

# Twitcident: Fighting Fire with Information from Social Web Streams

Fabian Abel, Claudia Hauff,  
Geert-Jan Houben, Ke Tao  
TU Delft, Web Information Systems  
PO Box 5031, 2600 GA Delft, the Netherlands  
{f.abel,c.hauff,g.j.p.m.houben,  
k.tao}@tudelft.nl

Richard Stronkman  
Twitcident.com  
Koning Wilheminaplein 400, Amsterdam, the  
Netherlands  
richard@twitcident.com

## ABSTRACT

In this paper, we present Twitcident, a framework and Web-based system for filtering, searching and analyzing information about real-world incidents or crises. Twitcident connects to emergency broadcasting services and automatically starts tracking and filtering information from Social Web streams (Twitter) when a new incident occurs. It enriches the semantics of streamed Twitter messages to profile incidents and to continuously improve and adapt the information filtering to the current temporal context. Faceted search and analytical tools allow users to retrieve particular information fragments and overview and analyze the current situation as reported on the Social Web.

Demo: <http://wis.ewi.tudelft.nl/twitcident/>

## Categories and Subject Descriptors

H.3.3 [Information Systems]: Information Search and Retrieval—*Information filtering*

## General Terms

Algorithms, Design, Experimentation

## Keywords

Social Web Streams, Filtering, Faceted Search, Semantics

## 1. INTRODUCTION

During crisis situations such as large fires, storms or other types of incidents, people nowadays report and discuss about their observations, experiences, and opinions in various Social Web streams. Recent studies show that data from the Social Web and particularly Twitter helps to detect incidents [4] or to analyze the information that people report about an incident [2]. However, (i) automatically filtering relevant information from Social Web streams and (ii) making the information accessible and findable in a given inci-

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2012 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

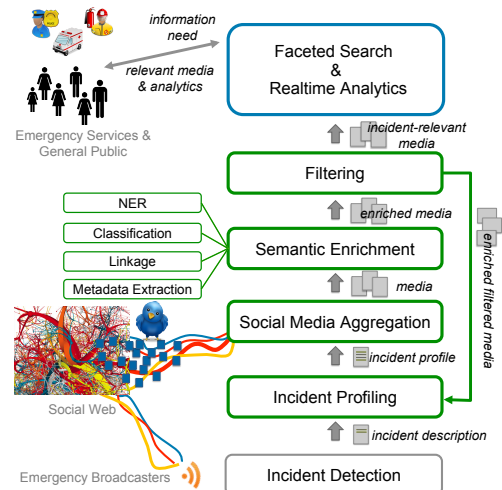


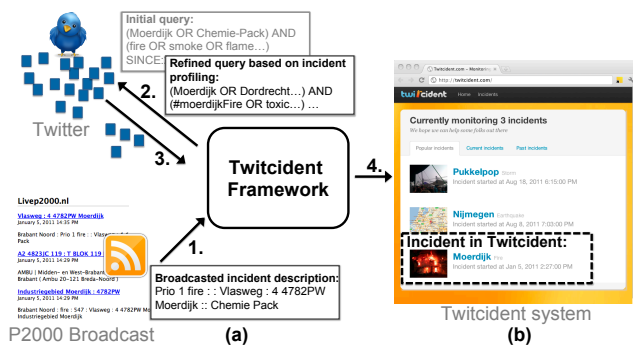
Figure 1: Twitcident architecture: (i) *incident profiling* and *filtering* of social media that is relevant to an incident (green boxes) and (ii) *faceted search* and *realtime analytics* functionality to explore and overview the media (blue box). Both types of components benefit from the *semantic enrichment* of the aggregated media.

dent context are two fundamental challenges that have not been researched sufficiently yet. In this paper, we demonstrate solutions to these challenges.

Filtering and search in Social Web streams differs from traditional Web search. For example, Teevan et al. [6] reveal significant differences in the search behavior on Twitter compared to Web search: Twitter users are specifically interested in information related to events and often use the rudimentary search functionality of Twitter to monitor information streams. With Twitcident, we introduce a framework that automates the process of monitoring relevant information that is published in Social Web streams. Faceted search capabilities [1] allow end-users to further filter and explore the information.

## 2. TWITCIDENT

In this section, we first overview the architecture of the Twitcident framework and then detail its key components (see Figure 1). The Web-based front-end of the Twitcident system is depicted in Figure 3.



**Figure 2: Incident detection:** (1) as soon as an incident is broadcasted via the P2000 network, the Twitcident framework (2) transforms the encoded P2000 message into an initial incident query to (3) collect Twitter messages that are possibly relevant for the incident so that (4) information about the incident can be accessed via the Twitcident system. Over time, the incident profiling effects refinements of the queries that are used to collect tweets. The screenshot shows the dashboard of popular incidents that are (and have been) monitored by Twitcident.

## 2.1 Architecture and Overview

The core framework functionality is triggered by an incident detection module that senses for incidents being broadcasted by emergency services. Whenever an incident is detected, Twitcident starts a new thread for profiling the incident and aggregating social media and Twitter messages from the Web. The collected messages are further processed by the semantic enrichment module which features named entity recognition (NER), classification of messages, linkage of messages to external Web resources and further metadata extraction. The semantic enrichment is one of the key enabling components of the Twitcident framework as it (i) supports semantic filtering of Twitter messages to identify those tweets that are relevant for a given incident, (ii) allows for faceted search on the filtered media and (iii) provides means for summarizing information about incidents and performing realtime analytics.

In the Twitcident system, both faceted search and realtime analytics are made available to client users via a graphical user interface that is displayed in Figure 3. The search functionality allows end-users to further filter messages about an incident while analytics deliver diagrams and gadgets that enable users to analyze and overview how people report about the incident on the Social Web. We now discuss each of the components of our architecture in detail.

## 2.2 Incident Detection

For detecting incidents, the Twitcident system relies on emergency broadcasting services. In the Netherlands, incidents which require the police, fire department or other public emergency services to take an action and which are moreover of interest to the general public, are immediately published via the *P2000 communication network* and describe what type of incident has happened, where and when it happened and also what scale the incident is classified as. Figure 2(a) shows an example P2000 message inform-

ing about a large fire incident that happened in the city of Moerdijk, the Netherlands, on January 5th 2011. The figure visualizes the automatic workflow that is triggered whenever a new incident is reported. If a new incident is detected then the Twitcident framework translates the broadcasted message into an initial incident profile that is applied as query to collect relevant messages from Twitter. All incidents that are monitored by the Twitcident system are listed on the dashboard that is depicted in Figure 2(b).

## 2.3 Incident Profiling and Filtering

While monitoring an incident, Twitcident continuously adapts the incident profiling to improve the filtering of messages. This process is realized via the following components (see Figure 1): (i) incident profiling, (ii) social media aggregation, (iii) semantic enrichment and (iv) filtering.

### 2.3.1 Incident Profiling

Based on the initial incident description and the collected, enriched Social Web messages, the incident profiling module generates a profile that is used to refine the media aggregation and the filtering. An incident profile is a set of weighted attribute-value pairs that describe the characteristics of the incident. Attributes may refer to concepts that are related to the incident (e.g. locations, persons) and the weights specify the importance of an attribute-value pair for the incident. Therefore, the aforementioned fire that happened in Moerdijk may have the following incident profile:  $P(i_{moerdijk}) = \{((location, Moerdijk), 1.0), ((location, Dordrecht), 0.73), ((type, Fire), 1.0), \dots\}$ . Incident profiles are continuously updated to adapt to topic changes that arise within an incident. To prevent topic drift, we combine the current profile with the initial incident profile following a classical mixture approach.

### 2.3.2 Social Media Aggregation

Based on the incident profiling, the Twitcident system exploits the social media aggregation component to collect Twitter messages as well as related pictures and videos that are posted on platforms such as Twitpic or Twitvid<sup>1</sup> respectively. Twitcident utilizes both the REST API and the Streaming API of Twitter to collect messages. The REST API enables Twitcident to collect also those incident-related tweets that have been posted before the emergency services reported about a given incident.

### 2.3.3 Semantic Enrichment

The aggregated Social Web content (Twitter messages) is processed by the semantic enrichment component of Twitcident which features the following functionality.

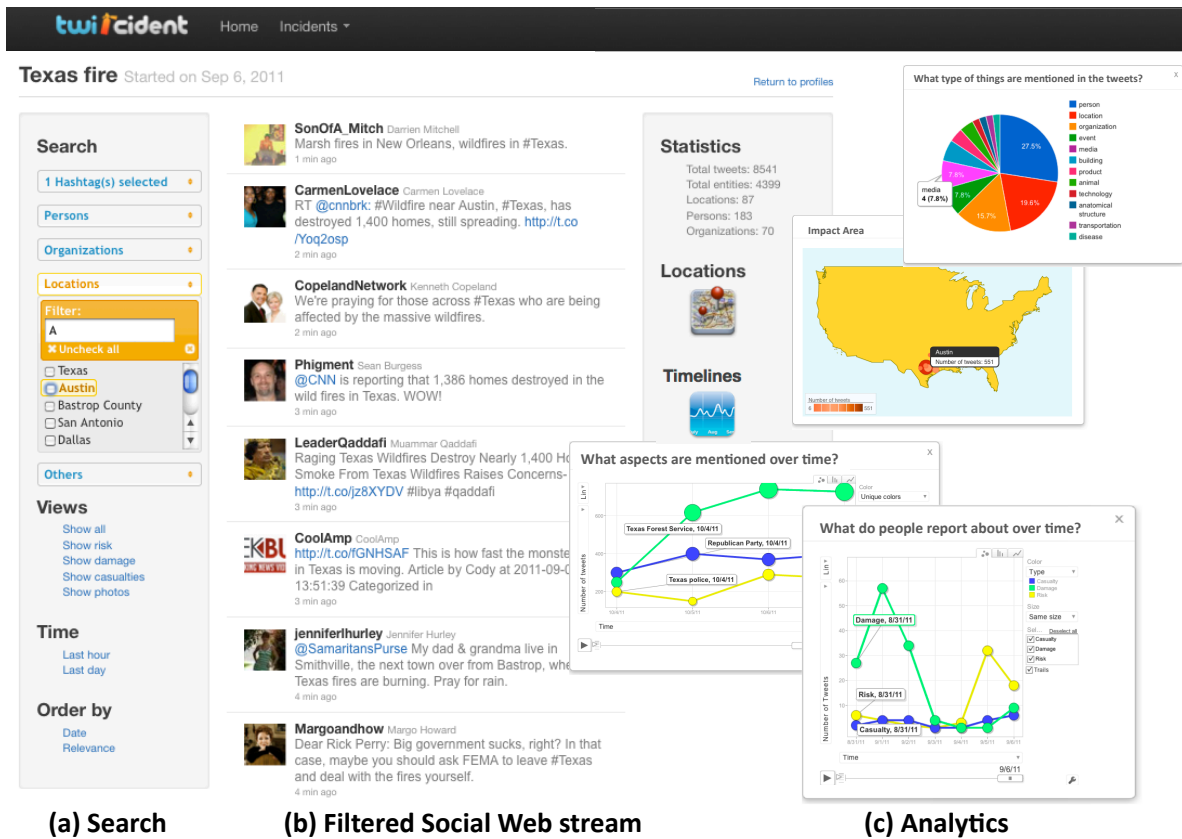
**NER.** The named entity recognition (NER) module assembles four different services<sup>2</sup> for detecting entities such as persons, locations or organizations that are mentioned in tweets. As these services only function for English texts, Twitcident translates non-English tweets to English<sup>3</sup>. The

<sup>1</sup><http://twitpic.com> and <http://twitvid.com>

<sup>2</sup><http://dbpedia.org/spotlight>, <http://alchemyapi.com>, <http://openalais.com>, <http://zemanta.com>

<sup>3</sup>Language detection:

<http://code.google.com/p/language-detection/>  
Translation: <http://code.google.com/apis/language-translate/overview.html>



(a) Search

(b) Filtered Social Web stream

(c) Analytics

Figure 3: Screenshot of the Twitcident system: (a) search and filtering functionality to explore and retrieve particular Twitter messages, (b) messages that are related to the given incident (here: fires in Texas) and match the given query of the user and (c) realtime analytics of the matching messages.

extracted entities are mapped to concepts in DBpedia<sup>4</sup>, the RDF representation of Wikipedia, and the type of an entity is utilized to specify the attribute of the corresponding attribute-value pair. For example, given a Twitter message such as “#txfire is approaching Austin, 50 houses destroyed <http://bit.ly/3r6fgt>”, the NER module allows for detecting “(location, dbpedia:Austin\_Texas)”.

**Classification.** Twitcident classifies the content of Twitter messages into reports about casualties, damages or risks and also categorizes the type of experience that is reported in a tweet, e.g. whether the publisher of a tweet is seeing, feeling, hearing or smelling something. The classification is done by means of hand-crafted rules that operate on both the attribute-value pairs and the plain words that are mentioned in a tweet.

**Linkage.** By following links that are posted within messages, Twitcident further contextualizes the semantics of a message. Therefore, the semantic enrichment module extracts the main content of the Web resource that is referenced from a tweet using Boilerpipe<sup>5</sup> and processes it via the NER module to further enrich the Twitter message with facet-value pairs that describe its content. For the aforementioned tweet which lists “<http://bit.ly/3r6fgt>”, one may extract additional attribute-value pairs.

**Metadata extraction.** Twitcident also collects additional metadata about Twitter messages such as background

information about the publisher of a tweet; for example, the profile picture, number of followers, number of tweets published during the incident or the location of the user when publishing her tweets. Such provenance data may support end-users in assessing the trustworthiness of a tweet and is moreover exploited by the Twitcident system when tweets that match the current query are ranked by relevance [5].

### 2.3.4 Filtering


The goal of the filtering step is to identify those tweets that are relevant to an incident. Twitcident features two core filtering strategies: (i) keyword-based filtering and (ii) semantic filtering.

**Keyword-based Filtering.** The keyword-based filtering strategy transforms an incident profile into a keyword query. The query is evaluated using language modeling with relevance model RM2 [3] to decide whether a tweet is relevant or not. Tweets with less than 100 characters and tweets with words that contain a single letter three or more times in sequence (e.g., “ooooooooo”) are removed independently from their relevance score.

**Semantic Filtering.** The semantic filtering operates on the semantically enriched Twitter messages. Both tweets and incident profiles are represented in a vector space model where each dimension corresponds to a certain attribute-value pair. Jaccard similarity between an incident profile and a tweet in vector representation is applied to decide whether a tweet is relevant for a given incident or not.

<sup>4</sup><http://dbpedia.org>

<sup>5</sup><http://code.google.com/p/boilerpipe/>



**Figure 4: Results of the (a) filtering and (b) faceted search evaluation.**

## 2.4 Faceted Search and Analytics

Incident detection, incident profiling, media aggregation, semantic enrichment and filtering are automatic processes that deliver information about an incident as reported by people on the Social Web. For enabling users to explore information in the filtered Social Web streams, Twitcident provides faceted search and analytics functionality.

### 2.4.1 Faceted Search Strategies

The faceted search functionality of the Twitcident framework allows users to further filter incident-related messages by selecting attribute-value pairs that should be featured by the retrieved messages (see Figure 3(a)). A faceted query may consist of several facet-value pairs. Only those tweets that match all the facet-value constraints will be returned to the user. Twitcident moreover provides different strategies for recommending faceted queries ranging from a frequency-based approach, which ranks the attribute-value pairs according to their occurrence frequency in the current result set, to strategies that adapt the ranking to the temporal context or to the personal preferences of the user. These strategies are described in [1].

The ranking of the tweets that match a query is a research problem of its own. Twitcident ranks the matching tweets according to their (i) creation time or (ii) relevance. The relevance is computed by exploiting various features including provenance information such as the authority score of the user who published a tweet [5].

### 2.4.2 Realtime Analytics

Based on the semantic enrichment, the Twitcident framework provides functionality to analyze the current Social Web stream about an incident. Figure 3(c) shows some of the graphical gadgets that are delivered to the users such as the evolution of topics over time or the geographical impact area of an incident. Twitcident exploits the incident and tweet profiles to generate these diagrams. For example, the impact area of an incident is deduced from the geographical location of Twitter messages that report about experiences of users, e.g. in which people state that they see, hear or smell something. The analytical tools adapt furthermore to the current context of a user: if a user further filters the Social Web stream by means of faceted search then the diagrams summarize and visualize only that fraction of the information that matches the filter.

## 3. EVALUATION

With Twitcident, we propose solutions to two fundamental research challenges: (a) the automated filtering of relevant information from Social Web streams and (b) search within Social Web streams. We evaluated our framework

regarding these two challenges based on the TREC Twitter corpus that consists of approximately 16 million tweets. The corpus also provides relevance judgements for 50 different topics.

As depicted in Figure 4(a), the semantic filtering clearly outperforms the keyword-based filtering regarding mean average precision (MAP), precision@k and recall. For example, the semantic filtering performs—with respect to MAP, P@10 and P@30—more than twice as good as the baseline and regarding recall it improves the filtering performance by 41.8%.

On top of the filtered Twitter streams, we also measured the quality of faceted search and keyword-based search. The setup of our evaluation is described in [1]. The results indicate that the faceted search functionality is beneficial for exploring information in Twitter streams and improves the search performance—measured by means of the mean reciprocal rank of the first relevant tweet—regarding keyword-based search significantly (see Figure 4(b)).

## 4. CONCLUSIONS

In this paper, we presented Twitcident, a framework for filtering, searching and analyzing information about incidents that people publish in their Social Web streams. Triggered by an incident detection module that monitors emergency broadcasting services, our framework automatically collects and filters relevant information from Twitter. It enriches the semantics of Twitter messages to adapt and improve the incident profiling and filtering over time. Semantic enrichment is also the foundation for faceted search and realtime analytics provided by the Twitcident framework. Our evaluation shows that semantic enrichment boosts the performance of both the filtering of Twitter messages for a given incident and improves search for relevant information about an incident within the filtered messages significantly.

**Acknowledgements.** This work is co-funded by the EU FP7 project ImREAL (<http://imreal-project.eu>).

## 5. REFERENCES

- [1] F. Abel, I. Celic, G.-J. Houben, and P. Siehndel. Leveraging the Semantics of Tweets for Adaptive Faceted Search on Twitter. In *Proc. of ISWC*, pages 1–17, 2011. Springer.
- [2] D. Gaffney. #iranelection: quantifying online activism. In *Proc. of Web Science*, 2010. ACM.
- [3] V. Lavrenko and W. B. Croft. Relevance based language models. In *Proc. of SIGIR*, pages 120–127, 2001.
- [4] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proc. of WWW*, pages 851–860, 2010. ACM.
- [5] R. Stronkman. Exploiting Twitter to fulfill information needs during incidents. Master thesis, TU Delft, 2011. <http://wis.ewi.tudelft.nl/twitcident/thesis.pdf>.
- [6] J. Teevan, D. Ramage, and M. R. Morris. #TwitterSearch: a comparison of microblog search and web search. In *Proc. of WSDM*, pages 35–44, 2011. ACM.