

A work project, presented as part of the requirements for the Award of a Master's degree in
Business Analytics from the Nova School of Business and Economics.

Segmenting Economic Agents on the XRP Ledger: An Unsupervised Learning
Approach
MARKUS GIESBRECHT

Work project carried out under the supervision of:

Prof. Leid Zejnilovic

17/12/2024

Abstract

XRP is among the top 5 most prevalent cryptocurrencies as of 2024. However, the pseudonymous nature of XRP Ledger transactions complicates the understanding of economic activities and regulatory oversight of fraud within the network. This study uses descriptive statistics to develop heuristics for categorizing distinct types of economic agent activities and leverages unsupervised machine learning to segment agents into clusters; successfully distinguishing accounts such as decentralized exchanges, gambling sites, and NFT-related entities. Additionally, supervised fraud detection models are trained with an off-chain dataset of accounts involved in spam, token theft, and Ponzi schemes, achieving fraud detection accuracies of over 90%.

Keywords

Behavioral Analysis, Blockchain, Fraud Detection, Heuristics, Ledgerlytics, Machine Learning, Ripple, Segmentation, Supervised, Unsupervised, XRP Ledger

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

1 Introduction

Since its introduction in 2009, Bitcoin has significantly influenced the global monetary system, leading to the emergence of almost ten thousand competing cryptocurrencies on the market (Statista, n.d.; Alahmad et al. 2023). The XRP token, as one of these currencies, was introduced by Ripple Inc. in 2012 and is widely recognized as one of the most prominent alternatives to Bitcoin. XRP ranks in the top 5 digital currencies worldwide with a market capitalization of over 100 billion USD as of November 2024 (Ahmadova and Ereik 2022; CoinMarketCap 2024).

The XRP network mainly aims to serve banks and financial institutions by providing a decentralized alternative to traditional cross-border payment systems such as the Society for Worldwide Interbank Financial Telecommunication (SWIFT), offering superior transaction speed and reduced costs for currency exchanges (Sergeenkov 2024). Additionally, the network supports key applications, including tokenization, central bank digital currencies, and stablecoins, as well as Non-Fungible Tokens (Ripple, n.d.-b).

While the XRP Ledger, as the decentralized ledger technology behind its token, operates as a publicly accessible database, the information recorded on-chain alone is insufficient to trace users back to their real-world counterparts. In blockchain, participants transact assets using unique alphanumeric addresses. These addresses lack a direct connection to real-world identities or individual agents, a trait known as pseudonymity (Hellwig, Karlic, and Huchzermeier 2020).

This characteristic inherently results in a lack of transparency and uncertainty regarding the types of economic activities in blockchain networks such as the XRP Ledger. The limited traceability of transactions also raises concerns about the potential misuse of the technology for illicit purposes. Illegal crypto activities such as scams, stolen funds, and illicit trade accounted for a total estimate of 39.6 billion USD in 2022 (Chainalysis 2024). As a result, stakeholders – including governments, industry leaders, and financial institutions – require a shared

understanding of the participants within blockchain ecosystems to foster sustainable growth and address the vulnerabilities inherent to the networks (Chainalysis, n.d.).

Addressing these challenges, this study seeks to analyze agents on the XRP Ledger to characterize the economic activities within the ecosystem and aims to detect accounts involved in selected fraud schemes. The main objectives of the study are:

1. To explore the applicability of account-based heuristics in characterizing different economic agent categories and their behavior on the XRP Ledger.
2. To evaluate the effectiveness of unsupervised machine learning in segmenting economic agents within the XRP Ledger.
3. To assess the ability of supervised machine learning to identify accounts that are involved in fraudulent activity on the XRP Ledger.

To the best of our knowledge, no similar research has been conducted on XRP, leading to uncertainties about the applicability of existing methods to gain insights into the network. While modeling techniques, such as machine learning models have been applied with similar aims to Bitcoin and Ethereum, their relevance and effectiveness on other ecosystems remains unclear, as ledger structure and use cases differ between different ledgers. By addressing this research gap, the primary goal of this work is to test the applicability of the selected research methods on XRP by adapting them to the specific properties of the networks' structure as well as the available off-chain data. The secondary goal is to provide insights into the implications of our results for the XRP ecosystem and potential implementation scenarios for our fraud detection classifiers.

The relevance to apply this research on the XRP Ledger arises from its increasing importance in financial markets. The lack of studies in this area creates ambiguity regarding the network's usage, activity types, and techniques to mitigate fraudulent behavior. This presents a challenge for regulators, compliance bodies, investors, financial institutions, and the broader

XRP ecosystem in assessing the network's reliability and its participation in economically significant activities within markets. The segmentation of agents and detection of fraud in the network has the potential to implement effective ecosystem monitoring and demonstrate XRP's utility and diverse applications to stakeholders, including investors, financial institutions, and researchers. The potential of the utilized methods also lies in monitoring regulatory compliance, especially in jurisdictions with oversight of typical high-risk activities. Ripple can leverage the findings of our work to optimize the network for its most valuable segments and manage funded partnerships through data-driven decisions. Eventually, the development of a model capable of detecting fraudulent accounts is expected to benefit regulatory and prosecutorial bodies by facilitating automated real-time detection of illicit agents, thereby strengthening compliance and enforcement mechanisms. The trust of users, particularly institutional stakeholders, is expected to depend on XRP being a secure payment environment where fraudulent activities are effectively identified and mitigated. Ultimately, we expect this work to contribute to greater transparency in a financially significant blockchain network, fostering trust, supporting informed decision-making, and enhancing confidence in its integrity.

2 The XRP Ledger

This section introduces the XRP ecosystem as the primary domain of this study. The XRP Ledger (XRPL) is a blockchain solution and distributed ledger technology (DLT) developed and first launched in 2012 by Ripple Inc (Ripple, n.d.-c). A blockchain is a decentralized public ledger that records transactions within a peer-to-peer network. It consists of linked data blocks, each referencing the previous one. The XRP network relies on trusted validator nodes to reach consensus on data validity, enabling secure and integer transactions without a central authority (Ripple 2024b; European Union Agency for Cybersecurity (ENISA), n.d.).

The XRPL leverages its native cryptocurrency, XRP. Unlike Bitcoin (BTC), which requires mining to create new coins, all 100 billion XRP tokens were created at its inception, ensuring a fixed supply from the beginning (Alahmad et al. 2023; Li and Whinston 2020; Ripple 2024b). As of November 2024, XRP's market capitalization exceeded \$100 billion, with each token valued at around 1.80 USD (CoinMarketCap 2024). Ripple Inc., is the company and issuer behind the network and token, holding about 48 billion XRP in escrow (2024), with a monthly release of up to 1 billion tokens to stabilize XRP supply and prices (Sergeenkov 2024). The company has partnered with over 200 financial institutions, including major and central banks, to integrate the XRPL network into their systems (Rella 2020; Ahmadova and Ereik 2022).

Designed to enable seamless global value transfers, the XRPL primarily functions as a distributed payment system aimed at improving cross-border payment efficiency, offering an alternative to traditional solutions like SWIFT (Chase and MacBrough 2018; Sergeenkov 2024). While widely adopted, SWIFT transactions can be slow and costly due to intermediary fees and the need for pre-funded foreign accounts. In contrast, the XRPL supports rapid, cost-effective transactions across multiple currencies through features like trust lines and an on-chain decentralized exchange (Rella 2020; Peduzzi, James, and Xu 2021). By eliminating intermediary fees and the need for pre-funded foreign accounts, adopting the XRPL can reduce operational costs for banks by up to 60% (Sergeenkov 2024).

Beyond payments, the XRPL network supports key applications, including tokenization, central bank digital currencies (CBDCs), and stablecoins, as well as Non-Fungible tokens (NFTs) (Ripple, n.d.-b). In comparison, other ecosystems like BTC primarily serve as digital assets, while networks such as Ethereum (ETH) focus on decentralized applications powered by smart contracts (Alahmad et al. 2023).

Ultimately, the XRPL also differs from other blockchain technologies in its consensus mechanism and fee system. The XRPL employs the Ripple Protocol Consensus Algorithm (RPCA), where trusted validators from a Unique Node List (UNL) vote on transaction validity. This process confirms transactions within seconds, enabling fast and efficient processing. In comparison, BTC's Proof of Work (PoW) mechanism relies on miners solving complex problems, resulting in an average block time of ten minutes. ETH transitioned from the PoW in 2022, to a less energy-intensive approach called Proof of Stake (PoS) that uses coin ownership for validation (Mauri, Cimoto, and Damiani 2018; Alahmad et al. 2023). Thereby, the XRPL outperforms other blockchains in terms of transaction speed and fees, processing around 1,500 transactions per second with costs of approx. 0.00001 XRP (depending on network load) per transaction, which represents only a fraction of a dollar cent. (Mauri, Cimoto, and Damiani 2018; Alahmad et al. 2023; Sergeenkov 2024)

The XRPL also differs in its data and ledger structure. It uses an account-based system that directly tracks each account's balance, unlike BTC's Unspent Transaction Output (UTXO) model, where balances are spread across multiple addresses and transactions involve multiple inputs and outputs. This simplifies transaction validation, as XRPL accounts can spend part of their balance while retaining the rest, using a single address repeatedly for both sending and receiving funds. (Mauri, Cimoto, and Damiani 2018; Alahmad et al. 2023; Akcora, Gel, and Kantarcioglu 2022)

Eventually, each XRP transaction includes detailed information such as the sender and receiver addresses, transaction amount, fee, transaction type (from over 40 types), and additional metadata such as source tags for user identification, optional memos for supplementary information, and flags or settings related to the transaction. (Ripple, n.d.-a; 2024a)

3 Research Context and Related Studies

This section reviews related works on characterizing economic activity types, segmenting accounts, and automating the detection of fraudulent users on DLTs. Section 3.1 provides an overview of agent segmentation, with Sections 3.1.1 and 3.1.2 focusing on the applications of heuristics and machine learning techniques, respectively. Section 3.2 examines fraud in blockchain ecosystems and highlights previous studies on the automated detection of fraudulent accounts with supervised machine learning techniques.

3.1 Segmenting Economic Agents on DLTs

The segmentation of agents on DLTs has been a prominent research topic in recent years. Thereby, economic agents can be defined as individuals, organizations, or entities engaged in economic activities, including decisions related to the production, distribution, and consumption of goods and services (United Nations, n.d.). Reviewing related works reveals that the definitions of economic agent segments in blockchain networks are highly adaptable, often influenced by the specific use cases of the ledger analyzed. It is reasonable to expect that the primary use cases of the XRPL align with its key applications, as outlined in Section 2. However, since no prior efforts have segmented users within its ecosystem, insights from other blockchains offer valuable guidance for identifying and segmenting XRPL agents. *Chainalysis* categorizes economic entities within blockchain ecosystems, such as services and organizations, based on their specific use cases of the technology (Table 1).

Category	Description
Merchant Services	Businesses that accept cryptocurrency as payment for goods or services, functioning similarly to traditional payment processors.
Hosted Wallets	Custodial services where the provider manages the security of users' public and private keys, making transactions easier but with less financial privacy.
Mining Pools	Groups that combine computing power to mine cryptocurrency, often considered low-risk as they primarily earn through mining rewards.
Exchanges	Platforms where users can buy, sell, or trade cryptocurrency. Includes both centralized (custodial) and decentralized (non-custodial) exchanges.
Nested Services	Services operating within larger exchanges, such as instant exchangers and OTC brokers, which can pose risks if compliance standards are lax.

Gambling Services	Platforms for betting and gaming. Risk level varies by jurisdiction; can be used for legitimate purposes or as a means for money laundering.
Cyber Infrastructure	Providers of web hosting, VPNs, and other internet-based services that accept cryptocurrency, sometimes offering anonymity features, leading to medium risk categorization.
Mixers	Services that enhance transaction privacy by mixing funds to obscure their origin and destination. Frequently scrutinized for potential use in money laundering.
Darknet Markets	Online platforms for trading illegal goods and services, accessible via anonymity networks like Tor. High-risk due to involvement in illicit activities.
Illicit Actor Organizations	Groups associated with activities such as ransomware, fraud, terrorism, or money laundering. Includes organizations indirectly linked to illegal activities.

Table 1: Chainalysis Blockchain Segments (Chainalysis, n.d.)

The categories include businesses accepting cryptocurrencies (merchant services), custodial wallets (hosted wallets), mining pools, centralized or decentralized trading platforms (exchanges), nested services within exchanges, gambling sites, cyber-infrastructure providers, privacy-enhancing mixers, darknet markets, and entities linked to illegal activities. Thereby, *Chainalysis* does not provide any ledger-specific information, it mainly focuses on categorizing actors without diving into specific mechanisms unique to individual ledger ecosystems.

However, the diversity of blockchain ecosystems necessitates segmentation approaches tailored to the specific characteristics of each ledger. Blockchain use cases and ledger mechanisms differ significantly as mentioned in Section 2. The following subsections will present related works and their approaches in characterizing and identifying agents specific to unique ledger technologies.

3.1.2 Machine Learning Approaches

This section examines related research on supervised and unsupervised agent segmentation on blockchain networks, specifically BTC and ETH, providing a foundational framework for our study. Supervised approaches rely on labeled datasets to train models for classifying agents, whereas unsupervised methods detect patterns and group agents without predefined labels.

Supervised Detection on BTC and ETH Networks

An extensive literature review reveals that current research on agent segmentation using machine learning approaches predominantly focuses on BTC and ETH. Whereby, early studies employed basic heuristic methods, other research leverages these heuristics for supervised machine learning models or applies unsupervised techniques.

Lin et al. (2019) leveraged eight different supervised models and a dataset of 26,308 labeled addresses aiming to improve address classification in the BTC ecosystem, achieving a Micro-F1 score of 87% and a Macro-F1 score of 86% using LightGBM. Addresses were labeled using heuristic methods and the classification categories were: Exchange, Faucet, Gambling, HYIP, Market, Mixer, and Pool. They reflect typical behaviors identified through transaction features such as transaction frequency, received/spent ratios, and balance variances. As each category was associated with distinct economic activities, their economic relevance could be highlighted within the ecosystem. Gambling sites and exchanges – for instance – accounted for the largest categories in terms of the number of transactions, indicating their significant role in the network.

Several other studies follow the same idea of utilizing supervised models by leveraging off-chain labeled data showing the importance of data triangulation and real-world verification of identified patterns and cluster characteristics. Harlev et al. (2018) employed data from *Chainalysis*, which had previously clustered and categorized BTC addresses. The domains were segmented into ten predefined categories, including exchanges, gambling services, hosted wallets, merchant services, mining, scams, and others. The authors stated that using a supervised learning approach could successfully categorize entities, with potential applications for crime investigations and financial compliance, enhancing the transparency of BTC transactions. They achieved a model accuracy of 78%. Yet, the authors highlight that the approach's success depends heavily on the quality and balance of the training data, particularly due to under-represented categories like hosted wallets and mixing services.

Other papers also integrate graph-based methodologies into their approaches, demonstrating their utility in analyzing relationships and transaction patterns (Yu et al. 2023). However, the details of these graph-based approaches fall beyond the scope of this section.

In conclusion, identifying economic agents and their activity and thereby enhancing transparency of blockchain networks seems partially doable through supervised machine learning algorithms. Yet, this implies access to labeled off-chain data that is often more accessible with larger networks like ETH or BTC. As far as we know, smaller ecosystems like XRP lack sufficient community or service-based data resources, posing a significant challenge for supervised learning approaches.

Unsupervised Detection on BTC and ETH Networks

A review of the literature reveals that studies on unsupervised agent segmentation are significantly less common than those focusing on supervised approaches. Payette, Schwager, and Murphy (2017) analyzed a dataset of 250,000 ETH addresses, including their balances and complete transaction histories. They applied three clustering techniques: K-means, hierarchical clustering (or agglomerative clustering), and Birch clustering. K-means was ultimately selected as the preferred method due to its computational efficiency and ability to handle large datasets effectively. The features used for clustering included transaction counts, average transaction values, and monthly activity metrics, capturing both incoming and outgoing transactions as well as their associated values in Ether and USD. The Calinski-Harabasz score was used to evaluate the quality of clustering, and the elbow method was used to estimate the optimal number of clusters. However, the authors did not label the clusters or verify them with external data, limiting the ability to draw definitive conclusions about the specific roles or identities of the addresses, as well as the overall reliability of the clustering results in a real-world context.

Birrane (2017) explored unsupervised analysis of the ETH network, utilizing K-means clustering to group over 31,000 ETH addresses with at least 100 transactions. This study

discovered 15 clusters, some dominated by high-activity which were assumed to be users like miners or smart contract operators, and others comprising low-activity accounts, illustrating the diverse ecosystem of ETH. The clustering was based on features that captured the transactional behavior of addresses, such as their in- and out-degrees, transaction counts, and the mean and standard deviation of transaction values, both overall and specifically for interactions with smart contracts. However, the lack of labeled datasets limited their ability to validate these clusters.

A different study by Ermilov, Panov, and Yanovich (2017) tackled this issue by utilizing off-chain data on BTC addresses to guide their clustering. Instead of machine learning, they employed a probabilistic model that optimizes cluster formation based on log-likelihood. The off-chain information came from more than 90 sources (e.g., *Twitter*, *walletexplorer.com*), resulting in over 20 million tags categorized into six groups: mining pools, exchanges, Darknet markets, mixers, gambling, and other services. This method effectively used external data to improve clustering accuracy, reducing errors from heuristic-based methods.

A more recent approach by Vlahavas, Karasavvas, and Vakali (2024) clustered BTC transactions rather than accounts or addresses. They used graph-based features and a trimmed K-means clustering algorithm, achieving a Silhouette score of 0.78. Their attempt to validate the results using external data (addresses categorized by type) was only partially successful, as they were unable to effectively differentiate between exchanges, gambling platforms, and service accounts.

In conclusion, while significant progress has been made in segmenting economic agents on blockchains like BTC and ETH, the XRPL remains unexplored in this regard. Most research relies on supervised methodologies using labeled datasets from off-chain verification, successfully identifying actors such as exchanges, gambling services, and mining pools (a detailed overview of the mentioned papers can be found in Table 6). Unsupervised approaches,

though less common, have shown potential by clustering economic agents using transaction patterns and graph-based features. However, verifying these clusters with ground truth data has been rare, limiting their reliability and broader applicability. This study closes this research gap by applying unsupervised clustering techniques to the XRPL and verifying the resulting clusters with off-chain data, ensuring greater accuracy and validation.

4 Methodology

This section provides a concise overview of the proposed methods and model setups, as well as the data utilized in this work, including preprocessing and feature engineering steps. Subsection 4.1 presents an overview of the main datasets as well as the features engineered. Subsections 4.2 to 4.4 provide detailed descriptions of the data subsets utilized, the labeled ground truth data, feature selection processes, and the proposed heuristics and machine learning models as well as their evaluation methods.

4.1 Data and Feature Engineering

The transactional dataset for this study comprises over 2 billion XRPL transactions spanning 36 million ledgers (ledger 50,000,000 to 86,000,000), collected in a decentralized manner over 1,618 days (from 2019-09-13 to 2024-02-16). The main source of this data is a publicly available dataset with transactions from a full history ripple node provided by XRPL Labs developer *Wietse Wind* (Wind 2024).

During preprocessing, only successful transactions were retained to capture actual value transfers and avoid double-counting, as failed transactions are often retried. The XRPL supports over 40 transaction types, many of which relate to account settings or niche applications (Ripple, n.d.-a). To ensure a focus on economically significant activities, the analysis included transaction types such as *Payment*, *OfferCreate*, *OfferCancel*, and *TrustSet*. Additionally,

escrow-related transactions (*EscrowCreate*, *EscrowCancel*, etc.) and specific NFT operations (*NFTokenCreateOffer*, *NFTokenAcceptOffer*, *NFTokenMint*, etc.) were included to account for value-transfer and contract-related activities. The selected transaction types account for over 99% of all transactions in the dataset. Other types, such as *AccountDelete*, *AccountSet*, and *TicketCreate*, were excluded as they do not represent asset movement or economic activity.

Feature Engineering

A total of 55 features were created on account level and are presented in Table 8 in the Appendix. The foundation for engineering these features are related studies discussed in Section 3, along with features specifically adapted to the characteristics of the XRPL. The features involve general transaction metrics that provide insights into overall account activity, including transaction count, active periods, and value movements, such as total amounts sent, received, and net balance changes. Payment ratios, divided into outgoing and ingoing, reflect how frequently accounts send or receive payments. Offer-related ratios and trust set & escrow ratios capture account involvement in creating offers or managing trust lines and escrows. NFT ratios, based on NFT related transaction types, describe activities related to NFT creation and trading. Tag presence ratios show how often tags are used, indicating interactions with platform-based entities that use tags to link transactions to end users. Finally, interaction metrics, such as unique transaction partners and ratio of interaction with the 20 largest exchanges, assess the diversity and characteristics of an account's interactions.

Aggregating XRP transaction data at the address level posed several challenges. The high volume of transactions made computations demanding, while the diversity of transaction types added complexity in defining consistent features. This was addressed by focusing on the most relevant transaction categories, as outlined earlier. Also, significant variance in activity levels between highly active accounts and smaller accounts with minimal transactions posed

challenges in drawing meaningful insights across all accounts equally. To address this, ratio-based features were introduced to normalize activity levels and ensure account comparability. Additionally, data incompleteness and missing values created complications, particularly due to missing information on transaction amounts for non-XRP token payments. To address this, token-specific features were developed to distinguish between XRP and non-XRP transactions, ensuring a clearer data representation.

Ultimately, segmenting agents and identifying fraud required slightly different feature configurations. Therefore, different feature combinations were selected for each method to achieve optimal results, as detailed in the following subsections of Section 4.

Ground Truth Data

To aid in developing agent heuristics and interpreting clusters, an additional dataset containing labeled addresses was created, referred to as the ‘known domains dataset’ in this work. The primary sources for this dataset included a comprehensive collection of account names along with their associated web domains and *X* (formerly *Twitter*) handles, obtained through the *XRPScan* API. Additionally, it incorporated a dataset of domains published by address holders on the XRPL in the *verified domain* transaction field, gathered from the *Bithomp* website. (<https://xrplexplorer.com/en/domains>, as of October 6, 2024). Most of the addresses provided by *XRPScan* were manually verified by sending small test transactions to confirm their authenticity. The domains were manually labeled according to the categories in Table 2. Domains with unclear distinctions, such as those combining exchange activities with wallet or NFT marketplace hosting, as well as domains that were untraceable online or apparently linked to fraudulent activities, were excluded from the data.

Category	Definition
Exchanges	Addresses associated with centralized or decentralized platforms used for trading digital assets.
Gambling/Gaming	Addresses linked to platforms that provide gambling or gaming services, including casinos, online betting, and gaming with potential financial stakes.

NFT	Addresses connected to Non-Fungible Token (NFT) creation, sales, or related platform activities.
Services	Addresses belonging to businesses that provide various services, such as financial services like payment providers, loan services, or wallet providers, as well as non-financial utilities like consulting services.
Bridging Services	Addresses associated with platforms or entities to transfer assets or tokens between different blockchains or networks.
Issuer	Addresses that issue tokens or act as a source for specific assets on the XRPL. (These addresses may belong to entities involved in exchange, gambling, NFT, or other services.)
Other	Addresses that do not fit into the above categories, such as donation accounts of non-profit organizations, communities, and sellers of physical goods.

Table 2: Categories of economic agents, used as labels in known domains dataset

A separate dataset containing fraudulent addresses, provided by *XRP Forensics*, was used for the supervised fraud detection models and will be referred to as the ‘fraud dataset’ in this work. The dataset includes 12,351 fraudulent XRP accounts, consisting of a combination of flagged and auto-traced cases. The accounts were involved in a wide range of fraudulent activities, primarily consisting of addresses related to *PlusToken* fraud cases (32.27%), spam (27.19%), theft (19.31%), and giveaway scams (17.32%) as well as a minority of other types, such as token scams and different *Ponzi* schemes. Definitions of these fraud types were presented in Section 3.2. The dataset does not include other forms of illegal activity, such as involvement in the trade of illegal goods, money laundering, or terrorist financing.

4.3 Unsupervised Agent Segmentation

The goal of this study's second part is to segment economic agents on the XRPL using unsupervised clustering techniques. Unlike supervised methods, which require extensive labeled data, unsupervised clustering detects patterns in data without prior labels, making it an ideal method for exploring agent activities on the XRPL. By uncovering hidden groupings, clustering provides insights into different types of agents and their behaviors, even in the absence of detailed prior knowledge about the network.

To date, no research – neither supervised nor unsupervised – has been conducted to segment agents on the XRPL. Our approach is not only the first to apply unsupervised clustering

on this ledger but also introduces a novel element: verifying the resulting clusters against off-chain labeled data. This validation ensures the reliability of our findings and allows for meaningful interpretations of the clusters.

This methodology was inspired by a combination of previous studies. Payette, Schwager, and Murphy (2017) utilized clustering to explore ETH transaction patterns, while Ermilov, Panov, and Yanovich (2017) incorporated off-chain data to improve heuristic-based clustering accuracy on the BTC blockchain. Building on these insights, we aim to adapt and expand these methods to address the unique challenges posed by the XRPL, ultimately contributing to a deeper understanding of agent behavior in decentralized ecosystems.

Data

To begin the analysis, the account-based feature dataset was filtered to include only addresses with more than 10 transactions. This filtering step ensured the focus remained on active accounts, reduced noise from infrequent users, and provided sufficient data for meaningful statistical analysis. A subsample of 40,000 addresses was then created, which included the 2,310 verified and labeled addresses of the known domains dataset. This subsample size was selected to balance computational feasibility, as larger datasets resulted in prohibitively long runtimes.

To refine the feature set, an exhaustive recursive feature elimination was conducted for each of the three algorithms, with the Silhouette score used as the evaluation metric. Recursive feature elimination is an iterative method that trains a model, evaluates feature importance, and eliminates the least important features, ultimately identifying the most informative subset.

Table 12 and Table 13 in the Appendix show the remaining features for each algorithm. Figure 1 shows the methodological framework for the unsupervised agent segmentation.

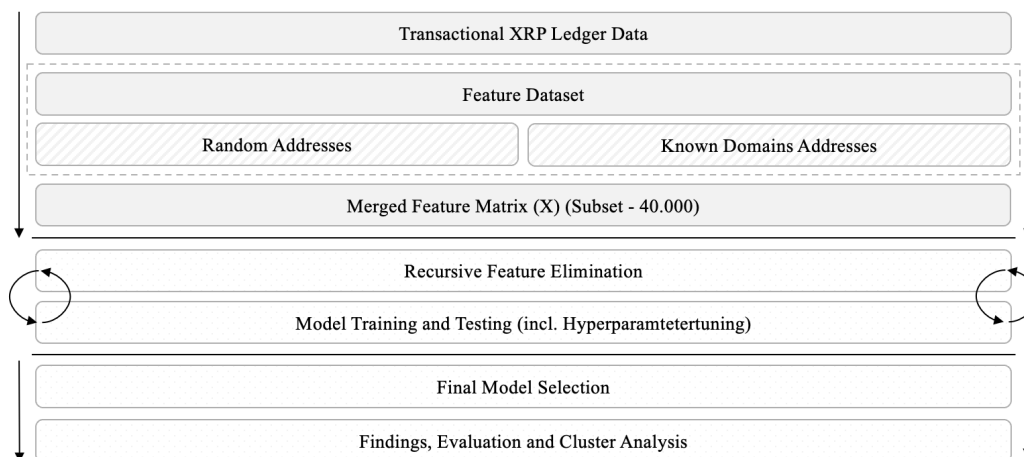


Figure 1: Methodological framework for unsupervised agent segmentation

Proposed Algorithms

The following section describes the algorithms that were employed to cluster addresses on the XRPL. Other algorithms like BIRCH clustering, CLARA, and DBSCAN were initially considered but not included in this work due to limited performances. The clusters resulting from these models failed to show meaningful and interpretable separations.

Model Types and Algorithms

1. Agglomerative Clustering

Our baseline model is an agglomerative clustering algorithm. In Scikit-Learn, this method adopts a bottom-up hierarchical approach, starting with each data point in its own separate cluster. Clusters are merged successively based on their pairwise distances until the desired number of clusters is formed. This process builds a hierarchy of clusters step by step, making it possible to identify groupings at different levels of granularity. The merging strategy is guided by linkage criteria, which determine how the distance between clusters is calculated. The hierarchical nature of agglomerative clustering enables visualization through dendrograms, providing a clear representation of how clusters are formed. This method is particularly suited

for exploring the natural clustering structure of the data, as it does not require the number of clusters to be predefined (Scikit-Learn, n.d.-e).

2. Gaussian Mixture

For our next model, we used a Gaussian Mixture Model (GMM), which is a probabilistic method for clustering that presumes the data originates from a combination of multiple Gaussian distributions. GMM supports flexible cluster shapes and uses soft clustering, where each observation is assigned a probability of belonging to each cluster. The implementation in Scikit-Learn relies on the Expectation-Maximization (EM) algorithm, which iteratively computes the parameters of the Gaussian distributions and determines probabilities. The number of clusters can be optimized using criteria like the Bayesian Information Criterion (BIC) or Akaike Information Criterion (AIC). Additionally, GMM supports various covariance types, which allow it to adapt to different cluster shapes and orientations (Scikit-Learn, n.d.-d).

3. K-Means

K-Means is a popular clustering technique that divides data into a specified number of clusters by minimizing the total squared distance between data points and the centroids of their assigned clusters. The algorithm iteratively updates the cluster centroids and reassigns observations until convergence, ensuring cluster assignments stabilize. In Scikit-Learn, users must define the number of clusters beforehand, and approaches like the Elbow Method or Silhouette score can assist in selecting the most appropriate number. K-Means is computationally efficient and straightforward to implement, with centroids providing a concise summary of cluster characteristics (Scikit-Learn, n.d.-c).

4. Principal Component Analysis

In addition to running these algorithms on the full feature space, Principal Component Analysis (PCA) is introduced as a modeling step before executing the clustering algorithm. PCA is a method for reducing dimensionality in datasets, designed to simplify high-dimensional data while retaining as much variability as possible. It works by transforming the data into a new set of variables, known as principal components, which are derived as linear combinations of the original variables. These components are ranked based on the amount of variance they capture, with the first principal component accounting for the most variance and each subsequent component capturing less. By focusing only on the most important components, it helps reduce complexity and computational costs while minimizing information loss (Scikit-Learn, n.d.-b).

Model Evaluation

To evaluate clustering quality, internal measures like the Silhouette score, Calinski-Harabasz index, and Davies-Bouldin index were used to assess cluster compactness and separation. These metrics helped identify the most suitable clustering model.

The Silhouette score quantifies how similar a data point is to the other points in its cluster compared to those in other clusters, producing values between -1 and 1. Higher scores indicate tightly grouped points and well-separated clusters. This metric is particularly useful for evaluating models with clearly defined cluster boundaries (Rousseeuw 1987).

The Calinski-Harabasz index, or Variance Ratio Criterion, examines the proportion of dispersion between clusters to the dispersion within clusters. Higher values suggest compact, well-separated clusters with distinct boundaries (Caliński and Harabasz 1974).

The Davies-Bouldin index estimates the degree of resemblance between each cluster and the one it most closely resembles, based on dispersion. Lower values indicate better-defined clusters with minimal overlap (Davies and Bouldin 1979).

To evaluate alignment with known labels from the known domains dataset, external metrics like the Adjusted Rand index (ARI) and Normalized Mutual Information (NMI) were used. ARI measures the similarity between clustering results and ground truth labels while adjusting for random agreement, with values ranging from -1 (no agreement) to 1 (perfect match). NMI quantifies the shared information between clusters and ground truth labels, with values between 0 and 1, where higher values indicate better alignment (Shannon 1948).

The final model and optimal number of clusters were selected using a combination of these metrics. High Silhouette and Calinski-Harabasz scores, along with low Davies-Bouldin scores, indicated compact and distinct clusters. For external validation, models with high ARI and NMI values were prioritized to ensure strong alignment with the labeled data.

The final cluster analysis examined patterns relative to the ground truth of the best-performing model, revealing whether agent types clustered together or were effectively separated. By calculating the mean and median of key features as well as visualizing the first three principal components, distinct behavioral patterns within each cluster were identified.

5. Results and Discussion

This section presents the findings of our study, summarizing the key results of each approach and discussing their implications within the XRP ecosystem and XRPL research. Thereby, the study's primary objectives are addressed: characterizing and segmenting economic agents on the XRPL and evaluating the predictability of fraudulent accounts within the network.

5.1 Findings and Implications: Segmenting Economic Agents

This section focuses on the analysis of economic agents on the XRPL, with 5.1.1 detailing our heuristics-based approaches and 5.1.2 exploring the unsupervised clustering models. Subsection 5.1.3 then compares the results and approaches of both methodologies.

5.1.2 Clustering Models for Agent Segmentation

This section presents the results of the unsupervised clustering models, including a detailed analysis of the selected model's clusters, and their alignment with the known domains dataset. Finally, the findings and their implications for the XRPL ecosystem are discussed.

Results of Unsupervised Clustering Models

Table 3 demonstrates the results of all models including their calculated metrics.

Agglomerative Clustering, using the "average" linkage method, initially produced moderate results with 4 clusters, achieving a Silhouette score of 0.41, a Calinski-Harabasz index of 3,001.35, and a Davies-Bouldin index of 0.48. However, when PCA was applied beforehand, the performance significantly improved, yielding a higher Silhouette score of 0.47, a substantially increased Calinski-Harabasz index of 40,290.92, and a lower Davies-Bouldin index of 0.63. The models' clusters exhibit some alignment with the known domains dataset, with an ARI of 0.4 and a NMI of 0.2.

The GMM displayed comparable trends. Without PCA, the "tied" covariance GMM with 5 clusters resulted in a relatively low Silhouette score of 0.31, a Calinski-Harabasz index of 13,882.33, and a Davies-Bouldin index of 1.4. After incorporating PCA, GMM performance improved notably with 6 clusters, achieving a Silhouette score of 0.50, a Calinski-Harabasz index of 50,715.88, and a Davies-Bouldin index of 0.74. However, its ARI of 0.31 and NMI of 0.25 indicate that while PCA improved the overall clustering structure, the alignment with labeled data remained limited.

K-Means demonstrated the strongest performance, particularly when combined with PCA. Without PCA, K-Means identified 4 clusters with a Silhouette score of 0.32, a Calinski-Harabasz index of 17,022.74, and a Davies-Bouldin index of 1.15. When PCA was applied, K-Means again identified 4 clusters but showed improvements across all metrics, with a Silhouette

score of 0.51, a Calinski-Harabasz index of 46,800.46, and a Davies-Bouldin index of 0.65. This configuration also achieved the highest ARI of 0.49 and an NMI of 0.29, highlighting its ability to generate clusters that are both well-separated and aligned with labeled data.

The inclusion of PCA prior to clustering improved the performance of all models by reducing the dimensionality of the data. This step likely removed noise and redundant features, enabling the algorithms to focus on the most informative components. PCA also simplified the clustering task by projecting the data into a lower-dimensional space where meaningful patterns could emerge more clearly.

K-Means with PCA was selected as the final model based on its superior performance across the presented metrics. With 4 clusters, this configuration achieved the highest Silhouette score (0.51) among all tested models, indicating moderately separated and cohesive clusters. It also produced the lowest Davies-Bouldin index (0.65), suggesting minimal overlap between clusters. Figure 4 in the Appendix provides a visual representation of the clustering using the first three principal components. Additionally, its external validation metrics, with an ARI of 0.49 and a NMI of 0.29, further emphasize the best alignment with externally labeled data.

		Agglomerative Clustering	Gaussian Mixture	K-Means	Agglomerative Clustering with PCA	Gaussian Mixture with PCA	K-Means with PCA
	<i>Number of clusters</i>	4	5	4	4	6	4
Internal metrics	<i>Silhouette score</i>	0.41	0.31	0.32	0.47	0.50	0.51
	<i>Calinski-Harabasz index</i>	3,001.35	13,822.33	17,022.74	40,290.92	50,715.88	46,800.46
	<i>Davies Bouldin index</i>	0.48	1.40	1.15	0.63	0.74	0.65
External metrics	<i>ARI</i>	-	-	-	0.4	0.31	0.49
	<i>NMI</i>	-	-	-	0.2	0.25	0.29

Table 3: Calculated metrics for unsupervised models

In terms of computational efficiency and speed, K-Means was the fastest algorithm among those tested, while Agglomerative Clustering was the slowest. This computational advantage

further supports the suitability of K-Means, especially when combined with PCA (Table 14 in the Appendix).

Cluster Analysis

The following clusters are the result of the selected model, K-Means with prior PCA, yielding 4 clusters. Table 4 shows a detailed overview of the median of key features for each cluster.

Cluster 0 is the second-largest cluster, consisting of 14,000 addresses, which account for 35% of the dataset. This cluster shows high activity, with the second-highest median *total_transaction* count (177 transactions) and the highest number of *active_days* (56 days). It also has the highest ratio of TrustSet transactions (median 48%). Accounts in Cluster 0 typically process small transaction amounts, with transactions exceeding 500 XRP accounting for less than 0.33% of outgoing payments and 0.34% of incoming payments. Only Cluster 3 has a lower *total_sum_received* (median 100 XRP) and a lower *unique_transaction_partner_ratio* (median 16%).

Cluster 1 contains 9,000 addresses, representing approximately 23% of the dataset. It has the lowest median *total_transaction* count (20 transactions) but the longest absolute *active_period* (median 636 days). Accounts in this cluster primarily act as receivers, with a median of 56% of incoming payments under 100 XRP, while only 0.42% of outgoing payments are below 100 XRP. Cluster 1 also exhibits the highest *unique_transaction_partner_ratio* (median 59%) and ranks second in *total_sum_received* (median 1591 XRP).

cluster	÷ 0	÷ 1	÷ 2	÷ 3	÷
total_transactions	177.0000	20.0000	21.0000	1041.0000	
absolute_active_period	495.0000	636.0000	433.0000	16.0000	
active_days	56.0000	16.0000	12.0000	8.0000	
transaction_frequency	0.6875	0.0329	0.0698	105.5000	
nftoken_create_offer_ratio	0.0000	0.0000	0.0000	0.0000	
trust_set_ratio	0.4828	0.0523	0.0000	0.0000	
payment_under_100_xrp_receiver_ratio	0.0330	0.5625	0.1026	0.0008	
payment_under_100_xrp_sender_ratio	0.0068	0.0000	0.0500	0.9983	
payment_over_500_xrp_receiver_ratio	0.0000	0.0337	0.2632	0.0000	
payment_over_500_xrp_sender_ratio	0.0000	0.0000	0.1379	0.0000	
average_sent_amount_xrp	18.0000	85.3524	1343.9485	0.0000	
average_received_amount_xrp	17.3037	105.8770	1291.0318	13.0000	
unique_transaction_partner_ratio	0.1663	0.5926	0.2692	0.0032	
total_sum_received	100.0000	1590.8989	19901.8383	14.8000	

Table 4: Median of key features per cluster for K-Means and 4 clusters

Cluster 2 is the largest cluster, consisting of 15,000 addresses and accounting for 38% of the dataset. This cluster handles the largest transaction volumes, with median amounts of 1,343 XRP sent and 1,291 XRP received per transaction. Median payments exceeding 500 XRP account for 26% of outgoing transactions and 26% of incoming transactions, the highest among all clusters. It also ranks second in *unique_transaction_partner_ratio* (median 27%) and has the highest *total_sum_received* (median 19902 XRP).

Cluster 3 contains 2,000 addresses, making up only 5% of the dataset. Despite its small size, it has the highest median *total_transaction* count (1,041 transactions) but the shortest *active_period* (median 16 days). This results in the highest *transaction_frequency* (median 105.5 transactions per day). Accounts in this cluster predominantly make small payments, with 90% of all transactions involving amounts under 10 XRP, the highest proportion across all clusters. Cluster 3 also has the lowest *total_sum_received* (median 14.8 XRP) and the lowest *unique_transaction_partner_ratio* (median 0.3%).

Alignment with known domains

Verifying the clusters with the known domains dataset reveals how well the identified groupings capture distinct behavioral patterns and account types within the XRPL ecosystem

(Figure 2). Cluster 0 is characterized by a strong concentration of NFT-related addresses, with 90% of all NFT addresses from the labeled dataset falling within this cluster. These accounts exhibit behaviors commonly associated with NFT activities, such as frequent minting and trading of NFTs and offer placing. Additionally, 85% of token issuer addresses are grouped in this cluster, suggesting a shared behavioral pattern between these two account types. Both exhibit smaller payment amounts, lower connectivity to major exchanges, and a similar ratio of unique counterparties. Another pattern they share is the use of trust lines. On the XRPL, receiving payments in assets other than XRP requires a trust line to the issuing account of that asset (Ripple, n.d.-d). This mechanism may contribute to the observed clustering of token issuers and NFT traders, as both rely on such trust lines to facilitate the creation and transfer of tokenized assets. This overlap indicates that while NFT addresses and token issuers share common features, they are distinguishable from other account types in the dataset. However, consistent with NFT addresses and token issuers is the low diversity in unique transaction partners and the transfer of small amounts as interactions are often confined to specific counterparties, such as marketplaces, traders, or platforms facilitating these activities. While these accounts play a significant role in driving economic activity on the ledger through their frequent transactions, their overall contribution in terms of transaction volume and connectivity is comparatively moderate.

Cluster 1 contains 51% of the addresses labeled as gaming/gambling platforms, making it the dominant category within this group. Other labeled categories, such as NFT traders, exchanges or token issuers, are rarely represented, apart from the "other" category, which includes addresses that are difficult to define. This strong representation suggests that the selected model shows potential in isolating gambling addresses effectively.

A key trait of this cluster is its predominant role as a receiver of small payments, aligning closely with the transactional patterns of gambling platforms. These small incoming

payments likely represent deposits made by individuals who gamble on these platforms. Additionally, Cluster 1 has the highest ratio of unique transaction partners, reflecting its broad network of interactions. Its significant rate of interaction with known exchanges further emphasizes its role in the XRPL ecosystem. On the one hand, these addresses contribute significantly to the network by processing large amounts of XRP in total and maintaining a high level of connectivity within the ecosystem. Their interactions with a wide range of transaction partners and known exchanges highlight their integration into the broader XRPL economy. On the other hand, the inherently speculative and risk-driven nature of gambling platforms introduces complexities. While their activity supports network usage and liquidity, it also

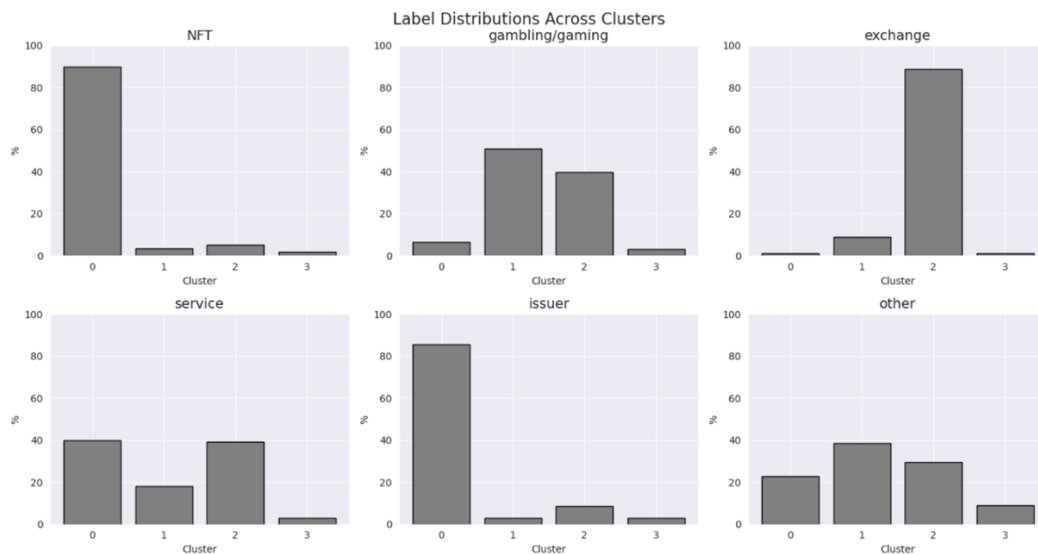


Figure 2: Label distribution across clusters

carries potential downsides, such as promoting behaviors that are not universally seen as positive or sustainable for the ecosystem's long-term health. Cluster 2 stands out as the largest cluster, with 89% of all known exchange addresses grouped within it, demonstrating that this cluster effectively captures the typical behavior of exchanges. Interestingly, 40% of all gambling and gaming addresses are also found in this cluster, likely due to shared characteristics such as frequent high-volume transactions. While exchanges are not grouped into the predominantly gambling-focused Cluster 1, some gambling addresses are classified within this exchange-dominated cluster. This overlap suggests similarities in their transactional

patterns, particularly in terms of volume and network connectivity. Incorporating time-based features could further differentiate these two account types by capturing temporal dynamics unique to their behaviors.

Cluster 2 plays a crucial role in the network. As the largest cluster, it provides essential liquidity and facilitates a substantial portion of the network's transaction volume. Its high level of connectivity and interactions with other major entities underscore its importance in maintaining the ledger's overall functionality.

Cluster 3 is the smallest group within the dataset, representing only 5% of all addresses. Unlike other clusters, it does not exhibit a major concentration of any specific labels, making it a more heterogeneous grouping. However, its distinct characteristics suggest a specialized and narrow role within the XRPL ecosystem. Accounts in this cluster are defined by their high transaction frequency within a short active period, indicating a focus on rapid and repetitive transactions. They also exhibit the lowest average sent transaction amounts and a low ratio of unique transaction partners, suggesting limited diversity in counterparties and repetitive interactions with the same accounts. Additionally, the absence of significant relationships with major exchanges further highlights the isolated and niche behavior of these addresses. These observations indicate that the known domains dataset may not contain a category that fully explains the behavior of this cluster, suggesting the presence of an unidentified type within the ecosystem.

The chosen approach, which incorporated external data in form of known domains, enabled a verification process, making it superior to other methods that lacked the ability to validate their clusters effectively. While other studies (Payette, Schwager, and Murphy 2017; Birrane, 2017) can only assess whether their models form distinct clusters, our approach allows for an additional evaluation step: determining whether these clusters accurately reflect real-world economic segments. Moreover, this approach also outperforms the work of Vlahavas,

Karasavvas, and Vakali (2024) who encountered challenges in distinguishing between clusters, as gambling, exchanges, and services were often grouped together. Presumably due to the additional information from the XRP Ledger, our model was able to differentiate effectively between gambling/gaming, exchanges, and NFT/issuers. This distinction underscores the suitability of the XRP Ledger for this type of research and highlights its potential to generate insights that could be transferred to other blockchains.

Practical Implications

The clustering analysis provided valuable insights into the structure and behavior of the XRPL ecosystem. The majority of gambling/gaming addresses formed a distinct cluster, reflecting their unique transactional patterns as recipients of small payments. Additionally, the two largest clusters were dominated by exchange addresses and NFT/token issuers, collectively accounting for 75% of all addresses. These clusters highlight the critical roles of exchanges in providing liquidity and managing high-volume transactions, as well as NFT/token issuers' activities focused on minting and trading. A small group of addresses (5%) displayed no significant role, characterized by repetitive transactions, low diversity, and limited connectivity to major entities.

The insights derived from the unsupervised clustering models provide actionable implications for the XRPL ecosystem, demonstrating the utility and effectiveness of the selected model in uncovering distinct behavioral patterns. By segmenting economic agents based on their transactional activities, the chosen clustering model reveals how the ledger facilitates diverse use cases, such as high-value exchanges, NFT trading, and gambling. This understanding is critical not only for ongoing ecosystem monitoring but also for communicating the ledger's utility to external parties, such as investors or regulatory bodies like the Security Exchange Commission (SEC), potentially strengthening the case for XRP by showcasing its

legitimate and varied applications. The model's ability to isolate specific categories, such as gambling-related addresses, demonstrates its potential for addressing regulatory concerns. This is particularly relevant in jurisdictions where gambling or other high-risk activities are closely scrutinized. Finally, Ripple and other ecosystem stakeholders can leverage the model to gain strategic insights into user behavior. By understanding how different segments engage with the network, decisions can be informed by data to optimize the ledger for its most active and economically valuable participants.

5.1.2 Comparison of Heuristics and Clustering Models

The methodologies of the heuristics and clustering approaches differ significantly in their focus and execution. The heuristics approach aims to identify clear rules and thresholds that distinguish specific account types based on the previously observed patterns. In contrast, the clustering approach forms distinct groups in the data without predefined labels, comparing these clusters against labeled data to validate their alignment with known account categories.

One key difference between both approaches is the feature selection. The heuristics approach identifies relevant features for each type using logistic regression and Kruskal-Wallis test. The clustering approach employs recursive feature elimination to select features for each model type.

While the heuristics approach focuses solely on the labeled addresses to characterize predefined segments, the clustering approach identifies clusters based on similarities in the dataset, additionally including randomly selected addresses and leveraging the labeled domains for cluster verification.

The insights derived from both approaches reveal overlapping trends: The heuristics approach performed best in characterizing exchanges, NFTs, and issuers, achieving F1 scores of 0.93, 0.75, and 0.72, respectively. The clustering approach similarly identified NFTs/token

issuers and exchanges as the easiest types to segment. NFTs and issuers were grouped together due to similar patterns and clearly distinguishable from other types. Despite some overlap with gambling/gaming entities exchanges were successfully grouped together.

Gambling/gaming entities (F1 score: 0.38) and services (F1 score: 0.46) were harder to distinguish through heuristics, suggesting that their behaviors may overlap more with other types or lack distinct patterns easily captured by threshold-based rules. No heuristics were defined for the “other” category, due to the lack of distinct behavioral patterns. Clustering of gambling/gaming entities showed moderate success, with 51% grouped almost exclusively into a dedicated cluster. However, a significant portion of gambling entities were mixed with exchanges, indicating overlaps in transaction behaviors. Services and "other" were the hardest to segment, as they were distributed across multiple clusters without forming distinct groups.

While *payment_without_XRP_ratio*, *payment_with_XRP_ratio*, and *unique_transaction_partner_ratio* were frequently utilized in best-performing heuristics, two of these three features were excluded during the recursive feature selection process for clustering models (*payment_without_XRP_ratio* and *payment_with_XRP_ratio*). The differences in feature selection show how the methodologies’ differing approaches can lead to variations in results.

On a broader level, both approaches provide valuable insights into usage patterns within the XRPL ecosystem, but from different perspectives. The heuristics approach, with its manual, hypothesis-driven methodology, integrates labeled data from the start. The results are focused outputs that highlight distinct patterns within this data, making it particularly effective for understanding the specific differences between the predefined agent categories. Additionally, its simplified selection process ensures clarity about the features driving distinctions between the categories.

The clustering approach employs a more automated and exploratory methodology. By including random accounts in addition to labeled data, it captures a broader range of behaviors, offering a more complete picture of account activity across the XRPL. Its scalability and efficiency make it particularly useful for scenarios where larger dataset are available.

Each approach demonstrates its own implications for the XRP ecosystem. Since the XRPL remains relatively underexplored, heuristics provide a valuable first step for researchers to understand economic activities on the network, following the typical process of early ledger exploration (Section 3.1.1). In the future, heuristics could aid in creating larger labeled datasets for supervised machine learning, as shown by researchers on other DLTs (e.g. Lin et al., 2019). The clustering approach provides a more comprehensive understanding of account patterns on the XRPL and offers a framework for potential applications, such as ecosystem monitoring by regulatory authorities. Additionally, the novel approach of verifying clusters with off-chain data presents opportunities for future research to apply this methodology to other DLTs.

6 Limitations

Our study is subject to several limitations that impact the scope and depth of the analysis.

The feature dataset was restricted to a specific time frame due to computational constraints, preventing the examination of transaction patterns beyond this period. The random sampling of data used for different approaches introduces variability that could impact the consistency of results. Crucial metadata, such as details of non-XRP token transactions, is absent due to the limitations of the public data source used. The study is also limited by its feature set. Time-based variables could not be fully computed due to computational resource constraints, and account-specific metrics, such as total address balances, were unavailable. Eventually, the inclusion of graph-based features could have enhanced the analysis by providing deeper insights into relationships and interactions between accounts.

Additionally, self-labeling the known domains data might have introduced inaccuracies or inconsistencies, potentially biasing the results. Furthermore, entities often possess multiple addresses, which they presumably use for different purposes, leading to varied transaction patterns across these addresses. This makes characterizing and verifying segments challenging.

Ultimately, the dynamic nature of blockchain usage presents a hurdle. Methodologies effective today may become less applicable over time as user behaviors and network structures evolve (Victor 2020).

Further limitations specific to our distinct approaches include:

2. Unsupervised Clustering Limitations

For our clustering models, a subsample size of 40,000 addresses was selected for computational feasibility. Expanding the subset to include more addresses could enhance the results by identifying additional patterns with greater reliability. External validation using labeled data is constrained by the limited and potentially biased ground truth dataset, particularly for underrepresented classes, as highlighted by the limitations of the heuristics.

Additionally, clustering models like K-Means and GMM require the number of clusters to be predefined, which can sometimes result in suboptimal groupings. For instance, clusters like Cluster 3 remain difficult to fully interpret, likely due to missing label categories, such as individuals or automated agents, which limits overall interpretability.

Furthermore, access to more granularly labeled data could enhance classification accuracy, addressing challenges such as overlapping categories during the validation of clusters with external data.

7 Conclusion

By analyzing XRPL transaction data on account-level, this study has shown that the segmenting XRPL agents is partially doable and can inform on the characteristics of the economic activity inherent in the XRP network. Furthermore, the study demonstrated the capability of supervised machine learning in effectively detecting fraudulent XRP accounts.

The need for greater transparency on the XRPL arises from its role in facilitating large-scale cross-border transactions, token issuance, decentralized exchange activities, NFT operations, all conducted in a pseudonymous manner. This lack of transparency presents challenges for regulators, compliance bodies, investors, financial institutions, and the broader XRP ecosystem in assessing the network's reliability and its participation in economically significant activities within markets.

Despite the scarcity of ground truth off-chain data and the lack of research on XRPL network activity, our study successfully uncovered novel insights and introduced a framework for distinguishing economic agent categories on the XRPL.

Our account-based heuristics establish a foundational framework for characterizing different economic agent categories on the ledger, providing a hypothesis-driven methodology that highlights distinct patterns within the ground truth data. This makes it particularly valuable for simplifying complex data structures and gaining initial insights into the XRP network.

The clustering models introduced an automated method capable of capturing a broader range of behaviors, providing a more comprehensive view of agent activity on the XRPL. This methodology leverages a labeled dataset ensuring the interpretability for most clusters, enabling an off-chain validation of the results. The scalability of this approach makes it applicable to a wide range of use cases, including ecosystem monitoring, regulatory compliance, and XRPL stakeholder-informing.

Promising results were observed for both approaches. Heuristics help in identifying distinct characteristics, especially for exchanges, NFTs and issuers. The clustering models

effectively segmented NFT-related addresses and issuers, which, although grouped together, are clearly distinguishable from other account types, while exchanges also formed a distinct cluster apart from a small overlap with gambling addresses. However, differentiating addresses labeled as service and other from the rest remained a significant challenge in both approaches, highlighting an area for future research.

Ultimately, our fraud detection models demonstrated that identifying illicit accounts on the XRPL is feasible using supervised learning algorithms. The final model accurately detected nearly all fraud cases in the dataset and exhibited robust generalization capabilities, making it a strong candidate for real-time detection systems as well as for historical analyses of past XRP fraud cases. The insights gained into fraud patterns and the key features driving the classification showed that XRPL fraudsters tend to use non-XRP tokens and exhibit transactions with a large base of counterparties in high frequencies. Eventually, our insights offer valuable guidance for tailoring regulatory frameworks to the XRPL. To enhance our findings, future research could build on our approach by incorporating additional fraud types not represented in the current dataset.

Eventually, certain limitations of our study must be acknowledged. Constraints in data availability, including limited metadata and imbalances in the ground truth datasets, affected the generalizability of results. Computational restrictions limited the inclusion of more complex features and larger data subsets, which could have enhanced insights into account behaviors and interactions. Additionally, our self-labeled data might have introduced some bias as well as the assumptions made during data sampling processes may influence the reliability of the findings. Future work should address these limitations by leveraging more extensive datasets, incorporating advanced feature engineering, and adapting the methodologies to the ever evolving blockchain dynamics.

Ultimately, the insights and methodologies developed in this study extend beyond the XRPL, offering a framework that can be adapted to other blockchain ecosystems to address similar challenges in transparency, fraud detection, and user segmentation. By bridging on-chain analysis with off-chain validation, reliable insights can be generated to enhance trust and accountability in decentralized systems. To build on our findings, further research should prioritize collaborations between blockchain developers, regulatory bodies, and academic institutions to refine data collection processes and standardize methodologies for addressing fraud and economic activity monitoring. Strengthening transparency and mitigating risks in blockchain ecosystems will be key to enforce broader adoption and acceptance of decentralized technologies in financial markets, and beyond.

List of References

- Ahmadova, Sevinj, and Mustafa Salim Ere. 2022. 'A Review on Ripple, a Financial Intermediary Coin'. *Journal of Academic Projection* 7 (2): 117–30.
- Akcora, Cuneyt Gurcan, Yulia R. Gel, and Murat Kantarcioglu. 2022. 'Blockchain Networks: Data Structures of Bitcoin, Monero, Zcash, Ethereum, Ripple, and Iota'. *Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery* 12 (1): e1436. <https://doi.org/10.1002/widm.1436>.
- Alahmad, Mohammed, Adel Alfouderi, Ahmad Alonaizi, and Meshal Aldhamen. 2023. 'Comparison Study of the Top 5 Leading Cryptocurrencies Based on General Consensus Protocol: Bitcoin, Ethereum, Tether, XRP and Bitcoin Cash'. *WSEAS TRANSACTIONS ON COMPUTER RESEARCH* 11 (April):23–32. <https://doi.org/10.37394/232018.2023.11.3>.
- Androulaki, E, G. O. Karame, M. Roeschlin, T. Scherer, and S. Capkun. 2013. 'Evaluating User Privacy in Bitcoin'. In , Revised Selected Papers 17:34–51. Okinawa, Japan: Springer Berlin Heidelberg.
- Birrane, Kieran Daniel. n.d. 'An Exploration of Blockchain: An Unsupervised Analysis of the Ethereum Network'.
- Breiman, Leo. 2001. 'Random Forests'. *Machine Learning* 45 (1): 5–32. <https://doi.org/10.1023/A:1010933404324>.
- Caliński, T., and J Harabasz. 1974. 'A Dendrite Method for Cluster Analysis'. *Communications in Statistics* 3 (1): 1–27. <https://doi.org/10.1080/03610927408827101>.
- Chainalysis. 2019. 'PlusToken Scammers Didn't Just Steal \$2+ Billion Worth of Cryptocurrency. They May Also Be Driving Down the Price of Bitcoin.' *Chainalysis* (blog). 16 December 2019. <https://www.chainalysis.com/blog/plustoken-scam-bitcoin-price/>.
- . 2024. 'The 2024 Crypto Crime Report'. <https://www.chainalysis.com/wp-content/uploads/2024/06/the-2024-crypto-crime-report-release.pdf>.
- . n.d. 'Key Players In Crypto Report'. Chainalysis.
- Chase, Brad, and Ethan MacBrough. 2018. 'Analysis of the XRP Ledger Consensus Protocol'. arXiv. <https://doi.org/10.48550/arXiv.1802.07242>.
- Chawla, N. V., K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. 2002. 'SMOTE: Synthetic Minority Over-Sampling Technique'. *Journal of Artificial Intelligence Research* 16 (June):321–57. <https://doi.org/10.1613/jair.953>.
- Chen, Tianqi, and Carlos Guestrin. 2016. 'XGBoost: A Scalable Tree Boosting System'. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–94. San Francisco California USA: ACM. <https://doi.org/10.1145/2939672.2939785>.
- CoinMarketCap. 2024. 'Cryptocurrency Prices, Charts And Market Capitalizations'.

CoinMarketCap. 2024. <https://coinmarketcap.com/>.

Davies, David, and Don Bouldin. 1979. 'A Cluster Separation Measure'. *Pattern Analysis and Machine Intelligence, IEEE Transactions On PAMI-1* (May):224–27. <https://doi.org/10.1109/TPAMI.1979.4766909>.

Ermilov, Dmitry, Maxim Panov, and Yury Yanovich. 2017. 'Automatic Bitcoin Address Clustering'. In , 461–66. <https://doi.org/10.1109/ICMLA.2017.0-118>.

European Union Agency for Cybersecurity (ENISA). n.d. 'Blockchain'. Page. ENISA. Accessed 28 November 2024. <https://www.enisa.europa.eu/topics/incident-response/glossary/blockchain>.

Farrugia, Steven, Joshua Ellul, and George Azzopardi. 2020. 'Detection of Illicit Accounts over the Ethereum Blockchain'. *Expert Systems with Applications* 150 (July):113318. <https://doi.org/10.1016/j.eswa.2020.113318>.

Gigerenzer, Gerd, and Wolfgang Gaissmaier. 2011. 'Heuristic Decision Making'. *Annual Review of Psychology* 62 (Volume 62, 2011): 451–82. <https://doi.org/10.1146/annurev-psych-120709-145346>.

Harlev, Mikkel Alexander, Haohua Sun Yin, Klaus Christian Langenheldt, Raghava Rao Mukkamala, and Ravi Vatrapu. 2018. 'Breaking Bad: De-Anonymising Entity Types on the Bitcoin Blockchain Using Supervised Machine Learning'. In . <https://core.ac.uk/download/pdf/143481278.pdf>.

Harrigan, Martin, and Christoph Fretter. 2016. *The Unreasonable Effectiveness of Address Clustering*. <https://doi.org/10.48550/arXiv.1605.06369>.

Hellwig, Daniel, Goran Karlic, and Arnd Huchzermeier. 2020. *Build Your Own Blockchain: A Practical Guide to Distributed Ledger Technology*. Management for Professionals. Cham: Springer International Publishing. <https://doi.org/10.1007/978-3-030-40142-9>.

Jourdan, Marc, Sebastien Blandin, Laura Wynter, and Pralhad Deshpande. 2018. *Characterizing Entities in the Bitcoin Blockchain*. <https://doi.org/10.48550/arXiv.1810.11956>.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. 'LightGBM: A Highly Efficient Gradient Boosting Decision Tree'. In *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html.

Kılıç, Baran, Alper Sen, and Can Özturan. 2022. 'Fraud Detection in Blockchains Using Machine Learning'. In *2022 Fourth International Conference on Blockchain Computing and Applications (BCCA)*, 214–18. <https://doi.org/10.1109/BCCA55292.2022.9922045>.

Kraken. n.d. 'Beware of Crypto Giveaway Scams | Kraken'. Accessed 27 November 2024. <https://support.kraken.com/hc/en-us/articles/360057159411-Beware-of-crypto-giveaway-scams>.

- Li, Xiaofan, and Andrew Whinston. 2020. 'Analyzing Cryptocurrencies'. *Information Systems Frontiers* 22 (February). <https://doi.org/10.1007/s10796-019-09966-2>.
- Lin, Yu-Jing, Po-Wei Wu, Cheng-Han Hsu, I-Ping Tu, and Shih-wei Liao. 2019. 'An Evaluation of Bitcoin Address Classification Based on Transaction History Summarization'. In *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 302–10. <https://doi.org/10.1109/BLOC.2019.8751410>.
- Lorenz, Joana, Maria Inês Silva, David Aparício, João Tiago Ascensão, and Pedro Bizarro. 2021. 'Machine Learning Methods to Detect Money Laundering in the Bitcoin Blockchain in the Presence of Label Scarcity'. arXiv. <http://arxiv.org/abs/2005.14635>.
- Lundberg, Scott, and Su-In Lee. 2017. 'A Unified Approach to Interpreting Model Predictions'. arXiv. <http://arxiv.org/abs/1705.07874>.
- Mauri, Lara, Stelvio Cimato, and Ernesto Damiani. 2018. 'A Comparative Analysis of Current Cryptocurrencies': In *Proceedings of the 4th International Conference on Information Systems Security and Privacy*, 127–38. Funchal, Madeira, Portugal: SCITEPRESS - Science and Technology Publications. <https://doi.org/10.5220/0006648801270138>.
- Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2013. 'A Fistful of Bitcoins: Characterizing Payments among Men with No Names'. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, 127–40. Barcelona Spain: ACM. <https://doi.org/10.1145/2504730.2504747>.
- Molnar, Christoph. 2024. *Interpretable Machine Learning A Guide for Making Black Box Models Explainable*. <https://christophm.github.io/interpretable-ml-book/feature-importance.html>.
- Monaco, John V. 2015. 'Identifying Bitcoin Users by Transaction Behavior'. In , edited by Ioannis A. Kakadiaris, Ajay Kumar, and Walter J. Scheirer, 945704. Baltimore, Maryland, United States. <https://doi.org/10.1117/12.2177039>.
- Morgia, Massimo La, Alessandro Mei, Francesco Sassi, and Julinda Stefa. 2024. 'Pump and Dumps in the Bitcoin Era: Real Time Detection of Cryptocurrency Market Manipulations'. arXiv. <http://arxiv.org/abs/2005.06610>.
- Nakamoto, Satoshi. 2008. 'Bitcoin: A Peer-to-Peer Electronic Cash System'.
- Nerurkar, Pranav, Sunil Bhirud, Dhiren Patel, Romaric Ludinard, Yann Busnel, and Saru Kumari. 2021. 'Supervised Learning Model for Identifying Illegal Activities in Bitcoin'. *Applied Intelligence* 51 (June):1–20. <https://doi.org/10.1007/s10489-020-02048-w>.
- Nick, Jonas David. 2015. 'Data-Driven De-Anonymization in Bitcoin'. Application/pdf, Online-Ressource. <https://doi.org/10.3929/ETHZ-A-010541254>.
- Ostapowicz, Michał, and Kamil Żbikowski. 2019. 'Detecting Fraudulent Accounts on Blockchain: A Supervised Approach'. In , edited by Reynold Cheng, Nikos Mamoulis,

- Yizhou Sun, and Xin Huang, 11881:18–31. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-34223-4_2.
- Ostertagová, Eva, Oskar Ostertag, and Jozef Kováč. 2014. ‘Methodology and Application of the Kruskal-Wallis Test’. *Applied Mechanics and Materials* 611:115–20. <https://doi.org/10.4028/www.scientific.net/AMM.611.115>.
- Payette, James, Samuel Schwager, and Joseph Murphy. 2017. ‘CHARACTERIZING THE ETHEREUM ADDRESS SPACE’.
- Peduzzi, Gaspard, Jason James, and Jiahua Xu. 2021. ‘Jack the Rippler: Arbitrage on the Decentralized Exchange of the XRP Ledger’. In *2021 3rd Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*, 1–2. <https://doi.org/10.1109/BRAINS52497.2021.9569833>.
- Ramezan, Gholamreza, and Cyril Leung. 2020. ‘Analysis of Proof-of-Work-Based Blockchains Under an Adaptive Double-Spend Attack’. *IEEE Transactions on Industrial Informatics* 16 (11): 7035–45. <https://doi.org/10.1109/TII.2020.2977689>.
- Reid, Fergal, and Martin Harrigan. 2011. ‘An Analysis of Anonymity in the Bitcoin System’. *Security and Privacy in Social Networks* 3 (July). <https://doi.org/10.1109/PASSAT/SocialCom.2011.79>.
- Rella, Ludovico. 2020. ‘Steps towards an Ecology of Money Infrastructures: Materiality and Cultures of Ripple’. *Journal of Cultural Economy* 13 (2): 236–49. <https://doi.org/10.1080/17530350.2020.1711532>.
- Ripple. 2024a. ‘Transaction Metadata’. 9 October 2024. <https://xrpl.org/docs/references/protocol/transactions/common-fields>.
- . 2024b. ‘What Is the XRP Ledger?’ <https://xrpl.org/docs/introduction/what-is-the-xrp-ledger>.
- . n.d.-a. ‘Transaction Types’. XRPL. Accessed 6 November 2024. <https://xrpl.org/docs/references/protocol/transactions/types>.
- . n.d.-b. ‘XRP Ledger - Use Cases & Featured Projects’. Accessed 28 November 2024. <https://xrpl.org/about/uses>.
- . n.d.-c. ‘XRP Ledger History’. Accessed 28 November 2024. <https://xrpl.org/about/history>.
- . n.d.-d. ‘XRP Ledger Home | XRPL.Org’. Accessed 30 November 2024. <https://xrpl.org/docs/concepts/tokens/fungible-tokens/authorized-trust-lines>.
- Rousseeuw, Peter J. 1987. ‘Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis’. *Journal of Computational and Applied Mathematics* 20 (November):53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7).
- Scikit-Learn. 2024. ‘Scikit Learn Documentation: Decision Trees’. 2024. <https://scikit-learn.org>.

learn/stable/modules/tree.html.

———. n.d.-a. ‘Ensembles: Gradient Boosting, Random Forests, Bagging, Voting, Stacking’. Scikit-Learn. Accessed 7 November 2024. <https://scikit-learn/stable/modules/ensemble.html>.

———. n.d.-b. ‘Scikit Learn Documentation: PCA’. Scikit-Learn. Accessed 30 November 2024. <https://scikit-learn/stable/modules/generated/sklearn.decomposition.PCA.html>.

———. n.d.-c. ‘Scikit-Learn Documentation: K-Means’. Scikit-Learn. Accessed 8 November 2024. <https://scikit-learn/stable/modules/clustering.html>.

———. n.d.-d. ‘Scikit-Learn Documentation: GaussianMixture’. Scikit-Learn. Accessed 30 November 2024. <https://scikit-learn/stable/modules/generated/sklearn.mixture.GaussianMixture.html>.

———. n.d.-e. ‘Scikit-Learn Documentation: Hierarchical Clustering’. Scikit-Learn. Accessed 8 November 2024. <https://scikit-learn/stable/modules/clustering.html>.

SEC. 2017. ‘Investor Alert: Ponzi Schemes Using Virtual Currencies’. https://www.sec.gov/files/ia_virtualcurrencies.pdf.

Sergeenkov, Andrey. 2024. ‘What Is Ripple (XRP)?’ *Forbes*. 27 October 2024. <https://www.forbes.com/sites/digital-assets/article/what-is-ripple-xrp/>.

Shannon, C. E. 1948. ‘A Mathematical Theory of Communication’. *The Bell System Technical Journal* 27 (3): 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.

Shayegan, Mohammad Javad. n.d. ‘A Collective Anomaly Detection Method Over Bitcoin Network’. <https://arxiv.org/pdf/2107.00925>.

Shayegan, Mohammad Javad, Hamid Reza Sabor, Mueen Uddin, and Chin-Ling Chen. 2022. ‘A Collective Anomaly Detection Technique to Detect Crypto Wallet Frauds on Bitcoin Network’. *Symmetry* 14 (2): 328. <https://doi.org/10.3390/sym14020328>.

Statista. n.d. ‘Number of Cryptocurrencies 2013-2024’. Statista. Accessed 30 November 2024. <https://www.statista.com/statistics/863917/number-crypto-coins-tokens/>.

Subashi, Roland. 2024. ‘Cryptocurrencies and Money Laundering’. *Balkan Journal of Interdisciplinary Research* 10 (1): 55–62. <https://doi.org/10.2478/bjir-2024-0005>.

Sun Yin, Hao Hua, Klaus Langenheldt, Mikkel Harlev, Raghava Rao Mukkamala, and Ravi Vatrapu. 2019. ‘Regulating Cryptocurrencies: A Supervised Machine Learning Approach to De-Anonymizing the Bitcoin Blockchain’. *Journal of Management Information Systems* 36 (1): 37–73. <https://doi.org/10.1080/07421222.2018.1550550>.

Toyoda, Kentaroh, Tomoaki Ohtsuki, and P. Takis Mathiopoulos. 2018. ‘Multi-Class Bitcoin-Enabled Service Identification Based on Transaction History Summarization’. In *2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, 1153–60.

https://doi.org/10.1109/Cybermatics_2018.2018.00208.

Ul Hassan, Muneeb, Mubashir Husain Rehmani, and Jinjun Chen. 2023. 'Anomaly Detection in Blockchain Networks: A Comprehensive Survey'. *IEEE Communications Surveys & Tutorials* 25 (1): 289–318. <https://doi.org/10.1109/COMST.2022.3205643>.

United Nations. n.d. 'Economic Agent'. UNTERM - The United Nations Terminology Database. Accessed 22 October 2024. <https://unterm.un.org/unterm2/en/view/bc6e1185-6da5-49ec-b4c7-7bf612a49236>.

Victor, Friedhelm. 2020. 'Address Clustering Heuristics for Ethereum'. In *Financial Cryptography and Data Security*, edited by Joseph Bonneau and Nadia Heninger, 12059:617–33. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-51280-4_33.

Victor, Friedhelm, and Andrea Marie Weintraud. 2021. 'Detecting and Quantifying Wash Trading on Decentralized Cryptocurrency Exchanges'. arXiv. <https://doi.org/10.48550/arXiv.2102.07001>.

Vlahavas, George, Kostas Karasavvas, and Athena Vakali. 2024. 'Unsupervised Clustering of Bitcoin Transactions'. *Financial Innovation* 10 (1): 25. <https://doi.org/10.1186/s40854-023-00525-y>.

Wind, Wietse. 2024. 'WietseWind/Fetch-Xrpl-Transactions'. <https://github.com/WietseWind/fetch-xrpl-transactions>.

Wu, Jiajing, Jieli Liu, Yijing Zhao, and Zibin Zheng. 2021. 'Analysis of Cryptocurrency Transactions from a Network Perspective: An Overview'. *Journal of Network and Computer Applications* 190 (September):103139. <https://doi.org/10.1016/j.jnca.2021.103139>.

Yu, Congcong, Chen Yang, Zheng Che, and Liehuang Zhu. 2023. 'Robust Clustering of Ethereum Transactions Using Time Leakage from Fixed Nodes'. *Blockchain: Research and Applications* 4 (1): 100112. <https://doi.org/10.1016/j.bcra.2022.100112>.

APPENDIX

Study	Data	Description
<i>Reid and Harrigan (2011)</i>	An Analysis of Anonymity in the Bitcoin System	<p>Approach: Analyzed Bitcoin's transaction history to study anonymity through network structures.</p> <p>Methodology: Combined on-chain transaction analysis (topological network structure) with off-chain data (context discovery and flow analysis) to link addresses to real-world identities.</p> <p>Key Insights: Demonstrated that Bitcoin's pseudo-anonymity can be compromised by combining transaction patterns with external information, as shown in a case study investigating Bitcoin theft.</p>
<i>Androulaki et al. (2013)</i>	Evaluating User Privacy in Bitcoin	<p>Approach: Heuristic clustering and re-identification attacks applied to Bitcoin.</p> <p>Methodology: Used transaction features such as timing, sender/receiver indices, and transaction amounts to cluster users and evaluate privacy measures.</p> <p>Key Insights: Demonstrated that nearly 40% of user profiles could be recovered, showing significant privacy vulnerabilities even when recommended privacy measures are applied.</p>
<i>Meiklejohn et al. (2013)</i>	A fistful of bitcoins: characterizing payments among men with no names	<p>Approach: Heuristic clustering and re-identification attacks applied to the Bitcoin blockchain.</p> <p>Methodology: Grouped Bitcoin wallets using evidence of shared authority (e.g., multi-input heuristic) and empirically linked clusters to real-world entities through purchasing experiments.</p> <p>Key Insights: Demonstrated the visibility of Bitcoin's transaction flows despite pseudo-anonymity.</p>
<i>Nick (2015)</i>	Data-Driven De-Anonymization in Bitcoin	<p>Approach: Analyzed clustering strategies to group Bitcoin addresses belonging to the same wallet.</p> <p>Methodology: Used ground truth data from a Connection Bloom Filtering vulnerability. Evaluated clustering techniques for well-known multi-input heuristic and two newly proposed heuristics, combining them for improved results.</p> <p>Key Insights: Demonstrated that even modern Bitcoin wallets fail to fully protect users; the multi-input heuristic alone reveals 68.59% of addresses on average, with further improvements achieved through advanced heuristics.</p>
<i>Monaco (2015)</i>	Identifying Bitcoin users by transaction behavior	<p>Approach: Long-term transactional behavior analysis applied to Bitcoin to identify and verify account holders.</p> <p>Methodology: Analyzed transaction features such as timestamps, coin flow, transaction intervals, and network connectivity using a dynamical systems approach on users with 100–1000 monthly transactions for at least six months.</p> <p>Key Insights: Exploiting patterns in long-term transactional behavior reduces anonymity, revealing identifiable trends across outgoing and incoming transactions.</p>
<i>Harrigan and Fretter (2016)</i>	The Unreasonable Effectiveness of Address Clustering	<p>Approach: Address clustering using transaction micro-structures in the Bitcoin system.</p> <p>Methodology: Analyzed clustering heuristics using transaction micro-structures to analyze cluster growth and centrality.</p>

		Key Insights: Identified key factors driving clustering effectiveness, including address reuse, incremental cluster growth, and super-clusters with high centrality.
<i>Ermilov, Panov, and Yanovich (2017)</i>	Automatic Bitcoin Address Clustering	Approach: Combined blockchain-based heuristics and off-chain data for Bitcoin address clustering. Methodology: Integrated common behavior patterns (e.g., common spending and one-time change) with off-chain information as additional “votes” for address separation to improve clustering precision. Key Insights: Demonstrated that combining blockchain and off-chain data enhances clustering accuracy, enabling more advanced de-anonymization and helping identify insecure Bitcoin usage patterns.
<i>Jourdan et al. (2018)</i>	Characterizing Entities in the Bitcoin Blockchain	Approach: Analyzed user-identifying patterns on the Bitcoin network to demonstrate unintended exposure of user information. Methodology: Investigated transaction patterns and network activity to uncover identifying information, highlighting widely available data-mining techniques. Key Insights: Bitcoin's pseudonymity can be compromised through transaction behaviors, as patterns and surrounding activity can reveal user identities.
<i>Toyoda, Ohtsuki, and Mathiopoulos (2018)</i>	Multi-Class Bitcoin-Enabled Service Identification Based on Transaction History Summarization	Approach: Multi-class service identification in Bitcoin using transaction history summarization from 2009-2017. Methodology: Extracted transaction history features from Bitcoin addresses and classified them into seven service categories using a supervised machine learning model. Key Insights: Achieved 72% accuracy in identifying services such as exchanges, gambling, mixers, and scams.
<i>Victor (2020)</i>	Address Clustering Heuristics for Ethereum	Approach: Developed address clustering heuristics tailored to Ethereum’s account-based model, addressing limitations of UTXO-based methods. Methodology: Proposed heuristics based on deposit address reuse, multiple airdrop participation, and token authorization patterns, applied over 4 years of Ethereum data. Key Insights: Identified 17.9% of active externally owned addresses as clusters, representing over 340,000 entities, with the deposit address heuristic being the most effective.
<i>Wu et al. (2021)</i>	Analysis of cryptocurrency transactions from a network perspective: An overview	Approach: Comprehensive review of existing de-anonymization methods on different blockchains Key Insights: De-anonymization methods are grouped into three main types: transaction property-based, behavior-based, and off-chain information-based approaches.

Table 5: Overview of key work of heuristics

	Authors	Results
Supervised Agent Classification	<i>Lin et al. 2019</i>	Utilized a labeled dataset of 26,313 BTC addresses categorized into six groups to train several classification models. They achieved a Micro-F1 score of 0.87 and a Macro-F1 score of 0.86 using LightGBM.
	<i>Harlev et al. 2018</i>	Worked with a labeled dataset of 434 BTC addresses categorized into ten groups to train classification models. They achieved an accuracy of 77% and an F1 score of 75% using Gradient Boosting.
Unsupervised Agent Segmentation/Clustering	<i>Payette, Schwager, and Murphy 2017</i>	Used several unsupervised models to cluster 250,000 Ethereum addresses. Found K-Means as their best performing model using the Calinski-Harabasz score as a metric and identified 4 clusters. They didn't link the clusters to real world categories.
	<i>Birrane 2017</i>	Clustered 31,006 Ethereum users using K-Means into 15 clusters. Used the Elbow method to determine the optimal number of clusters. 99,59% of the records were grouped in one cluster. The clusters were not linked to real world categories.
	<i>Ermilov, Panov, and Yanovich 2018</i>	They used a probabilistic model and off-chain data to cluster around 95 million BTC addresses into 6 groups. No metrics that test cluster separation have been used.
	<i>Vlahava, Karasavvas, and Vakali 2024</i>	Clustered transactions (instead of addresses) into five clusters using trimmed K-Means achieving a Silhouette score of 0.78. They used off-chain data to verify their results. Gambling, exchanges, and services were almost completely grouped together.

Table 6: Overview of key work of agent classification

Study	Data	Methods	Key Results
<i>Ostapowicz & Żbikowski (2019)</i>	ETH (2,200 fraudulent wallets, 349,999 non-fraudulent wallets)	Random Forest , SVM, XGBoost	Random Forest: Best recall (84.92%), FPR (9.69%); XGBoost: similar to Random Forest, RF slightly better; SVM: High recall, but high FPR (less practical).
<i>Farrugia et al. (2020)</i>	ETH network (4,681 accounts, 2,179 flagged, 2,502 normal accounts)	XGBoost	Best accuracy: 0.96, F1 score: 0.96. Generalization and resilience towards overfitting.
<i>Lorenz et al. (2021)</i>	BTC transactional dataset (203,769 transactions), comparing supervised vs. unsupervised models	Supervised: Random Forest Unsupervised: KNN, PCA, Isolation Forest	Unsupervised models performed below supervised baseline. Fraudulent transactions not detected as outliers.
<i>Nerurkar et al. (2021)</i>	BTC (non-binary classification task: darknet markets, exchanges, gambling, Ponzi, unclassified)	Log Regression, SVM, Random Forest, XGBoost	Random Forest: best accuracy (0.92)

Table 7: Overview of key work of fraud detection

Feature	Description
<i>payment_with_xrp_ratio</i>	Ratio of payment transactions with a specified XRP amount to total transactions.
<i>payment_without_xrp_ratio</i>	Ratio of payment transactions without a specified XRP amount to total transactions.
<i>offer_create_with_xrp_ratio</i>	Ratio of offercreate transactions with a specified XRP amount (TakerGets/TakerPays) to total transactions.
<i>offer_create_without_xrp_ratio</i>	Ratio of offercreate transactions without a specified XRP amount to total transactions.
<i>offer_cancel_ratio</i>	Ratio of offercancel transactions to total transactions.
<i>trust_set_ratio</i>	Ratio of trustset transactions to total transactions.
<i>escrow_ratio</i>	Ratio of escrow transactions (all escrow-related types) to total transactions.
<i>nftoken_create_offer_ratio</i>	Ratio of nftokencreateoffer transactions to total transactions.
<i>nftoken_cancel_offer_ratio</i>	Ratio of nftokencanceloffer transactions to total transactions.
<i>nftoken_accept_offer_ratio</i>	Ratio of nftokenacceptoffer transactions to total transactions.
<i>nftoken_mint_ratio</i>	Ratio of nftokenmint transactions to total transactions.
<i>total_transactions</i>	Count of all transactions (incoming and outgoing) involving the account.
<i>average_sent_amount_xrp</i>	Average XRP amount sent in payment transactions where the address is the sender.
<i>average_received_amount_xrp</i>	Average XRP amount received in payment transactions where the address is the receiver.
<i>median_sent_amount_xrp</i>	Median XRP amount sent in payment transactions where the address is the sender.
<i>median_received_amount_xrp</i>	Median XRP amount received in payment transactions where the address is the receiver.
<i>stddev_sent_amount_xrp</i>	Standard deviation of XRP amounts sent in payment transactions where the address is the sender.

<i>stddev_received_amount_xrp</i>	Standard deviation of XRP amounts received in payment transactions where the address is the receiver.
<i>payment_as_account_ratio</i>	Ratio of payment transactions where the address is the sender (account) to total transactions.
<i>payment_as_destination_ratio</i>	Ratio of payment transactions where the address is the receiver (destination) to total transactions.
<i>source_tag_present_ratio</i>	Ratio of transactions with a non-null SourceTag when the address is the sender to total transactions.
<i>destination_tag_present_ratio</i>	Ratio of transactions with a non-null DestinationTag when the address is the receiver to total transactions.
<i>unique_transaction_partner_ratio</i>	Ratio of unique transaction partners in payment transactions to total transactions.
<i>payment_small_amounts_sender_ratio</i>	Ratio of payment transactions under specific amounts (10, 20, 50, 100, 200, 300, 500 XRP) sent to total transactions.
<i>payment_small_amounts_receiver_ratio</i>	Ratio of payment transactions under specific amounts (10, 20, 50, 100, 200, 300, 500 XRP) received to total transactions.
<i>payment_large_amounts_sender_ratio</i>	Ratio of payment transactions over specific amounts (500, 1,000, 10,000, 100,000, 1,000,000 XRP) sent to total transactions.
<i>payment_large_amounts_receiver_ratio</i>	Ratio of payment transactions over specific amounts (500, 1,000, 10,000, 100,000, 1,000,000 XRP) received to total transactions.
<i>absolute_active_period</i>	Difference in days between the first and last transaction timestamps for the account, indicating the span of activity.
<i>active_days</i>	Count of unique days during which the account conducted at least one transaction.
<i>transaction_frequency</i>	Number of transactions over the active period of the account.
<i>unique_destination_partners</i>	Count of unique transaction partners where the address is the sender.
<i>unique_account_partners</i>	Count of unique transaction partners where the address is the receiver.
<i>transaction_ratio_20_largest_exchanges</i>	Ratio of transactions associated with the 20 largest exchanges to total transactions.
<i>total_sum_sent</i>	Sum of all outgoing transaction values for the account over the observed period.
<i>total_sum_received</i>	Sum of all incoming transaction values for the account over the observed period.

Table 8: Account level feature set

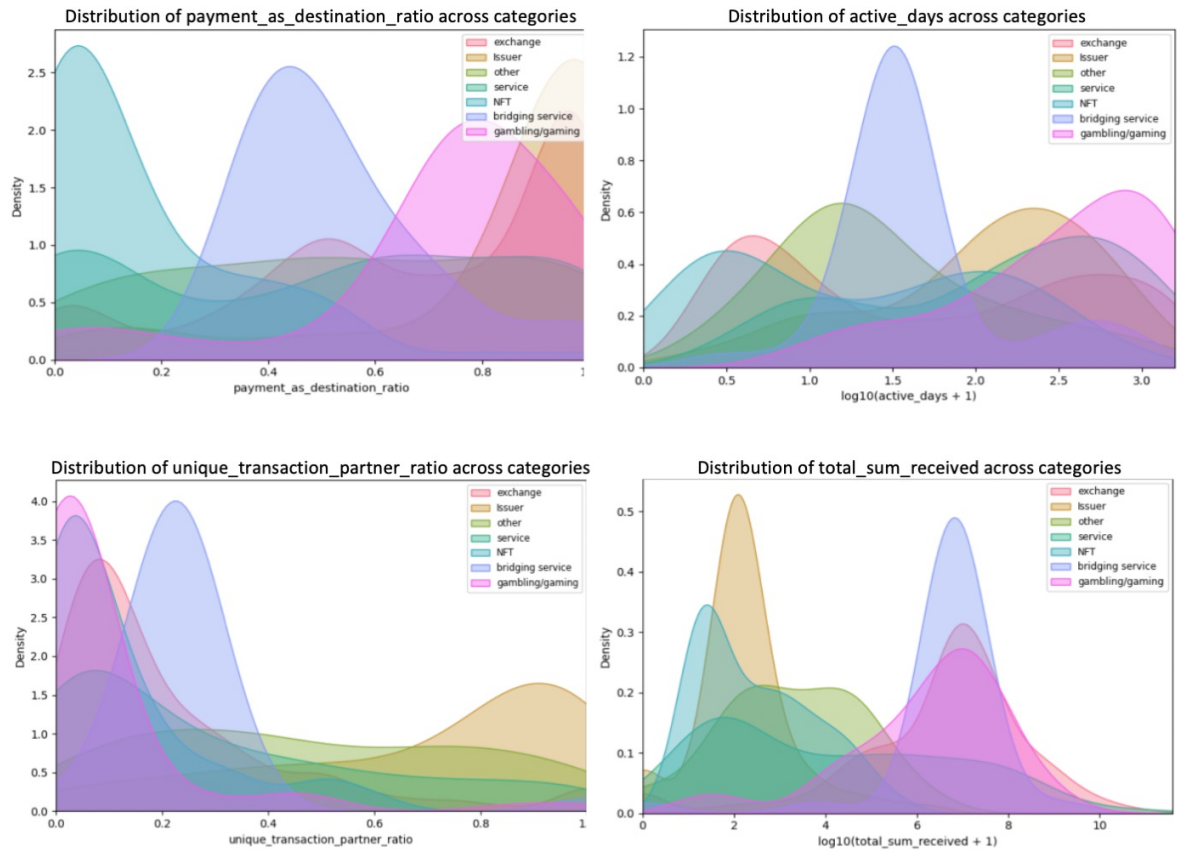


Figure 3: Distribution of various features across categories

Feature	H-statistic	p-value
<i>nftoken_mint_ratio</i>	1285.589967	1.43E-274
<i>payment_without_XRP_ratio</i>	788.167695	5.55E-167
<i>payment_over_1000_xrp_receiver_ratio</i>	630.5956361	5.85E-133
<i>nftoken_create_offer_ratio</i>	612.526185	4.63E-129
<i>payment_over_10000_xrp_receiver_ratio</i>	560.8910019	6.34E-118
<i>nftoken_cancel_offer_ratio</i>	538.617142	4.01E-113
<i>nftoken_accept_offer_ratio</i>	480.5705936	1.29E-100
<i>total_sum_received</i>	470.6342545	1.77E-98
<i>payment_over_1000_xrp_sender_ratio</i>	470.1915373	2.21E-98
<i>payment_over_500_xrp_sender_ratio</i>	443.8323629	1.04E-92
<i>payment_over_10000_xrp_sender_ratio</i>	437.7274626	2.15E-91
<i>trust_set_ratio</i>	410.7263861	1.38E-85
<i>payment_over_100000_xrp_receiver_ratio</i>	402.3527507	8.72E-84
<i>total_sum_sent</i>	394.8461681	3.58E-82
<i>destination_tag_present_ratio</i>	347.7990786	4.58E-72
<i>payment_over_100000_xrp_sender_ratio</i>	335.6828905	1.83E-69
<i>payment_over_1000000_xrp_sender_ratio</i>	287.5388858	3.82E-59
<i>offer_create_with_XRP_ratio</i>	257.1853796	1.19E-52
<i>unique_transaction_partner_ratio</i>	235.2377555	5.84E-48
<i>payment_over_1000000_xrp_receiver_ratio</i>	177.5972987	1.10E-35
<i>unique_account_partners</i>	162.461934	1.78E-32
<i>offer_cancel_ratio</i>	162.419805	1.82E-32
<i>payment_under_500_xrp_receiver_ratio</i>	136.9778122	4.35E-27
<i>ratio_known_counterparty_transactions</i>	136.1662466	6.45E-27
<i>total_transactions</i>	114.1729806	2.72E-22
<i>unique_destination_partners</i>	110.2187033	1.83E-21
<i>source_tag_present_ratio</i>	107.8585669	5.72E-21
<i>active_days</i>	103.4768496	4.72E-20
<i>payment_under_500_xrp_sender_ratio</i>	40.40788751	3.79E-07
<i>offer_create_without_XRP_ratio</i>	36.59046804	2.12E-06
<i>escrow_ratio</i>	14.77972275	0.022040934
<i>median_sent_amount_xrp</i>	NaN	NaN
<i>median_received_amount_xrp</i>	NaN	NaN
<i>stddev_received_amount_xrp</i>	NaN	NaN

Table 9: Kruskal-Wallis test results

Feature	NFT	Exchange	Gambling	Bridging Service	Issuer	Others	Service
<i>absolute_active_period</i>	-0.657	-0.046	-1.026	-0.122	0.061	0.060	0.124
<i>active_days</i>	0.294	-0.146	0.815	-0.038	0.008	-0.066	0.044
<i>average_received_amount_xrp</i>	-0.136	0.084	-0.406	0.003	-0.030	-0.016	-0.025
<i>destination_tag_present_ratio</i>	-1.338	0.213	1.001	-0.141	0.035	-0.103	-0.210
<i>escrow_ratio</i>	-0.220	-0.086		-0.011	-0.024	0.042	0.030
<i>median_received_amount_xrp</i>	-0.154	0.058	-0.949	-0.004	-0.028	-0.007	-0.009
<i>nftoken_accept_offer_ratio</i>	0.084	-0.057		-0.016	-0.152	0.038	-0.001
<i>nftoken_cancel_offer_ratio</i>	-0.030	0.003		-0.014	0.071	0.008	0.002
<i>nftoken_create_offer_ratio</i>	0.491	-0.102		-0.026	-0.050	-0.091	-0.022
<i>nftoken_mint_ratio</i>	0.844	-0.102		-0.044	-0.050	-0.104	-0.080
<i>offer_cancel_ratio</i>	0.083	-0.030		-0.039	0.054	0.031	0.005
<i>offer_create_with_XRP_ratio</i>	-0.049	-0.119		-0.045	-0.029	0.079	0.111
<i>offer_create_without_XRP_ratio</i>	-0.114	-0.019		-0.018	-0.020	-0.034	0.056
<i>payment_over_1000000_xrp_receiver_ratio</i>	-0.268	0.167	-1.109	-0.016	-0.049	-0.027	-0.033
<i>payment_over_1000000_xrp_sender_ratio</i>	-0.503	0.229	-2.319	-0.035	-0.081	-0.037	-0.042
<i>payment_over_100000_xrp_sender_ratio</i>	-0.685	0.193	-2.657	-0.036	-0.089	-0.042	-0.042
<i>payment_under_10_xrp_sender_ratio</i>	0.011	-0.016	0.304	0.054	0.006	0.091	-0.081
<i>ratio_known_counterparty_transactions</i>	-1.290	0.265	0.093	-0.040	-0.201	-0.050	-0.099
<i>source_tag_present_ratio</i>	-0.132	-0.005	-0.110	-0.039	-0.089	-0.010	0.078
<i>stddev_received_amount_xrp</i>	0.020	-0.064	-0.207	0.000	-0.029	-0.016	0.055
<i>total_sum_received</i>	-0.069	0.023	-0.462	0.000	-0.031	0.010	0.000
<i>total_sum_sent</i>	-0.118	0.117	-0.957	0.010	-0.042	0.017	-0.052
<i>total_transactions</i>	-0.389	-0.017	0.087	-0.010	-0.054	0.014	-0.040
<i>transaction_frequency</i>	-0.549	-0.020	0.034	-0.023	-0.055	-0.024	0.039
<i>trust_set_ratio</i>	0.157	-0.218		-0.029	0.058	0.052	0.115
<i>unique_account_partners</i>	-0.007	0.065	-0.355	0.020	-0.013	0.020	-0.027
<i>unique_destination_partners</i>	0.129	0.059	-0.118	-0.028	-0.071	-0.028	0.028
<i>unique_transaction_partner_ratio</i>	-0.268	-0.168	-0.976	-0.045	0.399	0.162	0.066

Table 10: Logistic regression results with features listed vertically and coefficients shown horizontally

Feature	NFT	Exchange	Gambling	Bridging Service	Issuer	Others	Service
<i>absolute_active_period</i>	0.518	0.955	0.358	0.885	1.063	1.062	1.132
<i>active_days</i>	1.342	0.864	2.259	0.963	1.009	0.936	1.046
<i>average_received_amount_xrp</i>	0.873	1.088	0.666	1.003	0.970	0.984	0.976
<i>destination_tag_present_ratio</i>	0.262	1.237	2.721	0.869	1.036	0.902	0.810
<i>escrow_ratio</i>	0.803	0.918	1.000	0.989	0.976	1.043	1.031
<i>median_received_amount_xrp</i>	0.857	1.059	0.387	0.996	0.972	0.993	0.991
<i>nftoken_accept_offer_ratio</i>	1.087	0.944		0.984	0.859	1.038	0.999
<i>nftoken_cancel_offer_ratio</i>	0.970	1.003		0.986	1.074	1.008	1.002
<i>nftoken_create_offer_ratio</i>	1.635	0.903		0.974	0.952	0.913	0.978
<i>nftoken_mint_ratio</i>	2.326	0.903		0.957	0.951	0.901	0.923
<i>offer_cancel_ratio</i>	1.087	0.970		0.962	1.056	1.032	1.005
<i>offer_create_with_XRP_ratio</i>	0.952	0.888		0.956	0.971	1.082	1.118
<i>offer_create_without_XRP_ratio</i>	0.892	0.981		0.982	0.980	0.967	1.057
<i>payment_over_1000000_xrp_receiver_ratio</i>	0.765	1.181	0.330	0.985	0.952	0.974	0.967
<i>payment_over_1000000_xrp_sender_ratio</i>	0.605	1.258	0.098	0.965	0.922	0.963	0.959
<i>payment_over_100000_xrp_sender_ratio</i>	0.504	1.213	0.070	0.965	0.915	0.959	0.959
<i>payment_under_10_xrp_sender_ratio</i>	1.011	0.984	1.356	1.056	1.006	1.095	0.922
<i>ratio_known_counterparty_transactions</i>	0.275	1.304	1.098	0.961	0.818	0.951	0.906
<i>source_tag_present_ratio</i>	0.876	0.995	0.896	0.962	0.914	0.990	1.081
<i>stddev_received_amount_xrp</i>	1.020	0.938	0.813	1.000	0.971	0.984	1.056
<i>total_sum_received</i>	0.934	1.023	0.630	1.000	0.969	1.010	1.000
<i>total_sum_sent</i>	0.889	1.124	0.384	1.011	0.959	1.017	0.949
<i>total_transactions</i>	0.678	0.983	1.091	0.990	0.947	1.014	0.961
<i>transaction_frequency</i>	0.577	0.980	1.034	0.977	0.946	0.977	1.039
<i>trust_set_ratio</i>	1.170	0.805	1.000	0.972	1.060	1.053	1.122
<i>unique_account_partners</i>	0.993	1.067	0.701	1.021	0.987	1.021	0.973
<i>unique_destination_partners</i>	1.137	1.061	0.889	0.973	0.931	0.972	1.029
<i>unique_transaction_partner_ratio</i>	0.765	0.845	0.377	0.956	1.490	1.176	1.068

Table 11: Logistic regression results with features listed vertically and odds ratios shown horizontally

Features for K-means and Gaussian Mixture
offer_create_with_xrp_ratio
offer_create_without_xrp_ratio
offer_cancel_ratio
trust_set_ratio
escrow_ratio
nftoken_create_offer_ratio
nftoken_cancel_offer_ratio
nftoken_accept_offer_ratio
nftoken_mint_ratio
total_transactions
average_sent_amount_xrp
average_received_amount_xrp
median_sent_amount_xrp
median_received_amount_xrp
stddev_sent_amount_xrp
stddev_received_amount_xrp
unique_transaction_partner_ratio
payment_small_amounts_sender_ratio
payment_small_amounts_receiver_ratio
payment_large_amounts_sender_ratio
payment_large_amounts_receiver_ratio
absolute_active_period
active_days
transaction_frequency
unique_destination_partners
unique_account_partners
total_sum_received
transaction_ratio_20_largest_exchanges

Table 12: Features used for K-Means clustering and Gaussian Mixture

Features for Agglomerative clustering
offer_create_with_xrp_ratio
offer_create_without_xrp_ratio
offer_cancel_ratio
trust_set_ratio
escrow_ratio
nftoken_create_offer_ratio
nftoken_cancel_offer_ratio
nftoken_accept_offer_ratio
nftoken_mint_ratio
total_transactions
average_sent_amount_xrp
average_received_amount_xrp

median_sent_amount_xrp
median_received_amount_xrp
stddev_sent_amount_xrp
stddev_received_amount_xrp
unique_transaction_partner_ratio
payment_small_amounts_sender_ratio
payment_small_amounts_receiver_ratio
payment_large_amounts_sender_ratio
payment_large_amounts_receiver_ratio
absolute_active_period
active_days
transaction_frequency
unique_destination_partners
unique_account_partners
transaction_ratio_20_largest_exchanges

Table 13: Features used for Agglomerative clustering

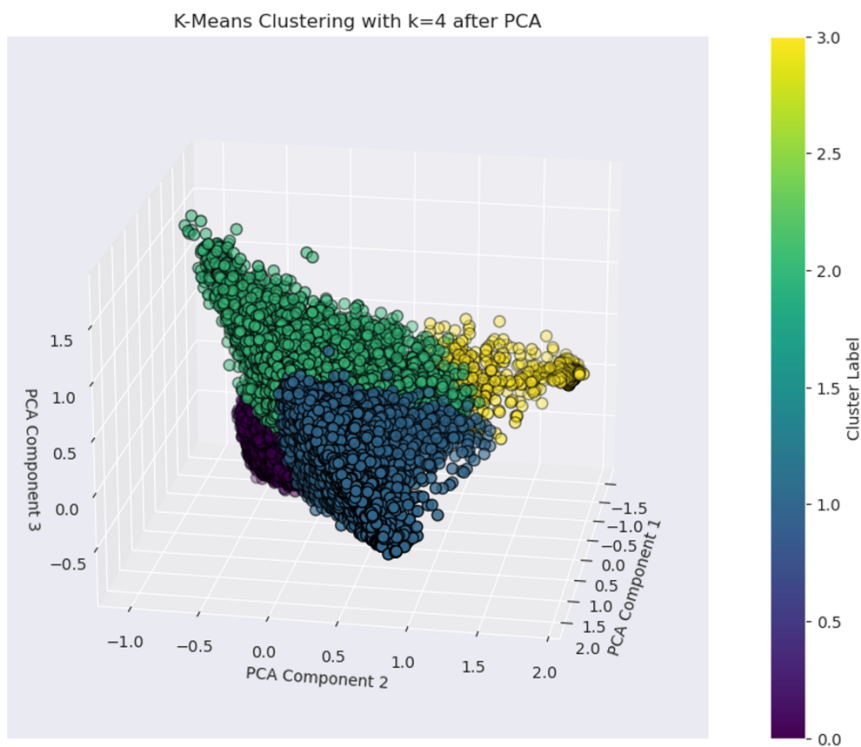


Figure 4: Visualization of the three principal components of K-Means with $k=4$

Metric	Agglomerative Clustering with PCA (k=5)	Gaussian Mixture with PCA (k=6)	K-Means with PCA (k=4)
Running Time (s)	452.97	6.16	5.67
Memory Usage (MB)	9.13	30.11	15.82

Table 14: Comparison of runtime and memory usage of the unsupervised models

Features for LightGBM

absolute_active_period
active_days
average_received_amount_xrp
average_sent_amount_xrp
destination_tag_present_ratio
median_received_amount_xrp
median_sent_amount_xrp
nitoken_create_offer_ratio
offer_cancel_ratio
offer_create_with_XRP_ratio
offer_create_without_XRP_ratio
payment_as_account_ratio
payment_over_100000_xrp_receiver_ratio
payment_over_100000_xrp_sender_ratio
payment_over_1000_xrp_receiver_ratio
payment_over_1000_xrp_sender_ratio
payment_over_10000_xrp_receiver_ratio
payment_over_10000_xrp_sender_ratio
payment_over_500_xrp_receiver_ratio
payment_over_500_xrp_sender_ratio
payment_under_10_xrp_sender_ratio
payment_under_100_xrp_receiver_ratio
payment_under_20_xrp_receiver_ratio
payment_under_200_xrp_receiver_ratio
payment_under_200_xrp_sender_ratio
payment_under_50_xrp_sender_ratio
payment_under_500_xrp_receiver_ratio
payment_without_XRP_ratio
ratio_known_counterparty_transactions
source_tag_present_ratio
stddev_received_amount_xrp
total_sum_received
total_sum_sent
transaction_frequency
trust_set_ratio
unique_account_partners
unique_destination_partners

unique_transaction_partner_ratio

Table 15: Final selection of features for LightGBM classifier

Metric	Best Random Forest	Best LightGBM
Training Time (s)	49.88	1.03
Inference Time per Instance (ms)	0.050	0.002
Peak Memory Usage (MB)	458.30	491.61

Table 16: Comparison of computational performance

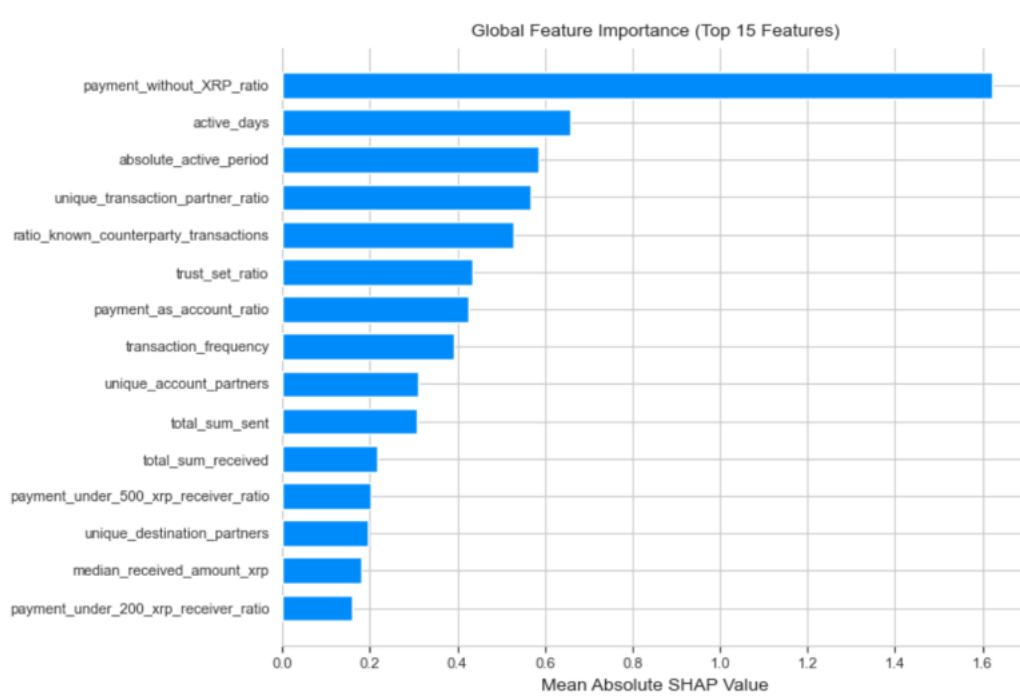


Figure 5: Global feature importance top 15 – LightGBM