFT repeatedly appears in the transaction history of a single address, this address is highly likely to exhibit the characteristics of an airdrop hunter. However, it is essential to note that some collectors might trade similar NFTs frequently, and certain top-tier NFTs naturally have a low circulation rate. Therefore, this indicator alone cannot conclusively identify an airdrop hunter.

To refine this filtering process, we set a threshold: if an identical NFT appears more than the threshold number of times in the transaction history of the same address, it will be flagged for further analysis. After testing, we determined a best threshold of 3 (Table 3). Consequently, we identified 782 addresses that met this criterion. Among these, 415 addresses were already marked as airdrop hunters in our dataset, validating the effectiveness of this filter.By monitoring the unique identifiers of NFTs and setting appropriate thresholds, we can effectively identify patterns indicative of airdrop hunting. This method, while not foolproof, provides a robust tool for filtering suspicious activity in the NFT trading space.

4.3. Transaction time patterns

In the second module of ARTEMIX feature extraction, we focus on uncovering behavioral characteristics hidden within the timestamps of transaction records across different addresses. Most professional airdrop hunters categorize and isolate their addresses to maximize their gains through mixed strategies. A simple example is that they typically manually interact with key addresses while using automated scripts for lower-quality duplicate accounts. These scripts often operate with fixed daily interaction strategies. By analyzing transaction time patterns, we can effectively identify such scripts. Additionally, manually operated addresses usually display specific time strategy habits over a longer timeline, requiring us to monitor transaction habits across multiple time dimensions.

To simulate the behavior of airdrop hunters, we observe two main characteristics in their transaction time patterns: (1)Abnormal Trading Behavior During Specific Time Periods: Airdrop hunters often show unusual trading activity during particular time frames; (2)Regular Daily or Weekly Trading Patterns: Airdrop hunter addresses typically exhibit consistent daily or weekly trading patterns. By focusing on timestamp patterns, we aim to design more effective detection mechanisms to identify and mitigate suspicious activities. In the following sections, we introduce methods for identifying suspicious addresses based on these two characteristics and conduct a preliminary evaluation of their effectiveness.

As shown in Figure 4, in the selected typical transaction frequency heatmaps, normal users tend to conduct more transactions during certain characteristic time periods. This behavior reveals the users' trading tendencies, and there are significant differences in these trading

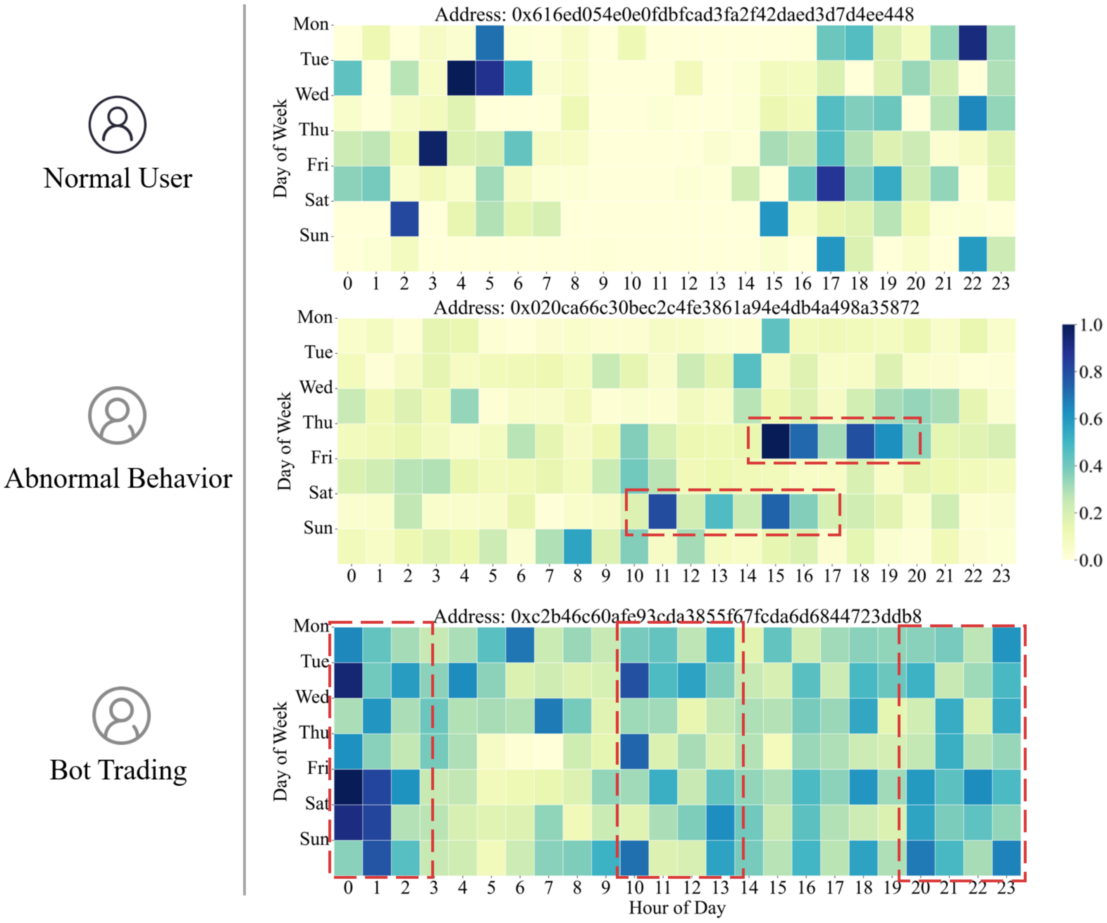

Figure 4. Heatmap analysis of different transaction time patterns.

patterns from day to day, demonstrating the inherent randomness of their trading activities. Typically, these high-frequency trading periods align with the users' daily routines, which is characteristic of genuine individuals interacting normally in the market. In contrast, entities known as Airdrop Hunters, who often control numerous virtual addresses, exhibit different trading behaviors. As illustrated in Figure 1, their transactions usually form one or more clusters, typically resulting in a central node that connects all sub-nodes. These central nodes are responsible for receiving and distributing the assets of the entire transaction cluster. Unlike normal users, Airdrop Hunters do not engage in frequent interactions; instead, they conduct large-scale transactions at specific predetermined times. A typical example of this transaction pattern is depicted in the second heatmap of Figure 4, which we term as Abnormal Trading Behavior. The address in question conducts a high volume of transactions concentrated on Thursday evenings and Saturday afternoons, with significantly lower transaction volumes at other times. Over an extended period, these characteristics deviate significantly from the trading behavior tendencies of normal users.

To identify addresses with abnormal trading behavior in our dataset, we employed the Isolation Forest algorithm [46], a classic technique using binary trees in anomaly detection. The algorithm has a linear time complexity and a low memory requirement, making it well-suited for high-volume data. Unlike decision tree algorithms, Isolation Forest uses only the path-length measure or approximation to generate the anomaly score, without relying on leaf node statistics on class distribution or target value. In our detection process, we initially extracted the timestamps of transaction records for each address. We then segmented these timestamps into four distinct time windows: 1 hour, 6 hours, 1 day, and 1 week. For each address, we calculated the number of transactions within each time window and combined these as time-based features. Next, we applied the Isolation Forest algorithm to the time features across the entire dataset, opting to use the default parameter settings after a comparative analysis. The detector identified 1,385 anomalous addresses during the dataset evaluation, with 1,252 of these being labeled as airdrop hunters. The precision reached an impressive 90.4%, underscoring the effectiveness of this feature detection approach in accurately identifying airdrop hunters and detecting anomalous transaction behaviors.

Figure 4's third heatmap illustrates the second typical trading time pattern—Regular Daily or Weekly Trading Patterns. The address shows highly regular trading activities during three specific time periods each day of the week: 0-2h, 10-13h, and 20-23h. Additionally, the address displays a similar trading frequency outside these concentrated time periods. This behavior is characteristic of automated airdrop trading scripts. The differences between automated scripts and real users mainly lie in two aspects: 1. Automated scripts follow identical daily or weekly trading patterns with little variation. 2. Automated scripts trade at uniformly distributed time intervals, whereas real users, due to their regular schedules, typically have sparse trading periods. Thus, trading patterns that adhere to such regularity can be easily identified as airdrop hunters. To identify periodic trading patterns indicative of automated behavior, we applied the Fast Fourier Transform (FFT) to transaction time series data. The FFT is an algorithm that computes the Discrete Fourier Transform (DFT) or its inverse (IDFT), transforming a time series from the time domain to the frequency domain to identify periodic components within the data[47]. The DFT of a sequence $x[n]$ of length N is given by:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-i\frac{2\pi}{N}kn} \tag{1}$$

where $X[k]$ represents the frequency components, $x[n]$ is the time-domain data, N is the total number of data points, and i is the imaginary unit. We began by loading and preprocessing transaction data, converting timestamps to datetime format, and aggregating them into hourly intervals to form a consistent time series for each address. The hourly transaction counts were centered by subtracting the mean:

$$T'[n] = T[n] - \mu \tag{2}$$

where μ is the mean of the transaction counts $T[n]$. The FFT was then applied to yield frequency components xf and their corresponding magnitudes yf:

$$|Y[k]| = \sqrt{\Re(Y[k])^2 + \Im(Y[k])^2} \tag{3}$$

where $\Re(Y[k])$ and $\Im(Y[k])$ are the real and imaginary parts of the FFT result, respectively. Significant periodic patterns were identified by applying a threshold of 20 to the magnitudes, with peaks exceeding this threshold indicating notable periodicity:

$$\text{Significant Peaks} = \{k \mid |Y[k]| > \text{threshold}\} \tag{4}$$

Addresses exhibiting such peaks were flagged for periodic trading patterns. This method effectively distinguishes the uniform, regular trading patterns of automated scripts from the irregular patterns typical of human traders, aiding in the detection of airdrop hunters and enhancing the integrity of the trading environment.

4.4. Characteristic trading features

In this subsection, we delineate the distinctive trading characteristics of airdrop hunters in comparison to regular users, as illustrated in Figure 5. Airdrop hunters exhibit particular trading strategies, such as high asset turnover rates and frequent repurchase behaviors. Their wallets typically interact with a substantial number of different NFT projects over extended periods. Additionally, airdrop hunters' addresses are often highly active, characterized by frequent transactions, a larger number of interaction addresses, and a tendency to use regular small-amount transactions to maximize the number of airdrops received.

To ensure that we do not overlook the fundamental characteristics of these airdrop hunter addresses, we have extracted and quantified four representative features. The visual analysis of these features is conducted through a series of density plots, providing a detailed comparison across various metrics:

(a) Collections of NFTs Purchased: The density plot reveals that airdrop hunters (blue) tend to purchase a significantly higher number of NFT collections compared to regular users (green). Airdrop hunters exhibit a wider spread, with many purchasing up to 300 collections, whereas regular users show a higher density at the lower end, purchasing fewer collections.

(b) Average NFT Holding Time: The holding time for NFTs indicates a clear distinction between the two groups. Airdrop hunters generally hold NFTs for a shorter duration, with the density decreasing sharply beyond 25 days. In contrast, normal users display a wider range of holding times, with a substantial portion holding NFTs for as long as 100-200 days, suggesting a more long-term investment behavior.

(c) Unique Address Interactions: The interaction with unique addresses highlights that airdrop hunters engage with a significantly larger number of unique addresses. The density plot shows a broad distribution up to 10,000 unique addresses for airdrop hunters. In contrast, regular users exhibit a higher density at lower interaction counts, engaging with fewer unique addresses overall.

(d) Number of Active Days: Airdrop hunters and regular users also differ in terms of

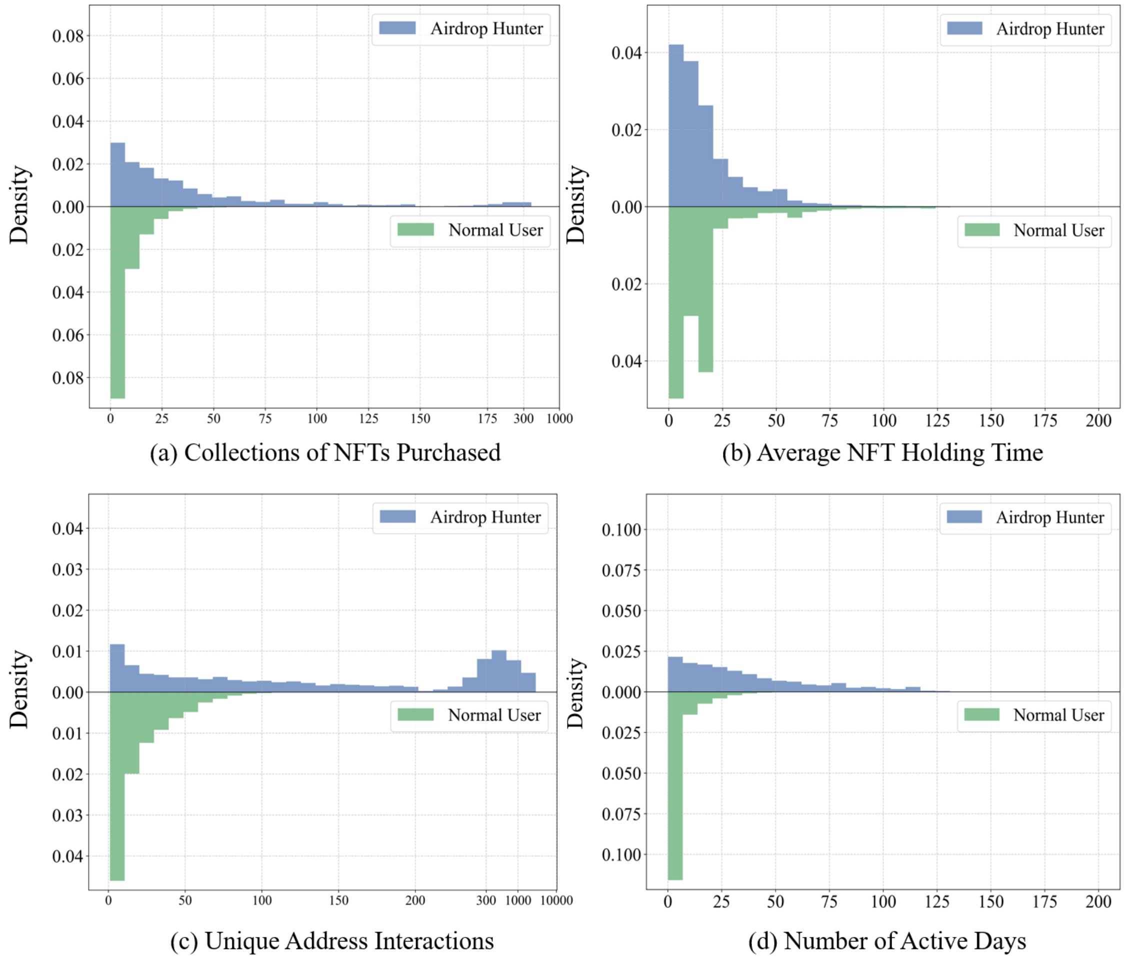

(a) Collections of NFTs Purchased

(b) Average NFT Holding Time

(c) Unique Address Interactions

(d) Number of Active Days

Figure 5. Comparison of airdrop hunters and regular users.

their activity levels. Airdrop hunters have a density that spans a broader range of active days, indicating sporadic but numerous engagements. Regular users tend to have a higher density concentrated at fewer active days, suggesting more consistent but less frequent activity.

To enhance the accuracy and robustness of detecting airdrop hunters, we employed a Random Forest classifier [48] to integrate characteristic trading features. Random forests are an ensemble of tree predictors, where each tree depends on the values of a random vector sampled independently and with the same distribution for all trees in the forest. The Random Forest model was trained using features extracted from our dataset, each contributing to distinguishing airdrop hunters from regular users by capturing unique trading behaviors. The model's reliability and generalizability were validated using cross-validation techniques.

In addition to the previously discussed trading characteristics, we introduce PageRank as an independent feature to evaluate the influence and importance of wallet addresses within the transaction graph of the NFT ecosystem. PageRank, a widely used algorithm for ranking nodes in a network, assigns higher values to addresses that are more central or influential based on their connectivity and transaction activity[49].

By analyzing the PageRank distribution across airdrop hunters and regular users, we gain further insights into the strategic positioning and behaviors of these groups. Airdrop hunters, who typically engage with a larger number of addresses and execute frequent transactions, tend to have higher PageRank values, highlighting their central role within the trading network. This feature, although not integrated into the Random Forest classifier, serves as a complementary measure to distinguish between airdrop hunters and regular users, providing a more comprehensive understanding of their trading patterns and network influence.

4.5. Community boosting fusion model

Common graph methods and heuristic algorithms often underperform regarding training cost and accuracy. Additionally, they lack model portability and flexibility for diverse airdrop initiatives. To address these issues, we propose a novel Community Boosting Fusion Model. This model integrates node feature extraction with classifier detection results using a boosting algorithm to enhance the detection of airdrop hunters. Our method is easily portable, straightforward to optimize, and suitable for various airdrop models.

4.5.1. Model input

The Community Boosted Fusion Model we developed integrates various essential elements derived from blockchain transaction data. Initially, we devised four filters to detect and exclude common transaction patterns associated with airdrop hunters by picking the most effective thresholds. These filters produce four separate lists of suspicious addresses, detecting potential malicious actors by analyzing identified suspicious transaction patterns. In addition, we have created two specialized detectors for identifying anomalous transaction timing patterns. These detectors examine the duration and arrangement of transaction timings, generating two lists of questionable addresses. Their main emphasis is on atypical transaction timing patterns, such as the clustering of high-frequency transactions within defined time intervals and the presence of extremely regular trading habits like those of bots. In order to improve the accuracy of detection, we employed a random forest model to combine four transaction parameters related to airdrop hunting. This resulted in a suspicion score being assigned to each address. This score quantifies the likelihood of each account engaging in airdrop hunting during the full transaction cycle. The random forest model integrates various factors, including transaction amount volatility, transaction counterparties, and address activity level, to generate a comprehensive suspicion assessment. By combining the information from the four transaction pattern filters, two transaction timing pattern detectors, and the random forest fusion model, our model is able to acquire a comprehensive picture of the blockchain network.

The ultimate community fusion model enhances the precision of identifying malevolent actors, successfully detecting those who are trying to avoid discovery by employing intricate techniques.

4.5.2. Community detection

Community detection is a fundamental step in our model, aimed at identifying clusters of nodes that exhibit dense interconnections. We employ the Louvain algorithm, a well-established method for detecting communities in large networks.The Louvain method for community detection is a method to extract non-overlapping communities from large networks created by Blondel *et al.*[50] This algorithm maximizes modularity, ensuring that the detected communities are highly cohesive. By identifying these communities, we can better understand the structure of the network and isolate regions where airdrop hunters may operate.

Our Louvain algorithm is based on a similarity graph, where each node represents an address, and each edge represents the similarity weight of transaction flows between two addresses. By partitioning the similarity graph, we identify potential communities. In addition to constructing the transaction network graph of addresses, we extract basic features of nodes from transaction data to compare similarities between nodes. Specifically, we extract 11 basic transaction features for the i-th address, including total asset value, number of NFT holdings, number of transactions, etc., and denote j as the index of the features. Through data processing, we obtain the feature vector of the i-th address $\mathbf{X}_i = \{x_1, x_2, \ldots, x_{11}\}$. For better similarity calculation, we normalize the feature data such that $\sum_{j=1}^{11} x_j = 1$. After extracting the address behavior features, we choose cosine similarity to calculate the similarity between nodes, which is generally more suitable for handling high-dimensional sparse data [51].

$$d(X_1, X_2) = \frac{\sum_{i=1}^{11} X_{1,i} \cdot X_{2,i}}{\sqrt{\sum_{i=1}^{11} X_{1,i}^2} \cdot \sqrt{\sum_{i=1}^{11} X_{2,i}^2}} \tag{5}$$

The cosine similarity value $d(X_1, X_2)$ ranges from 0 to 1, with smaller values indicating higher similarity between transaction flows. By establishing a similarity graph, we use the Louvain algorithm for community detection and adjust the resolution parameter to determine the granularity of the detected communities [51]. We use modularity to measure the effectiveness of the community partitioning, where modularity is defined as a value in the range $[-0.5, 1]$, used to measure the density of links inside communities compared to links between communities [52]. A modularity closer to 1 indicates a more pronounced community/cluster division. It is generally considered that a modularity above 0.3 can produce a relatively good partitioning effect.

Through experiments, we used the Louvain algorithm with a resolution of 1.0, detecting 2973 communities among 203,101 addresses, with a modularity of 0.31. We then isolate these communities and train them separately, combining them to further improve the accuracy of detecting airdrop hunters. By combining multiple detection modules (including transaction pattern filters, transaction time pattern recognizers, and a random forest fusion model), we can conduct a more comprehensive analysis, while community classification ensures more effective identification of different categories of airdrop hunters.

4.5.3. Community boosting fusion model

The primary innovation of our approach is encapsulated in the community boosting fusion Model. This model synergizes node features with initial classifier detection results using a

1. Total Asset Value. 2. Number of NFT Holdings. 3. Number of NFT Collections Holdings. 4. Number of Transactions. 5. Number of Internal Transactions. 6. Total Number of $ETH Transfers. 7. Total Number of $ETH Withdrawals. 8. Total Number of ERC-20 Token Transfers In. 9. Total Number of ERC-20 Token. Transfers Out. 10. Number of NFT Transfers In. 11. Number of NFT Transfers Out.

boosting algorithm. Boosting, an ensemble learning technique, amalgamates multiple weak classifiers to forge a robust final model. Within this framework, boosting incrementally enhances detection accuracy by rectifying misclassifications from preceding iterations. This iterative refinement augments the model's adaptability to various airdrop projects, thereby bolstering its portability. The integration of community detection with boosting not only amplifies detection accuracy but also ensures operational efficiency and optimization simplicity.

In scenarios where community detection yields optimal classification outcomes—where addresses within the same community are controlled by a single entity and demonstrate analogous trading strategies—a boosting algorithm can proficiently assign the most suitable classifier to each community. By aggregating the final weights, we can merge all features to construct an optimal model tailored to each specific community's characteristics. The low complexity of each module within our model enables the efficient training of optimal detection models every communities within a brief timeframe, utilizing minimal computational resources. Ultimately, these models are consolidated to conduct comprehensive detection across the entire dataset, thereby enhancing both accuracy and efficiency.

In our research, Decision trees [53] were selected as the base classifiers, with the Adaptive Boosting (AdaBoost) algorithm [54] serving as the boosting model. Decision Trees are particularly advantageous due to their simplicity and strong interpretability, which clearly delineate decision paths. The AdaBoost algorithm improves the performance of weak classifiers by iteratively adjusting sample weights, culminating in a more accurate and robust final model. Training individual models for each community allows us to capture unique characteristics and optimize classifier performance based on each community's specific attributes. We build these community-specific models in order to integrate them into the full dataset with the capability of being effective on large-scale data but still maintaining the level of community precision. This ensures that our system identifies trading patterns and realizes small anomalies in keeping the growth of the Web3 community safe.

5. Experiments

5.1. *Experimental setup*

Task Description: Our experiment utilizes the Blur dataset, introduced in Section 3, which comprises 203,101 addresses, including 4,808 airdrop hunter addresses. The airdrop hunter detection model outputs a binary classification, with airdrop hunters labeled as the positive class and regular users as the negative class. For the purpose of the experiment, the dataset is split into a training set and a validation set at a 7:3 ratio. The training tasks are divided into three main parts: constructing the detection module, performing community detection, and training the final ensemble model. In the validation set, we evaluate the model using precision, recall, and F1 score for positive samples. Precision measures the proportion of true positive predictions, recall measures the proportion of actual positives correctly identified, and the F1 score is the harmonic mean of these two metrics. Consequently, we primarily use the F1 score to compare the models' overall performance. Given the task's specific nature, where the priority is to avoid penalizing regular users, we focus on maximizing precision over recall (minimizing false positives).

Baselines: In this experiment, our model was first compared with the ARTEMIS model proposed by Zhou *et al.*[8], an optimized graph neural network system specifically designed to identify airdrop hunters in NFT transactions. The hyperparameters for ARTEMIS were set according to the optimal parameters described in their work. Additionally, we compared our model with three baseline approaches: (1)Structured Data Methods: This includes methods such as SVM [55] and LightGBM [56], which rely solely on node features for classification and cannot utilize edge information; (2)Graph Random Walk-Based Methods: This category includes methods like DeepWalk [57] and Node2Vec [58], which leverage both graph structure and node features; (3)GNN-Based Methods: This includes models such as GCN [59], GraphSAGE [60], GAT [61], and GIN [62]. We meticulously optimized the hyperparameters for each baseline model, including learning rate, batch size, and other key parameters, employing grid search to attain the best performance on the dataset.

Implementation: In the model feature integration, the implementation of the first module continues to follow the previous evaluation parameter selection method. For the three typical hunter trading filters that require threshold selection, the optimal performance thresholds were determined using ROC curve comparison. Based on the test set, the thresholds for the Back-and-Forth Trading filter, short cycles trading filter, and NFT Cycle Trading filter were selected as 4125, and 3, respectively. In the second module, the contamination rate of the Isolation Forest algorithm was set to 0.01. During the spectral analysis phase, we established a fixed threshold of 20 to identify significant periodic patterns. The Random Forest algorithm used in the Characteristic Trading Features was configured with default parameters. For the community detection part, the Louvain algorithm was employed with a resolution parameter set to 1.0. The optimal modularity community division was selected from the results of ten partitions. For the final ensemble model, we conducted multiple experiments on the parameters of the DecisionTreeClassifier and AdaBoostClassifier. The results showed that the fluctuations of the three indicators were within the normal range of different training iterations (<0.1). Therefore, we used the default parameter settings for the final evaluation. All models are trained and evaluated in five rounds and the results are averaged.

5.2. *Performance evaluation*

Table 1. Comparison for Airdrop Hunters Detection.

Method	Precision	Recall	F1
SVM [55]	0.744	0.544	0.629
LightGBM [56]	0.793	0.597	0.680
DeepWalk [57]	0.567	0.501	0.496
Node2Vec [58]	0.620	0.502	0.500
GCN [59]	0.648	0.896	0.752
GraphSAGE [60]	0.562	0.934	0.701
GAT [61]	0.464	0.873	0.579
GIN [62]	0.680	0.903	0.776
ARTEMIS[8]	0.820	0.833	0.826
ARTEMIX	0.928	0.869	0.898

The performance evaluation for airdrop hunter detection methods is summarized in Table 1. The table compares various models using precision, recall, and F1-score metrics. Among the evaluated methods, ARTEMIX outperforms all others, achieving the highest precision (0.928), recall (0.869), and F1-score (0.898). This indicates that ARTEMIX has a balanced performance in terms of accurately identifying airdrop hunters while maintaining a low rate of false positives and false negatives. A high precision value, as achieved by ARTEMIX, indicates that when the model identifies an airdrop hunter, it is very likely correct. This is crucial in scenarios where the cost of mislabeling legitimate users as hunters is high, as it can lead to user dissatisfaction and damage to the platform's reputation. In ARTEMIX, the recall is also strong (0.869), indicating that the model does not overly sacrifice its ability to detect airdrop hunters while maintaining high precision. The F1-score, which is the harmonic mean of precision and recall, further confirms ARTEMIX's balanced performance, showing that it effectively manages the trade-off between these two metrics.

The ARTEMIS model shows strong performance with an F1-score of 0.826, the second highest among all methods, while traditional machine learning models like SVM and Light-

GBM lag behind with F1-scores of 0.629 and 0.680, respectively. Graph-based neural networks such as GCN, GraphSAGE, and GIN improve recall rates—especially GraphSAGE with a recall of 0.934—but at the cost of precision, leading to lower overall F1-scores compared to ARTEMIX and ARTEMIS. Node embedding techniques like DeepWalk and Node2Vec perform even worse, with F1-scores of 0.496 and 0.500, respectively, indicating that while they capture structural graph features, they fall short for airdrop hunter detection. Overall, ARTEMIX's superior performance highlights its effectiveness in accurately and reliably detecting airdrop hunters, making it a valuable tool for combating fraud in blockchain ecosystems.

5.3. Ablation study

To comprehensively evaluate the contribution of each component in our proposed ARTEMIX model, we conducted an ablation study by systematically removing key components and observing the impact on performance metrics such as precision, recall, and F1 score. The results, summarized in Table 2, show that the full ARTEMIX model achieves the highest performance across all metrics, demonstrating the synergistic effect of all components. Specifically, removing time-based features results in a drop in Recall and F1 Score, indicating that temporal information is crucial for accurately identifying airdrop hunters, as it captures essential behaviors indicative of hunter activities. Excluding feature engineering components significantly reduces Recall and F1 Score, highlighting the importance of handcrafted features in capturing nuanced behaviors of airdrop hunters that raw data alone might miss. Omitting the pattern detection modules impacts both Precision and Recall, confirming that pattern detection is vital for identifying suspicious activities by capturing complex behaviors characteristic of airdrop hunters.

Table 2. Ablation Study for Different Modules.

Method	Precision	Recall	F1
ARTEMIX (full model)	0.928	0.869	0.898
Ablation study of different modules			
w/o Trading Patterns	0.917	0.849	0.882
w/o Trading Features	0.889	0.775	0.828
w/o Time Patterns	0.913	0.862	0.887

5.4. Model Efficiency

The ARTEMIX model is specifically designed to efficiently detect airdrop hunters in NFT transactions by integrating community detection with Boosting classifiers. By leveraging this unique combination, ARTEMIX first segments the graph data into multiple smaller communities, where each community's nodes share similar structural characteristics. This community detection approach allows ARTEMIX to break down large-scale graph data into more manageable subsets, significantly reducing the overall computational complexity.

Next, ARTEMIX trains boosting classifiers separately within each community, rather than training a single complex model across the entire graph. This method not only reduces the computational resources required for training but also avoids potential bottlenecks that can arise when processing global graph data. Specifically, the model utilizes AdaBoost, which iteratively trains a series of weak classifiers, such as decision trees, and adjusts the weights of the samples in each iteration to gradually enhance accuracy. This approach is computationally efficient as it avoids the need to process complex graph structures or perform extensive matrix operations. Additionally, AdaBoost's inherent parallelization capabilities enable rapid training on multi-core processors.

Compared to other models, ARTEMIX shows a marked improvement in computational efficiency. Traditional structured data methods like SVM and LightGBM are known for their speed but struggle to capture the intricate relationships within graph structures. Meanwhile, graph-based models such as DeepWalk and Node2Vec, though effective at preserving graph topology through random walks and embeddings, can become resource-intensive when dealing with very large graphs. Graph Neural Networks (GNNs) like GCN, GraphSAGE, GAT, and GIN offer powerful tools for capturing both node and edge information but are computationally expensive, especially models like GAT that incorporate attention mechanisms. ARTEMIS, another model that employs advanced GNN techniques, also faces challenges with its computational demands, despite its effectiveness. ARTEMIX breaks down the problem into community-level tasks, leveraging AdaBoost's lightweight computational advantages to reduce costs and time consumption significantly. This makes ARTEMIX not only more efficient at handling large-scale graph data but also better suited for resource-constrained environments or applications that demand quick responses.

Table 3. Comparison of AdaBoost and GNN.

Aspect	AdaBoost	GNN
Complexity	$O(m \cdot \log n + T \cdot k \cdot n')$	$O(L \cdot m \cdot d)$
Main Bottleneck	Low-dimensional feature training + community partitioning	Global message passing + embedding propagation
Complexity Characteristics	Reduced to local scale through community partitioning	Linear growth with the number of nodes/edges and feature dimensions

As outlined in the table 3, AdaBoost's complexity, $O(m \cdot \log n + T \cdot k \cdot n')$, is significantly smaller than GNN's $O(L \cdot m \cdot d)$ due to several key differences. Firstly, AdaBoost leverages community detection to partition the graph into smaller subgraphs, dramatically reducing the scale of operations to local communities with $n' \ll n$. In contrast, GNNs operate on the entire graph, requiring global message passing that scales linearly with the total number of edges m and feature dimensions d. Secondly, AdaBoost operates on low-dimensional, hand-crafted features (k), making each iteration lightweight, while GNNs process high-dimensional embeddings (d) across multiple layers (L), significantly increasing computational demands. Moreover, AdaBoost's iterative approach with a fixed number of weak classifiers (T) further limits its complexity, whereas GNNs must repeatedly perform matrix multiplications and aggregations for each layer. These distinctions underscore why AdaBoost, when combined with community-based partitioning, is computationally more efficient than GNNs, particularly for large-scale graph tasks.

6. Discussion

6.1. Generalizability and robustness

ARTEMIX is a highly adaptable and robust framework designed to analyze not only Blur NFT marketplace data but also various Web3 protocols and platforms. Its architecture integrates advanced transaction feature extraction, temporal pattern analysis, and community detection, enabling its application in diverse scenarios such as detecting complex incentive mechanisms in DeFi projects or analyzing virtual asset transactions within the GameFi sector. The framework's adaptability is further underscored by its ability to dynamically adjust feature extraction modules to suit different ecosystems, such as integrating unique behavioral patterns for ERC-20 and ERC-721 tokens or addressing varying airdrop strategies employed by platforms.

ARTEMIX employs a Boosting-based ensemble learning methodology, combining transaction patterns, temporal features, and community characteristics into a cohesive model that excels in detecting intricate behavioral patterns. This adaptability is enhanced through incremental learning and real-time data updates, enabling ARTEMIX to respond to evolving tactics used by airdrop hunters, such as mimicking genuine user behavior or leveraging cross-chain transactions. To future-proof its capabilities, ARTEMIX is set to incorporate adversarial testing and data augmentation techniques, which will not only bolster its ability to generalize across dynamic blockchain ecosystems but also enhance its defense mechanisms against increasingly sophisticated threats. This comprehensive and flexible design positions ARTEMIX as a reliable detection tool across a wide range of decentralized platforms and use cases.

6.2. Mitigating false positives in airdrop strategies

While the Blur airdrop strategy has been notably effective in driving community participation, it carries the inherent risk of misclassifying genuine participants as airdrop hunters. Addressing this challenge necessitates a deep investigation into the root causes of false positives and the continuous refinement of differentiation mechanisms.

One promising approach involves a detailed analysis of behavioral data associated with high-frequency traders. This could include manual verification of transaction tags provided by clients, which may reveal patterns characteristic of airdrop hunters, such as repetitive or automated transactions between related accounts. Incorporating case studies of common false-positive scenarios into research efforts, alongside quantifying the false-positive rate, would provide a foundation for improving classification methodologies.

To further enhance discriminatory capabilities, more sophisticated behavioral metrics should be integrated. Examples include tracking social network engagement scores, active participation in governance or voting, and sustained investment behaviors, such as longer holding periods for NFTs or tokens. A credibility scoring model could combine trading behaviors with community engagement metrics to yield a composite score reflecting a user's likelihood of being a genuine participant. Additionally, identifying post-airdrop behaviors—such as the immediate liquidation of assets—could offer valuable features for model training, enabling more precise classification.

Beyond refining analytical models, revisiting the structural design of airdrop rules is essential to address their influence on false-positive rates. Adjustments to incentive mechanisms could promote long-term community participation while discouraging exploitative behaviors by airdrop hunters. By aligning these rules more closely with overarching ecosystem objectives, greater fairness and equity can be achieved, fostering a healthier, more sustainable community dynamic.

6.3. Understanding the success of Web3 community

The impact of airdrop hunters and speculators on Web3 projects is indeed significant. They squeeze the space for healthy users to develop and thrive, causing them to lose confidence in the projects. However, despite the fact that airdrops serve as a crucial mechanism to incentivize users, we believe that in the context of the entire project design, the overall performance of tokenomics is only slightly affected by airdrop strategies [63]. For instance, TerraUSD (UST) stablecoin and its native cryptocurrency Luna (LUNA) were once among the most popular and successful projects in the cryptocurrency space. However, in May 2022, the mechanism that pegged UST to the US dollar broke down, causing both UST and LUNA prices to plummet to near zero, leading to significant turmoil in the crypto market [64]. There are many reasons behind the collapse of UST, but poor tokenomics is a major factor. Tokenomics encompasses the entire ecosystem of a cryptocurrency, from its distribution and supply to its demand and utility. The key to design lies in how tokenomics should nurture a thriving ecosystem.

From a broader community perspective, most initial members of current Web3 projects are attracted by the financial attributes of airdrops rather than the governance attributes. For example, 86.39% of initial members of Paraswap sold their PSP within six months and left the community [4]. Financial incentives are excellent for attracting users in the short term, but they do not help retain users and promote community engagement in the long run. A thriving Web3 community must first have a clear and compelling vision that resonates with its members and reflects core values such as decentralization, user control, and openness. Effectively communicating this vision helps community leaders attract like-minded individuals who are passionate about the project's goals. Identifying airdrop hunters is important, but it is not the sole determinant of a Web3 community's success. Instead, innovative strategies that promote engagement, trust, and growth are key. Aligning community goals with principles of decentralization, transparency, and user empowerment is crucial. Essential strategies include defining a clear vision and mission, implementing inclusive decision-making through decentralized governance, and incentivizing active participation through tokens or rewards. Understanding user behavior and values is vital for incentivizing and building a community. By summarizing their behavioral traits and the values and ideologies driving them [65], we might more easily find answers. We believe that by prioritizing these elements, Web3 communities can establish a strong, loyal, and highly engaged user base, driving sustainable growth and innovation.

7. Limitations and future work

7.1. Limitations

Despite the optimistic outcomes, our research has limits. The dataset utilized to evaluate ARTEMIX is predominantly derived from Blur, potentially lacking representation of the full spectrum of airdrop hunting behaviors across various platforms. The model's dependency on specific features identified within the Blur dataset raises concerns regarding its applicability to platforms with differing trading dynamics. Furthermore, our analysis concentrated solely on particular user behavior data within the dataset, employing relatively rudimentary methods. The absence of tailored algorithms for the comprehensive examination of user trading characteristics and behavior identification constrains its performance when faced with other complex and unanticipated trading behaviors. Although we have optimized the computational requirements for training and deploying the model, the necessity for real-time updates presents a challenge for smaller projects with limited resources. Additionally, ARTEMIX's reliance on historical trading data renders it a reactive method rather than a proactive detection mechanism, highlighting the need for real-time detection capabilities.

7.2. Future work

Future research should prioritize diversifying datasets to include multiple NFT marketplaces and a broader range of airdrop events, thereby improving the model's generalizability and adaptability across different contexts. The development of real-time data analysis capabilities is essential to shift detection mechanisms from reactive to proactive, enabling immediate interventions against airdrop hunters and reducing their impact. Advanced machine learning techniques, particularly deep learning models capable of adapting dynamically to emerging patterns, can further enhance the robustness and precision of detection systems. Collaborations with blockchain projects for real-world implementation and testing are crucial for validating these models in practical settings, while advancements in algorithmic efficiency or the integration of distributed computing resources can ensure these sophisticated detection tools remain accessible to smaller projects with limited resources.

Additionally, a deeper exploration of blockchain user behavior patterns—encompassing factors such as transaction frequency, timing, and interaction networks—can reveal nuanced characteristics associated with airdrop hunting, enabling the identification of more complex and subtle behaviors. These refinements will collectively enhance ARTEMIX's applicability and effectiveness across diverse blockchain platforms and environments, fostering a healthier and more resilient blockchain ecosystem.

8. Conclusion

In this paper, we present ARTEMIX, a new model for identifying airdrop hunters in NFT transactions. It outperforms previous models by using a mix of custom-engineered features and a boosted ensemble learning strategy. The system's capacity to precisely detect airdrop hunters while simultaneously ensuring effectiveness and scalability renders it a powerful instrument for augmenting the integrity of airdrop events and the wider Web3 ecosystem. The results of our study emphasize the significance of prioritizing crucial characteristics that have a substantial influence on the accuracy of detection, as well as the advantages of incorporating various detection modules into a cohesive framework. To further advance the field of airdrop hunter identification and ensure the sustainable evolution of decentralized communities, it is crucial to address the observed constraints and explore future research avenues.

Acknowledgments

This work was supported by Guangdong-Hong Kong-Macao Joint Laboratory for Emotional Intelligence and Pervasive Computing, Artificial Intelligence Research Institute, Shenzhen MSU-BIT University.

Conflicts of Interests

The authors declared that they have no conflicts of interests.

Authors contribution

Conceptualization & Investigation, Y.Q., H.C., H.D.; Writing-original draft, Y.Q., H.D.; Writing review & editing, T.M., H.C., H.D.; Project administration, Y.Q., H.D.; Supervision, H.D.; All authors have read and agreed to the published version of the manuscript.

References

[1] Finance. Crypto Airdrop: What is Airdrop in Web3 and How does it work? Available: https://medium.com/coinmonks/crypto-airdrop-what-is-airdrop-in-web3-and-how-does-it-work-7f036af235d7 (accessed on July 24, 2024).

[2] Messias J, Yaish A, Livshits B. Airdrops: Giving money away is harder than it seems. *arXiv preprint arXiv* 2023, 2312.02752.

[3] Atzori M. Blockchain Technology and Decentralized Governance: Is the State Still Necessary? *J. Gov. Regul.* 2017, 6:45–62.

[4] Fan S, Min T, Wu X, Cai W. Altruistic and Profit-oriented: Making Sense of Roles in Web3 Community from Airdrop Perspective. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, Hamburg, Germany,23–28 April, 2023 pp. 1–16.

[5] Malinova K, Park A. Tokenomics: When tokens beat equity. *Manag. Sci.* 2023, 69(11):6568–6583.

[6] Cong LW, Li Y, Wang N. Tokenomics: Dynamic adoption and valuation. *Rev. Financ. Stud.* 2021, 34(3):1105–1155.

[7] Liu Z, Zhu H. Fighting Sybils in Airdrops. *arXiv preprint arXiv* 2022, 2209.04603.

[8] Zhou C, Chen H, Wu H, Zhang J, Cai W. ARTEMIS: Detecting Airdrop Hunters in NFT

Markets with a Graph Learning System. In *Proceedings of the ACM on Web Conference 2024*, Singapore,13–17 May, 2024, pp. 1824–1834.

[9] Cong LW, Li X, Tang K, Yang Y. Crypto wash trading. *Manag. Sci.* 2023, 69(11):6427–6454.

[10] Victor F, Weintraud AM. Detecting and quantifying wash trading on decentralized cryptocurrency exchanges. In *Proceedings of the Web Conference 2021*, Ljubljana, Slovenia, 19–23 April, 2021, pp. 23–32.

[11] Han R, Yan Z, Liang X, Yang LT. How Can Incentive Mechanisms and Blockchain Benefit with Each Other? A Survey. *ACM Comput. Surv.* 2022, 55(7):1–38.

[12] Zhang Y, Tiňo P, Leonardis A, Tang K. A survey on neural network interpretability. *IEEE Trans. Emerg. Top. Comput. Intell.* 2021, 5(5):726–742.

[13] Blondel V, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of community hierarchies in large networks. *ArXiv* 2008, 0803.0476.

[14] Weyl EG, Ohlhaver P, Buterin V. Decentralized Society: Finding Web3's Soul. *SSRN Electron. J.* 2022 .

[15] Bambacht J, Pouwelse JA. Web3: A Decentralized Societal Infrastructure for Identity, Trust, Money, and Data. *ArXiv* 2022, 2203.00398.

[16] Chen H, Duan H, Abdallah M, Zhu Y, Wen Y, *et al.* Web3 Metaverse: State-of-the-Art and Vision. *ACM Trans. Multimedia Comput. Commun. Appl.* 2023, 20(4):1–42.

[17] Duan H, Li J, Fan S, Lin Z, Wu X, *et al.* Metaverse for Social Good: A University Campus Prototype. In *Proceedings of the 29th ACM International Conference on Multimedia*, New York, United States,20–24 October, 2021 pp. 153–161.

[18] Cai W, Wang Z, Ernst JB, Hong Z, Feng C, *et al.* Decentralized Applications: The Blockchain-Empowered Software System, 2018,.

[19] Yang Q, Zhao Y, Huang H, Xiong Z, Kang J, *et al.* Fusing Blockchain and AI With Metaverse: A Survey. *IEEE Open J. Comput. Soc.* 2022, 3:122–136.

[20] Huang H, Zhang Q, Li T, Yang Q, Yin Z, *et al.* Economic Systems in Metaverse: Basics, State of the Art, and Challenges. *ACM Comput. Surv.* 2023, 56(4):1–33.

[21] Duan H, El Saddik A, Cai W. Incentive Mechanism Design Toward a Win–Win Situation for Generative Art Trainers and Artists. *IEEE Trans. Comput. Social Syst.* 2024, .

[22] Qin R, Ding W, Li J, Guan S, Wang G, *et al.* Web3-Based Decentralized Autonomous Organizations and Operations: Architectures, Models, and Mechanisms. *IEEE Trans. Comput. Social Syst.* 2023, 53:2073–2082.

[23] Wood G, *et al.* Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper* 2014, 151(2014):1–32.

[24] Wu K, Ma Y, Huang G, Liu X. A first look at blockchain-based decentralized applications. *Softw. Pract. Exp.* 2021, 51(10):2033–2050.

[25] Raval S. *Decentralized applications: harnessing Bitcoin's blockchain technology*, Eds. O'Reilly Media, United States, 2016.

[26] Leiponen A, Thomas LD, Wang Q. The dApp economy: A new platform for distributed innovation? *Innovation* 2022, 24(1):125–143.

[27] Cong LW, Li Y, Wang N. Tokenomics: Dynamic Adoption and Valuation. *Rev. Financ. Stud.* 2021, 34(3):1105–1155.

[28] Allen DW, Berg C, Lane AM. Why airdrop cryptocurrency tokens? *J. Bus. Res.* 2023, 163:113945.

[29] Ledger. What is a Crypto Airdrop? Available: https://www.ledger.com/academy/what-is-an-airdrop (accessed on July 30, 2024).

[30] Hq G. Sybil Attack explained — Beginners guide for Airdrop farmers. Available: https://medium.com/@gamicHQ/sybil-attack-explained-beginners-guide-for-airdrop-farmers-62dfc34a10f5 (accessed on July 31, 2024).

[31] Zhang L, Ma X, Liu Y. Sok: blockchain decentralization. *arXiv preprint arXiv* 2022,

(2205.04256).

[32] Liu Y, Lu Y, Nayak K, Zhang F, Zhang L, *et al.* Empirical analysis of eip-1559: Transaction fees, waiting times, and consensus security. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. Los Angeles, United States, 7–11 November, 2022, pp. 2099–2113.

[33] Cong LW, Tang K, Wang Y, Zhao X. Inclusion and democratization through web3 and defi? initial evidence from the ethereum ecosystem. Tech. rep., National Bureau of Economic Research Cambridge, MA, USA, 2023.

[34] Singh V. How Sybil Addresses Plunge Aptos (APT) Price. Available: https://coingape.com/sybil-addresses-plunges-aptos-apt-price/ (accessed on July 24, 2024).

[35] Barthere A. "All for One and One for All" - An on-chain distribution model for the Arbitrum community. Available: https://www.nansen.ai/research/an-on-chain-distribution-model-for-the-arbitrum-community (accessed on July 27, 2024).

[36] Xyz D. Outsmarting the system: Sybil attacks in airdrop farming. Available: https://droppables.beehiiv.com/p/sybil (accessed on July 29, 2024).

[37] Kumar A. KYC Frauds Explained: Types, prevention, and impact - Decentro. Available: https://decentro.tech/blog/kyc-frauds/ (accessed on July 25, 2024).

[38] Victor F. Address clustering heuristics for Ethereum. In *Financial Cryptography and Data Security: 24th International Conference*, Kota Kinabalu, Malaysia, 10–14 February, 2020, pp. 617–633.

[39] Wu J, Liu J, Chen W, Huang H, Zheng Z, *et al.* Detecting Mixing Services via Mining Bitcoin Transaction Network With Hybrid Motifs. *IEEE Trans. Syst. Man Cybern. Syst.* 2020, 52:2237–2249.

[40] Gai Y, Zhou L, Qin K, Song D, Gervais A. Blockchain large language models. *arXiv preprint arXiv* 2023, 2304.12749.

[41] TrustaLabs. GitHub - TrustaLabs/Airdrop-Sybil-Identification: This is a repository for Trusta's AI and machine learning framework for robust Sybil identification in airdrops. Available: https://github.com/TrustaLabs/Airdrop-Sybil-Identification (accessed on July 31, 2024).

[42] Labs A. Point Systems: Blur's winning strategy in the NFT marketplace. Available: https://medium.com/@absinthelabs/point-systems-blurs-winning-strategy-in-the-nft-marketplace-265d168c2f85 (accessed on 13 August 2024).

[43] Udeji C. This is how Blur outperformed OpenSea | Blur NFT Marketplace review. Available: https://medium.com/@chinmacryptowriter/blur-outperforms-opensea-see-how-it-did-it-blur-nft-marketplace-review-b5ce6cf06c77 (accessed on July 28, 2024).

[44] Hayward A. How much wash trading is really happening on Blur? Available: https://decrypt.co/122369/wash-trading-blur-ethereum-nfts (accessed on July 27, 2024).

[45] Marsanic D. Blur Airdrop: Just 23 Users Received More Than $1 Million in BLUR Each. Available: https://dailycoin.com/blur-airdrop-23-users-got-more-than-1-million-in-blur/ (accessed on July 29, 2024).

[46] Liu FT, Ting KM, Zhou ZH. Isolation forest. In *2008 eighth ieee international conference on data mining*, Pisa, Italy, 15-19 December, 2008, pp. 413–422.

[47] Brigham EO. The fast Fourier transform and its applications. *Prentice Hall* 1988 .

[48] Breiman L. Random forests. *Mach. Learn.* 2001 45:5–32.

[49] Page L, Brin S, Motwani R, Winograd T. The PageRank citation ranking: Bringing order to the web. Tech. rep., Stanford infolab, 1999.

[50] Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *J. Stat. Mech. Theor. Exp.* 2008 2008(10):P10008.

[51] Lambiotte R, Delvenne JC, Barahona M. Laplacian dynamics and multiscale modular structure in networks. *arXiv preprint arXiv* 2008, 0812.1770.

[52] Newman ME. Modularity and community structure in networks. *Proc. Natl. Acad. Sci.*

2006 103(23):8577–8582.

[53] Song YY, Ying L. Decision tree methods: applications for classification and prediction. *Shanghai Arch. Psychiatry* 2015 27(2):130.

[54] Schapire RE. Explaining adaboost. In *Empirical inference: festschrift in honor of vladimir N. Vapnik*, Springer2013, pp. 37–52.

[55] Vapnik V. The nature of statistical learning theory. *Springer science & business media* 2013 .

[56] Ke G, Meng Q, Finley T, Wang T, Chen W, *et al.* Lightgbm: A highly efficient gradient boosting decision tree. *Adv. Neural Inf. Process. Syst.* 2017, 30.

[57] Perozzi B, Al-Rfou R, Skiena S. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, New York, United States, 24 - 27 August, 2014, pp. 701–710.

[58] Grover A, Leskovec J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, San Francisco California, United States, 13–17 August, 2016, pp. 855–864.

[59] Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv* 2016, 1609.02907.

[60] Hamilton W, Ying Z, Leskovec J. Inductive representation learning on large graphs. *Adv. Neural Inf. Process. Syst.* 2017, 30.

[61] Velickovic P, Cucurull G, Casanova A, Romero A, Lio P, *et al.* Graph attention networks. *stat* 2017, 1050(20):10–48550.

[62] Xu K, Hu W, Leskovec J, Jegelka S. How powerful are graph neural networks? *arXiv preprint arXiv* 2018, 1810.00826.

[63] Nolus. What makes “Good Tokenomics” - Nolus - Medium. Available: https://medium.com/nolusprotocol/what-makes-good-tokenomics-b6a35493fe5 (accessed on July 26, 2024).

[64] Exchange C. Tokenomics 101: How to differentiate good & bad Crypto. Available: https://coinwofficial.medium.com/tokenomics-101-how-to-differentiate-good-bad-crypto-1b99c7bd2487 (accessed on July 25, 2024).

[65] Knittel M, Pitts S, Wash R. " The Most Trustworthy Coin" How Ideological Tensions Drive Trust in Bitcoin. *Proc. ACM Hum. Comput. Interact.* 2019 3(CSCW):1–23.