A

Accident Impact Prediction

Cyrus Shahabi[1] and Bei (Penny) Pan[2]
[1]Integrated Media Systems Center, University of Southern California, Los Angeles, CA, USA
[2]Microsoft Corp., Redmond, WA, USA

Definition

For the first time, real-time high-fidelity spatiotemporal data on the transportation networks of major cities have become available. This gold mine of data can be utilized to learn about the behavior of traffic congestion at different times and locations, potentially resulting in major savings in time and fuel, the two important commodities of the twenty-first century. According to FASANA Motion report (Report 2012), approximately 50 % of the freeway congestions are caused by nonrecurring issues, such as traffic accidents, weather hazard, special events, and construction zone closures. Hence, it is fairly important to quantify and predict the impact of traffic incidents on the surrounding traffic. This quantification can alleviate the significant financial and time losses attributed to traffic incidents, for example, it can be used by city transportation agencies for providing evacuation plan to eliminate potential congested grid locks, for effective dispatching of emergency vehicles, or even for long-term policy-making. Moreover, the predictive information can be either used by a driver directly to avoid potential gridlocks or consumed by a predictive route-planning algorithm (e.g., Demiryurek et al. 2011) to ensure a driver to select the best route *from the start*.

The McKinsey report (McK 2011) predicts a worldwide consumer saving of more than $600 billion annually by 2020 for location-based services, where the biggest single consumer benefit will be from time and fuel savings from navigation services tapping into real-time traffic data. Therefore, let us consider a navigation system utilizing predictive route-planning algorithm as a next-generation consumer navigation system (in-car or on smartphone). We notate such systems as *ClearPath*, as a motivating application, which can help drivers to effectively plan their routes in real time by avoiding the incidents' impact areas. That is, suppose an accident is reported in real time (by crowdsourcing (WAZE 2014) or through agency reports or SIGALERTS (2013)) in front of a driver, but the accident is 20 min away. If we can effectively predict the impact of the accident, ClearPath would know that this accident would be cleared in the next 10 min. Thereby, ClearPath would guide the driver directly toward the accident because it knows that by the time the driver arrives in the area, there would be no accident.

Problem Definition: For a traffic accident e occurring at time t_0, given two transportation datasets: (1) traffic accident reports and (2) traffic sensor data collected from a historical time stamp

© Springer International Publishing Switzerland 2015
S. Shekhar et al. (eds.), *Encyclopedia of GIS*,
DOI 10.1007/978-3-319-23519-6_1568-1

until t_0, the following three sets of parameters must be predicted:

(a) The set of road segments that are impacted by the incident: $\{r_i\}$.
(b) For each impacted road segment r_i, the significance of the impact (i.e., scale of speed decrease): Δv_i.
(c) For each impacted road segment r_i, the time stamp when the impact starts: Δt_i.

In this definition, a sensor refers to a loop detector or any other sensing device built on a road segment. It continuously (e.g., every 30 s) reports readings (e.g., speed) to reflect traffic situation on road segments. In this problem setting, to quantify the traffic situation on a road segment (e.g., impacted or not), the readings collected from the sensors located on this segment are utilized. Other terms that are seen frequently in this entry are defined as follows:

Impacted Road Segment: For a road segment r_i equipped with a sensor s and time stamp t (e.g., 8:30 AM), if the speed readings reported by s presents an anomalous decrease (e.g., 40 % drop) compared with historical daily readings at time t (i.e., average of all readings collected at 8:30 AM in the dataset), we consider r_i as impacted by traffic events.

Backlog: For a particular accident ev, its backlog (b) refers to the total length of all impacted road segments between ev's location and the last impacted road segment, along the opposite direction of vehicle flow.

Propagation Behavior: Given a traffic accident (ev) occurred at time t_0, ev's propagation behavior is defined as a time series of backlog (b) after t_0 and until it propagate to the maximum backlog. Assuming ev reaches the maximum backlog after t time units, its *propagation behavior* is represented as $\vec{b}$ or $\{b_0, b_1, \ldots, b_t\}$, where the subscript i for b_i represents the time unit after t_0.

Historical Background

Several disciplines, such as transportation science, civil engineering, policy planning, and operations research have studied the traffic congestion problem through mathematical models, simulation studies, and field surveys. However, due to the recent sensor instrumentations of road networks in major cities as well as the vast availability of auxiliary commodity sensors from which traffic information can be derived (e.g., CCTV cameras, GPS devices), for the first time a large volume of real-time traffic data at very high spatial and temporal resolutions has become available. While this is a gold mine of data, the most popular utilization of this data is to simply visualize and utilize the *current* real-time traffic congestion on online maps, car navigation systems, sig-alerts, or mobile applications. However, the most useful application of this data is to predict the traffic ahead of you during the course of a commute to avoid traffic congestions, especially in the presence of traffic accidents.

In the last decade, most of the studies on accident impact prediction are based on theoretical modeling and simulations, which can be classified into three groups: (1) deterministic queuing theory or shock wave theory (e.g., Lawson et al. (1997) and Wirasinghe (1978)), (2) heuristic methods and simulations (e.g., Pal and Sinha 2002), and (3) microscopic modeling of driver's behavior (e.g., Daganzo (1994) and Wang and Murray-Tuite (2010)). However, the outcome of these studies relies on theoretical simulations of road network traffic instead of the real-world collected traffic data. Also, none of these studies uses a source of incident data with description variables and reporting techniques, and their spatial transferability is limited.

When working with real-world data, it is important to identify certain characteristics of traffic data, such as temporal patterns of *rush hours* or the spatial impacts of *accidents*, which need to be incorporated into a data-mining technique to make the prediction much more accurate. For example, for generic time series, the observations made in the immediate past are usually a good indication of the short-term future. However, for traffic time series, this is not true in the beginning of a traffic accident. Specifically, for accident impact prediction, it is necessary to predict the sudden speed changes caused by traffic accidents in a faraway future (e.g., the next 30 min). In fact, the occurrence of most accidents involves two phenomenon: (1) *abrupt* speed changes, for example, it is very common for the traffic speed to drop 60 % when an accident occurs on freeways in LA and (2) *long*-lasting propagation of the speed changes, for example, a closer sensor to the accident may report a speed decrease in the 3rd min after its occurrence, and a further sensor may report similar decrease in the 30th min. Since traditional prediction approaches rely on the immediate past data to predict the future, they cannot effectively predict the *abrupt* speed changes and how they propagate over a *long* term. Hence, the navigation systems relied on these approaches may hardly navigate drivers around the accident impact area.

Scientific Fundamentals

For the motivating navigation application, ClearPath, to be effective, it is essential to predict specific values of speed changes and backlog lengths over the lifetime (i.e., temporal) and impact area (i.e., spatial) of an accident. In particular, the following three aspects need to be considered:

First, the numeric values of speed changes and backlog lengths. There are two major approaches to measure the impact of accidents: (1) *qualitative* approaches (i.e., classify accident's impact into conceptual categories such as "severe" or "non-severe" and "significant delay" or "slight delay") and (2) *quantitative* approaches (i.e., providing numeric measurement such as 45 % speed decrease and 3.2 miles of congested backlog). In the past, most studies focused on *qualitative* approaches for measuring impact, which makes the impact harder to quantify (e.g., Ozbay and Kachroo 1999). The qualitative measurement may be sufficient for general decision-making or response analysis, however, not precise and informative enough for ClearPath. In section "Impact Parameters", the prediction of *quantitative* information, which provides numeric measurements of the impact to the surrounding areas, is introduced.

Second, the spatiotemporal behavior of the impact. In previous studies, it was sufficient to predict the impact of an accident as a single or a set of aggregate values. For example, in the literature by Pan et al. (2012), the impact is predicted as average speed decrease or average of the backlog length. Since the impact region of an accident evolves over time and space, the outcome of prediction approach should be the exact length of time varying backlogs (i.e., evolution of congested spatial span) with different scales of speed changes. The section "Impact Propagation" will explain the prediction strategy of the propagation behavior of traffic accident.

Third, the comprehensive area impacted by a traffic accident. Most of existing researches focused on predicting the impact with respect to the set of upstream road segments impacted from a traffic accident (Kwon et al. 2006). In reality, traffic incidents may cause surges in traffic demand that overwhelm the system in their vicinity with a radically different flow from typical patterns. Section "Impact on Other Streets" explains the algorithms to forecast the impact of incidents on the nearby streets and intersecting freeways, which can (1) identify a set of road segments that will be impacted given a new incident and, (2) for each impacted road segment, predict the spatiotemporal performance decrease, i.e., determine when and how the impact will occur in time and space.

Impact Parameters

In this entry, we utilize two real-world transportation datasets: (1) accident reports and (2) traffic sensor data. And we address the problem of predicting and quantifying the impact of traffic accidents. By analyzing historic *accident* data, the main idea is to classify accidents based on their features (e.g., time, location, type of accident). Subsequently, we model the impact of each accident class on its surrounding traffic by analyzing the archived *traffic* data at the time and location of the accidents. Consequently, if a similar accident (from real-time accident data) is observed, its impact can be predicted and quantified on the surrounding traffic in real time using the information from past accidents.

The impact of a traffic accident can be characterized in multiple ways. Three typical quantification impact parameters are (1) impact backlog, (2) speed decrease caused by the accident, and (3) congestion duration.

Based on the analysis of real-world data, it is observed that the impact parameters vary across accidents with different attributes. The accident reports normally contain (but not limited to) the following metadata: (1) accident date, (2) accident start time, (3) accident location (i.e., street name, latitude, longitude), (4) accident type (Note that the accident type usually refers to one of the following: Traffic collision+no/minor injuries, Traffic collision+major injuries/ambulance, Traffic collision-no details, Signal alert, Natural weather hazard, Lane closure and Fire, etc.)., (5) type of vehicles involved if incident is an accident, and (6) number of affected lanes. Let us consider one of the attributes "*start time*" as an example. The impact backlog of accidents that happen during daytime may be large compared with accidents happening at midnight, due to higher traffic flow during the daytime. Thereby, the key to predict impact parameters (e.g., impact backlog) is to investigate which accident attributes are correlated with them. It is likely that some accident attributes are irrelevant or redundant for inferring the impact backlog. In order to identify the most correlated subset, we first process the accident attributes as normalized features and impact backlog as numerical classes. Then we apply the Correlation-based Feature Selection (CFS) algorithm (Hall and Smith 1998) on top of this normalized data to select correlated features. From the result obtained from this procedure, the following accident attributes are selected as the most relevant: {*start time*, *location*, *direction*, *type*, *#. of affected lanes*}. We use the selected attributes to categorize the traffic accident according to the values of their attributes and utilize the average value of the impact parameters in each category to predict the impact of an accident with corresponding attributes.

Impact Propagation

For next-generation navigation systems to be beneficial, it is essential to predict specific values of speed changes, backlog lengths over the lifetime (i.e., temporal), and impact area (i.e., spatial) of an accident. This is in contrast to previous scenarios where forecasting abstract or aggregate impact parameters (e.g., backlog) was sufficient.

To calculate the propagation behavior for an accident, one naive way is to record the speed changes on all the possible upstream locations. However, this method requires a fairly dense placement of sensors. In most sensor networks, the sensors reporting traffic speed are always distributed with a certain distance interval (e.g., 0.5 mile) to each other. Therefore, due to the limited data availability, it is only feasible to derive impact backlog from the locations equipped with sensors. To solve this problem and create a continuous propagation behavior, interpolation can be used, which may be achieved by curve fitting or regression analysis. An example fitting strategy is summarized as follows:

1. Utilize the distance of a sensor from the incident location to represent the impact backlog at time t, at what time they start to get "impacted."
2. Subsequently, plot the derived impact backlogs into 2D space (e.g., the scatter points in Fig. 1a) using the information from all the impacted sensors (e.g., sensor S1 to S4 in Fig. 1a) and train a function (e.g., polynomial function) to fit the plotted discrete points with minimal error.
3. Finally, utilize the learned fitting function and interpolate the backlogs at missing time stamps and generate a complete propagation behavior. Figure 1b shows the propagation behavior for our running example, where the impact backlog $\{b_0, b_1, \ldots, b_{19}\}$ is plotted at each minute.

There are alternative modeling approaches, such as the use of learned parameters in the fitting function to represent the propagative curve. The superiority of using interpolation result compared

Accident Impact Prediction, Fig. 1 Sample propagation behavior. (**a**) Fitting result. (**b**) Interpolation result

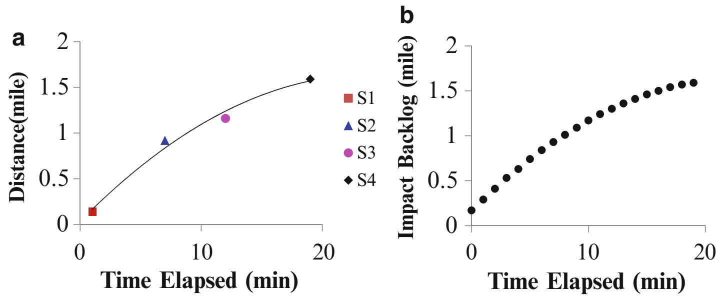

Accident Impact Prediction, Fig. 2 Impact of a traffic incident

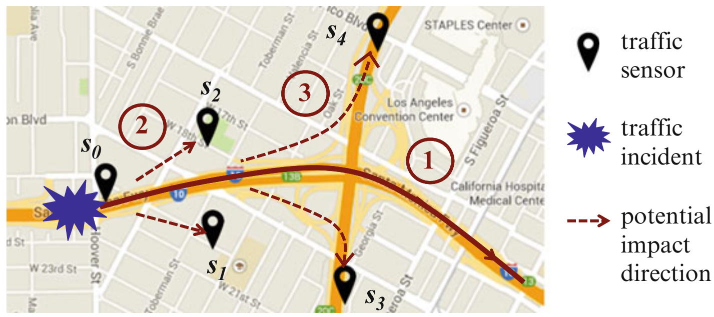

with using the parameters is as follows: (1) when the interpolation is constructed, this above strategy only uses the fitting function to interpolate the missing impact backlogs; for existing impact backlogs, it still uses the original data. However, if the coefficient vectors of the fitting function are directly used, additional fitting error might be introduced into the original data, which may result in inaccurate representation of the propagation behavior. And (2) when evaluating the prediction accuracy, the variation between the actual backlog vector and predicted backlog vector can be directly used as an error measurement, and it is also straightforward to interpret. However, the differences between the actual and predicted coefficient vectors cannot intuitively explain the prediction accuracy.

With the propagation behavior constructed, the same prediction strategy for impact parameters can be utilized to predict propagation behavior. However, in some particular cases, it is observed that although two incidents have similar attributes, their propagation behaviors are still highly different from each other. Therefore, it is important to incorporate more information such as traffic density measures (e.g., volume and occupancy) to improve the prediction accuracy. Moreover, the consideration of using a multistep prediction approach that takes into account the initial behavior (i.e., sub-pattern of propagation behavior) of an incident may further improve the prediction accuracy.

Impact on Other Streets

As illustrated in Fig. 2, the impact caused by a traffic accident on a freeway may affect the traffic flow in the following three types of locations:

(1) Upstream stretch of the occurrence freeway,
(2) Adjacent arterial streets, and
(3) Other surrounding freeways.

This section focuses on how to forecast the impact of incidents on the nearby streets and freeways (i.e., the locations (2) and (3)).

The intuitive way to predict the impact of accidents on the nearby streets and freeways is to identify the causal interactions among traffic at different road segments to address the aforementioned challenges. To identify the causality relationship, the main idea is to utilize archived traffic sensor datasets to train causality models to determine whether the time series data (e.g., collected from s_0 in Fig. 2) is *useful* for predicting

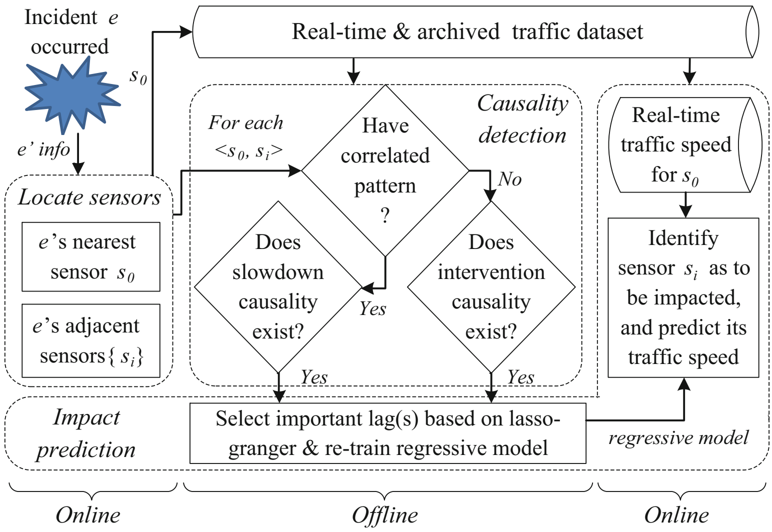

Accident Impact Prediction, Fig. 3 Flow chart for impact prediction

other time series data (e.g., collected from s_1). If the change in traffic performance (e.g., decrease or increase in traffic speed) at s_0 leads to a change in traffic performance at another location s_1, in the presence of a traffic accident near s_0, then s_1 could be identified as part of the impacted area. Consequently, given a traffic accident and its attributes, the detected causality can be utilized to predict the impact in the vicinity area of a traffic accident.

Given the strategy above, the challenge is how to detect the causality between the traffic speed time series. One straightforward idea is to use the traditional causality test (e.g.,Granger 1969) to detect the causality. However, with real-world traffic data, it is observed that hardly any Granger causality existed between any pair of traffic speed time series. This was a surprising observation and counterintuitive as it is expected strong causality relationship among traffic time series. With further investigation regarding the unique characteristics of traffic speed time series, two types of time-sensitive causalities that are unique to traffic speed time series are discovered. Specifically, for two traffic speed time series with correlated historical patterns, it is observed that sometimes the causality only exists during the beginning of rush hours when the traffic starts to become congested, named as *slowdown causality*. Such causality only exists between two road segments that have strong connectivity in the road network. Conversely, in other connectivity scenarios, especially when the two time series are not correlated, another type of causality is observed that only exists in the presence of traffic accidents and during non-rush hours, named as *intervention causality*. Consequently, the detected causalities can be utilized for predicting the impact of traffic accidents, with procedure illustrated in Fig. 3.

Given that a new incident e has just occurred, its closest upstream sensor s_0 is sent to the archived database to retrieve the relevant time intervals for causality detection. It is also utilized to search among the nearby sensors and retrieve a potential candidate sensor to be impacted. Then, the sensor pair $< s_0, s_i >$, together with the corresponding dataset and the causality detection model, is used to identify whether the slowdown causality or intervention causality exists from s_0 to s_i. If the slowdown causality exists, there is no need to examine the intervention causality because the impact of significant speed drops from traffic incidents is already covered in the definition of the slowdown causality. At the end of the causality detection, the sensor pairs ($< s_0, s_i >$) holding the causality relationship can proceed to the next step, and the sensor pairs ($< s_0, s_i >$) holding neither slowdown nor intervention causality are disregarded. In the former case, s_i is considered one of the impacted sensors that can be contributed to the spatial impact range. In the latter case, s_i is excluded from the spatial impact range caused by incident e. For sensor pairs ($< s_0, s_i >$) holding the causality relationship, the next step is to select the most important time stamps (i.e., $t + h$ given the accident occurs at t) to identify when s_i starts to become impacted. In the pipeline illustrated in Fig. 3, we resort to lasso-Granger (Arnold et al. 2007) approach to achieve this step. Note that after lasso, we need to retrain the regressive model for predicting s_i based on the selected lag in s_0. Finally, the real-time traffic speed data collected from s_0 and the learned regressive model are utilized to predict the speed of s_i.

To enable real-time impact prediction, in Fig. 3, the causality detection and important variable selection steps need to be implemented off-line for every sensor pair on the road networks. Because the training step in the regressive model and the lasso approach require access to large amounts of archived traffic time series data, the causality detection and important variable selection significantly delays the online prediction process due to a great deal of training time consumption. In this way, when a new incident occurs, the system will search within the off-line training results to identify whether causality exists between the corresponding sensor pairs and will further retrieve the learned regressive model for the online traffic speed prediction for the sensor to be impacted.

Note that in the domain of social science and economics, the causality models have already been widely applied (Pearl 1988; Glymour et al. 1987; Spirtes et al. 2001), many of which are superior to Granger causality in multivariate causality inference. However, for the impact problem, Granger causality model is a better candidate for causality detection for the following reasons. First, in this study, the ultimate goal is to enable the better *prediction of the traffic time series* in the presence of traffic incidents by taking advantage of the detected causality relationship. Revealing the complete causality relationship among all traffic data on the road network is not our focus. Compared with other causality inference models, the regressive model of Granger causality serves as a fairly effective predictor for time series data, such as traffic sensor data. Second, for accident impact prediction problem, it not only needs the identification of the causality but also the time lag of the causality (i.e., how much time needs to pass until a road segment's traffic gets impacted by a traffic incident). For the time series-based Granger causality model, such time lag can be effectively learned through the model's learning process. However, for the graph-based causal inference models (e.g., Pearl (1988) and Glymour et al. (1987)), it is fairly difficult to learn such a temporal dependency in the detected causality. Finally, the existing literatures focus on predicting the impact of one traffic incident at a time. In particular, the one-to-one causality relationship detection between the traffic at the incident location and one other location is the major focus in transportation networks. It is entirely possible that traffic at different locations is causally dependent, and they may have more than one cause. However, such cases are barely useful to this problem, which is predicting the impact by a single cause (i.e., a particular traffic incident). Thus, the Granger causality, even though it ignores multivariate dependencies, is particularly effective for this purpose.

Key Applications

Navigation Systems

The result of impact prediction can be applied in smart-routing applications in real time to help users avoid unexpected congestion. Specifically, when there is traffic accident, the prediction result of event impact including the backlogs and speed decrease caused by the traffic events can also be utilized for the purpose of avoiding traffic congestion. To be more specific, consider another example illustrated in Fig. 4. In this figure, the caution mark, the directed solid red lines, and the dashed blue lines represent the incident location, the congested region caused by the incident, and the route a driver plans to follow, respectively. Without prediction, but with the

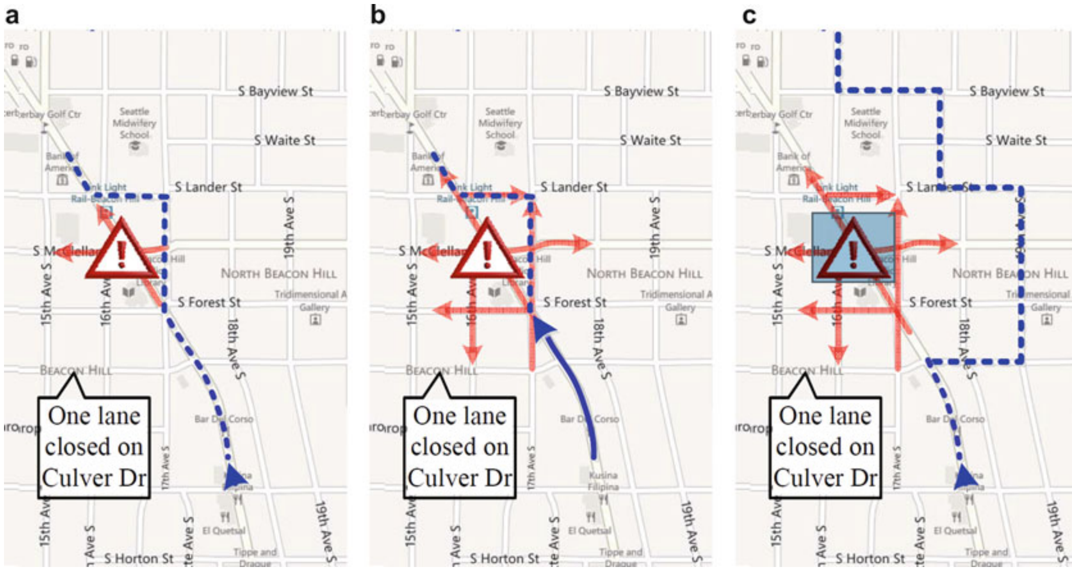

Accident Impact Prediction, Fig. 4 (**a**) Route calculated based on current incident's impact. (**b**) Time-varying expansion of impacted region as driver approaches the incident location. (**c**) Route calculated based on accurate prediction of impact

knowledge of the incident, a typical navigation application, such as Waze (WAZE 2014), may suggest the route shown in Fig. 4a to the drivers. If the driver follows this route, he would be stuck in the traffic congestion caused by the incident, as illustrated in Fig. 4b, due to the fact that the congested region has grown. On the other hand, if we can predict how the impacted spatial span (i.e., congested region) evolves over time, ClearPath could calculate the route that can effectively avoid the congestion from the beginning, as shown in Fig. 4c.

Public Policy and Decision-Making

The accident impact prediction cannot only benefit individual drivers through navigation system but also transportation authorities, e.g., by notifying drivers when they are approaching an accident and suggesting alternative routes, as well as implementing traffic jam diagnosis and dispersal. Moreover, the predicted result can be visualized through a web-based user interface, which provides the transportation authorities with a global view of all the traffic accidents in a city. Equipped with such a service, transportation authorities could efficiently monitor all the traffic accidents with detailed diagnoses of their impact regions for the purpose of better policy and decision-making.

Future Directions

The research on accident impact prediction can be extended in several directions. First, besides traffic accidents, more complex events causing congestions, such as large-scale parades or sporting events, can be studied, and their impacts can be predicted. Second, besides traffic sensor datasets, other modalities of data acquired from video cameras and/or mobile phones can also be utilized for better prediction of accident impacts. Finally, studying the long-term or online strategies to update the prediction models using streaming data is of great importance.

Cross-Reference

▸ Causality Detection
▸ Predictive Route Planning
▸ Time Series Prediction

A

References

Arnold A, Liu Y, Abe N (2007) Temporal causal modeling with graphical granger methods. In: Proceedings of the 13th ACM SIGKDD international conference on knowledge discovery and data mining, KDD'07. ACM, New York, pp 66–75

Daganzo CF (1994) The cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory. Transp Res Part B: Methodol 28:269–287

Demiryurek U, Banaei-Kashani F, Shahabi C, Ranganathan A (2011) Online computation of fastest path in time-dependent spatial networks. In: SSTD, Minneapolis

Glymour C, Scheines R, Spirtes P, Kelly K (1987) Discovering causal structure. Academic, Orlando

Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37:424–438

Hall MA, Smith LA (1998) Practical feature subset selection for machine learning. In: ACSC98, Perth. Springer, Berlin, pp 181–191

Kwon J, Mauch M, Varaiya PP (2006) Components of congestion: delay from incidents, special events, lane closures, weather, potential ramp metering gain, and excess demand. Transp Res Rec 1959:84–91

Lawson TW, Lovell DJ, Daganzo CF (1997) Using the input-output diagram to determine the spatial and temporal extents of a queue upstream of a bottleneck. Trans Res Rec 1572:140–147

Manyika J, Chui M, Brown B, Bughin J, Dobbs R, Roxburgh C, Byers AH (2011) Big data: the next frontier for innovation, competition, and productivity. McKinsey Global Institute, New York

Ozbay K, Kachroo P (1999) Incident management in intelligent transportation systems. Artech House, Norwood, MA

Pal R, Sinha KC (2002) Simulation model for evaluating and improving effectiveness of freeway service patrol programs. J Transp Eng 128:355–365

Pan B, Demiryurek U, Shahabi C (2012) Utilizing real-world transportation data for accurate traffic prediction. In: ICDM, Brussels

Pearl J (1988) Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan Kaufmann, San Mateo

Report FM (2012) Http://www.metro.net/board/Items/2012/03March/20120322RBMItem57.pdf

SIGALERT (2013) Http://www.sigalert.com. Last visited May 2013

Spirtes P, Glymour C, Scheines R (2001) Causation, prediction, and search. MIT, Cambridge

Wang Z, Murray-Tuite PM (2010) A cellular automata approach to estimate incident-related travel time on interstate 66 in near real time. Virginia Transportation Research Council, Charlottesville

WAZE (2014) Http://www.waze.com. Last visited May 2014

Wirasinghe SC (1978) Determination of traffic delays from shock-wave analysis. Trans Res 12:343–348