December 2019

Unified information processing for detecting, triaging, and responding to crisis events

Mor Schlesinger

Jeremy Moskovich

Unified information processing for detecting, triaging, and responding to crisis events

ABSTRACT

Timely detection of dynamically unfolding crisis events is challenging. Further, determining which information is accurate and whether reports represent disinformation attacks is challenging. Parsing available data and transforming it into complete information that can be presented to end users may also add overhead to crisis response processes. This disclosure describes a system to detect, triage, and respond to a variety of crisis events. System performance may improve over time, e.g., by evaluating prior events and corresponding output from the system using machine learning techniques.

KEYWORDS

- Crisis event
- Crisis detection
- Crisis response
- Crisis alert
- Information ingestion
- Information clustering
- SOS

BACKGROUND

Crisis events, such as natural disasters, accidents, terrorist attacks, etc. occur without prior notice and often unfold dynamically at a fast pace. Therefore, timely detection of such events is challenging. Moreover, information about such events spreads via a variety of information sources, such as news articles, official reports, social media posts, etc., especially in the early stages of the event. This may lead to delays in accurately identifying event occurrences, and challenges in determining best sources and accuracy of reports, reconciling conflicting reports, and filtering out disinformation. Additionally, data pertaining to crisis events varies in format which can often be rigidly defined by the underlying service that provides the data. Parsing the data and transforming it into information that can be presented to end users may therefore also add overhead to crisis response processes.

DESCRIPTION

This disclosure describes a system to detect, triage, and respond to a variety of crisis events. System response improves over time, e.g., by evaluating prior events and corresponding output from the system using machine learning techniques.

The system is designed to be embedded in the larger ecosystem of crisis response efforts that includes relevant authorities, response units, partners, researchers, content producers, and ordinary citizens. The goals of the system include providing support for people in times of crisis by provision of relevant high quality and trustworthy information as well as enabling people to generate helpful information themselves and take an active part in helping others affected by the crisis.

The system serves as a gateway that receives and processes information from a variety of sources, each of which can individually provide only partial and/or unverified information about an event. The processing involves combining data from disparate sources and cross-referencing the input information to develop an awareness of the crisis situation and determine its severity. Cross-referencing can be applied to confirm unverified information from a single source by checking it against relevant related information from other sources. For instance, an external report of a crisis occurring at a location with a contemporaneous spike in crisis-related public social media posts coming from the location of the event can boost confidence in the reliability of the information. Clustering is applied to information obtained from different sources about different events to derive more complete information about the crisis based on combining partial data across disparate sources. The situational awareness and corresponding outcomes are continually updated as the crisis develops and fresh information is obtained from one or more of the sources.

Current awareness of the crisis situation is used to suggest appropriate response measures such as issuing alerts, paging appropriate personnel, deploying rescue teams, etc. Suggested measures can involve disseminating appropriate information across various platforms and modes, such as text messages, emails, website alerts, street signs, etc.

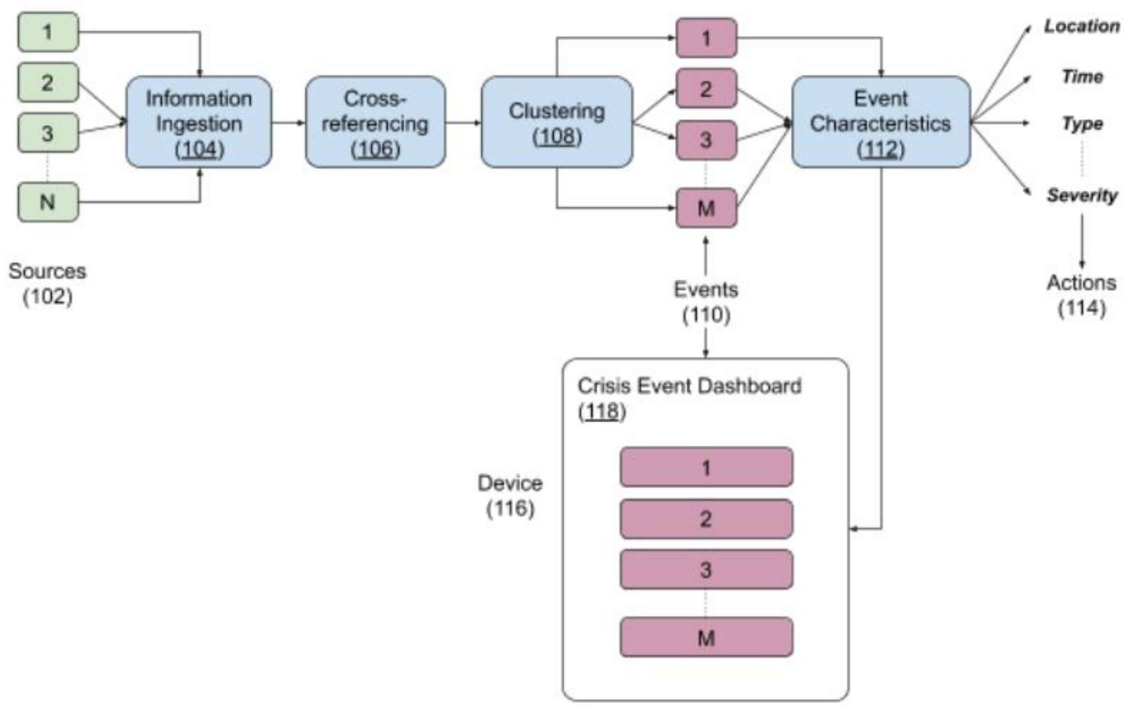

Fig. 1: Architecture of the crisis event handling system

Fig. 1 shows the various components of the described system as embedded in an overall crisis response ecosystem. An information ingestion module (104) aggregates information from a variety of sources (102). The aggregated information is processed by a cross-referencing module (106) to combine partial data and to verify unverified sources by checking across inputs. Next, the information is collected and categorized into discrete crisis events (110) by a clustering module (108).

Each event is further analyzed by an event characteristics extraction module (112) to determine various factors associated with the event, such as event location, start time, type, severity, etc. These factors determine the appropriate suggested actions (114) for responding to the event as it unfolds. For end users, the system can serve as the backend for a crisis event dashboard (118) available on a user's device (116). The dashboard provides users the ability to receive and provide relevant information about a crisis situation.

The ingested information can be obtained from various internal information sources of an organization that implements the system as well as from external partners. Where the information source is data from an individual user (e.g., location, social media post, device information, etc.), the corresponding user may be provided with options to enable or disable access to such data, choose the specific data that is provided, anonymize the data, and limit the use of such data to the purpose of crisis detection. A variety of information sources that can be used by the system and can include:

- Validated information from sources such as NC4 (https://www.nc4.com/).
- Public sources of natural and weather events such as the United States Geological Survey (USGS), National Oceanic and Atmospheric Administration (NOAA), National Weather Service, Japan Meteorological Agency (JMA), etc.
- Breaking news feeds.
- Public news feeds from relevant authorities.

- Web search queries and related trends, e.g., spikes in queries related to an event.
- News articles and social media posts, e.g., spikes in articles and/or posts related to the event.
- Locations and movements of people in specific areas, e.g., to identify severity, likely number of affected users, and evacuation areas.
- Information from mobile devices, e.g., battery level, connectivity status; etc., obtained with user permission.

The data obtained from each of the sources need not conform to a specific format, thus allowing the system to handle a diversity of data sources and making it possible to add information sources with minimal effort. Further, the data may contain partial information. Further still, some sources of data may be designated as verified and/or trusted, while others have no such verification. Data fusion is used to generate a complete picture from multiple pieces of partial information and cross-referencing across sources is performed to increase confidence in the veracity of the information, thus providing quality assurance regarding the data without requiring extensive manual evaluation.

Clustering mechanisms are used to separate information about different events, thus making it possible for the system to simultaneously support multiple ongoing crises at various stages of development. Further, relevant characteristics of each of the events are extracted and updated separately, providing comprehensive and reliable semantic understanding of the present state of the crisis situation.

The situational understanding is in turn applied to suggest one or more response actions based on the current state. Some examples of possible suggested actions include:

- Sending or updating a text alert.

- Paging an appropriate person for handling emergencies.
- Enriching the information as a crisis develops and new information becomes available.
- Inserting appropriate information across relevant websites and/or apps.

The actions can be automated with user permission and/or be controlled manually. For instance, an automated action can involve sending text messages to those that are likely to be affected, while a manual action can involve informing a human operator to examine certain information in more detail. The collected operational data and corresponding human decisions can serve as labeled training data for a machine learning classifier that can be leveraged to improve the level of automation of crisis response detection and response actions. The types, numbers, and sequences of suggested actions vary based on the determined situational characteristics and are updated as the situation unfolds.

The crisis event dashboard can be implemented as a website and/or a dedicated application for user devices. The dashboard can be used for viewing dynamically updated event related information and receiving alerts and notifications pertaining to events of interest. In addition, the dashboard can provide mechanisms for end users to serve as information sources by directly providing information known to them about an event. Further, the output of the system can be accessed via an application programming interface (API), thus enabling other parties to incorporate the functionality within their own systems.

CONCLUSION

This disclosure describes a semi-automatic system to detect, triage, and respond to a variety of crisis events. The system serves as a gateway that receives and processes information obtained from a variety of sources, each of which can individually provide only partial and/or unverified information about an event, verified or not. The processing involves cross-referencing the input information from disparate sources in order to develop an awareness of the crisis situation and determine its severity. Clustering is performed to collect information about different events and derive more complete information about the crisis based on combining partial data across disparate sources. The system can be embedded in the larger ecosystem of crisis response efforts involving relevant authorities, response units, partners, researchers, content producers, and ordinary citizens. Over time, system performance may improve, e.g., by evaluating prior events and corresponding output from the system using machine learning techniques.