



HAL
open science

Decoding Decentralized Finance Transactions through Ego Network Motif Mining

Natkamon Tovanich, Célestin Coquidé, Rémy Cazabet

► **To cite this version:**

Natkamon Tovanich, Célestin Coquidé, Rémy Cazabet. Decoding Decentralized Finance Transactions through Ego Network Motif Mining. The 13th International Conference on Complex Networks and their Applications, Dec 2024, Istanbul, Turkey. hal-04727895

HAL Id: hal-04727895

<https://hal.science/hal-04727895v1>

Submitted on 9 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Decoding Decentralized Finance Transactions through Ego Network Motif Mining

Natkamon Tovanich¹[0000-0001-9680-9282],
Célestin Coquidé²[0000-0001-8546-6587], and Rémy Cazabet²[0000-0002-9429-3865]

¹ CREST, CNRS, École Polytechnique, Institut Polytechnique de Paris,
91120 Palaiseau, France natkamon.tovanich@polytechnique.edu

² Univ Lyon, UCBL, CNRS, LIRIS, UMR 5205, 69622 Villeurbanne, France
celestin.coquide@liris.cnrs.fr, remy.cazabet@gmail.com

Abstract. Decentralized Finance (DeFi) is increasingly studied and adopted for its potential to provide accessible and transparent financial services. Analyzing how investors use DeFi is important for reaching a better understanding of their usage and for regulation purposes. However, analyzing DeFi transactions is challenging due to often incomplete or inaccurate labeled data. This paper presents a method to extract ego network motifs from the token transfer network, capturing the transfer of tokens between users and smart contracts. Our results demonstrate that smart contract methods performing specific DeFi operations can be efficiently identified by analyzing these motifs while providing insights into account activities.

Keywords: transaction network, ego network, network motifs, graph mining, blockchain data, decentralized finance

1 Introduction

Decentralized finance (DeFi) is a financial infrastructure implemented on smart contracts in blockchains [4]. DeFi replicates traditional financial services without centralized intermediaries. The building blocks of DeFi applications consist of tokens, protocols, and aggregators. Tokens are digital assets created on the blockchain that represent value or utility. Tokens can be transferred from any sender account (i.e., a user or a protocol) to a recipient account. In DeFi protocols, tokens provide specific financial services, such as borrowing, lending, and exchanging tokens. The user creates a “transaction” by calling a “method” (later, DeFi method) in the protocol’s smart contract to perform a specific DeFi operation. A single DeFi method can involve one or multiple token transfers. These protocols can be “composed” to create more complex financial products and strategies [13]. For instance, aggregators are platforms that integrate multiple DeFi protocols, typically allowing their clients to manage complex operations automatically to optimize their performance, e.g., automatically choosing the most favorable exchange rate or lending rate among multiple platforms.

As DeFi applications continue to grow and evolve rapidly, understanding the nature and usage of its protocols and analyzing user activities becomes increasingly important. Since DeFi trading involves tokens that an account spends or receives in a transaction, we use ego network motifs to represent these interactions. This approach also offers a more efficient way to collect transfer data related to specific accounts. This work addresses two research questions which, to the best of our knowledge, have not been analyzed in previous studies.

1. Can we infer DeFi methods (e.g., depositing, borrowing, swapping, etc.) from the ego token transfer network?
2. Can we identify users with similar activities using ego network motifs?

Analyzing DeFi methods that perform a specific operation within Ethereum transactions usually requires data in which the DeFi methods are labeled, obtained from Etherscan, private companies, or smart contract interfaces (Application Binary Interfaces, ABIs). While our study utilizes these DeFi method labels obtained through self-reported or manual labeling, we recognize this as a limitation, particularly when dealing with transactions that have missing or inaccurate labels. By leveraging ego network motifs, our approach aims not only to provide a more accurate and interpretable tool for inferring DeFi methods but also to analyze trading activities of accounts within DeFi ecosystems, even when labeled data is incomplete or unavailable.

2 Related Work

The study of blockchain networks has gained significant interest due to the rise of cryptocurrencies and DeFi applications. Many works study macroscopic properties, such as degree distribution or clustering coefficient, in various network representations [12,21]. Research on meso and microscopic behaviors is emerging, focusing on illegal activities in Bitcoin (e.g., [1,6]) and Ethereum (e.g., [7,8,19]).

Network motifs play a crucial role in characterizing different groups of behaviors within these networks. Kosyfaki et al. proposed an algorithm to identify significant flow motifs in the large Bitcoin user network [14], while Ba et al. defined triadic closure network motifs and explored their impact on the evolution mechanisms of platforms [5]. Temporal motifs have been employed in various tasks, including detecting mixing services in Bitcoin [20], analyzing dark web marketplaces and NFT wash tradings [3], detecting phishing scam accounts [18], and identifying account communities [2,17]. Building on these works, we introduce ego network motifs, demonstrating their utility in extracting DeFi method signatures and characterizing account behaviors.

In addition to behavioral analysis, our research contributes to understand DeFi trading strategies. Zhou et al. proposed two graph algorithms to identify cyclic arbitrage in DeFi tradings and calculate the optimal threshold of profitable transactions [22]. The majority of the work in this domain has focused on Maximal Extractable Value (MEV), which measures the maximum profit that can be extracted by reordering, including, or excluding transactions within a

block [9,16]. EigenPhi is a data platform that tracks these transactions in real time and provides graphs to analyze profitability.¹ Park et al. developed a graph neural network model to detect MEV transactions from token transfer graphs without relying on smart contract code or ABIs [15].

Our work extends this body of research by proposing a novel motif extraction technique to infer DeFi methods. These motifs serve as building blocks for analyzing more complex interactions, including arbitrage and MEV transactions.

3 Data Description

In this work, we focused on a subset of Ethereum accounts labeled on Etherscan as likely trading with decentralized finance (DeFi) protocols. We obtained the list of all “fund” accounts from Etherscan² as of March 2023, which consists of 65 entities. Among these, 25 are identified as associated with Alameda Research, a cryptocurrency trading firm linked to FTX that was active in DeFi before the collapse of both entities in November 2022 due to massive fraud. This subset was chosen to examine the DeFi interactions of notable trading accounts rather than to analyze the entire blockchain. Although this list does not cover all accounts, it provides a focused perspective on significant players in the DeFi space.

Token transfer dataset: A transfer is an action of moving Ether (ETH) or tokens from one account to another. Each transaction can trigger one or multiple transfers. An account can be either an address controlled by a user account (Externally Owned Account, EOA) or a smart contract (CA). The special address 0x00000..., known as NULL, is commonly used for minting and burning tokens.

To construct the ego transfer network (ETN), token transfer data can be collected from Transfer API from Alchemy³. This API provided detailed records of all token transfers sent or received by our selected accounts, including transaction hashes, sending and receiving accounts, token names, and transfer amounts. We collected token transfer data from the 65 selected accounts, resulting in 1,598,098 transfers across 1,131,723 transactions. Fig. 1 describes key elements of the dataset: Fig. 1 (a) shows, for each account, its number of transactions, number of different tokens used, and fraction of transactions of unknown nature. Fig. 1 (b) describe the number of token transfers per transaction. We observe that most transactions consist of a single token transfer (78%), while more than a hundred thousand transactions have more than one single token transfer.

Token list: Each blockchain has its native cryptocurrency, such as Bitcoin (BTC) on Bitcoin, Ether (ETH) on Ethereum, and MATIC on Polygon, allowing users to transfer tokens within the network. Besides its native token, Ethereum blockchain allows the creation of other tokens by implementing smart contracts that follow standards like ERC-20 (fungible), ERC-721 (non-fungible or NFT), and ERC-1155 (multi-token). Fungible tokens can be divided into

¹ <https://eigenphi.io/>

² <https://etherscan.io/accounts/label/fund>

³ <https://docs.alchemy.com/reference/transfers-api-quickstart>




Fig. 1. (a) Fund accounts in our dataset and (b) Token transfers histogram

varied-price cryptocurrencies (e.g., ETH, DODGE, SHIB) and stablecoins (e.g., USDC, USDT, DAI), which are pegged to a fiat currency like the US Dollar. Synthetic tokens are created to represent underlying assets, often through mechanisms such as wrapping, staking, or collateralization within DeFi protocols.

We listed all tokens in our token transfer dataset and applied the following data processing steps to filter out spam tokens and assign categories:

1. **Assigning token categories:** We manually verified the top 500 tokens in our dataset based on the number of transactions and assigned their category based on Messari’s classification⁴.
2. **Filtering out spam tokens:** Spam tokens are created to deceive the recipient into trading or interacting with valueless or harmful contracts. These spam tokens usually replicate the name and symbol of legitimate tokens or well-known brands. We checked if the token was on Etherscan’s list of verified tokens labeled⁵. If a token was not in this list, we removed transactions where such a spam token was transferred from our analysis.
3. **Unlabeled tokens:** Remaining tokens were annotated as “Unlabeled”. Although their specific purpose is unknown, we included them in our analysis since they are not considered as spam tokens.

Table 1 lists all token categories along with the number of tokens, transactions, and the top 5 tokens in each category. After excluding transactions involving spam tokens, we have 1,598,098 transfers across 1,095,374 transactions.

Labeled transaction dataset: So far, we know token transfers within the transaction but don’t know the DeFi method. Thus, we collected the transaction headers from Etherscan’s transaction list pages for each account⁶. The details of any transaction include the transaction hash, DeFi method, block number, timestamp, source and target accounts, ETH amount, and transaction fees.

⁴ List of tokens: <https://messari.io/assets?view=e88a221b-cb05-4995-96a7-cf2c5e217e79>
Classification scheme: <https://docs.messari.io/docs/messari-classification-system>

⁵ <https://etherscan.io/tokens>

⁶ e.g., <https://etherscan.io/txs?a=0x0f4ee9631f4be0a63756515141281a3e2b293bbe>

Table 1. Token Categories

Token Category	#Tokens (%)	#Txs (%)	Top-5 Tokens (with %Txs)
Cryptocurrency	17 (2.04%)	483,038 (35.94%)	ETH (17.70%), WETH (14.43%), WBTC (2.14%), SHIB (1.40%), PAXG (0.12%)
Stablecoin	27 (3.23%)	385,452 (28.68%)	USDT (15.03%), USDC (8.93%), BUSD (1.76%), DAI (1.14%), TUSD (0.46%)
Marketplace	99 (11.86%)	223,954 (16.66%)	COMP (1.50%), UNI (1.46%), SNX (1.19%), CRV (1.04%), CHI (1.01%)
Other	100 (11.98%)	98,927 (7.36%)	LINK (1.50%), ENS (0.97%), GRT (0.63%), NMR (0.52%), CVX (0.35%)
NFT & Metaverse	57 (6.83%)	59,482 (4.43%)	MANA (0.63%), APE (0.58%), CHZ (0.56%), AXS (0.53%), LOOKS (0.30%)
Network	29 (3.47%)	35,737 (2.66%)	MATIC (1.12%), FTM (0.65%), OMG (0.22%), INJ (0.20%), LRC (0.10%)
Financial Service	40 (4.79%)	19,794 (1.47%)	CEL (0.32%), BADGER (0.20%), SXP (0.20%), RAY (0.16%), MTA (0.12%)
Synthetic	339 (40.60%)	18,310 (1.36%)	stETH (0.17%), WNXM (0.13%), UNI-V3-POS (0.07%), XRP-BULL (0.06%), variableDebtUSDT (0.05%)
Bridge	12 (1.44%)	16,279 (1.21%)	REN (0.65%), SYN (0.15%), QNT (0.12%), T (0.10%), NU (0.08%)
Unlabeled	115 (13.77%)	3,146 (0.23%)	MIC (0.03%), AKRO (0.03%), XCN (0.01%), POLY (0.01%), EMB (0.01%)

However, we were able to retrieve DeFi method labels for 55% of transactions, due to technical constraints of the Etherscan web crawling (limiting access to the latest 100,000 transactions per account and excluding transactions not created by the selected accounts).

DeFi method list: The *method* column in the labeled transaction dataset specifies the DeFi method. We constructed the ground truth dataset from DeFi methods occurring in at least 100 transactions, manually grouping each DeFi method based on its label. Table 2 presents the method groups and their full names. We merged the “Exchange” and “Swap” method groups as well as the “Redeem” and “Withdraw” groups since they perform the same function.

Table 2. Method Groups

Group	# Txs	Full Method Names with ≥ 100 Txs
Transfer	404,130	Complete Transfer, Safe Transfer From, Transfer, Transfer From, Transfer Tokens
Swap	112,387	Swap: Any Swap Out Underlying, Batch Eth Out Swap Exact In, Exact Input Single, Simple Swap, Swap, Swap ETH For Exact Tokens, Swap Erc20, Swap Exact ETH For Tokens, Swap Exact Tokens For ETH, Swap Exact Tokens For Tokens, Uniswap V3Swap, Exchange: Exchange, Exchange underlying
Withdraw	3,619	Withdraw: Remove Liquidity ETH With Permit, Remove liquidity one coin, Withdraw, Withdraw Erc20, Redeem: Redeem, Redeem Underlying
Deposit	3,325	Add Liquidity, Add Liquidity ETH, Deposit, Deposit ETH, Deposit For
Claim Reward	2,881	Claim, Claim Comp, Claim Reward, Claim Rewards, Claim Token, Get Reward
Borrow	1,389	Borrow
Repay	1,256	Repay, Repay Borrow
Mint	1,189	Mint, Mint many
Exit	276	Exit
Burn	241	Burn
Stake	215	Stake

We chose only the 8 DeFi method groups with more than 1,000 transactions in our dataset, ensuring a sufficient sample size for our classification models.

The final dataset results in 530,170 transactions (33% of transactions in the token transfer dataset) with labeled DeFi method groups. The missing method data highlights the limits of relying only on labeled data, as is common in the literature, and thus the interest of our approach, which proposes to automatically infer functions from the transaction itself, even without the DeFi method label.

4 Ego Network Motifs

For each transaction, we construct the ego transfer network (ETN), represented as a directed heterogeneous graph. Each account is represented as a node with an attribute denoting its account type. The account types can be Ego (E), address (A), contract (C), or the NULL address (N). Each edge represents the token transfer directed from one account to another, with the edge attribute indicating the token category (see Table 1). We extract two features for each ego network: motif counts (denoted as M) and edge lists (E).

Motif counts: Ego motifs from the ETN are directed subgraphs consisting of 2 or 3 nodes, represented in Fig. 2. These motif sizes are chosen based on the average of 2.09 accounts ($SD = 0.40$) per transaction. Note that, as our dataset records ego transfers—thus without edges between non-ego nodes—there are only 8 possible directed ego network motifs (out of 15 motifs in full networks). The top node is always an ego node (E) connected to other nodes (denoted i, j).




Fig. 2. List of all possible ego network motifs

We use the VF2++ algorithm [11] in NetworkX [10], to perform subgraph isomorphism matching to identify and count each occurrence of the motifs within the ETN. We then check the node type for each motif found and represent it as a string with the pattern $\{motifID\}(E, i, j)$, where $\{motifID\} \in \{m1, m2, \dots, m8\}$, with motif IDs corresponding to those in Fig. 2, while $i, j \in \{A, C, N\}$.

Edge lists: Edge list indicates which token an ego received/spent from/to which account type. We obtain the edge list from the ego network with the token category. The edge count is represented as a vector $(S, T)\{token\ category\}$, where S is the source node type and T is the target node type. Three examples of ETNs are presented in Fig. 3, along with motif counts and edge list vectors.

5 DeFi Method Classification

We use supervised learning methods to test whether motif counting and network features can infer DeFi methods (e.g., deposit, borrow, swap). To determine




Fig. 3. Examples of transaction description by motifs and edge lists

which motif features can distinguish different DeFi methods, we tested four sets of training features: ego motifs with node types (M), edge list with token types (E), the combination of ego motifs and edge list ($M + E$), and the concatenation between motifs and edge list features within each motif ($M \times E$).

We trained three classification models: logistic regression (LR), decision tree (DT), and random forest (RF). For the LR model, we perform one-vs-all multi-class classifiers. To prevent overfitting while capturing patterns in the minority class, we set the minimum leaf node to 10 samples in both DT and RF. Given the dataset’s high skew towards the most common DeFi methods (e.g., transfers and swaps in Table 2), we assigned inversely proportional weights to the minority methods in the classification models to mitigate the class imbalance problem.

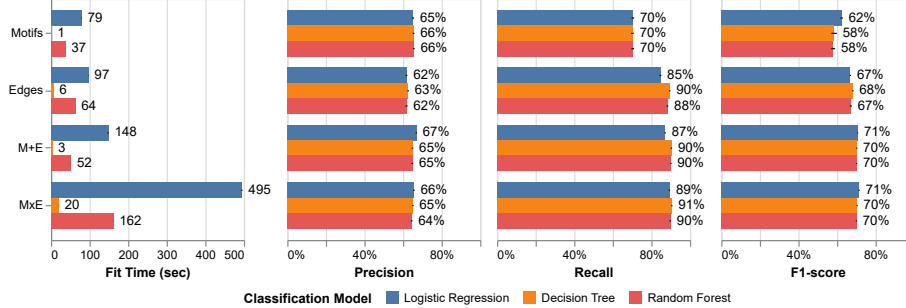


Fig. 4. Comparison of classification results with 10-fold CV scores

We evaluated the classification models using stratified 10-fold cross-validation (10-fold CV), ensuring that each fold contained the same distribution of classes as the training set. The average F1-score, precision, and recall are reported in Fig. 4, as accuracy is not a reliable metric for highly skewed classes. The results indicate that using motifs with node and edge types can infer DeFi methods with moderate precision and high recall and outperform models using only motifs or edges. The best model is LR with $M + E$ and $M \times E$ features, achieving an F-score of 71%, followed by the DT and RF model with 70%. The $M \times E$ models did not significantly improve the F1-score compared to $M + E$ models, despite having a larger number of features.

We chose to investigate the DT model with $M+E$ features further because it provides a simpler and more interpretable model while its performance is very close to that of the best-performing model. Focusing on $M+E$ models, LR has slightly higher precision and F1-Score, while DT and RF had 3% higher recall.

The confusion matrix in Fig. 5 shows the number of samples predicted compared to the actual DeFi method. The color scale indicates the percentage of the true method for each predicted label (row-wise), while the circle size is proportional to the percentage predicted compared to the true class (column-wise). DT model achieves high recall for most DeFi methods, although it tends to mispredict transfer as deposit and repay, and swap as mint and withdraw, indicating the similarity in motif features and resulting in lower precision for these methods.

The classification results confirm that ego network motifs with account and token types are important features for inferring DeFi methods.

6 Interpretable Method Signatures

While the previous section has shown that one can use machine learning to classify DeFi methods from ego network motifs, the models remain a black box. In this section, we study which specific motif can be associated with which DeFi method to better understand the nature of tradings and make decisions more interpretable. We used the decision tree (DT) model as a reference because it offered results comparable to the best-performing logistic regression (LR) model while achieving higher recall and maintaining interpretability. However, the full DT model is complex, making it challenging to interpret all the rules.

We performed post-pruning to reduce the tree’s complexity and obtain a manageable set of rules while retaining predictive performance, thus simplifying the model. Fig. 6 (a) shows post-pruning decision trees with varying cost-complexity pruning alphas and 10-fold CV scores for each number of leaf nodes. We chose a pruned tree with 18 leaf nodes ($\alpha = 0.007326$), depicted in Fig. 6 (b), while (c) shows the predicted class for each leaf node, with probabilities represented by bars and the number of samples (on a log scale) indicated by a gray line.

We adopt a frequent pattern mining approach to find a maximal set of motif features with a support of more than 80%, i.e., features that appear in over 80%

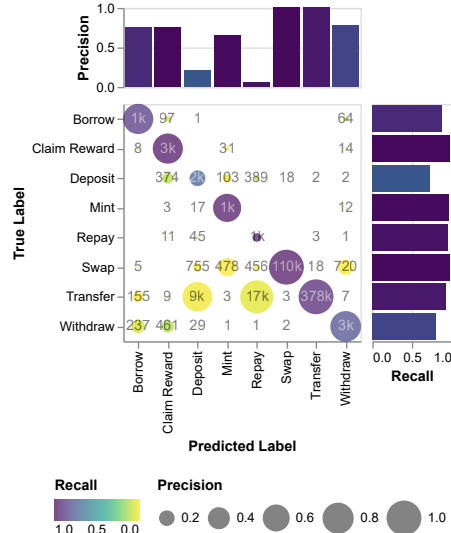


Fig. 5. Confusion metric from the decision tree model using $M + E$ features

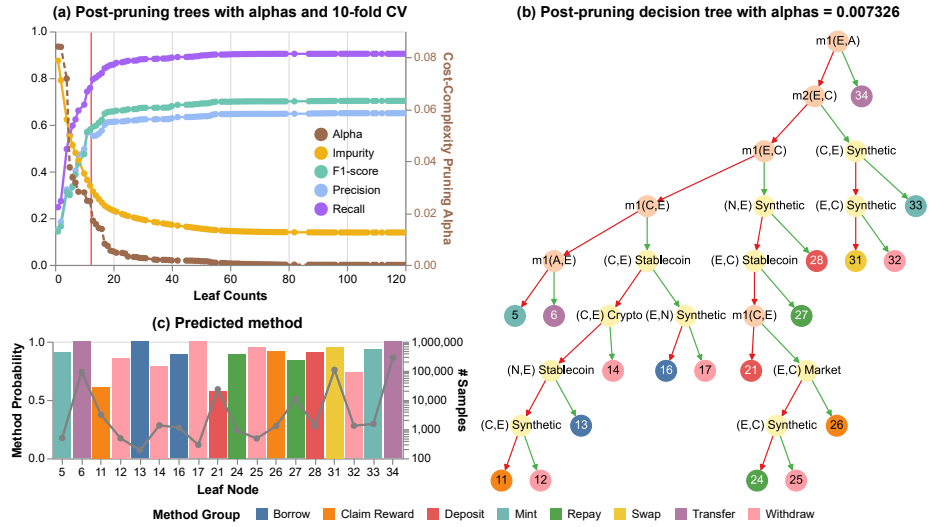


Fig. 6. (a) Post-pruning decision tree with varying alphas, leaf nodes, and 10-fold cross-validation metrics; (b) Pruned tree selected with $\alpha = 0.007326$; (c) Predicted DeFi method at each leaf node (color), along with probability (bars) and sample size (line).

of the instances at that leaf node. For each leaf node, we start by listing motif features with support greater than the threshold, incrementally adding more features until no additional features meet the support criteria. We selected the longest frequent itemset as the signature motif of the leaf and used it to analyze token transfer networks of unidentified transactions.

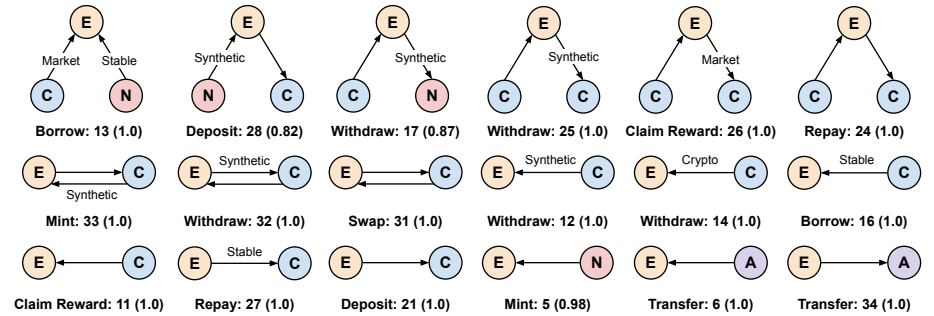


Fig. 7. List of DeFi method signature motifs on the leaf node with the predicted method. The support value for each motif is provided in parentheses. *Note: Crypto: Cryptocurrency, Stable: Stablecoin, Market: Marketplace tokens.*

Fig. 7 displays the DeFi method signatures in 18 leaves of the pruned tree. Most of the signature motifs reflect the functionality of the labeled methods. The deposit (28) pattern indicates that an ego deposits tokens into a contract and receives synthetic tokens as collateral. Conversely, the withdraw method (17,

25) patterns show that an ego must burn or return synthetic tokens to withdraw their deposited tokens from the contract. The swap method (31) pattern shows the exchange of tokens between the protocol’s contract and the account. However, if the user receives (or sends) synthetic tokens to the contract, it corresponds to the mint (33) or withdraw (32) methods, respectively. These signature motifs are interpretable, offering insights into transaction mechanisms while also being useful for investigating transactions with unknown DeFi methods. However, we found that some signature motifs are overly specific to particular token categories, even though they could potentially apply to other tokens as well. For instance, the borrow (13, 16) patterns suggest that accounts (egos) in this dataset mostly borrowed stablecoins. These results are not yet generalizable as we trained on a subset of Ethereum transactions. However, we expect the results to become more precise and applicable when we either filter for known contracts associated with DeFi protocols or use a larger dataset.

7 Profiling Ethereum Accounts

We demonstrate that extracted signature motifs can also be utilized to analyze DeFi activities at the account level. For each account, we create a profile by first counting the transactions associated with each DeFi method signature and then normalizing these counts by account. Fig. 8 shows a hierarchical cluster map of accounts and method-leaf. Hierarchical clustering helps us detect groups of accounts with similar patterns of activity. Since some DeFi methods, like transfers, are over-represented in all accounts, we show z-score normalization.




Fig. 8. Heat map and hierarchical clustering on fund account profiles

We use the silhouette score to determine the number of clusters ($k = 8$, silhouette = 0.42). These 8 clusters are represented by different colors in the first row of the heat map in Fig. 8, from Cluster 1 to Cluster 8 (left to right).

Cluster 1 ■ has a higher average z-score for transfer methods, typically receiving tokens from addresses (leaf 6), while Cluster 2 ■ shows activity in both receiving and spending tokens (leaves 6 and 34). These accounts likely trading with off-chain exchange companies. Cluster 8 ■ includes accounts primarily interacting with DeFi protocols, as indicated by a higher z-score for various methods, such as claim reward, withdraw, borrow, and repay. Smaller clusters, with fewer than five accounts, are considered specialized. For instance, Cluster 5 ■ is characterized by higher token swap activities, suggesting interaction with decentralized exchanges. Cluster 4 ■ is involved in token minting, while Cluster 6 ■ can be identified as liquidity providers, focusing on deposit activities.

This section shows how motif signatures can be used to characterize accounts in the ego transfer network (ETN), despite unlabeled DeFi methods. These insights help us better understand the different groups of activities in our dataset.

8 Conclusion and Future Work

We demonstrated the effectiveness of using ego network motifs to represent and decompose token transfers in Ethereum transactions. By leveraging these motifs as fundamental building blocks, we successfully inferred DeFi methods and extracted method signatures to investigate transactions lacking labeled data. These motifs provide a powerful tool for analyzing account activities, allowing us to characterize different groups of accounts based on their DeFi activities.

The implications of our findings extend beyond academic research. For companies operating in the DeFi space, this approach offers a robust method for monitoring account activities and detecting trading strategies and anomaly behaviors. Researchers can apply these techniques to analyze the evolution of DeFi protocols and account activities over time, investigating how they respond to market fluctuations and shock events.

In future work, we plan to apply our method to the full transaction graph to break down complex tradings, such as flash loans and arbitrage. We will extract larger motif sizes on an expanded dataset to improve accuracy and motif generalization. Additionally, we aim to incorporate temporal and higher-order network motifs to describe the dynamics and trading strategies of ego networks. This line of work will deepen our understanding of trading strategies and the evolution of the DeFi ecosystem, helping to identify emerging use cases without needing fully labeled data.

Acknowledgement This work was supported by the CHIST-ERA grant FairOn-Chain, the Agence Nationale de la Recherche (ANR-23-CHRO-0002), and UCBL’s grant BIT-STABLENET. We thank Stefan Kitzler for his constructive feedback.

References

1. Akcora, C.G., Li, Y., Gel, Y.R., Kantarcioglu, M.: BitcoinHeist: Topological data analysis for ransomware detection on the Bitcoin blockchain. arXiv:1906.07852 (2019)
2. Ao, X., Liu, Y., Qin, Z., Sun, Y., He, Q.: Temporal high-order proximity aware behavior analysis on Ethereum. *World Wide Web* **24**(5), 1565–1585 (2021)

3. Arnold, N.A., Zhong, P., Ba, C.T., Steer, B., Mondragon, R., Cuadrado, F., Lambiotte, R., Clegg, R.G.: Insights and caveats from mining local and global temporal motifs in cryptocurrency transaction networks. *arXiv:2402.09272* (2024)
4. Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., Victor, F.: The technology of decentralized finance (DeFi). *Digital Finance* **6**(1), 55–95 (2024)
5. Ba, C.T., Zignani, M., Gaito, S.: Characterizing growth in decentralized socio-economic networks through triadic closure-related network motifs. *Online Soc. Netw. Media* **37**, 100,266 (2023)
6. Bellei, C., Xu, M., Phillips, R., Robinson, T., Weber, M., Kaler, T., Leiserson, C.E., Chen, J.: The shape of money laundering: Subgraph representation learning on the blockchain with the Elliptic2 dataset. *arXiv:2404.19109* (2024)
7. Chen, H., Pendleton, M., Njilla, L., Xu, S.: A survey on Ethereum systems security: Vulnerabilities, attacks, and defenses. *ACM Comput. Surv.* **53**(3), 1–43 (2020)
8. Chen, L., Peng, J., Liu, Y., Li, J., Xie, F., Zheng, Z.: Phishing scams detection in Ethereum transaction network. *ACM Trans. Internet Technol.* **21**(1), 1–16 (2020)
9. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. *arXiv:1904.05234* (2019)
10. Hagberg, A.A., Schult, D.A., Swart, P.J.: Exploring network structure, dynamics, and function using networkx. In: *Proc. 7th Python in Science Conf.*, pp. 11 – 15. Pasadena, CA USA (2008)
11. Jüttner, A., Madarasi, P.: Vf2++—an improved subgraph isomorphism algorithm. *Discrete Appl. Math.* **242**, 69–81 (2018)
12. Khan, A.: Graph analysis of the Ethereum blockchain data: A survey of datasets, methods, and future work. In: *IEEE Int. Conf. Blockchain*, pp. 250–257 (2022)
13. Kitzler, S., Victor, F., Saggese, P., Haslhofer, B.: Disentangling decentralized finance (DeFi) compositions. *ACM Trans. Web* **17**(2), 1–26 (2023)
14. Kosyfaki, C., Mamoulis, N., Pitoura, E., Tsaparas, P.: Flow motifs in interaction networks. *arXiv:1810.08408* (2018)
15. Park, S., Jeong, W., Lee, Y., Son, B., Jang, H., Lee, J.: Unraveling the MEV enigma: ABI-free detection model using graph neural networks. *Future Gener. Comput. Syst.* **153**, 70–83 (2024)
16. Qin, K., Zhou, L., Gervais, A.: Quantifying blockchain extractable value: How dark is the forest? In: *IEEE Symp. Security Privacy*, pp. 198–214 (2022)
17. Tovanich, N., Cazabet, R.: Fingerprinting bitcoin entities using money flow representation learning. *Appl. Netw. Sci.* **8**(1), 63 (2023)
18. Wang, Y., Wang, H., Lu, X., Zhou, L., Liu, L.: Detecting Ethereum phishing scams with temporal motif features of subgraph. In: *IEEE Symp. Comput. Commun.*, pp. 631–636 (2023)
19. Wu, J., Lin, D., Fu, Q., Yang, S., Chen, T., Zheng, Z., Song, B.: Towards understanding asset flows in crypto money laundering through the lenses of Ethereum heists. *IEEE Trans. Inf. Forensics Security* (2023)
20. Wu, J., Liu, J., Chen, W., Huang, H., Zheng, Z., Zhang, Y.: Detecting mixing services via mining Bitcoin transaction network with hybrid motifs. *IEEE Trans. Syst., Man, Cybern. Syst.* **52**(4), 2237–2249 (2021)
21. Wu, J., Liu, J., Zhao, Y., Zheng, Z.: Analysis of cryptocurrency transactions from a network perspective: An overview. *J. Netw. Comput. Appl.* **190**, 103,139 (2021)
22. Zhou, L., Qin, K., Cully, A., Livshits, B., Gervais, A.: On the just-in-time discovery of profit-generating transactions in DeFi protocols. In: *IEEE Symp. Security Privacy*, pp. 919–936 (2021)