

# Batman or the Joker? The Powerful Urban Computing and its Ethics Issues

Kaiqun Fu, Abdulaziz Alhamadani, Taoran Ji, Chang-Tien Lu  
Department of Computer Science, Virginia Tech, Falls Church, USA  
E-mail: {fukaiqun, hamdani, jtr, ctlu}@vt.edu

## Abstract

*The exponential growth of the urban data generated by urban sensors, government reports, and crowdsourcing services endorses the rapid development of urban computing and spatial data mining technologies. Easier accessibility to such enormous urban data may be a double-bladed sword. On the one hand, urban data can be applied to solve a wide range of practical issues such as urban safety analysis and urban event detection. On the other hand, ethical issues such as biasedly polluted urban data, problematic algorithms, and unprotected privacy may cause moral disaster not only for the research fields but also for the society. This paper seeks to identify ethical vulnerabilities from three primary research directions of urban computing: urban safety analysis, urban transportation analysis, and social media analysis for urban events. Visions for future improvements in the perspective of ethics are addressed.*

## 1 Introduction

Ethics issues in data mining such as privacy and anonymity protection, algorithmic biases, and surveillances have been raised by some of the researchers in the recent decade [22, 14, 3]. With the ubiquitous deployment of the urban sensors and rapid growth of crowdsourcing technologies, urban computing and spatial data mining techniques begin to thrive in detecting and analyzing urban events. An enormous amount of urban data is generated from multiple sources such as traffic sensors, air quality meters, various government agency reports, and event social media as crowdsourcing. Such a rise of “big urban data” has boosted the development of the urban computing techniques in several important directions such as urban safety inference, intelligent transportation systems, and social media analysis for urban events. However, the ethical issues in urban computing and spatial data mining fields do not receive enough attention as the exponential growth of the big urban data.

This paper is dedicated to providing discussions of ethics for several research directions of urban computing in detail. The paper is structured into three directions:

- *Urban Safety Analysis.* Safety and security-related events detection and prediction are important research topics in urban computing and spatial data mining. With the rapid deployment of urban sensors, more accessible government reports, and social media platforms, there are more innovative approaches and applications to tackle urban safety analysis problems. However, the importance of ethical issues is rarely addressed by the newly proposed methods. We summarize the state-of-the-art works in urban safety analysis, point out the ethical issues in the existing practices, and provide potential improvements.
- *Urban Transportation Prediction and Analysis.* To provide and maintain benign mobility within urban areas is one of the most fundamental requirements of Intelligent Transportation Systems (ITS). The massive mounted traffic sensors nowadays are generating Gigabytes of data per hour. With effective spatial data

mining techniques and urban computing algorithms, various ITS related topics such as traffic forecasting, incident and disruption detection, and incident impact analysis have been proposed. However, such transportation-related urban data is mostly generated by government agencies, which leads to unscrupulous concerns such as intruding commuters’ privacy and discriminatory surveillance. We summarize the state-of-the-art works in urban transportation analysis, point out the ethical issues in the existing works, and provide potential improvements.

- *Social Media Analysis for Urban Events.* Social media and location-based services with user-posted content have generated a staggering amount of information that has potential applications in various areas such as urban event detection and incident analysis. The omnipresence of social media and location data is capable of representing people’s behavior, attitudes, feelings, and relationships. These features can be ethically problematic, even where such data exist in the public domain. Unfortunately, such issues are rarely addressed in the research fields of urban computing and spatial data mining. We summarize the state-of-the-art works in urban events detection from social media, point out the ethical issues in the existing works, and provide potential improvements.

## 2 Ethics of Urban Computing and Urban Safety

With the ubiquitous deployment of the urban sensors and rapid growth of crowdsourcing technologies, urban computing and spatial data mining techniques begin to thrive in detecting and analyzing urban events. Among all categories of urban events, safety and security-related events should be treated as one of the most important ones without a doubt. The urban computing community has addressed important problems such as urban safety and crime prediction [6, 26, 9], safe route recommendations [33, 12], and threats detection [21]. However, the convenience and accessibility of such abundant urban and spatial data generated by the urban sensors, end-users, and city administrators put a spotlight on unethical issues such as biased datasets, biased algorithms, biased results, and compromised privacy. Such problems are rarely addressed by the researchers in the urban safety analysis fields. In this section, we summarize some of the pioneering research works in the urban safety analysis field, address the potential ethical issues, and then provide our visions on how to tackle and improve or mitigate the current research status of the ethical issues in the urban computing and spatial data mining fields.




Figure 1: Spatial Distributions of Different Categories of Crime

### 2.1 Current Works for Urban Safety Analysis

**Urban Crime Perception with Convolutional Neural Network.** Nadai et al. [6] studied the relationship between a neighborhood’s appearance of safety and its levels of human activity in two main Italian cities. By combining the recognized safety scores predicted using a convolutional neural network trained on Google Street View images with mobile phone data, they found that there is a significant positive and negative correlation

depending on gender and age demographics. Urban recommendations could be given to any specific city if the data were available. Figure 1 shows the correlation between the crime distribution and the physical appearances of the city. In another study to explore the connection between urban perception and crime inferences, Liu et al. [26] present a unified framework to learn to quantify safety attributes of physical urban environments using crowd-sourced street-view photos without human annotations. A large-scale urban image dataset is collected in multiple major cities. Safety scores from the government’s criminal records are collected as objective safety indicators. A deep convolutional neural network is proposed to parameterize the instance-level scoring function. Figure 2 shows the structure of the proposed model. The method is capable of localizing interesting images and image regions for each place.




Figure 2: Representation of the Convolutional Neural Network Structure for Learning Urban Crime

**Safe Route Recommendation with Crowdsourcing.** Shah et al. [33] propose a travel route recommendation system that suggests safe travel routes in real-time by incorporating social media data resources and points of interest review summarization techniques. The system consists of an efficient route recommendation service that considers safety and user interest factors, a transportation-related tweets retriever with high accuracy, and a text summarization module that provides summaries of location-based Twitter data and Yelp reviews to enhance route recommendation service.

**Threats Identification from News and Social Media.** Airports are a prime target for terrorist organizations, drug traffickers, smugglers, and other nefarious groups. Homeland security professionals must rely on measures of the attractiveness of an airport as a target for attacks. Khandpur et al. [21] present an open-source indicators approach, using urban data sources such as social media and news articles, to conduct a relative threat assessment, i.e., estimating if one airport is under greater threat than another.

## 2.2 Ethics Issues for Urban Safety Analysis

For crime prediction and safety, analysis works in urban computing, predictive policing and similar urban safety AI models are subjects to problematical algorithms. The problematical algorithms can be generated in several ways:

**1) Biased data sources.** Existing works involve safety analysis [6, 12] utilize crime reports from the police departments. Such reports from the authorities may inevitably be biasedly polluted during recording, for example, based on stereotypes towards certain groups of people. The review “Policing Predictive Policing” revealed that the crime data is not just biased, it is “notoriously incomplete”. There are certain crimes which were under-reported to authorities such as sexual assault, domestic violence, and fraud. Not only that, some communities, frustrated with current policing practices, simply decline to report crimes [8]. The review, also, reported that half of the crimes with victims go unreported. With such incomplete data, it is not fair to employ PredPol in areas that do not have enough information or poor data. Another case where data can be biased in a different way happened when the GPS location data of potholes in Boston was collected by the smartphone app StreetBump to help patch the potholes in Boston. The idea was great but there was a major problem. Crawford reported that people in lower-income groups and elder residents groups are less likely to have smartphones, where smartphone access can be as low as 16%. This indicates that there is a significant amount of the population have not reported the potholes, therefore there is a crucial part missing from the data [4].

**2) Biased Data produces biased results (garbage in, garbage out):** Deep learning algorithms learn from the existing data, and when bias exists in the data, discrimination will exist in the results. The flaw in the data will allow these algorithms to “inherit the prejudices of prior decisions makers” [1]. For example, the primary goal of predictive policing is to inform better personnel deployment, then areas that are “less hot” as per the software would most definitely see less police deployment and hence less crime intervention by law enforcement in general. However, this effect would be seen in cases when the police department relied solely on the predictive policing tool over its network of police informants and other sources of ground-level reconnaissance [8]. The intuitive solution to mitigate discrimination in AI models is to equally represent everyone. However, it appears to be another problem. Nordling [30] points out that information from certain population groups usually is not representative because of missing data. Therefore, this lack of diversity is likely to result in biased algorithms. Another intuitive solution is “simply to throw more data at the AI” to expose the AI to wrong cases and correct the errors, another solution is to “change the first network’s inputs” [15]. However, such solutions are not yet enough to fix those errors but at least the errors are recognized.

**Vision:** An effort mentioned by Kamiran when there is bias in data because data is not accessible, missing or underrepresented is to compensate some of the bias by oversampling underrepresented communities [19]. Further, one of the several solutions for the problem of the discriminatory algorithm that Barocas et al. [1] mentioned is to educate the employees (researchers) to rectify the problem if they understand the causes of the problem, although it is hard to identify the problems sometimes because the resulting discrimination is unintentional. There has been a lot of research to solve the bias in data and bias outcomes. Yet, the problem still persists which requires more effort in future work.

### 3 Transportation Predictions with Urban Data Sources

Detecting and analyzing transportation-related incidents and forecasting transportation status is critical to the success of the research fields of Intelligent Transportation Systems (ITS) and smart cities. Predicting traffic on urban traffic networks using spatiotemporal models has become a popular research area in the past decade [34]. Transportation-related incident detection, analysis of transit systems [17] and road networks [10] have also gained increasing attention in recent years. Various previous works on traffic prediction and incident analysis apply traffic data sources from traffic sensor providers such as INRIX<sup>1</sup> and Regional Integrated Transportation Information System<sup>2</sup> (RITIS). However, most of the existing research works ignore ethical issues while utilizing the traffic data generated by urban traffic sensors. Ethical issues such as surveillance on urban activities with such an enormous amount of sensor-generated data and privacy issues on conducting experiments on big data that is generated based on people’s daily mobility should be brought under the spotlight by the researchers in the research areas of intelligent transportation systems and smart city.

#### 3.1 Current Works for Transportation Predictions

**Transit Service Disruption Detection from Social Media.** Transit agencies are seeking to move beyond traditional customer questionnaires and manual service inspections to leveraging open source indicators like social media for detecting emerging transit events. Inspired by the multi-task learning framework, Ji et al. [17] propose the Metro Disruption Detection Model (*MDDM*), which captures the semantic similarity between transit lines in the Twitter space. The *MDDM* model novel constraints on feature semantic similarity exploiting prior knowledge about the spatial connectivity and shared tracks of the metro network. Figure 3 shows disruption events for 2015 on the Orange, Silver, and Blue lines operated by the Washington Metropolitan Transit Authority. Disruption 1, disruption 2, and disruption 3 occurred on the Orange line. Disruption 4 and disruption 5 occurred on the

---

<sup>1</sup>INRIX: <http://inrix.com/>

<sup>2</sup>Regional Integrated Transportation Information System: <https://ritis.org/>

Silver line. Disruption 6, disruption 7 and disruption 8 occurred on the Blue line. *MDDM* model successfully detects disruptions 1, 4 and 6.




Figure 3: A timeline of metro disruptions on the Orange, Blue, and Silver metro lines in 2015. Events along these spatially interconnected lines often co-occur

**Traffic Incident Detection Analysis with Social Media Summarization.** Fu et al. [11] propose a social media-based traffic status monitoring system (*Steds*). The system is initiated by a transportation-related keyword generation process. Then an association rules based iterative query expansion algorithm is applied to extract real-time transportation-related tweets for incident management purposes. The feasibility of summarizing the redundant tweets to generate concise and comprehensible textual contents is confirmed.

**Traffic prediction in a bike-sharing system.** Li et al. [25] propose a hierarchical prediction model to predict the number of bikes that will be rented from/returned to each station cluster in a future period so that reallocation can be executed in advance. A bipartite clustering algorithm is proposed to cluster bike stations into groups, formulating a two-level hierarchy of stations. The total number of bikes that will be rented in a city is predicted by a Gradient Boosting Regression Tree (GBRT). A multi-similarity-based inference model is proposed to predict the rent proportion across clusters, and the inter-cluster transition, based on the number of bikes rented from/ returned to each cluster, can be easily inferred.

### 3.2 Ethics Issues for Transportation Predictions

Various types of urban sensors are deployed for recording transportation-related information such as vehicle travel speed, road occupancy, and volumes. The mobility of the commuters is monitored by multiple types of techniques such as loops, microwave, acoustic, video, and crowdsourcing such as Twitter and Waze. Collecting seemingly anonymous data that appear hard to identify an individual may in fact be problematic in many ways.

**1) The risk of compromising the privacy of an individual.** In the research [23], the MIT group combined two anonymized datasets of people in Singapore, one of mobile phone logs and the other of transit trip, each including “geolocation stamps” with time and place of each point. They succeeded to match up 17% of the users in one week and 11 weeks to get 95%. if GPS data from smartphones were added, it took less than a week to reach that number. The group acted as ethical or ‘white hat’ hackers to prove that someone acting in bad faith could do the same and compromise the privacy of many people.

**2) Governments’ surveillance on urban mobility.** Surveillance is infringing upon personal privacy. The times we live in, there’s a necessity for some surveillance in public spaces to maintain order, but as soon as we set into private spaces, it becomes an issue. The level of government surveillance varies from one government to another depending on the level of freedom practiced by the system. Most governments justify their unnecessary surveillance with words such as “National Threats”, “Security Stabilization”, or “Fighting Crime”. This will lead everyone (citizens) to feel that it is their duty to allow governments to collect their data even when it is not needed. As a result, the collected data from a government surveillance will have different impacts. Data can be collected from anywhere, GPS locations, car license plate scanners, phone calls, social media, closed-circuit television CCTV, and so on. Data can be used in “positive” and “negative” ways. This website presents what

happens when we are negatively being monitored by the government<sup>3</sup>. Surveillance in public spaces must be done without violating privacy. Some of the previous works in transportation analysis utilize GPS based data sources such as user check-in records and driver's travel paths. These works are vulnerable to violating privacy protection, which should be further addressed in the urban computing and spatial data mining fields.

**Vision:** The intelligent transportation system community can strengthen its ethics awareness by considering several additional aspects: 1) differential privacy to protect users' privacy, 2) users' awareness to share mobility data, 3) transparency between urban data collectors and users, and 4) comprehensive legislation. One of the advanced methods to preserve the privacy of users' information is implementing Differential privacy which was introduced by Dwork [7]. The method overlaps two areas: data analytics and statistics. Differential privacy uses techniques such as noise injection to keep the data of individual users completely private, and guarantees that the outcome prevents enabling anyone to learn anything about an individual. Meanwhile, it allows the researcher to query from the data and obtain good research results. By maintaining the balance between the implemented model and differential privacy, the privacy of an individual can be protected [18]. Users should have full awareness of the accessibility and availability of the mobility data generated by themselves. Therefore, the user should realize the potential risks of sharing their mobility data. Furthermore, the transparency between the data vendors and the users is not balanced: the vendors or institutions which collect and share data about people know a lot about the users but the users do not know as much about the data collectors. For that, users have the right to know how their data are used and have the right to withdraw from any experiment if asked. Finally, in the future developments of the urban computing industry, more comprehensive legislations should be proposed. Tighter regulations for collecting and sharing data should be established to prevent users' privacy from infringed such as General Data Protection Regulation known as GDPR which was applied in Europe.

## 4 Social Media based Event Detection and Analysis

Location-based service embedded social media sources such as Twitter and Foursquare have become popular data sources as surrogates for monitoring and detecting events. Targeted domains such as crime, election, and social unrest require the creation of algorithms capable of detecting events pertinent to these domains. Due to the unstructured language, short-length messages, dynamics, and heterogeneity typical of social media streams, it is technically difficult and labor-intensive to develop and maintain supervised learning systems. Recent studies in the research fields of event detection and analysis have applied both supervised and unsupervised learning methods to better modeling and forecasting the spatiotemporal events from social media data sources [36, 37, 20]. However, while utilizing the social media-based data sources, some of the ethical issues, such as built-in biases in dataset and algorithms, have not been paid with enough attention nowadays. In this section, we review some work on location-based social media analysis, state, and propose solutions to the ethical issues arising in the research fields of spatial data mining and social media analysis.

### 4.1 Current Works for Event Detection from Social Media

**Spatial Event Forecasting in Social Media.** Social media has become a significant surrogate for spatial event forecasting. The accuracy and discernibility of a spatial event forecasting model are two key concerns, which respectively determine how accurate and how detailed the model's predictions could be. Zhao et al. [37] propose a multi-resolution spatial event forecasting (*MREF*) model that concurrently addresses all the above challenges by formulating prediction tasks for different locations with different spatial resolutions, allowing the heterogeneous relationships among the tasks to be characterized.

**Cyber Attack Detection using Social Media.** Social media can also be viewed as a sensor into various societal events such as disease outbreaks, protests, and elections. Khandpur et al. [20] describe the use of social

---

<sup>3</sup><https://theyarewatching.org/>




Figure 4: Streamgraph showing normalized volume of tweets (September 2015 through October 2016) tagged with data breach (red), DDoS activity (grey) and account hijacking (blue) types of cyber-security events

media as a crowdsourced sensor to gain insight into ongoing cyber-attacks. The proposed social media-based approach detects a broad range of cyber-attacks in a weakly supervised manner using just a small set of seed event triggers and requires no training or labeled samples. A new query expansion strategy based on convolution kernels and dependency parses helps model semantic structure and aids in identifying key event characteristics. Figure 4 shows the wide range of events that the proposed system is able to detect. Notice the clear bursts in Twitter activity that the query expansion algorithm is able to detect.

**Real-Time Detection of Traffic Accidents from Twitter.** D'Andrea et al. [5] propose a real-time monitoring system for traffic event detection from Twitter stream analysis. The system fetches tweets from Twitter according to several search criteria, processes tweets, by applying text mining techniques, and finally performs the classification of tweets. The proposed SVM-based system aims to assign the appropriate class label to each tweet, as related to a traffic event or not. The traffic accident detection system is employed for real-time monitoring of several areas of the road networks, allowing for detection of traffic events almost in real-time, often before online traffic news sites.

## 4.2 Ethics Issues for Event Detection from Social Media

For the work on social media based event detection and analysis, several concerns may contribute to the ethical issues: **1) the built-in biases in social media data.** Bias in social media can happen when a sample is collected in such a way that some members of the intended population are not equally represented in the sample, resulting in a non-random sample. If such an error occurs in sampling, the results of the research could be mistakenly attributed to the study phenomenon instead of the sampling method [2]. In addition, due to the way humans generate their opinions, it is argued that misinformation in social media data can be produced in several ways, purposefully or accidentally. In a study experimented on Facebook's News Feed [16, 24], the group demonstrated how negative or positive posts disseminate by reducing positive or negative expressions. One case is that emotional contagion can play a massive role in generating biases in the social media platform, and this is demonstrated by the 2016 US elections when emotional contagion is used to foster negative viewpoints against specific politicians [35]. If such biases exist in Social media data of Urban planning projects, discriminatory results may occur.

**2) the consequences of flaws in data being manipulated by the users.** Relying on data collected from users in social media in order to use the model to recommend safer routes or detect threats is vulnerable to manipulation by malicious users. Recommendation systems, a similar system to the mentioned works, have been under attack since their appearance. An attack is carried out when fake users or certain legitimate malicious

users feed content to the system in a certain way that makes the system frequently favors a recommendation over what would the system should result [29, 13]. Leveraging human's tendency to use their peer's opinions to make their own decisions, large companies tend to use fake accounts and bots on social media for product promotion, posting positive comments, and silencing critics [28].

**3) Social media privacy concerns.** Social media is a beneficial source for real-time and historical data which makes it a great candidate to be surveilled positively and negatively. Social media has helped law enforcement to detect human trafficking activities, solve murder cases and other criminal activities [27]. In a negative way, in some governments where censoring is more strict than others. However, those governments do not censor many social media sources. Social media is where people tend to speak out about their local problems expecting solutions from the government. Yet, the government has escalated their efforts to monitor and suppress [32]. In a recent report published by the Brennan Center, Social Media Monitoring, there is proof that Department of Homeland Security DHS is exploiting personal data mined from social media to surveil protestors, religious and ethnic minorities [31]. The surveillance has expanded from positive (national security) to discriminatory surveillance.

**Vision:** While using social media data, researchers do need to be very careful and watchful for these issues, i.e., ensuring that the data's source and gathering methodologies mitigate bias, the cleaned data did not incidentally impose bias, and the data is used in a way that does not use or impose any bias. Monitored social media chatter could be used to inform law enforcement personnel of large potentially violent or non-violent protest events or other sociopolitical gatherings that might require a police presence to prevent untoward incidents. Such alerts, if provided to police departments with enough lead time, will allow for better planning of deployment logistics ensuring the safety of the officers and the civilians they are deployed to protect. EMBERS is once such an anticipatory intelligence system that forecasts population-level events in multiple countries in Latin America and provides not only the counts of the number of protests but also other details pertaining to the date, time and location of the potential protest event.

## 5 Conclusion

The rapid growth of the urban data for urban computing and spatial data mining methods raises new challenges along with new opportunities for various application fields such as urban safety analysis, intelligent transportation systems, and event detection. However, unfortunately, the ethical issues while obtaining, utilizing, and inferring from such enormous urban data are rarely addressed by the current researchers in urban computing communities. This paper reviews the most popular research branches of urban computing and spatial data mining, including urban safety analysis, intelligent transportation systems, and event detection. Ethical vulnerabilities for the existing urban computing and spatial data mining works such as predictive policing, biased data sources, compromised privacy, surveillance, and discrimination are revealed by this paper. Promising potential solutions and prospective visions towards such identified ethical vulnerabilities are directed. We offer these discussions of ethics in urban computing with the hope that they contribute highlight attention from the urban computing communities.

## References

- [1] S. Barocas and A. D. Selbst. Big data's disparate impact. *Calif. L. Rev.*, 104:671, 2016.
- [2] C. Cortes, M. Mohri, M. Riley, and A. Rostamizadeh. Sample selection bias correction theory. In *International conference on algorithmic learning theory*, pages 38–53. Springer, 2008.
- [3] J. Cranshaw. Whose city of tomorrow is it?: on urban computing, utopianism, and ethics. In *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing*, page 17. ACM, 2013.

- [4] K. Crawford. The hidden biases in big data. *Harvard Business Review*, 1:1, 2013.
- [5] E. D’Andrea, P. Ducange, B. Lazzerini, and F. Marcelloni. Real-time detection of traffic from twitter stream analysis. *IEEE transactions on intelligent transportation systems*, 16(4):2269–2283, 2015.
- [6] M. De Nadai, R. L. Vieriu, G. Zen, S. Dragicevic, N. Naik, M. Caraviello, C. A. Hidalgo, N. Sebe, and B. Lepri. Are safer looking neighborhoods more lively?: A multimodal investigation into urban life. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 1127–1135. ACM, 2016.
- [7] C. Dwork. Differential privacy. *Encyclopedia of Cryptography and Security*, pages 338–340, 2011.
- [8] A. G. Ferguson. Policing predictive policing. *Wash. UL Rev.*, 94:1109, 2016.
- [9] K. Fu, Z. Chen, and C.-T. Lu. Streetnet: preference learning with convolutional neural network on urban crime perception. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 269–278. ACM, 2018.
- [10] K. Fu, T. Ji, L. Zhao, and C.-T. Lu. Titan: A spatiotemporal feature learning framework for traffic incident duration prediction. In *Proceedings of the 27nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2019.
- [11] K. Fu, C.-T. Lu, R. Nune, and J. X. Tao. Steds: Social media based transportation event detection with text summarization. In *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pages 1952–1957. IEEE, 2015.
- [12] K. Fu, Y.-C. Lu, and C.-T. Lu. Treads: A safe route recommender using social media mining and text summarization. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 557–560. ACM, 2014.
- [13] I. Gunes, C. Kaleli, A. Bilge, and H. Polat. Shilling attacks against recommender systems: a comprehensive survey. *Artificial Intelligence Review*, 42(4):767–799, 2014.
- [14] S. Hajian, F. Bonchi, and C. Castillo. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 2125–2126. ACM, 2016.
- [15] D. Heaven. Why deep-learning ais are so easy to fool., 2019.
- [16] M. Ienca and E. Vayena. Cambridge analytica and online manipulation. *Scientific American*, 30, 2018.
- [17] T. Ji, K. Fu, N. Self, C.-T. Lu, and N. Ramakrishnan. Multi-task learning for transit service disruption detection. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 634–641. IEEE, 2018.
- [18] Z. Ji, Z. C. Lipton, and C. Elkan. Differential privacy and machine learning: a survey and review. *arXiv preprint arXiv:1412.7584*, 2014.
- [19] F. Kamiran and T. Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- [20] R. P. Khandpur, T. Ji, S. Jan, G. Wang, C.-T. Lu, and N. Ramakrishnan. Crowdsourcing cybersecurity: Cyber attack detection using social media. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 1049–1057. ACM, 2017.

- [21] R. P. Khandpur, T. Ji, Y. Ning, L. Zhao, C.-T. Lu, E. R. Smith, C. Adams, and N. Ramakrishnan. Determining relative airport threats from news and social media. In *Twenty-Ninth IAAI Conference*, 2017.
- [22] J. M. Kleinberg. Challenges in mining social network data: processes, privacy, and paradoxes. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 4–5. ACM, 2007.
- [23] D. Kondor, B. Hashemian, Y.-A. de Montjoye, and C. Ratti. Towards matching user mobility traces in large-scale datasets. *IEEE Transactions on Big Data*, 2018.
- [24] A. D. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [25] Y. Li, Y. Zheng, H. Zhang, and L. Chen. Traffic prediction in a bike-sharing system. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, page 33. ACM, 2015.
- [26] X. Liu, Q. Chen, L. Zhu, Y. Xu, and L. Lin. Place-centric visual urban perception with deep multi-instance regression. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 19–27. ACM, 2017.
- [27] A. Mateescu, D. Brunton, A. Rosenblat, D. Patton, Z. Gold, and D. Boyd. Social media surveillance and law enforcement. *Data Civ Rights*, 27:2015–2027, 2015.
- [28] J. McGregor. Reddit is being manipulated by big financial services companies. *Forbes Review*, 1:2, 2017.
- [29] B. Mobasher, R. Burke, R. Bhaumik, and C. Williams. Toward trustworthy recommender systems: An analysis of attack models and algorithm robustness. *ACM Transactions on Internet Technology (TOIT)*, 7(4):23, 2007.
- [30] L. Nordling. Mind the gap, 2019.
- [31] F. Patel, R. Levinson-Waldman, S. DenUyl, and R. Koreh. Social media monitoring. 2019.
- [32] B. Qin, D. Strömberg, and Y. Wu. Why does china allow freer social media? protests versus surveillance and propaganda. *Journal of Economic Perspectives*, 31(1):117–40, 2017.
- [33] S. Shah, F. Bao, C.-T. Lu, and I.-R. Chen. Crowdsafe: crowd sourcing of crime incidents and safe routing on mobile devices. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 521–524. ACM, 2011.
- [34] Y.-J. Wu, F. Chen, C.-T. Lu, and S. Yang. Urban traffic flow prediction using a spatio-temporal random effects model. *Journal of Intelligent Transportation Systems*, 20(3):282–293, 2016.
- [35] U. Yaqub, S. Chun, V. Atluri, and J. Vaidya. Sentiment based analysis of tweets during the us presidential elections. In *Proceedings of the 18th annual international conference on digital government research*, pages 1–10. ACM, 2017.
- [36] L. Zhao, F. Chen, J. Dai, T. Hua, C.-T. Lu, and N. Ramakrishnan. Unsupervised spatial event detection in targeted domains with applications to civil unrest modeling. *PloS one*, 9(10):e110206, 2014.
- [37] L. Zhao, F. Chen, C.-T. Lu, and N. Ramakrishnan. Multi-resolution spatial event forecasting in social media. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 689–698. IEEE, 2016.