

STATISTICAL ANALYSIS AND DATA MINING

Original Article

To catch a fake: Curbing deceptive Yelp ratings and venues

Mahmudur Rahman, Bogdan Carbunar , Jaime Ballesteros, Duen Horng (Polo) Chau

First published: 10 April 2015

<https://doi.org/10.1002/sam.11264>

Cited by: 2

This research was supported in part by DoD grant W911NF - 13 - 1 - 0142 and NSF grant EAGER - 1450619.



About



Access

als



ire

Abstract

The popularity and influence of reviews, make sites like Yelp ideal targets for malicious behaviors. We present Marco, a novel system that exploits the unique combination of social, spatial and temporal signals gleaned from Yelp, to detect venues whose ratings are impacted by fraudulent reviews. Marco increases the cost and complexity of attacks, by imposing a tradeoff on fraudsters, between their ability to impact venue ratings and their ability to remain undetected. We contribute a new dataset to the community, which consists of both ground truth and gold standard data. We show that Marco significantly outperforms state-of-the-art approaches, by achieving 94% accuracy in classifying reviews as fraudulent or genuine, and 95.8% accuracy in classifying venues as deceptive or legitimate. Marco successfully flagged 244 deceptive venues from our large dataset with 7,435 venues, 270,121 reviews and 195,417 users. Furthermore, we use Marco to evaluate the impact of Yelp events, organized for elite reviewers, on the hosting venues. We collect data from 149 Yelp elite events throughout the US. We show that two weeks after an event, twice as many hosting venues experience a significant rating boost rather than a negative impact.

Citing Literature 

About Wiley Online Library 

Help & Support 

Opportunities



Connect with Wiley



Copyright © 1999-2018 John Wiley & Sons, Inc. All rights reserved

WILEY