Spatio-temporal Event Forecasting and Precursor Identification

Yue Ning
Stevens Institute of Technology
Hoboken, NJ
yue.ning@stevens.edu

Liang Zhao
George Mason University
Fairfax, VA
lzhao9@gmu.edu

Feng Chen
University at Albany – SUNY
Albany, NY
fchen5@albany.edu

Chang-Tien Lu
Virginia Tech
Falls Church, VA
ctlu@vt.edu

Huzefa Rangwala
George Mason University
Fairfax, VA
rangwala@gmu.edu

ABSTRACT

Spatio-temporal societal event forecasting, which has traditionally been prohibitively challenging, is now becoming possible and experiencing rapid growth thanks to the big data from Open Source Indicators (OSI) such as social media, news sources, blogs, economic indicators, and other meta-data sources. Spatio-temporal societal event forecasting and their precursor discovery benefit the society by providing insight into events such as political crises, humanitarian crises, mass violence, riots, mass migrations, disease outbreaks, economic instability, resource shortages, natural disasters, and others. In contrast to traditional event detection that identifies ongoing events, event forecasting focuses on predicting future events yet to happen. Also different from traditional spatio-temporal predictions on numerical indices, spatio-temporal event forecasting needs to leverage the heterogeneous information from OSI to discover the predictive indicators and mappings to future societal events. While studying large scale societal events, policy makers and practitioners aim to identify precursors to such events to help understand causative attributes and ensure accountability. The resulting problems typically require the predictive modeling techniques that can jointly handle semantic, temporal, and spatial information, and require a design of efficient and interpretable algorithms that scale to high-dimensional large real-world datasets.

In this tutorial, we will present a comprehensive review of the state-of-the-art methods for spatio-temporal societal event forecasting. First, we will categorize the inputs OSI and the predicted societal events commonly researched in the literature. Then we will review methods for temporal and spatio-temporal societal event forecasting. Next, we will also discuss the foundations of precursor identification with an introduction to various machine learning approaches that aim to discover precursors while forecasting events. Through the tutorial, we expect to illustrate the basic theoretical and algorithmic ideas and discuss specific applications in all the above settings.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

ACM ISBN 978-1-4503-6201-6/19/08.
https://doi.org/10.1145/3292500.3332291

CCS CONCEPTS

• **Information systems → Spatial-temporal systems**; **Data analytics**.

KEYWORDS

Spatio-Temporal Modeling; Precursor Discovery

ACM Reference Format:
Yue Ning, Liang Zhao, Feng Chen, Chang-Tien Lu, and Huzefa Rangwala. 2019. Spatio-temporal Event Forecasting and Precursor Identification. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19), August 4–8, 2019, Anchorage, AK, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3292500.3332291

1 INTRODUCTION

Social events are often interwoven with various factors such as economics, politics, and public policies of our society. Nowadays more and more online data sources, such as social networks, official news websites, and personal blogs, chronicle social events. This strongly motivates the need to utilize open source data to model events for social science, public health care, and decision making. In recent years, a number of causality-learning-based, temporal-mining-based, and spatial-mining-based methods have been developed to address temporal and spatial event forecasting in different scenarios in both the data mining and machine learning communities. The developed techniques originate from several related areas, including statistical causal learning, spatial and temporal dependency mining, sparse feature learning, and predictive subgraph detection. By a comprehensive review of all these methods in this tutorial, the audience can gain a more complete understanding of this problem setting, which can then inspire new research within the community and bring practical value to practitioners.

While studying large scale societal events, policy makers and professionals aim to identify precursors to such events to help understand causative attributes. Given a document reporting an event of interest (e.g., a civil protest), precursors are other documents published earlier than reported incidents and can be viewed as influential in the lead up to the event. Predictive analysis for precursor identification is an interdisciplinary subject which brings data scientists and domain specialists addressing key challenges: i) obtaining precise annotations of precursors is infeasible in practice; ii) data and feature sparsity among different geolocations lead to tepid performance of several machine learning models; iii) feature selection from large-scale heterogeneous data sources is time consuming;

iv) structure learning for precursor representations becomes critical for social analysts while little research has been conducted for capturing semantic relationships or hidden contextual patterns for event precursors.

Workshops for predictive analytics have been organized in the past to educate participants regarding various challenges in event detection/prediction associated with finance, healthcare, and politics. This tutorial focuses on machine learning solutions to spatial-temporal event forecasting and precursor identification. Given a set of spatio-temporal training data with event annotations, the topics of this tutorial include: 1) methodologies of predicting the occurrence of target events at a location in a future time window from historical data streams; 2) highlights of recent progress of new promising domains on spatial and temporal events; 3) identification of precursors or prognostics for event analysis and decision making.

2 TUTORIAL OUTLINE

In this tutorial, we will discuss the foundations of event prediction and precursor identification. We will proceed with a discussion of various machine learning approaches that aim to forecast events along with identifying precursors. These approaches extend classical learning approaches within frameworks of causal dependency, temporal dependency, unsupervised learning, discriminative learning, generative models, multi-instance learning, multi-task learning, and deep learning. Finally we will discuss a few broad applications: (i) civil unrest movement detection, (ii) public health epidemic identification, (iii) crime events, and (iv) anomaly precursor detection and showcase a demo with case study for future use.

In the first part, we will provide an overview of event prediction and precursor identification problems. We will discuss different domains of events and open source data sources for modeling and analyzing events. We will also introduce the challenges in this area compared with event detection and spatial prediction.

In the second part, we will focus on temporal event forecasting which includes cauasal dependency mining (predefined causality[7, 17] and optimized causality [1, 12]) and temporal dependency mining (markov decision processes [11, 15] and deep neural networks [5, 6, 10, 18]).

In the third part, we discuss spatio-temporal event forecasting using unsupervised learning (distance based models [14] and scan statistics-based models [2, 16]), discriminative learning (spatial heterogeneity and dependency models[4, 19, 24, 25] and spatial multi-level models [22, 26]), generative models [20, 21], mechanistic-based models[23], and ensemble models [13].

In the fourth part, we will talk about precursor identification with multi-instance learning [8], multi-task learning [9], and deep neural networks [3]. In this tutorial, we will also discuss sequential models such as Recurrent Neural Networks (RNNs) on time series prediction problems and how the model generates predictions for anomaly events in different applications.

At last, we will discuss real world applications and future directions.

ACKNOWLEDGMENTS

We thank Naren Ramakrishnan for his help and contribution. This work was partially supported by NSF Grants No. 1755850 and No. 1841520.

REFERENCES

[1] Marta Arias, Argimiro Arratia, and Ramon Xuriguera. 2013. Forecasting with Twitter data. *TIST* 5, 1 (2013), 8.
[2] Feng Chen and Daniel B Neill. 2014. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *KDD 2014*. ACM, 1166–1175.
[3] Songgaojun Deng, Huzefa Rangwala, and Yue Ning. 2019. Learning Dynamic Context Graphs for Predicting Social Events. In *KDD (KDD 2019)*. ACM.
[4] Yuyang Gao and Liang Zhao. 2018. Incomplete Label Multi-Task Ordinal Regression for Spatial Event Scale Forecasting. In *AAAI Conference on Artificial Intelligence*. 2999–3006.
[5] Mark Granroth-Wilding and Stephen Clark. 2016. What happens next? event prediction using a compositional neural network model. In *AAAI Conference on Artificial Intelligence*.
[6] Linmei Hu, Juanzi Li, Liqiang Nie, Xiao-Li Li, and Chao Shao. 2017. What Happens Next? Future Subevent Prediction Using Contextual Hierarchical LSTM. In *AAAI Conference on Artificial Intelligence*.
[7] Sathappan Muthiah, Patrick Butler, Rupinder Paul Khandpur, and et al. 2016. EMBERS at 4 Years: Experiences Operating an Open Source Indicators Forecasting System. In *KDD*. ACM, New York, NY, USA, 205–214.
[8] Yue Ning, Sathappan Muthiah, Huzefa Rangwala, and Naren Ramakrishnan. 2016. Modeling Precursors for Event Forecasting via Nested Multi-Instance Learning. In *KDD*. ACM, 1095–1104.
[9] Yue Ning, Rongrong Tao, Chandan K. Reddy, Huzefa Rangwala, James C. Starz, and Naren Ramakrishnan. 2018. STAPLE: Spatio-Temporal Precursor Learning for Event Forecasting. In *SDM*. 99–107.
[10] Karl Pichotta and Raymond J Mooney. 2016. Learning Statistical Scripts with LSTM Recurrent Neural Networks.. In *AAAI*. 2800–2806.
[11] Fengcai Qiao, Pei Li, Xin Zhang, Zhaoyun Ding, Jiajun Cheng, and Hui Wang. 2017. Predicting Social Unrest Events with Hidden Markov Models Using GDELT. *Discrete Dynamics in Nature and Society* 2017 (2017).
[12] Kira Radinsky and Sagie Davidovich. 2012. Learning to predict from textual data. *Journal of Artificial Intelligence Research* 45, 1 (2012), 641–684.
[13] Naren Ramakrishnan, Patrick Butler, Sathappan Muthiah, Nathan Self, Rupinder Khandpur, Parang Saraf, Wei Wang, Jose Cadena, Anil Vullikanti, Gizem Korkmaz, et al. 2014. 'Beating the news' with EMBERS: forecasting civil unrest using open source indicators. In *KDD*. ACM, 1799–1808.
[14] Polina Rozenshtein, Aris Anagnostopoulos, Aristides Gionis, and Nikolaj Tatti. 2014. Event detection in activity networks. In *KDD*. ACM, 1176–1185.
[15] Philip A Schrodt. 2006. Forecasting conflict in the Balkans using hidden Markov models. In *Programming for Peace*. Springer, 161–184.
[16] Minglai Shao, Jianxin Li, Feng Chen, Hongyi Huang, Shuai Zhang, and Xunxun Chen. 2017. An Efficient Approach to Event Detection and Forecasting in Dynamic Multivariate Social Media Networks. In *WWW*. International World Wide Web Conferences Steering Committee, 1631–1639.
[17] Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welpe. 2010. Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. *ICWSM* 10 (2010), 178–185.
[18] Zhongqing Wang and Yue Zhang. 2017. DDoS event forecasting using Twitter data. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. AAAI Press, 4151–4157.
[19] Zhuoning Yuan, Xun Zhou, and Tianbao Yang. 2018. Hetero-ConvLSTM: A deep learning approach to traffic accident prediction on heterogeneous spatio-temporal data. In *KDD*. ACM, 984–992.
[20] Qian Zhang, Nicola Perra, Daniela Perrotta, Michele Tizzoni, Daniela Paolotti, and Alessandro Vespignani. 2017. Forecasting seasonal influenza fusing digital indicators and a mechanistic disease model. In *WWW*. International World Wide Web Conferences Steering Committee, 311–319.
[21] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Spatiotemporal Event Forecasting in Social Media. In *Proceedings of the 2015 SIAM international conference on data mining (SDM)*. SIAM, 963–971.
[22] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Multi-resolution Spatial Event Forecasting in Social Media. In *ICDM*. IEEE, 689–698.
[23] Liang Zhao, Jiangzhuo Chen, Feng Chen, Wei Wang, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Simnest: Social media nested epidemic simulation via online semi-supervised deep learning. In *Data Mining (ICDM), 2015 IEEE International Conference on*. IEEE, 639–648.
[24] Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Multi-task learning for spatio-temporal event forecasting. In *KDD*. ACM, 1503–1512.
[25] Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2017. Feature Constrained Multi-Task Learning Models for Spatiotemporal Event Forecasting. *IEEE Transactions on Knowledge and Data Engineering* 29, 5 (2017), 1059–1072.
[26] Liang Zhao, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Hierarchical Incomplete Multi-source Feature Learning for Spatiotemporal Event Forecasting. In *KDD*. ACM, New York, NY, USA, 2085–2094.