Fraud Detection through Graph-Based User Behavior Modeling

Alex Beutel*
Carnegie Mellon University
Pittsburgh, PA, USA
abeutel@cs.cmu.edu

Leman Akoglu
Stony Brook University
Stony Brook, NY, USA
leman@cs.stonybrook.edu

Christos Faloutsos
Carnegie Mellon University
Pittsburgh, PA, USA
christos@cs.cmu.edu

ABSTRACT

How do anomalies, fraud, and spam effect our models of normal user behavior? How can we modify our models to catch fraudsters? In this tutorial we will answer these questions - connecting graph analysis tools for user behavior modeling to anomaly and fraud detection. In particular, we will focus on three data mining techniques: subgraph analysis, label propagation and latent factor models; and their application to static graphs, e.g. social networks, evolving graphs, e.g. "who-calls-whom" networks, and attributed graphs, e.g. the "who-reviews-what" graphs of Amazon and Yelp.

For each of these techniques we will give an explanation of the algorithms and the intuition behind them. We will then give brief examples of recent research using the techniques to model, understand and predict normal behavior. With this intuition for how these methods are applied to graphs and user behavior, we will focus on state-of-the-art research showing how the outcomes of these methods are effected by fraud, and how they have been used to catch fraudsters.

1. TUTORIAL PERSPECTIVE

In this tutorial we focus on fraud, spam and anomaly detection through the lens of normal user behavior modeling. The data mining and machine learning communities have developed a plethora of models and methods for understanding user behavior. However, these methods generally assume that the behavior is that of real, honest people. On the other hand, fraud detection systems frequently use similar techniques as those used in modeling "normal" behavior, but are often framed as an independent problem. However, by focusing on the relations and intersections of the two perspectives we can gain a more complete understanding of the methods and hopefully inspire new research joining these two communities.

2. TARGET AUDIENCE

This tutorial is aimed at anyone interested in understanding user behavior data, from data miners to security researchers to practitioners from industry and government. For those new to these algorithms, the tutorial will cover the necessary background material to understand these systems and will offer a concise, intuitive overview of the state-of-the-art in user behavior modeling. Additionally, the tutorial aims to offer a new perspective that will be valuable and interesting even for researchers with more experience in these domains. In particular, for those researchers having worked on fraud detection systems, we hope to inspire new research directions through connecting with recent developments in modeling "normal" behavior.

*Alex will be available to present, and Professor Akoglu and Professor Faloutsos will also present if their schedules allow.

3. INSTRUCTORS

Alex Beutel is a fifth year Ph.D. candidate at Carnegie Mellon University in the Computer Science Department. He previously received his B.S. from Duke University. His Ph.D. research focuses on large scale user behavior modeling, covering both recommendation systems and fraud detection systems. He has interned at Facebook on both the Site Integrity and News Feed Ranking teams, at Microsoft in the Cloud and Information Services Laboratory, and at Google Research. His research is supported by a Facebook Fellowship and the National Science Foundation Graduate Research Fellowship Program. More details can be found at `http://alexbeutel.com`.

Leman Akoglu is an Assistant Professor in the Department of Computer Science at Stony Brook University. She received her Ph.D. from the Computer Science Department at Carnegie Mellon University in 2012. She also worked at IBM T. J. Watson Research Labs and Microsoft Research at Redmond during summers. Her research interests span a wide range of data mining and machine learning topics with a focus on algorithmic problems arising in graph mining, pattern discovery, social and information networks, and especially anomaly mining; outlier, fraud, and event detection. Dr. Akoglu's research has won 4 publication awards; Best Research Paper at SIAM SDM 2015, Best Paper at ADC 2014, Best Paper at PAKDD 2010, and Best Knowledge Discovery Paper at ECML/PKDD 2009. She also holds 3 U.S. patents filed by IBM T. J. Watson Research Labs. Dr. Akoglu is a recipient of the NSF CAREER award (2015) and Army Research Office Young Investigator award (2013). Her research is currently supported by the National Science Foundation, the US Army Research Office, DARPA, and a gift from Northrop Grumman Aerospace Systems. More details can be found at `http://www.cs.stonybrook.edu/~leman`.

Christos Faloutsos is a Professor at Carnegie Mellon University. He has received the Presidential Young Investigator Award by the National Science Foundation (1989), the Research Contributions Award in ICDM 2006, the Innovations award in KDD'10, 20 "best paper" awards, and several teaching awards. He has served as a member of the executive committee of SIGKDD; he has published over 200 refereed articles, 11 book chapters and one monograph. He holds five patents and he has given over 30 tutorials and over 10 invited distinguished lectures. His research in-

terests include data mining for graphs and streams, fractals, database performance, and indexing for multimedia and bio-informatics data. More details can be found at `http://www.cs.cmu.edu/~christos/`.

4. OUTLINE

I. **Introduction**

A. Graphs are a useful abstraction for a wide variety of domains: social networks, movie or product ratings and reviews, text in articles, medical diagnoses, financial transactions, etc.

B. How can we make sense of the data? What does normal behavior look like? For example, we can predict ratings on Netflix or friends on Twitter and fill in missing information on Facebook.

C. Fraud is rampant in nearly all of these applications - fake reviews on Yelp, purchased followers on Twitter, inflated trust on eBay, medical fraud, bank fraud. This activity deceives users and manipulates our prediction algorithms. Therefore it is important to understand how fraud influences our models and how we can isolate and catch anomalous behavior.

II. **Subgraph Analysis and Patterns**

A. **Background:** Graph clustering and partitioning

i. Local search and graph queries [4, 19]

ii. Co-clustering and cross associations [23, 13]

B. **Normal Behavior:** Subgraph Patterns

i. Ego-nets: [30, 5]

ii. Subgraph patterns in social networks: [41]

iii. Influence of subgraphs on recommendation: [42]

iv. Co-clustering for recommendation: ACCAMS [6]

C. **Abnormal Behavior:** What are anomalous or fraudulent subgraphs?

i. Ego-net analysis: COI [14], OddBall [3]

ii. Attributed subgraphs: FocusCO [36], SODA [20], CODA [16]

iii. Temporal lockstep behavior: CopyCatch [8]; extended in [11]

iv. Subsets of attributes: CrossSpot [25]

v. Graph queries: "volcanoes" and "blackholes" on static graphs [31], attributed subgraph [47]

vi. Detecting fraud with co-clustering [35]

vii. Graph cut for intrusion detection [15]

III. **Label Propagation Methods**

A. Random Walks and Eigenvectors

i. **Background:** Why do random walks "find" important parts of a graph?

ii. **Normal Behavior:** Power method, HITS [28], and PageRank [9]

B. Belief and Label Propagation

i. **Background:** What is semi-supervised learning? What are belief and label propagation?

ii. **Normal Behavior:** Predict attributes on nodes or why certain people are friends [12]

C. **Abnormal Behavior:** How can random walks help us find fraud?

i. Surprising patterns in HITS: CatchSync [26]

ii. Modifications of PageRank: TrustRank [22], SybilRank[10], CollusionRank[17], BadRank [44]

iii. Use belief propagation for "guilt-by-association:"

- Binary graphs: NetProbe [32]
- Attributed graphs: FraudEagle [2], [27]
- "Guilt-by-constellation" [43]

IV. **Latent Factor Models**

A. **Background:** What is the singular vector decomposition?

i. Generalization from Eigenvectors and HITS

ii. Why do latent factor models, like SVD, work well for relational data?

iii. What do the factors typically represent? Why? For example, the factorization of a user by movie ratings matrix gives genres; the decomposition of a word by document matrix gives topics.

B. **Normal Behavior:** Modifications for different settings:

i. Finding communities (binary matrices): MMSB [1], overlapping communities [45]

ii. Missing data and prediction: SVD++ [29], BPMF [38], CoBaFi [7]

iii. Multi-modal data: PARAFAC, Tensor factorization [34, 33]

iv. Coupled factorization [40, 21]

C. **Abnormal Behavior:** What happens if there is fraud in the data?

i. Surprising patterns in latent factors (binary matrices): EigenSpokes [37], Get-the-scoop [24], fBox [39]

ii. Surprising group patterns in ratings data: CoBaFi [7]

iii. Surprising pattern in coupled factorization for heterogeneous graphs [18]

iv. Group anomalies: GLAD [46]

V. **Looking forward**

A. How can we handle *multiple data sources and complex data*?

B. With more complex data and methods, how can maintain the *interpretability* of discovered fraud?

C. Can we provide stronger *provable limits and guarantees* on our systems?

5. REFERENCES

[1] Edoardo M Airoldi, David M Blei, Stephen E Fienberg, and Eric P Xing. Mixed membership stochastic blockmodels. In *Advances in Neural Information Processing Systems*, pages 33–40, 2009.

[2] Leman Akoglu, Rishi Chandy, and Christos Faloutsos. Opinion fraud detection in online reviews by network effects. In *ICWSM*, 2013.

[3] Leman Akoglu, Mary McGlohon, and Christos Faloutsos. Oddball: Spotting anomalies in weighted graphs. *PAKDD 2010*, 21-24 June 2010.

[4] Reid Andersen, Fan Chung, and Kevin Lang. Local graph partitioning using pagerank vectors. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 475–486. IEEE, 2006.

[5] Lars Backstrom and Jon Kleinberg. Romantic partnerships and the dispersion of social ties: a network analysis of relationship status on facebook. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 831–841. ACM, 2014.

[6] Alex Beutel, Amr Ahmed, and Alexander J Smola. ACCAMS: Additive Co-Clustering to Approximate Matrices Succinctly. In *Proceedings of the 24th international conference on World wide web*. International World Wide Web Conferences Steering Committee, 2015.

[7] Alex Beutel, Kenton Murray, Christos Faloutsos, and Alexander J Smola. CoBaFi: Collaborative Bayesian Filtering. In *Proceedings of the 23rd international conference on World wide web*, pages 97–108. International World Wide Web Conferences Steering Committee, 2014.

[8] Alex Beutel, Wanhong Xu, Venkatesan Guruswami, Christopher Palow, and Christos Faloutsos. Copycatch: stopping group attacks by spotting lockstep behavior in social networks. In *Proceedings of the 22nd international conference on World Wide Web*, pages 119–130. International World Wide Web Conferences Steering Committee, 2013.

[9] Sergey Brin and Larry Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of the Seventh International World Wide Web Conference*, 1998.

[10] Qiang Cao, Michael Sirivianos, Xiaowei Yang, and Tiago Pregueiro. Aiding the detection of fake accounts in large scale social online services. In *NSDI*, 2012.

[11] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. Uncovering large groups of active malicious accounts in online social networks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 477–488. ACM, 2014.

[12] Deepayan Chakrabarti, Stanislav Funiak, Jonathan Chang, and Sofus A Macskassy. Joint inference of multiple label types in large networks. *arXiv preprint arXiv:1401.7709*, 2014.

[13] Deepayan Chakrabarti, Spiros Papadimitriou, Dharmendra S Modha, and Christos Faloutsos. Fully automatic cross-associations. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 79–88. ACM, 2004.

[14] Corinna Cortes, Daryl Pregibon, and Chris Volinsky. *Communities of interest*. Springer, 2001.

[15] Qi Ding, Natallia Katenka, Paul Barford, Eric Kolaczyk, and Mark Crovella. Intrusion as (anti) social communication: characterization and detection. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 886–894. ACM, 2012.

[16] Jing Gao, Feng Liang, Wei Fan, Chi Wang, Yizhou Sun, and Jiawei Han. On community outliers and their efficient detection in information networks. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 813–822. ACM, 2010.

[17] Saptarshi Ghosh, Bimal Viswanath, Farshad Kooti, Naveen Kumar Sharma, Gautam Korlam, Fabricio Benevenuto, Niloy Ganguly, and Krishna Phani Gummadi. Understanding and combating link farming in the twitter social network. In *Proceedings of the 21st international conference on World Wide Web*, pages 61–70. ACM, 2012.

[18] Manish Gupta, Jing Gao, and Jiawei Han. Community distribution outlier detection in heterogeneous information networks. In *Machine Learning and Knowledge Discovery in Databases*, pages 557–573. Springer, 2013.

[19] Manish Gupta, Jing Gao, Xifeng Yan, Hasan Cam, and Jiawei Han. On detecting association-based clique outliers in heterogeneous information networks. In *Advances in Social Networks Analysis and Mining (ASONAM), 2013 IEEE/ACM International Conference on*, pages 108–115. IEEE, 2013.

[20] Manish Gupta, Arun Mallya, Subhro Roy, Jason HD Cho, and Jiawei Han. Local learning for mining outlier subgraphs from network datasets. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pages 73–81, 2014.

[21] Nitish Gupta and Sameer Singh. Collective factorization for relational data: An evaluation on the yelp datasets.

[22] Zoltán Gyöngyi, Hector Garcia-Molina, and Jan Pedersen. Combating web spam with trustrank. In *VLDB Endowment*, pages 576–587, 2004.

[23] J. A. Hartigan. Direct clustering of a data matrix. *Journal of the american statistical association*, 67(337):123–129, 1972.

[24] Meng Jiang. Peng cui, alex beutel, christos faloutsos and shiqiang yang. inferring strange behavior from connectivity pattern in social networks. In *The 18th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2014.

[25] Meng Jiang, Alex Beutel, Peng Cui, Bryan Hooi, Shiqiang Yang, and Christos Faloutsos. A general suspiciousness metric for dense blocks in multimodal data. In *ICDM*, 2015.

[26] Meng Jiang, Peng Cui, Alex Beutel, Christos Faloutsos, and Shiqiang Yang. Catchsync: catching synchronized behavior in large directed graphs. In *Proceedings of the 20th ACM SIGKDD international*

conference on Knowledge discovery and data mining, pages 941–950. ACM, 2014.

[27] Enric Junqué de Fortuny, Marija Stankova, Julie Moeyersoms, Bart Minnaert, Foster Provost, and David Martens. Corporate residence fraud detection. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pages 1650–1659, New York, NY, USA, 2014. ACM.

[28] J.M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.

[29] Yehuda Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434. ACM, 2008.

[30] Jure Leskovec and Julian J Mcauley. Learning to discover social circles in ego networks. In *Advances in neural information processing systems*, pages 539–547, 2012.

[31] Zhongmou Li, Hui Xiong, and Yanchi Liu. Detecting blackholes and volcanoes in directed networks. *arXiv preprint arXiv:1005.2179*, 2010.

[32] S. Pandit, D.H. Chau, S. Wang, and C. Faloutsos. Netprobe: a fast and scalable system for fraud detection in online auction networks. In *Proceedings of the 16th international conference on World Wide Web*, pages 201–210. ACM, 2007.

[33] Evangelos Papalexakis, Konstantinos Pelechrinis, and Christos Faloutsos. Spotting misbehaviors in location-based social networks using tensors. In *Proceedings of the companion publication of the 23rd international conference on World wide web companion*, pages 551–552. International World Wide Web Conferences Steering Committee, 2014.

[34] Evangelos E Papalexakis, Leman Akoglu, and D Ience. Do more views of a graph help? community detection and clustering in multi-graphs. In *Information Fusion (FUSION), 2013 16th International Conference on*, pages 899–905. IEEE, 2013.

[35] Evangelos E Papalexakis, Alex Beutel, and Peter Steenkiste. Network anomaly detection using co-clustering. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*, pages 403–410. IEEE Computer Society, 2012.

[36] Bryan Perozzi, Leman Akoglu, Patricia Iglesias Sánchez, and Emmanuel Müller. Focused clustering and outlier detection in large attributed graphs. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1346–1355. ACM, 2014.

[37] B. Aditya Prakash, Mukund Seshadri, Ashwin Sridharan, Sridhar Machiraju, and Christos Faloutsos. Eigenspokes: Surprising patterns and scalable community chipping in large graphs. *PAKDD 2010*, 21-24 June 2010.

[38] Ruslan Salakhutdinov and Andriy Mnih. Bayesian probabilistic matrix factorization using markov chain monte carlo. In *Proceedings of the 25th international conference on Machine learning*, pages 880–887. ACM, 2008.

[39] Neil Shah, Alex Beutel, Brian Gallagher, and Christos Faloutsos. Spotting suspicious link behavior with fBox: An adversarial perspective. In *ICDM*, 2014.

[40] Ajit P Singh and Geoffrey J Gordon. A unified view of matrix factorization models. In *Machine Learning and Knowledge Discovery in Databases*, pages 358–373. Springer, 2008.

[41] Johan Ugander, Lars Backstrom, and Jon Kleinberg. Subgraph frequencies: Mapping the empirical and extremal geography of large graph collections. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1307–1318. International World Wide Web Conferences Steering Committee, 2013.

[42] Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences*, page 201116502, 2012.

[43] Véronique Van Vlasselaer, Veronique VanVlasselaer, Leman Akoglu, Tina Eliassi-Rad, Monique Snoeck, and Bart Baesens. Guilt-by-constellation: Fraud detection by suspicious clique memberships. In *Proceedings of 48 Annual Hawaii International Conference on System Sciences*, 2015.

[44] Baoning Wu, Vinay Goel, and Brian D Davison. Propagating trust and distrust to demote web spam. *MTW*, 190, 2006.

[45] Jaewon Yang and Jure Leskovec. Structure and overlaps of ground-truth communities in networks. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(2):26, 2014.

[46] Rose Yu, Xinran He, and Yan Liu. Glad: group anomaly detection in social media analysis. In *SIGKDD*, pages 372–381, 2014.

[47] Honglei Zhuang, Jing Zhang, George Brova, Jie Tang, Hasan Cam, Xifeng Yan, and Jiawei Han. Mining query-based subnetwork outliers in heterogeneous information networks.