An In-depth Analysis on the Broken Ties on Twitter

Kiana Kheiri
Utah State University
kiana.kheiri@usu.edu

Thomas Kerby
Utah State University
thomas.kerby@usu.edu

Tyler Derr
Vanderbilt University
Tyler.Derr@vanderbilt.edu

Hamid Karimi
Utah State University
hamid.karimi@usu.edu

***Abstract*—Social ties are at the heart of online social networks, enabling users to exchange information, communicate, share content, and build communities. However, an under-explored aspect of these networks is the dissolution of online relationships, which complements the studies of tie maintenance and formation. A comprehensive understanding of these connections, including their formation, dissolution, and potential prediction of breakdowns, can provide a more detailed view of the network's dynamics and the evolution of interpersonal ties. However, a notable barrier to studying these broken ties is the lack of longitudinal, detailed data. This paper aims to address this gap by creating a large-scale dataset of more than 120K Twitter users over a period of 15 weeks (weekly snapshots). With this dataset, we undertake an extensive analysis of links on Twitter. Our investigation includes a range of features that span five distinct categories on the Twitter social graph. These include structural features such as centrality, content aspects like post polarity, user profile characteristics (e.g., verified status), egocentric network elements like reciprocity, and dense user representation, e.g., node2vec. Next, we conduct a thorough analysis of these features to identify meaningful patterns. Ultimately, through extensive experimentation, we employ several machine learning algorithms to discern the impact of the extracted features on the prediction of broken ties.**

***Index Terms*—Broken Ties, Dissolved Ties, Unfollow, Twitter, Social Media**

I. Introduction

In today's digital age, online social ties, a frequently under-appreciated component, form the cornerstone of Online Social Networks (OSN). Establishing these connections empowers users to communicate, build relationships, share and acquire information, promote businesses, and participate in numerous other activities. Consequently, OSNs have intricately woven themselves into human life, linking millions, or even billions, of users through myriad connections. However, an aspect of these ties that is equally significant but less examined is their dissolution, commonly known as 'unfollowing' or 'unfriending.'[1] There are several reasons why studying the broken ties in OSNs is crucial: Firstly, unlike tie formation, breakage is a conscious and rational action based on past interactions [1]. Secondly, the disintegration of these ties can cause substantial changes in the underlying online network structure, potentially disrupting information diffusion or escalating polarization [2]. Lastly, severed ties can affect interpersonal relationships beyond online social media [3], [4]. Therefore, it is essential to delve deeper into our understanding of social ties, mainly focusing on their breakage. In doing so, we acknowledge the importance of these ties in creating connections and in their termination, reflecting the dynamic nature of relationships in the digital world.

[1]In this paper, the terms 'broken tie', 'dissolved tie' and 'unfollow' interchangeably. The same is true for 'maintained/formed tie' and 'follow'.

Nonetheless, the investigation of broken ties introduces several obstacles. The first challenge lies in the identification of instances of broken ties, which requires longitudinal network data. To more accurately determine the time of tie dissolution, this data needs to be of a high temporal resolution. The second challenge is that the maintenance and dissolution of ties depend on various structural and behavioral factors. This complexity necessitates meticulous data collection and feature extraction. The third challenge arises from the necessity of conducting an exhaustive, large-scale data analysis to identify informative patterns about broken ties, an approach that existing studies have not fully implemented [5]. Finally, although predicting broken ties holds substantial value, current research has not sufficiently emphasized developing a robust model capable of predicting future tie dissolution with precision [6].

This paper analyzes and predicts Twitter breakups to address these issues. Twitter is ideal for this investigation because it is public and users can follow accounts without approval. We start with a dataset of over 120K Twitter accounts, 15 weekly snapshots of the network, posted content, and other details. To our knowledge, this Twitter dataset is the first with temporal precision and scale. This dataset is used to extract heterogeneous features from the evolving Twitter network. These include structural features like user centrality, content-related features like Tweet subjectivity, user profile attributes like self-declared bio information, egocentric traits like reciprocity, and dense user representation (node embeddings). These features help us analyze broken ties in established ties. We use multiple machine learning models to show that these features can predict future broken ties, even for users not in the training data. Finally, this paper makes three main contributions:

- We create a large-scale social network encompassing over 120K Twitter users. This network includes not only their content but also 15 weekly temporal snapshots, providing insight into users' unfollowing activities.
- We extract a multitude of features from various perspectives based on our dataset. We conduct analysis on broken ties from various perspectives based on extracted features.
- We use the extracted features and develop several machine learning models capable of predicting broken ties with high performance.

II. Related Work

A. Follow Prediction

Link prediction, a prominent area of research in social networks, seeks to forecast future relationships. Researchers have primarily employed supervised machine learning for this purpose, framing it as a classification task [7]. This approach leverages both topological and content-based characteristics, including explicit network edges or implicit ones formed due to one node's actions on another [8]. Topological features like Common Neighbors, preferential attachment, shortest path, or node degree have been particularly informative. These methods successfully predict links between tightly interconnected users [9]. Researchers have also fused explicit topological data and content-based features, outperforming conventional methods in link prediction tasks [10], [11]. Unsupervised machine learning methods and Global-based Probabilistic Approaches have also been explored, segmenting networks into communities or generating edge labels based on the likelihood of existence [12], [13]. However, these approaches often neglect edge attributes, focusing only on structural features. Other studies proposed that Twitter is a hybrid of a social network and a news source, enabling both information dissemination and the exchange of ideas among members [14]. In this line of research, models incorporating geographical information, user content networks, tweet credibility, and temporal data emerged, enabling more accurate forecasting and more nuanced relationship predictions [15].

B. Unfollow Prediction

While existing studies have delved into online social connections, most present significant limitations. Early research mainly centered on single static network snapshots, neglecting temporal information [16], [17], [18], [19]. While some studies using longitudinal data do exist, they have offered only minimal focus on social tie maintenance [20], [21], [22], [23]. Work on tie dissolution is even rarer, with some primarily computational studies centered on dyadic relationships or non-explanatory predictions/classifications [1], [24], [25], [5], and others building on existing social theories [26], [27]. However, these theory-grounded studies lack in-depth analysis and development of new theories on online social tie dissolution, are less experimental, and typically use smaller survey data.

Our work, to the best of our knowledge, stands as the first to combine 1) multiple fine-grained weekly network snapshots (15 weeks), 2) over 120K users with diverse online data, 3) a comprehensive analysis of broken ties highlighting various user and network features, and 4) accurate prediction of future broken ties using historical network snapshots.

III. Problem Statement

We operate under the assumption that there are $\mathcal{T}$ available Twitter social graphs (snapshots), each represented by $G^t = (U, E^t)$, where $1 \leq t \leq \mathcal{T}$. In this notation, U denotes the set of users (nodes), and E^t refers to the set of edges between the nodes in U. We also assume that the users remain constant, with only the edges varying across snapshots. A directed edge from user u_i to u_j ($u_i, u_j \in U$) is denoted as $(u_i, u_j) \in E^t$. Based on this notation, we define the following entities related to social ties.

- ❒ **Follower:** At timestamp t, user u_i is a *follower* of user u_j if $(u_i, u_j) \in E^t$.
- ❒ **Followee:** At timestamp t, user u_j is a *followee* of user u_i if $(u_i, u_j) \in E^t$.
- ❒ **Follow:** An edge (u_i, u_j) is labeled as *follow* if $(u_i, u_j) \notin E^t$ and $(u_i, u_j) \in E^{t+1}$.
- ❒ **Unfollower:** User u_i is an *unfollower* if $(u_i, u_j) \in E^t$ and $(u_i, u_j) \notin E^{t+1}$. Consequently, edge (u_i, u_j) signifies an **unfollow** event at time t.
- ❒ **Unfollowee:** User u_j becomes an *unfollowee* if $(u_i, u_j) \in E^t$ and $(u_i, u_j) \notin E^{t+1}$.

With the above notation and definitions, this paper focuses on two major tasks.

Broken Tie Analysis: Given $\mathcal{T}$ Twitter network snapshots and network-related and user-related features, we aim to uncover intriguing patterns and trends that enhance our understanding of the dynamics of tie formation and dissolution on social media.

Broken Tie Prediction: We aim to develop a machine learning model $\mathcal{M}$ trained on $\mathcal{T} - k$ Twitter network snapshots and network-related and user-related features. The model is designed to predict the status of edges (u_i, u_j) (follow or unfollow) for subsequent snapshots $\{\mathcal{T}-k+1, \mathcal{T}-k+2, \cdots, \mathcal{T}\}$.

IV. Dataset

TABLE I: Weekly Twitter dataset statistics (15 weeks)

Network Property	Value
Total users	123,829
Total ties	2,922,732
# Verified accounts	3,829
Avg weekly new followers	10,855
Avg weekly new unfollowers	465
Avg weekly new Tweets	1,175,846
Percentage Verified Users	1.687
Avg Followees Count per User	205
Avg Followers Count per User	150
Diameter (longest shortest path)	8
Avg new tweets (w/ mentions)	2,021

Twitter provided the large-scale, evolving social network and detailed online user data our study needed. Twitter's API academic access made data collection much easier. From 2019 to 2021, we collected weekly snapshot data for 123,829 Twitter users. This dataset includes users' social connections and tweets. An author's Twitter breadth-first search identified 130,000 users. We excluded users who changed their privacy settings or deactivated accounts during data collection to refine this set. Initial fundamental statistics were performed on 15-week data. Table reftab:twitterdataset shows these preliminary results. Importantly, Twitter and its APIs do not provide information about dissolved social ties. User retention and experience may be the reason. Thus, longitudinal data collection of the same users' social connections is the only way to obtain such data. Our study used 15-week data for computational efficiency.

V. Broken Tie Analysis

Twitter connections, or "ties," show how users interact. Online connections show influence, information spread, emotions, and user ideas. These interactions influence trends, thoughts, and even real-life events. Understanding these ties improves social media use, content creation, fighting misinformation, and online behavior. Different perspectives help us understand Twitter ties. Five categories highlight different aspects of these ties in our analysis. Our detailed broken tie analysis relies on these categories. Our Twitter ties study uses these characteristics. This helps us understand Twitter's complex social network and its implications.

A. Sociocentric Network Structure Analysis

The sociocentric network analysis provides insights into the broader social context within which individuals interact, allowing us to understand not only direct relationships but also the larger networks within which they are embedded. It moves beyond merely examining a user's immediate connections (i.e., whom they follow and who follows them) to also look at the connections between a user's connections and the overall structure of the network. Concerning broken ties, understanding the sociocentric structure can help shed light on the causes and implications of tie dissolution. For instance, tie dissolution may be influenced by indirect connections, such as shared followers or mutual friends, and the broader network structure, like the user's position within their community (e.g., their centrality). Therefore, sociocentric network analysis is vital for gaining a deeper understanding of the complex dynamics of tie formation and dissolution on social platforms like Twitter.

We examine two classes of sociocentric structural features: network centrality and the predicted community. Network centrality encompasses various measures, including in-degree and out-degree, Eigenvector, betweenness, hubs, authority, and PageRank. These measures serve to quantify the influence or prominence of a node within the network structure. As for the predicted community, we utilize Spectral Clustering to identify distinct communities within the network. This process begins with the calculation of the graph's normalized Laplacian matrix, followed by the use of its Eigenvectors to cluster the nodes. Through this methodology, we have identified five separate communities in the network data spanning 15 weeks.

Figure 1 presents the analysis of sociocentric structural features, comprising four centrality measures and predicted community. Centrality measures are categorized into three bins (i.e., high, medium, and low), while the predicted community feature corresponds to five extracted communities. We derive the following observations from the results illustrated in Figure 1.

- ➩ Unfollowers with high in-degrees and unfollowees with low in-degrees generally experience fewer unfollow events. This may be attributed to the likelihood that users with low in-degrees have fewer but stronger ties. Hence they are less likely to be unfollowed. Conversely, suppose an unfollower has a low in-degree, and the unfollowee has a high in-degree (e.g., an ordinary user follows a celebrity). In that case, the chance of unfollowing is higher due to the relatively low tie strength between the parties involved.
- ➩ Regarding out-degree binning, we observe that unfollowees with a low out-degree (fewer than 39 connections) account for the most significant increase in unfollow events. Such users may need more reciprocal connections and maintain high standards for whom they follow.
- ➩ Concerning eigenvector centrality, we notice that most unfollow events occur along the diagonal, reflecting the nature of eigenvector centrality where a user's centrality is passed onto her connections. This suggests that users of similar importance (according to eigenvector centrality) are more likely to be connected.
- ➩ For betweenness centrality, unfollowees with medium betweenness values tend to be unfollowed less frequently. However, if an unfollowee's betweenness value is low, the higher betweenness centrality of an unfollower increases the likelihood of unfollowing events. This trend is reversed when the betweenness of the unfollowee is high. Users with high betweenness centrality are more likely to bridge different communities (being on many shortest paths), making them more likely to be unfollowed by users from different communities.
- ➩ Considering the predicted community plots, we observe three larger and two smaller communities. Most unfollow events ($> 90\%$) occur within a community, likely due to the majority of user connections being within their respective communities. Thus, most opportunities to unfollow also exist within a community.

These observations suggest that the unfollower and unfollowee's structural centrality scores are associated with unfollowing behavior. Moreover, the community structure tends to have a localized effect on unfollowing.

B. Egocentric Network Structure Analysis

Egocentric structure network analysis looks at a user's immediate connections, or "friends," and the relationships between those friends. Features of the egocentric network, such as the number of mutual followers, the ratio of followers to followees, and the density of connections among a user's followers, can provide important clues about the user's position within the network and their potential for tie dissolution. Overall, egocentric structure network analysis enables a more nuanced understanding of a user's local network, allowing us to identify patterns that might predict tie dissolution more accurately. The features under this category include:

- ❒ # Lost followees: This feature tracks the count of accounts a user has ceased following within a certain time frame.
- ❒ # New followees: This feature quantifies the number of accounts a user has begun following within a given period.

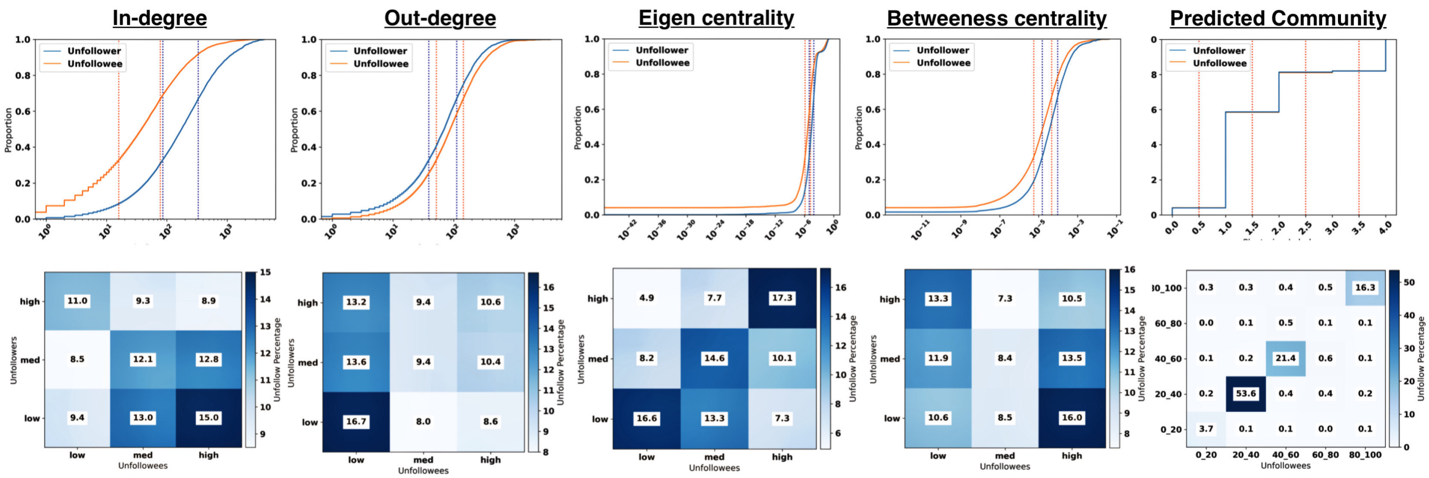

Fig. 1: CDF plots and binning for a subset of structural features

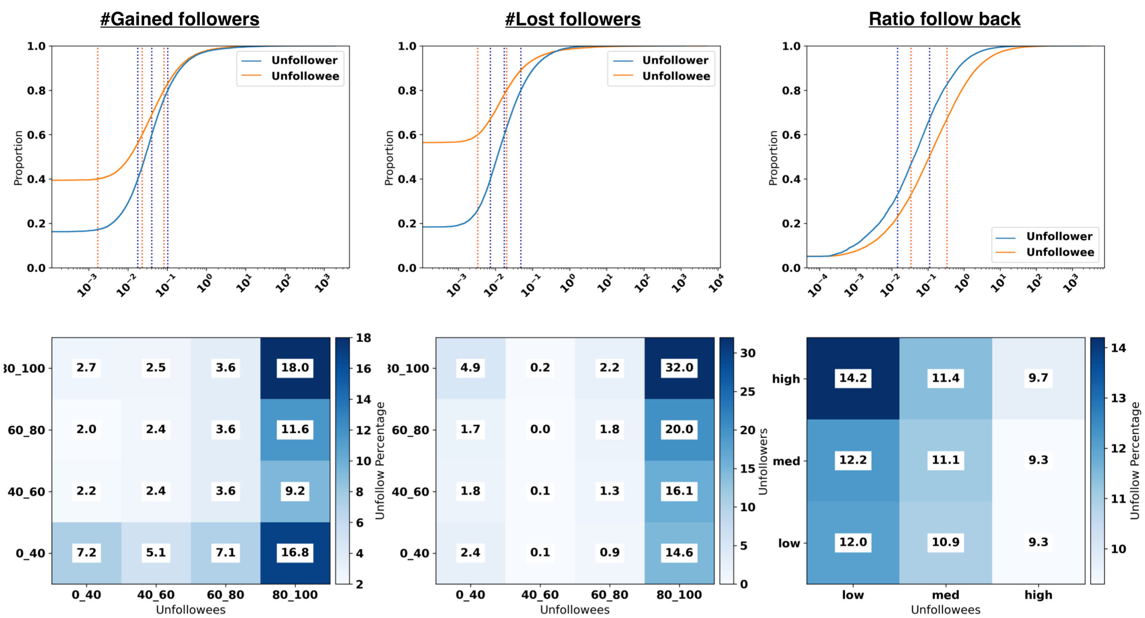

Fig. 2: CDF plots and binning for a subset of egocentric network features

- ❐ # Lost followers: This feature registers the number of users who have unfollowed a specific account over a certain time frame.
- ❐ # New followers: This feature logs the count of users who have started following a particular account within a set time period.
- ❐ Follow-back ratio: This feature calculates the ratio of follow-backs (following a user who follows you) to the total followees.
- ❐ # Followers / # Followees: Similar to the follow-back ratio, this feature measures the ratio of followers to followees, a metric found to correlate with social ties on Twitter [28].

Figure 2 presents the analysis results for several egocentric network features. Based on these results, we make the following observations:

- ➩ In the 'Gained Followers' plot, the x-axis represents the quantile range of the unfollowee's gained percentage of followers, and the y-axis represents the quantile range for the unfollower's gained percentage of followees (from one week to the next). Over 50% of unfollow events occur when the unfollowee's gained followers fall within the top 20% quantile (i.e., the last column corresponding to the 80-100 range). We surmise that this occurs because unfollowee users who are highly active in expanding their ego network (perhaps indiscriminately) are more likely to be subject to unfollowing events from other users. The most significant contribution comes from unfollowers who have also gained a substantial number of followers (i.e., the bin of 80-100 and 80-100).

- The next plot, 'Lost Followers', shows the quantile range of the unfollowee's lost % of followees on the x-axis, while the y-axis shows the unfollower's lost % of followers. Notably, over 80% of the unfollow incidents occur to unfollowees in the highest bin of the lost % of followees. This suggests that if a user loses a significant number of connections, they are likely to continue losing additional connections. This trend is similar to preferential attachment in evolving networks [29], but in this case, it applies to lost ties rather than gained ones.
- The 'Ratio Follow Back' plots in Figure 2 clearly show that reciprocal ties are stronger. Specifically, users with a higher proportion of reciprocal ties are less likely to be unfollowed. This finding aligns with the notion that reciprocal links in the network strengthen the user's position and reduce the likelihood of them being unfollowed.

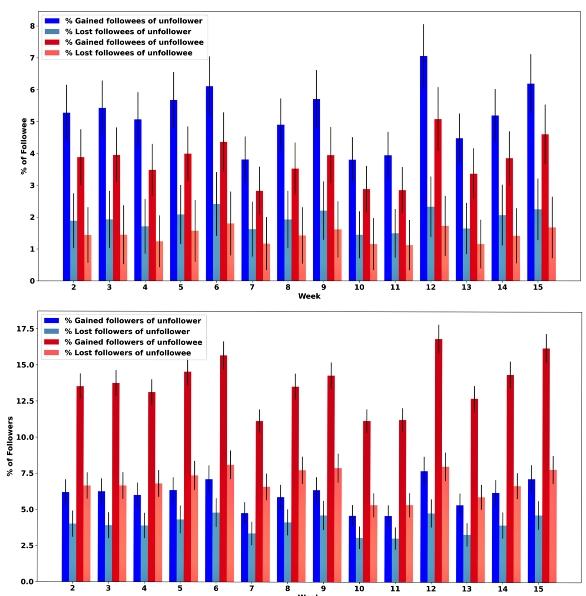

Fig. 3: The changes in % of gained/lost followers (top plot) and % of gained/lost followees (lower plot) of unfollower and unfollowee.

To delve deeper into the egocentric network, Figure 3 showcases the variations in the percentage of gained or lost followers and followees of unfollowers and unfollowees. These alterations are captured from weekly snapshot i to $i+1$. As demonstrated by Figure 3, the social graph is growing denser as the percentage of followers and followees gained persistently outstrips the percentage lost. Notably, the values for followers and followees gained or lost, particularly the ratios between these values, maintain remarkable consistency across various time snapshots. Consequently, even though severed connections momentarily reduce the size of the egocentric networks of unfollowers and unfollowees, users typically form new connections. As a result, both the individual egocentric networks and the overall social graph exhibit continuous expansion over time.

C. User Content Analysis

Understanding the nuances of online connections often requires analyzing user-generated content such as tweets, retweets, and hashtags [24]. These pieces of content can reveal user sentiment, topic of interest, level of engagement, and more. Therefore, the second category of features we consider for broken tie analysis focuses on this user-generated content. Below, we elaborate on these features:

- Hashtag similarity: This feature quantifies the degree of similarity between the hashtags utilized by two users. For this purpose, we employ Fuzzywuzzy [30], a tool that determines similarity ratios between text tokens based on the Jaccard similarity algorithm.
- Tweet polarity: This feature assesses the sentiment expressed in a tweet, with values ranging from -1 (indicating negativity) to 1 (indicating positivity).
- Tweet subjectivity: This metric gauges the subjectivity inherent in a tweet, with values extending from 0 (signifying objectivity) to 1 (indicating subjectivity). To derive polarity and subjectivity, we utilized the Textblob sentiment analysis library [31].
- Tweet word count: This feature represents the count of words contained in a tweet.
- # Mentions: This feature denotes the frequency of mentions of a user's handle in tweets.
- # Mentions from (un)follower to (un)followee: This feature measures the number of times an (un)follower has mentioned the (un)followee.
- # Mentions from (un)followee to (un)follower: This feature accounts for the frequency with which an (un)followee has mentioned the (un)follower.
- # Hashtags: This quantifies the total count of hashtags used by a user.
- # Tweets containing URL: This feature notes the number of tweets that incorporate a URL.
- # Tweets containing symbols: This feature tallies the number of tweets containing symbols.

Figure 4 displays the results pertaining to content polarity and subjectivity. The upper segment of this figure presents the cumulative distribution function (CDF) for content polarity and subjectivity. Based on these distributions, we establish a cut-off point to categorize polarity and subjectivity scores into two bins: high and low. The lower segment of Figure 4 depicts the percentage of broken ties associated with the four combinations of polarity and subjectivity, considering both parties involved in a broken tie: the unfollower and unfollowee. Note that we implemented a similar binning strategy for most feature analysis experiments. The observations from Figure 4 are as follows:

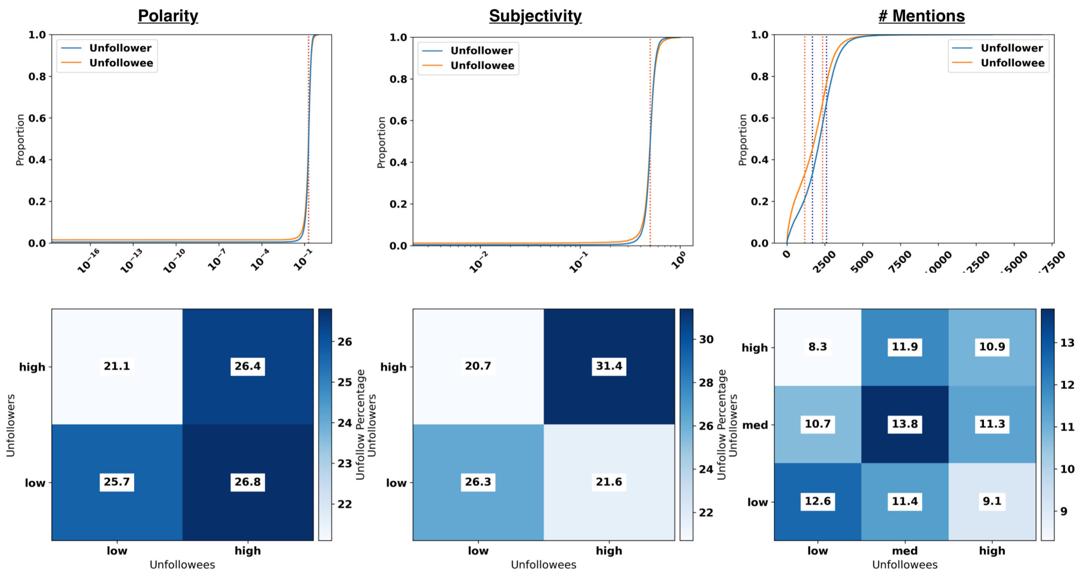

Fig. 4: CDF plots and binning for polarity, subjectivity, and mentions within the content

- ➩ Irrespective of the unfollowers' average tweet polarity, they exhibit a higher likelihood of unfollowing users possessing a high tweet polarity score.
- ➩ A surge in the count of unfollows is observed when the unfollower's average tweet polarity is low.
- ➩ The majority of unfollows occur between users whose average tweet subjectivity mirrors their own.

These observations indicate that the subjectivity and polarity of user-generated content significantly influence the dissolution of online ties on Twitter.

D. User Profile Analysis

A user's profile is a snapshot of their chosen online identity, showcasing information they have elected to share publicly. It typically includes elements like their username, bio, location, profile picture, cover photo, and, potentially, a website link. These elements can provide insightful details about users' interests, personality traits, and affiliations, which could influence their social ties. Verified accounts, for example, often belong to public figures or notable individuals in various fields. As a result, followers might be less likely to dissolve ties with such accounts, given their public standing and the unique, authoritative content they provide. The profile-based features are described as follows:

- ❒ Verified account: This feature indicates whether a user's account has received Twitter's verification.
- ❒ # Subscribed lists: This feature accounts for the number of lists a user is subscribed to.
- ❒ Favorite count: This feature denotes the count of tweets a user has marked as favorites.
- ❒ Presence of Profile Image: This feature signifies whether a user has uploaded a profile image.
- ❒ Bio similarity: This feature measures the degree of similarity between the bios of two users. For this, we employ Fuzzywuzzy [30], which calculates similarity ratios between text tokens based on the Jaccard similarity algorithm.
- ❒ # Tweets: This feature represents the count of tweets a user has posted.
- ❒ # Followers: This feature records the total number of a user's followers.
- ❒ # Followees: This feature indicates the count of users followed by a user. We have included the last three features in this category as they are easily visible on a Twitter user profile and can serve as distinguishing identifiers, particularly in the case of the follower count.

Figure 5 showcases the results of analyzing user profile features, specifically focusing on three distinct attributes. First, we only demonstrate the binning for verified and unverified categories for verification status. In terms of the number of tweets, Figure 5 displays both the CDF plot, three binning categories (low, medium, and high for exact cut-off values), and the weekly changes in the average number of tweets, normalized by the number of users in each unfollower and unfollowee category. A similar methodology is employed for calculating weekly favorites counts. Based on the analytical results reflected in Figure 5, we draw the following observations:

- ➩ As demonstrated in Table I, a significant majority of Twitter users possess unverified accounts, which explains why 98.1% of the unfollow incidents fall into the unverified-unverified bin. Nonetheless, an intriguing pattern emerges when considering the percentages in off-diagonal elements. Verified users unfollowing unverified accounts constitute five times more unfollow incidents than unverified accounts unfollowing verified accounts.
- ➩ In relation to favorites counts, both unfollowers and

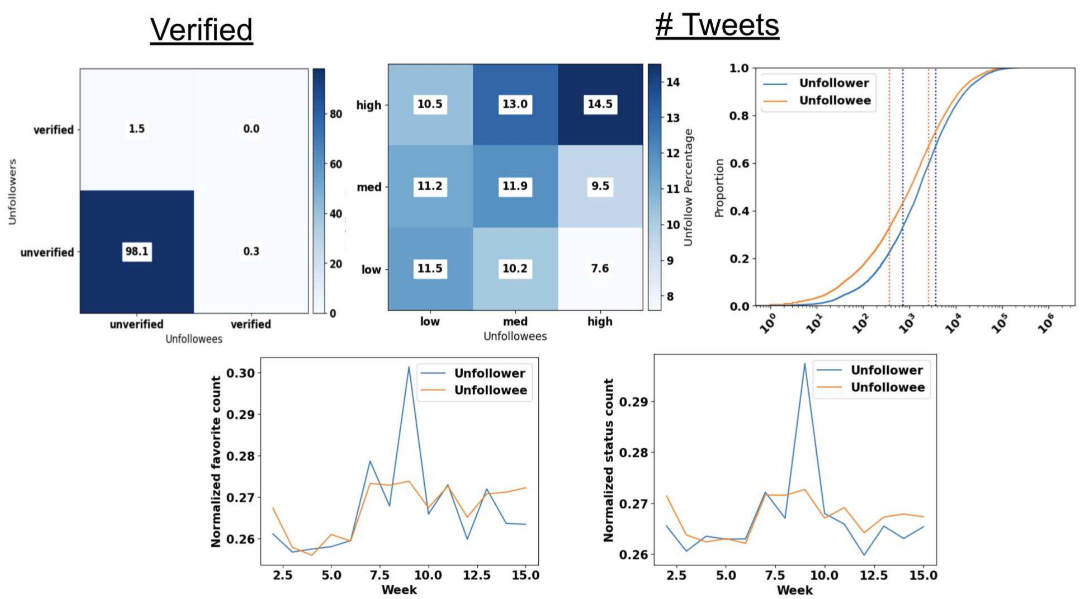

Fig. 5: CDF plots and binning for several user profile features

unfollowees exhibit similar behaviors for the most part, with the exception of week 9.

- ⇨ There is a negative correlation between the activity levels of unfollowers (measured by the number of tweets posted) and the likelihood of unfollowing. This could potentially be attributed to the increased chances of interaction and engagement between the unfollower and the unfollowee, along with their content. Conversely, a high tweet posting rate for the unfollowee correlates with an increased likelihood of being unfollowed. This is consistent with previous research in [32] where authors demonstrated that high tweet bursts by users leads to their unfollowing.

Drawing from these observations, we can conclude that profile features such as account verification and tweeting rate exhibit associations with broken ties on Twitter.

E. Dense User Representation Analysis

Building upon the progress in neural networks, particularly graph representation learning, we can obtain dense user representations, or 'node embeddings'. These embeddings encapsulate critical information about users, their social interactions, and their communities. Our hypothesis is that these node embeddings evolve over time, reflecting the dynamic nature of the graph structure due to the creation and dissolution of social ties. To extract these dense user representations, we employed the Node2vec method [33]. This technique allows for a flexible neighborhood sampling strategy, bridging the gap between Breadth-First Search (BFS) and Depth-First Search (DFS) by introducing a bias-random walk procedure.

To examine the dense user representations, we compute the change in Euclidean distance between the embeddings of (un)followers and (un)followees from week 1 to week 15. The results are presented in Figure 6. We study both follow and unfollow incidents and consider the number of directional ties between users at the time of the follow/unfollow event. For a follow event, this can be either 1 (a user follows another) or 2 (mutual following). For an unfollow event, there are three potential scenarios: 0 (no remaining ties), 1 (one user continues to follow the other), or 2 (a user unfollows another and then subsequently follows back again). In addition, we segregate users into two categories based on their in-degree, distinguishing between those with high in-degree (H) and low in-degree (L). This separation is crucial as users with high (low) in-degrees tend to be more well-known (less-known) and have larger (smaller) communities. This factor impacts the user representations generated by Node2vec, which aims to create representations for users within similar neighborhoods. The bottom part of Figure 6 displays the results across weeks without this in-degree segregation. Based on the results showcased in Figure 6, we make the following observations.

- ❐ Notably, users who end up without ties following an unfollow event generally experience an increased distance between their embeddings. This implies that the dissolution of a tie typically results in the two associated nodes moving farther apart within the user representation space. An exception occurs when both the unfollower and unfollowee have high in-degrees. This outcome is expected given that individuals with numerous connections, who were previously connected, likely share numerous common ties. Since user representations depend on common neighbors, two individuals with many shared connections would likely remain close in the user representation space even if their direct connection is dissolved.
- ❐ Since user representation is derived from an undirected graph rather than a directed graph, if at least one directional tie exists between two users, the user representation

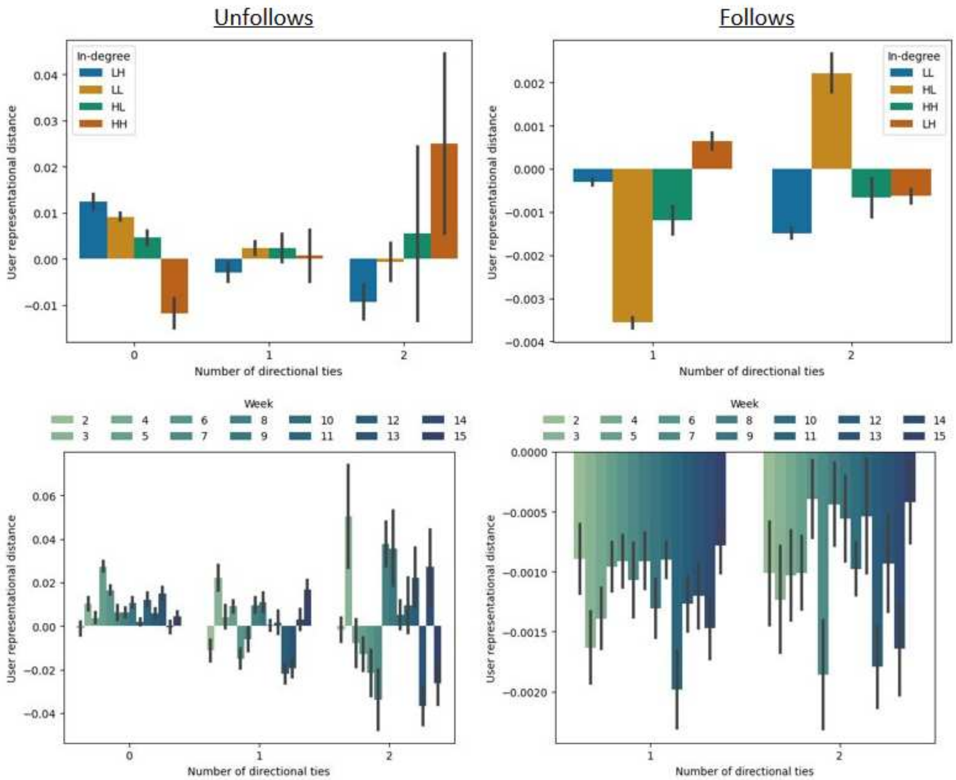

Fig. 6: Dense user representation plots

algorithm treats it as a reciprocal tie. Therefore, as shown on the left-hand side images, for pairs of users maintaining a single tie while experiencing an unfollow event, the average change in distance is nearly negligible. There is a bin for users maintaining two ties despite experiencing an unfollow event. However, such cases, where a user unfollows another only to refollow later, are rare and expectedly associated with higher standard error bars.

VI. Broken Tie Prediction

In the preceding section, we focused on the analysis of broken ties. Here, we extend the analysis by developing models capable of predicting broken ties. This approach serves two principal objectives. Firstly, it allows us to validate the features extracted as beneficial and informative for predicting broken ties. Secondly, it aids in understanding the future evolution of the social network. Such knowledge can be employed for several downstream applications, including friend recommendation, client retention, popularity maintenance, and business growth. The models developed take user features as input and predict whether a follow (maintained tie), or an unfollow (broken tie) relationship exists between them. Therefore, this task reduces to binary classification. The experimental settings and obtained results are elaborated upon in the subsequent subsections.

There are many published works on link prediction, and it would be worthwhile to discuss why we cannot use such techniques in broken tie prediction. Link prediction strategies are effective for follow prediction and cannot be directly applied to broken tie prediction due to different governing dynamics and factors. Follow prediction models usually consider features like shared interests, mutual acquaintances, and demographic similarities [34], [35]. In contrast, broken tie prediction must account for factors like posting frequency, content relevance or quality, and shifting interests over time [3]. The temporal dynamics also differ; follow events can happen anytime, while unfollow events are typically reactionary, instigated by specific incidents or posts [24]. Additionally, broken tie prediction encounters issues with data availability and bias. Unlike follow prediction, which benefits from plentiful historical data, unfollow events are less frequent, harder to observe directly and cause a class imbalance problem [36]. Lastly, the nature of the problem contributes to the complexity of broken tie prediction. Broken tie prediction is inherently more challenging because it is an imbalanced classification problem—the number of retained follows substantially outweighs the broken ties.

A. Experimental Settings

This section details the experimental settings, including train and test split, evaluation metrics, implementation details, and machine learning prediction models.

Training and Test Split. The proportion of unfollows in our dataset is significantly smaller than follows (constituting 1.08% of all ties). To reduce computational cost and facilitate model training, we constructed a balanced subset of data employing the NearMiss algorithm [37], with a ratio of 1 unfollow to 10 follows. To render our prediction practical and useful, we used historical weekly snapshots for training and future data for testing. Specifically, we trained models on data from weeks 1 to 11 and used weeks 12 to 15 for testing, denoted as the *Entire Future Snapshots* test set. We also created a *Disjoint Users Test* set, excluding any dyadic ties where users were present in the training set, ensuring the train and test user sets were utterly disjoint.

TABLE II: Performance of broken tie prediction

Model	Test set		Acc	AUC	Recall	Precision	F-1	Kappa
Light Gradient Boosting Machine	Entire Future Snapshots	Unfollow	0.955	0.983	0.909	0.911	0.910	0.880
		Follow	0.962	**0.985**	0.921	**0.929**	0.925	**0.900**
		Total	0.959	0.950	0.7023	0.860	0.818	0.808
	Disjoint Users Test	Unfollow	0.845	0.753	0.796	0.858	0.818	0.800
		Follow	**0.976**	0.965	0.985	0.854	0.981	0.794
		Total	0.959	0.950	0.702	0.860	0.818	0.800
Extreme Gradient Boosting Machine	Entire Future Snapshots	Unfollow	0.872	0.855	0.779	0.824	0.732	0.709
		Follow	0.965	0.899	0.834	0.827	0.745	0.724
		Total	0.952	0.938	0.747	0.868	0.803	0.785
	Disjoint Users Test	Unfollow	0.829	0.899	0.755	0.831	0.791	0.745
		Follow	0.975	0.970	**0.984**	0.865	**0.980**	0.838
		Total	0.967	0.956	0.746	0.877	0.806	0.789
Random Forest	Entire Future Snapshots	Unfollow	0.919	0.754	0.651	0.660	0.658	0.611
		Follow	0.943	0.898	0.739	0.676	0.753	0.604
		Total	0.946	0.928	0.718	0.670	0.635	0.608
	Disjoint Users Test	Unfollow	0.822	0.867	0.721	0.653	0.768	0.608
		Follow	0.972	0.984	0.984	0.856	0.978	0.880
		Total	0.957	0.926	0.768	0.766	0.767	0.744

Evaluation Metrics. We employed seven key indicators to evaluate the efficacy of our unfollow prediction, including accuracy, precision, F-1 score, recall, AUC (area under the curve), and the Kappa metric. The latter measures the model's predictive accuracy by comparing observed and expected accuracy.

Implementation Details. After extracting all features as explained in Section V, we normalized them and removed perfect collinearity. We employed the Pycaret library [38] to investigate eleven classification models, from which we selected the best three tuned models. We used 10-fold cross-validation on the training set for hyperparameter tuning.

Machine Learning Prediction Models. We employed three popular and traditional machine learning classification methods for the broken tie prediction: Random Forest, Extreme Gradient Boosting Machine, and Light Gradient Boosting Machine.

B. Experimental Results

Table II presents the performance of different approaches on two distinct train and test splits for predicting follow and unfollow instances. We observe the high performance of the models across the board, demonstrating the efficacy of our feature extraction for the task of broken tie prediction. Further, high performance on the *Disjoint Users Test* set indicates the generalizability of the models to unseen users in dyadic ties. Notably, the Light Gradient Boosting Machine exhibits exceptional performance, rendering it suitable for real-world applications interested in mapping the evolution of the Twitter network.

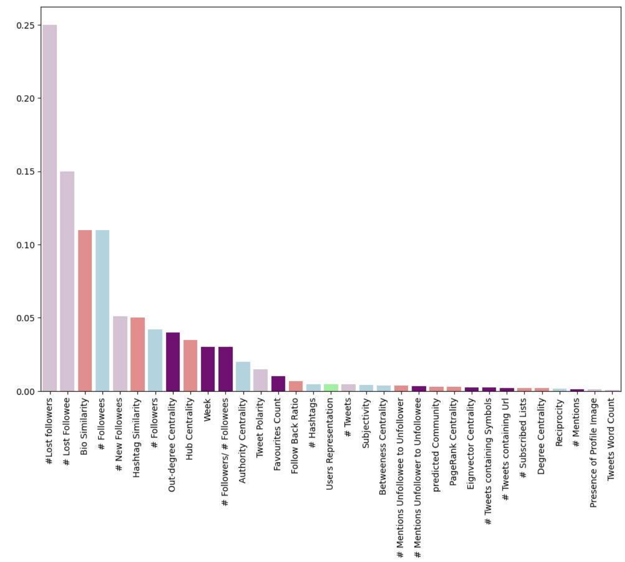

Fig. 7: Importance of each feature based on the Light Gradient Boosting Machine. Feature importance ratios are 0.1729 for Sociocentric Network Structure, 0.4054 for Egocentric Network Structure, 0.1349 for User Content, 0.2657 for User Profile, and 0.0172 for Dense User Representation features.

C. Feature Importance

Based on our best model, the Light Gradient Boosting Machine, we determined the importance of each feature and category, as demonstrated in Figure 7. Except for the dense user representation, all feature categories play a significant role in broken tie prediction. Notably, the two egocentric network features (# Lost followers and # Lost followees) lead in broken tie prediction, underscoring the potential of egocentric network dynamics to forecast future tie statuses.

VII. Conclusion and Future Works

In this paper, we introduced a large-scale and longitudinal dataset of Twitter social networks spanning over 15 weeks. Leveraging this data, we thoroughly analyzed broken ties on Twitter from multiple perspectives. This analysis revealed interesting patterns in the complex nature of evolving social ties on social media. Using our extracted rich and comprehensive feature sets, we employed several popular machine learning algorithms and predicted broken ties. Our predictions exhibited high performance according to several metrics (e.g., F1-score and AUC). Utilizing the dataset and prediction models, we can more accurately understand and predict the evolution of relationships between social media users, which has potential applications in understanding user behavior, political campaigns, and influencer marketing.

Looking towards the future, we aim to leverage the capabilities of deep learning models to predict impending unfollow events and to delve deeper into the correlation between the derived features and social unfollowing patterns. Furthermore, we also intend to explore the application of graph neural networks in this context, as these models can capture the complex interactions within social networks more effectively. In addition, we plan to investigate the causal relationships that drive the unfollowing phenomenon, which will involve creating sophisticated models that can handle temporal data and complex social interactions. These lines of investigation will not only provide a deeper understanding of unfollow behavior but could also improve Twitter engagement and user retention strategies for both individuals and organizations.

REFERENCES

[1] B. Xu, Y. Huang, H. Kwak, and N. Contractor, "Structures of broken ties: exploring unfollow behavior on twitter," in *Proceedings of the 2013 conference on Computer supported cooperative work*, 2013, pp. 871–876.

[2] K. Sasahara, W. Chen, H. Peng, G. L. Ciampaglia, A. Flammini, and F. Menczer, "Social influence and unfollowing accelerate the emergence of echo chambers," *Journal of Computational Social Science*, vol. 4, no. 1, pp. 381–402, 2021.

[3] F. Kivran-Swaine, P. Govindan, and M. Naaman, "The impact of network structure on breaking ties in online social networks: unfollowing on twitter," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2011, pp. 1101–1104.

[4] M. Choi, C. Budak, D. M. Romero, and D. Jurgens, "More than meets the tie: Examining the role of interpersonal relationships in social networks," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 15, 2021, pp. 105–116.

[5] S. K. Maity, R. Gajula, and A. Mukherjee, "Why did they # unfollow me? early detection of follower loss on twitter," in *Proceedings of the 2018 ACM Conference on Supporting Groupwork*, 2018, pp. 127–131.

[6] H. Wu, Z. Hu, J. Jia, Y. Bu, X. He, and T.-S. Chua, "Mining unfollow behavior in large-scale online social networks via spatial-temporal interaction," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, 2020, pp. 254–261.

[7] M. Hasan, V. Chaoji, S. Salem, and M. Zaki, "Link prediction using supervised learning," 01 2006.

[8] W.-H. Jang, M. Kwak, and D.-S. Han, "A network link prediction model based on object-object match method," 2014.

[9] M. Fire, L. Tenenboim, O. Lesser, R. Puzis, L. Rokach, and Y. Elovici, "Link prediction in social networks using computationally efficient topological features," in *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*. IEEE, 2011, pp. 73–80.

[10] C. Lei and J. Ruan, "A novel link prediction algorithm for reconstructing protein–protein interaction networks by topological similarity," *Bioinformatics*, vol. 29, no. 3, pp. 355–364, 2013.

[11] Z. Huang, "Link prediction based on graph topology: The predictive value of generalized clustering coefficient," *Available at SSRN 1634014*, 2010.

[12] C. P. Muniz, R. Goldschmidt, and R. Choren, "Combining contextual, temporal and topological information for unsupervised link prediction in social networks," *Knowledge-Based Systems*, vol. 156, pp. 129–137, 2018.

[13] T.-T. Kuo, R. Yan, Y.-Y. Huang, P.-H. Kung, and S.-D. Lin, "Unsupervised link prediction using aggregative statistics on heterogeneous social networks," in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013, pp. 775–783.

[14] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*, 2010, pp. 591–600.

[15] L. Backstrom and J. Leskovec, "Supervised random walks: predicting and recommending links in social networks," in *Proceedings of the fourth ACM international conference on Web search and data mining*, 2011, pp. 635–644.

[16] P. A. Grabowicz, J. J. Ramasco, E. Moro, J. M. Pujol, and V. M. Eguiluz, "Social features of online networks: The strength of intermediary ties in online social media," *PloS one*, vol. 7, no. 1, p. e29358, 2012.

[17] P. De Meo, E. Ferrara, G. Fiumara, and A. Provetti, "On facebook, most ties are weak," *Commun. ACM*, vol. 57, no. 11, p. 78–84, oct 2014. [Online]. Available: https://doi.org/10.1145/2629438

[18] L. Weng, F. Menczer, and Y.-Y. Ahn, "Virality prediction and community structure in social networks," *Scientific reports*, vol. 3, no. 1, pp. 1–6, 2013.

[19] J. J. Jones, J. E. Settle, R. M. Bond, C. J. Fariss, C. Marlow, and J. H. Fowler, "Inferring tie strength from online directed behavior," *PLOS ONE*, vol. 8, no. 1, pp. 1–6, 01 2013. [Online]. Available: https://doi.org/10.1371/journal.pone.0052168

[20] R. Toivonen, L. Kovanen, M. Kivelä, J.-P. Onnela, J. Saramäki, and K. Kaski, "A comparative study of social network models: Network evolution models and nodal attribute models," *Social Networks*, vol. 31, no. 4, pp. 240–254, 2009. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378873309000331

[21] K. Lewis, J. Kaufman, M. Gonzalez, A. Wimmer, and N. Christakis, "Tastes, ties, and time: A new social network dataset using facebook. com," *Social networks*, vol. 30, no. 4, pp. 330–342, 2008.

[22] G. Kossinets and D. J. Watts, "Origins of homophily in an evolving social network," *American journal of sociology*, vol. 115, no. 2, pp. 405–450, 2009.

[23] T. Y. Berger-Wolf and J. Saia, "A framework for analysis of dynamic social networks," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006, pp. 523–528.

[24] H. Kwak, H. Chun, and S. Moon, "Fragile online relationship: a first look at unfollow dynamics in twitter," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011, pp. 1091–1100.

[25] H. Kwak, S. Moon, and W. Lee, "More of a receiver than a giver: why do people unfollow in twitter?" in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 6, no. 1, 2012, pp. 499–502.

[26] N. Pennington, "An examination of relational maintenance and dissolution through social networking sites," *Computers in Human Behavior*, vol. 105, p. 106196, 2020.

[27] N. Pennington and J. A. Hall, "Does facebook-enabled communication influence weak-tie relationships over time? a longitudinal investigation into mediated relationship maintenance," *Communication Monographs*, vol. 88, no. 1, pp. 48–70, 2021.

[28] H. Oshimo, S. Hironaka, M. Yoshida, and K. Umemura, "Follower–followee ratio category and user vector for analyzing following behavior," in *2022 9th International Conference on Advanced Informatics: Concepts, Theory and Applications (ICAICTA)*. IEEE, 2022, pp. 1–6.

[29] H. Jeong, Z. Néda, and A.-L. Barabási, "Measuring preferential attachment in evolving networks," *EPL (Europhysics Letters)*, vol. 61, no. 4, p. 567, 2003.

[30] J. Wang, G. Li, and J. Fe, "Fast-join: An efficient method for fuzzy token matching based string similarity join," in *2011 IEEE 27th International Conference on Data Engineering*, 2011, pp. 458–469.

[31] s. Loria, "Textblob," https://github.com/charlespwd/project-titlehttps://github.com/sloria/TextBlob, 2013.

[32] H. Kwak, H. Chun, and S. B. Moon, "Fragile online relationship: a first look at unfollow dynamics in twitter," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2011.

[33] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.

[34] D. Liben-Nowell and J. Kleinberg, "The link prediction problem for social networks," in *Proceedings of the twelfth international conference on Information and knowledge management*, 2003, pp. 556–559.

[35] L. Tabourier, A. Libert, and R. Lambiotte, "Rankmerging: Learning-to-rank in large-scale social networks (extended version)," *CoRR*, vol. abs/1407.2515, 2014. [Online]. Available: http://arxiv.org/abs/1407.2515

[36] C. J. Hutto, S. Yardi, and E. Gilbert, "A longitudinal study of follow predictors on twitter," in *Proceedings of the sigchi conference on human factors in computing systems*, 2013, pp. 821–830.

[37] S.-J. Yen and Y.-S. Lee, "Under-sampling approaches for improving prediction of the minority class in an imbalanced dataset," in *Intelligent Control and Automation*. Springer, 2006, pp. 731–740.

[38] M. Ali, *PyCaret: An open source, low-code machine learning library in Python*, April 2020, pyCaret version 1.0.0. [Online]. Available: https://www.pycaret.org