Spotting suspicious reviews via (quasi-)clique extraction

Paras Jain paras@gatech.edu
Shang-Tse Chen schen351@gatech.edu
Mozhgan Azimpourkivi mazim003@cs.fiu.edu
Duen Horng Chau polo@gatech.edu
Bogdan Carbunar carbunar@cs.fiu.edu

Overview

- We construct a Reviewer Similarity Graph of >300K users based on 1.6M reviews.
- We extract cliques and quasi-cliques from these Reviewer Similarity Graphs
- Suspicious cliques of sizes up to 11 users were found
- Large populations of "Yelp Scouts" were discovered
  - Yelp Scouts are users paid by Yelp to review certain venues
  - Contriversial as Scouts are artifical influences on Yelp

Motivation

- 16% of Yelp reviews are fraudulent (data from Yelp)
- Community trust and the livelihood of legitimate businesses are threatened by venue review fraud
- Recent literature has explored finding fraud through analyzing link structures among fraudsters in eBay

Objectives

- Generate review similarity graphs based on data from the Yelp Dataset Challenge
- Mine reviewer similarity graphs for cliques and quasi-cliques
- Study if such cliques are indeed suspicious
  - Is the (quasi-)clique structure an indicator for suspicious behavior?
  - How to minimize false positives for flagged suspicious reviews?

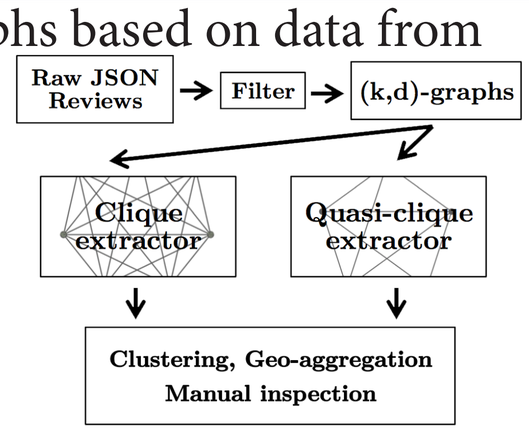

Figure 1: Data Processing Pipeline

Reviewer Similarity Graphs

- The entire Yelp social graph is massive with 3M+ edges
- Reviewer similarity graph is created by extracting a (k,d)-graph
  - Edge represents two users reviewing ≥k venues within d days of each other
  - Nodes represent Yelp users
- 42,153 venues and 1,125,458 reviews represented

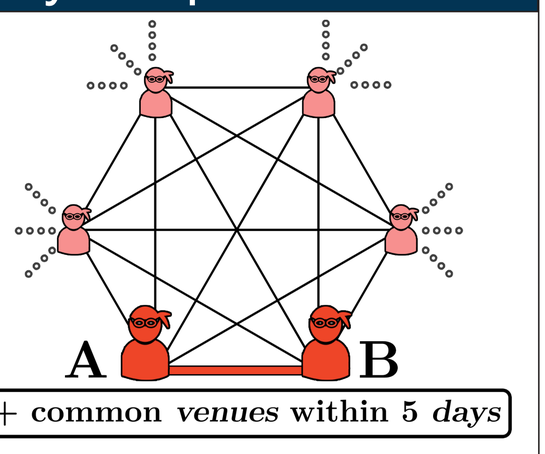

Figure 2: Review similarity graphs link users with similar reviewing patterns

Clique Extraction

- The maximum clique problem is NP-hard
- Cliques are extracted from the (k,d)-graph using Bron-Kerbosch
  - Sparse real-world graphs allow clique finding in reasonable time

Quasi-clique Extraction

- Quasi-cliques are sub-graphs with edge densities above a certain threshold
- Perfect cliques are not common in real social graphs
- Takeaki Uno presented a greedy method for searching for quasi-cliques through growing cliques

Results

- Cliques up to size 11 and quasi-cliques of size up to 12 were found

(k, d)-graph	3,5	3,6	3,8	4,5	4,6	5,5	5,6	6,5
9–clique	112	152	1040	29	73	13	28	10
10–clique	22	25	290	3	13	1	3	**1**
11–clique	**2**	**2**	**50**	—	**1**	—	—	—

(k, d)-graph	6,5	6,8	7,5	7,8	8,5	8,8	9,5
9–quasiclique	144	649	94	351	42	227	8
10–quasiclique	44	315	33	134	12	84	—
11–quasiclique	**7**	**100**	**4**	**33**	—	**15**	—
12–quasiclique	**1**	**20**	—	**4**	—	**1**	—

Table 1: Counts of (quasi-)cliques for various (k,d)- graphs. Graphs with larger *k* and smaller *d* values are more suspicious.

- Manual examination of suspicious accounts (members of large cliques) revealed large populations of "Yelp Scouts"
- Yelp Scouts tend to associate with each other. See Figure 2 where the red nodes are closely clustered.

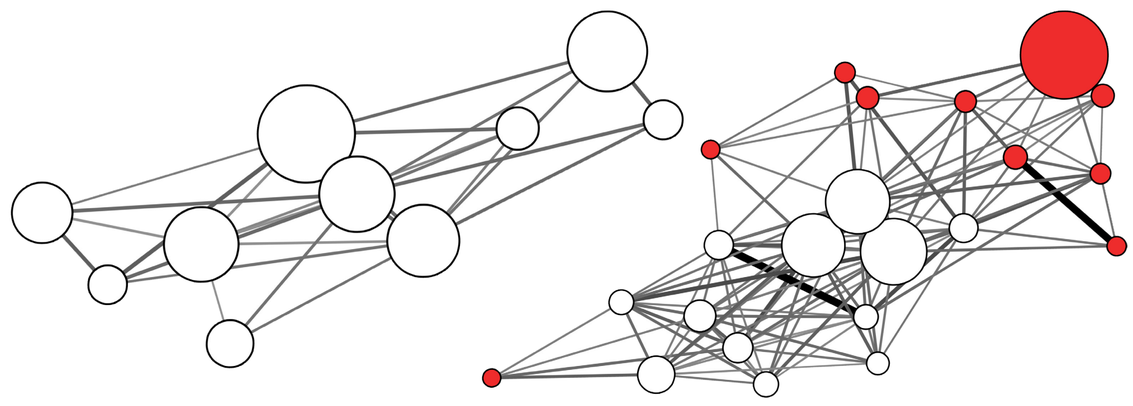

Figure 2: Scouts are tightly clustered and appear to associate with other Scouts. Graph of combined cliques in a weighted (6, 5)-graph. Larger nodes represent more reviews. Red nodes are Yelp Scouts while white nodes are regular users.

- In the weighted (6, 5)-graph, **23%** of users in a clique were Scouts
- Scouts are very tightly geographically clustered. For example, Edinburg has large populations of past Yelp scouts. One reason for this might be that Yelp pays locals to seed reviews in new regions.
- Reviews written by suspicious users are generally positive.

Figure 3: A geographic plot of Edinburgh with reviews (part of graph edges) from cliques in the (6, 5)-graph. The majority of reviews are located in a narrow region and are generaly positive.

Conclusion

- Our research yielded favorable results in finding suspicious reviews.
- As a side-effect of the clique and quasi-clique extraction process, Yelp Scouts were found alongside possibly fraudulent results due to their artificial reviewing structures.
- The research has extensions to other crowd-sourced review sites, such as Amazon, Walmart.com or Tripadvisor and in fraud detection for electronic commerce platforms.
- We plan to incorporate other rich signals from the Yelp data to help with this, such as by analyzing review text, and the spatial and temporal relationships among reviewed venues.

Georgia Tech | FIU FLORIDA INTERNATIONAL UNIVERSITY