Predicting Donor Behavior using the Dynamics of Event Co-Attendance Networks

Shwetha K. M. Srikanta, Katie L. Pierce, Joshua E. Introne, Chilukuri K. Mohan, Sucheta Soundarajan

Abstract This study investigates the dynamics of charitable donor co-attendance networks to help understand fundraising outcomes. We analyze a multi-year network of co-attendance at fundraising events, examining topological structure, node characteristics, and network properties. Key findings include a 76% increase in giving value for donors with increased centrality rank and a positive peer effect indicated by increased co-attendance with high-capacity donors, among donors who increased or maintained their giving levels. Additionally, we observe assortativity in giving patterns among co-attending donors, and identify interlinked cliques of communities with varying sizes, with larger communities having a higher proportion of wealthy donors. Furthermore, we develop a Long Short-Term Memory (LSTM) network to predict donor propensity, considering node-level features such as assortativity, centrality, clustering coefficient, and exposure to high-capacity donors. Our results demonstrate that a multi-variate model outperforms baseline statistical models, enabling outcome assessment before events. These findings have implications for budget allocation strategies and maximizing fundraising effectiveness.

Keywords Charitable donations, Co-attendance network, Deep Learning, Multi-variate forecasting model.

1 Introduction

For non-profit organizations, fundraising events serve as platforms for donors and staff to strengthen or form new connections. The outcomes of such events have been observed to be influenced by similarities between donors, their relationships with the recipient organization, and interactions with the soliciting individuals. [11, 27]. We propose a network-theoretic perspective on analyzing such events, exploring

Shwetha Srikanta, Katie Pierce, Joshua Introne, Chilukuri Mohan, Sucheta Soundarajan
Syracuse University, Syracuse, New York, USA, e-mail: smanchin@syr.edu, klpierce@syr.edu, jeintron@syr.edu, ckmohan@syr.edu, susounda@syr.edu

the extent to which the event co-attendance network influences giving behaviors of donors.

This chapter expands on our study presented in [31], analyzing multi-year data (from a large private research university) that includes donations made by individuals and events attended by those individuals (Section 2). Using this data, we construct a co-attendance network, where nodes represent donors and edges indicate whether donors have attended at least one fundraising event together. We recognize that other factors beyond the co-attendance network also influence giving behavior, such as undocumented personal relationships, mutual interactions, influence networks, and recommendations.

Our analysis reveals low-moderate positive assortativity, indicating a tendency for donors with similar giving behaviors to co-attend events. We observe a positive correlation between centrality and the amount of giving, consistent with prior studies [19] in this area. More importantly, our findings indicate that co-attendance with high-capacity donors correlates with increased donation amounts. We also explore other network properties that may influence giving behavior, discovering that larger communities have a greater proportion of above-median and high-giving donors.

In addition to the work presented in [31], this chapter includes the following additional results.

1. We enhance our formulation for updating edge weights in the network construction that captures the potential strength of interactions between individuals while co-attending events with emphasis on event size.
2. We predict donor propensity using node-level features extracted from the co-attendance network, including assortativity with neighbors, centrality rank, clustering coefficient, and exposure to donors within the local neighborhood and community.
3. We extend our analysis beyond donors retained between non-overlapped windows ($S = 456$) to include all (including irregular) donors, significantly increasing the sample size ($S = 99431$). This facilitates exploring a wider range of donor behaviors, so that we can identify patterns in donations across different donor segments.

We propose a multivariate forecasting model that significantly outperforms the baseline in predicting the giving class of an individual in the current time window based on the co-attendance network node properties in preceding time windows. The AUC for our model exceeds that of the baseline for both the donor segments. The node features are perfectly separable in our multi-year analysis, indicated by an AUC of 1.0 for all three classes. Learning from the network snapshots enhances the model prediction capability because it incorporates the dynamic nature of the network and its influence on donor behavior.

The proposed multivariate model outperforms all statistical baselines in terms of median absolute error (MAE) for both regular and irregular event attendee segments. Although the performance of statistical models in predicting the donation amount for regular donors is comparable to that of the proposed multivariate model (in terms of RMSE), the latter demonstrates superior performance for a filtered set of donors (i.e., excluding very high donors). In this segment, the multivariate model achieves a lower test RMSE of $16874.49 compared to the statistical model's RMSE of $22428.29.

Thus, our research presents a new approach to predicting donor propensity by incorporating network properties, specifically event co-attendance networks, unlike previous models that have focused on personal traits and other properties. Our analysis can be useful to nonprofit organizations in identifying potential donors and donor segments, understanding donation patterns, and optimizing their resource allocation to various donor segments such as retained donors, new donors, and donors who changed their behaviors. This helps in improved targeting and thus improves fundraising outcomes.

Section 2 describes the properties of the dataset. Section 3 discusses prior work related to network analysis in fund-raising efforts, as well as the baseline statistical approach and relevant neural network models. This is followed by Section 4, which describes our methodology and parameters of the multi-variate model. Next we present our analysis in Section 5 followed by discussion in Section 6. Conclusions are presented in Section 7.

2 Dataset Description and Statistics

The dataset used for our study contains over 46,175 (including staff) unique donors who contributed varying amounts to events and non-event-associated donations. The data spans over 3000 events organized by a large private research university from 2014 to 2021. These events include meetings, fundraising campaigns, alumni engagement programs, networking events, and lectures and donations received by the university are utilized for endowments, scholarships, and funding various university facilities.

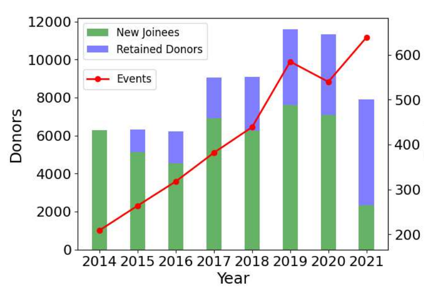

Fig. 1 Growth of donors and events organized in period 2014-2021 [31]

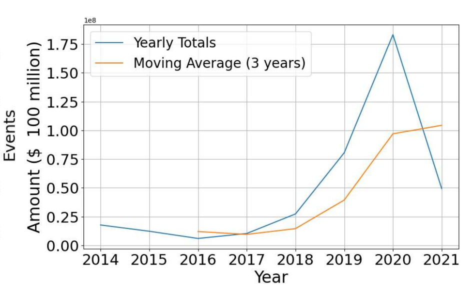

Fig. 2 Yearly totals and 3-year moving averages for all donors in period 2014-2021

Fig. 1 shows the retained (existing) donors and newly joined donors for each year, from 2014 to 2021. We see that each year the events draw a similar number of new donors and with time there is an increase in the number of retained donors. Fig. 1 also provides the number of unique events organized each year, which increases with time. The donors have varying primary constituencies with the University, including alumni, students, parents or spouses of students/alumni, and staff. For this analysis, we focus on two types of giving: promised lifetime giving (which includes amounts that have not yet been given), and the actual amount given each year.

Fig. 2 illustrates the yearly totals, the actual data points that represent the totals for each years. The moving average curve helps to smooth out short term fluctuations and shows an upward trend. To enhance the prediction capabilities, we have augmented our previous dataset with new fields, such as first campaign gifts.

Table 1 Sample size for dataset

Analysis	Attendee type	Individual donors	Sample size
Yearly	Regular attendee (all events)	57	456
Yearly	All attendees (small events)	175	5578
Multi-year	All attendees (all events)	44385	99431

We perform various analyses to forecast donor propensity. Initially, we utilize the annual co-attendance network to examine two distinct donor segments. The first segment comprises consistent donors who consistently attend events over the course of eight years. The second segment consists of both regular and irregular donors who attend small-sized events.

Additionally, we conduct a multi-year analysis to evaluate the multivariate model's performance on donor activities within a rolling window of 3 years. This investigates how the network dynamics potentially influence donor propensity. The sample sizes representing time series data of individuals attending events and contributing amounts for all dataset types are listed in Table 1.

3 Background

In this section, we discuss prior research in three areas related to our work. First, we summarize network analytics work addressing donation data. We then describe the baseline statistical approach used to evaluate our proposed approach. We then discuss the use of recent neural network models for forecasting.

3.1 Donation Data Analysis

Prior analyses have shown that effective communication with donors can increase charitable donations, with people more likely to donate when they perceive similarities with fundraisers [7, 30]. Donors also give more often to organizations to which they feel connected, such as their *alma mater* [4]. Overall, research suggests a positive social influence among similar individuals in donor-solicitor and donor-recipient dyads, while the effects of social influence on charitable giving in donor-donor dyads can be both positive and negative [25, 29]. Existing research has found that donor-donor dyads exhibit assortativity, meaning individuals tend to donate similar amounts as their peers [20, 15]. Studies have also examined the extent to which giving behavior is influenced by similarities in personal characteristics, such as gender, location, income, age, and relationships with others [32].

Network properties have been used to detect donor influence, particularly on social media [8, 27, 34]. Studies have found that a Twitter campaign network's size, decentralization, and cultural norms are positively correlated with the average donation amount per day [33]. Authority and centrality within internal social connections also impact success in fund-raising [12]. Additionally, social connection variables have been found to improve the performance of algorithms that predict funding outcomes, with external social connections having a stronger impact compared to internal connections.

Numerous studies have focused on constructing predictive models to enhance the effectiveness of fundraising efforts within charitable organizations [23]. These research endeavors have attempted to predict a donor's giving likelihood (propensity) [28], and to forecast the specific amount a donor might contribute, based on historical data. The influence of organizational connections on individuals' donation propensities has been studied based on the preceding 3 years of donations, considering age, gender, race/ethnicity, and relational characteristics as features [2].

Research has also focused on identifying individuals who change their giving behavior over time [26, 1]. Previous research has also explored donor segmentation based on various factors, including altruism, belief in mission clarity, and susceptibility to external influences [18].

In this study, we extract individual (node)-level features from the co-attendance network to predict donor propensity. We segment donors based on event attendance frequency and compare this against baselines. Our approach includes conducting a multi-year network analysis to study the potential influence of network dynamics on donor behavior, aiming to predict donor propensity. For classification tasks, we use simple time series-dependent baselines, while for regression tasks, we employ more complex statistical methods such as ARIMA modeling. This work contributes to the limited research on associating fundraising outcomes with longitudinal network analysis and on predicting donor propensity using such data. The following subsections describe our baselines in more detail.

3.2 Baseline

For our classification task we employ a baseline that predicts the giving class of each individual for the next time window to be the same as their most recent value. This is a simple approach based on the assumption that the future behavior of an entity will be similar to its past behavior.

We also use two baseline methodologies for regression tasks: Exponential smoothing [10] and Auto-Regressive Integrated Moving Average (ARIMA) models [22]. These approaches offer complementary techniques to tackle the complexities of forecasting. While exponential smoothing models focus on the seasonal patterns and trends in the data, the ARIMA model captures the autocorrelation structure inherent in the dataset. In an autoregressive model **(AR (p))** of order p, the variable of interest is forecast using a linear combination of past values of the variable. A moving average model of order q **(MA (q))**, on the other hand, uses past forecasting errors in a regression-like model. If we combine differencing with autoregression and a moving average model, we obtain a non-seasonal ARIMA model, expressed as follows.

$$y'_t = c + \phi_1 y'_{t-1} + \ldots + \phi_p y'_{t-p} + \phi_1 \epsilon_{t-1} + \ldots + \phi_{t-q} \epsilon_{t-q} + \epsilon_t . \tag{1}$$

3.3 Neural network models for forecasting

Artificial neural networks provide forecasting algorithms that use mathematical models inspired by biological neural networks and capture nonlinear relationships between the response variable and its predictors [21]. They were first used for forecasting in the early 1990s [16, 3] but the recent introduction of well-designed recurrent neural network (RNN) modules has resulted in extraordinary successes in forecasting for several practical applications [17]. For our multivariate time series model, we employ Long Short-Term Memory (LSTM) [14] and Gated Recurrent Unit (GRU) networks [5].

An LSTM is a type of recurrent neural network (RNN) module designed to learn long-term dependencies. Unlike traditional RNNs, LSTMs have a cell state to facilitate remembering information over extended periods, and use gates to control the flow of information into and out of the cell state, enabling learning of complex temporal relationships. GRU networks are similar, but use a simpler recurrent network module. The key differences between LSTM and GRU lie in their architectures and associated trade-offs. LSTM networks have more gates and parameters, providing greater flexibility and expressiveness, but requiring higher computational cost and overfitting risk. GRU networks are simpler and faster since they have fewer gates and parameters, but don't always perform as well as LSTM networks. LSTM networks maintain separate cell state and output channels, hence storing and outputting different information, unlike GRUs whose single hidden state may limit its capacity. Both models exhibit varying sensitivities to hyperparameters such as learning rate, dropout rate, and sequence length [24].

4 Methodology

Our primary objective is to understand the relationship between the network properties of an individual in the dataset and that individual's giving (propensity and level) by analyzing longitudinal network data. This information can be valuable in evaluating the potential effectiveness of fund-raising campaigns. To achieve this, we construct weighted network snapshots at the end of each specified time window and measure the co-attendance network properties, using them to predict donors' giving propensity for the subsequent time window.

4.1 Network Construction

The usefulness of events for predicting donor behavior declines with time, and donor participation in recent events is more important than in older events. One possible way to represent this fact in our updated co-attendance network is by halving the weight in every successive time slice [31].

$$W_{ij}^{t} = N_{ij}^{t} + \frac{W_{ij}^{t-1}}{2^{d}} \tag{2}$$

Event size matters: in smaller events, individuals have a higher chance of engaging in conversations and interacting with one another, while in larger gatherings, one-on-one interactions between two co-attendees is less likely. Equation 3 shows how we may use the relative strength of possible interactions between two co-attending individuals to update edge weights, using a slower decay function than Equation 2, capturing long term dependencies for our multivariate models (Section 5.5 to 5.7).

$$W_{ij}^{t} = \sum \frac{k}{|E^{t}(i,j)|} + \frac{W_{ij}^{t-1}}{d} \tag{3}$$

4.2 Node features

To predict donor propensity, we leverage network structural characteristics. Since we focus on individual donations, we consider properties at the node (individual) level. Our analysis reveals similarity in giving behavior among neighbors of a node, indicating the impact of assortativity on donor propensity. We observe a positive peer effect, with increased giving levels among donors who showed increased co-attendance with high-capacity donors. We hence investigate the influence of an individual's local neighborhood and their community (identified using the Louvain community detection algorithm [1]) on donor propensity. Additionally we consider centrality, as more central individuals tend to have higher giving levels supported by prior studies and our analysis. We also take into account the clustering coefficient, which measures the extent to which nodes cluster together into well-knit sub-networks.

4.2.1 Assortativity:

This property indicates the preference for nodes to attach to others that have similar characteristics, i.e., whether individuals with similar giving behaviors are more likely to co-attend the same events. To capture the local assortativity property, we use the simple and well-known *EI homophily index*, a measure of in-group (internal) and out-group (external) preference [6] whose value ranges from -1 (*complete homophily*) to 1 (*complete heterophily*), using Equation 4.

$$EI = \frac{External - Internal}{External + Internal} \tag{4}$$

4.2.2 Centrality:

The centrality of an individual in a network impacts their access to resources and opportunities. Central individuals are invited to more events, meet more people, and benefit from these resources, fostering a sense of gratitude and a desire to give back to the community. Also, highly influential donors may be specifically invited to larger gatherings, further increasing their centrality in the network [27]. Here we consider weighted Eigenvector centrality as well as degree centrality [13].

4.2.3 Clustering coefficient:

This measures the extent to which nodes from tightly knit clusters, providing insights into the structure of the co-attendance network [13].

4.2.4 Local Exposure:

This refers to the percentage of neighbors of a node that represent donors, and is important since the effect of co-attending donors may be diluted by the presence of many co-attending non-donors. Co-attending events with a higher number of donors can create a greater positive peer effect and increase the likelihood of donating.

4.2.5 Community Exposure:

Membership in a community (obtained using the Louvain method) with more donors may be more meaningful than local exposure which may be merely coincidental in a co-attendance network. We also expect giving behaviors to be positively influenced if the community to which a node belongs has a greater proportion of high-capacity donors, who may exert greater influence [11]. Hence we consider two variants of *Community Exposure*, respectively referring to the percentage of a node's community members that represent (i) donors, and (ii) high-capacity donors.

4.3 Multivariate forecasting model

We split the data into training and testing as described in Section 4.3.1. Then we train neural network models with the LSTM and GRU modules, using the node features extracted for each individual based on their previous time window of co-attendance. The models predict the giving class or donation amount for the test data.

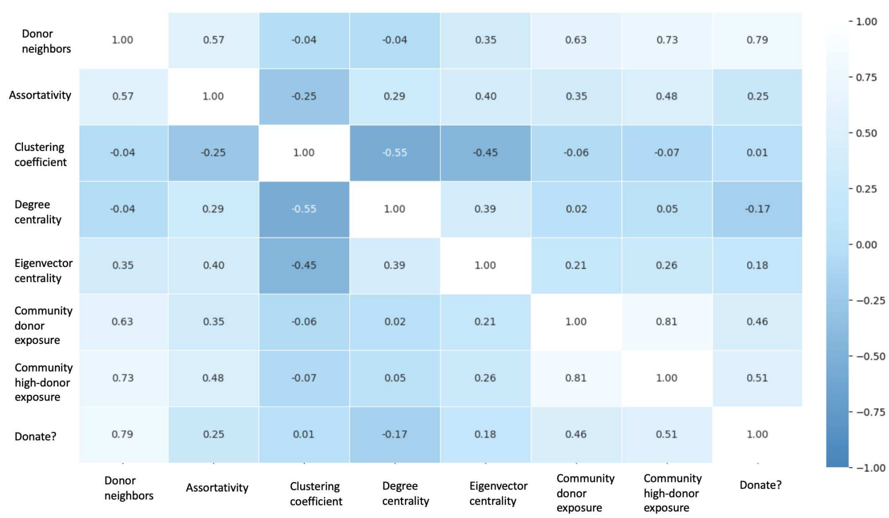

Fig. 3 Correlation heatmap

4.3.1 Model Parameters:

During model training, we incorporate lagged input features along with other node features. Specifically, we include lagged EI homophily index, which is positively correlated with donation likelihood, as shown in Fig. 3, and lagged exposure to high-capacity donors, considering the observed correlation between increased giving and increased co-attendance with high-capacity donors. By incorporating assortativity as a lagged feature, we capture the dynamic nature of this similarity and its influence on donor behavior. Additionally, we include a lagged output feature, i.e., giving class or amount.

We perform cross-validation on time series data, training the model for each time window and testing it on the subsequent window to prevent leakage of future data during training [9]. For binary classification, we employ *binary cross-entropy* loss with a *sigmoid* activation function, while for multi-class classification, we use *categorical cross-entropy* loss with a *softmax* activation function. For the regression task we minimize *mean squared error* using the *ADAM* optimizer.

4.3.2 Donor Segments:

Different parts of the donor dataset have different characteristics in event attendance and donor behavior, and are best understood separately. Hence we segment donors based on event attendance frequency (regular and irregular donors attending small events), identifying patterns related to co-attendance and giving.

4.3.3 Multi-year Analysis:

In order to identify network characteristics that drive changes in donor behavior, it is essential to analyze a multi-year network. To achieve this, we construct weighted network snapshots for a rolling window of three years using Equation 3. By extracting the aforementioned features from these network snapshots, we can test the multivariate model and investigate how network dynamics potentially influence donor propensity.

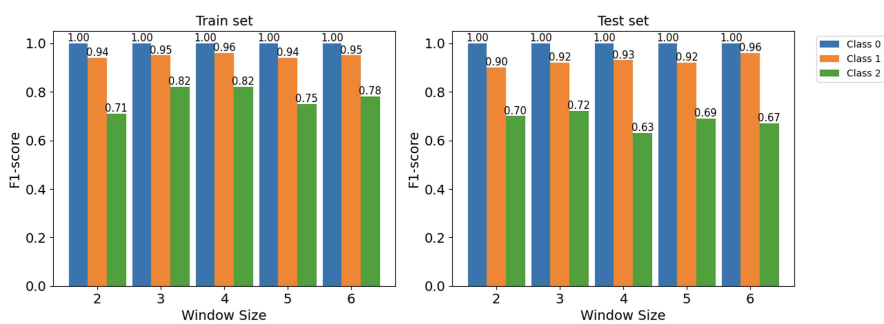

Fig. 4 Comparison of f1-score for the three classes for varying window sizes.

4.3.4 Choice of Time Window:

We perform sensitivity analysis for various time windows ranging from very short $w = 2$ to $w = 6$ and record the F1-score. The bar plot in Fig. 4 shows the optimal window (3 years). A small window size is appropriate for our dataset as it spans over only 8 years. A very short 2-year window is more sensitive to rapid changes in donor activity and captures the rapid fluctuations in yearly totals as seen in Fig. 2, whereas more smoothing occurs with a 3-year moving average window.

4.3.5 Balancing classes:

The time-series data for all donor segments exhibits class imbalance due to unequal numbers of donors and non-donors. The regular attendee segment has a larger set of donors (23% non-donors and 77% donors), while the irregular attendee segment shows a contrast (65% non-donors and 35% donors). To address this, we assign higher weights to the minority class compared to the majority class during loss function calculation using the inverse class frequency method indicated in Equation 5.

$$W_j = \frac{|S|}{(|N_c| * |S_j|)} \tag{5}$$

Here, the weight for each class w_j is computed using the sample size $|S|$, $|N_c|$ represents the number of classes for the target, and $|S_j|$ is the number of samples in class j. This approach prevents over-fitting and mitigates the impact of class imbalance, and is preferable to resampling or generating synthetic samples.

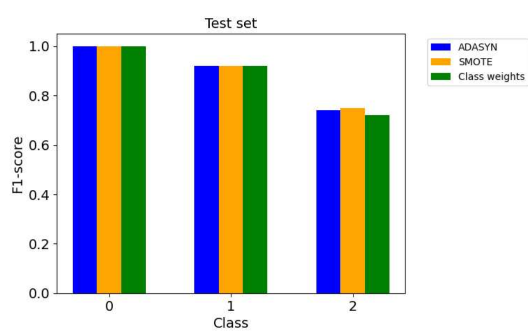

Fig. 5 Comparison of F1-scores obtained using various class balancing methods

Fig. 5 illustrates that our model's performance is invariant to the method used for balancing data of different classes.

4.3.6 Donor Propensity and Evaluation:

We employ LSTM and GRU models to make predictions for the current time window based on the data from the previous window. The binary classification models predict whether an individual will make a donation (Target=1) or not (Target=0), whereas multi-class classification models predict the giving class of that individual, and regression models predict the amount that individual is likely to donate. For classification tasks, we utilize precision, recall, and F1-score as evaluation metrics, while for regression we use root-mean-squared-error (RMSE) and median-absolute-error (MAE) to compare the predicted donation amounts with the actual donation amounts.

5 Analysis

In this section, we present key findings from our study that provide insights for assessing fundraising based on event participation and structural properties of the co-attendance network. We gain insights about factors that may influence increased giving behavior. Furthermore, we leverage node-level features extracted from the co-attendance network to predict donor propensity for different donor segments. Further, we utilize a multi-year network to study the influence of network dynamics on donor behavior.

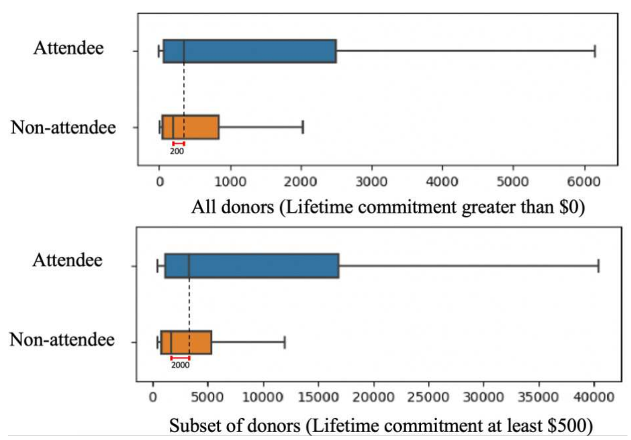

Fig. 6 Box plots for event attendees and nonattendees [31]: the upper subplot represents the distribution of donations > $0 and the lower subplot represents donations >= $500 in the second subplot.

5.1 Impact of event attendance

We first examine the impact of event attendance on the donor giving amount and compare it to the data for non-attendees. The box plot in Fig. 6 demonstrates that attendees have a higher giving average and median than non-attendees, particularly for larger donation amounts, as seen in the second subplot with donations of at least \$500. We confirm that our results are statistically significant by performing Mann-Whitney U test [18] on the bootstrapped samples.

Having established that participation has an impact on the effectiveness of fundraising, we identify similarities in donor characteristics in this network.

5.2 Evolving Assortativity at Network level

We search for similarities among event co-attendees by measuring the assortativity of the network, which indicates the preference for nodes to attach to others that have similar characteristics. We perform our tests on overlapped windows of 3 years, using the decay function described in [31] . In all cases, assortativity is consistently small but positive, indicating that individuals have a slight tendency to attend events with others of similar giving status.

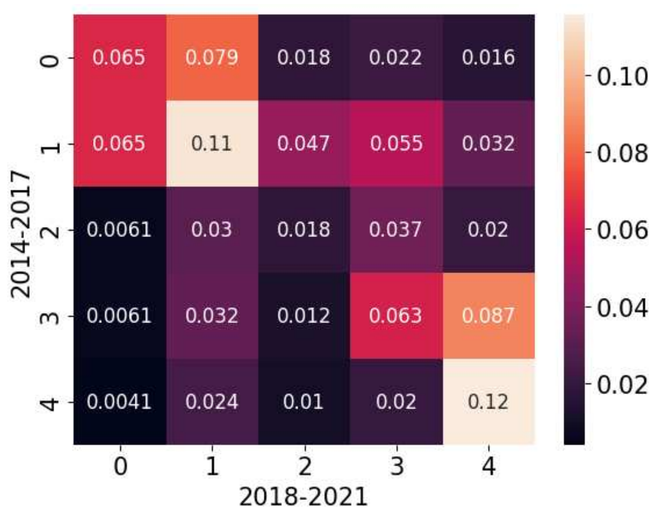

Fig. 7 Fraction of shift in donors between giving classes [31] listed in Section 5.2

5.3 Changes in behavior at node level with time

Next, we examine variations in donor behavior from the 2014-2017 period to the 2018-2021 period. We observe that more than half the donors retained between the two periods showed an increase in centrality and 76% of donors that had an increased centrality exhibited increased giving. Additionally, *more than half* of the donors (54%) that had a decreased or equal rank displayed a decrease in giving. Next, we examine the extent to which these donations amplify between the two periods. We label each node based on their donation band [$500, $1000], ($1000, $5000], ($5000, $10000], ($10000, $50000], ($50000, ...) discussed with the fundraising team [31].

From Fig. 7 we observe that 41% of people shifted from a lower class to a higher class. Moreover, 35% of those who shifted to a higher giving class or remained in the same class exhibited increased co-attendance with high-capacity donors. *None* of the donors that shifted to lower classes exhibited increased co-attendance.

5.4 Changes in group-level behavior with time

We apply the Louvain community detection algorithm [1] on the weighted co-attendance networks for a 5-year overlapping window. Fig. 8 shows a clique of communities that marks strong participation; we observe that community sizes are moderate to large. In the visualization, node size represents community size, green color intensity reflects average donation amounts (lighter shades indicate lower donations), while the edges are unweighted. We observe that community sizes increase over time, larger communities tend to have higher giving levels and the correlation becomes stronger over time [31]. Here community size represents the number of co-attendees who attend.

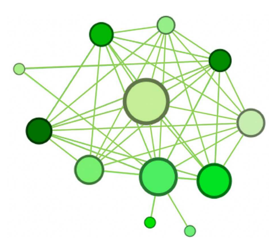

Fig. 8 Communities for a snapshot of 5 years (2017-2021) [31]; color intensity indicates average donation amounts, and circle size represents the community size (number of co-attendees).

Leveraging the structural attributes of the network, we develop our multivariate model and evaluate its performance on different donor segments and conduct a multi-year analysis. While binary classification yields poor results for both the baseline

Table 2 Equal size class segments for regular attendees

Donor Classes	Giving range	Class Label
Non donors	\$0	0
Low capacity donors	[\$10, \$145)	1
Medium capacity donors	[\$145, \$900)	2
High capacity donors	[\$900, \$45k]	3

and multivariate models, multi-class classification results in better performance for our model compared to the baseline. The regression results for multivariate model outperformed the baselines for both regular and irregular donors in terms of MAE. Although the test RMSE values were similar for the baseline and multivariate models, however, when trained on a filtered dataset excluding high donors, performance in terms of test RMSE was significantly better for the multivariate model than the baseline models.

In the multi-class prediction problem, we follow a separate class split for each case individually. This is necessary because the ratio of donors to non-donors differs between the sets, and a uniform approach of maintaining equal class sizes is not suitable. The irregular dataset consists of a higher proportion of smaller donors compared to very high donors, unlike the regular attendee segment. By adapting the class splits to the characteristics of each dataset, we can account for the variations in donor distribution and achieve more effective predictions.

5.5 Forecasting Donor Propensity: Yearly Predictions for Regular Attendees

Predicting yearly donations allows nonprofits to develop tailored fundraising strategies and identify the most effective approaches to engage donors and design campaigns that resonate with their giving patterns. Training the model on the segmented data gives accurate predictions.

We first predict the yearly donations for the regular attendees. These are donors who consistently attended events over the course of eight years. We further refine the dataset by removing outliers and ensuring that each class has an equal number of samples, thereby achieving balanced classes. As a result, the sample size decreases to $S = 456$.

The ROC curves in Fig. 9 for multi-class classification predict One-Vs-Rest classes, and demonstrate higher performance (greater AUC) for the LSTM model across all classes. The precision, recall, and F1-scores indicate that the classes are relatively separable for classes 0, 1, and 3 as seen in Table 3. The LSTM model outperforms the baseline model, with a significant difference in F1-score for class 0 and class 3. However, this difference is smaller for class 1 and 2. Notably, the precision and F1-score remain low for class 2 in both models. Meanwhile,the GRU

Table 3 Classification Report

Model	***Class***	***Precision***	***Recall***	***F1-score***
Regular attendee baseline	0	0.68	0.54	0.60
	1	0.54	0.62	0.58
	2	0.43	0.45	0.44
	3	0.71	0.81	0.76
Regular attendee LSTM	0	0.94	0.89	0.92
	1	0.75	0.46	0.57
	2	0.33	0.75	0.46
	3	0.83	0.88	0.86
Irregular attendee Baseline	0	0.78	0.46	0.58
	1	0.54	0.77	0.63
	2	0.56	0.77	0.65
Irregular attendee LSTM	0	**1.0**	**1.0**	**1.0**
	1	0.75	0.81	0.78
	2	0.74	0.67	0.70
Network snapshot Baseline	0	0.96	0.89	0.92
	1	0.60	0.78	0.68
	2	0.65	0.85	0.74
Network snapshot LSTM	0	**1.0**	**1.0**	**1.0**
	1	0.92	0.93	0.92
	2	0.73	0.71	0.72

Table 4 Regression Results

Type	***Model***	***Train RMSE***	***Test RMSE***	***Train MAE***	***Test MAE***
Regular event attendee	MA	3560.24	9221.99	200.0	78.12
	ARIMA	**3255.10**	7979.89	1053.31	280.46
	ES	3412.78	8438.86	365.39	280.46
	LSTM	6937.15	7873.73	209.59	211.05
	GRU	6922.49	**7858.51**	**111.79**	**90.03**
Irregular event attendee	MA	184480.33	33462.85	925.0	100.0
	ARIMA	**108452.34**	69808.79	20488.89	20994.67
	ES	118368.10	58727.36	16089.77	42077.67
	LSTM	110686.14	35487.12	**222.10**	**222.26**
	GRU	109775.15	**24239.65**	433.69	94.26
All attendee (Multi-year)	ARIMA	**73111.46**	150637.36	**0**	**0**
	ES	84824.18	145089.20	**0**	**0**
	LSTM	88213.66	103038.07	**0.20**	**0.16**
	GRU	88198.64	**103027.14**	**0.87**	**0.79**
Filtered attendee (Multi-year)	ARIMA	**9652.80**	22428.29	**0**	**0**
	ES	10857.28	20687.33	**0**	**0**
	LSTM	18891.85	16874.49	2.09	1.63
	GRU	18867.08	**16857.38**	4.12	4.29

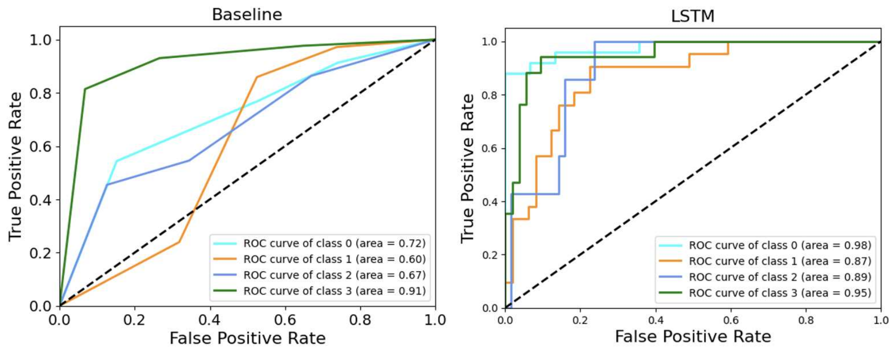

Fig. 9 Comparison of AUC for the Receiver Operating Characteristic for LSTM vs Baseline 3 for Regular event attendee

multivariate model performs better than the others in the cases with lowest test RMSE and MAE values, as seen in Table 4.

5.6 Forecasting Donor Propensity: Yearly Predictions for Small-Sized Event Attendees

We employ the same approach as Section 5.5 on a different segment that includes individuals who do not attend events consistently (irregular attendees). We focus on filtering donor attendees of small-sized events, as these events are indicative of their co-attendance behavior. Based on discussions with the fundraising team, events with 20 or fewer attendees are classified as small-sized events. Class 0 represents non-donors who attend events, Class 1 represents donors who donate between \$0 and \$1000, and Class 2 represents donors who donate more than \$1000. These particular bands are chosen because irregular donors typically contribute lower amounts, often influenced by spontaneous decisions or witnessing fellow donors' contributions. Given the significantly larger number of samples in classes 0 and 1, and the relatively smaller number in class 2, further dividing them would result in fewer samples for the model to learn from, inhibiting the model's ability to capture patterns and make accurate predictions. The LSTM model outperforms the baseline model because its network properties effectively handle the temporal irregularities present in the data.

When considering all donors ($S = 5578$) including those with more irregular attendance patterns, the baseline model's precision and F1-score decrease to 0.20 and 0.31 for class 1, and 0.22 and 0.33 for class 2, respectively. In contrast, the LSTM model significantly outperforms the baseline, achieving precision and F1-scores of 0.75 and 0.76 for class 1, and 0.74 and 0.72 for class 2. However, when we focus on a subset of donors who have attended events consecutively for over two years ($S = 589$), we observe an improvement in the baseline model's performance.

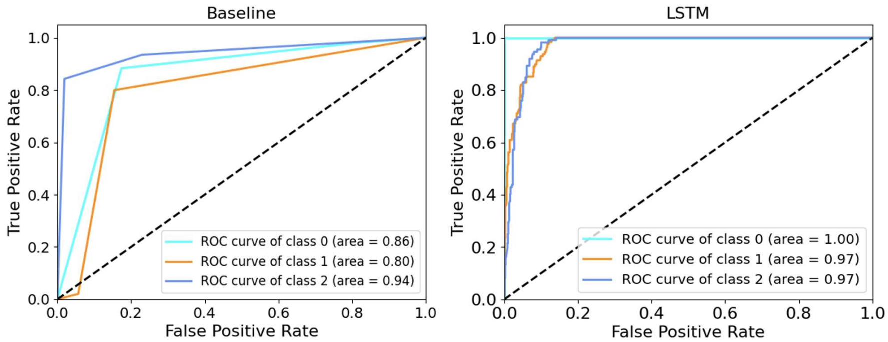

Fig. 10 Comparison of AUC for the Receiver Operating Characteristic for LSTM vs Baseline for irregular event attendees

Nevertheless, the LSTM model still outperforms the baseline for all classes as seen in Fig. 10 and Table 3. With respect to the regression task, LSTM multivariate model outperformed all the other models with lowest train and test MAE values (222.10 and 222.26).

5.7 Multi-Year Analysis: Learning from Network Snapshots for Long-Term Trends

Rolling window network characteristics provide insights into the evolving relationships and interactions among donors and nonprofits can identify long-term trends and seasonal trends in donor behavior. We observe improved model performance when trained with weighted network snapshots and an increased sample size (S = 99431). The LSTM model achieves complete separability for the classes, as illustrated in Fig. 11, outperforming the baseline model.

From Table 4 we observe lowest test RMSE for the multivariate GRU model, and the MAE values are fairly close to 0. We observe improved performance for the regression task, and the LSTM and GRU models outperform the statistical models in terms of test RMSE, despite having comparable training losses. This pattern persists when the donors are filtered by removing extremely high donors (those who donate more than a million dollars) with test RMSE significantly reduced.

6 Discussion

The data we examine suggests that the relevant fund-raising team successfully organized engaging events, as evidenced by the growth in donor participation over time (Fig. 1). We observe an upward trend in annual donations (Fig. 2), suggesting that the

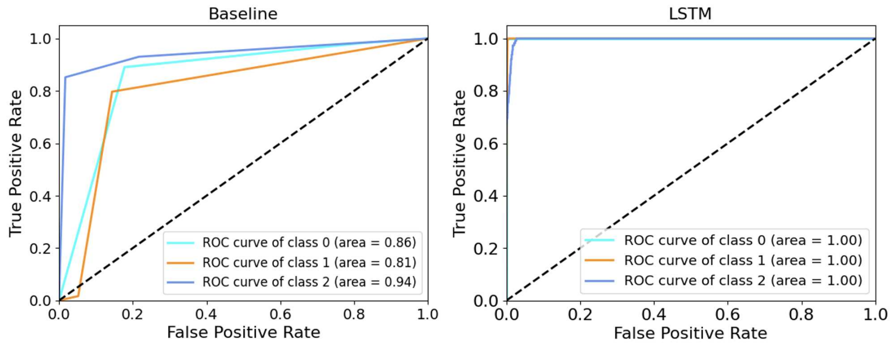

Fig. 11 Comparison of AUC for the Receiver Operating Characteristic for LSTM vs Baseline for Network snapshots for a rolling window of 3 years.

team effectively utilized their professional networks and attracted more high-giving donors, since our analysis indicates a statistically significant relationship between event attendance and donor giving.

As we seek to identify structural characteristics that correlate with fund-raising outcomes, we draw upon prior research indicating that donor-donor dyads can influence donor behavior. We identify similarities in individual characteristics for our dataset, observing positive assortativity of co-attendees with respect to their giving values. A low-to-moderate assortativity in co-attendance is beneficial for fund-raising networks and suggests the success of event organization and attendee placement. In contrast, a high assortativity value would hinder new connections and diminish the positive peer pressure effect. This is backed by studies that show individuals being influenced to match or exceed the generosity of their fellow attendees, resulting in higher overall donation amounts [30, 20].

It is important to assess how donor characteristics change with time. We find that 76% of donors that had increased centrality also showed increase in giving amounts, and more than half of the donors that displayed an equal or lesser centrality decreased their giving. Increased centrality indicates that donors are actively co-attending more events or are co-attending with other high-centrality donors. Fig. 7 shows the shift in giving class (of a donor) and its relationship with co-attendance. Overall, 41.37% donors increased their giving class, 37.53% maintained the same class, and the remaining decreased their giving level. Moreover, we observe that 35% that remained in the same class or shifted to an upper giving class showed an increased co-attendance with high-giving donors, while none of the donors that shifted to a lower class exhibited increased co-attendance.

In addition to individual characteristics associated with increased giving, we also analyze group-level characteristics that may contribute to this change. Fig. 8 provides a structural visualization of communities of co-attendees, showing a clique of communities that mark active participation of donors. The donor-attendee community analyses enable organizations to examine donor engagement and identify major donor prospects. This information can be used to strengthen relationships and maximize event impact. We further observe:

- Increased community sizes over the years;
- Larger and more tightly-knit participation of donors over time; and
- Increase (with time) in the correlation between fraction of above-median donors and community size.

Our results are partially explained by the fact that larger communities have many donors that benefit from resources such as business opportunities, and professional networks. Individuals who have gained from these resources may feel a sense of gratitude and a desire to give back to the community that provided them with such advantages [15].

Having gained insights into the network characteristics associated with fundraising outcomes, we construct a multivariate model by extracting features from the network. These features include assortativity (similarity) at the individual level, exposure to high-capacity donors, exposure to donors who can potentially influence non-donors, the structural positions of individuals in the network, and their levels of potential interactions with their local neighborhoods.

Our multivariate model outperforms baselines for both classification and regression tasks across various donor segments. Notably, the multivariate models' precision, recall, and F1-score are significantly better than the baselines when tested on irregular event attendees, and using multi-year analysis, as seen from Table 3. Figures 9, 10, and 11 show that the AUC for our model is consistently higher than the baseline for all three analysis. Moreover, the node features become perfectly separable when considering network snapshots for a window size of 3 years, as indicated by an AUC of 1.0 for all three classes.

When predicting the donation amount for regular donors, we find that the multivariate model performs better than the baselines, as evidenced by test RMSE and, train and test MAE(median absolute error) for the two donor segments in Table. 4. Significant differences in test RMSE and MAE values are observed for the irregular attendee segment and the multi-year network.

The enhanced performance of the multivariate model can be attributed to the following factors. First, statistical baselines encounter difficulties in handling missing data, which is prevalent in cases where users have donated in the past few years and may donate in the future. Second, there is a scarcity of known data about new donors, making it challenging for statistical models to draw accurate inferences. Third, donors who attend events regularly tend to be more loyal to the recipient organization. Their donation behavior may be influenced by external factors such as donor-solicitor and donor-recipient dyads, including communication with the fundraising team, beyond just donor-donor dyads. While these donors contribute annually, their giving potential may primarily reflect their long-term commitments.

Our analysis predicting donor propensity can be highly valuable for fund-raisers, as it enables them to identify potential donors and donor segments well in advance. By understanding the regularity in donation patterns and optimizing resource allocation to various donor segments, such as retained donors, new donors, and donors who have changed their behavior, fund-raisers can improve targeting and ultimately improve fund-raising outcomes.

One of the limitations of our work is the inability to model causal relationships linking network properties to giving outcomes. With respect to predicting donor propensity, our dataset appears to be insufficient for accurate binary classification as both the baseline model and the multivariate model failed to effectively identify the non-donors. The best performance was achieved using cross-validation, suggesting that the models do not necessarily benefit by learning from extensive historical data.

7 Conclusion

This study highlights the potential of utilizing donor co-attendance networks for understanding fund-raising outcomes and how network features can be harnessed for donor propensity prediction. Our analyses show that people who attended more events donated more often and greater amounts. We identified that there is a low-to-moderate similarity in the giving characteristics of co-attendees. We also observe that donors who occupied central positions had increased giving as well as increased co-attendance with high-capacity donors over the periods increased their giving class membership. The community structure also revealed strong connections that are beneficial for maximizing the effectiveness of their resource allocation strategies. By incorporating these network features into our multivariate model, we achieve accurate predictions of the giving class and donation amount, outperforming baseline models for all three analyses.

This study contributes to the body of knowledge on assessing fundraising outcomes by utilizing the dynamics of the co-attendance network. Our research demonstrates the enhanced predictive capability of incorporating event co-attendance networks in predicting donor propensity. This can be used by fundraisers in assessing outcomes, making informed decisions and develop tailored strategies to enhance fundraising outcomes.

Our results have relied on limited information readily available from the donor history database. With access to more data, we can consider interactions influenced by co-location, organizational co-affiliation, and other factors. Future work can also attempt to distinguish between correlation and causation in the relationships between network properties and fundraising outcomes. Finer-grained analysis is also needed to explain the divergence of behaviors of subsets of individuals that are initially similar, identifying the extent to which co-attendance may explain the divergence.

Acknowledgement

Srikanta and Soundarajan were funded by NSF 2047224.

References

1. Thomas Aynaud. Community detection for networkx's documentation, 2018.
2. Victor MH Borden, Genevieve G Shaker, and Brittany L Kienker. The impact of alumni status on institutional giving by faculty and staff. *Research in Higher Education*, 55:196–217, 2014.
3. Kanad Chakraborty, Kishan Mehrotra, Chilukuri K Mohan, and Sanjay Ranka. Forecasting the behavior of multivariate time series using neural networks. *Neural networks*, 5(6):961–970, 1992.
4. Cassandra M Chapman, Barbara M Masser, and Winnifred R Louis. The champion effect in peer-to-peer giving: Successful campaigns highlight fundraisers more than causes. *Nonprofit and Voluntary Sector Quarterly*, 48(3):572–592, 2019.
5. Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
6. James Samuel Coleman et al. Introduction to mathematical sociology. *Introduction to mathematical sociology.*, 1964.
7. Rod Duclos and Alixandra Barasch. Prosocial behavior in intergroup relations: How donor self-construal and recipient group-membership shape generosity. *Journal of Consumer Research*, 41(1):93–108, 2014.
8. Harmonie Farrow and Y Connie Yuan. Building stronger ties with alumni through facebook to increase volunteerism and charitable giving. *Journal of Computer-Mediated Communication*, 16(3):445–464, 2011.
9. Soledad Galli and Kishan Manani. Feature engineering for time series forecasting.
10. Everette S Gardner Jr. Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28, 1985.
11. Elizabeth M Gerber, Julie S Hui, Pei-Yi Kuo, et al. Crowdfunding: Why people are motivated to post and fund projects on crowdfunding platforms. In *Proceedings of the international workshop on design, influence, and social technologies: techniques, impacts and ethics*, volume 2, page 10, 2012.
12. Lihuan Guo, Wei Wang, Yenchun Jim Wu, and Mark Goh. How much do social connections matter in fundraising outcomes? *Financial Innovation*, 7(1):79, 2021.
13. Aric Hagberg and Drew Conway. Networkx: Network analysis with python. *URL: https://networkx. github. io*, 2020.
14. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
15. Nicola Lacetera, Mario Macis, and Angelo Mele. Viral altruism? charitable giving and social contagion in online networks. *Sociological Science*, 3, 2016.
16. Ming Li, K Mehrotra, Chilukuri Mohan, and Sanjay Ranka. Sunspot numbers forecasting using neural networks. In *Proceedings. 5th IEEE International Symposium on Intelligent Control 1990*, pages 524–529. IEEE, 1990.
17. Bryan Lim and Stefan Zohren. Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A*, 379(2194):20200209, 2021.
18. Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947.
19. Kevin R McClure, Leah Frierson, Adam W Hall, and Kara L Ostlund. Philanthropic giving by foundations to higher education institutions: A state-level social network analysis. *Philanthropy & Education*, 1(1):1–28, 2017.
20. Jonathan Meer. Brother, can you spare a dime? peer pressure in charitable solicitation. *Journal of public economics*, 95(7-8):926–941, 2011.
21. Kishan Mehrotra, Chilukuri K Mohan, and Sanjay Ranka. *Elements of artificial neural networks*. MIT press, 1997.
22. Brian K Nelson. Time series analysis using autoregressive integrated moving average (arima) models. *Academic emergency medicine*, 5(7):739–744, 1998.

23. Albert A Okunade and Robert L Berl. Determinants of charitable giving of business school alumni. *Research in higher education*, 38:201–214, 1997.
24. Anton Pudikov and Alexander Brovko. Comparison of lstm and gru recurrent neural network architectures. In *International Scientific and Practical Conference in Control Engineering and Decision Making*, pages 114–124. Springer, 2020.
25. Peter H Reingen. Test of a list procedure for inducing compliance with a request to donate money. *Journal of Applied Psychology*, 67(1):110, 1982.
26. Adrian Sargeant. Using donor lifetime value to inform fundraising strategy. *Nonprofit Management and Leadership*, 12(1):25–38, 2001.
27. Gregory D Saxton and Lili Wang. The social network effect: The determinants of giving through social media. *Nonprofit and voluntary sector quarterly*, 43(5):850–868, 2014.
28. George B Schreiber, UK Sharma, DJ Wright, SA Glynn, HE Ownby, Y Tu, G Garratty, J Piliavin, T Zuck, R Gilcher, et al. First year donation patterns predict long-term commitment for first-time donors. *Vox sanguinis*, 88(2):114–121, 2005.
29. Jen Shang, Rachel Croson, and Americus Reed. "i" give, but "we" give more: the impact of identity and the mere social information effect on donation behavior. *ACR North American Advances*, 2007.
30. Sarah Smith, Frank Windmeijer, and Edmund Wright. Peer effects in charitable giving: Evidence from the (running) field. *The Economic Journal*, 125(585):1053–1071, 2015.
31. Shwetha Koushik Manchinahalli Srikanta, Katie L Pierce, Joshua E Introne, Chilukuri K Mohan, and Sucheta Soundarajan. Structure and dynamics of a charitable donor co-attendance network. In *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining*, pages 514–521, 2023.
32. Yuan Tian and Sara Konrath. The effects of similarity on charitable giving in donor–donor dyads: A systematic literature review. *Voluntas: International Journal of Voluntary and Nonprofit Organizations*, 32:316–339, 2021.
33. Tijs A van den Broek, Ariana Need, Michel L Ehrenhard, Anna Priante, and Djoerd Hiemstra. The influence of network structure and prosocial cultural norms on charitable giving: A multilevel analysis of movember's fundraising campaigns in 24 countries. *Social networks*, 58:128–135, 2019.
34. Shengsheng Xiao and Qing Yue. The role you play, the life you have: Donor retention in online charitable crowdfunding platform. *Decision Support Systems*, 140:113427, 2021.