Traffic Accident Detection with Spatiotemporal Impact Measurement

Mingxuan Yue[1], Liyue Fan[2], and Cyrus Shahabi[1]

[1] University of Southern California, Los Angeles, CA
{mingxuay, shahabi}@usc.edu
[2] University at Albany SUNY, Albany, NY
liyuefan@albany.edu

Abstract. Traffic incidents continue to cause a significant loss in deaths, injuries, and property damages. Reported traffic accident data contains a considerable amount of human errors, hindering the studies on traffic accidents. Several approaches have been developed to detect accidents using traffic data in real time. However, those approaches do not consider the spatiotemporal patterns inherent in traffic data, resulting in high false alarm rates. In this paper, we study the problem of traffic accident detection by considering multiple traffic speed time series collected from road network sensors. To capture the spatiotemporal impact of traffic accidents to upstream locations, we adopt Impact Interval Grouping (IIG), which compares real-time traffic speed with historical data, and generates *impact intervals* to determine the presence of accidents. Furthermore, we take a multivariate time series classification approach and extract three novel features to measure the *severity* of traffic accidents. We use real-world traffic speed and accident datasets in our empirical evaluation, and our solutions outperform state-of-the-art approaches in multivariate time series classification.

Keywords: Traffic Accident, Multivariate Time Series Classification.

1 Introduction

Traffic accidents have been an essential concern in our society. In 2015 the total number of motor-vehicle deaths was 38,300, and the estimated cost of deaths, injuries, and property damage reached $412.1 billion [4]. Studying traffic accidents would help us understand the causes and potentially reduce the damage of such events. However, accident data is usually collected from various state and local agencies: these reports often contain duplicates, missing data, and/or inaccurate information as they are based on victim/witness estimates [15]. Therefore, we are in need of accurate characterization and detection of traffic accidents, which can be achieved utilizing the data from ubiquitous traffic sensors.

In the past decades, a plethora of research studied automatic incident detection (AID) [7, 8, 10, 16, 14, 9]. However, these approaches are prone to false alarms [6] due to improper calibration on insufficient and noisy data. On the other hand, video recordings by traffic cameras [9] and traffic flow data collected from probe cars [16] have also been used. But such data is expensive to acquire and does not provide a thorough coverage of the entire road networks. Traffic data, e.g., speed, has specific spatial and temporal patterns, which are unfortunately neglected in AID research. Such spatial and temporal pattern could be helpful in differentiating accidents from fluctuations that are often observed in ordinary traffic and attribute to false alarms. However, it is challenging to model and recognize such spatial and temporal patterns. Furthermore, real-world data could be noisy and has missing data or outliers. Hence it is also important to make the detection robust and not sensitive to such noise.

In this paper, we address the aforementioned challenges by considering traffic speed time series data collected from multiple sensors close to the accident location. We first adopt the Impact Interval Grouping (IIG) algorithm [15] to detect traffic accidents by recognizing the spatiotemporal patterns in traffic speed data. In addition, we propose a multivariate time series (MTS) classification technique, and define three novel features that measure the severity of traffic accidents. We propose two versions, i.e., Severity-I and Severity-T, to balance the algorithm's sensitivity to signals and noise in the traffic data. For evaluation, we utilize real-world traffic speed and accident dataset collected from Los Angeles and our proposed approaches are shown to outperform the state-of-the-art MTS classification methods. We conclude that our proposed severity features can sufficiently characterize the spatiotemporal impact of traffic accidents, and can be used for efficient and accurate traffic accident detection.

2 Related Work

Automatic Incident Detection In the field of transportation, automatic incident detection (AID) has been an active area of research since the end of the last century. The classical and widely used approach is California Algorithms [7, 8]. The basic idea is to use three levels of rules to raise an alarm, based on the difference between the upstream and downstream sensors. Other researchers also tried recognizing vehicles' behaviors from the camera videos. For example, Sadek et al.[9] estimated optical flow from the videos and constructed histograms of flow gradient to build a classifier. However, classical algorithms are not satisfactory due to high missing rate and false alarm rate. Moreover, the video data does not offer a sufficient coverage of most road networks. Yuan et al. [14] utilized SVM to classify accidents by traffic metrics but they do not consider the temporal and spatial correlations. We believe by using time series data from multiple sensors and utilizing multivariate time series classification approaches, the accidents could be modeled and recognized more accurately.

Multi-dimensional Dynamic Time Warping To the best of our knowledge, most of the current study on Multivariate Time Series Classification (MTSC) are extended from the Univariate Time Series Classification (TSC) methods. For TSC, the kNN approach is typically used with certain distance measures, such as the Dynamic Time Warping (DTW) distance. kNN-DTW is proven to be a reliable approach in TSC[2]. Moreover, for DTW distance in multivariate case (MD-DTW), researchers usually calculate DTW distance between two MTS by summing over individual DTW distances of each variate[1], or formulating values in multiple variates at the same timestamp as a vector and computing p-norm of vectors[11]. A kNN classifier based on 2-norm MD-DTW is used as a baseline in our empirical evaluation in Section 5.

Feature based MTS classification The feature-based representation of time series can be efficient in classification and can discover more latent patterns if the right set of features is extracted. Therefore, researchers developed various feature extraction approaches for time series. Yang et al. [13] proposed a feature selection technique for MTS based on Common Principal Component Analysis. Wang et al. [12] extracted statistical features such as trend and periodicity from MTS for clustering. Fulcher et al. [3] developed a comprehensive feature extraction approach which constructs features including statistics, correlations, and non-linear model fits, etc. Over 7000 features are extracted from the time series and forward selection is utilized to generate feature vectors for classification. This Feature Selection (FS) approach is also used as a baseline in Section 5.

3 Preliminaries

In this paper, the time series represents the sequence of traffic speed detected by a sensor every minute. Here, a sensor s is a loop detector located on major roads recording traffic metrics.

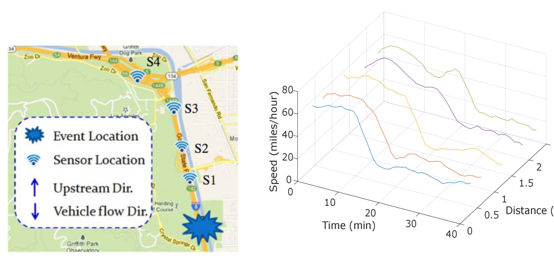

(a) Sample Accident [5] (b) Traffic Speed MTS

Fig. 1: Multivariate Time Series

Since time series collected from a single sensor is noisy and sensitive to abnormal behaviors of vehicles, it is important to study multiple sensors at different locations. Thus, we represent the time series collected from multiple sensors close to an accident as a single MTS and label it as "positive". Accordingly, we also generate MTS of regular traffic from random locations when there is no accidents as the negative class. Specifically, the MTS is generated from time series of each *upstream sensor*, as defined below.

Definition 1 (Multivariate Time Series (MTS)). *A Multivariate Time Series is a set of time series:* $X = \{x_1, x_2, ..., x_K\}$*. In this paper, a time series* x_k *is a sequence of speed, indexed by time:* $x_k = \{x_k(1), x_k(2), ..., x_k(n)\}$*. The* k_{th} *time series* x_k *is collected by the* k_{th} *sensor* s_k*.*

In the definition, n denotes the window size of the MTS, and K is the number of upstream sensors. The upstream sensors of a location l on a road are defined as a sequence of nearest sensors $S = \{s_1, s_2, ..., s_K\}$ ordered by the distance to l with the same direction as l, as shown in Figure 1a. Figure 1b presents the MTS of upstream sensors relative to the accident location/time. The time series are generated at different distances in different colors. Hence we define the problem of traffic incident detection based on MTS classification as follows.

Problem 1 (MTS based Traffic Accident Classification Problem (MTS-TACP)). Given a multivariate time series $X = \{x_1, x_2, ..., x_K\}$, in which each time series x_k is collected from the k_{th} upstream sensor s_k of location l, during the same time window $[t_1, t_n]$, the goal is to identify whether an accident happened at location l from t_1 to t_n, classify the MTS to be accident or normal instance.

4 Traffic Accident Classification with Traffic MTS

4.1 A Discrete Unsupervised Solution: IIG

In our previous study[15], we proposed an approach to realign the start of existing accidents based on the upstream traffic speed around the reported time. Though it has different assumptions and problems, the idea of identifying the propagation pattern in MTS can be extended to detect an accident.

First, *impact interval* in Definition 2 is defined to reduce the dimensionality while keeping the accident patterns. Basically, a time series will be compared to its historical average and the intervals below a given ratio are extracted. Such definition is derived from real-world intuitions. Essentially, accidents usually result in visible and consecutive speed drops at upstream sensors. However, rush hours and narrow roads could also have regular speed drops. We assume the accidents cause unusual speed drops, i.e. the speed should be observed at a lower value than regular speed if an accident happened. Therefore the real-time speed should be compared to the historical speed to quantify the impact[5]. Here historical speed is calculated using the average of speed at the same location and same time[5]. So the unusual speed drop is modeled in a discrete way by extracting impact intervals. Via such discretization, we can convert the complex time series into a concise formulation which is easier to model as in Figure 2a. The following paragraphs will describe the procedure of detecting accidents by the modified IIG method. The existence of accident in an MTS will be decided by the identification of propagated impact intervals via following steps.

Definition 2 (Impact Interval). *An impact interval is a tuple (t_s, t_e), s.t. $\forall t, t_s \leq t \leq t_e, \frac{|x(t) - \bar{x}(t)|}{\bar{x}(t)} \geq \theta$. Here $x(t)$ denotes the real-time speed at time t, and $\bar{x}(t)$ denotes the historical average speed of the same sensor, at time t. θ is a tuning parameter determining how strict the impact is measured.*

Discretization First, for each upstream sensor s_k, we calculate a set of impact intervals $I_k = \{(t_{1s}, t_{1e}), (t_{2s}, t_{2e}), ...\}$ as described above. For example, in Figure 2a, the time series (red line) will be compared to the historical time series (solid blue line). Then the impact intervals (solid bold segments) are generated by applying a threshold θ to the historical time series (dashed blue line). In this way, the MTS will be discretized into K sets of impact intervals $\{I_1, I_2, ...I_K\}$.

Smoothing Since fluctuations and outliers could generate useless and noisy impact intervals, smoothing and cleansing should be applied before and after discretization. We smooth the time series before creating impact intervals by moving average with windows size 5 minutes for both real-time and historical data. After discretization, we concatenate adjacent impact intervals with distance less than 2 minutes, similar to one-dimensional clustering. Because close impact intervals may result from the same event. Then isolated intervals not concatenated to any other intervals with a small length, i.e., 2 minutes will be eliminated. Those impact intervals are actually noises, i.e. speed fluctuations. After that, we have a clean formation of impact intervals from all sensors.

Grouping Intuitively, upstream sensors are affected by accident in the spatial order based on the traffic flow: a further upstream sensor is affected only after other upstream sensors closer to the accident. Temporally, the impacts of the same accident, observed at different upstream sensors, should be relatively close in time. So impact intervals can be grouped by their spatial and temporal distances, using heuristics such as maximum overlap and nearest center as in Figure 2b. Specifically, each impact interval in the closest sensor will form a group. Then other impact intervals in the adjacent sensors will be iteratively added to the groups they belong to using maximum overlap or nearest center. The grouping procedure finishes until all impact intervals are visited. Groups with too few intervals will be filtered out. Then if a group of impact intervals exists after this IIG procedure, we classify the MTS as an accident. Otherwise, we label the MTS as a normal instance.

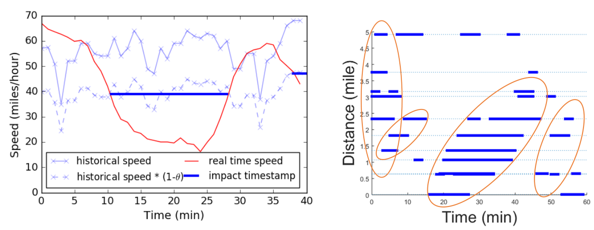

(a) Generate Impact Intervals (b) Impact Interval Groups

Fig. 2: Procedure of IIG

4.2 Severity Features based Solution (Severity-I)

Admittedly, the IIG approach models the spatiotemporal propagation pattern intuitively using impact interval groups. However, it could be too strict to identify an accident. A slight turbulence or malfunction could disqualify the entire impact interval group. Thus, we need more robust characteristics of accidents. In this section, three types of features are defined and extracted from a traffic speed MTS based on empirical observations. Then the features can be adopted by various classifiers to detect accidents.

Dropping Severity λ We first consider drops in traffic speed. Different from the IIG approach, the extent of speed drop should be essential in detecting accidents as opposed to the binary comparison to the threshold θ. Hence, we have the following observation: *An accident will cause a speed drop in upstream sensors to a certain extent.* A larger extent of speed drop can provide more confidence of an accident. For example, in Figure 2a, the red line depicts the time series of speed from the nearest sensor to an accident. As shown in the figure, the speed drops from 60 miles/hour to around 20 miles/hour, which can indicate a severe accident. The dropping severity can be estimated by the ratio of speed change from historical data to real-time data. As described in IIG approach, the comparison to historical speed is necessary to eliminate the rush hour or other periodical effects. Given an MTS of traffic $X = \{x_1, x_2...x_k\}$, the historical average of speed is denoted by $\overline{X} = \{\overline{x}_1, \overline{x}_2...\overline{x}_k\}$. We define the measurement of dropping severity in the following equation.

$$\lambda_{max} = \max_{i,k}(1 - x_k(i)/\overline{x}_k(i)); \lambda_{avg} = \underset{i,k}{avg}(1 - x_k(i)/\overline{x}_k(i)) \tag{1}$$

In the equation, we propose two options for this measurement for comprehensiveness. Dropping severity is first measured as the maximum speed drop ratio λ_{max}, which reflects the worst impact to all sensors of the accident. The other option is the average ratio λ_{avg}, aggregating the overall speed change in all sensors.

Figure 3a depicts the intermediate step in extracting dropping severity. The red line depicts the real-time speed reported by a certain sensor near an accident. The solid blue line represents the historical speed. Denoted by the red dot line, the maximum distance between real-time speed and historical speed is used for calculating λ_{max}, and such computation is applied to all upstream sensors.

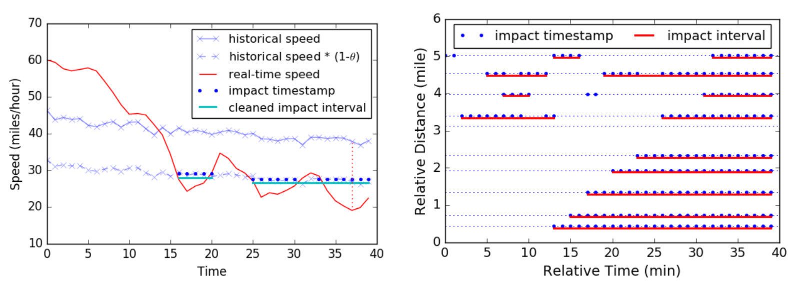

(a) Extract λ and τ

(b) Extract τ and σ

Fig. 3: Severity Features

Lasting Severity τ Not all accidents have grave speed drops and moreover normal traffic also has occasional drops, i.e., traffic MTS with moderate drops can be either normal traffic or accidents. To better differentiate these two conditions, we introduce the temporal criteria, lasting severity. For example, in Figure 2a, the speed drops at 8 minutes and resumes at 28 minutes, staying at a low value for 20 minutes. Then we can conclude such speed drops are not caused by normal traffic fluctuations. *After an accident happens, the drop of speed will last for a certain time.* Significant dropping time can be the evidence of an accident. Impact interval is used to measure lasting severity because the discretization provides an easy extraction of temporal patterns. A list of impact intervals I_k (horizontal cyan segments in Figure 3a) is generated. $|x_k|$ denotes the length of time series x_k. Then lasting severity is measured by the following formulation.

$$\begin{aligned} \tau_{max} &= \max_{i,k}(I_k(i)[1] - I_k(i)[0])/|x_k| \\ \tau_{avg} &= \underset{k}{avg}(\max_{i}(I_k(i)[1] - I_k(i)[0]))/|x_k| \end{aligned} \tag{2}$$

The lasting severity is provided with two options as well. We assume each impact interval after smoothing and concatenation is individually impacted by a single event, as supposed in IIG. So the maximum length of impact intervals should be the upper bound of all the events(the longer cyan segment in Figure 3a). Thus τ_{max} is the maximum of these maximum lengths in all sensors which indicates the longest impact. τ_{avg} is the average of the maximum lengths indicating the overall affected time in all sensors.

In addition, We also define a relaxed definition of τ, to overcome fluctuations in traffic speed which may prevent the formulation of impact intervals. Rather than extracting impact intervals, we compute impact timestamps(dark blue dots in Figure 3a), which are the set of time index at which the relative speed drop is below θ. Then the sizes of the impact timestamps sets are used instead of the lengths of impact intervals in calculating τ'_{max} and τ'_{avg}.

Distant Severity σ We also believe an accident usually influences a succession of cars rather than a single one, i.e. *an accident will affect a certain distance in the upstream traffic.* The longer stream of cars are affected, the traffic MTS is more likely to reflect an accident. To be consistent with lasting severity, the distant severity is measured based on the existence of impact intervals, as described in the following formulation. Here d_k denotes the distance of sensor s_k.

$$\begin{aligned} \sigma_{cons} &= d_k/d_K, k = arg\max_{k}\{I_1 \ to \ I_k \neq \emptyset\} \\ \sigma_{disc} &= d_k/d_K, k = arg\max_{k}\{I_k \neq \emptyset\} \end{aligned} \tag{3}$$

We provide two options for distant severity measurement to overcome the inconsistent behaviors among individual sensors. First we assume the impact at different locations be consecutive since the first sensor cannot affect the third sensor without influencing the second one. So σ_{cons} computes the furthest consecutively impacted distance. Optionally, σ_{disc} lifts such restriction to the furthest discrete impact location in case of noise or missing data.

Figure 3b shows the extraction of distant severity. The solid red lines depict impact intervals. The y-axis represents the distance to the detecting location l.

The distant severity is measured by the distance of furthest impacted sensor. The difference between σ_{cons} and σ_{disc} is the absence of impact interval in the 6th sensor s_6 (the bare-fine dotted blue line). σ_{cons} requires the impacts be consecutive so that it can separate events at different locations. σ_{disc} allows skipping some sensors to give more tolerance to malfunction of sensors. Similar to lasting severity, we also defined a relaxed definition of σ, to capture a loose distant severity. The relaxed impacted distant severity σ'_{cons} and σ'_{disc} will be computed based on the existence of the impact timestamps at different locations instead of the impact intervals. By replacing τ and σ with their relaxed version τ' and σ', we can derive a relaxed version of Severity-I, namely *Severity-T*.

Severity based Classification The advantage of measuring different types of severity (extent, time and space) is the capability of capturing enough variety of accidents. For example, an accident with low λ and large τ and σ could be caused by an emergency vehicle parked on the shoulder which does not reduce the traffic speed significantly but may last for 30 minutes and affects many upstream vehicles. An accident with low τ and large λ and σ may be a minor accident with no injury. So upstream drivers all brake and the speed drops a lot, but traffic is resumed soon after the accident is clear. Therefore, the combination of the three severity could describe accidents in a wide variety. After generating the three severity features from a traffic speed MTS, we can utilize various classifiers, e.g. Logistic Regression, SVM, Gradient Boosting Decision Tree (GBDT), etc. for classification. GBDT is adopted as the default classifier in Severity-I.

The complexity of extracting a severity-I feature is quite low. With n as the window length and K as the number of upstream sensors, it takes $O(Kn)$ time to traverse all historical and real-time speed to get Dropping Severity λ. Moreover, the extraction of impact intervals of one time series takes $O(n)$ time. So Lasting Severity τ and Distant Severity σ extractions also cost $O(Kn)$ time.

5 Empirical Evaluation

5.1 Experiments Settings

System Environment For all experiments, we implemented the algorithms in Python to ensure a fair comparison. All the experiments were conducted on a 64-bit Windows machine with a 2.60 GHz CPU and 8.00 GB memory.

Data Set In this paper, we use the accident reports in June 2012 as ground truths, which were reported by California Highway Patrol. Traffic metrics including speed, volume and occupancy, collected from more than 4,000 sensors during the same period are retrieved as real-time data. The historical average of traffic is generated from the sensor dataset during March to May 2012. 70% of real-time data is sampled as training data, and the rest 30% is testing data.

Baseline Approaches As baseline approaches, we implemented kNN-MD-DTW and the FS approach[2, 3]. In the implementation of kNN-MD-DTW, we set $k = 3$ and set the size of warping window in MD-DTW to 5 minutes. As for the extended FS approach using HTCSA[3], 4 features are selected from more than 7000 features. We generated 7000 features only using a 200 size sample because the extraction takes significant computation. The 4 feature extraction methods are SY_KPSStest_0_10.lagminstat (KPSS stationarity test), SY_LocalDistributions_5_each.meandiv (minimum divergence between two segments' distribution), MF_CompareAR_1_10_all.whereen4 (order of AR with little error), and CO_tc3_1.num (numerator of Normalized nonlinear autocorrelation function). We also tested the CA Algorithm [7] on real-world data, but the recall is lower than 5%. Therefore, it is not included in the following experiments.

5.2 Parameter Effects in Impact based Approaches

Effect of θ The impact threshold θ is a hyperparameter which determines how strict we are evaluating an impact. As shown in Figure 4, we can make the following observations as θ increases. 1) Severity-I and Severity-T have similar trends but Severity-T is less sensitive to large θ, because Severity-T uses impact timestamps which may still exist while impact intervals disappear when θ is large. 2) For severity based approaches, the precision increases first then drops. Since a higher threshold θ is more strict to impact generation, most detected accidents should have severe enough impacts and are always real accidents. However, as we use an overly strict threshold, only a small portion of impacts are captured in accidents, which will be mixed with the normal instances with drastic fluctuations. 3) On the other hand, the recall decreases first and then rise in severity approaches. 4) Though IIG can reach a very high precision as θ reaches 0.9, it fails to detect most real accidents. However, for the severity based approaches, the precision and recall are well-balanced, and $\theta = 0.3$ is chosen to strike a balance between the two metrics. Note that we can tune θ, e.g., choosing a high value for θ for detecting very severe accidents. In the following experiments, θ will be set to 0.3 as the impact threshold.

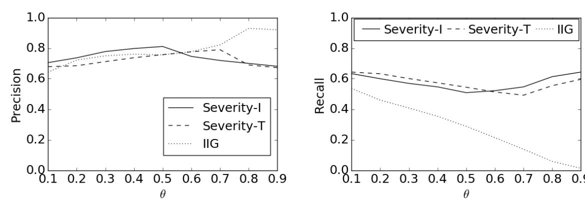

(a) Precision (b) Recall

Fig. 4: Vary θ

Effect of severity options While controlling all other parameters, we vary the selection of different severity options and compare their effects. As listed in Table 1, λ_{avg} has a lower recall than λ_{max} since the average option can smooth a significant speed drop, the algorithm may miss those accidents. The effect and trend of τ are very similar to λ. In addition, the use of consecutive option in σ will exclude fluctuating normal traffic speeds. So σ_{cons} has a better precision than σ_{disc}. In the end, we choose λ_{max}, τ_{max}, σ_{cons} as the default severity options.

Table 1: Vary Severity Options

λ	prec.	rec.	f-score
max	0.78	0.57	0.66
avg	0.79	0.52	0.63

τ	prec.	rec.	f-score
max	0.78	0.57	0.66
avg	0.79	0.55	0.65

σ	prec.	rec.	f-score
cons	0.78	0.57	0.66
disc	0.75	0.58	0.65

Effect of classifier After extracting severity features of Severity-I, we applied different approaches to make the classification. In Table 2, we can observe that decision tree is not so good as others because the overall severity is the combination of the three severity features, which is not easily separable by decisions. GBDT, Logistic Regression, SVM and AdaBoost have similar performance.

Table 2: Vary Classifiers

method	prec.	rec.	f-score	acc.
Logistic Regression	0.78	0.57	0.66	0.72
SVM	0.78	0.57	0.66	0.73
DecisionTree	0.66	0.58	0.62	0.67
GBDT	0.79	0.58	0.67	0.73
Neural Network	0.78	0.58	0.66	0.73
AdaBoost	0.78	0.59	0.67	0.73

Table 3: Comparison of Approaches

metric	prec.	rec.	f-score	acc.
Severity-I	**0.79**	0.58	**0.67**	**0.73**
FS	0.72	**0.61**	0.66	0.71
DTW	0.7	0.6	0.64	0.69
IIG	0.75	0.41	0.53	0.66

5.3 Comparison of Different Approaches

To evaluate the reliability of the proposed approaches, we compare them using 4 metrics: precision, recall, f-score and accuracy. As shown in Table 3, we can make following observations: 1) IIG has a high precision but low recall, f-score and accuracy. The reason is that IIG is the most strict among all approaches since it identifies accidents based on the existence of a propagation behavior across all upstream sensors; 2) kNN-MD-DTW is neither too bad nor too good in all metrics. Due to the fluctuations often observed in traffic data, the optimal alignment may not be accurate; 3)Severity-I has better precision, f-score and accuracy than any other approaches. The FS approach has the highest recall, but requires exhaustive computation to select from more than 7000 features. Our Severity-I method only relies on three features and achieves the highest precision, f-score and accuracy for accident detection.

Understand misclassification cases Real-world data always does not follow theoretical assumptions since some system noises and outliers are inevitable. Figure 5a shows a false negative case by Severity-I. Different colors of lines depict the speed at different sensors. We can observe that speed does not change much in all sensors, except a small drop at the 4th sensor. However this case is an accident in our record. Such case is difficult to be detected by any approach without any observable impact. It is possible that some reported accidents may not show much impact to the traffic because of a wrongly reported time or location. Figure 5b shows a false positive case of Severity-I. It is observable that there are many nontrivial propagated speed drops, which does not like a normal instance. For these instances, none of the proposed approaches can successfully differentiate accidents with normal cases.

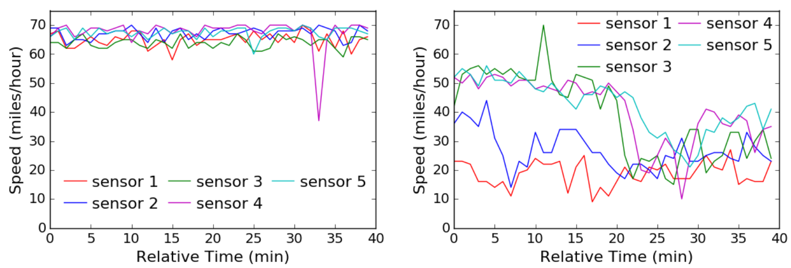

(a) A False Negative Case

(b) A False Positive Case

Fig. 5: Severity-I Misclassification Cases

6 Conclusions

We presented IIG and Severity-I, two techniques that detect traffic accidents from the traffic speed data by exploiting the spatiotemporal impact of traffic accidents. IIG detects the presence of propagation behavior by impact interval groups thus has a high precision with proper impact threshold. Severity-I takes an MTS classification approach and extracts three features i.e., drop ratio, lasting time, and impact distance, to measure the severity of an accident. As the three severity features capture different aspects of the accident impact, we are able to classify traffic speed MTS generated from nearby sensors into accidents and normal instances. Our proposed methods are evaluated and compared to state-of-the-art MTS classification methods with real-world accident and traffic data: Severity-I is shown to be superior to other methods. Future work includes integrating our algorithms into real-time traffic streaming systems. Our work can also be applied to accident severity evaluation for routing problems.

Acknowledgement

This research has been funded in part by NSF grants CNS-1461963, Caltrans-65A0533, the USC Integrated Media Systems Center (IMSC), the USC METRANS Transportation Center, and a UAlbany FRAP-A award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of the sponsors such as NSF.

References

1. M. Bashir and J. Kempf. Reduced dynamic time warping for handwriting recognition based on multidimensional time series of a novel pen device. *International Journal of Intelligent Systems and Technologies, WASET*, 3(4):194, 2008.
2. H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh. Querying and mining of time series data: experimental comparison of representations and distance measures. *Proceedings of the VLDB Endowment*, 1(2):1542–1552, 2008.
3. B. D. Fulcher and N. S. Jones. Highly comparative feature-based time-series classification. *IEEE Transactions on Knowledge and Data Engineering*, 26(12):3026–3037, 2014.
4. National Safety Council. Nsc motor vehicle fatality estimates. `http://www.nsc.org/NewsDocuments/2016/mv-fatality-report-1215.pdf`, 2016. [Online; accessed 13-Febrary-2017].
5. B. Pan, U. Demiryurek, C. Shahabi, and C. Gupta. Forecasting spatiotemporal impact of traffic incidents on road networks. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 587–596. IEEE, 2013.
6. E. Parkany and C. Xie. A complete review of incident detection algorithm & their deployment : what works and what doesn't. Technical report, New England Transportation Consortium, 2005.
7. H. Payne, E. Helfenbein, and H. Knobel. Development and testing of incident detection algorithms, volume 2: research methodology and detailed results. Technical report, FHWA, 1976.
8. H. Payne and S. Tignor. Freeway incident-detection algorithms based on decision trees with states. Technical report, National Research Council, 1978.
9. S. Sadek, A. Al-Hamadi, B. Michaelis, and U. Sayed. A statistical framework for real-time traffic accident recognition. *Journal of Signal and Information Processing*, 1(1):77–81, 2010.
10. Y. J. Stephanedes and A. P. Chassiakos. Freeway incident detection through filtering. *Transp. Res. C, Emerg. Technol.*, 1(3):219–233, September 1993.
11. G. A. Ten Holt, M. J. Reinders, and E. Hendriks. Multi-dimensional dynamic time warping for gesture recognition. In *Thirteenth annual conference of the Advanced School for Computing and Imaging*, volume 300, 2007.
12. X. Wang, A. Wirth, and L. Wang. Structure-based statistical features and multivariate time series clustering. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 351–360. IEEE, 2007.
13. H. Yoon, K. Yang, and C. Shahabi. Feature subset selection and feature ranking for multivariate time series. *IEEE transactions on knowledge and data engineering*, 17(9):1186–1198, 2005.
14. F. Yuan and R. Cheu. Incident detection using support vector machines. *Transp. Res. Part C: Emerg. Technol.*, 11(3-4):309–328, June-August 2003.
15. M. Yue, L. Fan, and C. Shahabi. Inferring traffic incident start time with loop sensor data. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 2481–2484. ACM, 2016.
16. T. Zhu, J. Wang, and W. Lv. Outlier mining based automatic incident detection on urban arterial road. *Proceedings of the 6th International Conference on Mobile Technology, Application & Systems*, (29), September 2009.