A Survey on Societal Event Forecasting with Deep Learning

SONGGAOJUN DENG and YUE NING, Stevens Institute of Technology, USA

Population-level societal events, such as civil unrest and crime, often have a significant impact on our daily life. Forecasting such events is of great importance for decision-making and resource allocation. Event prediction has traditionally been challenging due to the lack of knowledge regarding the true causes and underlying mechanisms of event occurrence. In recent years, research on event forecasting has made significant progress due to two main reasons: (1) the development of machine learning and deep learning algorithms and (2) the accessibility of public data such as social media, news sources, blogs, economic indicators, and other meta-data sources. The explosive growth of data and the remarkable advancement in software/hardware technologies have led to applications of deep learning techniques in societal event studies. This paper is dedicated to providing a systematic and comprehensive overview of deep learning technologies for societal event predictions. We focus on two domains of societal events: *civil unrest* and *crime*. We first introduce how event forecasting problems are formulated as a machine learning prediction task. Then, we summarize data resources, traditional methods, and recent development of deep learning models for these problems. Finally, we discuss the challenges in societal event forecasting and put forward some promising directions for future research.

Additional Key Words and Phrases: Event Prediction, Deep Learning

1 INTRODUCTION

Societal events can be broadly categorized into offline events (e.g., strikes, protests, and robberies) and online events (e.g., activism, petitions, and hoaxes on online platforms [157]). This paper focuses on studying offline societal events that occur in specific places and times and affect society in different ways. Understanding such events and their recurring patterns is an urgent issue for many stakeholders such as investors and suppliers. Different from retrospective studies such as event detection [145] and summarization [25], event forecasting focuses on anticipating events in the future based on historical event patterns. Accurate prediction of future events is conducive to effective allocation of public resources, reducing economic loss and social damage. Generally, it can bring enormous benefits to society and individuals, such as effective disaster response and socio-economic growth. In this survey, we focus on summarizing the literature on societal event prediction, where events are initiated by humans and affect society on a large scale. Specifically, we organize related work on the prediction of civil unrest and criminal activities that have widespread social impacts from individuals to cities and nations [17].

With the availability of mass media sources, data-driven approaches have been widely studied in various fields, including civil unrest [34] and criminal activities [55]. In these studies, researchers attempt to develop predictive analytic techniques based on historical data of events and other relevant sources to predict future events. The methods studied are from a variety of fields, such as statistics, data mining, machine learning, and deep learning. Nowadays, the effective collection of good-quality data, the popularity of high-performance computing equipment [39], and the accelerated development of the field of machine learning and deep learning have led to substantial progress in event prediction research. However, there are still many challenges present in addressing this problem:

- **Leveraging Heterogeneous Data.** As the availability of various open-source data continues to increase, researchers have begun to resort to heterogeneous data to develop predictive models [86, 119, 136]. Societal events occur in a dynamic social environment, and their corresponding key information can publish in various forms, such as traditional news reports, social media, and government official reports. It brings unique challenges of efficiently processing and learning from heterogeneous data to make accurate event predictions.
- **Studying Dependencies in Societal Events.** Societal events exhibit geographical properties and also have a high degree of temporal dependency [158]. Modeling event contextual information requires an in-depth investigation of the spatiotemporal dependencies of events. Traditional methodologies have shown limitations in modeling complex event data [94, 117], which encourages the development of advanced models to address this challenge.
- **Interpreting Event Predictions.** Event prediction models enable human users to automatically forecast future events of interest without manually inspecting historical data. Along with event forecasting, interpreting the prediction results is an equally important task, as it can assist practitioners in understanding prediction results and making reasonable and practical decisions. Data-driven event prediction models with adequate interpretability are more conducive to supporting human users in event analysis and decision making than black box models.

Over the past few decades, a large amount of research has been devoted to the development and application of event prediction techniques in response to the above challenges. Deep learning has gained tremendous success in many applications such as computer vision, speech recognition, and natural language processing, given its superiority in improving accuracy when trained with sufficient data. Recently, there has been a surge of deep learning approaches in the societal event domain. The purpose of this survey is to summarize deep learning technologies for event forecasting and to highlight the open problems and new possibilities for societal event prediction.

1.1 Previous Work and Contributions

There are several other surveys on societal event predictions. Phillips et al. [97] provided a literature review that examines the problems and techniques for predictive analysis using social media data. Recently, Zhao [157] provided a systematic survey of existing data-driven event prediction methods, covering challenges, techniques, applications, evaluations, and open problems. As event prediction methods are typically motivated by specific application areas, some surveys instead focus on event technologies in particular domains. Ganar and Ardhapurkar [41] summarized studies of civil unrest prediction in social media analysis based on keyword filtering. For predictive studies of crime, Mookiah et al. [89] discussed crime research that considers various crime-related variables and examined the impact of these variables on crime prediction. Kounadi et al. [73] provided an overview and evaluation of the state-of-the-art technologies in spatial crime prediction, with a focus on research design and techniques. Butt et al. [19] investigated detection and prediction methods for crime hotspots, mainly covering papers published in the recent literature.

In contrast to previous work, this survey is structured around problems, data, and deep learning models for civil unrest and crime forecasting. Civil unrest and crime are influential societal events and share some common characteristics. Both civil unrest and crime are human-initiated events, largely depending on evolving social environments and affecting humans and their life to a large extent. We seek to provide a systematic survey that bridges the prediction of such influential societal events with the development of deep learning and point out possible directions for research and development. The main contributions of this work can be summarized as follows:

- **A Comprehensive Summary of Data Resources** Developing data-driven approaches for predicting societal events relies heavily on big data. In this work, we organize and categorize the event datasets that have been investigated in existing work. As data are the basis for data-driven modeling of societal events, this summary will help researchers and human experts in several ways, such as selecting critical features for societal event analysis.
- **A Systematic Review of Existing Deep Learning Methods.** This survey categorizes existing deep learning methods for predicting two types of societal events (i.e., civil unrest and crime). We analyze the characteristics of deep learning-based technologies for each event domain. Traditional statistical and machine learning methods are also summarized, with their limitations discussed. The methodological review aims to help researchers systematically understand the advantages and achievements of various deep learning techniques in societal event prediction.
- **A Discussion of Ongoing Challenges and Future Prospects.** This paper provides an overall outline of the current research status by discussing ongoing research challenges and future directions. With this discussion, we aim to provide comprehensive insights to practitioners in related fields and advance the studies of societal event prediction.

1.2 Outline

The rest of this paper is organized as follows. We first present the usage of mathematical notations in this survey in Section 1.3. Section 2 presents the problem formulation of civil unrest prediction, along with the introduction of data resources, traditional methods, and a comprehensive review of deep learning models developed for civil unrest prediction. Section 3 discusses crime prediction, followed by the same outline as in Section 2. In Section 4, we summarize the deep learning technologies studied in civil unrest and crime prediction. Section 5 lists open problems and suggestions for future research directions. This survey concludes with a summary in Section 6.

1.3 Notation

To ensure the readability of the survey, we use consistent symbols when describing variables. Calligraphy capital letters represent general variables. For example, $\mathcal{X}$ denotes input data which could be in any format, and $\mathcal{Y}$ is the generic target variable, which can be a binary or continuous value. Bold lowercase letters indicate vectors (e.g., $\mathbf{x}$), bold uppercase letters indicate matrices (e.g., $\mathbf{X}$), and heavier uppercase letters signify tensors/higher dimensional collections of data (e.g., $\mathbf{X}$). We use the letters W, b to represent learnable model parameters. For simplicity and consistency, we will use the letter f (e.g., $f(\cdot)$) in some cases to represent the components in a deep learning model such as a feedforward neural network. σ represents the sigmoid function, $\odot$ represents the Hadamard product, and $[;]$ signifies concatenation.

2 CIVIL UNREST PREDICTION

Civil unrest, also known as civil disorder, is an activity arising from a mass act of gathering. In such activities, participants express their protest against things that affect their lives and for which they assume that governments or other entities have a responsibility (e.g., cost of urban transportation, poor infrastructure, etc.). Civil unrest events include protests, strikes, and demonstrations for different reasons. Anticipating such events can help stakeholders (e.g., investors and suppliers) understand the dynamic patterns of these events and improve resource allocation. This section starts with the formal problem formulation of civil unrest prediction followed by an introduction of available data resources in this research area. In the following, we summarize traditional civil unrest forecasting methods and provide a more detailed description of advanced deep learning approaches categorized by the characteristics of model structures.

2.1 Problem Formulation

Event forecasting can be expressed as a supervised learning task in machine learning, which aims at learning a function that maps an input to an output based on input-output pairs. It infers a function from labeled training data consisting of a set of training instances, where the input is historical data collected before the timestamp of the output variable. Before going into the problem definition of societal event forecasting, we first introduce the preliminaries, including terminologies and mathematical notations.

Historical Data. Suppose there are L locations (e.g., cities, states) of interest, and each location l can be represented by a set of features used for prediction. We divide features into two categories, static and dynamic.[1] Static features such as population and political ideology remain constant or change slowly over a long time, while dynamic features such as frequency of events or number of tweets expressing "angry" emotions are updated for each time interval t (e.g., day, week). Let $\mathcal{S}_l$ denote the set of static features of location l, and $\mathcal{X}_{t,l}$ be the collection of dynamic features for location l at time t. The collection of dynamic features from location l within an historical window (i.e., observing time window) with size k up to time t can be represented as $\mathcal{X}_{t-k+1:t,l} = (\mathcal{X}_{t-k+1,l}, ..., \mathcal{X}_{t,l})$.

Ground-truth Event Occurrence. A target variable $\mathcal{Y}_{t^*,l}$ indicates the occurrence of a future civil unrest (e.g., protest) for each location l at time t^*. Note that t^* can be either a time point $(t + \Delta)$ or a time window in the future $(t + \Delta : t + \Delta + \delta)$. $\Delta \geq 1$ is the **lead time** that denotes the number of time steps in advance for a prediction. We use $\delta \geq 0$ to denote the **lead time window** that represents whether a civil unrest event will occur between time $t + \Delta$ and $t + \Delta + \delta$.

Definition 2.1. **Civil Unrest Prediction as Classification**. Learn a classifier $f(\mathcal{S}_l, \mathcal{X}_{t-k+1:t,l}) \rightarrow \mathcal{Y}_{t^*,l}$ that maps an input including static and dynamic features, to a categorical civil unrest variable $\mathcal{Y}_{t^*,l} \in \{0, 1, ..., M-1\}$ at the future time t^* for the target location l. If there is only one category of events to be predicted, $\mathcal{Y}_{t^*,l} = 1(0)$ indicates that events of interest will (not) occur at location l at time t^*. If there are multiple categories of events to be predicted at the same time $(M > 2)$, this problem becomes a multi-label classification problem where the target value indicates the occurrence of multiple classes (e.g., multiple subtypes of events), i.e., $\mathcal{Y}_{t^*,l} = \{0, 1\}^M$.

Definition 2.2. **Civil Unrest Prediction as Regression**. Given static features $\mathcal{S}_l$ and dynamic features $\mathcal{X}_{t-k+1:t,l}$, the goal is to learn a regressor $f(\mathcal{S}_l, \mathcal{X}_{t-k+1:t,l}) \rightarrow \mathcal{Y}_{t^*,l}$ that maps the input to a numerical civil unrest indicator $\mathcal{Y}_{t^*,l} \in \mathbb{R}$ at the future time t^* for the target location l. In general, $\mathcal{Y}_{t^*,l}$ refers to normalized number of event frequencies.

2.2 Data Resources

In this section, we summarize the event data sources utilized in civil unrest prediction studies. Event data are typically used as the ground-truth of event occurrence and historical input data. We also discuss some open-source indicators that have been used as historical input data for civil unrest prediction.

2.2.1 Civil Unrest Event Related Data. Since the last century, various event data projects with different data collection, coding, and analysis processes have emerged. We organize event data related to civil unrest in Table 1. According to the coding process of event data, we divide these data into human-coded and machine-coded data. Human-coded data depend on human research teams with specific knowledge of the local context, while machine-coded data rely entirely on automated event encoding systems.

Human-Encoded Events. Many human-coded event datasets have been developed and maintained, which allow researchers to build forecasters at specific sub-state geographic units [86, 113, 140]. In the earliest days, due to technological limitations (i.e., the lack of electronic articles and computational power), the World Event Interaction Survey (WEIS) [84] and the Conflict and Peace Data Bank (COPDAB) [9] projects hire human analysts

---

[1]Dynamic features are usually involved in civil unrest prediction studies, while static features are sometimes not considered.

Table 1. A summary of event datasets used for civil unrest prediction. All listed datasets are labeled with geolocation information. The start/end time in Temporal Coverage indicates the earliest/latest time of data collection. "/" means unavailable. We use a horizontal line to separate the human-encoded data (top) from the machine-encoded data (bottom).

Dataset	Mendate	Temporal coverage	Spatial coverage	Event definition	Coding process	Open sourced
COPDAB [9]	Political events	1948-1978	Near global	/	Human	Yes
WEIS [84]	Political events	1966-1978	Global	/	Human	Yes
SPEED [91]	Social, political and economic events	1945-2008	Global	/	Human	Yes
SCAD [109]	Social conflict events	1990-2017	Africa and Latin America	/	Human	Yes
ACLED [101]	Political violence and demonstrations	1997-	Near global	A single altercation where force is often used by one or more groups to a political end	Human	Yes
Daly [32]	Violent events/Rebellion	1964-1984	Colombia	The entire repertoire of violence and threats of violence that suggest the presence of a militarized organization	Human	No
Urdal and Hoelscher [123]	Civil unrest events	1960-2009	Asia and Sub-Saharan Africa	/	Human	No
KOSVED [112]	One-sided violence	1991-2008	Africa and Europe	Violent acts perpetrated by an organized group, directed against a group of unarmed non-combatants during, shortly before, or after a conflict	Human	Yes
GTD [74]	Terrorism	1970-2019	Global	The threatened or actual use of illegal force and violence by a non-state actor to attain a political, economic, religious, or social goal through fear, coercion, or intimidation	Human	Yes
UCDP GED [120]	Organized violence	1989-2020	Near global	An incident where armed force was used by an organised actor against another organized actor, or against civilians, resulting in at least 1 direct death	Human	Yes
GSR [102]	Civil unrest events	/	Latin America	Organized by an independent evaluation team (MITRE) by surveying newspapers for reportings of civil unrest	Human	No
KEDS [114]	Political events	1979-1997	Middle East, Balkans, and West Africa	Use dictionaries of nouns and verb phrases to code the actors and events	Machine	Yes
Phoenix [2]	Political events	1945-2019	Global	CAMEO methodology	Machine	Yes
GDELT [78]	Various physical activities, including political and non-political events	1979-	Global	CAMEO methodology	Machine	Yes
ICEWS [16]	Political events	1995-	Global	CAMEO methodology	Machine	Yes
AutoGSR [111]	Civil unrest events	/	Latin Americ	Use the developed automatic event coding system to classify articles as protests or non-protests	Machine	No

to physically collect newspaper clippings, press reports, and summary accounts from Western news sources to obtain news stories. These projects focus on daily international and domestic events or interactions. Analysts coded events in a multi-column format based on their subjective judgment. The event record mainly includes the date of the event, actors involved in the event, sources from which the information about the event was gathered, issue area(s), and textual description.

Later, there were more event data projects focused on specific areas. The Social, Political, Economic Event Database (SPEED) Project [91] is a technology-intensive effort to extract event data from a global archive of news reports covering the Post WWII era. This project develops more than one hundred attributes to encode each event, excluding text descriptions. Salehyan et al. developed an event dataset, Social Conflict in Africa Database (SCAD) [109], which contains instances of protests, riots, strikes, government repression, communal violence, and other forms of unrest that happened mainly in Africa. The Gold Standard Report (GSR) is a collection of human classified civil unrest news reports from the most influential newspaper outlets in Latin America [102]. The Armed Conflict Location and Event Data Project (ACLED) [101] is a disaggregated data collection, analysis, and crisis mapping project. ACLED collects the dates, actors, locations, fatalities, and types of all reported political events (e.g., violence, protests) around the world. Urdal and Hoelscher [123] introduced an event dataset on urban unrests at the city level that covers 55 major cities in Asia and Sub-Saharan African countries.

In addition to political events, some datasets focus on violent events. Daly [32] collected violent events at the municipality-month level in Colombia. The Global Terrorism Database (GTD) [74] provides information on domestic and international terrorist attacks around the world. The Konstanz One-Sided Event Dataset (KOSVED) [112] provides detailed information on the magnitude and locations of one-sided violent events in 20 civil wars. The Uppsala Conflict Data Program Georeferenced Event Dataset (UCDP GED) [120] is an event dataset that disaggregates three types of organized violence (state-based conflict, non-state conflict, and one-sided violence) both spatially and temporally.

Machine-Encoded Events. Manual approaches began to be replaced with automated coding with the first iteration of the Kansas Event Data Set (KEDS) [114] project in the late 1980s. KEDS uses the automated coding of English-language news reports to generate political event data. These data are used in statistical early warning models to predict political changes. The Integrated Crisis Early Warning System (ICEWS) [16] includes a database of political events with global coverage. Similar to ICEWS, the Global Dataset of Events, Location, and Tone (GDELT) [78] has been developed and compiled a comprehensive list of electronic news sources. Both GDELT and ICEWS are active automatic systems that identify and classify events from public data following the Conflict and Mediation Event Observations (CAMEO) [43] which is a framework for coding event data. These two datasets have been extensively studied in various fields. Some researchers compared ICEWS with GDELT [7, 133, 137] and pointed out limitations on local conflict processes that rely too heavily on machine-coded data [47, 48]. The Historical Phoenix Event Data [2] includes events extracted from The New York Times, BBC Monitoring's Summary of World Broadcasts, and the CIA's Foreign Broadcast Information Service. It also uses the CAMEO methodology to encode events. Given the large scale and extensive spatial coverage of machine-coded data, many researchers have used machine-coded event data to build forecasts for civil unrest [34, 42, 93, 98, 99, 151, 159–162].

2.2.2 Open Source Indicators. In civil unrest prediction, researchers typically define ground truth data for learning tasks. In general, historical civil unrest or related events can be used as basic features for modeling future event patterns of interest. Recent research has incorporated public media data as features for civil unrest forecasting. This type of data is known as Open Source Indicators (OSI). OSI includes traditional media data such as digital newspapers, blogs, and social media data such as posts from Twitter and Facebook. These data provide a wealth of background information that helps one understand the social context and public opinion of civil unrest events. Economic indicators and other meta-data sources have also been explored in this line of research. Studies show that exogenous political and economic variables can serve as the necessary underlying drivers of social unrest

besides social media [86, 141]. Social and economic features derived from the World Development Indicators (WDI) [13] and Worldwide Governance Indicators (WGI) [66] have been investigated [95]. Researchers have also utilized Google Trends (GT) to uncover social dynamics associated with behavior that precedes episodes of civil unrest [83]. Google Trends analyzes the popularity of top search queries in Google Search across various regions and languages.

2.3   Traditional Methods

Over the years, researchers have leveraged various predictive techniques for civil unrest predictions ranging from simplest threshold-based methods [83] to more sophisticated methods such as logistic regression and deep neural network. We outline the development of traditional predictive approaches in this field before introducing deep learning-based methods.

As a pioneering study, Radinsky and Horvitz [100] mined chains of events from massive news archives and proposed a probabilistic method that predicts the likelihoods of future worldwide events of interest. Jin et al. [60] characterized mass protest propagation using a bispace diffusion model comprising a Twitter mentions network and a latent space. Chen and Neill [27] explored a nonparametric graph scan algorithm to the problem of civil unrest event detection and forecasting using heterogeneous social media graphs. Subsequently, people began to apply traditional machine learning models to civil unrest prediction. Kallus [64] studied the power of massive online-accessible public data to predict crowd behavior such as protests utilizing random forests. Logistic regression as a classification technique has been extensively explored in combining heterogeneous data sources for civil unrest forecasting [20, 72, 102, 142, 160]. Early Model Based Event Recognition using Surrogates (EMBERS) [90, 102] is an automated system developed for generating forecasts about civil unrest from massive and multiple data sources. This system consists of individual models that leverage specific data sources, including logistic regression, lexical pattern matching, dynamic query expansion, etc. More advanced methodologies based on logistic regression have emerged. Multi-task learning frameworks were incorporated in forecasting spatial-temporal civil unrest events [42, 93, 159–162]. Ning et al. [92] introduced a multi-instance learning method for forecasting and modeling precursors for civil unrest. More recently, Zhao et al. [163] presented a group-Lasso based hierarchical feature learning model to characterize feature dependence, feature sparsity, and interactions among missing values. Temporal models such as Hidden Markov Models (HMM) were proposed to leverage temporal burst patterns from large-scale digital history-coded events [98, 99] or Twitter streams [158] to reveal the underlying developmental mechanism of civil unrest events. Other work considered civil unrest forecasting as a regression task and investigated time series models such as autoregressive model [113, 140, 151].

Although substantial progress has been made, traditional approaches usually require extensive efforts in feature engineering, such as analyzing news reports to obtain keywords as feature variables. Moreover, these models often consider only simple features and cannot model complex dependencies in the data, which limits the predictive power of such models. Next, we discuss deep learning-based models that are capable of capturing complex data dependencies.

2.4   Deep Learning Models

A number of deep learning-based approaches have been proposed to predict civil unrest events and have demonstrated their advantages in terms of improved predictive capability or interpretability. We classify existing work based on their methodology into Recurrent Neural Network (RNN), Attention, and Graph Neural Network (GNN). We provide an overview of deep learning-based models that focus on the prediction of civil unrest in Table 2.

2.4.1   Recurrent Neural Network Based Approaches. Predicting civil unrest events can be considered as a time series prediction problem. Recurrent neural networks (such as RNN, LSTM [50] and GRU [31]) are proposed to

Table 2. Deep learning based civil unrest prediction techniques. We use abbreviations "class." and "reg." to indicate classification and regression, respectively. We use horizontal lines to partition the models with different structural characteristics, corresponding to section 2.4.

| Technique | Space | | Time | Dataset | | Forecasting | | | Evaluation |
	Study area	Scale		Event data	External data	Inference	Task	Temporal unit	
LSTM [117]	Afghanistan	Country	2001-2012	GDELT	/	# of material conflict events	Reg.	Month	MAE, RMSE, MAPE
GRU [95]	158 countries	Country	Mar 2001-Mar 2014	GDELT	WDI, WGI	Occurrence of disruptive events	Binary class.	Month	AUC, ROC
LSTM [47]	Libya, Sudan, Egypt, Maldives, Nicaragua	Country	1989-2019	ICEWS, GDELT	/	# of material conflict events	Reg.	Month	RMSE
Cov-LSTM [86]	Egypt, Jordan	Country, city	May 2015-Jan 2019	GSR, ACLED	WDI, Trading Economics, Wikipedia, Twitter, Google Trends	# of civil unrest events	Class.	Day	Mercury score[a], ROC
CA-LSTM [135]	USA	State	Jan 2012-Jan 2018	GDELT	/	# of material conflict events	Reg.	Day	MSE, MAE
ActAttn [38]	USA	State	2014, 2017	Charlottesville rally, Ferguson protests I, Ferguson protests II	Twitter	Occurrence of protests	Binary class.	Day	AUC, F1
DynamicGCN [34]	Thailand, India, Egypt, Russia	City	2010-2016	ICEWS	News reports	Occurrence of protests	Binary class.	Day	Precision, Recall, F1
Glean [35]	Nigeria, India, Afghanistan, Russia	City	2012-2016	ICEWS	News reports, Knowledge graphs	Occurrence of multiple political events	Class.	Day	F1, F2[b], Recall
CMF [36]	Thailand, India, Egypt, Russia	City	2012-2016; 2017-2019	ICEWS, GDELT	News reports, Knowledge graphs	Occurrence of protests	Binary class.	Day	F1

[a] The evaluation metric is defined in the IARPA Mercury Challenge.
[b] F2 score is similar to F1 score, but assigns less weight to precision and more weight to recall.

capture short- and long-term dependencies in time series data, and they have been proven to be more expressive and powerful than traditional methods such as autoregressive models. Researchers have employed recurrent neural networks to model temporal dependencies on time series event data for civil unrest prediction.

Suppose that each sample in the time series data for a location consists of a sequence $\mathbf{X} = (\mathbf{x}_{t-k+1}, ..., \mathbf{x}_t)$ with length k where the element $\mathbf{x}_t$ is the feature vector at time t, e.g., event frequency. The sequence is fed into the RNN architecture for predicting the occurrence of the target event y_{t^*} at the future time t^*. A vanilla RNN model is defined as belows:

$$\mathbf{h}_t = \tanh(\mathbf{W}[\mathbf{x}_t; \mathbf{h}_{t-1}] + \mathbf{b}), \tag{1}$$

where $\{\mathbf{h}, \mathbf{x}\}_t$ are the hidden states and input feature vector at time t and $\mathbf{h}_{t-1}$ is the hidden state at time $t-1$. $[;]$ is concatenation. $\{\mathbf{W}, \mathbf{b}\}$ are model parameters. The time-series event prediction using vanilla RNN can be written in the following form:

$$\hat{y}_{t^*} = f(\mathbf{h}_\tau), \quad \tau \leq t. \tag{2}$$

The prediction $\hat{y}_{t^*}$ can be obtained from a feedforward network that takes the hidden states as input, either the last or all hidden states. The vanilla RNN model typically fails to store information for a longer period of time and has gradient vanishing and exploding problems. Existing work usually utilizes variants of RNN (e.g., LSTM and GRU) to model time series data, especially when the input sequence is not short. Smith et al. [117] and Halkia et al. [47] applied LSTM models to event prediction, focusing on material conflict events such as armed attacks and destruction of property. Unlike RNN models, LSTM includes cell states that can remove or add information to the cell. Thanks to this structure, LSTM is more accurate on datasets with very long sequences. GRU is another recurrent neural network that uses fewer training parameters and therefore uses less memory leading to faster training than LSTM. Parrish et al. [95] studied a GRU-based multi-feature driven approach to predict disruptive events. Instead of considering the sequence of event counts as input, this work consists of a feature selection process to select social and economic features derived from open-source Indicators.

More recently, Meng and Srihari [86] proposed a model, Cov-LSTM [86], that combines convolutional layers and LSTM layers. The convolutional layers aim to extract high-level representations of event time series data. The transformed input features $\mathbf{x}'_t = \text{Conv1D}(\mathbf{x}_t)$ are fed to the LSTM model to predict civil unrest events.

2.4.2 Attention Based Approaches. Several challenges have emerged in RNN-based approaches in civil unrest prediction, including limited ability to accurately predict unrest events and difficulty for users to understand the model behavior. Attention mechanisms were first introduced into a neural machine translation model based on encoder-decoder RNNs [11]. Attention mechanisms enable dynamically highlighting relevant features of the input data, mimicking cognitive attention in humans. This method enhances the important parts of input data and fades out the rest. Formally, the attention mechanism takes several hidden vectors $(\mathbf{h}_{t-k+1}, ..., \mathbf{h}_t)$, and a context vector $\mathbf{c}$ as input, and outputs an attention weight $\alpha_\tau, t - k + 1 \leq \tau \leq t$ for each hidden vector $\mathbf{h}_\tau$ and an attention vector $\mathbf{a}_t$:

$$\alpha_\tau = \text{softmax}\big(\text{score}(\mathbf{h}_\tau, \mathbf{c})\big), \quad \mathbf{a}_t = \sum_{t-k+1 \leq \tau \leq t} \alpha_\tau \cdot \mathbf{h}_\tau, \tag{3}$$

where $\text{score}(\cdot)$ is a scoring function that can be implemented in various forms (e.g., a feed-forward neural network, dot product, etc.) leading to different attention mechanisms, such as additive [11], dot product [81] or self-attention mechanisms [125]. The softmax function normalizes the output of the scoring function to the attention weights. Removing the context vector $\mathbf{c}$ results in a location-based attention, where the attention weight depends only on the target location.

Wang et al. proposed a context-aware attention-based LSTM framework named CA-LSTM [135] to study different contributions of data points in the time series and to improve the accuracy in predicting civil unrest events. They applied an attention layer on top of the LSTM layer to obtain an attention vector. To further include

trends in the occurrence of events, they explicitly incorporated target event variables in previous historical steps to the attention vector for civil unrest prediction, i.e., $\hat{y}_{t^*} = f(\mathbf{a}_t, (y_{t-k'}, ..., y_t))$.

Ertugrul et al. introduced a hierarchical attention-based spatiotemporal learning approach, ActAttn [38] for predicting the occurrence of future protests and explaining feature importance. Specifically, the model takes historical data (i.e., Twitter data) from the target (predicted) location $l \in L$, as well as historical data from all locations L, as input. The model consists of a two-level attention module built on LSTM models, which calculates intra-regional (temporal) and inter-regional (spatial) contributions, respectively. The proposed model can be briefly expressed as bellows:

$$\hat{y}_{t^*,l} = f(\mathcal{S}_l, \alpha^{\text{temporal}} \mathbf{h}_{t,l}^{\text{temporal}} + \alpha^{\text{spatial}} \mathbf{a}_t^{\text{spatial}}), \tag{4}$$

where $\mathbf{a}_t^{\text{spatial}} = \sum_{l \in L} \alpha_l \cdot \mathbf{h}_{t,l}$ is the spatial attention vector, obtained from applying attention on hidden states for all locations. A spatiotemporal attention is applied on the hidden states of the target location $\mathbf{h}_{t,l}^{\text{temporal}}$ and the spatial attention vector $\mathbf{a}_t^{\text{spatial}}$. All hidden states are obtained from LSTM models. The network can help interpret what features, from which places, have significant contributions to a prediction of protest.

2.4.3 Graph Neural Network Based Approaches. Recently, graph neural networks have achieved significant development in the machine learning community and have been applied to a wide range of fields such as computer vision [115], natural language processing [146], bioinformatics [104], etc. These methods have demonstrated superior performance in modeling latent embedding in graph-structured data.

Graph neural networks (GNNs) are a class of neural networks for prediction tasks on graph-represented data. GNNs learn embeddings/hidden features for each node in a graph by using topological information. In general, one GNN layer utilizes the node embeddings from last layer and the graph structure of a graph:

$$\mathbf{H}^{(l)} = f_{\text{GNN}}(\mathbf{A}, \mathbf{H}^{(l-1)}), \tag{5}$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix of the graph. $\mathbf{H}^{(l-1)} \in \mathbb{R}^{N \times d^{l-1}}$ denotes the node embeddings at the $(l-1)$-th GNN layer and $\mathbf{H}^{(l)}$ is the updated node embeddings. d^{l-1} is the dimension of the node embedding vectors at layer $l-1$. At the first layer, the node embeddings are usually predefined node attributes $\mathbf{H}^{(0)} = \mathbf{X}$.

Graph convolutional networks (GCN) [71] are a representative set of GNN models that have been widely studied on undirected (weighted or unweighted) graphs: The generic framework of a GCN layer can be represented as follows:

$$\mathbf{H}^{(l)} = \phi(\hat{\mathbf{A}} \mathbf{H}^{(l-1)} \mathbf{W}^{(l-1)}), \tag{6}$$

where $\hat{\mathbf{A}}$ is the normalized adjacency matrix by applying the renormalization trick to $\mathbf{A}$ [71]. $\mathbf{W}^{(l-1)}$ is a layer-specific trainable weight matrix, and ϕ is the activation function, e.g., ReLU.

Based on the GCN model, Deng et al. proposed a dynamic graph convolutional network (DynamicGCN) [34] for forecasting protest events and identifying context graphs for the predicted events. This paper introduced an encoding method to encode historical news articles into a sequence of word graphs $(\mathbf{A}_{t-k+1}, ..., \mathbf{A}_t)$. The adjacency matrix $\mathbf{A}_t \in \mathbb{R}^{N \times N}$ represents a word graph with N nodes at time t where each node is a word, and the weighted edge between two words is calculated by point-wise mutual information (PMI). The pre-trained word embeddings are used as node features. The proposed dynamic graph convolutional layer is defined as:

$$\mathbf{H}_t = \phi(\hat{\mathbf{A}}_{t-1} \tilde{\mathbf{H}}_{t-1} \mathbf{W}_{t-1}), \quad \tilde{\mathbf{H}}_t = \begin{cases} \mathbf{H}_0 & \text{if } t = 0 \\ f([\mathbf{H}_t; \mathbf{H}_0]) & \text{otherwise,} \end{cases} \tag{7}$$

where $f(\cdot)$ is a nonlinear layer that re-encodes the input features of graph convolutional layers, including the semantic information of the words ($\mathbf{H}_0$) and the learned graph embeddings ($\mathbf{H}_t$). The prediction is obtained by first aggregating node embeddings at the last layer to generate graph-level embeddings (e.g., linear transformation or pooling). The graph-level embeddings are fed to an output layer for prediction, i.e., $\hat{y}_{t^*} = f_{\text{output}}(f_{\text{agg}}(\mathbf{H}_t))$. The

model was demonstrated to be effective in the prediction of protest events in multiple countries. The authors also proposed a heuristic subgraph extraction to help explain event prediction results.

In addition to unstructured text data, knowledge graphs, as another data source, have been explored for temporal event prediction. Knowledge graphs are multi-relational graphs consisting of tuples of elements (i.e., subject entity, relation, object entity), denoted as $\mathcal{G} = \{(s, r, o)\}$. Nodes are entities, and edges represent relationship types. For example, {(*Citizen*, Criticizes, *Government*), (*Citizen*, Appeals, *Police*)} can be represented as a simple knowledge graph, where *Citizen*, *Government* and *Police* are entities and Criticizes and Appeals are relationship/event types. In general, graph convolution networks for knowledge graphs can be described in terms of an update mechanism of node and edge embeddings. For instance, Composition-based Multi-Relational Graph Convolutional Networks (CompGCN) [124] update node and edge embeddings as follows:

$$\mathbf{h}_o^{(l+1)} = \phi\Big(\sum_{(s,r)\,\exists\,(s,r,o)\in\mathcal{G}} f_{\text{node}}\big(\mathbf{h}_s^{(l)}, \mathbf{o}_r^{(l)}\big)\Big), \quad \mathbf{o}_r^{(l+1)} = f_{\text{edge}}(\mathbf{o}_r^{(l)}), \tag{8}$$

where $\mathbf{h}$, $\mathbf{o}$ denote node and edge embeddings, respectively. This model jointly embeds both nodes (subjects and objects) and edges (relations) in a relational graph, and allows for more complex information modeling than GCN models.

Utilizing event-based knowledge graphs, Glean [35], was proposed to predict concurrent events of multiple types, including civil unrest events. This paper investigates temporal event knowledge graphs which are built upon a sequence of event sets in ascending time order $(\mathcal{G}_{t-k+1}, ..., \mathcal{G}_t)$, as well as a sequence word graphs $(\mathbf{A}_{t-k+1}, ..., \mathbf{A}_t)$. The model includes a graph aggregation module that learns node embeddings from both knowledge graphs and semantic word graphs at each historical timestamp. A context-aware embedding fusion module was introduced to enhance representations of nodes and edges in event knowledge graphs by blending embeddings of contextual word nodes. The procedures can be expressed as bellows:

$$\mathbf{H}_t^{\text{event}}, \mathbf{O}_t^{\text{event}} = f_{\text{CompGCN}}(\mathcal{G}_t, \mathbf{H}_0^{\text{event}}, \mathbf{O}_0^{\text{event}}), \quad \mathbf{H}_t^{\text{word}} = f_{\text{GCN}}(\mathbf{A}_t, \mathbf{H}_0^{\text{word}}), \tag{9}$$

$$\tilde{\mathbf{H}}_t^{\text{event}} = f_{\text{Attn}}(\mathbf{H}_t^{\text{event}}, \mathbf{H}_t^{\text{word}}), \quad \tilde{\mathbf{O}}_t^{\text{event}} = f_{\text{Attn}}(\mathbf{O}_t^{\text{event}}, \mathbf{H}_t^{\text{word}}), \tag{10}$$

where $\mathbf{H}_0^{\text{event}}, \mathbf{O}_0^{\text{event}}$ are randomly initialized learnable embeddings of nodes and edges in event knowledge graphs. $\mathbf{H}_0^{\text{word}}$ denotes node initial features of word graphs, i.e., word embeddings. $f_{\text{Attn}}(\cdot)$ is an attention module that generates an attention vector for each node and edge in event knowledge graphs. Specifically, given a node in the event knowledge graph, its graph embedding is used as the context vector. The graph embeddings of the words that are semantically related to the node (obtained from $f_{\text{GCN}}(\cdot)$) are treated as the hidden states. The attention vector of the node is obtained according to Eq. 3. The enhanced representations of nodes and edges, as well as node embeddings of word graphs ($\{\tilde{\mathbf{H}}_t^{\text{event}}, \tilde{\mathbf{O}}_t^{\text{event}}, \mathbf{H}_t^{\text{word}}\}$) are then aggregated (e.g., pooling is applied) and fed into a GRU model for final event prediction. This approach can also infer the potential participants of the event of interest, providing additional information to practitioners.

Interpretation is crucial in societal event prediction because it can produce supporting evidence for reliable decision-making. More recently, Deng et al. [36] introduced an interpretable deep learning framework for forecasting civil unrest events and providing multilevel explanations for predictions. The approach aims to model three types of features including event frequencies, documents, and event graphs for a historical window, denoted as $\{\mathbf{x}, \mathcal{D}, \mathcal{G}\}_{t-k+1:t}$. The prediction model consists of multilevel feature learning that hierarchically models heterogeneous data. It captures the dependencies between different types of data by encouraging signals to propagate from higher-level features to lower-level features:

$$\mathbf{h}_t^{\text{freq}} = f_{\text{freq}}(\mathbf{x}_t), \quad \mathbf{H}_t^{\text{doc}} = f_{\text{doc}}(\mathcal{D}_t, \mathbf{h}_t^{\text{freq}}), \quad \mathbf{H}_t^{\text{event}} = f_{\text{event}}\big(f_{\text{CompGCN}}(\mathcal{G}_t, \mathbf{H}_t^{\text{doc}})\big). \tag{11}$$

Each type of feature is modeled at one level and then integrated to the next level. Similar to Glean [35], it uses a recurrent neural network to model temporal information for final event prediction.

The authors also proposed an event explainer module to provide temporal and multi-level explanations for the predictor model. The explainer learns temporal masks for edges and node features of temporal graphs to explain the predictions via mask optimization [150].

3 CRIME PREDICTION

Crime is an unlawful event that affects the harmony of humanity. It can harm individuals physically and mentally, and communities that experience higher crime rates are also adversely affected. Studying crime trends and patterns has been a top priority for law enforcement agencies, which use historical data to develop effective predictive policies to make a peaceful community [96]. A wealth of research has emerged on the prediction of crime incidences. The predictive capability can assist in crime prevention by facilitating the effective implementation of police patrols. In this section, we first formally introduce the problem of crime prediction, followed by a review of crime data resources. We then summarize traditional crime prediction methods and deep learning-based approaches.

3.1 Problem Formulation

We first introduce some mathematical notations and the data used in this research area. Then, we formulate some widely studied crime prediction problems.

Crime Occurrence. Assume there are I regions in an area (e.g., city) and J crime categories (e.g., burglary, assault, and criminal mischief) over T time slots (e.g., day, week, or month). The region can be a longitude-latitude grid cell or a community. We use a three-order tensor $\mathbf{X} \in \mathbb{R}^{I \times J \times T}$ to denote the crime data, where x_{ijt} is the number of j-th category crime incidents committed at the i-th region during the t-th time slot. Given a time window of length k, a historical crime tensor $\mathbf{X}_{t-k+1:t} = (\mathbf{X}_{t-k+1}, ..., \mathbf{X}_t)$ is defined to represent the number of occurrences of different categories of crimes over a historical window starting from the $(t - k + 1)$-th to t-th time slot. $\mathbf{X}_t$ is the crime matrix at time t, where the rows indicate the region and the columns denote the category of crime.

External Features. External features (e.g., meteorological data, traffic flow, and social media data) have been demonstrated to be tightly linked to crime incidents to facilitate the prediction of crime [134]. External features can be specified for each region i, each time slot j, or both. Simply put, we denote the external features of the whole area as $\mathcal{E}$. The external feature can be an empty set when the prediction is purely based on crime occurrence data.

Definition 3.1. **Crime Prediction as Classification**. Given the historical crime occurrence data $\mathbf{X}_{t-k+1:t}$ of window size k and external features $\mathcal{E}$, learn a classifier $f(\mathbf{X}_{t-k+1:t}, \mathcal{E}) \rightarrow \mathcal{Y}_{t+\Delta}$ that infers the future crime at each region as a categorical value. The target variable for region j can be written as an integer $\mathcal{Y}_{t+\Delta, j} \in \{0, 1, ..., M - 1\}$. The most common setting is a binary classification when M is 2, known as crime hotspot prediction. A hotspot is defined as a region that has a higher concentration of crime events. In some cases, the target value may indicate the range or severity of the number of crimes of interest, usually having $M > 2$.

Definition 3.2. **Crime Prediction as Regression**. Given the historical crime occurrence data $\mathbf{X}_{t-k+1:t}$ of window size k and external features $\mathcal{E}$, the goal is to learn a regression model $f(\mathbf{X}_{t-k+1:t}, \mathcal{E}) \rightarrow \mathcal{Y}_{t+\Delta}$ that predicts a continuous value of crime in the future time slot. The target variable $\mathcal{Y}_{t+\Delta}$ is usually defined as the (normalized) crime number , crime rate, or crime density matrix, i.e., $\mathbf{X}_{t+\Delta}$.

Definition 3.3. **Crime Prediction as Sequence Regression**. Given the historical crime occurrence data $\mathbf{X}_{t-k+1:t}$ of window size k and external features $\mathcal{E}$, the sequence regression task is to learn a prediction model $f(\mathbf{X}_{t-k+1:t}, \mathcal{E}) \rightarrow \{\mathcal{Y}_{t+\Delta}, ..., \mathcal{Y}_{t+\Delta+k'-1}\}$ that predicts the crime in the next k' steps. The target variables usually refer to the sequence of crime occurrences, i.e., $\{\mathbf{X}_{t+\Delta}, ..., \mathbf{X}_{t+\Delta+k'-1}\}$.

Definition 3.4. **Crime Prediction as Next-Location Prediction**. This problem aims to predict the location where an offender will commit a crime according to the offender's historical trajectories or other information. Formally, given an offender $u \in \mathcal{U}$ associated with their historical crime occurrence data $\mathbf{X}_{t-k+1:t}$ of window size k, external features $\mathcal{E}$, an underlying road network $\mathcal{R}$, and historical road trajectories information $\mathcal{RT}^u_{t-k+1:t}$, the next-location prediction is to develop a model $f(\mathbf{X}_{t-k+1:t}, \mathcal{E}, \mathcal{R}, \mathcal{RT}^u_{t-k+1:t}) \rightarrow \mathcal{Y}^u_{t+\Delta}$ that predict the next location where an offender u will commit a crime with highest likelihood. The target variable refers to a road or a road segment in the road network.

Note that crime occurrence data and target variables depend on specific studies, some of which consider each category of crime, and some combine several categories of crime into one group.

3.2　Data Resources

3.2.1　Crime Data. In recent years, crime data have been published online by governmental efforts, and fine-grained crime records are available for a few big cities, such as New York [2], Chicago [3], San Francisco [4], Atlanta [5], Philadelphia [6], Baltimore [7], Vancouver [8], and London [9], etc. Some are collected by online research communities, such as the crime data for Sao Paulo. [10] These data are incident-based and typically contain the incident date and time, crime type, demographic details, location data, etc. Crime can be divided into (1) violent crime such as homicide, rape, robbery, and (2) property crimes such as burglary, theft, and arson. In general, crime datasets in different regions follow different event definitions. Many of the crime predictive studies are based on public datasets from the cities mentioned above [12, 14, 51, 52, 55, 82, 103, 127, 156]. Some cities do not publish fine-grained criminal records but provide crime statistics, such as the number of crimes broken down by incident type, time, and location. For example, London provides statistics on homicides [11], business [12] crimes and hate crimes [13], and other crimes [14], etc. Researchers have utilized these statistical data as ground-truth data in crime prediction [3]. UCI machine learning repository [8] published a crime dataset for machine learning study, Communities and Crime [15], which combines socio-economic data from the 1990 US Census, law enforcement data from the 1990 US LEMAS survey, and crime data from the 1995 FBI UCR. It includes various attributes of cities in the United States and a predictor variable (i.e., crime rate). A set of studies were conducted on this dataset [6, 6, 57, 85, 134]. In addition, FBI's Crime Data Explorer (CDE) provides transparent and easy-access data in order to make awareness of criminal and noncriminal for law enforcement data sharing. It provides select datasets in the United States for download. [16] We organize some crime datasets in Table 3.

3.2.2　Open Source Indicators. We introduce some external data sources that have been shown to be helpful in crime prediction. Environmental context is often considered in crime predictive studies, such as meteorological

---

[2] https://data.cityofnewyork.us/Public-Safety/NYC-crime/qb7u-rbmr

[3] https://data.cityofchicago.org/Public-Safety/Crimes-2019/w98m-zvie

[4] https://data.sfgov.org/Public-Safety/Police-Department-Incident-Reports-2018-to-Present/wg3w-h783

[5] https://www.atlantapd.org/i-want-to/crime-data-downloads

[6] https://www.opendataphilly.org/dataset/crime-incidents

[7] https://data.baltimorecity.gov/datasets/part1-crime-data/explore

[8] https://geodash.vpd.ca/opendata/

[9] https://data.london.gov.uk/dataset/mps-homicide-dashboard-data

[10] https://www.kaggle.com/inquisitivecrow/crime-data-in-brazil

[11] https://data.london.gov.uk/dataset/mps-homicide-dashboard-data

[12] https://data.london.gov.uk/dataset/mps-business-crime-dashboard-data

[13] https://data.london.gov.uk/dataset/recorded_crime_rates

[14] https://data.london.gov.uk/dataset/recorded_crime_summary

[15] https://archive.ics.uci.edu/ml/datasets/Communities+and+Crime

[16] https://crime-data-explorer.fr.cloud.gov/pages/downloads#nibrs-downloads

Table 3. A summary of crime data used for crime prediction. All data listed are open source. Links to these data are provided in section 3.2.1.

Dataset	Data level	Temporal coverage	Spatial coverage	Geo encoding	Date encoding	Type	Summary	Provided by
New York City (NYC) Crime Data	Incident	2017-	New York City	Coordinates	Exact time	All	Date, location, type, and other description of the crime	New York Police Department
Chicago Crime Data	Incident	2001-2019	Chicago	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Chicago Police Department
San Francisco Crime Data	Incident	2018-	San Francisco	Coordinates	Exact time	All	Date, location, type, and other description of the crime	San Francisco Police Department
San Francisco Crime Data (Kaggle)	Incident	2003-2015	San Francisco	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Kaggle
Atlanta Crime Data	Incident	2009-	Atlanta	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Atlanta Police Department
Philadelphia Crime Data	Incident	2006-	Philadelphia	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Philadelphia Police Department
Baltimore Crime Data	Incident	1963-(few records in early years)	Baltimore	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Baltimore Police Department
Vancouver Crime data	Incident	2003-	Vancouver	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Vancouver Police Department
Brazil Crime Data	Incident	2007-2016	Sao Paulo	Coordinates	Exact time	All	Date, location, type, and other description of the crime	Kaggle
London Homicide Data	Incident	2003-	London boroughs	Borough	Month	Homicide	Date, location, type, and other description of the crime	Metropolitan Police Service
London Business Crime Data	Statistics	2011-	London boroughs	Borough	Month	Bussiness crime	Number of crimes by type, month and borough	Metropolitan Police Service
London Crime Rates Data	Statistics	1999-2017	London boroughs	Borough	Year	All	Numbers of crimes, and crime rates per thousand population, by type, year and borough.	Metropolitan Police Service
London Hate Crime and Special Crime Data	Statistics	2012-	London boroughs	Borough	Month	Hate crime, special crime	Number of crimes by type, month and borough	Metropolitan Police Service
London Crime Data With Different Geographic Breakdowns	Statistics	2001-	London boroughs	Borough, ward, and LSOA	Month	All	Number of crimes at three different geographic levels (Borough, ward, and LSOA) by type and month	Metropolitan Police Service
Communities and Crime Dataset [8]	/	-1995	USA	County	/	/	A set of attributes of a county derived from socio-economic and law enforcement data, and its crime rate	UCI Machine Learning Repository

data (e.g., temperature and weather) [127, 136, 139], demographic data (e.g., median age and race ratio) [14, 15, 134], geographic data (e.g., longitude and latitude) [134], Point-of-Interests (POI) data (e.g., shopping, sports and education) [14, 55, 136], urban environmental data (e.g., noise, traffic flow, taxi trip) [14, 103, 136, 139], and human behavior data (i.e., mobile data) [15, 103, 156]. A POI is a record of a place on a map that someone finds useful or interesting, typically defined by its geographical coordinates and a few additional attributes like name and category. Foursquare is a POI database from which some researchers have collected relevant data for crime prediction [103, 156]. For human behavior-related data, some studies utilized 311 public service complaint data to assist crime prediction [37, 55]. Such data are collected from 311 Service that documents urban complaint reports of different categories from citizens through a mobile app or phone calls. OpenStreetMap is a collaborative project that creates a free editable geographic database of the world. [17] In fine-grained crime analysis such as hot spot prediction, geographical information such as road network can be obtained from such data source [14, 147, 156]. Some researchers have explored visual data in crime prediction [156], such as Google Street View data, which provide more than 10 million miles of street view imagery across 83 countries. Social media data have also been investigated in this line of study. For instance, Twitter posts with rich and event-based context were leveraged for predicting criminal incidents [134, 156] and next-location of crime activities [131].

3.3 Traditional Methods

Geographic information systems (GIS) were the first and most common analytical tools for spatial data. GIS is useful for mapping and retrospectively finding links between criminal structures and various spatial and social conditions [24, 46, 59, 67], but itself does not provide much predictive power. With the increasing availability of fine-grained urban data, such as public service data, meteorological data, POI data, and human mobility data, data-driven crime prediction problems have received extensive attention from researchers for decades. Many data-driven approaches have emerged that enable more accurate and fine-grained predictions.

Early statistical methods for crime prediction were time series based analytical models. Gorr et al. [45] focused on predicting crime 30 days ahead for small areas, like police precincts. In a case study of Pittsburgh, they demonstrated that simple univariate time series models including Random Walk, Brown's Simple Exponential Smoothing [18], and Holt's Two-Parameter Linear Exponential Smoothing [18] are more accurate than naive methods commonly used by many police departments (i.e., Lag 12). Several statistical time series models have been investigated in this area, such as binomial regression [129, 130] and autoregressive models. Autoregressive Integrated Moving Average (ARIMA) was applied in one-week-ahead [29] and two-year-ahead [23] crime occurrence prediction. Some researchers studied autoregressive models in more complex scenarios by incorporating spatial analysis [21, 22, 116] and network analytic techniques [33] as well as developing hybrid models [4]. In addition to time series models, statistical Bayesian models were developed to forecast next crime locations [80] and to model spatial [12] or spatio-temporal [53] patterns of urban crime.

Nonparametric approaches have also been studied. Mohler et al. [88] implemented self-exciting point process models in the context of urban crime. A personalized random walk-based approach was introduced for spatial crime analysis and crime location prediction [122]. Andresen et al. [5] analyzed property crime through the use of a nonparametric spatial point pattern test that identifies the stability in spatial point patterns in pairwise and longitudinal contexts.

For traditional machine learning models, linear regression [85, 107, 131, 156], ridge and lasso regression [14, 87], support vector regression (SVR) [4, 14, 156] have been proposed to analyze the crime dynamics in regression problems. Classification models have been extensively studied in crime hotspot prediction, including Support Vector Machine (SVM) [68, 69], Logistic Regression [106], Naive Bayes [57, 144], shallow neural network [106, 144, 153], k-nearest neighbors algorithm [155], clustering based model [26, 56, 152], decision tree [1, 110], and

---

[17]https://www.openstreetmap.org/

ensemble models [3, 10, 15, 51, 57, 75, 106, 108, 144, 153]. Some researchers studied more advanced approaches. Zhao and Tang [164] modeled temporal-spatial correlations into a coherent optimization framework for crime prediction. Kadar et al. [63] developed an imbalance-aware hyper-ensemble for spatio-temporal crime prediction, focusing on low population density areas. Yi et al. [149] proposed a Clustered Continuous Conditional Random Field (Clustered-CCRF) model that exploits both spatial and temporal factors for crime prediction in an integrated way.

Most traditional methods rely on feature engineering and have limitations when modeling complex crime data. Recently, deep learning models have proven to be effective in crime prediction. We introduce the deep learning-based approach in the next section.

3.4   Deep Learning Models

Criminal events are often closely related to time and space, showing complex reoccurring patterns. Researchers have studied various deep learning-based models to capture spatial and temporal dependencies to predict crime and have made considerable progress. We present existing work according to model characteristics, mainly in the categories of Feedforward Neural Network (FNN), Recurrent Neural Networks (RNN), Attention, Autoencoder and Deep Generative approach, and Graph Neural Network (GNN). We provide a summary of deep learning-based models for crime prediction in Table 4.

3.4.1   Feedforward Neural Network Based Approaches. A feed-forward neural network (FNN) is a network without recurrent connections. One typical network is the fully-connected feed-forward neural network (FFNN), also known as multi-layered perceptron (MLP). Chun et al. [30] applied an FFNN model on individuals' criminal charge history to predict the severity of the crime at the individual level over the next 5 years. Kang and Kang [65] introduced an FFNN based model with a feature-level data fusion method for predicting daily crime hotspots. The model consists of several feature layers for multi-level feature representation and abstraction, where each feature layer operates independently using a set of features. They are then concatenated for joint feature representation learning. The authors introduced three sets of features, including spatial, temporal features, and environmental context features extracted by a CNN using image data.

CNN is a class of feedforward neural networks that can extract features from data with convolution kernels/filters. CNN has been extensively studied in crime prediction and has shown advantages in improving prediction accuracy by learning spatial features in crime data. The convolutional and pooling layers are two main building blocks of CNNs. Fig. 1 shows a simple example of convolution. In this example, the input is a two-dimensional tensor of size 3×3, and the kernel/convolution window is a two-dimensional tensor of size 2×2. The convolution window slides over the input tensor and performs element-wise multiplication with corresponding elements. The resulting tensor is summed up yielding a single scalar value. Convolution operation usually involves several other techniques such as padding, stride, and dilation, which control the size of the output. Pooling is a downsampling method that reduces the spatial size of the representation as well as the number of parameters and computation in the network. Fig. 2 shows an example of maximum pooling.
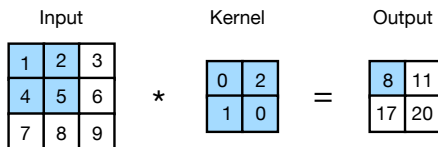

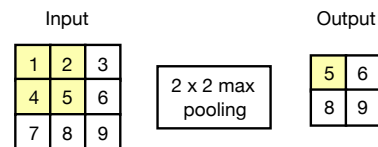Fig. 1. Two-dimensional convolution operation with a kernel of size 2×2.

Fig. 2. Maximum pooling with a pooling window of size 2×2.

Table 4. Deep learning based crime prediction techniques. Note that the intensity or rate of crime derives from the number of crimes, so we do not distinguish them in the inference column. We use horizontal lines to partition the models with different structural characteristics, corresponding to section 3.4.

Technique	Study area	Time	Dataset		Forecasting				Evaluation
			Crime type	External data	Inference	Task	Temporal unit	Spatial unit	
FFNN [30]	/	1997-2017	/	/	Severity of personal crime	Multi-label class.	5 years	/	Accuracy, F1
FFNN [65]	Chicago	2013-2014	All crime	Demographic, housing, economic, education, weather, and image data	Hotspots	Binary class.	Day	0.001 latitude×0.001 longitude grid cells	Accuracy, Precision, Recall, AUC
STCN [37]	NYC	2010-2015	Felony	311 data	Hotspots	Binary class.	Day	Grid cells of $0.18km^2$	F1, AUC
ST-ResNet [128]	Los Angeles	2015	All crime	Weather, holiday	# of crimes	Reg.	Hour	Grid cells	RMSE, Top-N-Accuracy
DIRNet [147]	NYC	2010-2015	3 theft crime types	311 data	Hotspots	Binary class.	Day	Grid cells of $0.28km^2$	F1
DuroNet [52]	Chicago, NYC	2016-2017, 2015-2016	All crime	/	# of crimes	Reg.	Day	Communities, Police precincts	RMSE, MAE
Sacked LSTM [138]	Poland	2008-2014, 2013-2016	12 crime types	/	# of crimes	Reg.	Day, Week, Year	Regions	MSE
STNN [165]	Portland, Oregon	2012-2016	3 crime types	/	Hotspots	Binary class.	2 weeks	Grid cells of $0.34km^2$	Accuracy, Precision, Recall, F1
RCN [119]	Chicago, Portland	2001-2018, 2012-2017	All crime	Weather, public transportation, and census data	# of crimes (in bins)	Class.	Day	Police beats, grid cells	MASE, Accuracy
CNN and RNN [118]	Seattle, Minneapolis, Philadelphia, San Fransisco, Washington, D.C.	1996-2016, 2010-2016, 2006-2017, 2003-2015, 2008-2017	All crime	/	Hotspots, # of crimes	Binary class., Reg.	Day	Grid cells of $0.2km^2$ and $0.64km^2$	F1, AUROC, AUPR, PAI
GSRNN [127]	Chicago, Los Angeles	2015, 2014-2015	All crime	/	# of crimes	Reg.	Hour	Zip code regions	RMSE, Precision Matrix
DeepCrime [55]	NYC	2014	4 crime types	POIs and 311 data	Hotspots	Binary class.	Day	Police precincts	F1
MiST [54]	NYC, Chicago	2015	4 crime types, 8 crime types	/	Hotspots	Binary class.	Day	2km×2km grid cells	F1, AUC
NN-CCRF [148]	NYC, Chicago	2015-2016, 2013-2015	2 crime types	/	# of crimes	Reg.	Day	1km×1km grid cells	RMSE, Hits@k
Crime-GAN [61]	San Francisco	2003-2018	4 crime types	/	# of crimes	Sequence reg.	Week	Grid cells	RMSE, MSPE, JS
CSAN [132]	San Francisco	2003-2018	4 crime types	/	# of crimes	Sequence reg.	Week	Grid cells	RMSE, MSPE, JS
TGCN [62]	San Francisco	2003-2019	4 crime types	/	# of crimes	Sequence reg.	Week	Grid cells	RMSE, MSPE, JS
DT-MGCN [136]	Chicago	2001-2014	10 crime types	Census, Taxi flow, traffic violations, POIs data	# of crimes	Sequence reg.	Month	Communities	MAE, RMSE, MRE
CrimeSTC [139]	NYC	2014	4 crime types	Meteorology, urban anomaly, demographic and POIs data	# of crimes	Reg.	Day	Police precincts	RMSE, MAE
ST-SHN [143]	NYC, Chicago	2014-2015, 2016-2017	4 crime types	/	Hotspots, # of crimes	Binary class., Reg.	Day	3km×3km grid cells	F1, RMSE, MAE

Note mAP (Mean average precision); PAI (Prediction Accuracy Index); JS (Jensen-Shannon Divergence); MRE (Mean Relative Error)

Fu et al. [40] developed a CNN model with stacked convolutional layers to extract latent features of street view images to infer the ranking of crime types. Instead of stacking convolution layers deeper and deeper in a CNN model, many effective and computationally efficient architectures have been developed. Some of them were adopted for crime prediction, such as Inception Network [121] and Residual Neural Network (ResNet) [49]. An inception network is a deep neural network with an architectural design that consists of repeating components referred to as inception modules, as shown in Fig. 3. A residual network, as shown in Fig. 4, consists of residual blocks which have skip connections, also called identity connections. The identity mapping helps in tackling the vanishing gradient problem in networks with a large number (even thousands) of layers.
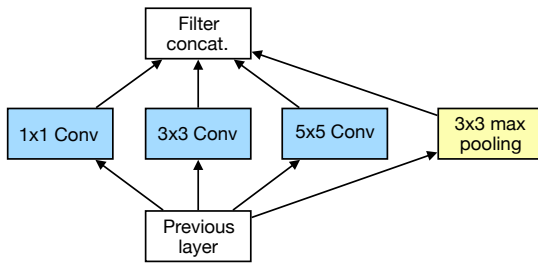
Fig. 3. Inception module, naive version [121].   Fig. 4. A residual block [49].
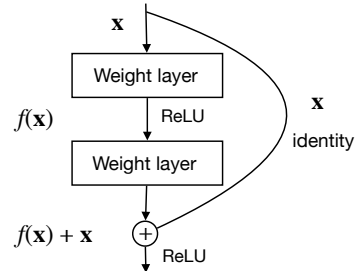
Based on the achievements of CNN models, Duan et al. proposed a CNN-based Spatiotemporal Crime Network (STCN) [37], to predict the occurrence of felonies on the next day. They employed crime event counts and 311 data and converted them into two 2D image-like arrays as input features. The STCN consists of convolutional layers, an inception block, a fractal block [76] (i.e., an alternative of the residual block that can achieve excellent performance as standard residual networks), pooling layers, and fully connected layers. Convolutional layers capture low-level spatio-temporal dependencies of criminal events, and the inception and fractal blocks intend to abstract high-level spatio-temporal features. Wang et al. [128] applied a popular deep CNN architecture, Spatio-Temporal Residual Networks (ST-ResNet) [154], to predict crime distributions over the Los Angeles area on an hourly timescale. The ST-ResNet captures crime dynamics through convolutional layers and employs residual neural networks to model the temporal features of crimes. Due to the low regularity of crime data in both space and time, the authors preprocessed the crime data to select appropriate spatio-temporal scales for prediction and proposed different data regularization methods for spatial and temporal dimensions. Furthermore, to address the problem of the sparsity of crime events in space and time, Ye et al. [147] proposed a deep inception-residual network (DIRNet) to conduct theft-related crime prediction based on crime and 311 data. This framework consists of convolutional layers and inception layers to extract spatiotemporal dependencies from crime and 311 data, respectively. The learned features are then combined and fed to stacked residual layers to capture high-level feature interactions for final prediction.

3.4.2   Recurrent Neural Network Based Approaches. RNNs have been developed to model the temporal dependencies of time-series crime data. In the context of crime prediction, Wawrzyniak et al. [138] studied short-term prediction for a selected crime type at daily, weekly, and annually levels by using a stacked LSTM architecture. The model does not handle spatial information and performs prediction for each region separately. Since crime is not randomly distributed geographically, considering the status of neighboring regions is beneficial for predicting the target region. Zhuang et al. [165] formulated the problem of crime forecasting as a space-time series prediction problem and introduced a Spatio-Temporal neural network (STNN), where RNN models are fed with spatial

embeddings. The spatial embeddings are the accumulation of the short-term crime numbers of a space-time window.

Some researchers combine CNN and RNN models to handle both the temporal and spatial features of crime patterns. Stec and Klabjan [119] used a joint recurrent and convolutional neural network (RCN) to predict crime while combining crime data with additional weather, public transportation, and census data. In this network, spatial features are passed through the CNN layer, and the output is combined with other features and then fed to the RNN model. Another work investigated the capability of several deep learning architectures, including CNNs and RNNs, to forecast crime hotspots in an urban environment [118]. This work examined three deep learning architecture configurations for crime prediction based on the encoding order of spatial and temporal patterns, where spatial features are modeled using CNN models such as ResNet, and RNNs learn temporal features.

Researchers have begun to investigate diverse representations of data for crime analysis. To address the challenges of forecasting spatio-temporal distributions at hourly or even finer temporal scales for crime data, Wang et al. [127] developed a graph-based multiscale framework to model sparse and unstructured spatial-temporal data. The flow chart of the framework is shown in Fig. 5. The framework contains two components. The first part is spatial-temporal weighted graph (STWG) inference. It results in a graph representation for the spatio-temporal evolution of the data. Each graph node is associated with a time series of crime intensity in a region, and the graph topology is inferred from the self-exciting point process model. The graph representation captures versatile spatial partitioning that enables forecasting at different spatial scales. The second component is a deep learning model that approximates the temporal evolution of the data. They introduced a graph-structured RNN (GSRNN) on the inferred graph based on the structural-RNN (SRNN) [58] architecture. The GSRNN arranges RNNs in a feedforward manner. It first assigns a cascaded LSTM to fit the time series on each node in the graph. Simultaneously, it associates each edge in the graph with a cascaded LSTM that receives the output from neighboring nodes along with the weights learned from the Hawkes process. The specially designed deep neural network takes advantage of the RNN's ability to learn time series patterns, capturing real-time interactions of each node to its connected neighbors.
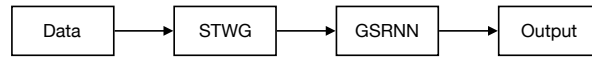

Fig. 5. The flow chart of the GSRNN framework. [127].

3.4.3 Attention Based Approaches. Many studies have shown that the combination of attention mechanisms in RNNs allows them to focus on parts of the input sequence when predicting the output sequence, making learning easier and of higher quality. With the broad exploration of attention in various research fields, more advanced attention mechanisms have been studied to model complex and rich features and improve the predictive ability of crime prediction. In this regard, Huang et al. proposed an attentive hierarchical recurrent network,
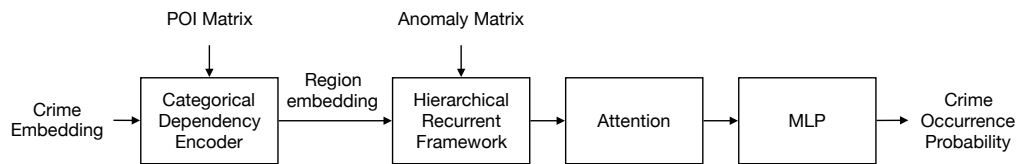

Fig. 6. The DeepCrime Framework [55].

DeepCrime [55], that uncovers dynamic crime patterns and explores the evolving inter-dependencies between

crimes and other ubiquitous data in urban space. The framework, as shown in Fig. 6, first models region-category interactions to generate a region embedding vector. It then includes a three-level GRU architecture that encodes the temporal dynamics of crime patterns and their interrelations with urban anomalies (i.e., 311 anomalies). The model further consists of an attention layer to capture the unknown temporal relevance and automatically assign importance weights to the learned hidden states at different time frames. The MLP is then applied to the learned attention features to output crime predictions.
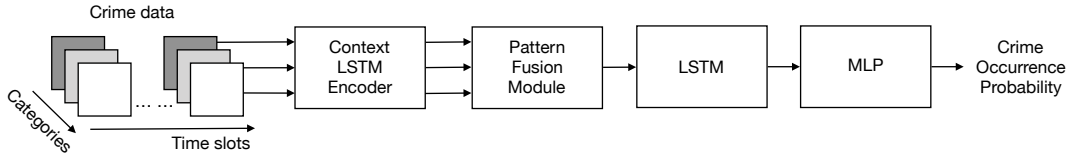
Fig. 7. The MiST Framework [54].

More recently, Huang et al. studied a Multi-View and Multi-Modal Spatial-Temporal learning framework (MiST) [54], to explicitly model the dynamic patterns of citywide abnormal events from spatial-temporal-categorical views. Fig. 7 presents the framework of the MiST model. The framework first models dynamic intra-region correlations from temporal views with a context-LSTM encoder. Then, a pattern fusion module with the attention mechanism captures complex inter-region and cross-categorical correlations from spatial-categorical views. Additionally, it incorporates an LSTM to capture the sequential patterns of cross-modal correlations between locations, times, and event categories. The proposed approach was evaluated on real-world crime and urban anomaly data and compared to state-of-the-art baselines to demonstrate its superior performance.

Instead of combining recurrent neural networks with attention, Hu et al. incorporated attention into convolutional networks. They studied the effect of noises (i.e., local outliers and irregular waves) on crime data and proposed a Dual-robust Enhanced Spatial-temporal Learning Network (DuroNet) [52] with an encoder-decoder architecture to capture deep crime patterns. The encoder consists of (1) a CNN based locality enhanced module that employs local temporal context information to smooth the deviation of outliers and uses gate mechanism to enhance spatial-temporal representation, and (2) a self-attention based pattern representation module to weaken the effect of irregular waves by learning attentive weights. A self-attention module compares every element in the sequence to every other element in the sentence, including itself, and reweigh the embeddings of each element to include contextual relevance. Finally, a feed-forward prediction network with convolutional layers serves as the decoder for crime prediction.

3.4.4 Autoencoder and Deep Generative Approaches. Autoencoders are a type of self-supervised learning model that can learn a compressed/reconstructed representation of input data, which are commonly used for feature selection and extraction. An autoencoder is composed of an encoder and a decoder sub-models, as shown in Fig. 8.

Yi et al. [148] adopted autoencoder to learn dynamic spatial correlations in a deep learning framework. To achieve fine-grained crime prediction at the daily level, this paper leveraged the mean-field approximation theory to simplify the inference of continuous conditional random field (CCRF) [105] model. They proposed an end-to-end neural network NN-CCRF[148] to model temporal and spatial correlations, respectively. The traditional CCRF model is one of the probabilistic graphical models that can explore sequential relationships in time series data. It consists of two parts, unary potential, and pairwise potential. The NN-CCRF model applies LSTM as the unary potential and leverages Stacked Denoising AutoEncoder (SDAE) to learn spatial correlations across regions used in pairwise potential. The Denoising AutoEncoder [126] attempts to obtain robust latent representations by introducing a stochastic noise to the original data.
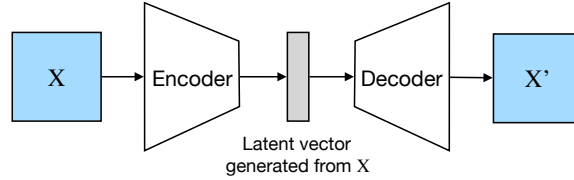
Fig. 8. Illustration of the autoencoder model architecture.

Recently, sequence generative networks have gained considerable attention in various fields such as computer vision and natural language processing. Researchers have found it promising in capturing the spatio-temporal dynamics of massive data. Some studies provided new insights into understanding complex phenomena in crime data from sequence generation. Wang et al. proposed a sequence generative neural network, named Crime Situation Awareness Network (CSAN) [132], for crime sequence prediction. The task is illustrated in Fig. 9. The
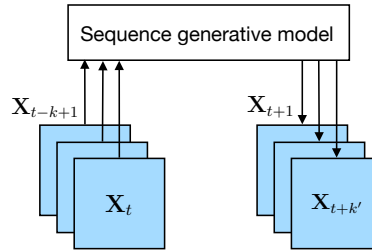

Fig. 9. Illustration of the crime sequence prediction task. A sequence generative model generates the next k' crime distributions conditioned on a sequence of previous k crime distributions.

CSAN includes two components: (1) a compact representation model with Multiple-VAEs attempts to embed high-dimensional sparse crime data into some compact latent spaces; (2) a crime situation generation model with GRU and attention layers captures potential dynamics in latent space, leading to sequence prediction. Variational Autoencoder (VAE) [70] consists of an encoder that compresses the input data into a constrained multivariate latent distribution, and a decoder that reconstructs the data given the latent distribution.

There is another effort focusing on crime sequence generation. Jin et al. developed a context-based sequence generative model in a generative adversarial network, Crime-GAN [61]. Generative adversarial network (GAN) [44] is a deep generative model that has shown tremendous success over the last few years. GAN is built on the principle of adversarial loss. It includes a Generator and a Discriminator, which follows adversarial nature to fool each other. The framework is presented in Fig. 10. In the Crime-GAN model, a CNN-based VAE first extracts latent variables from the crime data. Then, the latent variables are fed into a GRU-based Seq2Seq model to generate latent variables for the next k' time step. The model uses the GRU model as a discriminator to examine whether the generated sequence of crime latent variables follows the previous sequence of crime latent variables.

3.4.5 Graph Neural Network Based Approaches. Graph neural networks have shown the ability to capture spatially structured features beyond the Euclidean distance. There have been attempts to incorporate such networks in crime prediction. Jin et al. proposed a graph deep learning approach, Temporal Graph Convolutional Network (TGCN) [62] to predict crime sequences. The authors first constructed a sequence of crime spatial correlation graphs using the Pearson coefficient. Each node in the graph represents a region. TGCN consists of three main components: (1) a GCN [71] to extract latent representations from a defined crime spatial correlation graph;
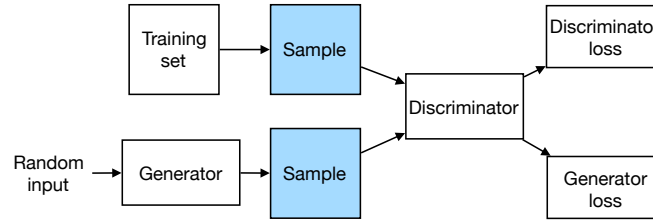
Fig. 10. The framework of the generative adversarial network (GAN) model.

(2) an LSTM to capture temporal dynamics of the latent space; and (3) a CNN-based model to output crime predictions. Researchers have incorporated more sophisticated graph information in crime sequence prediction. Wang et al. proposed a Deep Temporal Multi-Graph Convolutional Network (DT-MGCN) [136] that integrates a graph generation component and a spatial-temporal component to forecast spatial-temporal crime rate. The authors encoded various external factors into multiple graphs, including distance graph, interaction graph, and correlation graph, to capture the Euclidean and non-Euclidean correlations among communities, and generated a fused graph by combining the aforementioned graphs. The spatial-temporal component simultaneously employs spectral graph convolutions to capture spatial patterns and uses an encoder-decoder temporal convolutional network (EDTCN) [77] to model temporal features.
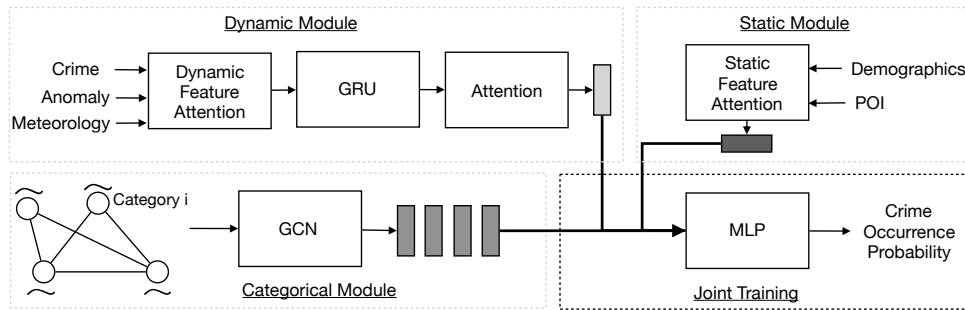

Fig. 11. The CrimeSTC Framework [139].

Graph neural networks have also been used to capture the dependencies among different crime categories to help predict the number of crimes. In particular, Wei et al. proposed a crime prediction framework, CrimeSTC [139] to jointly learn intricate spatial-temporal categorical correlations hidden inside crime and big urban data. The model framework is shown in Fig. 11. The model consists of four parts: (1) a dynamic module that processes temporal data through local CNN, GRU, and attention; (2) a static module handling static data through fully connected layers; (3) a categorical module that captures categorical dependencies through the GCN model; and (4) a joint training module concatenating dynamic and static representations to forecast crime numbers. More recently, Xia et al. proposed a Spatial-Temporal Sequential Hypergraph Network (ST-SHN) [143] architecture that achieves the state-of-the-art performance in crime occurrence prediction. ST-SHN collectively encodes complex crime spatial-temporal patterns as well as the underlying category-wise crime semantic relationships. The model framework is provided in Fig. 12. Specifically, the authors proposed an attention-based multi-channel routing mechanism to learn cross-type crime influences under the graph neural network framework. They designed a graph message passing architecture that integrates a hypergraph learning paradigm to enhance the cross-region

relation learning without the limitation of adjacent connections. To model the temporal dependencies of crime, a graph temporal shift mechanism is developed to inject the time-varying spatial-temporal crime patterns into the representation process.
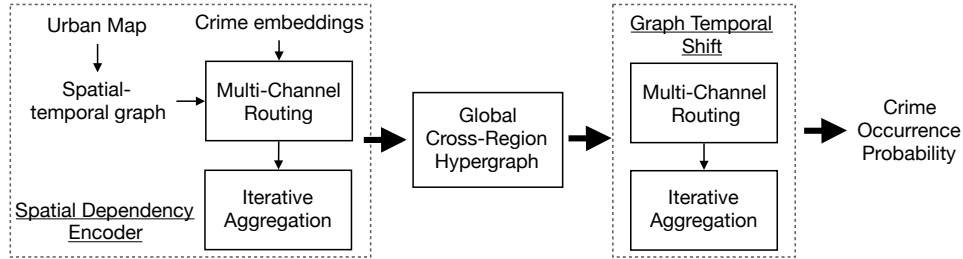

Fig. 12. The ST-SHN Framework [143].

4 DEEP LEARNING TECHNIQUES

Deep learning is a family of machine learning methods based on artificial neural networks and representation learning. In this section, we summarize the utilization of deep neural networks in civil unrest and crime prediction.

Convolutional Neural Networks. CNN is designed to automatically learn spatial hierarchies of features through backpropagation by using multiple building blocks, such as convolution layers, pooling layers. Compared to limited research on CNNs in civil unrest prediction more studies have explored CNNs in crime prediction [37, 52, 128, 147, 156]. This is mainly because that the spatial granularity of criminal events is usually smaller than civil unrest. Civil unrest such as demonstrations and riots can occur throughout a city. It usually produces a massive impact, spreading online or offline to influence people's behavior in surrounding areas or other cities. In contrast, crimes such as burglary and robbery may be more specific to a street or community. In some studies that focus on fine-grained crime prediction, researchers divide a region (e.g., a city) into grid cells and forecast crime for each cell. CNN provides an adaptive and effective way to learn spatial correlations among cells.

Recurrent Neural Networks. A recurrent neural network (RNN) is a class of artificial neural networks with recurrent connections. It uses the internal state (memory) to process input sequences of various lengths. Such models have been widely involved in societal event prediction and often outperform statistical time series models such as autoregressive models. RNN based methods capture nonlinearities and more complex dependencies in sequential data. In dynamic social environments, societal events usually do not occur independently. They are are often influenced by past events. This explains why historical time series data are used as important indicators when forecasting civil unrest and crime. To better capture temporal dependencies, researchers have used RNNs and their variants (i.e., LSTM, GRU) as the main components in temporal feature learning for civil unrest [47, 86, 95, 117, 135] and crime [54, 55, 118, 119, 127, 138, 148, 165] prediction.

Attention. Attention models, or attention mechanisms, are techniques in neural networks that allow a network to focus on specific aspects of a complex input, guided by a context vector (i.e., query). Attention is usually applied together with RNN models and is able to show the importance of features through attention weights. In the context of civil unrest prediction, Wang et al. [135] introduced the attention mechanism to discover crucial historical data points. Ertugrul et al. [38] studied a more advanced attention-based approach, where attention mechanisms are used to explain the importance of activities across different historical time steps and different regions. Some

researchers applied attention to enhance representation learning with contextual information [35]. With similar motivations, more sophisticated models have introduced attention mechanisms in crime prediction [54, 55, 139].

Autoencoder and Deep Generative Approach. Autoencoder is a type of artificial neural network used to learn efficient data representations in an unsupervised manner. Yi et al. [126] leveraged the Denoising AutoEncoder to learn latent representations to capture spatial correlations across regions for crime prediction. Others utilized Variational AutoEncoder (VAE) to learn a compressed latent representation of crime matrices [61, 132]. VAE can be regarded as a deep generative model. Deep generative models provide a powerful way to approximate complicated and high-dimensional probability distributions using a large number of samples. Some researchers have incorporated generative adversarial networks (GAN) to estimate future crime distributions across a region in a sequential predict task [61].

Graph Neural Networks. Graph Neural Networks (GNNs) are a class of deep neural networks that can be directly applied to graphs and provide an easy way for node-level, edge-level, and graph-level prediction tasks. In civil unrest and crime prediction, researchers have explored network structures for different purposes. Deng et al. [34] proposed a dynamic graph neural network to predict civil unrest, which takes a sequence of dynamic word graphs as input. They further discovered key contextual subgraphs as supporting evidence for event prediction. The subgraphs provide structured and concise data representations which can streamline the process of event analysis. Other studies have applied GNN models to knowledge graph-based data to enhance contextual information for event prediction by learning multi-relational features [35, 36]. For crime prediction, some work employed GNN models to enhance representation learning via capturing spatial correlations among pre-defined graph nodes, thereby improving forecasting accuracy [62, 136, 139, 143].

5 OPEN CHALLENGES AND FUTURE DIRECTIONS

Despite some achievements in societal event prediction in recent years, there are still open problems. In what follows, we discuss some challenges and future directions from three aspects.

5.1 Data Dynamics, Sufficiency, and Reliability

Data-driven approaches for societal event prediction depend heavily on data quality, which makes them subject to a number of data challenges. The dynamic nature of data is one of the main challenges. For text data, language, vocabulary, and mainstream slang are constantly evolving. In geographic data, location names and area boundaries may change due to major political events. Proposing advanced deep learning methods that can overcome the problem of data dynamics is of great significance for practical applications. The sufficiency of data is another challenge. As seen in existing work, researchers have investigated various external data sources in addition to historical event occurrences to improve the accuracy of predictions. Collecting external data from multiple sources and distinguishing correlated data from noisy data is expensive in terms of time, material, and computational costs. Moreover, the spatial scarcity of societal events will also hinder event prediction studies in underrepresented areas. A further fundamental issue in forecasting problems is data reliability. Missing or incorrect data often occurs during manual or automated data collection. For example, automated event collection systems may lose events due to unexpected network failures. Social media posts provide valuable resources for tracking user behavior and social activities. However, such data include typos, chit-chat, and misinformation that can mislead predictive models. Hence, over-reliance on data can make prediction models vulnerable to real-world applications.

5.2 Transparency and Interpretability of Deep Learning Models

Deep learning models have become increasingly sophisticated with increased predictive power. However, as model structures become more complex, the underlying mechanisms of predictive models become more opaque

to humans, making model outputs more difficult to understand. There is a growing awareness of the need for improving model transparency and interpretability. Relying solely on model predictions to analyze societal events and make decisions can lead to economic loss or other unwanted consequences due to misjudgments. One can mitigate such problems by providing reasonable explanations for model predictions so that practitioners can understand the behavior of the models and thus assess the reliability of the output produced by these models. For societal event prediction, it is beneficial to develop explanatory models that can identify critical information from copious amounts of historical data. Providing explanations for event prediction can improve the robustness of efficient event modeling supported by machine learning.

5.3 Causality Study in Societal Events

Machine learning methods generally focus on discovering correlations between input features and target variables rather than understanding causality. Causality research or causal inference aims to study the cause and effect relationships (i.e., causality) between two or more variables. To better understand societal events and forecast these events with greater confidence, introducing causality into event modeling has potential benefits. Analyzing probable causes of future events may allow us to develop causality-enhanced event prediction models that are not susceptible to data problems, such as noise and sparsity. Investigating the causality of societal events can also help us reveal the underlying mechanisms or critical factors behind event occurrences in dynamic social environments. A growing number of studies have absorbed the advantages of deep learning models, as well as causal studies [28, 79]. The incorporation of causality in deep-learning based event predictions is becoming an indispensable research topic.

6 CONCLUSION

In this article, we present a comprehensive survey of existing methods for societal event forecasting, focusing on deep learning-based approaches. We outline the challenges in societal event prediction. We summarize research papers published in recent years and discuss data resources, research problems, traditional and deep learning predictive techniques for civil unrest and crime events, respectively. In addition, we provide an overview of deep learning models that have been employed in prediction models, point out their characteristics and advantages in addressing different problems, and discuss open challenges and promising directions for future work.

REFERENCES

[1] Tahani Almanie, Rsha Mirza, and Elizabeth Lor. 2015. Crime prediction based on crime types and using spatial and temporal criminal hotspots. *arXiv preprint arXiv:1508.02050* (2015).

[2] Scott Althaus, Joseph Bajjalieh, John Carter, Buddy Peyton, and Dan Shalmon. 2020. Cline Center Historical Phoenix Event Data. Cline Center for Advanced Social Research. v1.3.0. May 4. https://doi.org/10.13012/B2IDB-0647142_V3

[3] Luiz GA Alves, Haroldo V Ribeiro, and Francisco A Rodrigues. 2018. Crime prediction through urban metrics and statistical learning. *Physica A: Statistical Mechanics and its Applications* 505 (2018), 435–443.

[4] Razana Alwee, Siti Mariyam Shamsuddin, and Roselina Sallehuddin. 2013. Hybrid support vector regression and autoregressive integrated moving average models improved by particle swarm optimization for property crime rates forecasting with economic indicators. *The Scientific World Journal* 2013 (2013).

[5] Martin A Andresen, Shannon J Linning, and Nick Malleson. 2017. Crime at places and spatial concentrations: Exploring the spatial stability of property crime in Vancouver BC, 2003–2013. *Journal of quantitative criminology* 33, 2 (2017), 255–275.

[6] Syahid Anuar, Ali Selamat, and Roselina Sallehuddin. 2015. Hybrid artificial neural network with artificial bee colony algorithm for crime classification. In *Computational Intelligence in Information Systems*. Springer, 31–40.

[7] Bryan Arva, John Beieler, Bejamin Fisher, Gustavo Lara, Philip A Schrodt, Wonjun Song, Marsha Sowell, and Sam Stehle. 2013. Improving forecasts of international events of interest. In *EPSA 2013 Annual General Conference Paper*, Vol. 78.

[8] Arthur Asuncion and David Newman. 2007. UCI machine learning repository.

[9] Edward E Azar. 1980. The conflict and peace data bank (COPDAB) project. *Journal of Conflict Resolution* 24, 1 (1980), 143–152.

[10] Maria Jeseca C Baculo, Charlie S Marzan, Remedios de Dios Bulos, and Conrado Ruiz. 2017. Geospatial-temporal analysis and classification of criminal data in manila. In *2017 2nd IEEE International Conference on Computational Intelligence and Applications (ICCIA)*. IEEE, 6–11.

[11] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).

[12] Cecilia Balocchi and Shane T Jensen. 2019. Spatial modeling of trends in crime over time in Philadelphia. *The Annals of Applied Statistics* 13, 4 (2019), 2235–2259.

[13] World Bank. 2016. World development report 2016: Digital dividends.

[14] Alexandros Belesiotis, George Papadakis, and Dimitrios Skoutas. 2018. Analyzing and predicting spatial crime distribution using crowdsourced and open data. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 3, 4 (2018), 1–31.

[15] Andrey Bogomolov, Bruno Lepri, Jacopo Staiano, Nuria Oliver, Fabio Pianesi, and Alex Pentland. 2014. Once upon a crime: towards crime prediction from demographics and mobile data. In *Proceedings of the 16th international conference on multimodal interaction*. 427–434.

[16] Elizabeth Boschee, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz, and Michael Ward. 2015. ICEWS Coded Event Data.

[17] Dan Braha. 2012. Global civil unrest: contagion, self-organization, and prediction. *PloS one* 7, 10 (2012), e48596.

[18] Robert Goodell Brown. 2004. *Smoothing, forecasting and prediction of discrete time series.* Courier Corporation.

[19] Umair Muneer Butt, Sukumar Letchmunan, Fadratul Hafinaz Hassan, Mubashir Ali, Anees Baqir, and Hafiz Husnain Raza Sherazi. 2020. Spatio-Temporal Crime HotSpot Detection and Prediction: A Systematic Literature Review. *IEEE Access* 8 (2020), 166553–166574.

[20] 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.

[21] Charlie Catlett, Eugenio Cesario, Domenico Talia, and Andrea Vinci. 2018. A data-driven approach for spatio-temporal crime predictions in smart cities. In *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 17–24.

[22] Charlie Catlett, Eugenio Cesario, Domenico Talia, and Andrea Vinci. 2019. Spatio-temporal crime predictions in smart cities: A data-driven approach and experiments. *Pervasive and Mobile Computing* 53 (2019), 62–74.

[23] Eugenio Cesario, Charlie Catlett, and Domenico Talia. 2016. Forecasting crimes using autoregressive models. In *2016 IEEE 14th Intl Conf on Dependable, Autonomic and Secure Computing, 14th Intl Conf on Pervasive Intelligence and Computing, 2nd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech)*. IEEE, 795–802.

[24] Spencer Chainey and Jerry Ratcliffe. 2013. *GIS and crime mapping.* John Wiley & Sons.

[25] Deepayan Chakrabarti and Kunal Punera. 2011. Event summarization using tweets. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 5.

[26] B Chandra, Manish Gupta, and MP Gupta. 2008. A multivariate time series clustering approach for crime trends prediction. In *2008 IEEE International Conference on Systems, Man and Cybernetics*. IEEE, 892–896.

[27] Feng Chen and Daniel B Neill. 2014. Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1166–1175.

[28] Jin Chen, Xinxiao Wu, Yao Hu, and Jiebo Luo. 2021. Spatial-temporal Causal Inference for Partial Image-to-video Adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 1027–1035.

[29] Peng Chen, Hongyong Yuan, and Xueming Shu. 2008. Forecasting crime using the arima model. In *2008 Fifth International Conference on Fuzzy Systems and Knowledge Discovery*, Vol. 5. IEEE, 627–630.

[30] Soon Ae Chun, Venkata Avinash Paturu, Shengcheng Yuan, Rohit Pathak, Vijayalakshmi Atluri, and Nabil R. Adam. 2019. Crime prediction model using deep neural networks. In *Proceedings of the 20th Annual International Conference on Digital Government Research*. 512–514.

[31] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555* (2014).

[32] Sarah Zukerman Daly. 2012. Organizational legacies of violence: Conditions favoring insurgency onset in Colombia, 1964–1984. *Journal of Peace Research* 49, 3 (2012), 473–491.

[33] Saroj Kumar Dash, Ilya Safro, and Ravisutha Sakrepatna Srinivasamurthy. 2018. Spatio-temporal prediction of crimes using network analytic approach. In *2018 IEEE International Conference on Big Data (Big Data)*. IEEE, 1912–1917.

[34] Songgaojun Deng, Huzefa Rangwala, and Yue Ning. 2019. Learning Dynamic Context Graphs for Predicting Social Events. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1007–1016.

[35] Songgaojun Deng, Huzefa Rangwala, and Yue Ning. 2020. Dynamic Knowledge Graph Based Multi-Event Forecasting. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1585–1595.

[36] Songgaojun Deng, Huzefa Rangwala, and Yue Ning. 2021. Understanding Event Predictions via Contextualized Multilevel Feature Learning. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 342–351.

[37] Lian Duan, Tao Hu, En Cheng, Jianfeng Zhu, and Chao Gao. 2017. Deep convolutional neural networks for spatiotemporal crime prediction. In *Proceedings of the International Conference on Information and Knowledge Engineering (IKE)*. The Steering Committee of The World Congress in Computer Science, Computer . . . , 61–67.

[38] Ali Mert Ertugrul, Yu-Ru Lin, Wen-Ting Chung, Muheng Yan, and Ang Li. 2019. Activism via attention: interpretable spatiotemporal learning to forecast protest activities. *EPJ Data Science* 8, 1 (2019), 1–26.

[39] Ioannis Flouris, Nikos Giatrakos, Antonios Deligiannakis, Minos Garofalakis, Michael Kamp, and Michael Mock. 2017. Issues in complex event processing: Status and prospects in the big data era. *Journal of Systems and Software* 127 (2017), 217–236.

[40] Kaiqun Fu, Zhiqian Chen, and Chang-Tien Lu. 2018. Streetnet: preference learning with convolutional neural network on urban crime perception. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 269–278.

[41] Ruchika Parmarth Ganar and Shrikant Ardhapurkar. 2016. Prediction of civil unrest by analysing social network using keyword filtering: A survey. In *2016 Online International Conference on Green Engineering and Technologies (IC-GET)*. IEEE, 1–4.

[42] Yuyang Gao and Liang Zhao. 2018. Incomplete label multi-task ordinal regression for spatial event scale forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

[43] Deborah J Gerner, Philip A Schrodt, Omür Yilmaz, and Rajaa Abu-Jabr. 2002. Conflict and mediation event observations (CAMEO): A new event data framework for the analysis of foreign policy interactions. *International Studies Association, New Orleans* (2002).

[44] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf

[45] Wilpen Gorr, Andreas Olligschlaeger, and Yvonne Thompson. 2003. Short-term forecasting of crime. *International Journal of Forecasting* 19, 4 (2003), 579–594.

[46] Elizabeth R Groff and Nancy G La Vigne. 2002. Forecasting the future of predictive crime mapping. *Crime Prevention Studies* 13 (2002), 29–58.

[47] Matina Halkia, Stefano Ferri, Michail Papazoglou, Marie-Sophie Van Damme, and Dimitrios Thomakos. 2020. Conflict Event Modelling: Research Experiment and Event Data Limitations. In *Proceedings of the Workshop on Automated Extraction of Socio-political Events from News 2020*. 42–48.

[48] Jesse Hammond and Nils B Weidmann. 2014. Using machine-coded event data for the micro-level study of political violence. *Research & Politics* 1, 2 (2014), 2053168014539924.

[49] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[50] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.

[51] Sohrab Hossain, Ahmed Abtahee, Imran Kashem, Mohammed Moshiul Hoque, and Iqbal H Sarker. 2020. Crime prediction using spatio-temporal data. In *International Conference on Computing Science, Communication and Security*. Springer, 277–289.

[52] Kaixi Hu, Lin Li, Jianquan Liu, and Daniel Sun. 2021. DuroNet: A Dual-robust Enhanced Spatial-temporal Learning Network for Urban Crime Prediction. *ACM Transactions on Internet Technology (TOIT)* 21, 1 (2021), 1–24.

[53] Tao Hu, Xinyan Zhu, Lian Duan, and Wei Guo. 2018. Urban crime prediction based on spatio-temporal Bayesian model. *PloS one* 13, 10 (2018), e0206215.

[54] 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 *The World Wide Web Conference*. 717–728.

[55] Chao Huang, Junbo Zhang, Yu Zheng, and Nitesh V Chawla. 2018. DeepCrime: Attentive hierarchical recurrent networks for crime prediction. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 1423–1432.

[56] Varvara Ingilevich and Sergey Ivanov. 2018. Crime rate prediction in the urban environment using social factors. *Procedia Computer Science* 136 (2018), 472–478.

[57] Rizwan Iqbal, Masrah Azrifah Azmi Murad, Aida Mustapha, Payam Hassany Shariat Panahy, and Nasim Khanahmadliravi. 2013. An experimental study of classification algorithms for crime prediction. *Indian Journal of Science and Technology* 6, 3 (2013), 4219–4225.

[58] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. 2016. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the ieee conference on computer vision and pattern recognition*. 5308–5317.

[59] Ákos Jakobi and Andrea Pődör. 2020. GIS-based statistical analysis of detecting fear of crime with digital sketch maps: A Hungarian multicity study. *ISPRS International Journal of Geo-Information* 9, 4 (2020), 229.

[60] Fang Jin, Rupinder Paul Khandpur, Nathan Self, Edward Dougherty, Sheng Guo, Feng Chen, B Aditya Prakash, and Naren Ramakrishnan. 2014. Modeling mass protest adoption in social network communities using geometric brownian motion. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1660–1669.

[61] Guangyin Jin, Qi Wang, Xia Zhao, Yanghe Feng, Qing Cheng, and Jincai Huang. 2019. Crime-gan: A context-based sequence generative network for crime forecasting with adversarial loss. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 1460–1469.

[62] Guangyin Jin, Qi Wang, Cunchao Zhu, Yanghe Feng, Jincai Huang, and Jiangping Zhou. 2020. Addressing Crime Situation Forecasting Task with Temporal Graph Convolutional Neural Network Approach. In *2020 12th International Conference on Measuring Technology and Mechatronics Automation (ICMTMA)*. IEEE, 474–478.

[63] Cristina Kadar, Rudolf Maculan, and Stefan Feuerriegel. 2019. Public decision support for low population density areas: An imbalance-aware hyper-ensemble for spatio-temporal crime prediction. *Decision Support Systems* 119 (2019), 107–117.

[64] Nathan Kallus. 2014. Predicting crowd behavior with big public data. In *Proceedings of the 23rd International Conference on World Wide Web*. 625–630.

[65] Hyeon-Woo Kang and Hang-Bong Kang. 2017. Prediction of crime occurrence from multi-modal data using deep learning. *PloS one* 12, 4 (2017), e0176244.

[66] Daniel Kaufmann, Aart Kraay, and Massimo Mastruzzi. 2011. The Worldwide Governance Indicators: Methodology and Analytical Issues1. *Hague journal on the rule of law* 3, 2 (2011), 220–246.

[67] Pranav Kedia. 2016. Crime mapping and analysis using GIS. *International Institute of Information Technology* 1, 1 (2016), 1–15.

[68] Keivan Kianmehr and Reda Alhajj. 2006. Crime Hot-spots prediction using support vector machine. In *IEEE International Conference on Computer Systems and Applications, 2006*. IEEE Computer Society, 952–959.

[69] Keivan Kianmehr and Reda Alhajj. 2008. Effectiveness of support vector machine for crime hot-spots prediction. *Applied Artificial Intelligence* 22, 5 (2008), 433–458.

[70] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).

[71] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*.

[72] Gizem Korkmaz, Jose Cadena, Chris J Kuhlman, Achla Marathe, Anil Vullikanti, and Naren Ramakrishnan. 2015. Combining heterogeneous data sources for civil unrest forecasting. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. 258–265.

[73] Ourania Kounadi, Alina Ristea, Adelson Araujo, and Michael Leitner. 2020. A systematic review on spatial crime forecasting. *Crime science* 9 (2020), 1–22.

[74] Gary LaFree and Laura Dugan. 2007. Introducing the global terrorism database. *Terrorism and political violence* 19, 2 (2007), 181–204.

[75] Yasmine Lamari, Bartol Freskura, Anass Abdessamad, Sarah Eichberg, and Simon de Bonviller. 2020. Predicting Spatial Crime Occurrences through an Efficient Ensemble-Learning Model. *ISPRS International Journal of Geo-Information* 9, 11 (2020), 645.

[76] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. 2016. Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648* (2016).

[77] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. 2017. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 156–165.

[78] Kalev Leetaru and Philip A Schrodt. 2013. Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, Vol. 2. Citeseer, 1–49.

[79] Jia Li, Xiaowei Jia, Haoyu Yang, Vipin Kumar, Michael Steinbach, and Gyorgy Simon. 2020. Teaching deep learning causal effects improves predictive performance. *arXiv preprint arXiv:2011.05466* (2020).

[80] Renjie Liao, Xueyao Wang, Lun Li, and Zengchang Qin. 2010. A novel serial crime prediction model based on Bayesian learning theory. In *2010 international conference on machine learning and cybernetics*, Vol. 4. IEEE, 1757–1762.

[81] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025* (2015).

[82] Nick Malleson and Martin A Andresen. 2016. Exploring the impact of ambient population measures on London crime hotspots. *Journal of Criminal Justice* 46 (2016), 52–63.

[83] Pedro Manrique, Hong Qi, Ana Morgenstern, Nicolas Velasquez, Tsai-Ching Lu, and Neil Johnson. 2013. Context matters: Improving the uses of big data for forecasting civil unrest: Emerging phenomena and big data. In *2013 IEEE International Conference on Intelligence and Security Informatics*. IEEE, 169–172.

[84] Charles McClelland. 1978. World event/interaction survey, 1966-1978. *WEIS Codebook ICPSR* 5211 (1978).

[85] Lawrence McClendon and Natarajan Meghanathan. 2015. Using machine learning algorithms to analyze crime data. *Machine Learning and Applications: An International Journal (MLAIJ)* 2, 1 (2015), 1–12.

[86] Lu Meng and Rohini K Srihari. [n.d.]. Leveraging Heterogeneous Data Sources for Civil Unrest Prediction. ([n. d.]).

[87] Michail Misyrlis, Chung Ming Cheung, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2017. Spatio-temporal modeling of criminal activity. In *Proceedings of the 2nd International Workshop on Social Sensing*. 3–8.

[88] George O Mohler, Martin B Short, P Jeffrey Brantingham, Frederic Paik Schoenberg, and George E Tita. 2011. Self-exciting point process modeling of crime. *J. Amer. Statist. Assoc.* 106, 493 (2011), 100–108.

[89] Lenin Mookiah, William Eberle, and Ambareen Siraj. 2015. Survey of crime analysis and prediction. In *The Twenty-Eighth International Flairs Conference*.

[90] Sathappan Muthiah, Patrick Butler, Rupinder Paul Khandpur, Parang Saraf, Nathan Self, Alla Rozovskaya, Liang Zhao, Jose Cadena, Chang-Tien Lu, Anil Vullikanti, et al. 2016. Embers at 4 years: Experiences operating an open source indicators forecasting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 205–214.

[91] Peter F Nardulli, Kalev H Leetaru, and Matthew Hayes. 2011. Event data, civil unrest and the social, political and economic event database (SPEED) project: post World War II trends in political protests and violence. In *Annual meeting of the International Studies Association, Quebec, Canada*.

[92] Yue Ning, Sathappan Muthiah, Huzefa Rangwala, and Naren Ramakrishnan. 2016. Modeling precursors for event forecasting via nested multi-instance learning. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1095–1104.

[93] 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 the 2018 SIAM International Conference on Data Mining*. SIAM, 99–107.

[94] Andreas M Olligschlaeger. 1997. Artificial neural networks and crime mapping. *Crime mapping and crime prevention* 1 (1997), 313.

[95] Nathan H Parrish, Anna L Buczak, Jared T Zook, James P Howard, Brian J Ellison, and Benjamin D Baugher. 2018. Crystal Cube: Multidisciplinary Approach to Disruptive Events Prediction. In *International Conference on Applied Human Factors and Ergonomics*. Springer, 571–581.

[96] Walt L Perry. 2013. *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation.

[97] Lawrence Phillips, Chase Dowling, Kyle Shaffer, Nathan Hodas, and Svitlana Volkova. 2017. Using social media to predict the future: a systematic literature review. *arXiv preprint arXiv:1706.06134* (2017).

[98] Fengcai Qiao, Pei Li, Xin Zhang, Zhaoyun Ding, Jiajun Cheng, and Hui Wang. 2017. Predicting social unrest events with hidden Markov models using GDELT. *Discrete Dynamics in Nature and Society* 2017 (2017).

[99] Fengcai Qiao, Xin Zhang, and Jinsheng Deng. 2020. Learning Evolutionary Stages with Hidden Semi-Markov Model for Predicting Social Unrest Events. *Discrete Dynamics in Nature and Society* 2020 (2020).

[100] Kira Radinsky and Eric Horvitz. 2013. Mining the web to predict future events. In *Proceedings of the sixth ACM international conference on Web search and data mining*. 255–264.

[101] Clionadh Raleigh and Caitriona Dowd. 2015. Armed conflict location and event data project (ACLED) codebook. *Find this resource* (2015).

[102] Naren Ramakrishnan, Patrick Butler, Sathappan Muthiah, Nathan Self, Rupinder Khandpur, Parang Saraf, Wei Wang, Jose Cadena, Anil Vullikanti, Gizem Korkmaz, et al. 2014. 'Beating the news' with EMBERS: Forecasting civil unrest using open source indicators. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1799–1808.

[103] Yeasir Rayhan and Tanzima Hashem. 2020. AIST: An Interpretable Attention-based Deep Learning Model for Crime Prediction. *arXiv preprint arXiv:2012.08713* (2020).

[104] Sungmin Rhee, Seokjun Seo, and Sun Kim. 2017. Hybrid approach of relation network and localized graph convolutional filtering for breast cancer subtype classification. *arXiv preprint arXiv:1711.05859* (2017).

[105] Kosta Ristovski, Vladan Radosavljevic, Slobodan Vucetic, and Zoran Obradovic. 2013. Continuous conditional random fields for efficient regression in large fully connected graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 27.

[106] Shakila Khan Rumi, Ke Deng, and Flora Dilys Salim. 2018. Crime event prediction with dynamic features. *EPJ Data Science* 7, 1 (2018), 43.

[107] Shakila Khan Rumi, Phillip Luong, and Flora D Salim. 2019. Crime rate prediction with region risk and movement patterns. *CoRR abs/1908.02570* (2019).

[108] Anneleen Rummens, Wim Hardyns, and Lieven Pauwels. 2017. The use of predictive analysis in spatiotemporal crime forecasting: Building and testing a model in an urban context. *Applied geography* 86 (2017), 255–261.

[109] Idean Salehyan, Cullen S Hendrix, Jesse Hamner, Christina Case, Christopher Linebarger, Emily Stull, and Jennifer Williams. 2012. Social conflict in Africa: A new database. *International Interactions* 38, 4 (2012), 503–511.

[110] Ginger Saltos and Mihaela Cocea. 2017. An exploration of crime prediction using data mining on open data. *International Journal of Information Technology & Decision Making* 16, 05 (2017), 1155–1181.

[111] Parang Saraf and Naren Ramakrishnan. 2016. EMBERS autogsr: Automated coding of civil unrest events. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 599–608.

[112] Gerald Schneider and Margit Bussmann. 2013. Accounting for the dynamics of one-sided violence: Introducing KOSVED. *Journal of Peace Research* 50, 5 (2013), 635–644.

[113] Gerald Schneider, Margit Bussmann, and Constantin Ruhe. 2012. The dynamics of mass killings: Testing time-series models of one-sided violence in the Bosnian civil war. *International Interactions* 38, 4 (2012), 443–461.

[114] Philip A Schrodt, Shannon G Davis, and Judith L Weddle. 1994. Political science: KEDS—a program for the machine coding of event data. *Social Science Computer Review* 12, 4 (1994), 561–587.

[115] Yantao Shen, Hongsheng Li, Shuai Yi, Dapeng Chen, and Xiaogang Wang. 2018. Person re-identification with deep similarity-guided graph neural network. In *Proceedings of the European conference on computer vision (ECCV)*. 486–504.

[116] Gary L Shoesmith. 2013. Space–time autoregressive models and forecasting national, regional and state crime rates. *International journal of forecasting* 29, 1 (2013), 191–201.

[117] Emmanuel M Smith, Jim Smith, Phil Legg, and Simon Francis. 2017. Predicting the occurrence of world news events using recurrent neural networks and auto-regressive moving average models. In *UK Workshop on Computational Intelligence*. Springer, 191–202.

[118] Panagiotis Stalidis, Theodoros Semertzidis, and Petros Daras. 2018. Examining deep learning architectures for crime classification and prediction. *arXiv preprint arXiv:1812.00602* (2018).

[119] Alexander Stec and Diego Klabjan. 2018. Forecasting crime with deep learning. *arXiv preprint arXiv:1806.01486* (2018).

[120] Ralph Sundberg and Erik Melander. 2013. Introducing the UCDP georeferenced event dataset. *Journal of Peace Research* 50, 4 (2013), 523–532.

[121] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.

[122] Mohammad A Tayebi, Martin Ester, Uwe Glässer, and Patricia L Brantingham. 2014. Crimetracer: Activity space based crime location prediction. In *2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014)*. IEEE, 472–480.

[123] Henrik Urdal and Kristian Hoelscher. 2012. Explaining urban social disorder and violence: An empirical study of event data from Asian and sub-Saharan African cities. *International Interactions* 38, 4 (2012), 512–528.

[124] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha Talukdar. 2019. Composition-based multi-relational graph convolutional networks. *arXiv preprint arXiv:1911.03082* (2019).

[125] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.

[126] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* 11, 12 (2010).

[127] Bao Wang, Xiyang Luo, Fangbo Zhang, Baichuan Yuan, Andrea L Bertozzi, and P Jeffrey Brantingham. 2018. Graph-based deep modeling and real time forecasting of sparse spatio-temporal data. *arXiv preprint arXiv:1804.00684* (2018).

[128] Bao Wang, Penghang Yin, Andrea Louise Bertozzi, P Jeffrey Brantingham, Stanley Joel Osher, and Jack Xin. 2019. Deep learning for real-time crime forecasting and its ternarization. *Chinese Annals of Mathematics, Series B* 40, 6 (2019), 949–966.

[129] Hongjian Wang, Daniel Kifer, Corina Graif, and Zhenhui Li. 2016. Crime rate inference with big data. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 635–644.

[130] Hongjian Wang, Huaxiu Yao, Daniel Kifer, Corina Graif, and Zhenhui Li. 2017. Non-stationary model for crime rate inference using modern urban data. *IEEE transactions on big data* 5, 2 (2017), 180–194.

[131] Mingjun Wang and Matthew S Gerber. 2015. Using twitter for next-place prediction, with an application to crime prediction. In *2015 IEEE Symposium Series on Computational Intelligence*. IEEE, 941–948.

[132] 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. *Knowledge-Based Systems* 189 (2020), 105120.

[133] Wei Wang, Ryan Kennedy, David Lazer, and Naren Ramakrishnan. 2016. Growing pains for global monitoring of societal events. *Science* 353, 6307 (2016), 1502–1503. https://doi.org/10.1126/science.aaf6758 arXiv:https://science.sciencemag.org/content/353/6307/1502.full.pdf

[134] Xiaofeng Wang, Donald E Brown, and Matthew S Gerber. 2012. Spatio-temporal modeling of criminal incidents using geographic, demographic, and Twitter-derived information. In *2012 IEEE International Conference on Intelligence and Security Informatics*. IEEE, 36–41.

[135] Xiuling Wang, Hao Chen, Zhoujun Li, and Zhonghua Zhao. 2018. Unrest news amount prediction with context-aware attention lstm. In *Pacific Rim International Conference on Artificial Intelligence*. Springer, 369–377.

[136] Yaqian Wang, Liang Ge, Siyu Li, and Feng Chang. 2020. Deep Temporal Multi-Graph Convolutional Network for Crime Prediction. In *International Conference on Conceptual Modeling*. Springer, 525–538.

[137] Michael D Ward, Andreas Beger, Josh Cutler, Matthew Dickenson, Cassy Dorff, and Ben Radford. 2013. Comparing GDELT and ICEWS event data. *Analysis* 21, 1 (2013), 267–297.

[138] Zbigniew M Wawrzyniak, Stanisław Jankowski, Eliza Szczechla, Zbigniew Szymański, Radosław Pytlak, Paweł Michalak, and Grzegorz Borowik. 2018. Data-driven models in machine learning for crime prediction. In *2018 26th International Conference on Systems Engineering (ICSEng)*. IEEE, 1–8.

[139] Yue Wei, Weichao Liang, Youquan Wang, and Jie Cao. 2020. CrimeSTC: A Deep Spatial-Temporal-Categorical Network for Citywide Crime Prediction. In *2020 The 3rd International Conference on Computational Intelligence and Intelligent Systems*. 75–79.

[140] Nils B Weidmann and Michael D Ward. 2010. Predicting conflict in space and time. *Journal of Conflict Resolution* 54, 6 (2010), 883–901.

[141] Gadi Wolfsfeld, Elad Segev, and Tamir Sheafer. 2013. Social media and the Arab Spring: Politics comes first. *The International Journal of Press/Politics* 18, 2 (2013), 115–137.

[142] Congyu Wu and Matthew S Gerber. 2017. Forecasting civil unrest using social media and protest participation theory. *IEEE Transactions on Computational Social Systems* 5, 1 (2017), 82–94.

[143] Lianghao Xia, Chao Huang, Yong Xu, Peng Dai, Liefeng Bo, Xiyue Zhang, and Tianyi Chen. [n.d.]. Spatial-Temporal Sequential Hypergraph Network for Crime Prediction with Dynamic Multiplex Relation Learning. ([n. d.]).

[144] Dingqi Yang, Terence Heaney, Alberto Tonon, Leye Wang, and Philippe Cudré-Mauroux. 2018. CrimeTelescope: crime hotspot prediction based on urban and social media data fusion. *World Wide Web* 21, 5 (2018), 1323–1347.

[145] Yiming Yang, Tom Pierce, and Jaime Carbonell. 1998. A study of retrospective and on-line event detection. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*. 28–36.

[146] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 7370–7377.

[147] Xinyue Ye, Lian Duan, and Qiong Peng. 2021. Spatiotemporal Prediction of Theft Risk with Deep Inception-Residual Networks. *Smart Cities* 4, 1 (2021), 204–216.

[148] Fei Yi, Zhiwen Yu, Fuzhen Zhuang, and Bin Guo. 2019. Neural Network based Continuous Conditional Random Field for Fine-grained Crime Prediction.. In *IJCAI*. 4157–4163.

[149] Fei Yi, Zhiwen Yu, Fuzhen Zhuang, Xiao Zhang, and Hui Xiong. 2018. An integrated model for crime prediction using temporal and spatial factors. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 1386–1391.

[150] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems* 32 (2019), 9240.

[151] James E Yonamine. 2013. Predicting future levels of violence in afghanistan districts using gdelt. *Unpublished manuscript* (2013).

[152] Chung-Hsien Yu, Wei Ding, Melissa Morabito, and Ping Chen. 2015. Hierarchical spatio-temporal pattern discovery and predictive modeling. *IEEE Transactions on Knowledge and Data Engineering* 28, 4 (2015), 979–993.

[153] Chung-Hsien Yu, Max W Ward, Melissa Morabito, and Wei Ding. 2011. Crime forecasting using data mining techniques. In *2011 IEEE 11th international conference on data mining workshops*. IEEE, 779–786.

[154] Junbo Zhang, Yu Zheng, and Dekang Qi. 2017. Deep spatio-temporal residual networks for citywide crowd flows prediction. In *Thirty-first AAAI conference on artificial intelligence*.

[155] Qiang Zhang, Pingmei Yuan, Qiyun Zhou, and Zhiming Yang. 2016. Mixed spatial-temporal characteristics based crime hot spots prediction. In *2016 IEEE 20th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. IEEE, 97–101.

[156] Yihong Zhang, Panote Siriaraya, Yukiko Kawai, and Adam Jatowt. 2019. Analysis of street crime predictors in web open data. *Journal of Intelligent Information Systems* (2019), 1–25.

[157] Liang Zhao. 2021. Event Prediction in the Big Data Era: A Systematic Survey. *ACM Computing Surveys (CSUR)* 54, 5 (2021), 1–37.

[158] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2015. Spatiotemporal event forecasting in social media. In *Proceedings of the 2015 SIAM international conference on data mining*. SIAM, 963–971.

[159] Liang Zhao, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2016. Multi-resolution spatial event forecasting in social media. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 689–698.

[160] 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 the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1503–1512.

[161] Liang Zhao, Qian Sun, Jieping Ye, Feng Chen, Chang-Tien Lu, and Naren Ramakrishnan. 2017. Feature constrained multi-task learning models for spatiotemporal event forecasting. *IEEE Transactions on Knowledge and Data Engineering* 29, 5 (2017), 1059–1072.

[162] Liang Zhao, Junxiang Wang, and Xiaojie Guo. 2018. Distant-supervision of heterogeneous multitask learning for social event forecasting with multilingual indicators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

[163] 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 the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2085–2094.

[164] Xiangyu Zhao and Jiliang Tang. 2017. Modeling temporal-spatial correlations for crime prediction. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 497–506.

[165] Yong Zhuang, Matthew Almeida, Melissa Morabito, and Wei Ding. 2017. Crime hot spot forecasting: A recurrent model with spatial and temporal information. In *2017 IEEE International Conference on Big Knowledge (ICBK)*. IEEE, 143–150.