

NFTDisk: Visual Detection of Wash Trading in NFT Markets

Xiaolin Wen
Sichuan University
Chengdu, China
Singapore Management University
Singapore
wenxiaolin@stu.scu.edu.cn

Yong Wang*
Singapore Management University
Singapore
yongwang@smu.edu.sg

Xuanwu Yue
Baiont Technology
Nanjing, China
Sinovation Ventures AI Institute
Beijin, China
xuanwu.yue@connect.ust.hk

Feida Zhu
Singapore Management University
Singapore
fdzhu@smu.edu.sg

Min Zhu*
Sichuan University
Chengdu, China
zhumin@scu.edu.cn

ABSTRACT

With the growing popularity of Non-Fungible Tokens (NFT), a new type of digital assets, various fraudulent activities have appeared in NFT markets. Among them, *wash trading* has become one of the most common frauds in NFT markets, which attempts to mislead investors by creating fake trading volumes. Due to the sophisticated patterns of wash trading, only a subset of them can be detected by automatic algorithms, and manual inspection is usually required. We propose *NFTDisk*, a novel visualization for investors to identify wash trading activities in NFT markets, where two linked visualization modules are presented: a radial visualization module with a disk metaphor to overview NFT transactions and a flow-based visualization module to reveal detailed NFT flows at multiple levels. We conduct two case studies and an in-depth user interview with 14 NFT investors to evaluate *NFTDisk*. The results demonstrate its effectiveness in exploring wash trading activities in NFT markets.

CCS CONCEPTS

• **Human-centered computing** → **Visual analytics; Information visualization.**

KEYWORDS

Non-Fungible Token; Wash Trading; Visual Analytics; Fintech

ACM Reference Format:

Xiaolin Wen, Yong Wang, Xuanwu Yue, Feida Zhu, and Min Zhu. 2023. *NFTDisk: Visual Detection of Wash Trading in NFT Markets*. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3544548.3581466>

*The corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI '23, April 23–28, 2023, Hamburg, Germany
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9421-5/23/04...\$15.00
<https://doi.org/10.1145/3544548.3581466>

1 INTRODUCTION

With the popularity of digital assets, *Non-Fungible Tokens (NFTs)* have captured increasing attention in recent years [2, 3, 42]. As a special type of cryptocurrency, NFTs represent tradeable ownerships of digital assets (e.g., images, music, videos, and virtual creations), where the ownership is recorded in smart contracts on a blockchain [15]. The huge potential of NFT markets attracts a large number of investors, and at the same time, many fraudulent activities also follow [13]. *Wash trading* is one of the most common fraudulent activities in NFT markets [36], which attempts to mislead investors by creating fake trading volumes [40]. Wash trading is rampant in NFT markets due to the lack of effective regulation. Das et al. [13] have detected 9,393 instances of wash trading that generated more than 96 million dollars in trading volume across 5,297 NFT collections involving 17,821 users. To keep investors from getting ripped off by wash trading, it is important and urgent to figure out how to find and analyze wash trading in NFT markets.

Recent work has presented various strategies to detect wash trading activities automatically by recognizing specific network structures in NFT transaction networks [13, 41]. To evade detection, wash traders are creating increasingly sophisticated transaction patterns by trading different NFTs among multiple anonymous addresses of their own, which makes it possible to avoid some features of wash trading detection in traditional financial markets, such as closed trading loops, and results in some false negatives for automatic algorithms [41]. Hence, automatic methods can just identify a lower bound for the actual extent of wash trading and evaluate the risk of NFT markets at a relatively high level, and manual inspection is usually required [13]. Also, the genuine NFT investors can only get limited help directly from the labels provided by automatic algorithms and the high-level analysis reports, since it is difficult for them to understand what was going on in the historical trading of the NFT collection of their interests. A new crucial concern is how to inform users of the NFT transactions on the blockchain in an intuitive way to support manually detecting the complicated wash trading behaviors and evaluating the impact by using their domain experience.

Visualization and visual analytics techniques have been widely used to enhance people's understanding of blockchain transaction data [37]. Some existing work has leveraged visualization to detect fraudulent activities, such as money laundering [27]. However,

these studies focus on visualizing transaction data for traditional cryptocurrencies, such as bitcoin [14]. Unlike traditional cryptocurrencies, each NFT is distinct and non-fungible, making it critical to pay attention not only to the number of NFTs traded, but also to which NFTs they are. Therefore, it is non-trivial to leverage visualizations to help users recognize and understand wash trading activities effectively. The challenge is two-fold. First, NFTs are often released based on NFT collections. Each NFT collection has a significant number of unique NFTs that are traded between different NFT holders. This creates a large amount of transaction data. How to design visualizations to support quick identification of suspicious transactions and addresses on such a large data scale requires further consideration. Second, verifying wash trading activities requires observing detailed NFT transaction patterns across a set of addresses. It is challenging to visualize NFT transactions to support analyzing NFT flows between groups of addresses at different levels, such as the whole group, individual addresses, or individual NFTs.

In this paper, we introduce *NFTDisk*, an interactive visualization for detecting and analyzing wash trading behaviors in NFT markets, which contains two visualization modules: the *Disk Module* and the *Flow Module*. To be specific, the *Disk Module* is a novel radial visual design with a disk metaphor to provide a quick overview of transactions in one NFT collection to help users recognize suspicious transactions and addresses, as well as evaluate their influence on price and trade volume. An address reordering strategy is integrated in the *Disk Module* to make addresses with more transactions closer together to reduce the visual clutter. The *Flow Module* is a novel flow-based design to visualize NFT transactions at multiple levels and allow users to further check the patterns of wash trading in detail. To evaluate the effectiveness and usability of *NFTDisk*, we conducted two case studies and an in-depth user interview with 14 NFT investors. The results demonstrate that *NFTDisk* is useful and effective for visually detecting and analyzing wash trading activities in NFT markets. In summary, our contributions are listed as follows:

- We formulate the design requirements for visually detecting and analyzing wash trading in NFT markets, together with NFT investment experts, and propose a top-down workflow to help investors filter out unwanted information from a large amount of NFT transactions and analyze wash trading activities effectively.
- We propose *NFTDisk*, a novel interactive visualization, which consists of two novel visualization modules: a radial visualization module with a disk metaphor enabling users to quickly identify wash trading activities in NFT transactions, and a flow-based module to display detailed NFT flows between addresses at multiple levels.
- We conduct two case studies and in-depth user interviews with 14 genuine NFT investors to demonstrate the effectiveness and usability of *NFTDisk*.

2 RELATED WORK

In this section, we summarize the related work in the most relevant fields, including wash trading analysis in NFT markets and visualization for blockchain fraud detection.

2.1 Wash Trading Analysis in NFT Markets

Wash trading is a well-known phenomenon in traditional financial markets and refers to the activity of repeatedly trading assets for the purpose of feeding misleading information to the market. Cao et al. [7] were among the first to analyze wash trading by specifying trading patterns. They conceptualized the basic structures of wash trading using a directed graph of traders and proposed a dynamic-programming algorithm to recognize suspicious transactions and traders. Wash trading in cryptocurrency markets has recently received a lot of attention [11, 16, 18]. Victor et al. [40] presented a heuristic method to detect wash trading and provided a systematic analysis of wash trading behavior on two popular decentralized exchanges on the Ethereum blockchain. Cong et al. [12] introduced systematic tests exploiting robust statistical and behavioral trading patterns to detect wash trading on 29 cryptocurrency exchanges. Chen et al. [9] presented a data mining-based method for detecting wash trading using both off-chain and on-chain data, and the results indicated that several cryptocurrency exchanges have obviously faked trading volume.

As a new type of cryptocurrency, NFTs are also suffering from the wash trading problem. Several recent publications have analyzed wash trading in the NFT market. Das et al. [13] presented a systematic overview of the security issues in the NFT ecosystem, where they detected wash trading by recognizing strongly connected components in the NFT trade networks. Tariq et al. [36] conducted multiple statistical tests and provided some evidence of anomalous price patterns and suspicious automated trading in a substantial portion of the NFT market. Von et al. [41] detected and quantified wash trading behaviors in NFT markets and summarized some features of NFT wash trading. They utilized the Deep-First-Search algorithm to find closed cycles in the NFT transaction graphs and further detect suspicious trade paths with both a high trade speed and low price deviation. CryptoSlam¹, an NFT analysis tool, detects the NFTs with wash trading by checking whether they have repeated holders and shows the label “Wash Sale” when displaying NFTs on their website. In a word, existing work on NFT wash trading focuses on designing automatic detection strategies based on specific features of wash trading and then analyzing suspicious transactions at a relatively high level.

Wash trading in NFT markets is protean since users can hold multiple addresses anonymously, and wash traders are creating increasingly sophisticated patterns to evade existing heuristic detection. Automatic detection can find a subset of wash trading behaviors [41]. However, informing investors about the number of suspicious transactions or labeling NFTs is insufficient to provide investors with enough information to understand what is happening in NFT transactions. For these reasons, we present an interactive visualization technique that allows investors to understand how the NFTs traded between addresses over time and analyze wash trading behavior in depth.

2.2 Visualization for Blockchain Fraud Detection

In recent years, various blockchain data visualization approaches have been proposed to help people understand data recorded on

¹<https://blog.cryptoslam.io/>

the blockchain in a more intuitive way [10, 29, 35, 38, 48]. Fraud detection is one of the most common tasks of blockchain data visualization [8, 31]. Most of the existing visualization tools for blockchain fraud detection rely on analyzing the *cryptocurrency value flow* and the *transaction network* [37].

The *value flow* depicts how the value of cryptocurrency is propagated across different entities over time, which can be used to differentiate some fraudulent activities. For example, Di Battista et al. [14] presented BitConeView, a system for the visual analysis of Bitcoin flows on the blockchain, to assist analysts in gaining an immediate understanding of when and how Bitcoins were mixing in a suspicious way. Bistarelli et al. [5] used node-link diagrams to represent the Bitcoin flows to help recognize frauds like money laundering. Ahmed et al. [1] used a node-link tree visualization to track the stolen coins and reveal the operational techniques of criminals. Xia et al. [44] presented a visual analytics system for monitoring the mining pool by visualizing the Bitcoin flows between miners.

The *transaction networks* emphasize the relationship between various entities on the blockchain, such as addresses [19], clusters [23, 34], and exchanges [46]. McGinn et al. [27] combined both the transaction and address graphs in one high-fidelity visualization of associations to overview Bitcoin transaction activity and support the discovery of unexpected high-frequency transaction patterns. Sun et al. [33] presented BitVis, an interactive system to visualize relationships between Bitcoin addresses and transactions, which can help users conveniently analyze transaction behaviors and track suspicious accounts. Further, Sun et al. [34] presented BitAnalysis, a system for effective interactive investigation of Bitcoin wallets (i.e., a cluster of bitcoin addresses), which can monitor interested wallets in real time and alert the user to new suspicious transactions at a wallet level.

Overall, previous work mostly focuses on some traditional cryptocurrencies such as Bitcoin [5, 27] and Ether [6, 28], but NFT is a new type of cryptocurrency with its own intrinsic features. Each NFT is unique, and all historical transactions of each individual NFT can be tracked on the blockchain. It is a big challenge to track unique NFTs within a large amount of transactions in the blockchain and detect fraudulent activities such as wash trading from the complex transaction network, which is the focus of this paper.

3 BACKGROUND

In this section, we first provide the data description, and then introduce that how we formulate our research problem. Finally, we summarize the design requirements of our method.

3.1 Data Description

Non-Fungible Token (NFT) is a type of cryptocurrency that is derived by the smart contracts of Ethereum [43], which was first proposed in Ethereum Improvement Proposal (EIP)-721². Unlike traditional cryptocurrencies such as Bitcoin and Ether, NFT is unique and can not be exchanged like-for-like (i.e., non-fungible), making it suitable for identifying something in a unique way. Therefore, NFT can prove the existence and ownership of digital assets, such as videos, images, or arts. NFT can be traded in various *NFT marketplaces*,

such as OpenSea³ and LooksRare⁴, as well as transferred directly on the blockchain. All previous holders of each NFT can be tracked since all NFT transactions are recorded on the blockchain, whether traded in the marketplaces or directly transferred.

In NFT markets, NFTs are listed based on their *NFT collections*. An *NFT collection* is an assortment of digital assets released by an artist (or group of artists) containing a limited number of individual NFTs with unique *token id*. Typically, most NFT collections consist of numerous NFTs that all conform to the same artistic style, with slight variations across each individual token. Most investors trade NFTs based on their interests in NFT collections, and the level of suspicious activity varies significantly across NFT collections [41]. For these reasons, our method detects and analyzes wash trading behaviors on a collection basis.

In this paper, we collected all transactions of NFTs in a given NFT collection from the Ethereum blockchain using EtherScan⁵, a block explorer for the Ethereum blockchain. After pre-processing, such as filtering out irrelevant information, each *transaction* used in our method includes:

- *Timestamp*: the timestamp that the transaction occurred at;
- *TokenID*: a unique number representing the token id of an NFT in one collection;
- *Value*: the number of Ether (the standard cryptocurrency on the Ethereum blockchain) contained in the transaction, which can be seen as the sale price of NFT and transferred from the buyer to the seller;
- *Status*: a label of *sale* or *transfer* according to whether the *Value* of the transaction is zero, where *transfer* implies the NFT is transferred without any Ether while *sale* represents that the NFT is sold at the price of *Value*.
- *FromAddress*: the address that sold or transferred NFT to another address;
- *ToAddress*: the address that bought or got NFT from another address.

3.2 Problem Formulation

Financial Conduct Authority⁶ defined wash trading as, “a *sale or purchase of a qualifying investment where there is no change in beneficial interest or market risk, or where the transfer of beneficial interest or market risk is only between parties acting in concert or collusion, other than for legitimate reasons.*” According to the definition, wash trading has two important features: a group of users colluded and pretended to create fake transactions (Feature one), and they traded NFTs without market risk (Feature two).

For Feature one, we need to know the characteristics of address collusion. In NFT markets, NFTs are traded between various addresses, but wash traders can create an arbitrary number of addresses anonymously, which makes it more convenient to trade NFTs between their own addresses. All transactions recorded on the blockchain really occurred and required a fee called *Gas Fee*, and the fees are also needed for wash traders transferring money among fake addresses. Therefore, the number of colluding addresses is

³<https://opensea.io/>

⁴<https://looksrare.org/>

⁵<https://etherscan.io/>

⁶<https://www.handbook.fca.org.uk/handbook/MAR/1/6.html?date=2016-03-07>

²<https://eips.ethereum.org/EIPS/eip-721>

usually limited [13]. Previous studies have reported that most wash trading in the cryptocurrency markets operated as a set of users heavily traded on only small number of NFTs [13]. Therefore, we need to find a limited number of suspicious addresses that heavily trade several NFTs between themselves.

For Feature two, trading without market risk means that there are no changes in the assets held by the colluding addresses after a series of transactions [12], which can make many fake transactions seem like real deals without taking risk. Due to the possibility of multiple addresses controlled by a same person, verify the changes on held assets should be made both at the individual address level and at the address group level. Furthermore, since NFT is non-fungible, we should not only focus on the change in the NFTs held by suspicious addresses, like the wash trading detection for other assets, but also need to further check the transfer process of each unique NFT, called *NFT flow* in this work, to confirm whether the wash trading behaviors really occurred.

In summary, our research problem is how to design a novel visualization to support the identification of addresses with a possibility of collusion, i.e., addresses heavily trading a small set of NFTs, as well as display the NFT flows between these addresses at multiple levels to check the features of wash trading.

3.3 Design Requirements

To collect the requirements that can guide our visualization design, we conducted a preliminary study with two experts (E1, E2). Both experts have conducted NFT investments for at least one year and have been using a few online analysis tools to guide their investments. E1 has experience analyzing the NFT transactions on the blockchain by developing some programs, and E2 has participated in the release of some new NFT projects. We interviewed E1 and E2 one by one and each interview lasted about 40 minutes. To confirm the importance of our study, we first asked the experts about the current state of wash trading in the NFT market, their perspectives on wash trading, and the limitations of existing tools. Then, we collected their feedback on their general workflow for their analysis of wash trading and the possible functions that they wanted our visual analytics system to support. The specific design requirements are summarized as follows:

R1 Analyze wash trading in the scope of NFT collection.

Experts said that NFTs in the same collection usually have similar price trends, and investors are concerned about the prospects of the whole NFT collection, rather than specific NFTs. Our method should support analyzing wash trading based on individual NFT collections.

R2 Recognize suspicious transactions and addresses from the overview.

Our method should provide a quick overview of the transaction history of the whole NFT collection, where the undesired information should be filtered out and the suspicious transactions and addresses should be easily distinguished in an effective way. Experts pointed out that the more recent transactions are the more important for analyzing wash trading, so our method should support the time selection and make the recent transactions more clear in the overview. Flexible interactions are also needed to be

integrated in the overview to help users select a group of suspicious addresses for further investigation.

R3 Reveal wash trading features at multiple levels. Our method should support analyzing the changes in NFTs held by the selected suspicious addresses in detail for verifying whether the wash trading really occurred. To be specific, it should be allowed to analyze the NFTs held by both the whole address group and individual addresses over time, as well as track the previous holders of each individual NFT.

R4 Display the detailed transaction patterns of wash trading. Experts said that investors care not only about whether there are any wash trading behaviors in the collection but also about how the colluding address group conducted the wash trading to avoid being cheated by them. Our method should display correct and comprehensive information of NFT transactions, such as the *Timestamp*, *Status*, and *TokenID*.

R5 Enable the evaluation of wash trading influence. Experts said that it is necessary to provide evidence to evaluate the extent to which the given NFT collection has been impacted by wash trading. Our method should answer the following questions: when do the wash trading behaviors appear more frequently, how many NFTs and addresses are involved in wash trading, and what is the impact on the price and trading volume of the given NFT collection? To this end, more information such as average price and trading volume should be included.

4 NFTDISK

According to the design requirements above, we present *NFTDisk*, an interactive visualization that supports detecting and analyzing wash trading behaviors at multiple levels, as shown in Figure 1. *NFTDisk* consists of two visualization modules: the *Disk Module* (Figure 1A) and the *Flow Module* (Figure 1B), and follows a drill-down workflow to help users filter out undesired information to concentrate on the wash trading behaviors. First, the *Disk Module* visualizes the NFT transactions between addresses with heavy transactions as well as the average price and trade volume of the whole NFT collection (R1). Users are allowed to find out the period and the group of addresses with suspicious transaction patterns from the overview (R2) and select them for further exploration in *Flow Module* by using a circular brush interaction (Figure 1A3). The *Flow Module* displays the detailed NFT transfer process among addresses at the levels of address group, individual addresses, and individual NFTs, respectively (R3). With the *Flow Module*, users can identify the features and transaction patterns of wash trading behaviors at multiple levels (R4). After that, users can think about both the wash trading activities and the average price or trading volumes shown in the background of the *Disk Module*, and evaluate the extent to which the wash trading impacts this collection in a reasonable way (R5). Following this workflow, users can easily detect and analyze the wash trading activities of a given NFT collection. The specific visual design and interactions are described in the following subsections.

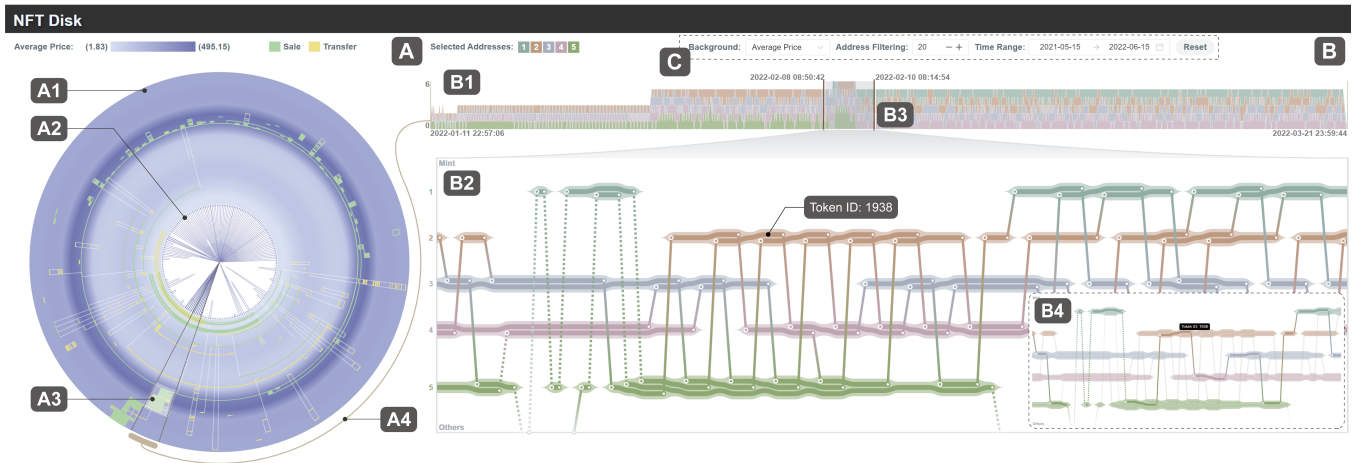


Figure 1: The NFTDisk interface consists of the Disk Module (A), the Flow Module (B), and a tool bar (C) to provide some necessary interactions. The Disk Module contains an Outer Ring (A1) to overview NFT transactions and Inner Circle (A2) to display the suspicious score of address pairs. Users can brush an area (A3) in the Disk Module to select a group of addresses. In the Flow Module, a stacked area chart (B1) shows NFT flows at the group level, where users can brush a period (B3) to see detailed NFT flows in a flow-based chart (B2). Users can hover over the paths of each NFT to highlight them (B4).

4.1 Disk Module

The Disk Module is designed to provide an overview of NFT transactions on the blockchain, which is presented in a radial layout. It contains two parts: the Outer Ring (Figure 1A1) to show the transaction details and the Inner circle (Figure 1A2) to show the possibility of collusion of address pairs, where the Outer Ring encloses the Inner Circle.

4.1.1 Outer Ring. In the Outer Ring, NFT transactions are displayed in the radial layout, which is extended from the circular Massive Sequence Views (MSV) [39]. As shown in Figure 2A, the addresses are placed evenly around the circle, with each address encoded by a radius line at a specific angle. The time is represented by the concentric circles, with the smaller circles denoting an earlier timestamp and a larger circle representing a later timestamp. The exact time range can be interactively configured by users. NFT transactions are encoded by the arcs between the two radius lines at different angles that indicate the two involved addresses, as shown in Figure 2A1. The arcs have two colors encoding the status of transactions: green for “Sale” status and yellow for “Transfer” status. We use the white lines along the radius (Figure 2A2) to link the first and last transaction of the addresses representing this radius. The white lines can help users check whether the addresses are from new investors in this collection or not, because new investors with suspicious behaviors have a greater possibility of being the fake addresses created by wash traders. Furthermore, a radial gradient blue background (Figure 2A3) is used to show the monthly average price or trading volume (R5), where the depth of the blue indicates the specific value and the change over time is represented by the inner circle to the outer circle. In the Outer Ring, we use absolute timestamps rather than the relative time sequence, as they help confirm the exact time when the wash trading occurred to assess the scale and impact. But the use of absolute timestamps

can result in overlapping arcs when several transactions occur at the same time or in a relatively short time period. We deal with the overlapping arcs by allowing users to zoom into a shorter time span or brushing this period to observe details in the Flow Module (Figure 1B). With the Outer Ring, users can overview the historical transaction patterns and analyze their influence on the average price or trading volume.

The reasons for the visual design are manifold. First, according to R2, our design should make the colluding address groups with heavy transactions easier to distinguish. The circular Massive Sequence Views (MSVs) have been shown able to visualize repeated patterns by utilizing Gestalt principles of closure and proximity [39]. Inspired by it, we designed the Outer Ring (Figure 2A) to show the repeated transactions within address groups and imply wash trading behaviors. Specifically, the repeated transactions between two addresses can form an obvious green block in the Outer Ring (Figure 2A4). Second, a circular layout can effectively reduce visual clutter by drawing the smaller of the two arcs between two addresses. In Section 3.3, the experts pointed out that the more recent transactions are the more important for analyzing wash trading. Thus, our visualization is designed to put the more recent transactions closer to the outer circle, where there is more space, and it can make the recent transaction patterns more clear. Third, to facilitate the impact analysis (R5), we added the gradient background to help compare the transactions, prices, and trade volumes based on the same timeline.

4.1.2 Inner Circle. To further help users recognize suspicious addresses quickly, we present the Inner Circle in the center of the Outer Ring. According to Das et al. [13], wash traders typically trade a small set of NFTs heavily, so it is suspicious when two addresses generate a large number of transactions but only involve a small number of NFTs. To this end, we propose a measurement, Suspicious Score, to indicate the possibility of colluding address




Figure 2: Visual encoding of *Disk Module*: In the Outer Ring (A), addresses are placed around the circle, while the timeline is from the inner circle to the outer circle. The green and yellow arcs (A1) represent the NFT sale and transfer between two addresses. The white line (A2) shows the first and last transactions at one address. The gradient blue color of the background (A3) shows the average price or trading volume over time. The repeated transactions between two addresses can form a green block (A4). In the Inner Circle (B), the inner path (B1) shows the suspicious scores (B4) of address pairs by their height (B2). The closer the inner path is to the center, the more suspicious the address pair is. (B3) shows some examples of suspicious addresses and a normal address.

pairs:

$$S = 1 - \frac{N}{M}, \quad (1)$$

where M is the number of transactions between the two addresses, and N is the number of unique NFTs involved in these transactions. The higher the suspicious score, the more likely the address pair is to collude. If each transaction from a pair of addresses includes a different NFT, then their suspicious score is zero.

The Inner Circle is designed to display the suspicious score of address pairs in an intuitive way to make the suspicious address pairs easier to identify. As shown in Figure 2B, each address is represented as a point around the circle, which is consistent with the layout in the *Disk Module*. The points are connected by curves, called the *inner paths* (Figure 2B1). As shown in Figure 2B2, we define the height of the inner paths as the distance from the outer circle to its end on the side of the circle center, which represents the suspicious score of this address pair. In other words, circles of different radius from the outer contour to the circle center represent suspicious scores from zero to one, as indicated by the dotted circle in Figure 2B4. An inner path representing a pair of addresses always begins at one of them and travels through the dotted circle representing their suspicious score to the other, as shown in Figure 2B2. It can also be regarded as the closer the inner path is to the center of the circle, the higher the suspicious score for this pair of addresses. If the suspicious score is zero, the inner path will not be drawn. When users observe the inner circle, they can quickly identify the addresses with a high suspicious score according to the

heights of the inner paths. Figure 2B3 shows examples of suspicious addresses and a normal address, respectively.

4.1.3 Reordering Strategy. For the layout based on MSV, the node order is of vital importance, and a bad address order can seriously degrade the readability of the *Disk Module*, as shown in Figure 3A and Figure 3C. Previous studies have proposed various reordering strategies to emphasize different features of MSV, such as reordering based on the node degree, edge length [39], or neighbor relationships [26]. In this work, we propose a reordering strategy to make the addresses with more transactions closer. Inspired by the reordering of rows or columns in matrix visualization [17], we regard address reordering as a *Traveling Salesman Problem* [21] to get an order so that the address pairs with many transactions are as close to each other as possible. Specifically, we convert the number of transactions between addresses into the distance between them, where two addresses with more transactions have a shorter distance. Then, we use the *optimal leaf ordering* algorithm [4] to get an optimal address order to minimize the sum of distances between adjacent addresses. The optimal leaf ordering algorithm will first compute hierarchical clustering according to the distance matrix between addresses to get a cluster tree, and then traverse the leaf nodes of the cluster tree according to the hierarchical structure to get the final order. Figure 3 shows the *Disk Module* of two NFT collections named *Meebits* and *BAYC* using the random order or reordering by our strategy, which demonstrates that our strategy can effectively make the addresses with more transactions closer and make colluding groups look like blocks composed of many arcs.




Figure 3: The figures showcase the *Disk Module* with different address orders to show the effectiveness of our reordering strategy. (A) and (C) show the *Disk Modules* of two NFT collections named *Meebits* and *BAYC* with a random address order, while (B) and (D) show them after using our reordering strategy. Compared with (A) and (C), (B) and (D) have less visual clutter.

4.2 Flow Module

The *Flow Module* (Figure 1B) is designed to visualize the detailed NFT flows at three different levels (R3): the group, address, and NFT levels. The group and address levels are intended to display the NFTs held by the entire address group and the individual addresses over time, respectively, while the NFT level is intended to demonstrate the trading process of each individual NFT. Specifically, the *Flow Module* contains two parts: a stacked area chart (Figure 1B1) and a flow-based chart (Figure 1B2), which cover the three levels above, as shown in Figure 4. The specific visual designs are described as follows.

4.2.1 Stacked Area Chart. In the *Flow Module*, a stacked area chart is used to demonstrate the number of NFTs held by the selected addresses over time at the group level, as shown in Figure 4A. The X-axis is a timeline, indicating the relative time sequence of the transactions rather than the absolute timestamp. The height of the whole area (Y-axis) represents the number of NFTs held by this group after each transaction over time, as shown in Figure 4A1. We use a categorical color scheme to represent different addresses, and the whole area is divided into parts in different colors by the individual addresses to show the number of each address in this group. Users can identify temporal changes in the number of NFTs held by the selected group and roughly deduce the number of NFTs held by each address from the heights of different colored areas using the stacked area chart. If the height of the stacked area chart remains constant over time (Figure 4A2), it indicates that NFTs flow only between this group of addresses, which is probably wash trading activity. Users are also allowed to brush a period in the stacked area chart to further analyze the detailed NFT flows in the flow-based chart below.

4.2.2 Flow-based Chart. The flow-based chart (Figure 1B2) is the core of the *Flow Module*, which is designed to visualize the number of NFTs over time at the address level as well as show the detailed transaction patterns at the NFT level. As shown in Figure 4B, the Y-axis places the addresses that have transactions in the selected period, and the X-axis represents the brushed time period in the stacked area chart. The order of addresses is the same as the order they are stacked in the stacked area chart as well as their order in

the *Disk Module*. The NFTs held by each address are represented by a horizontal flow graph in the color representing this address, where the height of the flow graph represents the number of NFTs held by this address, as shown in Figure 4B1. Users can identify the temporal evolution of NFTs held by each address by the height change of the flow graph along the X-axis.

For each individual NFT, we use a path to visualize its transition process among its previous holders, which is called an NFT trade path (Figure 4C1) in this work. If one address is holding an NFT in a time period, the path of this NFT will be packed in the flow chart of this address in this time period, where the path is in the same color as the holder (Figure 4C). When the NFT is transferred or sold to another address, the path will link from the *fromAddress* to the *toAddress* with a dotted line (Figure 4C2) or a solid line (Figure 4C3), respectively. In addition, if one party in the transaction is an address outside the selected address group, the path will link to the bottom border (Figure 4C4). And if the NFT is got from the owner of the NFT collection, called *Mint* in NFT markets, the path will link from the top border (Figure 4C5). The segment of the path that represents transactions is filled in a linear gradient color from the color of the *fromAddress* to the color of the *toAddress*. When an address holds multiple NFTs, the paths of these NFTs are placed side by side and just fill the flow graph of this address. To reduce the crossing of paths, the NFT is inserted from the side of *fromAddress* (above or below this address) when getting a new NFT. In this way, users can find out not only how many NFTs these addresses held by the contour of these paths and flow graphs in background, but also the detailed NFT transaction patterns by these trade paths.

4.3 Interactions

NFTDisk enables rich interactions to allow users to smoothly explore the wash trading behaviors in NFT transactions.

Visualization Configuration. In Figure 1C, *NFTDisk* allows users to configure the visualization to filter out some undesired information. The *Background Selector* allows users to select between monthly average price or trading volume to be shown in the background of the *Disk Module*. The *Time Selector* allows users to set the time range of their interest. The *Address Filter* allows users to filter the address pairs by the number of transactions between

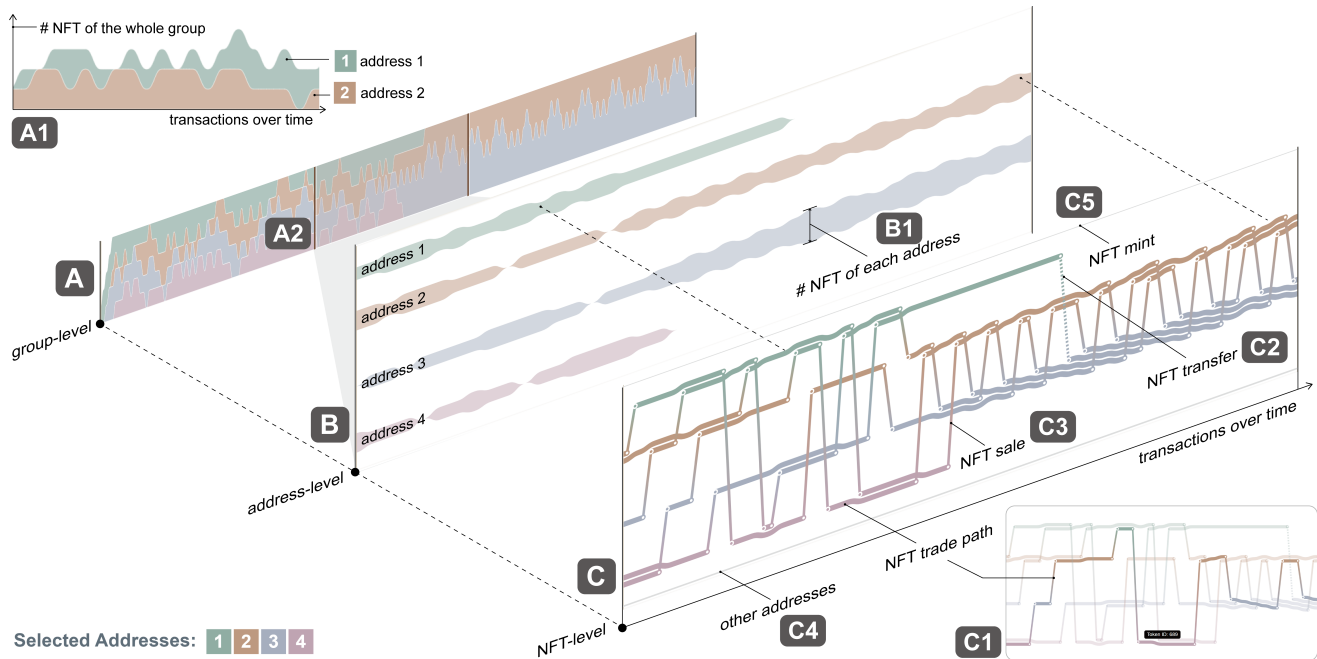


Figure 4: The multiple levels of *Flow Module*: The group level (A) shows the number of NFTs held by the whole selected address group using a stacked area chart (A1), where users are allowed to brush an area (A2) to check the details. The address level (B) shows the number of NFTs held by individual addresses, where each address is encoded with a flow graph (B1). The NFT level (C) superimposes on the address level, which demonstrates the trade paths (C1) of each NFT.

them to select suspicious addresses that are possible to collude with each other. Users set the threshold according to their experience, and then the address pairs with more than the threshold will be visualized in the *Disk Module*.

Multi-level Brush. *NFTDisk* provides two types of brush interactions to help users drill down to the details of their interests. Firstly, a circular brush is used in the *Disk Module* (Figure 1A3), where users can brush an arc area to select a group of suspicious addresses as well as a time period for further analysis in the *Flow Module*. For an arc area, the inner radius and outer radius determine the time range, while the start angle and end angle determine a set of addresses between the two angles. After brushing, a curve is drawn to link the selected area to the flow module, as shown in Figure 1A4. The second brush is integrated in *Flow Module* at the group level (i.e., the stack area chart) to allow users to select a time period, and then the detailed transactions of this period will be displayed below.

5 CASE STUDIES

We describe two case studies in this section to demonstrate the effectiveness of *NFTDisk* in analyzing wash trading activities in real NFT collections. The two cases were conducted by two users (U1, U2) during the user interview that will be introduced in Section 6, where users were asked to learn *NFTDisk* and use it to analyze NFT collections. To demonstrate the generalization of *NFTDisk*, four popular NFT collections with different features were used in our study, and all historical transactions of these collections from their

launch time to June 30, 2022, were collected. These collections are described as follows:

Bored Ape Yacht Club (BAYC)⁷ is a collection of 10,000 Bored Ape NFTs—unique digital collectibles that was launched on April 15, 2021. It is one of the most popular and expensive NFT collections in NFT markets, held by 6,460 unique owners, and has a total market cap of more than one billion dollars as of June 30, 2022.

Meebits⁸ is an NFT collection of 20,000 unique 3D voxel characters, created by a custom generative algorithm, then registered on the Ethereum blockchain. It was launched on May 4, 2021, was held by 6,572 unique owners, and had a total market cap of more than 110 million dollars as of June 30, 2022.

Azuki⁹ is a collection of 10,000 generative avatar NFTs launched on January 12, 2022. These avatars have unique characteristics based on anime-themed drawings, which have caught the attention of NFT enthusiasts around the world. As of June 30, 2022, it was held by 5,062 unique owners and had a total market cap of more than 118 million dollars.

Loot¹⁰ is a collection of 7,779 NFTs that is randomized adventurer gear generated and stored on blockchain. It is a typical collection based on the NFT gaming platform and was launched on August 27, 2021. As of June 30, 2022, it was held by 2,567 unique owners, and had a total market cap of more than 18 million dollars.

⁷<https://boredapeyachtclub.com/>

⁸<https://meebits.app/>

⁹<https://www.azuki.com/>

¹⁰<https://lootnft.io/>

5.1 Case 1: In-depth Investigating of Wash Trading within an NFT Collection

U1 is the builder of an NFT community and often analyzes NFT transactions to help investors in the community. When analyzing the *Meebits*, one of the collections that U1 analyzed in the user interview, he first set the time range from October 1, 2021, to June 30, 2022, and set the *Address Filter* to 20 according to his experience. Then, the *Disk Module* provided an overview of NFT transactions during this period, where the gradient blue background represented the average price, as shown in Figure 5A. In the Inner Circle, he found that most of the inner paths were so close to the circle center, i.e., with a high height, which meant that addresses connected by these inner paths had a high suspicious score of collusion. In the Outer Ring, he found that there were a lot of obvious shapes consisting of many consecutive arcs, of which a large green block (Figure 5A1) and a circle of small green blocks (Figure 5A2) caught his attention. These green blocks indicated some groups of addresses had traded heavily in NFTs within a relatively short time, which was a sign of wash trading. Additionally, the areas around these green blocks in the gradient background were much darker than the rest, indicating that these suspicious transactions might have an impact on the price of this collection.

A Big Group of Colluding Addresses. To further confirm the NFTs involved in wash trading and the trading patterns, U1 first selected the area of the big green block shown in Figure 5A1 with the circular brush, and then the stacked area chart showed the number of NFTs held by 11 selected addresses at the group level from January 11, 2022, to June 17, 2022 (Figure 5C). Overall, U1 found that the height of the stacked area chart increased over the whole time period, while most of the time the height stayed the same. It indicated that these addresses were holding an increasing number of NFTs at the group level, but they were always trading them within the group. Furthermore, U1 discovered that the colors in the stacked area chart change in a sophisticated way over time, and he wondered how this group of addresses conducted the wash trading.

By brushing and dragging the brushed area on the stacked area chart, he analyzed the detailed NFT flows of multiple time periods with the flow-based chart, as shown in Figure 5(D1-D6). During the period of Figure 5D1, U1 found four obvious flow graphs of address 7,8,9,10, and each of them kept the height of one or two NFTs. It implied that each address still held one or two NFTs during the heavily traded process, which was exactly feature two of wash trading, trading NFTs without taking market risk. By observing the trade paths of each unique NFT, U1 discovered that each NFT had many repeated transaction patterns, i.e., each NFT was traded in the exact same order. Then, he deduced that these addresses are trading robots that executed the wash trading automatically based on pre-programmed procedures. Following that, in Figure 5D2, U1 found that the flow graphs of these four addresses disappeared while the flow graphs of another four addresses appeared and were connected with dotted lines, indicating that these four addresses directly transferred all NFTs to address 3,4,5,6. U1 guessed that wash traders created four new addresses to conduct the wash trading to avoid detection. Similarly, U1 found that this group of addresses heavily traded NFTs in various ways after that period. For example,

they sold all NFTs to only addresses 4 and 5 (Figure 5D3), and then address 4 sent all NFTs back to address 6 (Figure 5D4). In Figure 5D5, all NFTs were sent to six addresses separately and then traded within two pairs of addresses (Figure 5D6). According to the findings above, U1 believed that it was a sophisticated wash trading activity by a big group of colluding addresses, and that they were attempting to create fake transactions that seemed normal with various trading patterns.

Simultaneous Appearance of Wash Trading. U1 then wanted to investigate the area in Figure 5A2, where many small green blocks were distributed on a ring area in dark blue. The majority of the green blocks contained numerous transactions between two addresses, which U1 can roughly attribute to wash trading of address pairs according to their suspicious scores presented by inner paths. He changed the background to show the monthly trading volume (Figure 5B), and found that a light blue ring appeared on the outer area of these green blocks (Figure 5B1), which was in dark blue when the background displayed the price (Figure 5A2). It indicated that after the appearance of these green blocks, the price of this collection rapidly increased while trading volume abruptly decreased. U1 believed that most transactions during this time were high-priced and created by wash traders, so he brushed these green blocks to check the details. After brushing the areas in Figure 5(B2-B4), U1 observed several wash trading activities on a small scale, involving a small number of addresses and NFTs, as shown in Figure 5(E1-E3). Specifically, in Figure 5E1, wash trading was just between two addresses by heavily trading two NFTs. In Figure 5E2, two addresses first traded one NFT to each other many times before transferring it to another two addresses to perform wash trading. From Figure 5E3, the second address traded NFTs with others alternatively, which might be the main address of this group according to U1.

Overall, after using *NFTDisk*, U1 believed that the *Meebits* collection was filled with various wash trading activities, and wash trading had impacted the price and trading volume of this collection. U1 said that gaining these insights would have taken much longer without using *NFTDisk*.

5.2 Case 2: Wash Trading Enhanced by Trading Rewards but Discouraged by Royalties

U2 is an NFT investor with over six months of experience trading NFT in NFT markets, and also a product manager of NFT projects. During the user interview with U2, we first introduced a usage scenario for the *Meebits* collection and then asked U2 to analyze the other three collections (i.e., *BAYC*, *Azuki*, and *Loot*) by using *NFTDisk*.

Trading Rewards Promoted Wash Trading. U2 first analyzed the wash trading activities in the *Loot* collection, and its *Disk Module* is shown in Figure 6A. She found that there were several green blocks on a dark blue ring (Figure 6A1), which reminded her of a similar pattern she had noticed while watching the usage scenario on the *Meebits* collection. This pattern indicated that some small groups of addresses had a high volume of transactions during the same time period, and the average price was obviously higher than in other periods. U2 brushed these green blocks and observed the *Flow Module* to see if they were wash trading activities. In




Figure 5: NFTDisk demonstrates various wash trading activities in the Meebits collection. (A) and (B) show the Disk Module of Meebits with the background representing the average price and trading volume respectively. In Disk Module, a big green block (A1) and many small green blocks (A2) are obvious. (C) shows the Flow Module at the group level after brushing the green block in (A). (D1-D6) demonstrate the flow-based chart after brushing six different time periods in (C). (E1-E3) show the Flow Module after brushing areas in (B2-B4), respectively.

Figure 6A2, she discovered a group with four addresses and saw the detailed NFT flows between them in the Flow Module, as shown in Figure 6A4. She brushed the green block in Figure 6A3 and found there were only two flow graphs in the flow-based chart (Figure 6A5), where some dotted lines connected from the bottom border to the brown flow graph and all trade paths were between the two flow graphs with little change in height over time. This indicated that the brown address obtained some NFTs from other addresses and traded them with the green one without taking any market risk. U2 found that other green blocks in Figure 6A1 have similar patterns that represented small scale wash trading activities between two addresses. Furthermore, U2 discovered that the time periods of wash trading in *Loot* and *Meebits* were both in January 2022, and she speculated that some events might have occurred that resulted in the emergence of these wash trading activities. Then, she thought about the NFT marketplace, *LooksRare*⁴, was established in January 2022 and became popular rapidly. U2 said that *LooksRare* suggested trading rewards for addresses that traded frequently, which could be one of the factors promoting the simultaneous appearance of wash trading in January 2022.

Royalties Discouraged Wash Trading. Next, U2 used *NFTDisk* to explore *Azuki* and *BAYC* one by one. In the Disk Module of *Azuki* (Figure 6B), U2 discovered many yellow blocks on the Outer Ring and only one green block in Figure 6B1. The yellow blocks contained many NFT transfers without money, which could have been caused by some addresses belonging to the same person being used for different functions and can not be considered wash trading. Then, U2 brushed the only green blocks and saw that there were

three NFT trade paths in the flow-based chart (Figure 6B2), two of which are packed in the green flow graph, and one of which traveled between the two flow graphs. This indicated that the green address held three NFTs and always traded one of them to the brown address to perform the wash trading. *BAYC*'s Disk Module (Figure 6C) was similar to *Azuki*'s, with many yellow blocks and only one green block. Figure 6C4 showed the Flow Module after brushing the yellow block in Figure 6C2, where U2 found two flow graphs connected with many dotted lines and the NFTs held by the two addresses changing irregularly, and she believed that it was not wash trading activities. By checking the Flow Module of the green block (Figure 6C1), U2 found that two addresses were heavily trading one NFT whose token ID is 4171. Overall, there were few wash trading activities in both the two collections. Since U2 is a NFT product manager, she really wondered why there have not been so many wash trading activities in January 2022, unlike the other two collections, i.e., *Meebits* and *Loot*. She discovered that trading NFTs of *BAYC* and *Azuki* required a certain amount of money, called royalties, to be paid to the owner of the collection or marketplace, whereas trading the other two collections did not. In this way, U2 got a rough insight that trading rewards might promote wash trading while royalties restrict it, and said that she is more likely to purchase NFTs with royalties in the future.

6 USER INTERVIEW

We conducted semi-structured user interviews with 14 NFT investors to demonstrate the effectiveness and usability of *NFTDisk*. Specifically, we asked the investors to analyze the wash trading

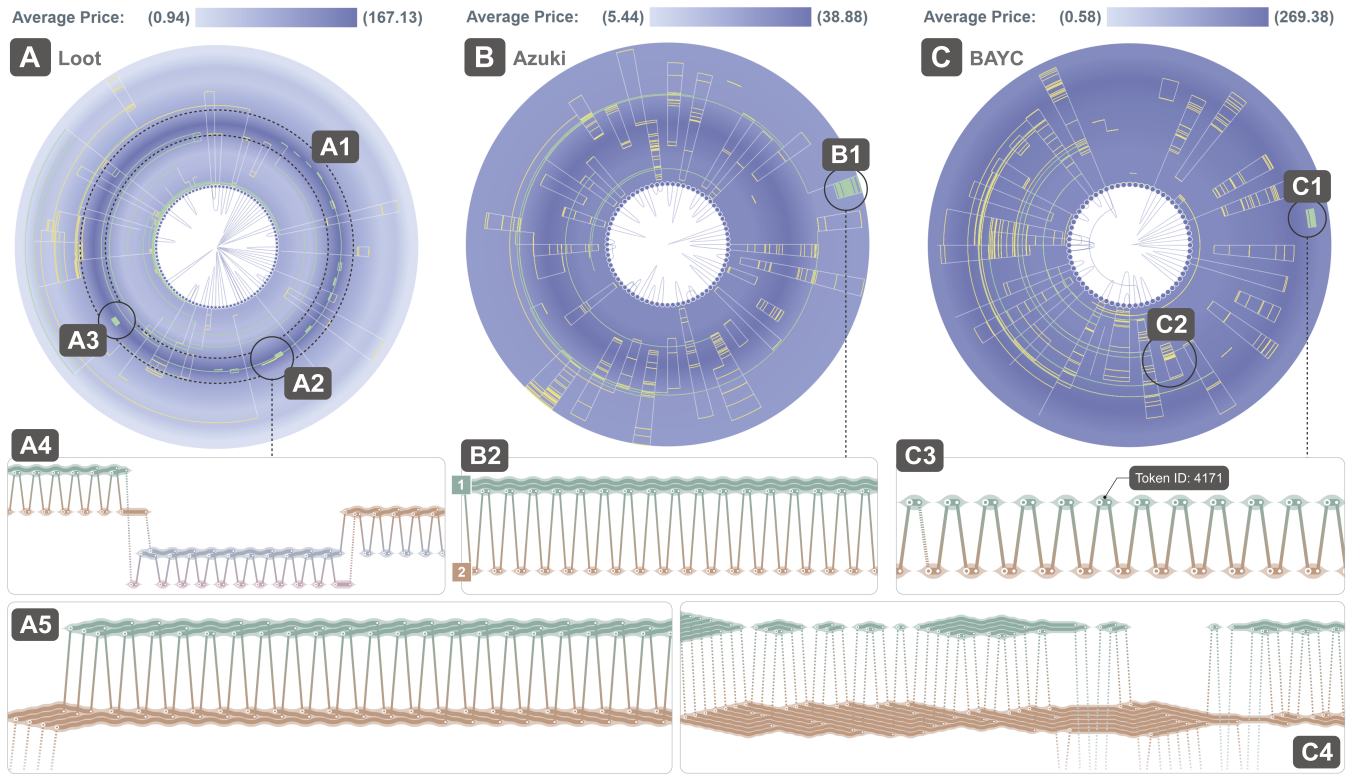


Figure 6: *NFTDisk* shows the wash trading activities in three different NFT collections. (A), (B), and (C) show the *Disk Module* of *Loot*, *Azuki*, and *BAYC*. (A1) shows a dark blue ring containing many green blocks in the *Disk Module* of *Loot*. After brushing green blocks in (A2) and (A3), the *Flow Module* is shown in (A4) and (A5). (B2) shows the NFT flows of the area in (B1). (C3) and (C4) show the *Flow Module* of the green block in (C1) and the yellow block in (C2).

activities in real NFT collections with *NFTDisk* and further collected their feedback on *NFTDisk*. This section first introduces the setup of our interviews, like participants’ backgrounds, interview procedures, and tasks, and then summarizes the feedback collected from participants.

6.1 Participants and Apparatus

We recruited 14 target users (U1-U14) from various NFT communities for the user interviews (4 females, 10 males, $age_{mean} = 28$, $age_{sd} = 5.45$, with normal vision and no color-blindness). All the participants are experienced cryptocurrency investors with a experience of at least two years, and they have been creating or investing in the NFT market. The detailed information of our participants is summarized in Table 1. U12-U14 are novices with less than six months of NFT trading experience, and U1-U11 have all traded NFTs for more than six months and are proficient with existing NFT data analysis tools, with U3 being the second expert (E2) that attended our preliminary study. U1, U4-U6 are the creators of some NFT communities that include thousands of investors, with U1 and U4 being key opinion leaders (KOLs) on Twitter with over ten thousand followers. U2, U7, and U8 have engaged in the issuance of NFT projects, and U13 is also a professor who studies the digital economy. Due to the COVID-19 pandemic, our interviews were

conducted online via Zoom. We launched the prototype system of *NFTDisk* on the server and allowed participants to assess it via their own laptops or desktops. Each interview lasted about one hour, and we paid a compensation of \$15 to each participant for their time in our user interviews.

6.2 Tasks

We collected transactions of four different NFT collections that mentioned in Section 5 and used them in our user interview. Specifically, we chose one of them for a tutorial usage scenario and asked participants to analyze the other three collections using *NFTDisk*. The collections are selected sequentially to ensure that each collection can be analyzed by participants. During the interviews, the participants were asked to perform the following four tasks sequentially on each collection, which guided them to go through *NFTDisk*’s workflow.

- T1. Initialize the visualization by using interactions components to filter out undesired information.
- T2. Observe the *Disk Module* to find suspicious addresses and time periods and brush to select them.
- T3. Analyze the NFT flows at the group level by the stacked area chart of the *Flow Module*.

Table 1: The detailed information of the user interview participants. All participants have a cryptocurrency investment experience for at least two years, and their detailed NFT investment experiences are listed in the third column.

ID	Gender	Age	NFT Experience	Description
U1	Male	23	13 months	A creator of an NFT community and a key opinion leader on Twitter.
U2	Female	25	8 months	A product manager for multiple NFT projects.
U3	Male	30	12 months	An NFT investor who is good at using NFT analysis tools.
U4	Female	26	12 months	A creator of an NFT community and a key opinion leader on Twitter.
U5	Male	29	10 months	A creator of an NFT community and a leader of an NFT project.
U6	Male	25	12 months	A creator of an NFT community and a leader of three NFT projects.
U7	Female	23	7 months	An NFT investor engaged in the issuance of NFT projects.
U8	Male	27	10 months	An NFT investor engaged in the issuance of NFT projects.
U9	Male	27	6 months	An NFT investor who is good at using NFT analysis tools.
U10	Male	30	12 months	An NFT investor investing in cryptocurrencies for five years.
U11	Male	25	7 months	An NFT investor investing in cryptocurrencies for two years.
U12	Male	28	5 months	An NFT investor investing in cryptocurrencies for two years.
U13	Male	46	4 months	A professor whose research focus is digital economy.
U14	Female	28	5 months	A PhD student with two-year research experience in cryptocurrencies.

- T4. Brush a period in the stacked area chart and check the detailed NFT flows in the flow-based chart.

Participants were allowed to repeat these tasks until they had a thorough understanding of the wash trading activities in this collection.

6.3 Procedure

During the interview, we first introduced the background, visual design, interactions, and workflow of *NFTDisk* to the participants. To show how to use *NFTDisk* to explore the wash trading of NFT collections, we went through an example usage scenario on one of the four collections. Then, the participants were allowed to explore this collection freely to make themselves familiar with *NFTDisk*. The tutorial above lasted about 15 minutes. After that, they were asked to follow the tasks above by using *NFTDisk* to analyze the other three datasets sequentially. The task phase had no hard time limit and lasted until participants fully understood the wash trading activities in these collections. In our interview, this phase usually lasted 30 minutes, and their comments and suggestions were recorded. After the task phase, we invited them to finish a post-study questionnaire with 14 questions (Q1-Q14), as shown in Figure 7. Q1-Q12 are closed-end questions that should be answered on a 7-point Likert scale and are designed to measure *NFTDisk*'s workflow effectiveness (Q1-Q5), visual design and interactions (Q6-Q8), and usability (Q9-Q12). By following the question design of prior studies [32, 45], Q1-Q5 are designed to assess the workflow effectiveness in terms of different concrete tasks, and Q6-Q8 emphasize the detailed visual designs and interactions. Q9-Q12 are usability evaluation questions we selected from the PSSUQ (Post-Study System Usability Questionnaire) [25]. Q13-Q14 are open-ended questions that aim to collect participants' feedback on the advantages and possible areas for further improvement of *NFTDisk*. Overall, each user interview session took about 60 minutes. All the data collected from the participants are recorded anonymously with their permission.

6.4 Results

Figure 7 summarizes the responses of participants to our post-study questionnaire. Overall, *NFTDisk* was highly rated by participants. Most of them agreed that *NFTDisk* was more useful and effective than previous methods for exploring wash trading behaviors in NFT collections. They stated that they will use it to benefit their NFT investments in the future and recommend it to others. However, two of the participants gave a rating of less than four because they found it difficult to learn such a complicated visual design and thought that *NFTDisk* was unsuitable for novice investors who had never learnt about the data on blockchain. The detailed comments and improvements from participants can be summarized as follows:

Workflow Effectiveness. Figure 7 shows that most participants provided a positive rating for Q1-Q5, which confirms the workflow effectiveness of *NFTDisk* in visually identifying and analyzing wash trading activities. Specifically, U1-U11, who are more experienced NFT investors, were able to quickly complete the tasks and interpret the patterns shown in the visualizations using their domain knowledge and investment experience, whereas novice participants (U12-U14) needed more time to interpret the visual discoveries and further explain the possible underlying reasons. U14 scored "3" for Q1 and Q3, because it took her some time and effort to fully understand the visual encodings of the *Flow Module* and *Disk Module* when analyzing the first NFT collection. But once she understood the visual designs of *NFTDisk*, she was able to easily finish the exploration and analysis of the remaining two NFT collections. Overall, with the help of *NFTDisk*, all the participants can efficiently identify the wash trading behaviors in the NFT collections and further assess the impact of them.

Visual Design and Interactions. All the ratings on the visual design and interactions are greater than a neutral score (i.e., four), as shown in Figure 7. U3 liked the design of *Disk Module* because it can help to quickly overview the suspicious transactions of the entire collection and assess their impact on prices or trading volume. U6 and U7 phrased that the *Flow Module* was quite intuitive for showing both the NFT number over time and the trading history of individual NFT, and they emphasize that the *Flow Module* can be

Workflow Effectiveness:

- Q1: I can identify the wash trading behaviors of NFT transactions using the disk module.
- Q2: I can see the temporal evolution of the NFTs of the selected addresses via the flow module.
- Q3: I can identify the NFTs of individual addresses over time via the flow module.
- Q4: I can track the trade path of individual NFTs using the flow module.
- Q5: Overall, the proposed visual analytics system help me analyze wash trading behaviors.

Visual Design and Interactions:

- Q6: The visual design of the disk module is intuitive.
- Q7: The visual design of the flow module is intuitive.
- Q8: The interactions are smooth and easy to use for analyzing NFT transactions.

Usability:

- Q9: NFT Disk is easy to learn.
- Q10: NFT Disk is easy to use.
- Q11: I would like to use NFT Disk to analyze wash trading in the future.
- Q12: I would like to recommend NFT Disk to others.

Open-ended Questions:

- Q13: What are the advantages and disadvantages of the system?
- Q14: How can the system be improved in the future?

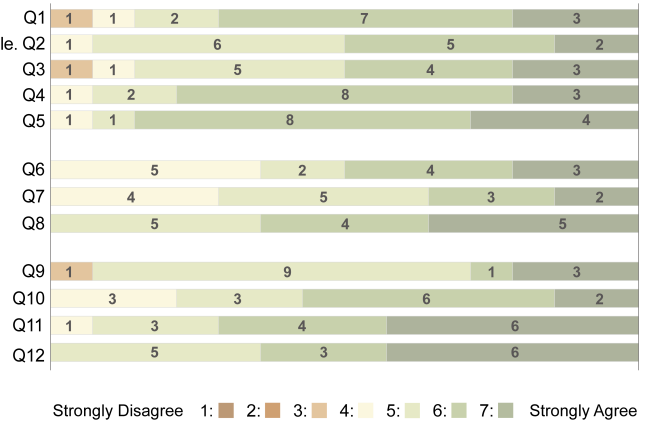


Figure 7: The results of the questionnaire in the user interview. Q1-Q12 are close-ended questions designed to evaluate the workflow effectiveness (Q1-Q5), visual design and interactions (Q6-Q8), and usability (Q9-Q12) of NFTDisk. Q13-Q14 are open-ended questions to gain the pros and cons of NFTDisk from participants. Q1-Q12 should be rated by 1-7 representing “strongly disagree” to “strongly agree” on each statement. The number of participants on each rating is represented by a horizontally stacked bar chart.

useful for other behavior analysis besides wash trading. However, five participants gave a rating of four for the design of *Disk Module* and *Flow Module* since such visualizations were more sophisticated than simple diagrams such as bar charts or line charts and took time to understand. These five participants also mentioned that it took some time (about 10 minutes) to learn visual encodings when using *NFTDisk* for the first time, but once they are familiar with the visual designs, they can perform analysis quickly. All participants agreed that our interactions were flexible and smooth, and two brush interactions can effectively filter out unwanted information. The interactions were highly rated by all participants, but some of them pointed out that the circular brush selected unwanted addresses sometimes. U8 added, “Sometimes, the circular brush will selected some undesired middle addresses if the distance between the target addresses is great.”

Usability. Overall, most participants believed that *NFTDisk* was easy to learn and use (Q9-Q10 of Figure 7). After learning the usage of *NFTDisk*, the majority of the participants would like to use it in the future and recommend it (Q11-Q12 of Figure 7). But some participants, like U12 and U13, also stated that it took them some time for novices like them to learn and use *NFTDisk*. But once they mastered the usage, they can easily analyze wash trading data and would like to use it in the future. U6 also mentioned, “*NFTDisk* is more helpful for people like NFT investors who are interested in analyzing the behaviors of addresses using blockchain data.”

Improvements. For the open-ended questions, participants provided some insightful suggestions for future improvements. U13 suggested that it may be helpful to further add some text descriptions in the user interface of *NFTDisk* to explain visual encoding and data information used by visualizations. Such information can help novice users without a background in data visualization master the usage of *NFTDisk* more easily. U4 remarked, “*NFTDisk* can facilitate the analysis of wash trading behaviors in NFT transactions. But if

we want to further analyze the detailed impact of the wash trading behaviors, we often need to include more information, such as other token transactions and account balances at these addresses.” Thus, U4 recommended us to further incorporate such information in *NFTDisk*, which is left as future work. Finally, U3 pointed out that the address filtering can be more expressive and go beyond filtering by only the amount of transactions. Instead, it can enable the filtering by more address attributes such as account balance and trading frequency, which has been implemented in the latest prototype of *NFTDisk*.

7 DISCUSSION

Our case studies and user interviews have demonstrated the effectiveness and usability of *NFTDisk*. In this section, we further discuss the lessons we learned, the specific design considerations for novice users, generalizability and limitations of our approach.

Lessons learned. During the development of *NFTDisk*, we also learned a few important lessons. First, when we are trying to detect wash trading in NFT transactions, it is usually not sufficient to focus on the transaction behaviors of only individual addresses, which is the general practice of most existing wash trading detection methods [13, 41]. Instead, we discovered that wash trading of NFTs often occurred among a group of addresses, where the NFTs held by a group of addresses remained the same while the NFTs held by each individual address were changing, such as in Figure 5D2. Thus, we have designed the *Flow Module*, which can intrinsically enable the convenient analysis of transaction behaviors within and across different groups of addresses. Second, with the help of *NFTDisk*, we also noticed that different addresses in a colluding address group for NFT wash trading may be responsible for different tasks. Some addresses were in charge of collecting NFTs, and some were dedicated to performing wash trading. Existing automatic algorithms [13, 41] can only find addresses performing wash trading

but are not able to reveal the addresses that are responsible for collecting NFTs for an NFT wash trading address group, which are both achieved in *NFTDisk* by visualizing their detailed transaction histories. Third, not all the wash trading behaviors are “harmful” to investors. When the price of NFTs in a NFT project is suddenly dropping, sometimes the owner of NFT projects may also frequently buy and sell a large number of NFTs in a short period to stabilize the NFT price. Such kind of “wash trading” behaviors by the NFT project owners are often welcomed by investors. *NFTDisk* can intuitively reveal such wash trading transactions to investors, assisting them in their decision making for NFT investments.

Design considerations for novice users. Our target users are NFT investors, but they are not necessarily familiar with data visualizations or even the technical knowledge of blockchain. During the development of *NFTDisk*, we paid special considerations into such issues and tried to make the overall visual designs and workflow as intuitive as possible. First, straightforward visualizations are proposed to represent the complex features of wash trading behaviors. For instance, we proposed the suspicious scores and explicitly encoded them by the height of inner paths, and we used the change in flow height to imply market risk. We also translated the professional technical terms in the blockchain data into concepts that general investors are more familiar with, e.g., the technical term “addresses” is rephrased as user-friendly concepts like “buyers” and “sellers”. Second, the overall design and workflow of *NFTDisk* are guided by the well-known visualization mantra of Overview-First, Details-On-Demand [30]. The *Disk Module* provides an overview of NFT transactions and reveals the suspicious addresses, and the *Flow Module* enables an in-depth investigation into these suspicious addresses with three levels of details. Such a workflow can assist even novice investors in easily identifying and analyzing the wash trading of NFT markets in a well-organized manner. In summary, target users’ knowledge background and their familiarity with visualizations should be carefully considered in the development of visual designs. Such design considerations can also apply to other real-world applications beyond the analysis of NFT wash trading.

Generalizability. *NFTDisk* is designed for visual detection of wash trading of NFT transactions. However, the workflow can be used as a reference for analyzing other frauds of cryptocurrency that show abnormal transaction patterns, such as Pump and Dumps [24] and money laundering [20]. Also, *NFTDisk* can be easily extended to the wash trading detection of stocks and funds in the traditional investment market [7], if the detailed transaction data between investors are available. Moreover, *NFTDisk* has the potential to work for the analysis of other abnormal online activities involving different participants. One possible application of *NFTDisk* can be the identification and analysis of political astroturfing on Twitter [22, 47], where a group of Twitter accounts are manipulated by their employer to distribute disinformation in a coordinated manner. The *Disk Module* can be helpful for overviewing the coordinated disinformation distribution activities of all the involved Twitter accounts, and the *Flow Module* can be used to reveal the detailed activity evolution of each account.

Limitations. In *NFTDisk*, we proposed a variety of solutions to enhance its visual scalability, such as the address ordering strategy, and the filtering and brushing interactions, which make *NFTDisk* work well for the majority of collections. When there are a huge

number of suspicious addresses and transactions that need to be checked, it will be impossible to show all of the addresses simultaneously due to the limited screen space. More processing like clustering similar addresses may be needed to further enhance the scalability. Another issue originates from the quality of the data on the blockchain. Even though all the transactions that occurred were recorded on the blockchain, there was also some information that is difficult to be identified in the transactions. For example, some NFT trading tools have emerged in the market to support bulk trading of NFTs, which can hide important information like price and cause some sales of NFTs to be identified as NFT transfers. It can be solved by combining the data record on the blockchain with the off-chain data provided by the NFT marketplaces. Moreover, a quantitative time cost comparison between using *NFTDisk* and other baseline approaches can further demonstrate the effectiveness of our approach in a more convincing way. However, we did not perform such a time cost comparison, and the major dilemma is that there are no such existing tools for investors to investigate wash trading activities. The current practice is to simply check the original transaction records, which is obviously less efficient than using our approach.

8 CONCLUSION

We propose a top-down workflow for investors to visually detect and analyze wash trading activities in NFT markets, and present *NFTDisk*, an interactive visualization comprised of two novel visualization modules: the *Disk Module* and the *Flow Module*. Specifically, the *Disk Module* is a radial visualization module with a disk metaphor to overview NFT transactions and help quickly identify potential wash trading behaviors, while the *Flow Module* is designed to demonstrate detailed NFT flow at multiple levels. A diverse set of interactions is also included to enhance the workflow. For evaluation, we describe two case studies and conduct an in-depth user interview with 14 real NFT investors. The results demonstrate that *NFTDisk* is useful and effective in exploring wash trading in NFT markets.

In future work, we plan to involve automatic approaches to facilitate the initial filtering of suspicious addresses, reducing the manual efforts of wash detection and further enhancing the effectiveness of *NFTDisk*. Moreover, we would like to incorporate more information, like the money flow of addresses, to support the identification of addresses involved in wash trading and assess its detailed impact.

ACKNOWLEDGMENTS

This work was done during Xiaolin Wen’s internship at the Singapore Management University (SMU) under the supervision of Dr. Yong Wang. This work was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant (Grant number: 20-C220-SMU-011), Lee Kong Chian Fellowship awarded to Dr. Yong Wang by SMU, and the National Natural Science Foundation of China (Grant number: 62172289). We would like to thank the participants in our user interviews and anonymous reviewers for their feedback.

REFERENCES

- [1] Mansoor Ahmed, Iliia Shumailov, and Ross Anderson. 2018. Tendrils of Crime: Visualizing the Diffusion of Stolen Bitcoins. In *Proceedings of International Workshop*

- on *Graphical Models for Security*. Springer, 1–12.
- [2] Lennart Ante. 2022. Non-fungible Token (NFT) Markets on the Ethereum Blockchain: Temporal Development, Cointegration and Interrelations. *Economics of Innovation and New Technology* (2022), 1–19.
 - [3] Hong Bao and David Roubaud. 2022. baoNon-Fungible Token: A Systematic Review and Research Agenda. *Journal of Risk and Financial Management* 15, 5 (2022), 215.
 - [4] Ziv Bar-Joseph, David K Gifford, and Tommi S Jaakkola. 2001. Fast Optimal Leaf Ordering for Hierarchical Clustering. *Bioinformatics* 17, suppl_1 (2001), S22–S29.
 - [5] Stefano Bistarelli and Francesco Santini. 2017. Go with the-Bitcoin-Flow, with Visual Analytics. In *Proceedings of the 12th International Conference on Availability, Reliability and Security*. Association for Computing Machinery, 1–6.
 - [6] Andreas Bogner. 2017. Seeing is Understanding: Anomaly Detection in Blockchains with Visualized Features. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*. 5–8.
 - [7] Yi Cao, Yuhua Li, Sonya Coleman, Ammar Belatreche, and Thomas Martin McGinnity. 2015. Detecting Wash Trade in Financial Market Using Digraphs and Dynamic Programming. *IEEE Transactions on Neural Networks and Learning Systems* 27, 11 (2015), 2351–2363.
 - [8] Sudarshan Chawathe. 2018. Monitoring Blockchains with Self-organizing Maps. In *Proceedings of 2018 17th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/12th IEEE International Conference on Big Data Science and Engineering (TrustCom/BigDataSE)*. IEEE, 1870–1875.
 - [9] Jialan Chen, Dan Lin, and Jiajing Wu. 2022. Do Cryptocurrency Exchanges Fake Trading Volumes? An Empirical Analysis of Wash Trading based on Data Mining. *Physica A: Statistical Mechanics and its Applications* 586 (2022), 126405.
 - [10] Pietro Manganelli Conforti, Matteo Emanuele, Pietro Nardelli, Giuseppe Santucci, and Marco Angelini. 2022. CryptoComparator: A Visual Analytics Environment for Cryptocurrencies Analysis. In *Proceedings of EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 43–47.
 - [11] Lin William Cong, Xi Li, Ke Tang, and Yang Yang. 2021. Crypto Wash Trading. *arXiv preprint arXiv:2108.10984* (2021).
 - [12] Lin William Cong, Xi Li, Ke Tang, and Yang Yang. 2022. Detecting and Quantifying Crypto Wash Trading. In *Proceedings of 3rd International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2021)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 10:1–10:6.
 - [13] Dipanjan Das, Priyanka Bose, Nicola Ruaro, Christopher Kruegel, and Giovanni Vigna. 2021. Understanding Security Issues in the NFT Ecosystem. *arXiv preprint arXiv:2111.08893* (2021).
 - [14] Giuseppe Di Battista, Valentino Di Donato, Maurizio Patrignani, Maurizio Pizzonia, Vincenzo Roselli, and Roberto Tamassia. 2015. Bitcoveview: Visualization of Flows in the Bitcoin Transaction Graph. In *Proceedings of 2015 IEEE Symposium on Visualization for Cyber Security (VizSec)*. IEEE, 1–8.
 - [15] Michael Dowling. 2022. Is Non-fungible Token Pricing Driven by Cryptocurrencies? *Finance Research Letters* 44 (2022), 102097.
 - [16] Felix Eigelshoven, Andre Ullrich, and Douglas A Parry. 2021. Cryptocurrency Market Manipulation: A Systematic Literature Review. In *Proceedings of International Conference on Information Systems*.
 - [17] Jean-Daniel Fekete. 2015. Reorder.js: A Javascript Library to Reorder Tables and Networks. In *Proceedings of IEEE VIS 2015*.
 - [18] Peter Fratrič, Giovanni Sileno, Sander Klous, and Tom van Engers. 2022. Manipulation of the Bitcoin Market: an Agent-based Study. *Financial Innovation* 8, 1 (2022), 1–29.
 - [19] Petra Isenberg, Christoph Kinkeldey, and Jean-Daniel Fekete. 2017. Exploring Entity Behavior on the Bitcoin Blockchain. In *Proceedings of IEEE Conference on Visualization (VIS)*. 1–2.
 - [20] Aleksandra Jordanoska. 2021. The Exciting World of NFTs: a Consideration of Regulatory and Financial Crime Risks. *Butterworths Journal of International Banking and Financial Law* 10 (2021), 716.
 - [21] Michael Jünger, Gerhard Reinelt, and Giovanni Rinaldi. 1995. The Traveling Salesman Problem. *Handbooks in Operations Research and Management Science* 7 (1995), 225–330.
 - [22] Franziska B Keller, David Schoch, Sebastian Stier, and JungHwan Yang. 2020. Political Astroturfing on Twitter: How to Coordinate a Disinformation Campaign. *Political Communication* 37, 2 (2020), 256–280.
 - [23] Christoph Kinkeldey, Jean-Daniel Fekete, and Petra Isenberg. 2017. Bitconduite: Visualizing and Analyzing Activity on the Bitcoin Network. In *Proceedings of EuroVis 2017 Eurographics Conference on Visualization, Posters Track*. 3.
 - [24] Massimo La Morgia, Alessandro Mei, Francesco Sassi, and Julinda Stefa. 2020. Pump and Dumps in the Bitcoin Era: Real Time Detection of Cryptocurrency Market Manipulations. In *Proceedings of 2020 29th International Conference on Computer Communications and Networks (ICCCN)*. IEEE, 1–9.
 - [25] James R Lewis. 1992. Psychometric Evaluation of the Post-study System Usability Questionnaire: The PSSUQ. In *Proceedings of the human factors society annual meeting*, Vol. 36. Sage Publications Sage CA: Los Angeles, CA, 1259–1260.
 - [26] Claudio DG Linhares, Bruno AN Travençolo, Jose Gustavo S Paiva, and Luis EC Rocha. 2017. DyNetVis: A System for Visualization of Dynamic Networks. In *Proceedings of the Symposium on Applied Computing*. Association for Computing Machinery, 187–194.
 - [27] Dan McGinn, David Birch, David Akroyd, Miguel Molina-Solana, Yike Guo, and William J Knottenbelt. 2016. Visualizing Dynamic Bitcoin Transaction Patterns. *Big Data* 4, 2 (2016), 109–119.
 - [28] Robert Norvill, Beltran Borja Fiz Pontiveros, Radu State, and Andrea Cullen. 2018. Visual Emulation for Ethereum’s Virtual Machine. In *Proceedings of NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 1–4.
 - [29] Patrick Ocheja, Brendan Flanagan, Hiroaki Ogata, and Solomon Sunday Oyeler. 2022. Visualization of Education Blockchain Data: Trends and Challenges. *Interactive Learning Environments* (2022), 1–25.
 - [30] Ben Shneiderman. 2003. The Eyes have It: A Task by Data Type Taxonomy for Information Visualizations. In *The Craft of Information Visualization*. Elsevier, 364–371.
 - [31] Jinha Song, Jongho Nang, and Juwook Jang. 2018. Design of Anomaly Detection and Visualization Tool for IOT Blockchain. In *Proceedings of 2018 International Conference on Computational Science and Computational Intelligence (CSCI)*. IEEE, 1464–1465.
 - [32] Dong Sun, Zezheng Feng, Yuanzhe Chen, Yong Wang, Jia Zeng, Mingxuan Yuan, Ting-Chuen Pong, and Huamin Qu. 2020. DFSeer: A Visual Analytics Approach to Facilitate Model Selection for Demand Forecasting. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [33] Yujing Sun, Hao Xiong, Siu Ming Yiu, and Kwok Yan Lam. 2019. Bitvis: An Interactive Visualization System for Bitcoin Accounts Analysis. In *Proceedings of 2019 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, 21–25.
 - [34] Yujing Sun, Hao Xiong, Siu Ming Yiu, and Kwok Yan Lam. 2022. BitAnalysis: A Visualization System for Bitcoin Wallet Investigation. *IEEE Transactions on Big Data* (2022), 1–16.
 - [35] Tri Sundara, Ideva Gaputra, and Siska Aulia. 2017. Study on Blockchain Visualization. *International Journal on Informatics Visualization (JOIV)* 1, 3 (2017), 76–82.
 - [36] Syed Ahzam Tariq and Imtiaz Sifat. 2022. Suspicious Trading in Nonfungible Tokens (Nfts): Evidence from Wash Trading. *Available at SSRN 4097642* (2022).
 - [37] Natkamon Tovanich, Nicolas Heulot, Jean-Daniel Fekete, and Petra Isenberg. 2019. Visualization of Blockchain Data: A Systematic Review. *IEEE Transactions on Visualization and Computer Graphics* 27, 7 (2019), 3135–3152.
 - [38] Natkamon Tovanich, Nicolas Soulié, Nicolas Heulot, and Petra Isenberg. 2021. MiningVis: Visual Analytics of the Bitcoin Mining Economy. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 868–878.
 - [39] Stef van den Elzen, Danny Holten, Jorik Blaas, and Jarke J van Wijk. 2013. Dynamic Network Visualization Withextended Massive Sequence Views. *IEEE Transactions on Visualization and Computer Graphics* 20, 8 (2013), 1087–1099.
 - [40] Friedrich Victor and Andrea Marie Weintraud. 2021. Detecting and Quantifying Wash Trading on Decentralized Cryptocurrency Exchanges. In *Proceedings of the Web Conference 2021*. Association for Computing Machinery, 23–32.
 - [41] Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, and Omri Ross. 2022. NFT Wash Trading: Quantifying Suspicious Behaviour in NFT Markets. *arXiv preprint arXiv:2202.03866* (2022).
 - [42] Qin Wang, Rujia Li, Qi Wang, and Shiping Chen. 2021. Non-fungible Token (NFT): Overview, Evaluation, Opportunities and Challenges. *arXiv preprint arXiv:2105.07447* (2021).
 - [43] Gavin Wood et al. 2014. Ethereum: A Secure Decentralised Generalised Transaction Ledger. *Ethereum Project Yellow Paper* 151, 2014 (2014), 1–32.
 - [44] Jia-zhi Xia, Yu-hong Zhang, Hui Ye, Ying Wang, Guang Jiang, Ying Zhao, Cong Xie, Xiao-yan Kui, Sheng-hui Liao, and Wei-ping Wang. 2020. SuPoolVisor: A visual Analytics System for Mining Pool Surveillance. *Frontiers of Information Technology & Electronic Engineering* 21, 4 (2020), 507–523.
 - [45] Laixin Xie, Ziming Wu, Peng Xu, Wei Li, Xiaojuan Ma, and Quan Li. 2022. Role-Seer: Understanding Informal Social Role Changes in MMORPGs via Visual Analytics. In *Proceedings of CHI Conference on Human Factors in Computing Systems*. 1–17.
 - [46] Xuanwu Yue, Xinhuan Shu, Xinyu Zhu, Xinnan Du, Zheqing Yu, Dimitrios Papadopoulos, and Siyuan Liu. 2018. Bitextract: Interactive Visualization for Extracting Bitcoin Exchange Intelligence. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 162–171.
 - [47] Jerry Zhang, Darrell Carpenter, and Myung Ko. 2013. Online astroturfing: A theoretical perspective. (2013).
 - [48] Zengsheng Zhong, Shuirun Wei, Yeting Xu, Ying Zhao, Fangfang Zhou, Feng Luo, and Ronghua Shi. 2020. SilkViser: A visual Explorer of Blockchain-based Cryptocurrency Transaction Data. In *Proceedings of 2020 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 95–106.