Event Prediction in the Big Data Era: A Systematic Survey

LIANG ZHAO, Emory University

Events are occurrences in specific locations, time, and semantics that nontrivially impact either our society or the nature, such as earthquakes, civil unrest, system failures, pandemics, and crimes. It is highly desirable to be able to anticipate the occurrence of such events in advance to reduce the potential social upheaval and damage caused. Event prediction, which has traditionally been prohibitively challenging, is now becoming a viable option in the big data era and is thus experiencing rapid growth, also thanks to advances in high performance computers and new Artificial Intelligence techniques. There is a large amount of existing work that focuses on addressing the challenges involved, including heterogeneous multi-faceted outputs, complex (e.g., spatial, temporal, and semantic) dependencies, and streaming data feeds. Due to the strong interdisciplinary nature of event prediction problems, most existing event prediction methods were initially designed to deal with specific application domains, though the techniques and evaluation procedures utilized are usually generalizable across different domains. However, it is imperative yet difficult to cross-reference the techniques across different domains, given the absence of a comprehensive literature survey for event prediction. This article aims to provide a systematic and comprehensive survey of the technologies, applications, and evaluations of event prediction in the big data era. First, systematic categorization and summary of existing techniques are presented, which facilitate domain experts' searches for suitable techniques and help model developers consolidate their research at the frontiers. Then, comprehensive categorization and summary of major application domains are provided to introduce wider applications to model developers to help them expand the impacts of their research. Evaluation metrics and procedures are summarized and standardized to unify the understanding of model performance among stakeholders, model developers, and domain experts in various application domains. Finally, open problems and future directions are discussed. Additional resources related to event prediction are included in the paper website: http://cs.emory.edu/~lzhao41/projects/event_prediction_site.html.

CCS Concepts: • **Information systems** → **Data mining**; • **Computing methodologies** → **Machine learning**;

Additional Key Words and Phrases: Event prediction, big data, artificial intelligence

ACM Reference format:

Liang Zhao. 2021. Event Prediction in the Big Data Era: A Systematic Survey. *ACM Comput. Surv.* 54, 5, Article 94 (May 2021), 37 pages.
<https://doi.org/10.1145/3450287>

1 INTRODUCTION

An event is a real-world occurrence that takes place in a specific location and time that relates to a particular topic. Events can range from large-scale (e.g., civil unrest events and earthquakes), to medium-scale (e.g., system failures and crime incidents), to small-scale (e.g., authentication events and individual actions) occurrences [10]. Event analytics are important in domains as different as healthcare, business, cybersphere, politics, and entertainment, influencing almost every corner of

Authors' address: L. Zhao, Emory University, 201 Dowman Dr, Atlanta, GA 30322; email: liang.zhao@emory.edu.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

0360-0300/2021/05-ART94

<https://doi.org/10.1145/3450287>

our lives [187]. The analysis of events has thus been attracting huge attention over the past few decades and can be categorized in terms of their timeliness for various research directions, such as event summarization, detection, and prediction. Unlike retrospective analyses such as event summarization and detection [30], event prediction focuses on anticipating events in the future and is the focus of this survey. Accurate anticipation of future events enables one to maximize the benefits and minimize the losses associated with some event in the future, bringing huge benefits for both society as a whole and individual members of society in key domains such as disease prevention, disaster management, business intelligence, and economics stability.

“Prediction is very difficult, especially if it’s about the future.” —Niels Bohr, 1970

Event prediction has traditionally been prohibitively challenging across different domains, due to the lack or incompleteness of our knowledge regarding the true causes and mechanisms driving event occurrences in most domains. With the advent of the big data era, however, we now enjoy unprecedented opportunities that open up many alternative approaches for dealing with event prediction problems, sidestepping the need to develop a complete understanding of the underlying mechanisms of event occurrence. Based on large amounts of data on historical events and their potential precursors, event prediction methods typically strive to apply predictive mapping to build on these observations to predict future events, utilizing predictive analysis techniques from domains such as machine learning, data mining, pattern recognition, statistics, and other computational models [20]. Event prediction is currently experiencing extremely rapid growth, thanks to advances in sensing techniques (physical sensors and social sensors), prediction techniques (Artificial Intelligence, especially Machine Learning), and high performance computing hardware [63].

Event prediction in big data is a difficult problem that requires the invention and integration of related techniques to address the serious challenges caused by its unique characteristics, including: **(1) Heterogeneous multi-output predictions.** Event prediction methods usually need to predict multiple facets of events including their time, location, topic, intensity, and duration, each of which may utilize a different data structure [144]. This creates unique challenges, including how to jointly predict these heterogeneous yet correlated facets of outputs. Due to the rich information in the outputs, label preparation is usually a highly labor-intensive task performed by human annotators, with automatic methods introducing numerous errors in items such as event coding. So, how can we improve the label quality as well as the model robustness under corrupted labels? The multi-faceted nature of events make event prediction a multi-objective problem, which raises the question of how to properly unify the prediction performance on different facets. It is also challenging to verify whether a predicted event “matches” a real event, given that the various facets are seldom, if ever, 100% accurately predicted. So, how can we set up the criteria needed to discriminate between a correct prediction (“true positive”) and a wrong one (“false positive”)?

(2) Complex dependencies among the prediction outputs. Beyond conventional isolated tasks in machine learning and predictive analysis, in event prediction the predicted events can correlate to and influence each other [121]. For example, an ongoing traffic incident event could cause congestion on the current road segment in the first 5 min but then lead to congestion on other contiguous road segments 10 min later. Global climate data might indicate a drought in one location, which could then cause famine in the area and lead to a mass exodus of refugees moving to another location. So, how should we consider the correlations among future events?

(3) Real-time stream of prediction tasks. Event prediction usually requires continuous monitoring of the observed input data to trigger timely alerts of future potential events [155]. However, during this process the trained prediction model gradually becomes outdated, as real world events continually change dynamically, concepts are fluid and distribution drifts are inevitable. For example, in September 2008 21% of the United States population were social media users, including 2% of

those over 65. However, by May 2018, 72% of the United States population were social media users, including 40% of those over [29]. Not only the data distribution but also the number of features and input data sources can also vary in real time. Hence, it is imperative to periodically upgrade the models, which raises further questions concerning how to train models based on non-stationary distributions, while balancing the cost (such as computation cost and data annotation cost) and timeliness? **(4) Challenges in the event big data.** The above-mentioned challenges and nature in event prediction tasks further cause more difficulties in collecting and utilizing event data. First, noisiness is a very common issue in event data. There are three major types including missing (e.g., due to device outage), erroneous (e.g., caused by malfunction in physical sensor and typos in social media), and adversarial data (e.g., rumors). Second, the high complexity of event occurrence mechanism encourages researchers to use multiple data sources to better cover all the indicators. However, different data sources can be highly heterogeneous, coming with different modalities, spatial resolutions, qualities, and temporal duration, leading to extreme difficulties in an effective fusion of them. Moreover, many events are rare, such as terrorism, natural disasters, and system failures, leading to an imbalance between the positive (event occurrence) and negative (no occurrence) data samples, which are well-known to be highly challenging especially classic supervised learning techniques.

In recent years, a considerable amount of research has been devoted to event prediction technique development and applications, to address the aforementioned challenges. Recently, there has been a surge of research that both proposes and applies new approaches in numerous domains, though event prediction techniques are generally still in their infancy. Most existing event prediction methods have been designed for a specific application domains, but their approaches are usually general enough to handle problems in other application domains. Unfortunately, it is difficult to cross-reference these techniques across different application domains serving totally different communities. Moreover, the quality of event prediction results require sophisticated and specially-designed evaluation strategies due to the subject matter's unique characteristics, for example its multi-objective nature (e.g., accuracy, resolution, efficiency, and lead time) and heterogeneous prediction results (e.g., heterogeneity and multi-output). As yet, however, we lack a systematic standardization and comprehensive summarization approaches with which to evaluate the various event prediction methodologies that have been proposed. This absence of a systematic summary and taxonomy of existing techniques and applications in event prediction causes major problems for those working in the field who lacks clear information on the existing bottlenecks, traps, open problems, and potentially fruitful future research directions.

To overcome these hurdles and facilitate the development of better event prediction methodologies and applications, this survey article aims to provide a comprehensive and systematic review of the current state of the art for event prediction in the big data era. The major contributions include:

- **A systematic categorization and summarization of existing techniques.** Existing event prediction methods are categorized according to their event aspects (time, location, and semantics), problem formulation, and corresponding techniques to create the taxonomy of a generic framework. Relationships, advantages, and disadvantages among different sub-categories are discussed, along with details of the techniques under each subcategory. The proposed taxonomy could help domain experts locate the most useful techniques for their targeted problem settings.
- **A comprehensive categorization and summarization of major application domains.** The first taxonomy of event prediction application domains is provided. The practical significance and problem formulation are elucidated for each application domain or

subdomain, enabling it to be easily mapped to the proposed technique taxonomy. This will help data scientists and model developers to search for additional application domains and datasets to evaluate their proposed methods, and expand their advanced techniques to encompass new application domains.

- **Standardized evaluation metrics and procedures.** Due to the nontrivial structure of event prediction outputs, which can contain multiple fields such as time, location, intensity, duration, and topic, this article proposes a set of standard metrics with which to standardize existing ways to pair predicted events with true events. Then additional metrics are introduced and standardized to evaluate how close the predicted events are to the real ones.
- **An insightful discussion of the current status of research in this area and future trends.** Based on the comprehensive and systematic survey and investigation of existing event prediction techniques and applications presented here, an overall picture and the shape of the current research frontiers are outlined. The article concludes by presenting fresh insights into the bottleneck, traps, and open problems, as well as a discussion of possible future directions.

1.1 Related Surveys

This section briefly outlines previous surveys in various domains that have some relevance to event prediction in big data in three categories, namely, (1) event detection, (2) predictive analytics, and (3) domain-specific event prediction.

Event detection has been an extensively explored domain with over many years. Its main purpose is to detect historical or ongoing events rather than to predict as yet unseen events in the future [154]. Event detection typically focuses on pattern recognition [20], anomaly detection [75], change footprint pattern discovery [216], and clustering [75], which are very different from those in event prediction. There have been several surveys of research in this domain in the last decade [9, 13, 52, 125]. For example, Deng et al. [52] and Atefeh and Khreich [13] provided overviews of event extraction techniques in social media, while Michelioudakis et al. [125] presented a survey of event recognition with uncertainty. Alevizos et al. [9] provided a review of event recognition methods using probabilistic methods.

Predictive analysis covers the prediction of target variables given a set of dependent variables. These target variables are typically homogeneous scalar or vector data for describing items such as economic indices, housing prices, or sentiments. The target variables may not necessarily be values in the future. Larose [103] provides a good tutorial and survey for this domain. Predictive analysis can be broken down into subdomains such as structured prediction [20], spatial prediction [88], and sequence prediction [71], enabling users to handle different types of structure for the target variable. Fülöp et al. [65] provided a survey and categorization of applications that utilize predictive analytics techniques to perform event processing and detection, while Jiang [88] focused on spatial prediction methods that predict the indices that have spatial dependency. Baklr et al. [14] summarized the literature on predicting structural data such as geometric objects and networks, and Arias et al. [12] Phillips et al. [138], and Yu and Kak [196] all proposed for predictive analysis using social data.

As event prediction methods are typically motivated by specific application domains, there are a number of surveys event predictions for domains such as flood events [37], social unrest [22], wind power ramp forecasting [61], tornado events [55], temporal events without location information [70], online failures [155], and business failures [6]. However, in spite of its promise and its rapid growth in recent years, the domain of event prediction in big data still suffers from the lack of a comprehensive and systematic literature survey covering all its various aspects, including relevant techniques, applications, evaluations, and open problems.

Fig. 1. Various types of times in event prediction.

1.2 Outline

The remainder of this article is organized as follows. Section 2 presents generic problem formulations for event prediction and the evaluation of event prediction results. Section 3 then presents a taxonomy and comprehensive description of event prediction techniques, after which Section 4 categorizes and summarizes the various applications of event prediction. Section 5 lists the open problems and suggests future research directions and this survey concludes in Section 6.

2 PROBLEM FORMULATION AND PERFORMANCE EVALUATIONS

This section begins by examining the generic denotation and formulation of the event prediction problem (Section 2.1) and considers way to standardize event prediction evaluations (Section 2.2).

2.1 Problem Formulation

An event refers to a real-world occurrence that happens at some specific time and location with specific semantic topic [187]. We can use $y = (t, l, s)$ to denote an event where its time $t \in \mathcal{T}$, its location $l \in \mathcal{L}$, and its semantic meaning $s \in \mathcal{S}$. Here, $\mathcal{T}$, $\mathcal{L}$, and $\mathcal{S}$ represent the time domain, location domain, and semantic domain, respectively. Notice that these domains need to have very general meanings that cover a wide range of types of entities. For example, the location $\mathcal{L}$ can include any features that can be used to locate the place of an event in terms of a point or a neighborhood in either Euclidean space (e.g., coordinate and geospatial region) or non-Euclidean space (e.g., a vertex or subgraph in a network). Similarly, the semantic domain $\mathcal{S}$ can contain any type of semantic features that are useful when elaborating the semantics of an event's various aspects, including its actors, objects, actions, magnitude, textual descriptions, and other profiling information. For example, ("11am, June 1, 2019," "Hermosillo, Sonora, Mexico," "Student Protests") and ("June 1, 2010," "Berlin, Germany," "Red Cross helps pandemics control") denote the time, location, and semantics, for two events, respectively. Lastly, there are different types of 'times' in event prediction problems as shown in Figure 1.

An event prediction system requires inputs that could indicate future events, called event indicators, and these could contain both critical information on events that precede the future event, known as precursors, as well as irrelevant information [144]. Event indicator data can be denoted as $X \subseteq \mathcal{T} \times \mathcal{L} \times \mathcal{F}$, where $\mathcal{F}$ is the domain of the features other than location and time. If we denote the current time as t_{now} and define the past time and future time as $\mathcal{T}^- \equiv \{t | t \leq t_{\text{now}}, t \in \mathcal{T}\}$ and $\mathcal{T}^+ \equiv \{t | t > t_{\text{now}}, t \in \mathcal{T}\}$, respectively, then the event prediction problem can be formulated as follows:

Definition 2.1 (Event Prediction). Given the event indicator data $X \subseteq \mathcal{T}^- \times \mathcal{L} \times \mathcal{F}$ and historical event data $Y_0 \subseteq \mathcal{T}^- \times \mathcal{L} \times \mathcal{S}$, event prediction is a process that outputs a set of predicted future events $\hat{Y} \subseteq \mathcal{T}^+ \times \mathcal{L} \times \mathcal{S}$, such that for each predicted future event $\hat{y} = (t, l, s) \in \hat{Y}$ where $t > t_{\text{now}}$.

Not every event prediction method necessarily focuses on predicting all three domains of time, location, and semantics simultaneously, but may instead predict any part of them. For example, when predicting a clinical event such as the recurrence of disease in a patient, the event location might not always be meaningful [141], but when predicting outbreaks of seasonal flu, the semantic meaning is already known and the focus is the location and time [21] and when predicting political

events, sometimes the location, time, and semantics (e.g., event type, participant population type, and event scale) are all necessary [144].

2.2 Event Prediction Evaluation

Event Prediction Evaluation essentially investigates the goodness of fit for a set of predicted events $\hat{Y}$ against real events Y . Unlike the outputs of conventional machine learning models such as the simple scalar values used to indicate class types in classification or numerical values in regression, the outputs of event prediction are entities with rich information. Before we evaluate the quality of prediction, we need to first determine the pairs of predictions and the labels that will be used for the comparison. Hence, we must first optimize the process of matching predictions and real events (Section 2.2.1) before evaluating the prediction error and accuracy (Section 2.2.2).

2.2.1 Matching Predicted Events and Real Events. The following types of matching are typically used:

- Prefixed matching.** The predicted events will be matched with the corresponding ground-true real events if they share some key attributes. For example, for event prediction at a particular time and location, we can evaluate the prediction against the ground truth for that time and location. This type of matching is most common when each of the prediction results can be uniquely distinguished along the predefined attributes (for example, location and time) that have a limited number of possible values, so that one-on-one matching between the predicted and real events are easily achieved [206]. For example, to evaluate the quality of a predicted event on June 1, 2019 in San Francisco, USA, the true event occurrence on that date and location can be used for the evaluation.

- Optimized matching.** In situations where one-on-one matching is not easily achieved for any event attribute, then the set of predicted events might need to assess the quality of the match achieved with the set of real events, via an optimized matching strategy [144]. For example, consider two predictions, **Prediction 1**: (“9am, June 4, 2019,” “Nogales, Sonora, Mexico,” “Worker Strike”), and **Prediction 2**: (“11am, June 1, 2019,” “Hermosillo, Sonora, Mexico,” “Student Protests”). The two ground truth events that these can usefully be compared with are **Real Event 1**: (“9am, June 1, 2019,” “Hermosillo, Sonora, Mexico,” “Teacher Protests”), and **Real Event 2**: (“June 4, 2019,” “Navojoa, Sonora, Mexico,” “General-population Protest”). None of the predictions are an exact match for any of the attributes of the real events, so we will need to find a “best” matching among them, which in this case is between **Prediction 1** and **Real Event 2** and **Prediction 2** and **Real Event 1**. This type of matching allows some degree of inaccuracy in the matching process by quantifying the distance between the predicted and real events among all the attribute dimensions. The distance metrics are typically either Euclidean distance [88] or some other distance metric [75]. Some researchers have hired referees to manually check the similarity of semantic meanings [142], but another way is to use event coding to code the events into an event type taxonomy and then consider a “match” if the event type matches [39]. Based on the distance between each pair of predicted and real events, the optimal matching will be the one that results in the smallest average distance [133]. However, suppose there are m predicted events and n real events, then there can be as many as $2^{m \cdot n}$ possible ways of matching, making it prohibitively difficult to find the optimal solution. Moreover, there could be different rules for matching. For example, the “multiple-to-multiple” rule shown in Figure 2(a) allows one predicted (real) event to match multiple real (predicted) events [143], while “Bipartite matching” only allows one-to-one matching between predicted and real events (Figure 2(b)). “Non-crossing matching” requires that the real events matched by the predicted events follow the same chronological order (Figure 2(c)). With specific matching types, researchers suggested using event matching optimization to learn the optimal set of “(real event, predicted event)” pairs [128].

Fig. 2. Generic framework for hierarchical RNN-based event forecasting.

2.2.2 Metrics of Effectiveness. The effectiveness of the event predictions is evaluated in terms of two indicators: (1) Goodness of Matching, which evaluates performance metrics such as the number and percentage of matched events [20], and (2) Quality of Matched Predictions, which evaluates how close the predicted event is to the real event for each pair of matched events [144].

•**Goodness of Matching.** A *true positive* means a real event has been successfully matched by a predicted event; if a real event has not been matched by any predicted event, then it is called a *false negative* and a *false positive* means a predicted event has failed to match any real event, which is referred to as a *false alarm*. Assume the total number of predictions is N , the number of real events is $\hat{N}$, the number of true positives is N_{TP} , the number of false negatives is N_{FN} and the number of false positives is N_{FP} . Then, the following key evaluation metrics can be calculated: $\text{Precision} = N_{TP} / (N_{TP} + N_{FP})$, $\text{Recall} = N_{TP} / (N_{TP} + N_{FN})$, $\text{F-measure} = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall})$. Other measurements such as the area under the ROC curves are also commonly used [20]. This approach can be extended to include other items such as multi-class precision/recall, and Precision/Recall at Top K [2, 86, 105, 213].

•**Quality of Matched Predictions.** If a predicted event matches a real one, then it is common to go on to evaluate how close they are. This reflects the quality of the matched predictions, in terms of different aspects of the events. Event time is typically a numerical values and hence can be easily measured in terms of metrics such as mean squared error, root mean squared error, and mean absolute error [20]. This is also the case for location in Euclidean space, which can be measured in terms of the Euclidean distance between the predicted point (or region) and the real point (or region). Some researchers consider the administrative unit resolution. For example, a predicted location (“New York City,” “New York State,” “USA”) has a distance of 2 from the real location (“Los Angeles,” “California,” “USA”) [203]. Others prefer to measure multi-resolution location prediction quality as follows: $(1/3)(l_{country} + l_{country} \cdot l_{state} + l_{country} \cdot l_{state} \cdot l_{city})$, where l_{city} , l_{state} , and $l_{country}$ can only be either 0 (i.e., no match to the truth) or 1 (i.e., completely matches the truth) [144]. For a location in non-Euclidean space such as a network [158], the quality can be measured in terms of the shortest path length between the predicted node (or subgraph) and the real node (or subgraph), or by the F-measure between the detected subsets of nodes against the real ones, which is similar to the approach for evaluating community detection [75]. For event topics, in addition to conventional ways of evaluating continuous values such as population size, ordinal values such as event scale, and categorical values such as event type, actors, and actions, as well as more complex semantic values such as texts, can be evaluated using Natural Language Process measurements such as edit distance, BLEU score, Top- K precision, and ROUGE [210].

3 EVENT PREDICTION TECHNIQUES

This section focuses on the taxonomy and representative techniques utilized for each category and subcategory. Due to the heterogeneity of the prediction output, the technique types depend on the type of output to be predicted, such as time, location, and semantics. As shown in Figure 3, all the event prediction methods are classified in terms of their goals, including time, location, semantics,

Fig. 3. Taxonomy of event prediction problems and techniques.

and the various combinations of these three. *Time prediction* focuses on the tasks where time is the most critical information of interest while location and semantics are either non-existent (e.g., in stock market events), not of interest, or already given. Similarly, *location prediction* concentrates on inferring the event position for specified time and event genres, such as the crime hotspot and air pollutant map for tomorrow. *Semantic prediction* focuses on the future event profiles such as actors, type, and scale where the event time and location can be either not of interest or trivially obtained from the semantics, such as news event prediction. The aforementioned goals are then further categorized in terms of the output forms of the goals the corresponding techniques normally used.

3.1 Time Prediction

Event time prediction focuses on predicting when future events will occur. Based on their time granularity, time prediction methods can be categorized into three types: (1) event Occurrence: Binary-valued prediction on whether an event does or does not occur in a future time period; (2) discrete-time prediction: in which future time slot will the event occur; and (3) continuous-time prediction: at which precise time point will the future event occur.

3.1.1 Occurrence Prediction. Occurrence prediction is arguably the most extensive, classical, and generally simplest type of event time prediction task [12]. It focuses on identifying whether there will be event occurrence (positive class) or not (negative class) in a future time period [206]. This problem is usually formulated as a binary classification problem, although a handful of other methods instead leverage anomaly detection or regression-based techniques.

(1) Binary classification. Binary classification methods have been extensively explored for event occurrence prediction. The goal here is essentially to estimate and compare the values of $f(y = \text{"Yes"}|X)$ and $f(y = \text{"No"}|X)$, where the former denotes the score or likelihood of event occurrence given observation X while the latter corresponds to no event occurrence. If the value of the former is larger than the latter, then a future event occurrence is predicted, but if not, there is no event predicted. To implement f , the methods typically used rely on discriminative models, where dedicated feature engineering is leveraged to manually extract potential event precursor features to feed into the models. Over the years, researchers have leveraged various binary classification

techniques ranging from the simplest threshold-based methods [212], to more sophisticated methods such as logistic regression [209], Support Vector Machines [85], (Convolutional) Neural Networks [117], and decision trees [47]. In addition to discriminative models, generative models [11] have also been used to embed human knowledge for classifying event occurrences using Bayesian decision techniques. Specifically, instead of assuming that the input features are independent, prior knowledge can also be directly leveraged to establish Bayesian networks among the observed features and variables based on graphical models such as (semi-)hidden Markov models [43] and autoregressive logit models [171]. The joint probabilities $p(y = \text{"Yes"}, X)$ or $p(y = \text{"No"}, X)$ can thus be estimated using graphical models, and then utilized to estimate $f(y = \text{"Yes"}|X) = p(y = \text{"Yes"}|X)$ and $f(y = \text{"No"}|X) = p(y = \text{"No"}|X)$ using Bayesian rules [20].

(2) Anomaly detection. Alternatively, anomaly detection can also be utilized to learn a “prototype” of normal samples (typical values corresponding to the situation of no event occurrence), and then identify if any newly-arriving sample is close to or distant from the normal samples, with distant ones being identified as future event occurrences. Notice that here the “detection” is executed on the signals indicative to future events so it has predictive power. Such methods are typically utilized to handle “rare event” occurrences, especially when the training data is highly imbalanced with little to no data for “positive” samples. Anomaly detection techniques such as one-classification [160] and hypotheses testing [83] are often utilized here.

(3) Regression. In addition to simply predicting the occurrence or not, some researchers have sought to extend the binary prediction problem to deal with ordinal and numerical prediction problems, including event count prediction based on (auto)regression [59], event size prediction using linear regression [194], and event scale prediction using ordinal regression [68].

3.1.2 Discrete-time Prediction. In many applications, practitioners want to know the approximate time (i.e., the date, week, or month) of future events in addition to just their occurrence. To do this, the time is typically first partitioned into different time slots and the various methods focus on identifying which time slot future events are likely to occur in. Existing research on this problem can be classified into either direct or indirect approaches.

(1) Direct Approaches. These methods discretize the future time into discrete values, which can take the form of some number of time windows or time scales such as near future, medium future, or distant future. These are then used to directly predict the integer-valued index of future time windows of the event occurrence using (auto)regression [126, 131] and time series methods such as ARIMA [162], or ordinal values of future time scales using ordinal regression or classification [172].

(2) Indirect Approaches. These methods adopt a two-step approach, with the first step being to place the data into a series of time bins and then perform time series forecasting using techniques such as autoregressive [20] based on the historical time series $x = \{x_1, \dots, x_T\}$ to obtain the future time series $\hat{x} = \{x_{T+1}, \dots, x_{\hat{T}}\}$. The second step is to identify events in the predicted future time series $\hat{x}$ using either unsupervised methods such as burstness detection [24] and change detection [92], or supervised techniques based on learning event characterization function. For example, existing works first represent the predicted future time series $\hat{x} \in \mathbb{R}^{\hat{T} \times D}$ using time-delayed embedding, into $\tilde{x} \in \mathbb{R}^{\hat{T} \times D'}$ where each observation at time t can be represented as $\tilde{x}_t = \{x_{t-(D'-1)\tau}, \dots, x_{t-2\tau}, x_{t-\tau}, x_t\}$ and $t = T, T+1, \dots, \hat{T}$ [146]. Then an event characterization function $f_c(\tilde{x}_t)$ is established to map $\tilde{x}_t$ to the likelihood of an event, which can be fitted based on the event labels provided in the training set intuitively. Overall, the unsupervised method requires users to assume the type of patterns (e.g., burstiness and change) of future events based on prior knowledge but do not require event label data. However, in cases where the event time series pattern is difficult to assume but the label data is available, supervised methods are usually used.

3.1.3 Continuous-time Prediction. Discrete-time prediction methods, although usually simple to establish, also suffer from several issues. First, their time-resolution is limited to the discretization granularity; increasing this granularity significantly increases the computational resources required, which means the resolution cannot be arbitrarily high. Moreover, this trade-off is itself a hyperparameter that is sensitive to the prediction accuracy, rendering it difficult and time-consuming to tune during training. To address these issues, a number of techniques work around it by directly predicting the continuous-valued event time [163], usually by leveraging one of three techniques.

(1) Simple Regression. The simplest methods directly formalize continuous-event-time prediction as a regression problem [20], where the output is the numerical-value future event time [179] and/or their duration [111]. Common regressors such as linear regression and recurrent neural networks have been utilized for this. Despite their apparent simplicity, this is not straightforward as simple regression typically assumes Gaussian distribution [111], which does not usually reflect the true distribution of event times. For example, the future event time needs to be left-bounded (i.e., larger than the current time), as well as being typically non-symmetric and usually periodic, with recurrent events having multiple peaks in the probability density function along time dimension.

(2) Point Processes. As they allow more flexibility in fitting true event time distributions, point process methods [141] are widely leveraged and have demonstrated their effectiveness for continuous time event prediction tasks. They require a conditional intensity function, defined as follows:

$$\lambda(t|X) = \mathbb{E}[N(t, t + dt)/dt|X] = g(t|X)/(1 - G(t|X)), \quad (1)$$

where $g(t|X)$ is the conditional density function of the event occurrence probability at time t given an observation X , and whose corresponding cumulative distribution function, $G(t|X)$, $N(t, t + dt)$, denotes the count of events during the time period between t and $t + dt$, where dt is an infinitely-small time period.

Hence, by leveraging the relation between density and accumulative functions and then rearranging Equation (1), the following conditional density function is obtained:

$$g(t|X) = \lambda(t|X) \cdot \exp \left(- \int_{t_0}^t \lambda(u|X) du \right). \quad (2)$$

Once the above model has been trained using a technique such as maximal likelihood [20], the time of the next event in the future is predicted as $\hat{t} = \int_{t_0}^{\infty} t \cdot g(t|X) dt$. Although existing methods typically share the same workflow as that shown above, they vary in the way they define the conditional intensity function $\lambda(t|X)$. Traditional models typically utilize prescribed distributions such as the Poisson distribution, Gamma distribution, Hawks process, Weibull process, and other distributions [184]. For example, Damaschke et al. [45] utilized a Weibull distribution to model volcano eruption events, while Ertekin et al. [58] instead proposed the use of a non-homogeneous Poisson process to fit the conditional intensity function for power system failure events. However, in many other situations where there is no information regarding appropriate prescribed distributions, researchers must start by leveraging nonparametric approaches to learn sophisticated distributions from the data using expressive models such as neural networks. For example, Simma and Jordan [163] utilized of **Recurrent Neural Network (RNN)** to learn a highly nonlinear function of $\lambda(t|X)$.

(3) Survival Analysis. Survival analysis [51] is related to point processes in that it also defines an event intensity or hazard function, but in this case based on survival probability considerations:

$$H(t|X) = (\xi(t - dt|X) - \xi(t|X)) / \xi(t|X), \text{ where } \xi(t|X) = p(\hat{t} > t), \quad (3)$$

Fig. 4. Raster-based and Point-based event location predictions [174, 201].

where $H(t|X)$ is the so-called Hazard function denoting the hazard of event occurrence between time $(t - dt)$ for a t for a given observation X . Either $H(t|X)$ or $\xi(t|X)$ could be utilized for predicting the time of future events. For example, the event occurrence time can be estimated when $\xi(t|X)$ is lower than a specific value. Also, one can obtain $\xi(t|X) = \exp\left(-\int_0^t H(u|X)du\right)$ according to Equation (3) [114]. Here $H(t|X)$ can adopt any one of several prescribed models, such as the well-known Cox hazard model [50]. To learn the model directly from the data, some researchers have recommended enhancing it using deep neural networks [101]. Vahedian et al. [175] suggest learning the survival probability $\xi(t|X)$ and then applying the function $H(\cdot|X)$ to indicate an event at time t if $H(t|X)$ is larger than a predefined threshold value. A classifier can also be utilized.

Instead of using the raw sequence data, the conditional intensity function can also be projected onto additional continuous-time latent state layers that eventually map to the observations [51]. These latent states can then be extracted using techniques such as hidden semi-Markov models [20], which ensure the elicitation of the continuous time patterns.

3.2 Location Prediction

Event location prediction focuses on predicting the location of future events. Location information can be formulated as one of two types: **(1) Raster-based.** Here, a continuous space is partitioned into a grid of cells, each of which represents a spatial region, as shown in Figure 4(a). This type of representation is suitable for situations where the spatial size of the event is non-negligible. **(2) Point-based.** In this case, each location is represented by an abstract point with infinitely-small size, as shown in Figure 4(b). This type of representation is most suitable for the situations where the spatial size of the event can be neglected, or the location regions of the events can only be in discrete spaces such as network nodes.

3.2.1 Raster-based Location Prediction. There are three types of techniques used for raster-based event location prediction, namely, spatial clustering, spatial embedding, and spatial convolution.

(1) Spatial clustering. In raster-based representations, each location unit is usually a regular grid cell with the same size and shape. However, regions with similar spatial characteristics typically have irregular shapes and sizes, which could be approximated as composite representations of a number of grids [88]. The purpose of spatial clustering here is to group the contiguous regions who collectively exhibit significant patterns. The methods are typically agglomerative style. They typically start from the original finest-grained spatial raster units and proceed by merging the spatial neighborhood of a specific unit in each iteration. But different research works define different criteria for instantiating the merging operation. For example, Wang and Ding [178] merge neighborhoods if the unified region after merging can maintain the spatially frequent patterns. Xiong et al. [185] chose an alternative approach by merging spatial neighbor locations into the current locations sequentially until the merged region possesses event data that is sufficiently statistically

significant. These methods usually run in a greedy style to ensure their time complexity remains smaller than quadratic. After the spatial clustering is completed, each spatial cluster will be input into the classifier to determine whether or not there is an event corresponding to it.

(2) Spatial interpolation. Unlike spatial clustering-based methods, spatial interpolation-based methods maintain the original fine granularity of the event location information. The estimation of event occurrence probability can be further interpolated for locations with no historical events and hence achieve spatial smoothness. This can be accomplished using commonly-used methods such as kernel density estimation [76] and spatial Kriging [88]. Kernel density estimation is a popular way to model the geo-statistics in numerous types of events such as crimes [5] and terrorism [76]:

$$k(s) = 1/(n \cdot \gamma) \sum_{i=1}^n K((s - s_i)/\gamma), \quad (4)$$

where $k(s)$ denotes the kernel estimation for the location point s , n is the number of historical event locations, each s_i is a historical event location, γ is a tunable bandwidth parameter, and $K(\cdot)$ is a kernel function such as Gaussian kernel [69].

More recently, Ristea et al. [148] further extended **Kernel Density Estimation (KDE)**-based techniques by leveraging Localized KDE and then applying spatial interpolation techniques to estimate spatial feature values for the cells in the grid. Since each cell is an area rather than a point, the center of each cell is usually leveraged as the representative of this cell. Finally, a classifier will take this as its input to predict the event occurrence for each grid [148].

(3) Spatial convolution. In the past few years, **Convolutional Neural Networks (CNNs)** have demonstrated significant success in learning and representing sophisticated spatial patterns from image and spatial data [71]. A CNN contains multiple convolutional layers that extract the hierarchical spatial semantics of images. In each convolutional layer, a convolution operation is executed by scanning a feature map with a filter, which results in another smaller feature map with a higher level semantic. Since raster-based spatial data and images share a similar mathematical form, it is natural to leverage CNNs to process it.

Existing methods in this category typically formulate a spatial map as input to predict another spatial map that denotes future event hotspots [139]. Such a formulation is analogous to the “image translation” problem popular in recent years in the computer vision domain [36]. Specifically, researchers typically leverage an encoder-decoder architecture, where the input images (or spatial map) are processed by multiple convolutional layers into a higher-level representation, which is then decoded back into an output image with the same size, through a reverse convolutional operations process known as transposed convolution [71].

(4) Trajectory destination prediction. This type of method typically focuses on population-based events whose patterns can be interpreted as the collective behaviors of individuals, such as “gathering events” and “dispersal events” [93, 174, 215]. These methods share a unified procedure that typically consists of two steps: (1) predict future locations based on the observed trajectories of individuals, and (2) detect the occurrence of the “future” events based on the future spatial patterns obtained in Step 1. Complex footprints of events such as those with flow direction information in a graph setting could also be considered [93, 174]. The specific methodologies for each step are:

•**Step 1:** Here, the aim is to predict each location an individual will visit in the future, given a historical sequence of locations visited. This can be formulated as a sequence prediction problem. For example, Wang and Gerber [182] sought to predict the probability of the next time point $t + 1$'s location s_{t+1} based on all the preceding time points: $p(s_{t+1}|s_{\leq t}) = p(s_{t+1}|s_t, s_{t-1}, \dots, s_0)$, based on various strategies including a historical volume-based prior model, Markov models, and multi-class classification models. Vahedian et al. [174] adopted Bayesian theory $p(s_{t+1}|s_{\leq t}) = p(s_{\leq t}|s_{t+1}) \cdot p(s_{t+1})/p(s_{\leq t})$, which requires the conditional probability $p(s_{\leq t}|s_{t+1})$ to be stored.

However, in many situations, there is huge number of possible trajectories for each destination. For example, with a 128×64 grid, one needs to store $(128 \times 64)^3 \approx 5.5 \times 10^{11}$ options. To improve the memory efficiency, this can be limited to a consideration of just the source and current locations, leveraging a quad-tree style architecture to store the historical information. To achieve more efficient storage and speed up $p(s_{\leq t} | s_{t+1})$ queries, Vahedian et al. [174] extended quad-tree into a new technique called VIGO, which removes the duplicate intermediate locations for the same destination in leaves.

•**Step 2:** The aim in this step is to forecast future event locations based on the future visiting patterns predicted in Step 1. The most basic strategy here is to consider each grid cell independently. For example, Wang and Gerber [182] adopted supervised learning strategies to build predictive mapping between the visiting patterns and the event occurrence. A more sophisticated approach is to consider the spatial outbreaks composited by multiple grids. Scalable algorithms have also been proposed to identify regions containing statistically significant hotspots [94], such as spatial scan statistics [98]. Khezerlou et al. [94] proposed a greedy-based heuristic tailored for the grid-based data formulation, which extends the original “seed” grid containing statistically-large future event densities to four directions until the extended region is no longer a statistically-significant outbreak.

3.2.2 Point-based Prediction. Unlike the raster-based formulation, which covers the prediction of a contiguous spatial region, point-based prediction focuses specifically on locations of interest, which can be distributed sparsely in a Euclidean (e.g., spatial region) or non-Euclidean space (e.g., graph topology). These methods can be categorized into supervised and unsupervised approaches.

(1) *Supervised approaches.* In supervised methods, each location will be classified as either “positive” or “negative” with regard to a future event occurrence. The simplest setting is based on the **independent and identically distributed (i.i.d.)** assumption among the locations, where each location is predicted by a classifier independently using their respective input features. However, given that different locations usually have strong spatial heterogeneity and dependency, further research has been proposed to tackle them based on different locations’ predictors and outputs, resulting in two research directions: (1) Spatial multi-task learning and (2) Spatial auto-regressive methods.

•**Spatial multi-task learning.** Multi-task learning is a popular learning strategy that can jointly learn the models for different tasks such that the learned model can not only share their knowledge but also preserve some exclusive characteristics of the individual tasks [206]. This notion coincides very well with spatial event prediction tasks, where combining the outputs of models from different locations needs to consider both their spatial dependency and heterogeneity. Zhao et al. [206] proposed a spatial multi-task learning framework as follows:

$$\min_W \sum_i^m \mathcal{L}(Y_i, f(W_i, X_i)) + \alpha \cdot \mathcal{R}(\{W_i\}_i^k, M), \quad s.t., \quad C(\{W_i\}_i^m, D) \in \mathbb{C}, \quad (5)$$

where m is the total number of locations (i.e., tasks), W_i and Y_i are the model parameters and true labels (event occurrence for all time points), respectively, of task i . $\mathcal{L}(\cdot)$ is the empirical loss, $f(W_i, X_i)$ is the predictor for task i , and $\mathcal{R}(\cdot)$ is the spatial regularization term based on the spatial dependency information $M \in \mathbb{R}^{m \times m}$, where $M_{i,j}$ records the spatial dependency between location i and j . $C(\cdot)$ represents the spatial constraints imposed over the corresponding models to enforce them to remain within the valid space $\mathbb{C}$. Over recent years, there have been multiple studies proposing different strategies for $\mathcal{R}(\cdot)$ and $C(\cdot)$. For example, Gao et al. [69] further differentiated the strength of the correlation between different locations’ tasks according to the spatial distance between them. This research has been further extended this approach to tree-structured multitask

learning to handle the hierarchical relationship among locations at different administrative levels (e.g., cities, states, and countries) [207] in a model that also considers the logical constraints over the predictions from different locations who have hierarchical relationships. Instead of evenly similar, Zhao, et al. [205] further estimated spatial dependency D utilizing inverse distance using Gaussian kernels, while Ning et al. [132] proposed estimating the spatial dependency D based on the event co-occurrence frequency between each pair of locations.

•**Spatial auto-regressive methods.** Spatial auto-regressive models have been extensively explored in domains such as geography and econometrics, where they are applied to perform predictions where the i.i.d. assumption is violated due to the strong dependencies among neighbor locations. Its generic framework is $\hat{Y}_{t+1} = \rho M \hat{Y}_{t+1} + X_t \cdot W + \varepsilon$, where $X_t \in \mathbb{R}^{m \times D}$ and $\hat{Y}_{t+1} \in \mathbb{R}^{m \times m}$ are the observations at time t and event predictions at time $t + 1$ over all the m locations, and $M \in \mathbb{R}^{m \times m}$ is the spatial dependency matrix with zero-valued diagonals. This means the prediction of each location $\hat{Y}_{t+1,i} \in \hat{Y}_{t+1}$ is jointly determined by its input $X_{t,i}$ and neighbors $\{j | M_{i,j} \neq 0\}$ and ρ is a positive value to balance these two factors. Since event occurrence requires discrete predictions, simple threshold-based strategies can be used to discretize $\hat{Y}_i$ into $\hat{Y}'_i = \{0, 1\}$ [25]. Moreover, due to the complexity of event prediction tasks and the large number of locations, sometimes it is difficult to define the whole M manually. Zhao et al. [205] proposed jointly learning the prediction model and spatial dependency from the data using graphical LASSO techniques. Yi et al. [193] took a different approach, leveraging conditional random fields to instantiate the spatial autoregression, where the spatial dependency is measured by Gaussian kernel-based metrics. Yi et al. [192] then went on to propose leveraging the neural network model to learn the locations' dependency.

(2) *Unsupervised approaches.* Without event labels, unsupervised methods must first identify potential precursors and determinant features in different locations. They can then detect anomalies characterized by specific feature selection and location combinatorial patterns (e.g., spatial out-breaks and connected subgraphs) as the future event indicators [31] as an optimization problem:

$$(F, R) = \arg \max_{F, R} q(F, R) \quad s.t., \quad \text{supp}(F_i) \in \mathbb{M}(\mathbb{G}, \beta), \quad \text{supp}(R_i) \in \mathbb{C}, \quad (6)$$

where $q(\cdot)$ denotes scan statistics that score the significance of each candidate pattern, represented by both a candidate location combinatorial pattern R and feature selection pattern F . Specifically, $F \in \{0, 1\}^{D \times n}$ denotes the feature selection results (where “1” means selected; “0,” otherwise) and $R \in \{0, 1\}^{m \times n}$ denotes the m involved locations for the n events. $\mathbb{M}(\mathbb{G}, \beta)$ and $\mathbb{C}$ are the set of all the feasible solutions of F and R , respectively. $q(\cdot)$ can be instantiated by scan statistics such as Kulldorff's scan statistics [98] and the Berk-Jones statistic [31], which can be applied to detect and forecast events such as epidemic outbreaks and civil unrest events [144]. Depending on whether the embedding space is an Euclidean region (e.g., a geographical region) or a non-Euclidean region (e.g., a network topology), the pattern constraint C can be either constrained to predefined geometric shapes such as a circle, rectangle, or an irregular shape or subgraphs such as connected, cliques, and k -cliques. The problem in Equation (6) is nonconvex and sometimes even discrete, and hence difficult to solve. A generic way is to optimize F using sparse feature selection; there is a useful survey provided in Reference [109] and R can be defined using the two-step graph-structured matching method detailed in Reference [32]. More recently, new techniques have been developed that are capable of jointly learning both feature and location selection [32].

3.3 Semantic Prediction

Event semantics prediction addresses the problem of forecasting topics, descriptions, or other meta-attributes in addition to future events' times and locations. Unlike time and location prediction, the data in event semantics prediction usually involves symbols and natural languages in addition to numerical quantities, which means different types of techniques may be utilized. The

data are categorized into three types based on how the historical data are organized and utilized to infer future events. The first of these categories covers rule-based methods, where future event precursors are extracted by mining association or logical patterns in historical data. The second type is sequence-based, considering event occurrence to be a consequence of temporal event chains. The third type further generalizes event chains into event graphs, where additional cross-chain contexts need to be modeled. These are discussed in turn below.

3.3.1 Association-based Prediction. Association rule-based methods are amongst the most classic approaches in data mining domain for event prediction, typically consisting of two steps: (1) learn the associations between precursors and target events, and then (2) utilize the learned associations to predict future events. For the first step, for example, an association could be $x = \{\text{"election," "fraud"}\} \rightarrow y = \text{"protest event,"}$ which indicates that serious fraud occurring in an election process could lead to future protest events. To discover all the significant associations from the ocean of candidate rules efficiently, frequent set mining [75] can be leveraged. Each discovered rule needs to come with both sufficient *support* and *confidence*. Here, *support* is defined as the number of cases where both “ x ” and “ y ” co-occur, while *confidence* means the ratio indicating that “ y ” occurs once “ x ” happens. To better estimate these discrimination rules, further temporal constraints can be added that require the occurrence time of “ x ” and “ y ” to be sufficiently close to be considered “co-occurrences.” Once the frequent set rules have been discovered, pruning strategies may be applied to retain the most accurate and specific ones, with various strategies for generating final predictions [75]. Specifically, given each new observation x' , one of the simplest strategies is to output the events that are triggered by any of the association rules starting from event x' [176]. Other strategies first rank the predicted results based on their confidence and then predict just the top r events [213]. More sophisticated and rigorous strategies tend to build a decision list where each element in the list is an association rule, so once a generative model has been built for the decision process, the maximal likelihood can be leveraged to optimize the order of the list [106].

3.3.2 Causality-based Prediction. This type of research leverages the causality inferred among the historical events to achieve future event predictions. The data here typically shares a generic framework consisting of the following procedures: (1) event representation, (2) event graph construction, and (3) future event inference.

Step 1: Event semantic representation. This approach typically begins by extracting the events from the target texts using natural language processing techniques such as sanitization, tokenization, POS tag analysis, and name entity recognition. Several types of objects can be extracted to represent the events: (i) Noun Phrase-based [77], where the noun-phrase corresponds to each event (for example, “2008 Sichuan Earthquake”); (ii) Verbs and Nouns [142], where an event is represented as a set of noun-verb pairs extracted from news headlines (for example, “<capture, people>,” “<escape, prison>,” or “<send, prison>”); and (iii) Tuple-based [210], where each event is represented by a tuple consisting of objects (such as actors, instruments, or receptors), a relationship (or property), and time. An RDF-based format has also been leveraged in some works [46].

Step 2: Event causality inference. The goal here is to infer the cause-effect pairs among historical events. Due to its combinatorial nature, narrowing down the number of candidate pairs is crucial. Existing works usually begin by clustering the events into event chains, each of which consist of a sequence of time-ordered events under the relevant semantics, typically the same topics, actors, and/or objects [2]. The causal relations among the event pairs can then be inferred in various ways. The simplest approach is just to consider the likelihood that y occurs after x has occurred throughout the training data. Other methods utilize NLP techniques to identify causal mentions such as causal connectives, prepositions, and verbs [142]. Some formulate causal-effect relationship identification as a classification task where the inputs are the cause

and effect candidate events, often incorporating contextual information including related background knowledge from web texts. Here, the classifier is built on a multi-column CNN that outputs either “1” or “0” to indicate whether the candidate has an effect or not [97]. In many situations, the cause-effect rules learned directly using the above methods can be too specific and sparse, with low generalizability, so a typical next step is to generalize the learned rules. For example, “Earthquake hits China” $\rightarrow$ “Red Cross help sent to Beijing” is a specific rule that can be generalized to “Earthquake hits [A country]” $\rightarrow$ “Red Cross help sent to [The capital of this country].” To achieve this, some external ontology or a knowledge base is typically needed to establish the underlying relationships among items or provide necessary information on their properties, such as Wikipedia (<https://www.wikipedia.org/>), YAGO [168], WordNet [60], or ConceptNet [119]. Based on these resources, the similarity between two cause-effect pairs (c_i, ε_i) and (c_j, ε_j) can be computed by jointly considering the respective similarity of the putative cause and effect: $\sigma((c_i, \varepsilon_i), (c_j, \varepsilon_j)) = (\sigma(c_i, c_j) + \sigma(\varepsilon_i, \varepsilon_j))/2$. An appropriate algorithm can then be utilized to apply hierarchical agglomerative clustering to group them and hence generate a data structure that can efficiently manage the task of storing and querying them to identify any cause-effect pairs. For example, Reference [161] leverage an abstraction tree, where each leaf is an original specific cause-effect pair and each intermediate node is the centroid of a cluster. Instead of using hierarchical clustering, Reference [210] directly uses the word ontology to simultaneously generalize cause and effect (e.g., the noun “violet” is generalized to “purple,” the verb “kill” is generalized to “murder-42.1¹”) and then leverage a hierarchical causal network to organize the generalized rules.

Step 3: Future event Inference. Given an arbitrary query event, two steps are needed to infer the future events caused by it based on the causality of events learned above. First, we need to retrieve similar events that match the query event from historical event pool. This requires the similarity between the query event and all the historical events to be calculated. To achieve this, Lei et al. [105] utilized context information, including event time, location, and other environmental and descriptive information. For methods requiring event generalization, the first step is to traverse the abstraction tree starting from the root that corresponds to the most general event rule. The search frontier then moves across the tree if the child node is more similar, culminating in the nodes that are the least general but still similar to the new event being retrieved [142]. Similarly, Reference [35] proposed another tree structure referred to as a “circular binary search tree” to manage the event occurrence pattern. We can now apply the learned predicate rules starting from the retrieved event to obtain the prediction results. Since each cause event can lead to multiple events, a convenient way to determine the final prediction is to calculate the support [142], or conditional probability [191] of the rules. Radinsky et al. [142] took a different approach, instead ranking the potential future events by their similarity defined by the length of their minimal generalization path. For example, the minimal generalization path for “London” and “Paris” is “London” $\xrightarrow{\text{capital-of}}$ “Great Britain” $\xrightarrow{\text{in-continent}}$ “Europe” $\xleftarrow{\text{in-continent}}$ “France” $\xleftarrow{\text{capital-of}}$ “Paris.” Alternatively, Zhao et al. [210] proposed embedding the event causality network into a continuous vector space and then designing an energy function to rank potential events, where true cause-effect pairs should have low energies.

3.3.3 Semantic Sequence. These methods share a very straightforward problem formulation. Given a temporal sequence for a historical event chain, the goal is to predict the semantics of the next event using sequence prediction [20]. Current methods come from two major categories: (1) Whole sequence classification and (2) Next event generation.

¹The form of verb class in VerbNet [210].

(1) Whole-sequence classification-based methods. These methods formulate event semantic prediction as a multi-class classification problem, where a finite number of candidate events are ranked and the top-ranked event is treated as the future event semantic. The objective is $\hat{C} = \arg \max_{C_i} u(s_{T+1} = C_i | s_1, \dots, s_T)$, where s_{T+1} denotes the event semantic in time slot $T + 1$ and $\hat{C}$ is the optimal semantic among all the semantic candidates C_i ($i = 1, \dots$). Multi-class classification problems can be split into events with different topics/semantic meaning. Three types of sequence classification methods have been utilized for this purpose, namely, feature-based methods, prototype-based methods, and model-based methods such as Markov models.

•**Feature-based.** One of the simplest methods is to ignore the temporal relationships among the events in the chain, by either aggregating the inputs or the outputs. Tama and Comuzzi [170] formulated historical event sequences with multiple attributes for event prediction, testing multiple conventional classifiers. Another type of approach based on this notion utilizes compositional-based methods [72] that typically leverage the assumption of independency among the historical input events to simplify the original problem $u(s_{T+1} | s_1, s_2, \dots, s_T) = u(s_{T+1} | s_{\leq T})$ into $v(u(s_{T+1} | s_1), u(s_{T+1} | s_2), \dots, u(s_{T+1} | s_T))$, where $v(\cdot)$ is simply an aggregation function that represents a summation operation over all the components. Each component function $u(s_{T+1} | s_i)$ can then be calculated by estimating how likely it is that event semantic s_{T+1} and s_i ($i \leq T$) co-occur in the same event chain. Granroth-Wilding and Clark [72] investigated various models ranging from straightforward similarity scoring functions through bigram models and word embedding combined with similarity scoring functions to newly developed composition neural networks that jointly learn the representation of s_{T+1} and s_i and then calculate their coherence. Some other researchers have gone further to consider the dependency among the historical events. For example, Letham et al. [107] proposed to optimizing the correct ordering among the candidate events as follows:

$$\sum_{i \in \mathcal{I}, j \in \mathcal{J}} \mathbb{1}_{[u(s_{T+1}=C_i | s_{\leq T}) > u(s_{T+1}=C_j | s_{\leq T})]} \Rightarrow \sum_{i \in \mathcal{I}, j \in \mathcal{J}} e^{u(s_{T+1}=C_i | s_{\leq T}) - u(s_{T+1}=C_j | s_{\leq T})} + \rho \|W\|_2^2,$$

where the semantic candidate in the set $\mathcal{I}$ should be ranked strictly to be lower than those in $\mathcal{J}$, with the goal being to penalize the “incorrect ordering.” Here, $\mathbb{1}_{[\cdot]}$ is an indicator function that is discrete, such that $\mathbb{1}_{[b \geq a]} \leq e^{b-a}$ and can thus be utilized as the upper-bound for minimization, as can be seen in the right-hand-side of the above equation. W is the set of parameters of the function $u(\cdot)$. This can now be relaxed to an exponential-based approximation for effective optimization using gradient-based algorithms [71]. Other methods focus on first transferring the sequential data into sequence embeddings that can encode the latent sequential context. For example, Fronza et al. [64] apply random indexing to represent the words in terms of their vector representations by embedding the information from neighboring words into each word before utilizing conventional classifiers such as SVM to identify the future events.

•**Model-based.** Markov-based models have also been leveraged to characterize temporal patterns [189]. These typically use E_i to denote each event under a specific type and $\mathcal{E}$ denotes the set of event types. The goal here is to predict the event type of the next event to occur in the future. In Reference [7], the event types are modeled using the Markov model so given the current event type, the next event type can be inferred simply by looking up the state with the highest probability in the transition matrix. A tool called Wayeb [8] has been developed based on this method. Laxman et al. [104] developed a more complicated model, based on a mixture of Hidden Markov models and introducing new assumptions and the concept of episodes composed of a subsequence of event types. They assumed different event episodes should have different transition patterns so started by discovering the frequent episodes for events, each of which they modeled by a specific hidden Markov model over various event types. This made it possible to establish the

generative process for each future event type s based on the mixture of the above episode Markov models. When predicting, the likelihood of an observed event sequence over each possible generative process, $p(X|\Lambda_Y)$ is evaluated, after which a future event type can be considered as either being over some threshold (as in Reference [104]) or the largest among all the different Y values (as in Reference [202]).

•**Prototype-based.** Adhikari et al. [3] took a different approach, utilizing a prototype-based strategy that first clusters the event sequences into different clusters in terms of their temporal patterns. When a new event sequence is observed, its closest cluster's centroid will then be leveraged as a "reference event sequence" whose sub-sequential events will be referred to when predicting future events for this new event sequence.

(2) **Next event generation-based methods.** Approaches in this category can be classified into two types: (1) Attribute-based models and (2) Descriptive-based models. The attribute-based models, ingest feature representation of events as input, while the descriptive-based models typically ingest unstructured information such as texts to directly predict future events.

•**Attributed-based Methods.** Here, each event $y = (t, l, s)$ at time t is recast and represented as $e_t = (e_{t,1}, e_{t,2}, \dots, e_{t,k})$, where $e_{t,i}$ is the i th feature of the event at time t . The feature here can include location and other information such as event topic and semantics. Each sequence $e = (e_1, \dots, e_t)$ is then input into the standard RNN architecture for predicting next event e_{t+1} in the sequence at time point $t + 1$ [116]. Various types of RNN components and architecture have been utilized for this purpose [26, 27], but a vanilla RNN [71] for sequence-based event prediction can be written in the following form:

$$h_i = \text{tanh}(a_i), a_i = b + W \cdot h_{i-1} + U \cdot e_i, o_i = c + V \cdot h_i, \psi(i + 1) = \text{softmax}(o_i), i \leq t,$$

where h_i , o_i , and a_i are the latent state, output, and activation for the i th event, respectively, and W , U , and V are the model parameters for fitting the corresponding mappings. The prediction $e_{t+1} := \psi(t + 1)$ can then be calculated in a feedforward way from the first event and the model training can be done by back-propagating the error from the layer of $\psi(t)$. Existing work typically utilizes the variants of vanilla RNN to handle the gradient vanishing problem, especially when the event chain is not short. The most commonly used methods for event prediction are LSTM and **Gated Recurrent Units (GRU)** [71]. For example, the architecture and equation for LSTM are as follows:

$$a_i = \sigma(W_j \cdot [h_{i-1}, e_i] + b_j), \tilde{C}_i = \tanh(W_C \cdot [h_{i-1}, h_i] + b_C), C_i = \zeta_i C_{i-1} + a_i \tilde{C}_i, \quad (7)$$

$$\zeta_i = \sigma(W_\zeta \cdot [h_{i-1}, e_i] + b_\zeta), o_i = \sigma(W_o[h_{i-1}], e_i) + b_o), h_i = o_i * \text{tanh}(C_i),$$

where the additional components C_{i-1} and ζ_i are introduced to keep tracking the previous "history" and gating the information for forgetting to handle longer sequences. For example, some researchers opt to leverage a simple type LSTM architecture to extend the RNN-based sequential event prediction [26], while others leverage variants of LSTM, such as bi-directional LSTM instead [95] and yet others prefer to leverage GRU [56].

Moving beyond considering just the chain relationships among events, Li et al. [113] generalized this into graph-structured relationships to better incorporate the event contextual information via the **Narrative Event Evolutionary Graph (NEEG)**. An NEEG is a knowledge graph where each node is an event and each edge denotes the association between a pair of events, enabling the NEEG to be represented by a weighted adjacency matrix A . The basic architecture can be denoted by the following, as detailed in the article [113]:

$$a_i = A^T h_{i-1} + b, z_i = \sigma(W_z a_i + U_z h_{i-1}), r_i = \sigma(W_r a_i + U_r h_{i-1}), \quad (8)$$

$$c_i = \tanh(W a_i + U(r_i h_{i-1})), h_i = (1 - z_i) h_{i-1} + z_i c_i.$$

Fig. 5. Generic framework for hierarchical RNN-based event forecasting.

Here, the current activation a_i is not only dependent on the previous time point but also influenced by its neighbor nodes in NEEG.

•**Descriptive-based Methods.** Attribute-based methods require extra effort during pre-processing to convert the unstructured raw data into feature vectors, a process that is not only computationally labor intensive but also not always feasible. Therefore, multiple architectures have been proposed to directly process the raw (textual) event descriptions to enable them to be used to predict future event semantics or descriptions. These models share a similar generic framework [80, 167, 186], which begins by encoding each sequence of words into event representations, utilizing an RNN architecture, as shown in Figure 5. The sequence of events must then be characterized by another higher-level RNN to predict future events. Under this framework, some works begin by decoding the predicted future candidate events into event embedding, after which they are compared with each other and the one with the largest confidence score is selected as the predicted event. These methods are usually constrained by the known list of event types, but sometimes we are interested in open set predictions where the predicted event type can be a new appearance of a type that has not previously been seen in the training set. To achieve this, other methods focus on directly generating future events' descriptions that characterize event semantics that may or may not have appeared before by designing an additional sequence decoder that decodes the latent representation of future events into word sequences. More recent research has enhanced the utility and interpretability of the relationship between words and relevant events, and all the previous events for the relevant future event, by adding a hierarchical attention mechanisms. For example, Yu et al. [195] and Su and Jiang [167] both proposed word-level attention and event-level attention, while Hu [79] leveraged word-level attention in the event encoder as well as in the event decoder.

3.4 Joint Prediction

This section discusses the research into ways to jointly predict the time, location, and semantics of future events. Existing work can be categorized into three types: (1) joint time and semantics prediction; (2) joint time and location prediction; and (3) joint time, location, and semantic prediction.

3.4.1 Time and Semantics. For joint time and semantic prediction, there are three popular types of methods, discussed in turn below:

Temporal association rule. A temporal association rule can be developed from the vanilla association rule $LHS \rightarrow RHS$ by embedding additional temporal information into either LHS , RHS , or both, thus redefining the meaning of co-occurrence and association with temporal constraints. For example, Vilalta and Ma [176] defined LHS as a tuple (E_L, τ) , where τ is the time window before the target in RHS predefined by the user. Only the events occurring within a time window before the event in RHS will satisfy the LHS . Similar techniques have also been leveraged by other researchers [35, 166]. However, τ is difficult to define beforehand and it is preferable to

Fig. 6. Tensor decomposition and forecasting for complex time-stamped events.

be flexible to suit different target events. To handle this challenge, Yang et al. [190] proposed a way to automatically identify information on a continuous time interval from the data. Here, each transaction is composed of not only items but also continuous time duration information. *LHS* is a set of items (e.g., previous events) while *RHS* is a tuple $(E_R, [t_1, t_2])$ consisting of a future event semantic representation and its time interval of occurrence. To automatically learn the time interval in *RHS*, Reference [190] proposed the use of two different methods. The first is called the *confidence-interval-based* method, which leverages a statistical distribution (e.g., Gaussian and student-t [20]) to fit all the observed occurrence times of events in *RHS*, and then treats the statistical confidence interval as the time interval. The other method is *minimal temporal region selection*, which finds the temporal region with the smallest interval and covers historical occurrences of the event in *RHS*.

Time expression extraction. In contrast to the above statistical-based methods, another way to achieve event time and semantics joint prediction comes from the pattern recognition domain, aiming to directly discover time expressions that mention the (planned) future events. As this type of technique can simultaneously identify time, semantics, and other information such as locations, it is widely used and will be discussed in more details later as part of the discussion of “Future event expression identification methods” in Section 3.4.3.

Time series forecasting-based methods. The methods based on time series forecasting can be separated into direct methods and indirect methods. *Direct methods* typically formulate the event semantic prediction problem as a multi-variate time series forecasting problem, where each variable corresponds to an event type C_i ($i = 1, \dots$) and hence the predicted event type at future time $\hat{t}$ is calculated as $\hat{s}_t = \arg \max_{C_i} f(s_t = C_i | X)$. For example, in Reference [110], a longitudinal support vector regressor is utilized to predict multi-attribute events, where n support vector regressors, each of which corresponds to an attribute, is built to achieve the goal of predicting the next time point’s attribute value. Weiss and Page [184] took a different approach, leveraging multiple point process models to predict multiple event types. To further estimate the confidence of their predictions, Biloš et al. [19] first leveraged RNN to learn the historical event representation and then input the result into a Gaussian process model to predict future event types. To better capture the joint dynamics across the multiple variables in the time series, Brandt et al. [23] extended this to Bayesian vector autoregression. Utilizing *indirect-style* methods, they focused on learning a mapping from the observed event semantics down to low-dimensional latent-topic space using tensor decomposition-based techniques. Similarly, Matsubara et al. [121] proposed a 3-way topic analysis of the original observed event tensor $Y_0 \in \mathbb{R}^{D_o \times D_a \times D_c}$ consisting of three factors, namely, actors, objects, and time. They then went on to decompose this tensor into latent variables via three corresponding low-rank matrices $P_o \in \mathbb{R}^{D_o \times D_k}$, $P_a \in \mathbb{R}^{D_a \times D_k}$, and $P_c \in \mathbb{R}^{D_c \times D_k}$, respectively, as shown in Figure 6. Here D_k is the number of latent topics. For the prediction, the time matrices P_c are predicted into the future $\hat{P}_c$ via multi-variate time series forecasting, after which a future event tensor are estimated by recovering a “future event tensor” $\hat{Y}$ by the multiplication among the predicted time matrix $\hat{P}_c$ as well as the known actor matrix P_a and object matrix P_o .

Fig. 7. Generic framework for spatiotemporal event prediction using CNN+RNN-based framework.

3.4.2 Time and Location Prediction. This category of methods focuses on jointly predicting the location and time of future events. These methods can be classified into two subtypes: (1) raster-based, which focus on predictions for individual time slots and location regions, and (2) point-based, which predicts continuous time and location points.

Raster-based. These methods usually formulate data into temporal sequences consisting of spatial snapshots. Over the last few years, various techniques have been proposed to characterize the spatial and temporal information for event prediction.

The simplest way to consider spatial information is to directly treat location information as one of the input features, and then feed it into predictive models, such as linear regression [211], LSTM [147] and Gaussian processes [100]. During model training, Zhao and Tang [211] leveraged the spatiotemporal dependency to regularize their model parameters.

Most of the methods in this domain aim to jointly consider the spatial and temporal dependency for predictions [53]. At present, the most popular framework is the CNN+RNN architecture, which implements sequence-to-sequence learning problems such as the one illustrated in Figure 7. Here, the multi-attributed spatial information for each time point can be organized as a series of multi-channel images, which can be encoded using convolution-based operations. For example, Huang et al. [82] proposed the addition of convolutional layers to process the input into vector representations. Other researchers have leveraged variational autoencoders [183] and CNN autoencoders [87] to learn the low-dimensional embedding of the raw spatial input data. This allows the learned representation of the input to be input into the temporal sequence learning architecture. Different recurring units have been investigated, including RNN, LSTM, convLSTM, and stacked-convLSTM [71]. The resulting representation of the input sequence is then sent to the output sequence as input. Here, another recurrent architecture is established. The output of the unit for each time point will be input into a spatial decoder component, which can be implemented using transposed convolutional layers [197], transposed convLSTM [87], or a spatial decoder in a variational autoencoder [183]. A conditional random field is another popular technique often used to model the spatial dependency [88].

Point-based. The spatiotemporal point process is an important technique for spatiotemporal event prediction, which models the event occurrence rate in both spatial and time points. It is defined as

$$\lambda(t, l|X) = \lim_{|dt| \rightarrow 0, |dl| \rightarrow 0} \mathbb{E}[N(dt \times dl)|X] / (|dt||dl|). \quad (9)$$

Various models have been proposed to instantiate the model of the framework illustrated in Equation (9). For example, Liu and Brown et al. [118] began by assuming there to be a conditional independence among spatial and temporal factors and hence achieved the following

decomposition:

$$\lambda(t, l|X) = \lambda(t, l|L, T, F) = \lambda_1(l|L, T, F, t) \cdot \lambda_2(t|T), \quad (10)$$

where X , L , T , and F denotes the whole input indicator data as well as its different facets, including location, time, and other semantic features, respectively. Then the term $\lambda_1(\cdot)$ can be modeled based on the Markov spatial point process while $\lambda_2(\cdot)$ can be characterized using temporal autoregressive models. To handle situations where explicit assumptions for model distributions are difficult, several methods have been proposed to involve the deep architecture during the point process. Most recently, Okawa et al. [134] have proposed the following:

$$\lambda(t, l|X) = \int g_\theta(t', l', \mathcal{F}(t', l')) \cdot \mathcal{K}((t, l), (t', l')) dt' dl', \quad (11)$$

where $\mathcal{K}(\cdot, \cdot)$ is a kernel function such as a Gaussian kernel [20] that measures the similarity in time and location dimensions. $\mathcal{F}(t', l') \subseteq F$ denotes the feature values (e.g., event semantics) for the data at location l' and time t' . $g_\theta(\cdot)$ can be a deep neural network that is parameterized by θ and returns a nonnegative scalar. The model selection of $g_\theta(\cdot)$ depends on the specific data types. For example, these authors constructed an image attention network by combining a CNN with the spatial attention model proposed by Lu et al. [120].

3.4.3 Time, Location, and Semantics. In this section, we introduce the strategies that jointly predict the time, location, and semantics of future events, which can be grouped into either crowd wisdom-based methods, future event expression detection methods, and tensor-decomposition-based methods.

Crowd wisdom-based. The first type of the system-based methods considered here is the model-fusion system. The most intuitive approach is to leverage and integrate the aforementioned techniques for time, location, and semantics prediction into an event prediction system. For example, a system named EMBERS [144] is an online warning system for future events that can jointly predict the time, location, and semantics including the type and population of future events. This system also provides information on the confidence of the predictions obtained. Using an ensemble of predictive models for time [135], location, and semantic prediction, this system achieves a significant performance boost in terms of both precision and recall. The trick here is to first prioritize the precision of each individual prediction model by suppressing their recall. Then, due to the diversity and complementary nature of the different models, the fusion of the predictions from different models will eventually result in a high recall. A Bayesian fusion-based strategy has also been investigated [78]. Another system named *Carbon* [91] also leverages a similar strategy.

The second type of model involves crowd-sourced systems that implement fusion strategies to generate the event predictions made by the human predictors. For example, to handle the heterogeneity and diversity of the human predictors' skill sets and background knowledge under limited human resources, Rostami et al. [149] proposed a recommender system for matching event forecasting tasks to human predictors with suitable skills to maximize the accuracy of their fused predictions. Li et al. [108] took a different approach, designing a prediction market system that operates like a futures market, integrating information from different human predictors to forecast future events. In this system, the predictors can decide whether to buy or sell the "tokens" (using virtual dollars, for example) for each specific prediction they have made according to their confidence in it. They typically make careful decisions as they will obtain corresponding awards (for correct predictions) or penalties (for erroneous predictions).

Identification of future event expression. These methods focus on detecting the planned future events, usually from various media such sources as social media and news and typically relying on NLP techniques and linguistic principles. Existing methods typically follow a workflow

Fig. 8. Generic framework for the detection of the planning of future events.

similar to the one shown in Figure 8, consisting of four main steps: (1) *Content filtering*. Methods for content filtering are typically leveraged to retain only the texts that are relevant to the topic of interest. Existing works utilize either supervised methods (e.g., textual classifiers [99]) or unsupervised methods (e.g., querying techniques [129]); (2) *Time expression identification* is then utilized to identify future reference expressions and determine the *time to event*. These methods either leverage existing tools such as the Rosetta text analyzer [44] or propose dedicated strategies based on linguistic rules [84]; (3) *Future reference sentence extraction* is the core of the detection of the planning of future events, and is implemented either by designing regular expression-based rules [130] or by textual classification [99]; and (4) *Location identification*. The expression of locations is typically highly heterogeneous and noisy. Existing works have relied heavily on geocoding techniques that can resolve the event location accurately. To infer the event locations, various types of locations are considered, such as article locations, authors' profile locations, locations mentioned in the articles, and authors' neighbors' locations [90]. Multiple locations have been selected using a geometric median [39] or fused using logical rules such as probabilistic soft logic [129].

Tensor-decomposition-based methods. Some methods formulate the data into tensor-form, with dimensions including location, time, and semantics. Tensor decomposition is then applied to approximate the original tensors as the product of multiple low-rank matrices, each of which is a mapping from latent topics to each dimension. Finally, the tensor is extrapolated toward future time periods by various strategies. For example, Mirtaheri [127] extrapolated the time dimension-matrix only, which they then multiplied with the other dimensions' matrices to recover the estimated extrapolated tensor into the future. Zhou et al. [214] took a different approach, choosing instead to add "empty values" for the entries in future time to the original tensor, and then use tensor completion techniques to infer the missing values corresponding to future events.

4 APPLICATIONS OF EVENT PREDICTIONS

4.1 Healthcare

This category generally consists of two types of event prediction: (1) population level, for disease epidemics and outbreaks, and (2) individual level, for to clinical longitudinal events.

4.1.1 Population-level. There has been extensive research on disease outbreaks for many different types of diseases and epidemics, including seasonal flu, Zika, H1N1, Ebola, and COVID-19 [137]. These predictions target both the location and time of future events, while the disease type is usually fixed to a specific type for each model. Compartmental models such as SIR models are among the classical mathematical tools used to analyze, model, and simulate the epidemic dynamics [200]. More recently, individual-based computational models have begun to be used to perform network-based epidemiology based on network science and graph-theoretical models, where an epidemic is modeled as a stochastic propagation over an explicit interaction network among people [42]. Thanks to the availability of high-performance computing resources, another option is to construct a "digital twin" of the real world, by considering a realistic representation of a population, including members' demographic, geographic, behavioral, and social contextual information, and then using individual-based simulations to study the spread of epidemics within each network [21]. The above techniques heavily rely on the model assumptions regarding how the disease progresses individually and is transmitted from person to person [21]. The rapid growth

of large surveillance data from the government agencies (e.g., CDC flu map, www.cdc.gov) and other public dataset (e.g., COVID-19 dataset [188]) data and social media data sets (e.g., Twitter and Google flu trends [69]) in recent years has led to a massive increase of interest in using data-driven approaches to directly learn the predictive mapping [3]. These methods are usually both more time-efficient and less dependent on assumptions, while the aforementioned computational models are more powerful for longer-term prediction due to their ability to consider the specific disease mechanisms [204]. Finally, there are also synergistic research works combining both techniques' complementary strengths [81].

4.1.2 Individual-level. This research thread focuses on the longitudinal predictive analysis of individual health-related events, including death occurrence [51], adverse drug events [157], sudden illnesses such as strokes [106] and cardiovascular events [18], as well as other clinical events [51] and life events [48] for different groups of people, including elders and people with mental disease. The goal here is usually to predict the time before an event occurs, although some researchers have attempted to predict the type of event. The data sources are usually the electronic health records of individual patients [145] and other user-generated health-related reports [181]. Recently, social media, forum, and mobile data has also been utilized for predicting drug adverse events [157] and events that arise during chronic disease (e.g., chemical radiation and surgery) [51].

4.2 Media

This category focuses on predicting events based on information held in various types of media including: video-based, audio-based, and text-based formats. The core issue is to retrieve key information related to future events utilizing semantic pattern recognition from the data.

4.2.1 Video- and Audio-based. While event detection has been extensively researched for video data [111] and audio mining [164], event prediction is more challenging and has been attracting increasing attention in recent years. The goal here is usually to predict the future status of the objects in the video, such as the next action of soccer players [49] or basketball players [131], or the movement of vehicles [198].

4.2.2 Text- and Script-based. A huge amount of news data has accumulated in recent decades, such as Sina News [80], RSS News [144], and New York Times [143] datasets. Many of them have been used for big data predictive analytics among news events. A number of researchers have focused on predicting the location, time, and semantics of various events. To achieve this, they usually leverage the immense historical news and knowledge base to learn the association and causality among events, which is then applied to forecast events when given current events. Some studies have even directly generated textual descriptions of future events by leveraging NLP techniques such as sequence to sequence learning [80, 142, 143].

4.3 Transportation

This category can be classified into: (1) population-based events, including dispersal events, gathering events, and congestion; and (2) individual-level events, which focus on fine-grained patterns such as human mobility behavior prediction.

4.3.1 Group Transportation Pattern. Here, researchers typically focus on transportation events such as congestion [33], large gatherings [174], and dispersal events [175]. The goal is thus to forecast the future time period and location [174] of such events. Data from traffic meters, GPS, and mobile devices are usually used to sense real-time human mobility patterns. Popular datasets

include NYU Yellow Taxi datasets [175] and Didi Gaia dataset.² Transportation and geographical theories are usually considered to determine the spatial and temporal dependencies for predicting these events.

4.3.2 Individual Transportation Behavior. Another research thread focuses on individual-level prediction, such as predicting an individual's next location [112] or the likelihood or time duration of car accidents [15]. Sequential and trajectory analyses are usually used.

4.4 Engineering Systems

Different types of engineering systems have begun to routinely apply event forecasting methods, including: (1) civil engineering, (2) electrical engineering, (3) energy engineering, and (4) other engineering domains. Despite the variety of systems in these widely different domains, the goal is essentially to predict future abnormal or failure events to support the system's sustainability and robustness. Both the location and time of future events are key factors for these predictions. The input features usually consist of sensing data relevant to specific engineering systems. **(1) Civil engineering.** This covers various a wide range of problems in diverse urban systems, such as smart building fault adverse event prediction [17], emergency management equipment failure prediction [54], manhole event prediction [151]. **(2) Electrical engineering.** This includes teleservice systems failures [50] and unexpected events in wire electrical discharge machining operations [156]. **(3) Energy engineering.** Event prediction is also a hot topic in energy engineering, as such systems usually require strong robustness to handle the disturbance from the nature environments. Active research domains here include wind power ramp prediction [67], solar power ramp prediction [1], and adverse events in low carbon energy production [40]. **(4) Other engineering domains.** Other domains are also active, such as irrigation event prediction in agricultural engineering [136] and mechanical fault prediction in mechanical engineering [169]. The event datasets in this domain are typically extracted from the system log files or sensing data, which are not usually publicly available. Some exceptions are some hot topics such as solar power (e.g., Global Energy Forecasting Competition dataset and Probabilistic Solar Power Forecasting 2014 datasets [1]).

4.5 Cyber Systems

Here, the prediction models proposed generally focus on either network-level events or device-level events. For both types, the general goal is essentially to predict the likelihood of future system failure or attacks based on various indicators of system vulnerability. So far these two categories have essentially differed only in their inputs: the former relies on network features, including system specifications, web access logs and search queries, mismanagement symptoms, spam, phishing, and scamming activity, although some researchers are investigating the use of social media text streams to identify semantics indicating future potential attack targets of DDoS [121]. For device-level events, the features of interest are usually the binary file appearance logs of machines [135]. Some work has been done on micro-architectural attacks [73] by observing and proactively analyzing the observations on speculative branches, out-of-order executions and shared last level caches [159]. The event datasets in this domain are not commonly available to the public. Few benchmark datasets are available such as Phoronix benchmark suite [73] and ARCS datasets.³

²Didi Gaia dataset: <https://outreach.didichuxing.com/research/opendata/en/>.

³ARCS datasets: <https://csr.lanl.gov/data/>.

4.6 Political Events

Political event prediction has become a very active research area in recent years, largely thanks to the popularity of social media. They can be categorized into (1) offline events and (2) online activism.

4.6.1 Offline Events. This includes civil unrest, conflicts, violence, and riots [144]. This type of research usually targets the future events' geo-location, time, and topics by leveraging the social sensors that indicate public opinions and intentions. Utilization of social media has become a popular approach for these endeavors, as social media is a source of vital information during the event development stage [144]. Specifically, many aspects are clearly visible in social media, including complaints from the public (e.g., toward the government), discussions about their intentions regarding specific political events and targets, as well as advertisements for the planned events. Further information on future events such as the type of event [69], the anticipated participants population [144], and the event scale [68] can also be discovered in advance. Many datasets are available for political events such as ICEWS [144], GDELTS [194], OSI [128], TERRIER [74].

4.6.2 Online Events. Due to the major impact of online media such as online forums and social media, many events such as online activism, petitions, and hoaxes in such online platform also involve strong motivations for achieving some political purpose [180]. Beyond simple detection, the prediction of various types of events have been studied to enable proactive intervention to sidetrack the events such as hoaxes and rumor propagation [89]. Other researchers have sought to foresee the results of future political events to benefit a particular group of practitioners, for example by predicting the outcome of online petitions or presidential elections [180]. The datasets are popularly directly crawled from social media and forum datasets. Readers who are interested in detailed crawling techniques can refer to the book [153].

4.7 Natural Disasters

Different types of natural disasters have been the focus of a great deal of research. Typically, these are rare events, but mechanistic models, long historical records (often extending back dozens or hundreds of years), and domain knowledge are usually available. The input data are typically collected by sensors or sensor networks and the output is the risk or hazard of future potential events. Since these event occurrences are typically rare but very high-stakes, many researchers strive to cover all event occurrences and hence aim to ensure high recalls. They can be categorized according to subdomains, and geophysics-related, atmospheric science-related, and astrophysics-related are arguably the three most intensively researched areas in natural disaster event prediction. The geophysics-related domain focuses on the events on the surface of the earth, including earthquakes (e.g., occurrence and magnitude [150]), fire events (e.g., urban fires [160] and wildfires [198]), and rarer events like volcanic eruptions [34]. Atmospheric science-related domain concentrates on the events related to the atmosphere, such as tornado [55], floods [37], droughts and ozone events [62]. Finally, astrophysics-related domain is also a very active research area that predict events in the space, such as solar flares [16] and solar eruptions [4], with imagery data from satellites and telescopes. Natural event datasets are available in government websites such as USGS earthquake datasets (www.usgs.gov), China environment data (<http://www.wen.ipe.org.cn/about/about.aspx>), NCEP-NCAR Reanalysis and Solar Data Services datasets from NOAA (www.noaa.gov).

4.8 Business

Business intelligence can be grouped into company-based events and customer-based events.

4.8.1 Customer Activity Prediction. The most important customer activities in business is whether a customer will continue doing business with a company and how long a customer will be willing to wait before receiving the service? A great deal of research has been devoted to these topics, which can be categorized based on the type of business entities, namely, enterprises, social media, and education, who are primarily interested in churn prediction, site migration, and student dropout, respectively. The first of these focuses on predicting whether and when a customer is likely to stop doing business with a profitable enterprise [57]. The second aims to predict whether a social media user will move from one site, such as Flickr, to another, such as Instagram, a movement known as site migration [199]. While site migration is not popular, attention migration might actually be much more common, as a user may “move” their major activities from one social media site to another. The third type, student dropout, is a critical domain for education data mining, where the goal is to predict the occurrence of absenteeism from school for no good reason for a continuous number of days; a comprehensive survey is available in Reference [122]. For all three types, the procedure is first to collect features of a customer’s profile and activities over a period of time and then conventional or sequential classifiers or regressors are generally used to predict the occurrence or time-to-event of the future targeted activity. There are a few publicly available datasets, including Telecom Churn datasets [57] and MOOC dataset [122].

4.8.2 Business Process Events. Financial event prediction has been attracting a huge amount of attention for risk management, marketing, investment prediction and fraud prevention. Multiple information resources, including news, company announcements, and social media data could be utilized as the input, often taking the form of time series or temporal sequences. These sequential inputs are used for the prediction of the time and occurrence of future high-stack events such as company distress, suspension, mergers, dividends, layoffs, bankruptcy, and market trends (rises and falls in the company’s stock price) [28, 123, 191]. There are a few publicly available datasets, including CRSP/Compustat Merged Database [28], AAERs datasets [28], and BPI Challenge datasets [123].

4.9 Crime

It is difficult to deduce the precise location and time for individual crime incidences. Therefore, the focus is instead estimating the risk and probability of the location, time, and types of future crimes. This field can be naturally categorized based on the various crime types:

4.9.1 Political Crimes and Terrorism. This type of crime is typically highly destructive, and hence attracts huge attention in its anticipation and prevention. Terrorist activities are usually aimed at religious, political, iconic, economic or social targets. The attacker typically targets larger numbers of people and the evidences related to such attacks is retained in the long run. Though it is extremely challenging to predict the precise location and time of individual terrorism incidents, numerous studies have shown the potential utility for predicting the regional risks of terrorism attacks based on information gathered from many data sources such as geopolitical data, weather data, and economics data. The Global Terrorism Database [102] and UCDP dataset [165] are the most widely recognized datasets that record the descriptions of world-wide terrorism events of recent decades. In addition to terrorism/war events, other similar events such as mass killings [173] and armed-conflict events [165] have also been studied using similar problem formulations.

4.9.2 Crime Incidents and Hotspots. Most studies on this topic focus on predicting the types, intensity, hotspot, count, and probability of crime events across defined geo-spatial regions. The local police department is typically the venue to provide crime datasets. The geospatial characteristics of the spatial areas, their demographics, and temporal data such as news, weather, economics,

and social media data are usually used as inputs, datasets such as Foursquare datasets have been used to track the individual check-in behaviors [152]. The geospatial dependency and correlation of the crime patterns are usually leveraged during the prediction process using techniques originally developed for spatial predictions, such as kernel density estimation and conditional random fields, which are leveraged especially when predicting the crime hotspots [96]. Some works focus on specific types of crimes such as theft, robbery, and burglary [41, 152].

4.9.3 Organized and Serial Crimes. Unlike the above research on regional crime risks, some recent studies strive to predict the next incidents of criminal individuals or groups. This is because different offenders may demonstrate different behavioral patterns, such as targeting specific regions (e.g., wealthy neighborhoods), victims (e.g., women), for specific benefits (e.g., money). The goal here is thus to predict the next crime site and/or time, based on the historical crime event sequence of the targeted criminal individual or group. Models such as point processes [112] or Bayesian networks [115] are usually used to address such problems.

4.10 Summary and Remarks

This section has systematically introduced event prediction works and datasets in various applications. Notice that the utilization of event indicators from multiple data sources is very common, especially in the domains where the event occurrence is influenced by various social aspects such as civil unrest, business events, and transportation. Multi-source event prediction have been extensively researched to handle various challenges including multi-resolution (e.g., via multitask learning [203]), multi-modal (e.g., via ensemble models [78, 91, 135, 144]), different time durations (e.g., via multi-task learning [209]), multilingual (e.g., by semantic latent space learning [208]), multi-view (e.g., via convLSTM [15, 197] and multi-column CNN [97]), and multi-scale (e.g., via multi-level models [209]). Readers are welcome to refer to the above papers for technique details.

5 OPEN CHALLENGES AND FUTURE DIRECTIONS

Despite the major advances in event prediction in recent years, there are still a number of open problems and potentially fruitful directions for future research, as follows:

5.1 Model Transparency, Interpretability, and Accountability

Increasingly sophisticated forecasting models have been proposed to improve the prediction accuracy, including those utilizing approaches such as ensemble models, neural networks, and the other complex systems mentioned above. However, although the accuracy can be improved, the event prediction models are rapidly becoming too complex to be interpreted by human operators. The need for better model accountability and interpretability is becoming an important issue; as big data and Artificial Intelligence techniques are applied to ever more domains this can lead to serious consequences for applications such as healthcare and disaster management. Models that are not interpretable by humans will find it hard to build the trust needed if they are to be fully integrated into the workflow of practitioners. A closely related key feature is the accountability of the event prediction system. For example, disaster managers need to thoroughly understand a model's recommendations if they want to explain the reason for a decision to displace people in a court of law. An ever increasing number of laws in countries around the world are beginning to require adequate explanations of decisions from prediction models. For example, Articles 13-15 in the European Union's General Data Protection Regulation [177] require algorithms that make decisions that "significantly affect" individuals to provide explanations ("right to explanation") by May 28, 2018. The United States [38] and China [140] have also established similar laws.

5.2 Vulnerability to Noise and Adversarial Attacks

The massive popularity of the proposal, development, and deployment of event prediction is stimulating a surge interest in developing ways to counter-attack these systems. It will therefore not be a surprise when we begin to see the introduction of techniques to obfuscate these event prediction methods in the near future. As with many state-of-the-art AI techniques applied in other domains such as object recognition, event prediction methods can also be very vulnerable to noise and adversarial attacks. The famous failure of Google Flu trends, which missed the peak of the 2013 flu season by 140 percent due to low relevance and high disturbance affecting the input signal, is a vivid memory for practitioners in the field [66]. Many predictions relying on social media data can also be easily influenced or flipped by injecting scam messages. Event prediction models also tend to over-rely on low-quality input data that can be easily disturbed or manipulated, lacking sufficient robustness to survive noisy signals and adversarial attacks. Similar problems threaten to other application domains such as business intelligence, crime, and cyber systems.

5.3 Integration of Mechanistic Knowledge and Data Driven-models

Over the years, many domains have accumulated a significant amount of knowledge and experience about event development occurrence mechanisms, which can thus provide important clues for anticipating future events, such as epidemiology models, socio-political models, and earthquake models. All of these models focus on simplifying real-world phenomena into concise principles to grasp the core mechanism, discarding many details in the process. In contrast, data-driven models strive to ensure the accurate fitting of large historical data sets, based on sufficient model expressiveness but cannot guarantee that the true underlying principle and causality of event occurrence modeled accurately. There is thus a clear motivation to combine their complementary strengths, and although this has already attracted great deal of interest [204], most of the models proposed so far are merely ensemble learning-based and simply merge the final predictions from each model. A more thorough integration is needed that can directly embed the core principles to regularize and instruct the training of data-driven event prediction methods. Moreover, existing attempts are typically specific to particular domains and are thus difficult to develop further as they require in-depth collaborations between data scientists and domain experts. A generic framework developed to encompass multiple different domains is imperative and would be highly beneficial.

5.4 Prescriptive and Counterfactual Analysis

The ultimate purpose of event prediction is usually not just to anticipate the future, but to change it, for example by avoiding a system failure and flattening the curve of a disease outbreak. However, it is difficult for practitioners to determine how to act appropriately and implement effective policies to achieve the desired results in the future. This requires a capability that goes beyond simply predicting future events based on the current situation, requiring them instead to also take into account the new actions being taken in real time and then predict how they might influence the future. One promising direction is the use of counterfactual event [124] prediction that models what would have happened if different circumstances had occurred. Another related direction is prescriptive analysis where different actions can be merged into the prediction system and future results anticipated or optimized. Related works have been developed in few domains such as epidemiology. However, as yet these lack sufficient research in many other domains that will be needed if we are to develop generic frameworks that can benefit different domains.

5.5 Multi-objective Training

Existing event prediction methods mostly focus primarily on accuracy. However, decision makers who utilize these predicted event results usually need much more, including key factors such as

event resolution (e.g., time resolution, location resolution, description details), confidence (e.g., the probability a predicted event will occur), efficiency (whether the model can predict per day or per second), lead time (how many days the prediction can be made prior to the event occurring), and event intensity (how serious it is). multi-objective optimization (e.g., accuracy, confidence, resolution). There are typically trade-offs among all the above metrics and accuracy, so merely optimizing accuracy during training will inevitably mean the results drift away from the overall optimal event-prediction-based decision. A multi-objective optimization system that can flexibly balance the trade-off between these metrics based on decision makers' needs is imperative.

6 CONCLUSION

This survey has presented a comprehensive survey of existing methodologies developed for event prediction methods in the big data era. It provides an extensive overview of the event prediction challenges, techniques, applications, evaluation procedures, and future outlook, summarizing the research presented in over 200 publications, most of which were published in the last five years. Event prediction challenges, opportunities, and formulations have been discussed in terms of the event element to be predicted, including the event location, time, and semantics, after which, we went on to propose a systematic taxonomy of the existing event prediction techniques according to the formulated problems and types of methodologies designed for the corresponding problems. We have also analyzed the relationships, differences, advantages, and disadvantages of these techniques from various domains, including machine learning, data mining, pattern recognition, natural language processing, information retrieval, statistics, and other computational models. In addition, a comprehensive and hierarchical categorization of popular event prediction applications has been provided that covers domains ranging from natural science to the social sciences. Based upon the numerous historical and state-of-the-art works discussed in this survey, the article concludes by discussing open problems and future trends in this fast-growing domain.

REFERENCES

- [1] Mohamed Abuella and Badrul Chowdhury. 2019. Forecasting of solar power ramp events: A post-processing approach. *Renew. Energy* 133 (2019), 1380–1392.
- [2] Saurav Acharya, Byung Suk Lee, and Paul Hines. 2017. Causal prediction of top-k event types over real-time event streams. *Comput. J.* 60, 11 (2017), 1561–1581.
- [3] Bijaya Adhikari, Xinfeng Xu, Naren Ramakrishnan, and B. Aditya Prakash. 2019. Epideep: Exploiting embeddings for epidemic forecasting. In *Proceedings of SIGKDD*. 577–586.
- [4] Ashna Aggarwal, Nicole Schanche, Katharine K. Reeves, Dustin Kempton, and Rafal Angrzyk. 2018. Prediction of solar eruptions using filament metadata. *Astrophys. J. Suppl. Ser.* 236, 1 (2018), 15.
- [5] Mohammad Al Boni and Matthew S. Gerber. 2016. Area-specific crime prediction models. In *Proceedings of ICMLA*. IEEE, 671–676.
- [6] Hafiz A. Alaka, Lukumon O. Oyedele, Hakeem A. Owolabi et al. 2016. Methodological approach of construction business failure prediction studies: A review. *Construct. Manage. Econ.* 34, 11 (2016), 808–842.
- [7] Elias Alevizos, Alexander Artikis, and George Paliouras. 2017. Event forecasting with pattern Markov chains. In *Proceedings of DEBS*. 146–157.
- [8] Elias Alevizos, Alexander Artikis, and Georgios Paliouras. 2018. Wayeb: A tool for complex event forecasting. Retrieved from <https://arXiv:1901.01826>.
- [9] Elias Alevizos, Anastasios Skarlatidis, Alexander Artikis, and Georgios Paliouras. 2017. Probabilistic complex event recognition: A survey. *ACM Comput. Surveys* 50, 5 (2017), 1–31.
- [10] James Allan, Ron Papka, and Victor Lavrenko. [n.d.]. On-line new event detection and tracking. In *Proceedings of SIGIR*. 37–45.
- [11] M. Antunes, M. A. Amaral Turkman, and K. F. Turkman. 2003. A Bayesian approach to event prediction. *J. Time Ser. Anal.* 24, 6 (2003), 631–646.
- [12] Marta Arias, Argimiro Arratia, and Ramon Xuriguera. 2014. Forecasting with Twitter data. *Trans. Intell. Syst. Technol.* 5, 1 (2014), 1–24.

- [13] Farzindar Atefeh and Wael Khreich. 2015. A survey of techniques for event detection in Twitter. *Comput. Intell.* 31, 1 (2015), 132–164.
- [14] Gökhan Bakır, Thomas Hofmann et al. 2007. *Predicting Structured Data*. MIT Press.
- [15] Jie Bao, Pan Liu, and Satish V. Ukkusuri. 2019. A spatiotemporal deep learning approach for citywide short-term crash risk prediction with multi-source data. *Accident Anal. Prevent.* 122 (2019), 239–254.
- [16] Graham Barnes, K. D. Leka, C. J. Schrijver et al. 2016. A comparison of flare forecasting methods. I. Results from the “all-clear” workshop. *Astrophys. J.* 829, 2 (2016), 89.
- [17] Aniruddha Basak, Ole Mengshoel, Stefan Hosein, and Rodney Martin. 2016. Scalable causal learning for predicting adverse events in smart buildings. In *Proceedings of AAAI*.
- [18] Juan-Jose Beunza, Enrique Puertas, Ester García-Ovejero et al. 2019. Comparison of machine learning algorithms for clinical event prediction. *J. Biomed. Info.* 97 (2019), 103257.
- [19] Marin Biloš, Bertrand Charpentier, and Stephan Günnemann. 2019. Uncertainty on asynchronous time event prediction. In *Proceedings of NeurIPS*. 12831–12840.
- [20] Christopher M Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer.
- [21] Keith R. Bisset, Jiangzhuo Chen, Xizhou Feng, V. S. Anil Kumar, and Madhav V. Marathe. 2009. EpiFast: A fast algorithm for large scale realistic epidemic simulations on distributed memory systems. In *Proceedings of SC*. 430–439.
- [22] Robert A. Blair and Nicholas Sambanis. 2020. Forecasting civil wars: Theory and structure in an age of “big data” and machine learning. *J. Conflict Resolut.* 64, 10 (2020), 1885–1915.
- [23] Patrick T. Brandt, John R. Freeman, and Philip A. Schrodtt. 2011. Real time, time series forecasting of inter-and intra-state political conflict. *Conflict Manage. Peace Sci.* 28, 1 (2011), 41–64.
- [24] Jose Cadena, Gizem Korkmaz, Chris J. Kuhlman, Achla Marathe, Naren Ramakrishnan, and Anil Vullikanti. 2015. Forecasting social unrest using activity cascades. *PLoS One* 10, 6 (2015), e0128879.
- [25] Raffaella Calabrese and Johan A. Elkind. 2016. Estimating binary spatial autoregressive models for rare events. In *Advances in Econometrics*, vol. 37. 145–166.
- [26] Flávia Dias Casagrande, Jim Tørresen, and Evi Zouganeli. 2018. Sensor event prediction using recurrent neural network in smart homes for older adults. In *Proceedings of IS. IEEE*, 662–668.
- [27] Flávia Dias Casagrande and Evi Zouganeli. 2019. Prediction of next sensor event and its time of occurrence using transfer learning across homes. In *Proceedings of AICCSA. IEEE*, 1–8.
- [28] Mark Cecchini, Haldun Aytug, Gary J. Koehler, and Praveen Pathak. 2010. Making words work: Using financial text as a predictor of financial events. *Decis. Supp. Syst.* 50, 1 (2010), 164–175.
- [29] Pew Research Center. 2017. Social media fact sheet. *Pew Research Center, Washington, DC*.
- [30] Deepayan Chakrabarti and Kunal Punera. 2011. Event summarization using tweets. In *Proceedings of ICWSM*.
- [31] Feng Chen and Daniel B. Neill. 2014. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *Proceedings of SIGKDD*. 1166–1175.
- [32] Feng Chen, Baojian Zhou, Adil Alim, and Liang Zhao. 2017. A generic framework for interesting subspace cluster detection in multi-attributed networks. In *Proceedings of ICDM. IEEE*, 41–50.
- [33] Meng Chen, Xiaohui Yu, and Yang Liu. 2018. PCNN: Deep convolutional networks for short-term traffic congestion prediction. *IEEE Trans. Intell. Transport. Syst.* 19, 11 (2018), 3550–3559.
- [34] Seong-Pyo Cheon, Sungshin Kim, So-Young Lee, and Chong-Bum Lee. 2009. Bayesian networks based rare event prediction with sensor data. *Knowl.-Based Syst.* 22, 5 (2009), 336–343.
- [35] Chung-Wen Cho, Ying Zheng et al. 2008. A tree-based approach for event prediction using episode rules over event streams. In *Proceedings of ICDESA. Springer*, 225–240.
- [36] Yunjei Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of CVPR*. 8789–8797.
- [37] H. L. Cloke and Florian Pappenberger. 2009. Ensemble flood forecasting: A review. *J. Hydrol.* 375, 3–4 (2009), 613–626.
- [38] Cary Coglianese and David Lehr. 2016. Regulating by robot: Administrative decision making in the machine-learning era. *Geo. L. J.* 105 (2016), 1147.
- [39] Ryan Compton, Craig Lee, Jiejun Xu, Luis Artieda-Moncada, Tsai-Ching Lu, Lalindra De Silva, and Michael Macy. 2014. Using publicly visible social media to build detailed forecasts of civil unrest. *Secur. Info.* 3, 1 (2014), 4.
- [40] Stefano Coniglio, Anthony J. Dunn, and Alain B. Zemkoho. 2020. Infrequent adverse event prediction in low carbon energy production using machine learning. Retrieved from <https://arXiv:2001.06916>.
- [41] Bitzel Cortez, Berny Carrera, Young-Jin Kim, and Jae-Yoon Jung. 2018. An architecture for emergency event prediction using LSTM recurrent neural networks. *Expert Syst. Appl.* 97 (2018), 315–324.
- [42] Meggan E. Craft, Erik Volz, Craig Packer, and Lauren Ancel Meyers. 2011. Disease transmission in territorial populations: The small-world network of Serengeti lions. *J. Roy. Soc. Interface* 8, 59 (2011), 776–786.

- [43] Alessandro Daidone, Felicita Di Giandomenico, Andrea Bondavalli, and Silvano Chiaradonna. 2006. Hidden Markov models as a support for diagnosis: Formalization of the problem and synthesis of the solution. In *Proceedings of SRDS*. IEEE, 245–256.
- [44] Robert Dale. 2018. Text analytics APIs. *Nat. Lang. Eng.* 24, 5 (2018), 797–803.
- [45] Magret Damaschke, Shane J. Cronin, and Mark S. Bebbington. 2018. A volcanic event forecasting model for multiple tephra records, demonstrated on Mt. Taranaki, New Zealand. *Bull. Volcanol.* 80, 1 (2018), 9.
- [46] Sina Dami, Ahmad Abdollahzadeh Barforoush, and Hossein Shirazi. 2018. News events prediction using Markov logic networks. *J. Info. Sci.* 44, 1 (2018), 91–109.
- [47] Arno De Caigny, Kristof Coussement, and Koen De Bock. 2018. A new hybrid classification algorithm for customer churn prediction based on logistic regression and decision trees. *Eur. J. Operation. Res.* 269, 2 (2018), 760–772.
- [48] Arno De Caigny, Kristof Coussement, and Koen De Bock. 2020. Leveraging fine-grained transaction data for customer life event predictions. *Decis. Supp. Syst.* 130 (2020), 113232.
- [49] Tom Decroos, Vladimir Dzyuba, Jan Van Haaren, and Jesse Davis. 2017. Predicting soccer highlights from spatio-temporal match event streams. In *Proceedings of AAAI*.
- [50] Akash Deep, Dharmaraj Veeramani, and Shiyu Zhou. 2019. Event prediction for individual unit based on recurrent event data collected in teleservice systems. *IEEE Trans. Reliabil.* 69, 1 (2019), 216–227.
- [51] Walter H. Dempsey, Alexander Moreno, Christy K. Scott et al. 2017. iSurvive: An interpretable, event-time prediction model for mHealth. In *Proceedings of ICML*. 970–979.
- [52] Jingsheng Deng, Fengcai Qiao et al. 2015. An overview of event extraction from Twitter. In *Proceedings of CyberC*. IEEE, 251–256.
- [53] Xiaolei Di, Yu Xiao, Chao Zhu, Yang Deng, Qinpei Zhao, and Weixiong Rao. 2019. Traffic congestion prediction by spatiotemporal propagation patterns. In *Proceedings of MDM*. IEEE, 298–303.
- [54] Zuohua Ding, Yuan Zhou, Geguang Pu, and MengChu Zhou. 2018. Online failure prediction for railway transportation systems based on fuzzy rules and data analysis. *IEEE Trans. Reliabil.* 67, 3 (2018), 1143–1158.
- [55] Charles A. Doswell III, Steven J. Weiss, and Robert H. Johns. 1993. Tornado forecasting: A review. *Tornado: Struct. Dynam. Predict. Haz. Geophys. Monogr.* 79 (1993), 557–571.
- [56] Huilong Duan, Zhoujian Sun, Wei Dong et al. 2019. On clinical event prediction in patient treatment trajectory using longitudinal electronic health records. *IEEE J. Biomed. Health Info.* 24, 7 (2019), 2053–2063.
- [57] Kamy Eria and Booma Poolan Marikannan. 2018. Systematic review of customer churn prediction in the telecom sector. *J. Appl. Technol. Innovat.* 2, 1 (2018).
- [58] Şeyda Ertekin, Cynthia Rudin, Tyler H. McCormick et al. 2015. Reactive point processes: A new approach to predicting power failures in underground electrical systems. *Ann. Appl. Stat.* 9, 1 (2015), 122–144.
- [59] Ali Mert Ertugrul, Yu-Ru Lin, Christina Mair, and T. Taskaya Temizel. 2018. Forecasting heroin overdose occurrences from crime incidents. In *Proceedings of SBP-BRIMS*.
- [60] Christiane Fellbaum. 2012. WordNet. *The Encyclopedia of Applied Linguistics*.
- [61] Carlos Ferreira, Joao Gama, L. Matias, Audun Botterud, and J. Wang. 2011. *A Survey on Wind Power Ramp Forecasting*. Technical Report. Argonne National Laboratory, Argonne, IL.
- [62] Christopher B. Field, Vicente Barros et al. 2012. *Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation: Special Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- [63] Ioannis Flouris, Nikos Giatrakos, Antonios Deligiannakis et al. 2017. Issues in complex event processing: Status and prospects in the big data era. *J. Syst. Softw.* 127 (2017), 217–236.
- [64] Ilenia Fronza, Alberto Sillitti, Giancarlo Succi, Mikko Terho, and Jelena Vlasenko. 2013. Failure prediction based on log files using random indexing and support vector machines. *J. Syst. Softw.* 86, 1 (2013), 2–11.
- [65] Lajos Jenő Fülöp, Gabriella Tóth, Róbert Rácz, János Pánczél et al. 2010. Survey on complex event processing and predictive analytics. In *Proceedings of BCI*. Citeseer, 26–31.
- [66] Kaiser Fung. 2014. Google Flu Trends' failure shows good data > big data. *Harvard Bus. Rev.* 25 (2014), 2014.
- [67] Cristobal Gallego-Castillo, Alvaro Cuerva-Tejero, and Oscar Lopez-Garcia. 2015. A review on the recent history of wind power ramp forecasting. *Renew. Sustain. Energy Rev.* 52 (2015), 1148–1157.
- [68] Yuyang Gao and Liang Zhao. 2018. Incomplete label multi-task ordinal regression for spatial event scale forecasting. In *Proceedings of AAAI*.
- [69] Yuyang Gao, Liang Zhao, Lingfei Wu, Yanfang Ye, Hui Xiong, and Chaowei Yang. 2019. Incomplete label multi-task deep learning for spatio-temporal event subtype forecasting. In *Proceedings of AAAI*, Vol. 33. 3638–3646.
- [70] Fatma Gmati, Salem Chakhar, Wided Lejouad Chaari, and Mark Xu. 2019. A taxonomy of event prediction methods. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*. 12–26.
- [71] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press.

- [72] Mark Granroth-Wilding and Stephen Clark. 2016. What happens next? Event prediction using a compositional neural network model. In *Proceedings of AAAI*. 2727–2733.
- [73] Berk Gulmezoglu, Ahmad Moghimi, Thomas Eisenbarth, and Berk Sunar. 2019. Fortuneteller: Predicting microarchitectural attacks via unsupervised deep learning. Retrieved from <https://arXiv:1907.03651>.
- [74] Andrew Halterman, Jill Irvine et al. 2017. Adaptive scalable pipelines for political event data generation. In *Proceedings of Big Data*. IEEE, 2879–2883.
- [75] Jiawei Han, Jian Pei, and Micheline Kamber. 2011. *Data Mining: Concepts and Techniques*. Elsevier.
- [76] Mengmeng Hao, Dong Jiang, Fangyu Ding, Jingying Fu, and Shuai Chen. 2019. Simulating spatio-temporal patterns of terrorism incidents on the Indochina Peninsula with GIS and the random forest method. *Int. J. Geo-Info*. 8, 3 (2019), 133.
- [77] Chikara Hashimoto, Kentaro Torisawa, Julien Kloetzer et al. 2014. Toward future scenario generation: Extracting event causality exploiting semantic relation, context, and association features. In *Proceedings of ACL*. 987–997.
- [78] Andrew Hoegh, Scotland Leman et al. 2015. Bayesian model fusion for forecasting civil unrest. *Technometrics* 57, 3 (2015), 332–340.
- [79] Linmei Hu. 2019. Integrating hierarchical attentions for future subevent prediction. *IEEE Access* 8 (2019), 3106–3114.
- [80] Linmei Hu, Juanzi Li, Liqiang Nie, Xiao-Li Li, and Chao Shao. 2017. What happens next? Future subevent prediction using contextual hierarchical LSTM. In *Proceedings of AAAI*.
- [81] Ting Hua, Chandan K. Reddy, Lei Zhang, Lijing Wang, Liang Zhao, Chang-Tien Lu, and Naren Ramakrishnan. 2018. Social media based simulation models for understanding disease dynamics. In *Proceedings of IJCAI*. 3797–3804.
- [82] Chao Huang, Chuxu Zhang, Jiashu Zhao, Xian Wu, Dawei Yin, and Nitesh Chawla. 2019. Mist: A multiview and multimodal spatial-temporal learning framework for citywide abnormal event forecasting. In *Proceedings of WWW*. 717–728.
- [83] Gordon F. Hughes, Joseph F. Murray, Kenneth Kreutz-Delgado, and Charles Elkan. 2002. Improved disk-drive failure warnings. *IEEE Trans. Reliabil.* 51, 3 (2002), 350–357.
- [84] Ali Hürriyetoğlu, Nelleke Oostdijk, and Antal van den Bosch. 2018. Estimating time to event of future events based on linguistic cues on Twitter. In *Intelligent Natural Language Processing: Trends and Applications*. 67–97.
- [85] Fadil Inceoglu, Jacob H. Jeppesen, Peter Kongstad et al. 2018. Using machine learning methods to forecast if solar flares will be associated with CMEs and SEPs. *Astrophys. J.* 861, 2 (2018), 128.
- [86] Bram Jans, Steven Bethard, Ivan Vulić, and Marie Francine Moens. 2012. Skip n-grams and ranking functions for predicting script events. In *Proceedings of EACL*. 336–344.
- [87] Renhe Jiang, Xuan Song, Dou Huang, Xiaoya Song, Tianqi Xia et al. 2019. Deepurbanevent: A system for predicting citywide crowd dynamics at big events. In *Proceedings of SIGKDD*. 2114–2122.
- [88] Zhe Jiang. 2018. A survey on spatial prediction methods. *Trans. Knowl. Data Eng.* 31, 9 (2018), 1645–1664.
- [89] Fang Jin, Edward Dougherty, Parang Saraf, Yang Cao, and Naren Ramakrishnan. 2013. Epidemiological modeling of news and rumors on Twitter. In *Proceedings of SNAKDD*. 1–9.
- [90] David Jurgens. 2013. That’s what friends are for: Inferring location in online social media platforms based on social relationships. In *Proceedings of ICWSM*.
- [91] Wei Kang, Jie Chen, Jiuyong Li, Jixue Liu, Lin Liu, Grant Osborne, Nick Lothian, Brenton Cooper, Terry Moschou, and Grant Neale. 2017. Carbon: Forecasting civil unrest events by monitoring news and social media. In *Proceedings of ADMA*. Springer, 859–865.
- [92] Ahmed Kattan, Shaheen Fatima, and Muhammad Arif. 2015. Time-series event-based prediction: An unsupervised learning framework based on genetic programming. *Info. Sci.* 301 (2015), 99–123.
- [93] Amin Vahedian Khezerlou, Xun Zhou, Lufan Li, Zubair Shafiq, Alex X. Liu, and Fan Zhang. 2017. A traffic flow approach to early detection of gathering events: Comprehensive results. *Trans. Intell. Syst. Technol.* 8, 6 (2017), 1–24.
- [94] Amin Vahedian Khezerlou, Xun Zhou, Ling Tong, Yanhua Li, and Jun Luo. 2019. Forecasting gathering events through trajectory destination prediction: A dynamic hybrid model. *Trans. Knowl. Data Eng.* (2019).
- [95] Hirokazu Kiyomaru, Kazumasa Omura et al. 2019. Diversity-aware event prediction based on a conditional variational autoencoder with reconstruction. In *Proceedings of COIN*. 113–122.
- [96] Ourania Kounadi et al. 2020. A systematic review on spatial crime forecasting. *Crime Sci.* 9, 1 (2020), 1–22.
- [97] Canasai Kruengkrai, Kentaro Torisawa et al. 2017. Improving event causality recognition with multiple background knowledge sources using multi-column convolutional neural networks. In *Proceedings of AAAI*. 3466–3473.
- [98] Martin Kulldorff. 1997. A spatial scan statistic. *Commun. Stat. Theory Methods* 26, 6 (1997), 1481–1496.
- [99] Florian A. Kunneman and A. P. J. van den Bosch. 2012. Leveraging unscheduled event prediction through mining scheduled event tweets. In *the 24th Benelux Conference on Artificial Intelligence (BNAIC’12)*. 147.
- [100] Matthew Kupilik and Frank Witmer. 2018. Spatio-temporal violent event prediction using Gaussian process regression. *J. Comput. Soc. Sci.* 1, 2 (2018), 437–451.
- [101] Håvard Kvamme, Ørnulf Borgan, and Ida Scheel. 2019. Time-to-event prediction with neural networks and Cox regression. *J. Mach. Learn. Res.* 20, 129 (2019), 1–30.

- [102] Gary LaFree and Laura Dugan. 2007. Introducing the global terrorism database. *Terror. Politic. Viol.* 19, 2 (2007), 181–204.
- [103] Daniel T. Larose. 2015. *Data Mining and Predictive Analytics*. John Wiley & Sons.
- [104] Srivatsan Laxman, Vikram Tankasali, and Ryen W White. 2008. Stream prediction using a generative model based on frequent episodes in event sequences. In *Proceedings of SIGKDD*. 453–461.
- [105] Lei Lei, Xuguang Ren, Nigel Franciscus, Junhu Wang, and Bela Stantic. 2019. Event prediction based on causality reasoning. In *Proceedings of ACIIDS*. Springer, 165–176.
- [106] Benjamin Letham, Cynthia Rudin et al. 2015. Interpretable classifiers using rules and Bayesian analysis: Building a better stroke prediction model. *Ann. Appl. Stat.* 9, 3 (2015), 1350–1371.
- [107] Benjamin Letham, Cynthia Rudin, and David Madigan. 2013. Sequential event prediction. *Mach. Learn.* 93, 2-3 (2013), 357–380.
- [108] Eldon Y. Li, Chen-Yuan Tung, and Shu-Hsun Chang. 2016. The wisdom of crowds in action: Forecasting epidemic diseases with a web-based prediction market system. *Int. J. Med. Info.* 92 (2016), 35–43.
- [109] Jundong Li, Kewei Cheng, Suhang Wang, Fred Morstatter, Robert P. Trevino, Jiliang Tang, and Huan Liu. 2017. Feature selection: A data perspective. *ACM Comput. Surveys* 50, 6 (2017), 1–45.
- [110] Shengzhi Li, Jianzhong Qiao, and Shukuan Lin. 2017. Multi-attribute event modeling and prediction over event streams from sensors. In *Proceedings of ICPADS*. IEEE, 796–799.
- [111] Yang Li, Nan Du, and Samy Bengio. 2017. Time-dependent representation for neural event sequence prediction. Retrieved from <https://arXiv:1708.00065>.
- [112] Yunyi Li and Tong Wang. [n.d.]. Next hit predictor-self-exciting risk modeling for predicting next locations of serial crimes. In *Proceedings of NIPS*.
- [113] Zhongyang Li, Xiao Ding, and Ting Liu. 2018. Constructing narrative event evolutionary graph for script event prediction. Retrieved from <https://arXiv:1805.05081>.
- [114] Zhiguo Li, Shiyu Zhou, Suresh Choubey, and Crispian Sievenpiper. 2007. Failure event prediction using the Cox proportional hazard model driven by frequent failure signatures. *IEE Trans.* 39, 3 (2007), 303–315.
- [115] Renjie Liao, Xueyao Wang, Lun Li, and Zengchang Qin. 2010. A novel serial crime prediction model based on Bayesian learning theory. In *Proceedings of ICMLC*. 1757–1762.
- [116] Li Lin, Lijie Wen, and Jianmin Wang. 2019. Mm-pred: A deep predictive model for multi-attribute event sequence. In *Proceedings of ICDM*. SIAM, 118–126.
- [117] Ying-Lung Lin, Meng-Feng Yen, and Liang-Chih Yu. 2018. Grid-based crime prediction using geographical features. *ISPRS Int. J. Geo-Inform.* 7, 8 (2018), 298.
- [118] Hua Liu and Donald E. Brown. 2004. A new point process transition density model for space-time event prediction. *IEEE Trans. Syst. Man Cybernet., Part C* 34, 3 (2004), 310–324.
- [119] Hugo Liu and Push Singh. 2004. ConceptNet—A practical commonsense reasoning tool-kit. *BT Technol. J.* 22, 4 (2004), 211–226.
- [120] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. 2017. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Proceedings of CVPR*. 375–383.
- [121] Yasuko Matsubara, Yasushi Sakurai, Christos Faloutsos, Tomoharu Iwata, and Masatoshi Yoshikawa. 2012. Fast mining and forecasting of complex time-stamped events. In *Proceedings of SIGKDD*. 271–279.
- [122] Neema Mduma, Khamisi Kalegele, and Dina Machuve. 2019. A survey of machine learning approaches and techniques for student dropout prediction. *Data Sci. J.* 18, 1 (2019).
- [123] Nijat Mehdiyev, Joerg Evermann, and Peter Fettke. 2017. A multi-stage deep learning approach for business process event prediction. In *Proceedings of CBI*. IEEE, 119–128.
- [124] Peter Menzies. 2001. Counterfactual theories of causation. In *The Stanford Encyclopedia of Philosophy*. Stanford University, 1–7.
- [125] Evangelos Michelioudakis, Anastasios Skarlatidis, Elias Alevizos et al. 2016. Event recognition and forecasting technology. *Technical Report*.
- [126] Bryan Minor and Diane Cook. 2017. Forecasting occurrences of activities. *Pervas. Mobile Comput.* 38 (2017), 77–91. DOI: <https://doi.org/10.1016/j.pmcj.2016.09.010>
- [127] Mehnoosh Mirtaheri, Sami Abu-El-Haija, Tozammel Hossain et al. [n.d.]. Tensor-based method for temporal geopolitical event forecasting. In *Proceedings of ICML*.
- [128] Sathappan Muthiah, Patrick Butler, Rupinder Paul Khandpur, Parang Saraf et al. 2016. EMBERS at 4 years: Experiences operating an open source indicators forecasting system. In *Proceedings of SIGKDD*. 205–214.
- [129] Sathappan Muthiah, Bert Huang, Jaime Arredondo, David Mares, Lise Getoor, Graham Katz, and Naren Ramakrishnan. 2016. Capturing planned protests from open source indicators. *AI Mag.* 37, 2 (2016), 63–75.
- [130] Yoko Nakajima, Michal Ptaszynski, F. Masui, and H. Hirotschi. 2017. A prototype method for future event prediction based on future reference sentence extraction. In *Proceedings of IJCAI*. 42–49.

- [131] Lukas Neumann, Andrew Zisserman, and Andrea Vedaldi. 2019. Future event prediction: If and when. In *Proceedings of CVPR*.
- [132] 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 *Proceedings of SDM*. SIAM, 99–107.
- [133] Yue Ning, Liang Zhao, Feng Chen, Chang-Tien Lu, and Huzefa Rangwala. 2019. Spatio-temporal event forecasting and precursor identification. In *Proceedings of SIGKDD*. 3237–3238.
- [134] Maya Okawa, Tomoharu Iwata, Takeshi Kurashima, Yusuke Tanaka et al. 2019. Deep mixture point processes: Spatio-temporal event prediction with rich contextual information. In *Proceedings of SIGKDD*. 373–383.
- [135] Motoyuki Oki, Koh Takeuchi, and Yukio Uematsu. 2018. Mobile network failure event detection and forecasting with multiple user activity data sets. In *Proceedings of AAAI*.
- [136] R. González Perea, E. Camacho Poyato, Pilar Montesinos, and J. A. Rodríguez Díaz. 2019. Prediction of irrigation event occurrence at farm level using optimal decision trees. *Comput. Electron. Agric.* 157 (2019), 173–180.
- [137] Fotios Petropoulos and Spyros Makridakis. 2020. Forecasting the novel coronavirus COVID-19. *PloS One* 15, 3 (2020), e0231236.
- [138] Lawrence Phillips, Chase Dowling, Kyle Shaffer, Nathan Hodas, and Svitlana Volkova. 2017. Using social media to predict the future: A systematic literature review. Retrieved from <https://arXiv:1706.06134>.
- [139] Freddy Piraján, Andrey Fajardo, and Miguel Melgarejo. 2019. Towards a deep learning approach for urban crime forecasting. In *Proceedings of WEA*. Springer, 179–189.
- [140] Aimin Qi, Guosong Shao, and Wentong Zheng. 2018. Assessing China’s cybersecurity law. *Comput. Law Secur. Rev.* 34, 6 (2018), 1342–1354.
- [141] Zhi Qiao, Shiwan Zhao, Cao Xiao, Xiang Li, Yong Qin, and Fei Wang. 2018. Pairwise-ranking based collaborative recurrent neural networks for clinical event prediction. In *Proceedings of IJCAI*. 3520–3526.
- [142] Kira Radinsky, Sagie Davidovich, and Shaul Markovitch. 2012. Learning causality for news events prediction. In *Proceedings of WWW*. 909–918.
- [143] Kira Radinsky and Eric Horvitz. 2013. Mining the web to predict future events. In *Proceedings of WSDM*. 255–264.
- [144] Naren Ramakrishnan, Patrick Butler, Sathappan Muthiah, Nathan Self et al. 2014. “Beating the news” with EMBERS: Forecasting civil unrest using open source indicators. In *Proceedings of SIGKDD*. 1799–1808.
- [145] Jonathan Rebane, Isak Karlsson, and Panagiotis Papapetrou. 2019. An investigation of interpretable deep learning for adverse drug event prediction. In *Proceedings of CBMS*. 337–342.
- [146] D. Reid, A. Jaafar Hussain, H. Tawfik, R. Ghazali, and D. Al-Jumeily. 2018. Forecasting natural events using axonal delay. In *Proceedings of CEC*. 1–6.
- [147] Honglei Ren, You Song, Jingwen Wang, Yucheng Hu, and Jinzhi Lei. 2018. A deep learning approach to the citywide traffic accident risk prediction. In *Proceedings of ITSC*. 3346–3351.
- [148] Alina Ristea, Mohammad Al Boni, Bernd Resch, Matthew S. Gerber, and Michael Leitner. 2020. Spatial crime distribution and prediction for sporting events using social media. *Int. J. Geogr. Info. Sci.* (2020), 1–32.
- [149] Mohammad Rostami, David Huber, and Tsai-Ching Lu. 2018. A crowdsourcing triage algorithm for geopolitical event forecasting. In *Proceedings of RecSys*. 377–381.
- [150] Bertrand Rouet-Leduc, Claudia Hulbert, Nicholas Lubbers, Kipton Barros, Colin J. Humphreys, and Paul A. Johnson. 2017. Machine learning predicts laboratory earthquakes. *Geophys. Res. Lett.* 44, 18 (2017), 9276–9282.
- [151] Cynthia Rudin, Rebecca J. Passonneau, Axinia Radeva, Haimonti Dutta, Steve Jerome, and Delfina Isaac. 2010. A process for predicting manhole events in Manhattan. *Mach. Learn.* 80, 1 (2010), 1–31.
- [152] Shakila Khan Rumi, Ke Deng, and Flora D Salim. 2018. Theft prediction with individual risk factor of visitors. In *Proceedings of SIGSPATIAL*. 552–555.
- [153] Matthew A. Russell. 2011. *Mining the Social Web: Analyzing Data from Facebook, Twitter, LinkedIn, and Other Social Media Sites*. O’Reilly Media.
- [154] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. 2010. Earthquake shakes Twitter users: Real-time event detection by social sensors. In *Proceedings of WWW*. 851–860.
- [155] Felix Salfner, Maren Lenk, and Mirosław Malek. 2010. A survey of online failure prediction methods. *ACM Comput. Surveys* 42, 3 (2010), 1–42.
- [156] Jose A. Sanchez, Aintzane Conde, Ander Arriandiaga, Jun Wang, and Soraya Plaza. 2018. Unexpected event prediction in wire electrical discharge machining using deep learning techniques. *Materials* 11, 7 (2018), 1100.
- [157] Sara Santiso, Arantza Casillas, and Pothers. 2014. Adverse drug event prediction combining shallow analysis and machine learning. In *Proceedings of LOUHI*. 85–89.
- [158] 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 *Proceedings of WWW*. 1631–1639.
- [159] Yun Shen, Enrico Mariconti, Pierre Antoine Vervier, and Gianluca Stringhini. 2018. Tiresias: Predicting security events through deep learning. In *Proceedings of SIGSAC*. 592–605.

- [160] Seung Yeop Shin and Han-joon Kim. 2019. Autoencoder-based one-class classification technique for event prediction. In *Proceedings of CCIOT*. 54–58.
- [161] Prajwal Shrestha, Byung Suk Lee, and James P. Bagrow. 2017. Predicting an effect event from a new cause event using a semantic web based abstraction tree of past cause-effect event pairs. In *Proceedings of SIMBig*. 198–207.
- [162] Robert H. Shumway and David S. Stoffer. 2017. *Time Series Analysis and Its Applications: With R Examples*. Springer.
- [163] Aleksandr Simma and Michael I. Jordan. 2012. Modeling events with cascades of Poisson processes. Retrieved from <https://arXiv:1203.3516>.
- [164] Hagen Soltau, Hank Liao, and Haşim Sak. 2017. Neural speech recognizer: Acoustic-to-word LSTM model for large vocabulary speech recognition. *Proc. Interspeech 2017* (2017), 3707–3711.
- [165] Michael Spagat, Neil F. Johnson, and Stijn van Weezel. 2018. Fundamental patterns and predictions of event size distributions in modern wars and terrorist campaigns. *PLoS One* 13, 10 (2018), e0204639.
- [166] Narayan Srinivasa, Qin Jiang, and Leandro G Barajas. 2008. High-impact event prediction by temporal data mining through genetic algorithms. In *Proceedings of ICNC*. 614–620.
- [167] Zichun Su and Jialin Jiang. 2020. Hierarchical gated recurrent unit with semantic attention for event prediction. *Future Internet* 12, 2 (2020), 39.
- [168] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. YAGO: A core of semantic knowledge. In *Proceedings of WWW*. 697–706.
- [169] Gian Antonio Susto, Andrea Schirru, Simone Pampuri, Seán McLoone et al. 2014. Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Trans. Industr. Info.* 11, 3 (2014), 812–820.
- [170] Bayu Adhi Tama and Marco Comuzzi. 2019. An empirical comparison of classification techniques for next event prediction using business process event logs. *Expert Syst. Appl.* 129 (2019), 233–245.
- [171] James W Taylor. 2017. Probabilistic forecasting of wind power ramp events using autoregressive logit models. *Eur. J. Operation. Res.* 259, 2 (2017), 703–712.
- [172] Hannah Tops, Antal Van Den Bosch, and Florian Kunneman. 2013. Predicting time-to-event from Twitter messages. In *Proceedings of BNAIC*. 207–214.
- [173] Jay Ulfelder. 2013. A multimodel ensemble to forecast onsets of state-sponsored mass killing. In *Proceedings of APSA*.
- [174] Amin Vahedian, Xun Zhou, Ling Tong, Yanhua Li, and Jun Luo. 2017. Forecasting gathering events through continuous destination prediction on big trajectory data. In *Proceedings of SIGSPATIAL*. 1–10.
- [175] Amin Vahedian, Xun Zhou, Ling Tong, W Nick Street, and Yanhua Li. 2019. Predicting urban dispersal events: A two-stage framework through deep survival analysis on mobility data. In *Proceedings of AAAI*, Vol. 33. 5199–5206.
- [176] Ricardo Vilalta and Sheng Ma. 2002. Predicting rare events in temporal domains. In *Proceedings of ICDM*. IEEE, 474–481.
- [177] Paul Voigt and Axel Von dem Bussche. 2017. *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed. Springer International Publishing, Cham.
- [178] Dawei Wang and Wei Ding. 2015. A hierarchical pattern learning framework for forecasting extreme weather events. In *Proceedings of ICDM*. IEEE, 1021–1026.
- [179] Dawei Wang, Wei Ding, Kui Yu, Xindong Wu et al. 2013. Towards long-lead forecasting of extreme flood events: A data mining framework for precipitation cluster precursors identification. In *Proceedings of SIGKDD*. 1285–1293.
- [180] Junxiang Wang, Yuyang Gao, Andreas Züfle, Jingyuan Yang, and Liang Zhao. 2018. Incomplete label uncertainty estimation for petition victory prediction with dynamic features. In *Proceedings of ICDM*. IEEE, 537–546.
- [181] Junxiang Wang and Liang Zhao. 2018. Multi-instance domain adaptation for vaccine adverse event detection. In *Proceedings of WWW*. 97–106.
- [182] Mingjun Wang and Matthew S. Gerber. 2015. Using Twitter for next-place prediction, with an application to crime prediction. In *Proceedings of SSCI*. IEEE, 941–948.
- [183] Qi Wang, Guangyin Jin, Xia Zhao, Yanghe Feng, and Jincai Huang. 2020. CSAN: A neural network benchmark model for crime forecasting in spatio-temporal scale. *Knowl.-Based Syst.* 189 (2020), 105120.
- [184] Jeremy C. Weiss and David Page. 2013. Forest-based point process for event prediction from electronic health records. In *Proceedings of ECML PKDD*. Springer, 547–562.
- [185] Chuanxiu Xiong, Ajitesh Srivastava, Rajgopal Kannan, Omkar Damle et al. 2019. On predicting crime with heterogeneous spatial patterns: Methods and evaluation. In *Proceedings of SIGSPATIAL*. 43–51.
- [186] Cong Xue, Zehua Zeng, Yuanye He, Lei Wang, and Neng Gao. 2018. A MIML-LSTM neural network for integrated fine-grained event forecasting. In *Proceedings of Big Data*. 44–51.
- [187] Kazuo Yamaguchi. 1991. *Event History Analysis*. Sage.
- [188] Chaowei Yang, Dexuan Sha et al. 2020. Taking the pulse of COVID-19: A spatiotemporal perspective. *Int. J. Dig. Earth* 13, 10 (2020), 1186–1211.
- [189] Jaewon Yang, Julian McAuley, Jure Leskovec, Paea LePendur, and Nigam Shah. 2014. Finding progression stages in time-evolving event sequences. In *Proceedings of WWW*. 783–794.

- [190] Qiang Yang, Hui Wang, and Wei Zhang. 2002. Web-log mining for quantitative temporal-event prediction. *IEEE Comput. Intell. Bull.* 1, 1 (2002), 10–18.
- [191] Yiyang Yang, Zhongyu Wei, Qin Chen, and Libo Wu. 2019. Using external knowledge for financial event prediction based on graph neural networks. In *Proceedings of CIKM*. 2161–2164.
- [192] Fei Yi, Zhiwen Yu, Fuzhen Zhuang, and Bin Guo. 2019. Neural network based continuous conditional random field for fine-grained crime prediction. In *Proceedings of IJCAI*. 4157–4163.
- [193] Fei Yi, Zhiwen Yu, Fuzhen Zhuang, Xiao Zhang, and Hui Xiong. 2018. An integrated model for crime prediction using temporal and spatial factors. In *Proceedings of ICDM*. IEEE, 1386–1391.
- [194] James Yonamine. 2013. Predicting future levels of violence in afghanistan districts using Gdelt. Retrieved on April 22, 2021 from <http://data.gdeltproject.org/documentation/Predicting-Future-Levels-of-Violence-in-Afghanistan-Districts-using-GDELT.pdf>.
- [195] Shuqi Yu, Linmei Hu, and Bin Wu. 2019. DRAM: A deep reinforced intra-attentive model for event prediction. In *Proceedings of KSEM*. Springer, 701–713.
- [196] Sheng Yu and Subhash Kak. 2012. A survey of prediction using social media. Retrieved from <https://arXiv:1203.1647>.
- [197] Zhuoning Yuan, Xun Zhou, and Tianbao Yang. 2018. Hetero-convLSTM: A deep learning approach to traffic accident prediction on heterogeneous spatio-temporal data. In *Proceedings of SIGKDD*. 984–992.
- [198] Jenny Yuen and Antonio Torralba. 2010. A data-driven approach for event prediction. In *Proceedings of ECCV*. 707–720.
- [199] Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. 2014. *Social Media Mining: An Introduction*. Cambridge University Press.
- [200] 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 *Proceedings of WWW*. 311–319.
- [201] Liang Zhao, Feng Chen, Jing Dai, Ting Hua, Chang-Tien Lu, and Naren Ramakrishnan. 2014. Unsupervised spatial event detection in targeted domains with applications to civil unrest modeling. *PLoS One* 9, 10 (2014), e110206.
- [202] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Spatiotemporal event forecasting in social media. In *Proceedings of SDM*. SIAM, 963–971.
- [203] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Multi-resolution spatial event forecasting in social media. In *Proceedings of ICDM*. IEEE, 689–698.
- [204] 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 *Proceedings of ICDM*. IEEE, 639–648.
- [205] Liang Zhao, Olga Gkountouna, and Dieter Pfoser. 2019. Spatial auto-regressive dependency interpretable learning based on spatial topological constraints. *ACM Trans. Spatial Algor. Syst.* 5, 3 (2019), 1–28.
- [206] Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Multi-task learning for spatio-temporal event forecasting. In *Proceedings of SIGKDD*. 1503–1512.
- [207] Liang Zhao, Junxiang Wang, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2017. Spatial event forecasting in social media with geographically hierarchical regularization. *Proc. IEEE* 105, 10 (2017), 1953–1970.
- [208] Liang Zhao, Junxiang Wang, and Xiaojie Guo. 2018. Distant-supervision of heterogeneous multitask learning for social event forecasting with multilingual indicators. In *Proceedings of AAAI*.
- [209] Liang Zhao, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Hierarchical incomplete multi-source feature learning for spatiotemporal event forecasting. In *Proceedings of SIGKDD*. 2085–2094.
- [210] Sendong Zhao, Quan Wang, Sean Massung, Bing Qin, Ting Liu, Bin Wang, and ChengXiang Zhai. 2017. Constructing and embedding abstract event causality networks from text snippets. In *Proceedings of WSDM*. 335–344.
- [211] Xiangyu Zhao and Jiliang Tang. 2017. Modeling temporal-spatial correlations for crime prediction. In *Proceedings of CIKM*. 497–506.
- [212] Qiuzhen Zhong, Jingjing Wang, Xuejie Meng, Siqing Liu, and Jiancun Gong. 2019. Prediction model for solar energetic proton events: Analysis and verification. *Space Weather* 17, 5 (2019), 709–726.
- [213] Cheng Zhou, Boris Cule, and Bart Goethals. 2015. A pattern based predictor for event streams. *Expert Syst. Appl.* 42, 23 (2015), 9294–9306.
- [214] Lihua Zhou, Guowang Du, Ruxin Wang, Dapeng Tao, Lizhen Wang, Jun Cheng, and Jing Wang. 2019. A tensor framework for geosensor data forecasting of significant societal events. *Pattern Recogn.* 88 (2019), 27–37.
- [215] Xun Zhou, Amin Vahedian Khezerlou, Alex Liu, Zubair Shafiq, and Fan Zhang. 2016. A traffic flow approach to early detection of gathering events. In *Proceedings of SIGSPATIAL*. 1–10.
- [216] Xun Zhou, Shashi Shekhar, and Reem Y. Ali. 2014. Spatiotemporal change footprint pattern discovery: An interdisciplinary survey. *Data Min. Knowl. Discov.* 4, 1 (2014), 1–23.

Received July 2020; revised December 2020; accepted February 2021