

Where *Not* to Go? Detecting Road Hazards Using Twitter

Avinash Kumar, Miao Jiang, Yi Fang
Department of Computer Engineering
Santa Clara University
500 El Camino Real, Santa Clara, California, USA
yfang@scu.edu

ABSTRACT

Conventional approaches to road hazard detection involve manual inspections of roads by government transportation agencies. These approaches are usually expensive to execute, and sometimes are not able to capture the most recent hazards. Moreover, they often only focus on major highways due to a lack of sufficient manpower. Consequently, many hazards on minor roads get ignored, which may pose serious dangers to drivers. In this paper, we demonstrate an application of Twitter to automatically determining road hazards. By building language models based on Twitter users' online communication, our system aims at pinpointing potential road hazards that pose driving risks. The likelihood of poor driving conditions can then be exposed via map overlays to warn drivers about potentially dangerous driving conditions in their locale or on current routes, thereby significantly reducing the chances of an accident occurring. To the best of our knowledge, this is the first work demonstrating the utility of social media to automatically detect road hazards. We conduct experiments on a testbed of tweets discussing road conditions and the initial results demonstrate the effectiveness of our approach.

Categories and Subject Descriptors

H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing

Keywords

Twitter; Sentiment analysis; Language models; Road hazards

1. INTRODUCTION

Many people die from or are injured by road hazard-related accidents every day, according to the NHTSA [1]. Dangerous road conditions pose serious driving risks. Conventional approaches rely on manual inspection of road by government transportation agencies. However, these approaches cannot scale to the mass amount of roads given the enormous expense of manpower required to conduct such operations. Consequently, the road inspection is only limited to major highways, which cause many hazards on minor roads go unnoticed and unattended to.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '14, July 6–11, 2014, Gold Coast, Queensland, Australia.
Copyright © 2014 ACM 978-1-4503-2257-7/14/07...\$15.00.
<http://dx.doi.org/10.1145/2600428.2609550>

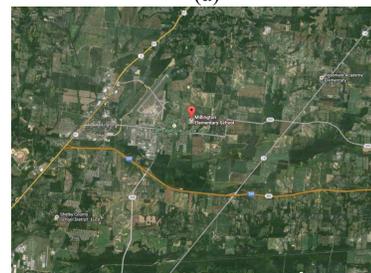
On the other hand, Twitter has become a prominent source of real-time news distribution amongst social networking sites. The inclusion of location-based services associates with each tweet a “where” as well as a “what”, with no added effort to the user. This presents a multitude of opportunities to leverage the vast amount of data offered by Twitter, as new information can be inferred from this data and shared amongst its user base.

In this paper, we explore the use of Twitter as a source of road hazard detection by aggregating hazard-related information posted by Twitter users. Figure 1 (a) gives an example of such tweets. It mentions that the bridge on highway TN-205 was so icy that the school bus could not cross, yet the district was still open, indicating that no government agencies issued any warning on the dangerous situation of this road. This example illustrates that while government agencies may not be able to detect road hazards rapidly, Twitter can disclose such information in a timely manner. Figure 1 (b) shows the geographical location of the tweet due to location-based services in many mobile phones. We can associate the location of the tweet with the event indicated by it, and thus pinpoint the exact location of the hazard. Moreover, by leveraging information retrieval and text mining techniques, we can potentially identify dangerous situations automatically with little human effort.

 Lucas Johnson @ljohnson0910 · 1m
Bridge @ TN205 is so icy @SCSK12Unified buses can't go across. The district opened despite icy road condition #Memisc
pic.twitter.com



(a)



(b)

Figure 1. (a) A tweet mentioning a road hazard was posted by the user along with a picture indicating the scene; (b) The geographical location associated with the tweet

I brake for deers literally! Saw about 13 deers leaping across the road 3 hr drive in dark dangerous!
[twitpic.com/aawx1](https://twitter.com/aawx1)

Figure 2. A tweet mentioning a dangerous road condition with deers across the road

Compared with other sources of road hazard information, Twitter is quite unique since it is able to cover a wide range of road hazards with its huge user base. Figure 2 depicts a tweet concerning a possible deer collision. This kind of hazard may often get ignored in a government-associated warning system, mainly because 1) it is hard to capture the information, and 2) it is considered insignificant. However, an incident such as this could cause serious accidents for drivers on the road.

To the best of our knowledge, this is the first work demonstrating the utility of Twitter for automatic road hazard detection. We collect tweets possibly discussing road hazards from Twitter’s Streaming APIs¹. We formulate the core task as a sentiment classification problem by analyzing the sentiment expressed in tweets. If the tweet is classified as having negative sentiment concerning a road condition, we generate an alarm based on that tweet. The sentiment classification component is a machine-learning model trained on a set of labeled data. Various machine-learning methods are explored in these experiments, demonstrating good classification and accurate detection.

2. RELATED WORK

Twitter has been used for various prediction tasks in recent years. Sakaki et. al [4] devised a classifier of tweets based on features including keywords and contexts to detect earthquakes. Subsequently, they produced a probabilistic spatiotemporal model for the target event that can find the center and the trajectory of the event location. In their work, they treat Twitter users as sensors, making assumptions that a user will detect a target event and generate a report probabilistically. In addition, they also associate each tweet with a time and location. Mathioudakis et. al [2] proposed a method of detecting trends over the Twitter stream to identify emerging topics in real-time. Twitter data has also been used in public health monitoring [6], and election prediction [7], among many other applications.

Many government agencies start sharing road information via Twitter such as California Highway News (@cahighways), KCBS traffic (@KCBSTraffic), and Bay Area Road Alert (@sfbayRoadAlerts). However, they mainly focus on traffic information for major highways. Our work aims to identify all the hazard related information available on Twitter, regardless of the types of hazards and roads, attempting to execute automatically and in a timely manner.

Table 1. A sample list of queries used to retrieve road hazard related tweets from Twitter

road hazard, hit animal, animal crossing, deer road, tires slipping, road blocked, road accident, road danger, highway closure, road foggy, road slippery, icy road

3. OUR APPROACH

The pipeline of our approach is illustrated in Figure 3. There are two major components: training and runtime detection. In training, a set of tweets associated with road hazards are collected by applying a search filter to Twitter historical tweet archives. This set of tweets is further filtered based on location, and is then preprocessed and normalized. The tweets are then manually classified as either road hazard or not road hazard. We make the claim that there is a relationship between negative sentiment and the mention of road hazards within a tweet. In other words, if a tweet containing road hazard information can be recognized as having negative sentiment, it indicates that a road hazard exists. Based on the labeled tweets, we train a sentiment classification model using language models. In the runtime detection phase, we receive new tweets related to road conditions in real time, and apply the same preprocessing steps with the training phase. We then apply the trained language model to classify the new tweets. If a tweet is predicted as having negative sentiment by the model, an alarm will be generated for hazard detection. The following subsections explain the individual components in more detail.

3.1. Retrieving Tweets from Twitter

We utilize the Twitter Streaming APIs to extract tweets, using search filters containing specific terms that relate to traffic and road hazards. To catch as many hazard-related road conditions as possible, the selection of queries is important. We divide possible hazardous situations into five categories: animals, emergency, weather, special events, and traffic. These allow us to generate a list of terms and phrases that attempt to cover a comprehensive set of hazardous events. Totally, we generated 103 queries. Table 1 shows a sample list of queries we used to retrieve road hazard related tweets.

3.2 Filtered by Location and Language Processing

We propose a method to label geographic locations as potentially hazardous based on tweets and the location information they provide. It is not necessary for a tweet to mention where an event occurred: so long as the user’s location services are enabled via Twitter, geographic coordinates can be extracted from the tweet and mapped to a specific road or road segment. We assume that the user is sending a tweet immediately or soon after following such an event: the GPS coordinates provided are then correlated to a road or intersection on a map. Any tweets for which location services have not been enabled are discarded. After we obtain a set of candidate tweets, we further process the tweets by removing a list of stop words and applying Porter stemming to reduce inflected words to their stem.

¹ <https://dev.twitter.com/docs/streaming-apis>

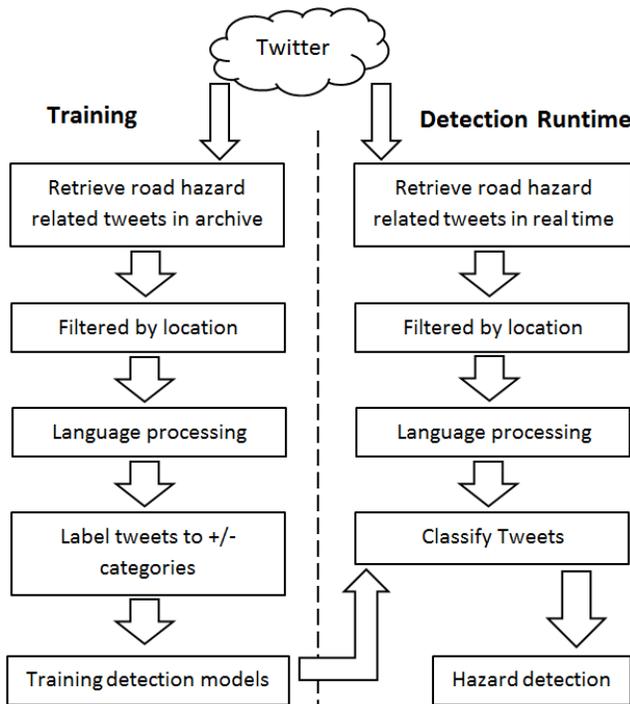


Figure 3. Architecture of our hazard detection system.

3.3 Sentiment Analysis and Language Models

Some prior studies have explored the classification of tweets from an opinion mining perspective. Pak and et.al. [3] built a sentiment classifier for Twitter to determine positive, negative and neutral sentiment. They also proposed a classifier to make judgments on whether a Twitter user gives positive or negative opinions across an entire issue. We extrapolate this concept to the theory that if sentiment from tweets can be classified and predicted, sentiments in tweets concerning road hazards can also be predicted, and vice versa. To train our classification models, we manually label a set of training data to classify each tweet into one of the two categories, *hazard* and *not hazard*.

While there exist various machine learning methods, the focus of our work is not on choosing the best learning model for sentiment classification, but on demonstrating the utility of sentiment classification in road hazard detection. In the experiments, we explore three machine learning methods: Naïve Bayes, K-nearest-neighbor, and a dynamic sentiment classification method proposed by Pang and Lee [5]. These three methods are essentially language model classifiers that model different usage patterns of language in respective negative and non-negative sentiment categories.

3.4 Hazard Detection at Runtime

Once we have the sentiment classification model trained, we apply it to predict and detect hazards at runtime. Specifically, we retrieve a stream of new tweets from Twitter in real time, and then apply the same location filter and language processing to preprocess the tweets. We then apply the trained language model to the new tweets to classify the tweets as having negative or non-negative sentiment. All the candidate tweets are assumed to be related to road hazards as they are collected by the procedure in

Table 2. The classification results (9-fold cross-validation) for K-nearest neighborhood (KNN), Naïve Bayes, and Dynamic Language Model (DLM). @- denotes the metric is with respect to the negative category (Hazard) and @+ is with respect to the positive category (Non-hazard).

	KNN	Naïve Bayes	Dynamic LM
Precision@-	0.672	0.775	0.757
Recall@-	0.534	0.515	0.498
Precision@+	0.813	0.817	0.809
Recall@+	0.890	0.938	0.934
Accuracy	0.782	0.812	0.800

Section 3.1. Thus, if the tweet expresses negative sentiment, we will regard it as the tweet with road hazard information.

4. EXPERIMENTS

In this section, we conduct some preliminary experiments and analysis to validate our proposed approach. To create a testbed of tweets, we crawled public tweets from Twitter Streaming API over the period from Feb 7 to Feb 10, 2014. We retrieved 30,876 tweets in total. For the purposes of this experiment, we restricted input to English and discarded all non-English tweets retrieved from the Stream APIs. We then labeled the tweets into hazard or non-hazard category. Once we obtain the labeled instances, we use them to train sentiment classification models. In these experiments, we use the LingPipe sentiment analysis tool² and compared three machine learning models: K-nearest neighborhood (KNN), Naïve Bayes, and the Dynamic Language Model (DLM). The evaluation metrics are precision, recall, and accuracy. We calculated precision and recall with respect to both the negative sentiment (hazard) category and non-negative sentiment (non-hazard) category. We randomly split the data into 9 slices and conduct 9-fold cross validation to evaluate the models. We chose 8-gram language models in these benchmarks. Table 2 shows the results.

In the application of road hazard detection, we care more about the negative sentiment (hazard) category than the non-negative sentiment category. As we can see from the table, Naïve Bayes yielded the best performance in precision for the negative category with a value of 0.775. This is quite an encouraging result, as it means out of 10 negative instances it found, and average of 7 or 8 are indeed negative and which are likely to contain useful hazard information. The Dynamic LM generated comparable results with Naïve Bayes. On the other hand, KNN gave the best recall result for the negative sentiment category, while the performance is not as good in precision. Methods of improving recall will be an important research direction in our future work. In overall accuracy, the Naïve Bayes and Dynamic LM approaches outperformed the instance-based KNN.

We also investigated the effect of n in n -gram on the performance of the models. Figure 4 shows the results in accuracy. We find that as n increases, the performance of the three models also increases at the beginning and flattens after $n=3$ or 4, indicating that $n=3$ or 4 is probably a good choice in this application.

²<http://alias-i.com/lingpipe/demos/tutorial/sentiment/>

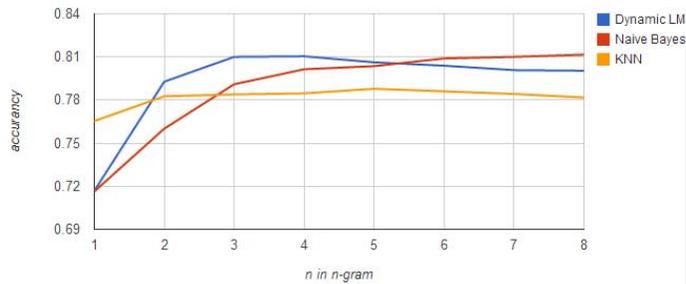


Figure 4. The effect of n in n -gram on the performance of the three models.

We further dig into the individual tweets we classified by the dynamic language model. Table 3 shows some sample tweets in four categories: true negative, true positive, false negative, false positive, in the terminology of a confusion matrix. In the context of our problem, a true positive instance means it is the tweet correctly classified as non-hazard and a true negative refers to a tweet that is correctly identified as hazard. From the table, we found that our approach can correctly identify “*It’s 2am and our driver, has been navigating insanely icy back road*” as a tweet with hazard information. This may be due to the fact it contains some words expressing negative sentiment such as “insanely icy”. On the other hand, the positive sentiment word “nice” in the false positive instance “*The house in London are so nice, could never live in them though traffic in London is too much*” may have confused the classifier.

The false negative example in the table demonstrates that search terms and phrases can be semantically ambiguous. At the time of this analysis, in recent football news the term “hazard” returns tweets related to Chelsea superstar Eden Hazard. Even if categorized as a non-hazard, the classifier was confused due to the inability to recognize such a common term as a proper noun. This false negative could be resolved by implementing a more advanced and customized natural language processing tool. We will explore this direction in the future work.

An interesting example of error is the second tweet in the false positive category of the table. It seems to contain some positive sentiment words such as “awesome” and “enjoy”, but the real sentiment is negative. This presence of sarcasm poses a big challenge to our specific application and to sentiment analysis in general.

5. CONCLUSION AND FUTURE WORK

We present a proof-of-concept for a novel utilization of Twitter as a real-time traffic and road hazard alert system. The ability to accurately correlate a filtered collection of tweets to road hazards establishes the foundation for several practical applications. For example, this information can be broadcasted via another Twitter account, to which users can subscribe and receive anonymously tweeted road hazard information.

The utilization of social networks as a reliable news source is becoming an increasingly popular phenomenon. Twitter is on the forefront of this paradigm shift, and is as useful practically as it is experimentally. Both users and researchers benefit from the seemingly instantaneous delivery time of tweets. While the character limit of a tweet is efficient in reading for users,

Table 3. Some sample tweets for the 4 categories of error analysis on our classified results

	Sample tweet
True negative	“ <i>It’s 2am and our driver, has been navigating insanely icy back road</i> ”
True positive	“ <i>Just posted a photo @ Pacific Coast Highway http://t.co/U1bXkmkclU</i> ”
False negative	“ <i>@MdFirdaus95 literally hazard IS danger hahaha anw lets shall see what happens next round of matches. Should be interesting ;)</i> ”
False positive	“ <i>The house in London are so nice, could never live in them though traffic in London is too much</i> ” “ <i>Everyone stuck in traffic on I-35 should turn off their AC, roll down the window, and enjoy this day because it’s awesome outside!</i> ”

researchers can also apply more complex text processing algorithms with near-negligible increases in latency. These improvements upon our methodology, combined with the provision of a user-friendly interface to which Twitter members can refer while driving, allow for an application that greatly reduces the potential for road accidents.

The proposed approach is an initial step towards a very promising research direction. In our future work, we plan on developing more sophisticated models to increase accuracy in distinguishing road hazards. By utilizing map overlays, hazard-classified tweets with location information can display pinpoints where the tweet was sent, which correspond approximately to the actual incident. This information can then be shared with other users. It is worth noting that the assumption that tweets are reliable is paramount to the success of this application, which needs to be further verified in the future work.

6. ACKNOWLEDGMENTS

We greatly appreciate the valuable comments by the anonymous reviewers. We thank Ching Lien for her help in querying road hazard related tweets.

7. REFERENCES

- [1] <http://www.nhtsa.gov/NCSA>.
- [2] Mathioudakis, M, and Koudas N. TwitterMonitor: Trend Detection over the Twitter Stream. In *SIGMOD*, 2010.
- [3] Pak, A, and Paroubek, P. Twitter as a Corpus for Sentiment Analysis and Opinion Mining. In *LREC*, 2010
- [4] Sakaki, T, Okazaki, M, and Matsuo, Y. Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors. In *WWW*, 2010.
- [5] Pang, B. and Lee, L. A Sentimental Education: Sentiment Analysis Using Subjectivity. In *ACL*, 2004.
- [6] Sadilek, A., Brennan, S., Kautz, H., and V. Silenzio. NEmesis: Which Restaurants Should You Avoid Today? In *AAAI*, 2013.
- [7] Tumasjan, A., Sprenger, T. O., Sandner, P. G., and Welpe, I. M. Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. In *ICWSM*, 2010.