# Examining the evolution of mobile social payments in Venmo

Clive J. Unger University of Texas at Austin Electrical and Computer Engineering Austin, TX 78712 U.S.A. clive.unger@utexas.edu

Ishank Arora University of Texas at Austin Department of Computer Science Austin, TX 78712 U.S.A. ishankarora1100@utexas.edu Dhiraj Murthy University of Texas at Austin Moody College of Communication Austin, TX 78712 U.S.A. dhiraj.murthy@austin.utexas.edu

Andy Yuan Chang University of Texas at Austin Electrical and Computer Engineering Austin, TX 78712 U.S.A. yuan.chang@utexas.edu Amelia Acker University of Texas at Austin School of Information Austin, TX 78701 U.S.A. aacker@ischool.utexas.edu

#### ABSTRACT

This study aims to understand the changes in behavior over time of users on Venmo, an American social payments platform. As there is a dearth of replication studies in social media studies, we chose to replicate an existing study of Venmo using new data we collected. This enabled us to track the growth of Venmo from the beginning of the platform, and to verify the robustness of the existing, very limited literature using Venmo data. To accomplish this, we studied how the structure of the transaction graph of Venmo transactions changed since 2016, the data endpoint of previous research. Additionally, we collected a much larger set of data and examined if new community structures or network-level features emerged since 2016. Although we found that Venmo's growth has maintained a similar pattern within its transaction graph and community structure, we discovered some changes, such as the existence of more communities of a smaller size and an increase towards users quitting the platform after one transaction versus becoming regular users. Ultimately, our study confirms both the methods and results of previous work on Venmo, opening up the possibility for future work to study if changes unfold in the next era of the platform. We argue that this replication study serves as an intervention in the field of social media research more broadly, and encourage others to take up replication studies to enhance the reliability and generalized knowledge of the field.

#### **CCS CONCEPTS**

• Human-centered computing • Collaborative and social computing • Collaborative and social computing design and evaluation methods • Social network analysis

#### **KEYWORDS**

Social network analysis, social payment platforms, social graphs, community detection, mobile apps, Venmo

# **1** Introduction

Mobile social payment systems continue to evolve and be used for payment transactions in new ways as mobile phones increasingly become primary communication technologies. More and more, we see mobile payments being used in existing social media platforms such as Facebook, Snapchat, or WeChat. Venmo in particular incorporates a layer of social features [2]. A number of studies examining payments in existing social media platforms exist [5, 8, 9, 16], but there is a deficiency of work that tracks changes in communities and behaviors within payment platforms over time.

Primarily this is due to changes in data access from Application Programming Interface (API) and increased regulation and protection of personally identifiable information collected by platforms. Often, researchers themselves move on to the next 'update' or killer app in terms of interests, with very few revisiting and replicating previous, early findings of social media platforms. Data sharing amongst scholars is also inhibited due to barriers such as incentives, policies, and long-term infrastructure support [12, 13]. Ultimately, replication studies do not always have the allure that new social media research has, though they provide immense scholarly benefit to those trying to understand longitudinal trends within communities who use mobile networks and apps to generate payments. Moreover, they confirm the science and knowledge upon which social media scholars are building sound theories, which are then used to inform their own empirical work.

In this study, we replicate some of the work of Zhang et al.'s [18] early and important investigation of the United States-based payment platform Venmo, and extend their vanguard methods using a larger data set and more contemporary data. This intervention is critically valuable as Venmo's transaction volume has almost doubled in size since May 2016 (when the original study's data ended). As the platform has become more mainstream, we particularly sought to evaluate whether analysis of older data and previous findings are consistent with current user behavior. Our study recreates and extends Zhang et al. [18], by using some of their methods for finding dataset metrics on a significantly larger dataset. Additionally, Zhang et al. [18] found that 'clustering all 5 million users is computationally challenging', leaving a gap for future work to tackle. However, it should be noted that we did not have access to the friend graph Zhang et al. [18] used as API features have changed [6] and we are reproducing the study five years later in 2020.

The Venmo platform is a subsidiary of PayPal that allows users to sign up and send each other money using a connected bank account, debit card or credit card [14]. Funds are initially deposited into a user's "Venmo Balance", where they can use that money for future Venmo transactions or deposit funds into a bank account [14]. What makes Venmo unique as a payment medium is its resemblance to a social network [1]. A user can add other users as friends within the app, leveraging their phone's address book, or bootstrap their friends list with Facebook friends if they connect their user accounts. Every Venmo payment transaction includes an open text message "memo" field, where users can specify what the purpose of the transactions is with words, emoji, and "fave" hearts. Often users will simply enter a single word such as "pizza" or even use a representative emoji, such as the pizza emoji, in this memo field. Users can opt for their payment transactions to be posted privately, posted amongst a designated list of "friends", or posted to the public Venmo feed accessible to all platform users. These memos of transaction messages can be collected through the public API if they are published to the public feed by users. Importantly, the user data made available through the Venmo API contains no financial information, including the amount of the payment, bank account, or financial data related to users.

The purpose of this study is to explore whether there are changes in the composition and behavior of communities discerned from Venmo financial transactions over the first six years of the Venmo platform in order to evaluate not only changes in data access and API features, but also how such changes impact the ways in which we conduct social media research. Using Zhang et al.'s "Cold Hard E-Cash: Friends and Vendors in the Venmo Digital Payments System" [18] as a starting point, we evaluate what has changed since that study's original analysis of social and transaction graphs in Venmo by replicating their initial analysis from 2016. Zhang et al. [18] sought to understand how Venmo's integration of a social network affects users' social and financial behavior, by first creating a financial transaction graph with users as nodes and with weighted edges to denote the number of transactions between two users. In their study, a social graph was also created from users and their list of friends [18]. Clustering methods such as K-core decomposition, which is a measure of sparsity, were used to compare Venmo communities to Twitter and Facebook communities, to show that Venmo users form dense communities with a higher than expected clustering coefficients [18].

Further, Zhang et al. [18] investigated the behavior of communities within the transaction graph, using the popular Louvain community detection algorithm [4]. Two main community types were identified: "friend-driven" and "business-driven". Then properties of communities were examined to show the existence of "niche groups" around a single type of transaction such as gambling on sporting events, business owners, or simply people paying rent. Finally, they examined what people use Venmo for by classifying payment types based on the transaction message. In addition,

their study provides a basic look at temporal dynamics of transactions amongst groups, by analyzing the periodic nature of certain payments.

For our study, we collected transactions from the Venmo public API from March 25th, 2012 to August 16, 2018. In this data spanning six years, there are 23,133,264 unique userIDs and over 341,309,788 total transactions. This represents a significant increase in data as Zhang et al.'s [18] data contained 7,091,915 unique userIDs and 91,355,414 transactions (also over a six-year period, though from April 15, 2010 to May 5, 2016). This API-collected data includes identification of the two parties involved, the transaction message, as well as various other metadata. The tremendous growth of Venmo in recent years means that our dataset is more than triple the size of what Zhang et al. [18] used to examine community behaviors using Venmo transactions. Studies examining the change in behavior on a social platform as it grows and emerges remain rare, so evaluating changes in the data landscape and how these access regimes impact social media research are important contributions to the social media literature. We wanted to understand how social graph metrics change in response to a larger, newer dataset from a payments platform that has matured since it has launched.

## 2 Methodology

Our methods seek to replicate and revisit some the findings of Zhang et al. [18] as social media researchers have not generally reckoned with questions of reproducibility. Our primary focus was to recreate the first half of their paper which focuses on the high level statistics and metrics of the Venmo dataset. The difference is that now, with our much larger dataset, we can examine changes over time. We believe examining the high level changes in the platform is a necessary first step, before deciphering more granular interactions. To do so, we first conducted preliminary data analysis to understand some properties of the data we collected. Here, we describe our data collection, processing, and coverage of the dataset. Since Zhang et al.'s [18] findings in 2016, Venmo has rolled back some elements of API access and, as a result, we do have some small gaps in data. For example, we have no access to friend list information, which they had access to in the initial study. Therefore, we only seek to replicate the analytics of the transactions in the dataset and are unable to replicate any work around friend list data.<sup>1</sup>

Next, we built a transaction graph. The transaction graph represents a weighted directed graph based on transactions between two parties. Key graph metrics are calculated from the transaction graph such as average node degree and clustering coefficient. We explain what these metrics mean in the context of Venmo then later compare them to the results found by Zhang et al. [18]. Finally, we run Louvain community detection [4] on the transaction graph to discover what community structures are present in the graph.

We sought approval from our university IRB to analyze and investigate this publicly available data. Each of our research team members has completed human subjects training. All data was stored securely using campus supercomputer infrastructure, to which only authorized personnel have access. With regard to the ethics of using publicly available data, we see this in the same vein of analysis that uses publicly available Twitter data which seeks to minimize risk by using aggregated data [3, 5].

#### 2.1 Data Collection and processing

Venmo transactions were derived directly from the public Venmo API endpoint using a custom designed Python script executed on a virtual Amazon EC2 server instance (via Amazon Web Services). Data was collected for a period of four months in the spring of 2018. Using the public endpoint, we were able to access six years of transaction data, from 2012 to 2018, which amounted to nearly 340 million U.S. transactions. These transactions were not filtered in any way. Each transaction includes such as transaction ID, sender, receiver, message, time created, amongst other kinds of metadata generated from public posts. Each transaction can also have likes or comments, but only 2.8 percent of transactions have comments and 12.3 percent have any likes. We used our campus cyberinfrastructure to process

<sup>&</sup>lt;sup>1</sup> We contacted Xinyi Zhang and Ben Zhao, co-authors of [18], via an intermediary to obtain their friend list data. However, we were told via personal correspondence received by the intermediary that the data has since been lost.

the 400GB dataset as this was beyond normal workstation processor and memory capacity. The system allows for scalable Hadoop clusters as well as Apache Spark to execute the codebase we developed for the project.

#### 2.2 Transaction Graph Construction

We created a directed transaction graph from each Venmo transaction, where nodes represent users and edges represent transactions. The node list consists of every unique user we discovered in the dataset along with selected metadata attributes, including username and account creation date. Since the dataset delivered by the Venmo API is a list of transactions, it is inherently an edge list. The edge list included other relevant fields for each transaction such as the message and time. However, it consists of duplicate edges for repeat payments between users. For example, if A pays B ten times, there will be ten unique entries in the dataset. As a result, we condense the data into a weighted edge list where the weight represents the number of times an interaction happens.

Next, since Venmo has both charges and payments we evaluated whether our graph should reflect interactions between "actors" and "targets" or flow of money more generally. Since Zhang et al. [18] use the flow of money model, we swap the source and destination for any transactions that were marked as "charges". This way the source is the user who is sending money and the destination is the user who is receiving money. Lastly, we found that some users were making payments/charges to nonexistent users, resulting in null values. We removed these transactions from the dataset, which accounted for 971,218 total dropped transactions. In many cases we did not need every field from the edge list, so to improve computation times we created a smaller edge list, which only included the source, destination, and weight.

#### 2.3 Coverage Estimation

Zhang et al. [18] discovered that the transaction ID and user ID's on Venmo are assigned sequentially, therefore we can estimate the total number of transactions and users by finding the maximum respective ID. By the end of our collection date, the estimated number of total transactions is 1,160,228,993 and the estimated number of total users is 41,493,705. Our collected data consists of 23,133,264 unique userIDs over 341,309,788 total transactions. Our dataset therefore covers 29.6 percent of all public transactions and 55.67 percent of all user accounts that existed on the last day of data collection. The remaining transactions are private. User accounts not represented in this dataset have either made no transactions or are exclusively private ones. It is also useful to note that users can make both "payments" or "charges", with 83.43 percent of transactions being payments and 16.57 percent being charges. This shows that most people use Venmo to send money rather than request payments.

## **3 Results**

To undertake our replication study, we analyzed our dataset and present some basic trends and properties of the dataset. We begin by taking a high-level look at Venmo's user growth since its rollout and how long people use the platform once they open an account. We then calculate key metrics of the transaction graph so they can later be compared to the results of Zhang et al. [18]. Finally, we explore results using Louvain community detection. While the Louvain analysis is a brief contribution, we see utility in providing an additional feature of comparison to past work and possible directions for future work.

## 3.1 Periodicity: Back to School and Home for the Holidays

As figures 1-2 illustrate, our data indicates the same superlinear growth (i.e., grows faster than a linear function) of Venmo, as seen in Zhang et al.'s [18] original study. There is some seasonality to the transaction count and user growth. Further, both counts indicate peaks around August and September, and less activity around Thanksgiving (a U.S. holiday in late November), the December holiday season, and the summer. There are various potential explanations to the periodicity of user behavior on Venmo. A high amount of the user base could be made up of students, who transact with peers their own age on Venmo. This would explain the high amount of new user signups

at the time that American universities start the academic school year, dips in transactions (when these college student users go home for holidays and the summer break) and peaks in transactions around spring break (March), when users could be planning leisure travel or other spending activities involving splitting expenses through Venmo amongst groups of friends.

Alternatively, the explanation for periodicity could be more mundane than college students going home for the holidays, and instead, the trends we see may be tied to other seasonal events such as the American football season. Moreover, local maximums in transactions around August could be caused by moving-related expenses for example. It is interesting to observe and theorize the cause of periodic behavior in transaction and user registration behavior, but further research is needed to understand the cause of these phenomena.



Figure 1: Monthly new users encountered in dataset.



Figure 2: Monthly transactions in dataset

Figure 2 indicates that weekly transactions also grow superlinearly. When we look at weekly transactions, we see superlinear growth as well as some periodicity across months. The periodicity appears to be lower transaction volume in the summer. This also supports the theory that Venmo is used by many students, who mainly transact when they are at school. It is unknown whether the low points at New Year's Eve 2015, 2016, and 2017 are missing data or true values. If they are true values, it is interesting to not see a similar low point at New Year's Eve 2018. Further analysis into periodicity across months or days of the week may reveal more trends, such as paying rent and bills, or splitting drinks with friends on the weekend.

## 3.2 Activity Ratio: Active users vs. Instant Quitters

If a user's first activity is seen on date X, their last activity is seen on date Y, and the last day in the dataset is date Z, we can define their activity ratio (Y-X)/(Z-X). As Figure 3 illustrates, Venmo users typically either use the platform for one transaction and then immediately stop, or they continue using it. This Figure does present somewhat different findings to Zhang et al. [18]. Specifically, the prior study found a higher proportion of "permanent users" to "instant quitters", and we found a higher proportion of instant quitters (so called "one-and-done" users) than permanent users. This could be because any platform over time will accumulate inactive users. However, all other signs indicate that Venmo as a platform is quite active, so this higher proportion of currently inactive users is striking and worthy of further study.

If a user immediately stops using Venmo, their data point is near 0.0 and if a user frequently uses Venmo, their data point is near 1.0. The bimodal distribution shows that most users either used it only once, or use it very regularly.



Figure 3: Histogram of activity ratio

## 3.3 Usage Length

We find that 22.10% of users used Venmo for less than one day, a very slight decrease from Zhang et al.'s [18] finding of 22.5%. This divergence is extremely minor and indicates similar usage length. We also found that 41.95% of users used Venmo for more than one year, a significant increase from Zhang et al.'s [18] finding of 30% for the same measure. Figure 4 illustrates the entire histogram of account age by year. This indicates that Venmo users with public profiles who use the platform past their first day will most likely become regular users.





#### 3.4 Transaction & Social Graphs: The longer you stay the more you pay

In total, our transaction graph consisted of 23,097,596 unique nodes and 131,705,390 edges. The nodes have an average node degree of 5.70. In the context of Venmo, the average node degree indicates the average number of individuals any given user has transacted with.

The degree distribution shows the number of nodes at each degree value. Figure 5 illustrates that there is a large number of nodes with a degree less than 10, after which the nodes with higher degrees steadily drops off. This suggests that most users have a small number of individuals with whom they transact, while there are a few heavy users who transact with many users (we imagine these are potentially businesses). The fact that node degree is not uniformly distributed therefore does make sense.



**Figure 5: Degree Distribution of Transaction Graph** 

We found tie strength to be 2.59, which is the sum of all edge weights divided by the number of edges. Since the edge weight represents the number of times a transaction between the same two parties has occurred (e.g., B paid A 5 times), the tie strength represents the average number of times a given transaction will happen.

The clustering coefficient of a graph indicates the level of local connectivity between users (i.e., nodes) in the graph. A higher clustering coefficient indicates a greater "cliquishness". We find the local clustering coefficient to be 0.113. In addition, 31.11 percent of users have a clustering coefficient greater than 0.1 and 50.52 percent have a clustering coefficient greater than 0. Moreover, when we calculated the clustering coefficient with a smaller sample of nodes, we did find a slightly higher clustering coefficient of 0.133. It is also noteworthy that nodes with higher degree values have cluster coefficients that are orders of magnitude lower than others, a finding that supports the continued existence of informal and formal 'businesses' in the Venmo dataset.

K-core decomposition is a measure of how sparse a graph is by finding the maximum subgraph that is left after all nodes of degree less than k have been removed. In other words, the k-core would be the maximal subgraph such that the minimum degree in the subgraph is greater than k. Figure 6 illustrates how the average k-core level grows with node degree. The max core number of the Venmo transaction graph is 52.



Figure 6: Average K-Core level by Node Degree

We found the average path length, which is the average shortest path length between all node pairs, to be 5.686. To calculate this metric, we sampled 1000 nodes and calculated the shortest path from the random node to every node in the graph. Many of the sampled nodes had a degree of 0 or 1, resulting in a 0 or near zero path length. If we exclude zeros only, the average path length is a bit higher at 6.79. If both zeros and near zeros are excluded, this value rises to 7.47. Average path length in the context of Venmo or other social networks corresponds to "degrees of separation" or the average number of people that connect any two people.

Reciprocity is a measure of how likely it is that interaction between a pair of users will be bidirectional. In other words, for every pair of users that has interacted (money has flowed one way or the other), how likely is it that the interaction has been bidirectional (money has been transacted both ways)? We found 24,283,077 edges that were bidirectional and 131,705,390 total edges (i.e., that were either uni- or bi-directional). Using the traditional definition of reciprocity, we found the reciprocity to be 24,283,077 / 131,705,390 = 0.184. This figure is surprisingly low. A number closer to 1 would suggest that almost all users reciprocate a payment. Given that we see Venmo as a social network, one might expect to find a result close to the value 1. However, the lower number seems to confirm reciprocity levels that a

payment platform should have, which accounts for businesses and other heavy users who only pay out or receive funds.

We found assortativity to be -0.003319 or nearly zero. Assortativity is the probability for nodes to connect to other nodes of similar degree. The near zero number is not unsurprising as payments happen both between nodes of similar degree and dissimilar degree. Users pay their friends, but also pay businesses or heavy users; this property has not changed since the inception of Venmo.

The size for the largest Strongly Connected Component (SCC) is 57.40 percent of the graph and the Largest Weakly Connected Component (WCC) is 95.54 percent of the graph. The strongly connected component represents the largest subgraph where this is a transaction path from every user to every other user. Meanwhile, the weakly connected component is simply the largest connected component in the graph; hence, it is a larger percentage of the graph.

# 4 Discussion

In this study, we found that while Venmo has grown significantly in the past few years, the transactions graph structure has largely remained the same. A major difference between our data set and that of Zhang et al. [18] is that we had over three times the volume of transactions, and that our date range extended until 2018, whereas their work terminates with data in 2016. Comparing the graph metrics of our new, larger transaction graph to the smaller transaction graph that Zhang et al. [18] used, we find that many of the structural properties have remained the same despite Venmo acquiring new users (see Table 1). Nearly all of the metrics we calculated are similar to Zhang et al.'s [18] findings. This was surprising as we had expected Venmo's growth to be reflected in changes in the graph. The fact that the metrics we calculated have not drastically changed does, however, confirm a successful replication of Zhang et al.'s [18] work. These results also indicate that as new users joined the platform at increased rates, the network structure largely remained consistent. In addition, we see that new users are not bridging pre-existing communities. If users were bridging or joining new communities in the social network, then we would see an increased clustering coefficient, and an increased average path length.

The average degree, tie strength, clustering coefficient, and average path length were the only metrics that decreased slightly. This suggests the transaction graph of Venmo is becoming slightly sparser or less clustered as the platform ages, grows, and users pay for different types of things. This might suggest that while Venmo's user base has grown, user retention may be struggling. Table 1 compares results between our study's results and those of Zhang et al. [18]. The following subsections discuss the significance of these comparisons.

	Zhang et al.'s [18] Transaction Graph Metrics	Our Transaction Graph Metrics
# of Nodes	7.08M	23.09M
# of Edges	35.0M	131.7M
Average Degree	9.89	5.70
Tie Strength	3.22	2.59
Clustering Coefficient	0.147	0.113
Average Path Length	6.98	5.686
Assortativity	-0.0022	-0.0033

 Table 1: Overall comparison of transaction graph metrics between Zhang et al. and our findings.

Average Reciprocity	0.147	0.184
Largest SCC	56.10%	57.40%
Largest WCC	95.50%	95.54%

## 4.1 Degree and Tie Strength

As Table 1 indicates, Zhang et al. [18] had a degree of 9.89 whereas our average node degree is almost half (5.70). This is likely due to an increase of instant quitters ("one-and-done" users) and it is expected that this figure will grow as more people try Venmo, thus creating a sparser network.<sup>2</sup> Following a lower degree, the tie strength has also decreased slightly (their 3.22 to our 2.59; see Table 1), which may also indicate a sparser network. Like Zhang et al. [18], we similarly find that the growth of one-and-done users has held steady over time. Ultimately, this indicates that individuals using Venmo and leaving has not changed over time, confirming that the platform has retained a similar usage pattern.

# 4.2 Clustering Coefficient

Zhang et al. [18] found the clustering coefficient to be 0.147 while we found 0.113. Though our finding is somewhat smaller, it does not substantively indicate any noteworthy change. Additionally, Zhang et al. [18] found 39.09 percent of Venmo users had a clustering coefficient more than 0.1, whereas we found 31.11 percent of users did. Since the clustering coefficient of the newer, larger dataset is smaller, it may suggest that Venmo users have become less densely clustered as the platform has grown. This hypothesis follows with the lower average degree of nodes found previously. While the metric values have shifted slightly since Zhang et al.'s [18] study, these results indicate that the Venmo graph has maintained the same level of local nodal connectivity as it has grown.

# 4.3 K-Core Decomposition

We compare the k-core decomposition to Zhang et al. [18] by examining the difference in their plot of average core level as a function of degree to our plot. We see similar results to their findings in the shape of the plot line (see Figure 6). However, in our study, the plot line approaches a k-core of 35 for nodes with degree 80 to 100, while Zhang et al. [18] leveled out around 25 for the same range. This is likely due to the increased size of our graph.

In Zhang et al. [18], the average core level as a function of degree was found to be substantially similar to that of Twitter, and Twitter activity was found to closely resemble real world social relationships from other prior research [3], implying Venmo activity resembles real world social relationships. The difference in our findings could indicate several things: an incorrectly calculated metric, a recent divergence from Venmo's similarity to Twitter, changing use cases of Venmo that no longer resemble real world social relationships, or a change in the structure of real world social relationships.

# 4.4 Average Path Length

The average path length we computed, 5.69, compares well with Zhang et al.'s [18] of 6.98. It is not known if they count 0-degree nodes or exclude them. Therefore, our value is likely similar to their findings. The smaller average path length suggests that the small-world effect of Venmo has become stronger as the platform has grown, suggesting

<sup>&</sup>lt;sup>2</sup> Another possible reason for the discrepancy in average degree is a calculation error. Average degree for a directed graph is computed by the number of edges divided by the number of nodes, whereas the formula for an undirected graph multiples the quotient by 2 since each edge is connected to two nodes thus counted twice. Zhang et al. [18] use a dataset with 35.0M edges and 7.08M nodes, resulting in an average degree of 4.94. We notice that multiplying this product by 2 gives 9.89 suggesting that Zhang et al. [18] may have used the formula for an undirected graph.

a smaller degree of separation between users. This finding is significant because it shows that Venmo's connectivity has moved closer to "six degrees of separation" [15], meaning that any Venmo user is roughly six users away from any other user.

# 4.5 Average Reciprocity

Zhang et al. [18] found an average reciprocity of 0.147 and we found 0.184. While our result is slightly higher than Zhang et al. [18], it is still quite similar in value, confirming comparable reciprocity levels as 2016. This indicates that people's use of Venmo as a peer-to-peer payment system has remained the same, simply with more users sending and receiving money from one another. In fact, this property has perhaps become stronger over time. Consequently, this comparison shows that the reciprocal payment patterns on the Venmo platform have remained the same.

# 4.6 Assortativity

We found almost the exact same assortativity as what Zhang et al. [18] found. This finding indicates that Venmo users still transact with a variety of users with diverse degree values. Moreover, since our resulting value is so close to that of Zhang et al., it also confirms a reproduction of their methods and confirms the validity of their work. Ultimately, this comparison also provides further evidence that Venmo users continue to pay or make requests to users who are not exclusively the same, degree-wise, to them.

# 4.7 Connected Components

Our results for strongly and weakly connected components are almost the exact same as Zhang et al. [18] found. This suggests that as Venmo's user base grew, new users were simply added to existing connected components, since they were most likely invited by friends. As a result, the relative connected component sizes remain the same and we are potentially seeing a continued level of chain migration of existing real-world social ties into the Venmo platform. Our finding also confirms that we have successfully reproduced calculations for SCC and WCC metrics as deployed by Zhang et al. [18]. It indicates that with newer and more voluminous data that this important set of graph metrics remains the same.

## 4.8 Users & Communities

We used the gold-standard Louvain community detection algorithm [4] on the transaction graph to further understand the interaction properties of Venmo users. We modify the graph to be undirected since Louvain is not well defined for directed graphs [7]. This results in 447,487 unique communities, with the smallest community size being 1 and the largest being 2,036,490. The average community size is 51, and the modularity is 0.77. A modularity above 0.3 indicates meaningful community structures [10].

We did find community detection results quite different from the results of Zhang et al. [18] as they only find 815 communities with a modularity of 0.836. We did expect our larger graph to have more communities than Zhang et al. [18] found, but not orders of magnitude more. This is quite likely a result of representing the graph as undirected instead of directed. We therefore discount this aspect and focus on the fact that we found a similar modularity value as Zhang et al. [18], which suggests the community structure has been maintained as Venmo has grown.

However, it should be noted that our community detection results revealed more communities on Venmo of a smaller size. This potentially follows face-to-face interactions where communities of people on Venmo would be smaller, restricted to around 50 people (e.g., imagine a group of friends exchanging transactions over a period of time). Meanwhile, Zhang et al. [18] only found 815 communities which covers 7 million unique users, which would result in much larger communities (~140,000 people per community). The individual users within each of these large communities do not really have any discernible relation. Therefore, we suggest that future work on Venmo as well as other payment platforms provide community detection based on an undirected graph.

# **5 Limitations and Future Work**

While our study employs and extends many of the methods used in Zhang et al. [18] with a larger set of data, future work could focus on analyzing the more granular interactions between users and how they have changed within a larger set of data. For example, clustering groups of users by payments or transaction memo behavior would be one application that would benefit from more data. Like many other researchers who use data extracted from platforms with APIs, the limitations to our study are largely related to data access and the new limits that Venmo has imposed on the public API since 2017 [6]. Because we do not have the friendship data that Zhang et al. [18] were able to collect at the time of their study (as the Venmo API formerly allowed developers access to these data), we could not recreate the statistics of a graph of transactions between friends. As a result, we were only to examine the transaction graph and reproduce the metrics for that graph without comparing it to any social graph. However, since Zhang et al. [18] primarily focused on the transaction graph in their study, we were still able to conduct a full-fledged replication and validation of their work. Ultimately, the API changes had a limited impact on our study.

Additional technical challenges restricted our ability to replicate aspects of community detection and clustering. Here, we were not limited by our access to libraries that could perform the analysis, but rather our inability to parallelize the given libraries. For example, we did implement a library to successfully perform Louvain community direction on a directed graph. However, due to the much larger size of our dataset than Zhang et al. [18] and the fact that the library was single threaded, the code ran for 48 hours and then timed out. Further details of the technical roadblocks are discussed at our study's GitHub page.<sup>3</sup>

Future work could explore topics such as the discrepancy between charges and payments. As mentioned previously, there are almost four "payment" transactions for every one "charge" transaction. This has a potential impact on how people discuss and facilitate Venmo transactions offline.

Another potential avenue for future work would be to examine transaction types according to more temporal properties of the dataset. For example, the graph metrics could be calculated on a month-by-month, week-by-week, or even dayby-day level of granularity. Moreover, future work can partition the data with a single year or a range of years of data and examine how various metrics differ based on comparisons. Potential structural-temporal graph properties might emerge that would explain more about how users interact on Venmo, such as paying rent and splitting utility bills with housemates, or sending graduation or wedding gifts amongst family members.

## **6 CONCLUSION**

The replication of social media research is important for validity, reliability, generalizability, application to other areas, discovering knowledge, applying novel methods, and inspiring future research. Computational social science on social media platforms remains challenging for researchers because platform data continues to be difficult to access, is constantly changing, contains sensitive Personal Identifying Information (PII) (sometimes with ethical constraints) [11], and relies on ever-changing and APIs that often lack thorough documentation. Despite these various challenges, we found ways to reproduce most of Zhang et al.'s [18] early work on Venmo. The platform has grown exponentially and processes many more transactions since their study. We found that the majority of the properties of transactions on the Venmo have remained consistent as it has grown. Any deviations were explained by the increase of one-and-done instant quitter users, which increased sparsity metrics and decreased clustering metrics. Our study also provides a basis for future work on even larger datasets of Venmo transactions without rechecking the underlying structure and assumptions of the transaction graph.

We recreated some of the methods and metrics from Zhang et al.'s [18] pioneering work and discovered a slight weakening in the structure of the graph. As new users are added to Venmo, they tend to interact with those who initiated them, and no one else. This is an important and noteworthy finding for social media research on early user behaviors when joining a platform. Importantly, new users do not actively engage in bridging unconnected communities using their connections. Rather, they tend to only use Venmo for transactions between themselves and the people who initiated them into the Venmo platform.

<sup>&</sup>lt;sup>3</sup> See: https://bit.ly/3ekVWG4

Ultimately, our study is valuable as we were able to replicate some of Zhang et al.'s [18] findings with far less API access but with much more data. There remains a severe paucity in replicability studies in social media research. Moreover, there are none that we discovered that have been done on social payments. Our work thus serves as an intervention - we encourage others to engage in replication studies to enhance the reliability, generalizability and impact of social media research because such contributions are valuable and necessary to the literature and the future of the field of social media research.

#### **REFERENCES:**

- [1] Acker, A. and Murthy, D. What is Venmo? A descriptive analysis of social features in the mobile payment platform. *Telematics and Informatics*, 101429 (2020). https://doi.org/10.1016/j.tele.2020.101429
- [2] Acker, A. and Murthy, D. Venmo: Understanding Mobile Payments as Social Media. In *Proceedings of the 9th International Conference on Social Media and Society* (Copenhagen, Denmark, 2018). Association for Computing Machinery.
- [3] Bild, D. R., Liu, Y., Dick, R. P., Mao, Z. M. and Wallach, D. S. Aggregate Characterization of User Behavior in Twitter and Analysis of the Retweet Graph. ACM Trans. Internet Technol., 15, 1 (2015), Article 4.
- [4] Blondel, V. D., Guillaume, J.-L., Lambiotte, R. and Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008, 10 (2008), P10008.
- [5] Caraway, M., Epstein, D. A. and Munson, S. A. Friends Don't Need Receipts: The Curious Case of Social Awareness Streams in the Mobile Payment App Venmo. *Proc. ACM Hum.-Comput. Interact.*, 1, CSCW (2017), Article 28.
- [6] Constine, J. Venmo Halts New Developer Access To Its API. TechCrunch, 2016. https://techcrunch.com/2016/02/26/how-not-to-run-a-platform/
- [7] Dugué, N. and Perez, A. Directed Louvain : maximizing modularity in directed networks. Université d'Orléans, 2015. http://dx.doi.org/10.13140/RG.2.1.4497.0328.
- [8] Han, B. and Windsor, J. User's Willingness to Pay on Social Network Sites. Journal of Computer Information Systems, 51, 4 (2011), 31-40.
- [9] Kow, Y. M., Gui, X. and Cheng, W. Special digital monies: The design of Alipay and WeChat wallet for mobile payment practices in china. Springer, 2017.
- [10] Kwak, H., Choi, Y., Eom, Y.-H., Jeon, H. and Moon, S. Mining communities in networks: a solution for consistency and its evaluation. In *Proceedings of the Proceedings of the 9th ACM SIGCOMM conference on Internet measurement* (Chicago, Illinois, USA, 2009). Association for Computing Machinery.
- [11] Sprague, R. and Ciocchetti, C. Preserving Identities: Protecting Personal Identifying Information Through Enhanced Privacy Policies and Laws. *Alb. LJ Sci. & Tech.*, 19 (2009), 91.
- [12] Tenopir, C., Allard, S., Douglass, K., Aydinoglu, A. U., Wu, L., Read, E., Manoff, M. and Frame, M. Data Sharing by Scientists: Practices and Perceptions. *PLoS One*, 6, 6 (2011), e21101.
- [13] Tenopir, C., Dalton, E. D., Allard, S., Frame, M., Pjesivac, I., Birch, B., Pollock, D. and Dorsett, K. Changes in Data Sharing and Data Reuse Practices and Perceptions among Scientists Worldwide. *PLoS One*, 10, 8 (2015), e0134826-e0134826.
- [14] Venmo Signing Up. 2020. https://help.venmo.com/hc/en-us/sections/201936967-Signing-Up
- [15] Watts, D. J. Six degrees : The science of a connected age. W. W. Norton & Company, New York, NY, 2003.
- [16] Webb, H., Jirotka, M., Stahl, B.C., Housley, W., Edwards, A., Williams, M., Procter, R., Rana, O. and Burnap, P. The ethical challenges of publishing Twitter data for research dissemination. In: *Proceedings of the 2017* ACM on Web Science Conference (2017), 339-348.
- [17] Wu, Z. and Ma, X. Money as a Social Currency to Manage Group Dynamics: Red Packet Gifting in Chinese Online Communities. In Proceedings of the Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (Denver, Colorado, USA, 2017). Association for Computing Machinery.
- [18] Zhang, X., Tang, S., Zhao, Y., Wang, G., Zheng, H. and Zhao, B. Y. Cold hard e-cash: Friends and vendors in the Venmo digital payments system. In *Proceedings of the Eleventh International AAAI Conference on Web and Social Media* (Montréal, Québec, Canada, 2017). Association for the Advancement of Artificial Intelligence.