

Spotting Collective Behaviour of Online Frauds in Customer Reviews

Sarthika Dhawan¹, Siva Charan Reddy Gangireddy¹,
Shiv Kumar² and Tanmoy Chakraborty^{1*}

¹Indraprastha Institute of Information Technology Delhi (IIITD), India

²Netaji Subhas University of Technology (NSUT), Delhi, India

{sarthika15170, sivag}@iiitd.ac.in, shivk.it.16@nsit.net.in, tanmoy@iiitd.ac.in

Abstract

Online reviews play a crucial role in deciding the quality before purchasing any product. Unfortunately, spammers often take advantage of online review forums by writing fraud reviews to promote/demote certain products. It may turn out to be more detrimental when such spammers collude and collectively inject spam reviews as they can take complete control of users' sentiment due to the volume of fraud reviews they inject. Group spam detection is thus more challenging than individual-level fraud detection due to unclear definition of a group, variation of inter-group dynamics, scarcity of labeled group-level spam data, etc. Here, we propose DeFrauder, an unsupervised method to detect online fraud reviewer groups. It first detects candidate fraud groups by leveraging the underlying product review graph and incorporating several behavioral signals which model multi-faceted collaboration among reviewers. It then maps reviewers into an embedding space and assigns a spam score to each group such that groups comprising spammers with highly similar behavioral traits achieve high spam score. While comparing with five baselines on four real-world datasets (two of them were curated by us), DeFrauder shows superior performance by outperforming the best baseline with 17.11% higher NDCG@50 (on average) across datasets.

1 Introduction

Nowadays, online reviews are becoming highly important for customers to take any purchase-related decisions. Driven by the immense financial profits from product reviews, several blackmarket syndicates facilitate to post deceptive reviews to promote/demote certain products. Groups of such fraud reviewers are hired to take complete control of the sentiment about products. Collective behaviour is therefore more subtle than individual behaviour. At individual level, activities might be normal; however, at the group level they might be substantively different from the normal behavior. Moreover,

it is not possible to understand the actual dynamics of a group by aggregating the behaviour of its members due to the complicated, multi-faceted, and evolving nature of inter-personal dynamics. Spammers in such collusive groups also adopt intelligent strategies (such as paraphrasing reviews of each other, employing a subset of their resources to one product, etc.) to evade detection.

Previous studies mostly focused on individual-level fraud detection [Lim *et al.*2010, Fei *et al.*2013, Akoglu *et al.*2013]. Few other studies which realized the detrimental effect of such collective activities detected groups simply based on Frequent Itemset Mining (FIM) [Mukherjee *et al.*2012, Xu *et al.*2013, Allahbakhsh *et al.*2013]. They thus focused more on ranking fraud groups, paying less attention to judge the quality of the detected groups.

[Wang *et al.*2016] pointed out several limitations of FIM for group detection – high computational complexity at low minimum support, absence of temporal information, unable to capture overlapping groups, prone to detect small and tighter groups, etc.

In this paper, we propose DeFrauder¹, a novel architecture for fraud reviewer group detection. DeFrauder contributes equally to (i) the detection of potential fraud groups by incorporating several coherent behavioral signals of reviewers, and (ii) the ranking of groups based on their degree of spamicity by proposing a novel ranking strategy. Experiments on four real-world labeled datasets (two of them were prepared by us) show that DeFrauder significantly outperforms five baselines – it beats the best baseline by 11.35% higher accuracy for detecting groups, and 17.11% higher NDCG@50 for ranking groups (averaged over all datasets).

In short, our contributions are fourfold: (i) two novel datasets, (ii) novel method for reviewer group detection, (iii) novel method for ranking groups, and (iv) a comprehensive evaluation to show the superiority of DeFrauder.

2 Related Work

Due to the abundance of literature [Kou *et al.*2004] on online fraud detection, we restrict our discussion to *fraud user detection on product review sites*, which we deem as pertinent to this paper.

¹DeFrauder: **D**etecting **F**raud Reviewer Groups, Code is available in [Dhawan *et al.*2019].

*Contact Author



Figure 1: (a) Workflow of DeFrauder. (b) Coherent review patterns of a fraud reviewer group on two products. To evade detection, four reviewers in a group paraphrase reviews of each other keeping the underlying sentiment same.

In user-level fraud detection, notable studies include [Lim *et al.*2010] which proposed scoring methods to measure the degree of spamicity of a reviewer based on rating behaviors; [Wang *et al.*2011] which developed a graph-based fraud reviewer detection model; [Fei *et al.*2013] which exploited burstiness in reviews to detect fake reviewers. SpEagle [Akoglu *et al.*2013] utilizes clues from all metadata as well as relational data and harnesses them collectively under a unified framework. ASM [Mukherjee *et al.*2013a] facilitates modeling spamicity as latent factor and allows to exploit various behavioral footprints of reviewers. [Wang *et al.*2012] argued that the dynamic behavioral signals can be captured through a heterogeneous reviewer graph by considering reviewers, reviews and products together. FraudEagle [Akoglu *et al.*2013] spots fraudsters and fake reviews simultaneously. [Dutta *et al.*2018, Chetan *et al.*2019] discussed about collusion in social media.

Few studies attempted to detect group-level fraud reviewer groups. GSRank [Mukherjee *et al.*2012] was the first method of this kind, which identifies spam reviewer groups using FIM, and ranks groups based on the relationship among groups, products, and individual reviewers. [Mukherjee *et al.*2012] argued that due to the scarcity of data, unsupervised method should be used to tackle this problem. Other studies largely focused on improving the ranking algorithm, ignoring the performance of the group detection method (they also used FIM for detecting groups) [Xu *et al.*2013, Allahbakhsh *et al.*2013, Mukherjee *et al.*2013b, Ott *et al.*2011]. [Xu and Zhang2015] proposed an Expectation Maximization algorithm to compute the collusive score of each group detected using FIM. [Wang *et al.*2016] argued that FIM tends to detect small-size and tighter groups. They proposed GSBP, a divide-and-conquer algorithm which emphasizes on the topological structure of the reviewer graph. GGSpam [Wang *et al.*2018b] directly generates spammer groups consisting of group members, target products, and spam reviews.

We propose DeFrauder that leverages temporal pattern, content similarity and underlying graph structure to detect collusive fraud groups in review forums (see Table 1 for a comparison). Unlike others, DeFrauder contributes equally to both the components – group detection and ranking.

	GSBP	GGSpam	GSRank	DeFrauder
Temporal	✓	✓	✓	✓
Content			✓	✓
Graph	✓	✓		✓

Table 1: DeFrauder captures all three signals obtained from users attributes and underlying graph structure.

3 Proposed DeFrauder Framework

In this section, we explain DeFrauder which is a three-stage algorithm – detecting candidate fraud groups, measuring different fraud indicators of candidate groups, and ranking groups based on the group spam score.

3.1 Group Fraud Indicators

Here we present six fraud indicators which measure the spamicity of a group [Mukherjee *et al.*2012, Wang *et al.*2018b, Wang *et al.*2016] by taking into account linguistic, behavioral, structural and temporal signals. All indicators are normalized to $[0, 1]$ – larger value indicates more spamming activity. Let R and P be the entire set of reviewers and products. Let $R(g)$ be the set of reviewers in a group g , and $P(g)$ be the set of target products reviewed by the reviewers in $R(g)$. Each reviewer i reviewed a set of products P_i .

To reduce the contingency of small groups, we use a penalty function $L(g)$ which is a logistic function [Wang *et al.*2016]: $L(g) = 1/(1 + e^{-(|R(g)|+|P(g)|-3)})$. We subtract 3 from the sum since we consider minimum number of reviewers and products within a group to be 2 and 1, respectively; therefore $L(g) \in [0.5, 1]$ [Wang *et al.*2018a].

Review tightness(RT). It is the ratio of the number of reviews in g (assuming each reviewer is allowed to write one review to each product) to the product of the size of reviewer set and product set in g .

$$RT(g) = \frac{\sum_{i \in R(g)} |P_i|}{|R(g)||P(g)|} \cdot L(g) \quad (1)$$

If significant proportion of people co-reviewed several products, then it may imply a spam group activity.

Neighbor tightness(NT). It is defined as the average value of the Jaccard similarity (JS) of the product sets of each reviewer pair.

$$NT(g) = \frac{\sum_{i,j \in R(g)} JS(P_i, P_j)}{\binom{|R(g)|}{2}} \cdot L(g) \quad (2)$$

If the product sets are highly similar, both the reviewers are then highly likely to be together in a collusive group.

Product tightness(PT). It is the ratio of the number of common products to the number of products reviewed by all the members in g [Wang *et al.*2016].

$$PT(g) = \frac{|\bigcap_{i \in R(g)} P_i|}{|\bigcup_{i \in R(g)} P_i|} L(g) \quad (3)$$

Group members reviewing certain number of products and not reviewing any other products are more likely to indulge in fraud activities.

Rating variance(RV). Group spammers tend to give similar ratings while reviewing any product. Let $S^2(p, g)$ be the variance of the rating scores of product p by reviewers in g . We take the average variance for all target products. The variance degree $[0.5, 1)$ is converted into spam score between $(0, 1]$.

$$RV(g) = 2L(g) \left(1 - \frac{1}{1 + e^{-\frac{\sum_{p \in P(g)} S^2(p, g)}{|P(g)|}}} \right) \quad (4)$$

Product reviewer ratio(RR). It is defined as the maximum value of the ratio of the number of reviewers in $R(g)$ who reviewed product $p \in P(g)$ to the number of all the reviewers of p , denoted by $|Rev(p)|$:

$$RR(g) = \max_{p \in P(g)} \frac{\sum_{i \in R(g)} \{1 : p \in P_i\}}{|Rev(p)|} \quad (5)$$

If a product is mainly reviewed by the reviewers in g , then the group is highly likely to be a spammer group.

Time window(TW). Fraudsters in a group are likely to post fake reviews during a short-time interval. Given a group g , and a product $p \in P_g$, we define the time-window based spamicity as

$$TW(g, p) = \begin{cases} 1 - \frac{SD_p}{T}, & \text{if } SD_p \leq T \\ 0, & SD_p > T \end{cases}$$

$$TW(g) = \frac{\sum_{p \in P_g} TW(g, p)}{|P(g)|} \cdot L(g), \quad (6)$$

where SD_p is the standard deviation of review time for a product p reviewed by reviewers in group g . T is a time threshold (set to 30 days in our experiments as suggested in [Wang *et al.*2018b]).

Algorithm 1 ExtractGroups

```

1: Initialize:
2:  $CCgroups \leftarrow \text{set}()$   $\triangleright$  Set of candidate groups
3:  $G(V, E) \leftarrow$  Edge attributed graph
4:  $CSet \leftarrow \{\}$   $\triangleright$  Potential merged groups
5:  $G, CCgroups \leftarrow \text{GroupDetector}(G, CCgroups)$ 
6: Iterate:
7:  $G'(V', E') \leftarrow$  Attributed line graph of  $G$ 
8:  $G', CCgroups = \text{GroupDetector}(G', CCgroups)$ 
9:  $G \leftarrow G'$ 
10: until  $|E| > 1$ 
11: return  $CCgroups$ 
12: function GROUPDETECTOR( $G, CCgroups$ )
13:   for each isolated node  $v_i$  in  $G$  do
14:      $CCgroups.add(v_i)$  and remove  $v_i$  from  $G$ 
15:   for each pair  $(e_i, e_j) \in E \times E$  do
16:     if  $a_i^e \subset a_j^e$  then
17:       if  $\frac{\bigcap_{m \in a_j^e} P_m}{\bigcup_{m \in a_j^e} P_m} > 0.5$  then
18:          $CSet(a_i^e) = CSet(a_i^e) \cup a_j^e$ 
19:       else
20:         if  $\frac{\bigcap_{m \in \{a_j^e \setminus a_i^e\}} P_m}{\bigcup_{m \in \{a_j^e \setminus a_i^e\}} P_m} > 0.5$  then
21:            $CCgroups.add(a_j^e \setminus a_i^e)$ 
22:       for each  $e_i \in E$  do
23:          $k = CSet(a_i^e)$ 
24:          $CCgroups.add(k)$  and remove  $k$  from  $G$ 
25:       for each connected component  $c$  with  $|c| > 2$  do
26:          $CCgroups.add(c)$  and remove  $c$  from  $G$ 
27:       for each group  $g \in CCgroups$  do
28:         if  $CollectiveScore(g) \leq \tau_{spam}$  then
29:            $CCgroups.remove(g)$ 
30:       return  $G, CCgroups$ 
31: end function
    
```

3.2 Detection of Candidate Fraud Groups

We propose a novel graph-based candidate group detection method based on the ‘‘coherence principle’’.

Hypothesis 1 (Coherence Principle) *Fraudsters in a group are coherent in terms of – (a) the products they review, (b) ratings they give to the products, and (c) the time of reviewing the products.*

We show that each component of this hypothesis is statistically significant (see Sec. 5). We incorporate these three factors into a novel **attributed product-product graph** construction – $G(V, E)$ such that each $v_{ij} \in V$ indicates a product-rating pair (p_i, r_j) and its attribute a_{ij}^v consists of the set of reviewers who rated p_i with r_j . An edge $e_{(ij, mn)} = \langle u_{ij}, v_{mn} \rangle \in E$ indicates the co-reviewing and co-rating patterns of two products p_i and p_m with rating r_j and r_n , respectively. The edge attribute $a_{(ij, mn)}^e$ indicates the set of co-reviewers $R_{(ij, mn)}$ who reviewed both p_i and p_m and gave same ratings r_j and r_n respectively within the same time τ_t (defined in Sec. 3.3). Note that edge $e_{(ij, in)}$ connecting same product with different ratings wont exist in G as we assume

that a reviewer is not allowed to give multiple reviews/ratings to a single product.

We then propose `ExtractGroups` (pseudocode in Algo 1, a toy example in Supplementary [Dhawan *et al.*2019]), a novel group detection algorithm that takes G as an input and executes a series of operations through `GroupDetector()` (Line 12) – isolated nodes are first removed (Lines 13-14), edge attributes are then merged and removed if Jaccard similarity (JS) of product sets that the corresponding reviewers reviewed is greater than a threshold (set as 0.5). Any group of reviewers eliminated due to the consideration of only common reviewers during the creation of edges is also checked through JS in order to avoid losing any potential candidate groups (Lines 15-24). Before proceeding to the next iteration, connected components containing more than two nodes are removed (Lines 25-26). We define *CollectiveScore* as the average of six group level indicators defined in Sec. 3.1, and consider those as potential groups whose *CollectiveScore* exceeds τ_{spam} (Lines 27-29, defined in Sec. 5.2).

It then converts the remaining structure of G into an attributed line graph (edges converted into vertices and vice versa) $G'(V', E')$ as follows: $v'_{(ij, mn)} \in V'$ corresponds to $e_{(ij, mn)}$ in G and $a'_{(ij, mn)} = a^e_{(ij, mn)}$; an edge $e'_{(ij, mn, ab)} = \langle v'_{(ij, mn)}, v'_{(ij, ab)} \rangle$ represents co-reviewing and co-rating patterns of products p_i, p_m and p_a ; the corresponding edge attribute is $a_{(ij, mn, ab)} = a^e_{(ij, mn)} \cap a^e_{(ij, ab)}$. G' is again fed into `GroupDetector` in the next iteration. Essentially, in each iteration, we keep clubbing reviewers together, who exhibit coherent reviewing patterns. The iteration continues until none of the edges remain in the resultant graph, and a set of candidate fraud groups are returned.

The worst case **time complexity** of `ExtractGroups` is $\mathcal{O}(k|E|^2)$, where k is the number of iterations, and $|E|$ is the number of edges in G (see Supplementary [Dhawan *et al.*2019] for more details).

Theorem 3.1 (Theorem of convergence)

`ExtractGroups` will converge within a finite number of iterations.

See Supplementary [Dhawan *et al.*2019] for the proof.

3.3 Ranking of Candidate Fraud Groups

Once candidate fraud groups are detected, `DeFrauder` ranks these groups based on their spamicity. It involves two steps – mapping reviewers into an embedding space based on their co-reviewing patterns, and ranking groups based on how close the constituent reviewers are in the embedding space.

Reviewer2Vec: Embedding Reviewers

Our proposed embedding method, `Reviewer2Vec` is motivated by [Wang *et al.*2018b]. Given two reviewers i and j co-reviewing a product $p \in P$ by writing reviews c_p^i and c_p^j with the rating r_p^i and r_p^j at time t_p^i and t_p^j respectively, we define the *collusive spamicity* of i, j w.r.t. p as:

$$Coll(i, j, p) = \begin{cases} 0, & |t_p^i - t_p^j| > \tau_t \vee |r_p^i - r_p^j| \geq \tau_r \\ \zeta(i, j, p), & \text{otherwise} \end{cases} \quad (7)$$

where,

$$\zeta(i, j, p) = s^p \left[\alpha \left(1 - \frac{|t_p^i - t_p^j|}{\tau_t} \right) + \beta \left(1 - \frac{|r_p^i - r_p^j|}{\tau_r} \right) + \gamma \cdot Cosine(c_p^i, c_p^j) \right]$$

Where $s^p = \frac{2}{1 + e^{-\frac{2}{q \in P} Rev(q) - Rev(p)} \theta + 2^\theta}} - 1$. Coefficients α, β and γ control the importance of time, rating and the similarity of review content, respectively ($0 \leq \alpha, \beta, \gamma \leq 1$ and $\alpha + \beta + \gamma = 1$). We set $\gamma > \alpha, \beta$ as same review content signifies more collusion as compared to the coherence in ratings and time [Wang *et al.*2018b].

s^p is the degree of suspicion of product p . τ_t, τ_r and θ are the parameters of the model. If the posting time difference among reviewers i and j on p is beyond τ_t or their rating on p deviates beyond τ_r (where $\tau_r = (\max - \min)x\%$, where \max and \min are the maximum and minimum ratings that a product can achieve respectively), we do not consider this co-reviewing pattern. `ExtractGroups` achieves best results with $\tau_t = 20$ and $\tau_r = (\max - \min)20\%$ (see Sec. 5.3).

$\zeta(i, j, p)$ is the collusiveness between two reviewers w.r.t. product p they co-reviewed; however there might be other reviewers who reviewed p as well. Lesser the number of reviewers of p , more is the probability that i and j colluded. This factor is handled by s^p after passing it through a sigmoid function. θ is a normalizing factor which ranges between $[0, 1]$ (set as 0.4 [Wang *et al.*2018b]). We take the cosine similarity of two reviews c_p^i and c_p^j after mapping them into embedding space using `Word2Vec` [Mikolov *et al.*2013].

Combining the collusive spamicity of a pair of reviewers across all the products they co-reviewed, we obtain the overall collusive spamicity between two reviewers:

$$\Phi(i, j) = \frac{2}{1 + e^{-\sigma(i, j)}} - 1$$

where $\sigma(i, j) = \left[\sum_{p \in \{P_i \cap P_j\}} Coll(i, j, p) \right] \frac{|P_i \cap P_j|}{|P_i \cup P_j|}$

We then create a reviewer-reviewer spammer graph [Wang *et al.*2018b] which is a bi-connected and weighted graph $G_s = (V_s, E_s)$, where V_s corresponds to the set of reviewers R , and two reviewers i and j are connected by an edge $e_{ij}^s \in E_s$ with the weight $W_{ij}^s = \Phi(i, j)$. Once G_s is created, we use state-of-the-art node embedding method to generate node (reviewer) embeddings (see Sec. 5.3).

Ranking Groups

For each detected group, we calculate the density of the group based on (Euclidean) distance of each reviewer in the group with the group's centroid in the embedding space. An average of all distances is taken as a measure of spamicity of the group. Let \vec{i} be the vector representation of reviewer i in the embedding space. The group spam score $Spam(g)$ of group g is measured as:

$$Spam(g) = \frac{1}{|R(g)|} \sum_{i \in R(g)} \left[\vec{i} - \left[\frac{1}{|R(g)|} \sum_{i \in R(g)} \vec{i} \right] \right]^2 \quad (8)$$

Dataset	# Reviews	# Reviewers	# Products
YelpNYC	359052	160225	923
YelpZIP	608598	260227	5044
Amazon	10261	1429	900
Playstore	325424	321436	192

Table 2: Statistics of four datasets.

4 Datasets

We collected four real-world datasets – **YelpNYC**: hotel/restaurant reviews of New York city [Rayana and Akoglu2015]; **YelpZip**: aggregation of reviews on restaurants/hotels from a number of areas with continuous zip codes starting from New York city [Rayana and Akoglu2015]; **Amazon**: reviews on musical instruments [He and McAuley2016], and **Playstore**: reviews of different applications available on Google Playstore. Fake reviews and spammers are already marked in both the Yelp datasets [Rayana and Akoglu2015]. For the remaining datasets, we employed three human experts² to label spammers based on the instructions mentioned in [Shojaee *et al.*2015, Mukherjee *et al.*2013a]³. The inter-rater agreement was 0.81 and 0.79 (Fleiss’ multi-rater kappa) for Amazon and Playstore, respectively. Table 2 shows the statistics of the datasets.

5 Experimental Results

Statistical validation. To measure the statistical significance of each component (say, products reviewed by the group members) of Hypothesis 1 (Sec. 3.2), we randomly generate pairs of reviewers (irrespective of the groups they belong to) and measure how their co-reviewing patterns (cumulative distribution) are different from the case if a pair of reviewers co-occur together in the same group. We hypothesize that both these patterns are different (whereas null hypothesis rejects our hypothesis). We observe that the difference is statistically significant as $p < 0.01$ (Kolmogorov-Smirnov test). See Supplementary [Dhawan *et al.*2019] for more discussion. We then perform evaluation in two stages – quality of the detected groups, and quality of the ranking algorithms.

5.1 Baseline Methods

We consider three existing methods (see Sec. 2 for their details) as baselines for both the evaluations – (i) **GSBP** [Wang *et al.*2016],(ii) **GGSpam** [Wang *et al.*2018b],and (iii) **GSRank** [Mukherjee *et al.*2012].

5.2 Evaluating Candidate Groups

Along with the three baselines mentioned above, we also consider two variations of DeFrauder as baselines for group detection: **DeFrauder_R** constructs the attributed product-product graph G based only on co-rating without considering the time of reviewing; and **DeFrauder_T** constructs G

²They were social media experts, and their age ranged between 25-35.

³<https://www.marketwatch.com/story/>

10-secrets-to-uncovering-which-online-reviews-are-fake-2018-09-21

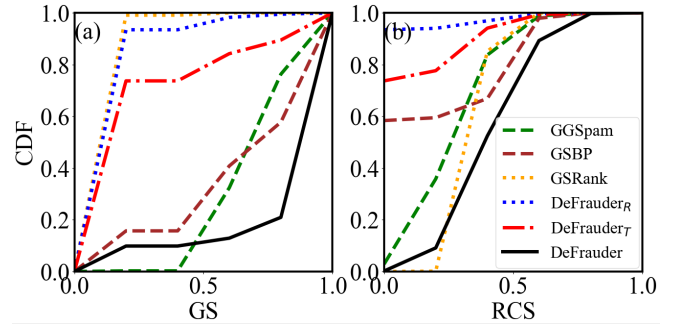


Figure 2: CDF of (a) GS and (b) RCS for YelpNYC. The more the distance of a CDF (corresponding to a method) from y-axis, the better the performance of the method.

based only on same co-reviewing time without considering co-rating. This also helps in justifying why both time and rating are important in constructing G for group detection. [Mukherjee *et al.*2012] suggested to use cumulative distribution (CDF) of *group size* (GS) and *review content similarity* (RCS) to evaluate the quality of the spam groups. Here we discard groups with size less than 2. **Group Size (GS)** favors large fraud groups as large groups are more damaging than smaller ones: $GS(g) = \frac{1}{1+e^{-(|R(g)|-2)}}$. **Review Content Similarity (RCS)** captures inter-reviewer review content similarity (as spammers copy reviews of each other): $RCS(g) = \max_{p \in P(g)} \left\{ \frac{1}{|R(g)|^2} \sum_{(i,j) \in R(g) \times R(g)} \cos(\theta(c_p^i, c_p^j)) \right\}$.

The larger the deviation of each distribution from the vertical axis (measured in terms of Earth Mover’s Distance (EMD)), the better the quality of the detected method [Mukherjee *et al.*2012].

Comparison. We choose the following parameter values as default based on our parameter selection strategy (Fig. 4): τ_t : 20 days, $\tau_{spam} = 0.4$. The number of groups we obtain from different datasets is reported in Table 3. Fig. 2 shows the CDF of GS and RCS for all the competing methods on YelpNYC⁴, which is further summarized quantitatively (in terms of EMD) in Table 3. DeFrauder outperforms the best baseline by 11.35% and 4.67% higher relative EMD (averaged over four datasets) for GS and RCS, respectively. We also notice that DeFrauder_T performs better than DeFrauder_R, indicating that temporal coherence is more important than rating coherence in detecting potential groups.

5.3 Evaluating Ranking Algorithm

We use **NDCG@k**, a standard graded evaluation metric. Since each reviewer was labeled as fraud/genuine, we consider the graded relevance value (ground-truth relevance score used in Normalized Discounted Cumulative Gain (NDCG) [Manning *et al.*2010]) for each group as the fraction of reviewers in a group, who are marked as fraud. The candidate groups are then ranked by each competing method, and top k groups are judged based on NDCG.

Comparison. We choose the following parameter values as default based on our parameter selection strategy (Fig. 4):

⁴Results are similar on the other datasets, and thus omitted.

Method	YelpNYC				YelpZIP				Amazon				Playstore			
	$ G $	GS	RCS	ND	$ G $	GS	RCS	ND	$ G $	GS	RCS	ND	$ G $	GS	RCS	ND
GGSpam	1218	0.574	0.218	0.567	1167	0.629	0.219	0.563	144	0.131	0.250	0.230	1213	0.749	0.010	0.464
GSBP	809	0.562	0.173	0.521	807	0.478	0.265	0.520	115	0.416	0.260	0.689	250	0.744	0.016	0.474
GSRank	998	0.102	0.313	0.569	1223	0.132	0.054	0.706	2922	0.293	0.309	0.144	994	0.577	0.018	0.476
DeFrauder _R	4399	0.124	0.021	—	6815	0.139	0.031	—	197	0.234	0.290	—	385	0.372	0.005	—
DeFrauder _T	152	0.237	0.069	—	3666	0.648	0.271	—	807	0.698	0.207	—	200	0.458	0.007	—
DeFrauder	1118	0.731	0.348	0.603	4574	0.667	0.287	0.602	713	0.718	0.314	0.768	940	0.841	0.018	0.789

Table 3: Performance of the competing methods: GGSpam [Wang *et al.*2018b], GSBP [Wang *et al.*2016], GSRank [Mukherjee *et al.*2012], DeFrauder_R, DeFrauder_T, and DeFrauder. Number of groups detected ($|G|$) are mentioned after removing groups of size less than 2 (cyan regions). Accuracy for the group detection (white regions) and ranking (gray regions) is reported in terms of EMD (the higher, the better), and NDCG@50 (ND) respectively. DeFrauder_R and DeFrauder_T are used only for group detection. The ranking methods of all baselines are run on the groups detection by DeFrauder.

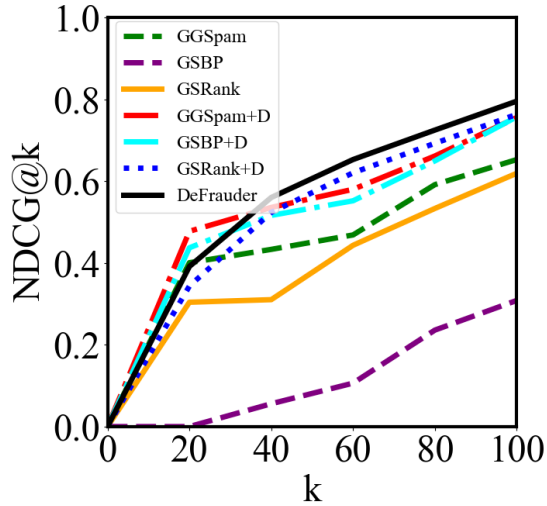


Figure 3: Performance on YelpNYC. Baselines are run with their detected groups as well as with the groups (+D) detected by DeFrauder (same naming convention as in Table 3).

$\alpha = \beta = 0.3, \gamma = 0.4, \tau_t = 20$ days, $\tau_r = (\max - \min)20\%$, $\tau_{spam} = 0.4$. We use Node2Vec [Grover and Leskovec2016] for embedding⁵. Fig. 3 shows the performance of the competing methods for different values of k (top k groups returned by each method). Since DeFrauder produces better groups (Sec. 5.2), we also check how the ranking method of each baseline performs on the groups detected of DeFrauder. Fig. 3 shows that with DeFrauder’s group structure, GGSpam and GSBP show better performance than DeFrauder till $k = 40$; after that DeFrauder dominates others. However, all the baselines perform poor with their own detected groups. This result also indicates the efficiency of our group detection method. Table 3 reports that DeFrauder beats other methods across all the datasets except YelpZIP on which GSRank performs better with DeFrauder’s detected groups. Interestingly, no single baseline turns out to be the best baseline across datasets. Nevertheless, DeFrauder outperforms the best baseline (varies across datasets) by 17.11% higher relative NDCG@50 (averaged over all the datasets).

⁵We tried with DeepWalk [Perozzi *et al.*2014] and LINE [Tang *et al.*2015] and obtained worse results compared to Node2Vec.

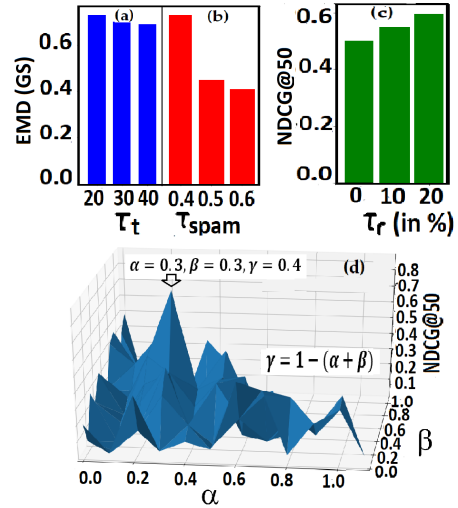


Figure 4: Parameter selection on YelpNYC. We vary each parameter keeping others as default. Since $\tau_t = 20$ produces best results for group detection, we kept this value for ranking as well.

6 Conclusion

In this paper, we studied the problem of fraud reviewer group detection in customer reviews. We established the principle of coherence among fraud reviewers in a group in terms of their co-reviewing patterns. This paper contributed in four directions: **Datasets:** We collected and annotated two new datasets which would be useful to the cybersecurity community; **Characterization:** We explored several group-level behavioral traits to model inter-personal collusive dynamics in a group; **Method:** We proposed DeFrauder, a novel method to detect and rank fraud reviewer groups; **Evaluation:** Exhaustive experiments were performed on four datasets to show the superiority of DeFrauder compared to five baselines.

Acknowledgements

The work was partially funded by Google Pvt. Ltd., DST (ECR/2017/001691) and Ramanujan Fellowship. The authors also acknowledge the support of the Infosys Centre of AI, IIT-Delhi, India.

References

- [Akoglu *et al.*, 2013] Leman Akoglu, Rishi Chandy, and Christos Faloutsos. Opinion fraud detection in online reviews by network effects. In *ICWSM*, pages 985–994, 2013.
- [Allahbakhsh *et al.*, 2013] Mohammad Allahbakhsh, Aleksandar Ignjatovic, Boualem Benatallah, Elisa Bertino, Norman Foo, et al. Collusion detection in online rating systems. In *Asia-Pacific Web Conference*, pages 196–207. Springer, 2013.
- [Chetan *et al.*, 2019] Aditya Chetan, Brihi Joshi, Hridoy Sankar Dutta, and Tanmoy Chakraborty. Corerank: Ranking to detect users involved in blackmarket-based collusive retweeting activities. In *WSDM*, pages 330–338, 2019.
- [Dhawan *et al.*, 2019] Sarthika Dhawan, Siva Charan Reddy Gangireddy, Shiv Kumar, and Tanmoy Chakraborty. Defrauder: Supplementary. <https://github.com/LCS2-IIITD/DeFrauder>, 2019.
- [Dutta *et al.*, 2018] Hridoy Sankar Dutta, Aditya Chetan, Brihi Joshi, and Tanmoy Chakraborty. Retweet us, we will retweet you: Spotting collusive retweeters involved in blackmarket services. In *ASONAM*, pages 242–249, 2018.
- [Fei *et al.*, 2013] Geli Fei, Arjun Mukherjee, Bing Liu, Meichun Hsu, Malu Castellanos, and Riddhiman Ghosh. Exploiting burstiness in reviews for review spammer detection. In *ICWSM*, 2013.
- [Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. Node2vec: Scalable feature learning for networks. In *SIGKDD*, pages 855–864, 2016.
- [He and McAuley, 2016] Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*, pages 507–517, 2016.
- [Kou *et al.*, 2004] Yufeng Kou, Chang-Tien Lu, Sirirat Sirwongwattana, and Yo-Ping Huang. Survey of fraud detection techniques. In *ICNSC*, pages 749–754, 2004.
- [Lim *et al.*, 2010] Ee-Peng Lim, Viet-An Nguyen, Nitin Jindal, Bing Liu, and Hady Wirawan Lauw. Detecting product review spammers using rating behaviors. In *CIKM*, pages 939–948. ACM, 2010.
- [Manning *et al.*, 2010] Christopher Manning, Prabhakar Raghavan, and Hinrich Schütze. Introduction to information retrieval. *Natural Language Engineering*, 16(1):100–103, 2010.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013.
- [Mukherjee *et al.*, 2012] Arjun Mukherjee, Bing Liu, and Natalie Glance. Spotting fake reviewer groups in consumer reviews. In *WWW*, pages 191–200, New York, April 2012.
- [Mukherjee *et al.*, 2013a] Arjun Mukherjee, Abhinav Kumar, Junhui Liu, Bing Wang, Meichun Hsu, Malu Castellanos, and Riddhiman Ghosh. Spotting opinion spammers using behavioral footprints. In *SIGKDD*, Chicago, USA, August 2013.
- [Mukherjee *et al.*, 2013b] Arjun Mukherjee, Vivek Venkataraman, Bing Liu, and Natalie Glance. What yelp fake review filter might be doing? In *ICWSM*, pages 1–12, 2013.
- [Ott *et al.*, 2011] Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *ACL-HLT*, pages 309–319, 2011.
- [Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *SIGKDD*, pages 701–710, 2014.
- [Rayana and Akoglu, 2015] Shebuti Rayana and Leman Akoglu. Collective opinion spam detection: Bridging review networks and metadata. In *SIGKDD*, pages 985–994, Sydney, NSW, Australia, August 2015.
- [Shojaee *et al.*, 2015] Somayeh Shojaee, Azreen Azman, Masrah Murad, Nurfadhline Sharef, and Nasir Sulaiman. A framework for fake review annotation. In *UKSIM-AMSS*, pages 153–158, 2015.
- [Tang *et al.*, 2015] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *WWW*, pages 1067–1077, 2015.
- [Wang *et al.*, 2011] Guan Wang, Sihong Xie, Bing Liu, and S Yu Philip. Review graph based online store review spammer detection. In *ICDM*, pages 1242–1247. IEEE, 2011.
- [Wang *et al.*, 2012] Guan Wang, Sihong Xie, Bing Liu, and Philip S Yu. Identify online store review spammers via social review graph. *ACM TIST*, 3(4):61:1–61:21, 2012.
- [Wang *et al.*, 2016] Zhuo Wang, Tingting Hou, Dawei Song, Zhun Li, and Tianqi Kong. Detecting review spammer groups via bipartite graph projection. *The Computer Journal*, 59(6):861–874, 2016.
- [Wang *et al.*, 2018a] Zhuo Wang, Songmin Gu, and Xiaowei Xu. Gsllda: Lda-based group spamming detection in product reviews. *Applied Intelligence*, 48(9):3094–3107, 2018.
- [Wang *et al.*, 2018b] Zhuo Wang, Songmin Gu, Xiangnan Zhao, and Xiaowei Xu. Graph-based review spammer group detection. *KIAS*, 55(3):571–597, Jun 2018.
- [Xu and Zhang, 2015] Chang Xu and Jie Zhang. Towards collusive fraud detection in online reviews. In *ICDM*, pages 1051–1056, 2015.
- [Xu *et al.*, 2013] Chang Xu, Jie Zhang, Kuiyu Chang, and Chong Long. Uncovering collusive spammers in chinese review websites. In *CIKM*, pages 979–988, 2013.