Mining the Disaster Hotspots – Situation-Adaptive Crowd Knowledge Extraction for Crisis Management

Andrea Salfinger, Wieland Schwinger, Werner Retschitzegger
Dept. of Cooperative Information Systems
Johannes Kepler University Linz
Altenbergerstr. 69
4040 Linz, Austria
{andrea.salfinger, wieland.schwinger, werner.retschitzegger}@cis.jku.at

Birgit Pröll
Inst. for Application Oriented Knowledge Processing
Johannes Kepler University Linz
Altenbergerstr. 69
4040 Linz, Austria
bproell@faw.jku.at

***Abstract*—When disaster strikes, emergency professionals rapidly need to gain *Situation Awareness* (SAW) on the unfolding crisis situation, thus need to determine what has happened and where help and resources are needed. Nowadays, platforms like Twitter are used as real-time communication hub for sharing such information, like humans' on-site observations, advice and requests, and thus can serve as a network of "human sensors" for retrieving information on crisis situations. Recently, so-called crowd-sensing systems for crisis management have started to utilize these networks for harvesting crisis-related social media content. However, up to now these mainly support their human operators in the visual analysis of retrieved messages only and do not aim at the automated extraction and fusion of semantically-grounded descriptions of the underlying real-world crisis events from these textual contents, such as providing structured descriptions of the types and locations of reported damage. This hampers further computational situation assessment, such as providing overall description of the on-going crisis situation, its associated consequences and required response actions. Consequently, this lack of semantically-grounded situational context does not allow to fully implement situation-adaptive crowd knowledge extraction, meaning the system can utilize already established (crowd) knowledge to correspondingly adapt its crowd-sensing and knowledge extraction process alongside the monitored situation, to keep pace with the underlying real-world incidents. In the light of this, in the present paper, we illustrate the realization of a situation-adaptive crowd-sensing and knowledge extraction system by introducing our $crowd^{SA}$ prototype, and examine its potential in a case study on a real-world Twitter crisis data set.**

I. Introduction

Crowd-Sensing for Crisis Situation Awareness. When disaster strikes, emergency professionals rapidly need to gain *Situation Awareness* (SAW) on the unfolding crisis situation, thus need to determine what has happened and where help and resources are needed. Nowadays, such information are often available first on social media. Platforms like Twitter have become a popular real-time communication hub for affected populations to share observations, advice and requests due to their ubiquitous availability on mobile devices, as examined in case studies across different crisis events [1]. To aid emergency professionals in the timely identification of these information within the plethora of social media chatter, dedicated *crowd-sensing* systems have been proposed, which support their human operators in the retrieval and analysis of crisis-related social media content (cf. surveys in [2], [3]).

Lack of Semantic Grounding. However, as these surveys reveal, currently available systems mainly support their human operators in the visual analysis of retrieved messages. These do not assist emergency managers in extracting and fusing semantically-grounded descriptions of the underlying real-world crisis events from these textual contents, such as providing structured descriptions of the types and locations of reported damage, like flooded areas and blocked roads, for further computational Information Fusion (IF) and situation assessment (SA) to provide the overall description of the on-going crisis situation, its associated consequences and required response actions. This lack of semantically-grounded situational context does not allow to fully implement situation-adaptive crowd-sensing and knowledge extraction, i.e., enable the system to use already established (crowd) knowledge to correspondingly refine its crowd-sensing and knowledge extraction process alongside the monitored situation, such as querying more information on affected areas, or using already established information on the on-going crisis event to aid the interpretation of social media messages, which are frequently lacking contextual information due to their required terseness.

Contributions. As a first step towards overcoming these limitations, we present the realization of a *situation-adaptive* crowd-sensing and knowledge extraction approach in $crowd^{SA}$, an IF architecture for crisis management [4]–[6]. First, we outline how $crowd^{SA}$ extracts and aggregates structured object descriptions from textual tweet content, yielding semantic descriptions of current crisis hotspots, which we assess in a proof-of-concept case study on a real-world Twitter crisis data set. Second, we demonstrate means how the system further on can utilize already established knowledge on the detected situational context to adapt its crowd-sensing and knowledge extraction alongside the detected events, to dynamically acquire further crisis knowledge by learning from the crowd.

Structure of the Paper. In Sec. II, we provide a brief recap of the devised $crowd^{SA}$ system and outline the scope addressed in the present work. Sec. III provides an implementation overview and initial results of our prototype on a real-world crisis data set obtained from Twitter. We contrast our approach to related work in Sec. IV, before concluding in Sec. V.

II. Supporting Crowd-sensed Situation Awareness for Crisis Management with $crowd^{SA}$

To aid emergency managers in maintaining SAW on evolving crisis situations, we proposed the architecture of $crowd^{SA}$, a situation-adaptive IF system which incorporates social media as additional

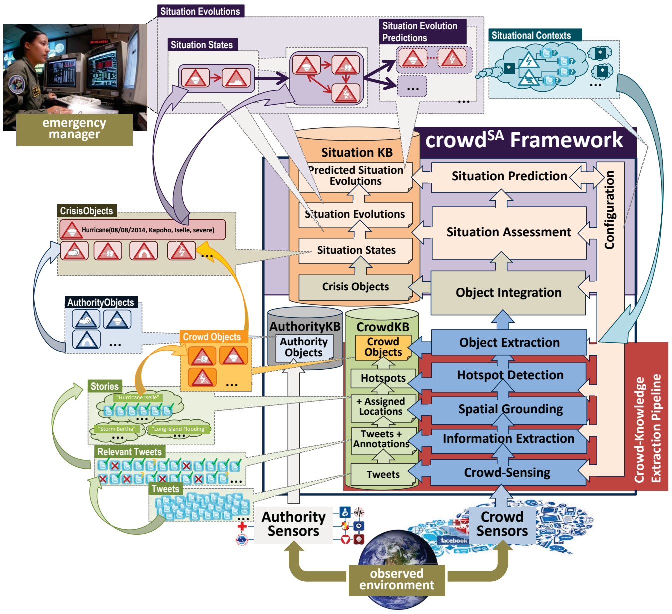

Figure 1: Overview of $crowd^{SA}$'s processing architecture.

data source complementing conventional data sources like physical sensors [4]–[6]. $crowd^{SA}$ aims at automatically detecting and tracking crisis situations by means of rule-based SA, i.e., matching descriptions of the monitored real-world objects against templates characterizing situations of interest. Thus, emergency managers are enabled to express their information need and consequently get timely alerts from the system whenever object and event constellations matching a situation template are detected. In order to obtain useable information from its crowd-sensing adapters tapping social media channels, $crowd^{SA}$ needs to provide the following functional blocks (cf. Fig. 1): Monitoring social media for messages containing potentially crisis-relevant information (`Crowd-Sensing`), extracting relevant information nuggets from these messages individually (`Information Extraction`), mapping these to their corresponding real-world location (`Spatial Grounding`), inferring the underlying real-world events described in these messages by aggregating multiple observations, i.e., fusing information from multiple messages to a single, coherent description of the monitored real-world event (`Hotspot Detection`), and subsequently determining the object-level crisis information within the determined hotspots (`Object Extraction`), i.e., so-called *Crowd Objects* (such as structured descriptions of the disaster agent, damaged buildings and affected individuals). These in turn can serve as input for its rule-based SA. In the present work, we will introduce the realization of this functionality within $crowd^{SA}$'s crowd-knowledge extraction pipeline, which fuses crisis-relevant information obtained from social media (currently focused at Twitter) to event-level information grounded in a domain ontology. We specifically examine how semantic Information Extraction (IE) on the messages' textual content enables a more precise description of the encountered situation(s) and supports situation-adaptive IF. As an alternative to commonly employed machine-learned based IE and its domain adaptation problems requiring continuous manual curation (cf. Sec. IV), we examine the applicability of utilizing a knowledge-based IE approach for extracting semantically-grounded crisis information from tweets and determining an overall, spatio-temporal-thematic event summarization. Whereas this demands an initial configuration effort with respect to (w.r.t.) specifying the general domain knowledge (in terms of an appropriate domain ontology describing entities of interest) and suitable extraction rules, further crisis-specific information can be obtained automatically: At runtime, new (instance-level) knowledge is acquired by learning from the crowd, to circumvent the knowledge acquisition problem commonly associated with such a top-down, declarative approach. Thus, to enhance the self-adaptivity of the system, we furthermore contribute means for a situation-adapative crowd-sensing and extraction refinement by introducing situational feedback loops, which enable the system to optimize its processing towards the encountered situation, thus distinguishing our IE approach from related approaches involving knowledge-based IE and ontology population, like [7]–[9].

III. Design and Implementation

In the present section, we discuss the realization of the $crowd^{SA}$ prototype's crowd-knowledge extraction pipeline, notably the following functional blocks (cf. Sec. II and Fig. 1):

1) *Information Extraction* — Analyzing the retrieved messages' textual content to extract crisis-relevant information.
2) *Spatial Grounding* — Determining the real-world locations the extracted information refers to.
3) *Hotspot Detection* — Detecting events based on analyzing densely covered geographical regions (so-called *hotspots*), and inferring the general event context by aggregating information extracted from tweets assigned to locations in the corresponding hotspot.

We examine how such a crowd-knowledge extraction pipeline can be assembled by adapting and integrating existing libraries and frameworks suitable for addressing each of these functional blocks. Furthermore, to enable *situation-adaptivity*, we describe how results obtained from one phase can be used to implement a *situational feedback loop* to refine the processing of its preceding phases and allow for dynamic knowledge acquisition. Fig. 2 presents an overview of employed components, which we will discuss based on the following examples from real-world crisis data sets recorded from Twitter (i.e., the *Iselle data set* on hurricanes Iselle and Julio affecting the Hawaiian islands in Aug. 2014[1], and the *Erica data set* on tropical storm Erica devastating Dominica in Aug. 2015[2]):

Tweet 1
@days_go_by29 hi Katie! No, #Iselle passed Oahu with mostly wind, rain, surf. The Big Island had some damage but luckily no casualties! Now,

Tweet 2
RT @reedtimmerTVN: 4 foot diameter Albezia Tree down across road near Pahoa, Hawaii from Tropical Storm Iselle http://[...]

Tweet 3
RT @nbc6: Threat of floods over for S. Fla., remnants of #Erika dissipate as Hurricane #Fred forms: http://[...] http://[...]

Tweet 4
BREAKING: Police in Dominica say at least 4 dead as Tropical Storm Erika hits Caribbean island... #PrayForDominica

A. Information Extraction

Motivation. As examined in [1], different types of information valuable for emergency responders can be obtained from Twitter. Based on their findings and our own empirical data analysis of recorded crisis data sets, we collated the following types of crisis information $crowd^{SA}$ should extract:

- severe weather formations (such as *hurricanes*, *typhoons*, *thunderstorms*, *tornadoes*) and their evolutionary states (a hurricane instance may *form*, *make landfall*, *turn*, *strengthen*, *weaken*, *reform*, *dissipate*), which determine the crisis management (CM) phase
- their weather-related consequences (e.g., *torrential rain*, *flooding*, *spring floods*, *storm surges*, *erosion*, *mudslides*, *high winds*)
- different types of entailed infrastructure damage (e.g., *downed trees*, *blocked roads*, *damaged buildings*, *power outages*), the states

[1]Recorded from the Twitter Streaming API by tracking the following keywords: Hurricane, #HurricaneIselle, #HurricanePrep, #updatehurricaneiselle, #hiwx, #HIGov, Iselle, #Genevieve, #Iselle, #Julio, #HIWX, #HIWx

[2]Recorded from the Twitter Streaming API by tracking the following keywords: #Erika, #TropicalStormErika, #Dominica, #PrayForDominica, #TSErika, #WestIndies, Roseau, #KeyWest, #flkeys, #FLwx, #florida

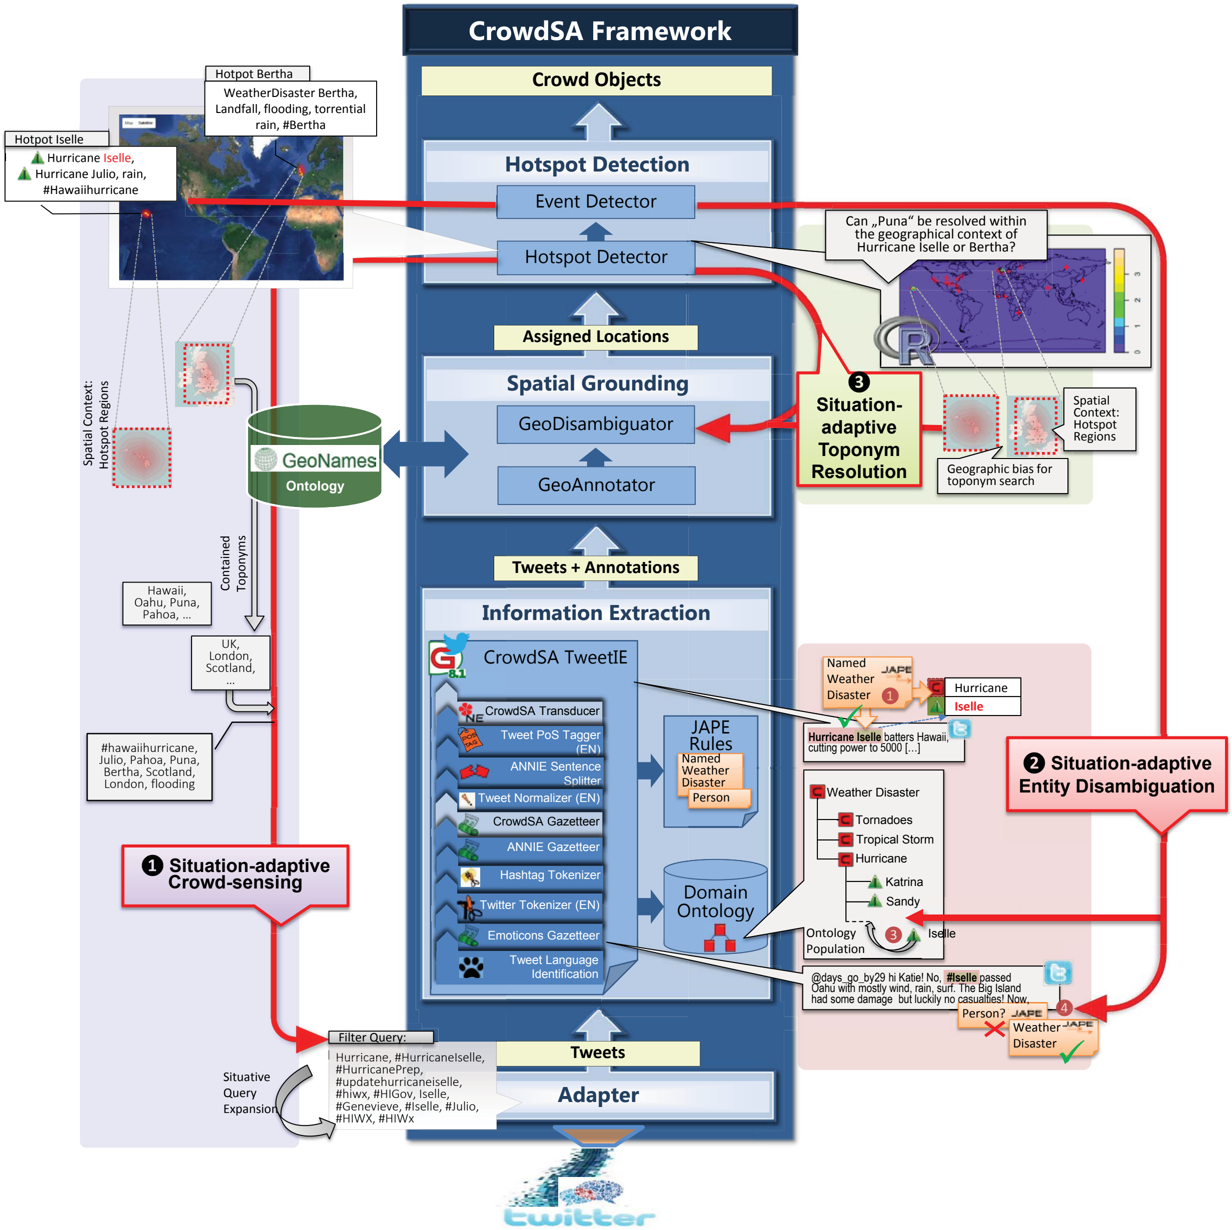

Figure 2: The interplay of the different processing components.

(*open*, *closed*, *blocked*, *reopened*, *evacuated*) of key infrastructure elements (*bridges*, *roads*, *schools* etc.)

- indication of needed or provided resources, such as emergency supplies (e.g., *food*, *free meals*, *drinking water*, *tarps*, *batteries*)
- information regarding affected individuals (e.g., *injured people*, *casualties*)

Whereas such information is also contained in our example Tweets 1-4, it also becomes apparent that extracting the corresponding factual information entails further requirements, due to the complex compositional semantics encountered in natural speech. Since factual information may be embedded in linguistic constructs such *negations*, *questions*, and references to *past*, *present* and *projected events*, this complicates the task of extracting the semantically coherent information transported in this message. For instance if we aim to extract the *DisasterEffects* contained in Tweet 1, simply extracting the word *casualties* as an instance of the concept (i.e., ontology *Class*) *DisasterEffect* obviously would deliver the wrong information (the correctly extracted *DisasterEffect* information would even correspond to the opposite, i.e., *"no casualties"*). Especially regarding the CM domain, confirmation of what has *not* happened, i.e., stated absence of disaster effects, is of equal importance as the acknowledged presence of disaster consequences and the key constituent of crisis SAW in order to adequately organize and prioritize emergency response actions. Considering such surrounding textual context is furthermore required to determine the CM phase, i.e., whether the situation is in the disaster preparation phase or the disaster's aftermath. This is essential information for crisis SAW indispensable for the adequate *comprehension* and *projection* of the crisis situation, as for instance provided in Tweet 3, which contains projection information ("threat of floods"), and furthermore describes the evolutionary state of two different disaster agents (one weather disaster is dissipating, i.e., corresponds to the crisis aftermath, whereas another one is in its formation, corresponding to the disaster preparation phase). Thus, a *semantic analysis* of the tweet texts is key for extracting information enhancing (instead of confounding) crisis SAW. We thus examine means for proper extraction and semantic-mapping of a crisis situation's evolutionary phase (e.g., *hurricane prior landfall* vs. *has made landfall*), representing the basis for tracking the crisis situation's evolution according to situation evolution models [10], [11].

Realization. For realizing $crowd^{SA}$'s IE, we base upon TwitIE [12], a plug-in for Natural Language Processing (NLP) on Twitter content for the popular NLP Framework GATE [13]. TwitIE already provides a *structural* analysis of textual tweet content, i.e., determines grammatical structure, and detects common types of Named Entities (NEs), such as person names, organizations and locations, and other domain-independent information types (e.g., dates, addresses, e-mails, emoticons), i.e., performs a lexical analysis. Since TwitIE is composed in a modular fashion, as it consists of a declaratively specified IE pipeline composed of individual components, so-called `Processing Resources` (PRs), such as `PRs` for language detection, tokenization, Gazetteer lookups, normalization, sentence splitting, part-of-speech tagging (PoS), and the splitting of multi-word hashtags into single-word tokens, its functionality can easily be extended by adding suitably configured `PRs`. In order to create an IE pipeline for extracting the CM-relevant information, we utilized existing `PRs` from the GATE library, which we provided with a crisis-specific configuration (as shown in Fig. 2). We devised a custom `crisis domain ontology` modeling the CM-relevant information types listed previously, which was thus based upon own empirical data analysis as well as findings presented in other work [1], [14], [15]. To detect these concepts in free-form text, we attached suitable lexicalizations (i.e., natural language descriptors for a specific concept) as class properties for each concept, and supplied custom gazetteer lists (i.e., word lists or dictionaries on a specific concept, such as a gazetteer on country names or currency units). These information can be made available to the TwitIE pipeline by adding a custom `Gazetteer PR` (which is configured with the devised gazetteer lists, currently we use 14 lists) and `OntoGazetteer PR` (which compiles custom gazetteer lists from an ontology's concepts and their lexicalizations). In order to extract matching text spans from free-form text, we specify corresponding extraction rules on sought-after information types, in form of so-called JAPE ("Java Annotation Patterns Engine") rules, GATE's annotation pattern description rule language. These rules are supplied to a `JAPE Transducer PR` (`CrowdSA Transducer` in Fig. 2), which compiles them into Finite-State-Transducers for matching the specified patterns in texts, and thus comprises the core IE logic. Simple rules can be based on *Gazetteer* look-ups (e.g., for annotating the text span "power outage" with the ontological class *PowerOutage*). However, such a structural and lexical analysis, as shown below for Tweet 1[3], does not allow to address composite semantics:

Tweet 1, Structurally & Lexically Analyzed & Annotated

@days_go_by29 hi Katie! No, #Iselle passed **[Oahu]**$_{Location}$ with mostly **[wind]**$_{DisasterEffect}$, **[rain]**$_{DisasterEffect}$, **[surf]**$_{DisasterEffect}$. The **[Big Island]**$_{Location}$ had some **[damage]**$_{DisasterEffect}$ but luckily **[no]**$_{Negation}$ **[casualties]**$_{DisasterEffect}$! Now

In order to resolve compositions such as negated concepts (e.g., "no casualties" or "no major damage"), we need to devise a *semantic analysis* by hierarchically combining annotations such as simple (lexicographic) annotations. Therefore, JAPE rules can utilize annotations created by preceding `PRs` and can be cascaded to detect composite patterns consisting of nested annotations. In conjunction with a `Domain Ontology` providing the generic, high-level concepts, this hierarchical composition of annotations enables us to codify domain information of interest in a high-level, abstract fashion, thereby only requiring a limited set of generic rules, such as for the detection of negated concepts, forecast events, or the association of specific damage types to specific disasters. Thus, we developed an initial set of 28 CM-specific JAPE rules for our proof-of-concept implementation, which base upon the basic (domain-independent) JAPE rules provided by TwitIE to realize our required semantic extraction of CM-relevant content, which analyze composites of annotations and thus are capable of producing the following annotations:

Tweet 1, Semantically Annotated

@days_go_by29 hi Katie! No, #Iselle passed **[Oahu]**$_{Location}$ with mostly **[wind]**$_{DisasterEffect}$, **[rain]**$_{DisasterEffect}$, **[surf]**$_{DisasterEffect}$. The **[Big Island]**$_{Location}$ had some **[damage]**$_{DisasterEffect}$ but luckily

$\Big[$ **[no]**$_{Negation}$ **[casualties]**$_{DisasterEffect}$ $\Big]_{DisasterEffect}$! Now,

Thus, the sequence of a negation expression followed by a *DisasterEffect* annotation is combined to yield a negated *DisasterEffect*, i.e., correctly results in the information "no casualties". For instance, for detecting a weather disaster's evolutionary phase (which we term *LifeCyclePhase*, e.g., *forms* vs. *dissipates*), as in Tweet 3, Listing 1 depicts how the corresponding domain concepts can be associated, i.e., a *WeatherDisaster* annotation followed by a *LifeCyclePhrase*, whereas Listing 2 shows a (simplified) example on the extraction of forecasts (e.g., indicated by phrases such as "threat of", "possible"), by detecting the sequence of a *Forecast* annotation and a *WeatherDisaster Consequence* annotation. Concluding, we need to note that whereas our lexicalizations and rules are currently targeted on English texts, this approach can be extended to other languages by enriching the ontology with corresponding lexicalizations and providing suitable IE rules, thereby enabling multi-lingual IE.

Listing 1: JAPE Rule Ex. 1

```
Rule: LifecycleDetection
(
 (
  {WeatherDisaster}
 ):disaster
 ({Token})[0,3]
 (
  {LifecyclePhrase}
 ):phase
):tag
-->
    { /* create annotations
          */}
```

Listing 2: JAPE Rule Ex. 2

```
Rule: ForecastEvents
(
 ({Lookup.majorType ==
                forecast})
 ({Token})[0,2]
 ({Lookup.majorType ==
                disaster,
   Lookup.minorType ==
        consequence}
 ):typeOfConsequence
):tag
 --> { /* create annotations
        */ }
```

Enabling Situation-Adaptivity. However, solely basing upon an a-priori specified `Domain Ontology` and `Rule`-base would entail the disadvantage of a static system that is not reactive towards novel, previously unseen events. Therefore, we need to equip the system with capabilities to dynamically derive new knowledge based on the extracted data, for which we suggest an approach that combines knowledge-based IE with data-driven aggregation strategies: Whereas the general domain knowledge of CM-relevant information can be specified a-priori, mainly corresponding to our `Domain Ontology` *classes*, each crisis' characteristics are *unique*, often corresponding to new *instance-level* information. In the tweets stated above, for instance, one may note that each hurricane event (i.e., *hurricane instance*) is given a proper name, a common practice by weather agencies, allowing to distinguish these different disaster event instances. Thus, hurricanes and related weather phenomena can be considered as *Named Entities* (NE). On the other hand, also the *omission of contextual information* becomes apparent, which is a crucial problem for the interpretation of social media content due to its length limitations: Given the full information, Tweet 3 would actually read as "... remnants of [tropical storm] #Erika...". Due to this omission of relevant context information, an IE component (lacking this instance level information) would thus interpret "Erika" as first name and create a *Person* annotation. If we were not able to dynamically detect new instances, however, our system would miss a large portion of relevant information, notably in cases were people were only referring to the disaster event's name, and omitting the class-level information

[3]We will further on utilize the following formalism for describing *annotations*, i.e., meta-information interpreting the original text's semantics: **[original tweet text]**$_{AnnotationType}$, i.e., the annotation's span is denoted by brackets, its type by the subscript.

"hurricane" (also termed "unmarked" information in literature [14]), which is actually the case in a large fraction of tweets. Therefore, we need to implement a means to dynamically derive such new "instance-level" information from data, and for this we can exploit the fact that we are analyzing large volumes of tweets, which may allow us to reconstruct omitted context from other messages, such as Tweet 4. In this tweet, we can also extract the instance-level information, i.e., an instance of a *TropicalStorm* referred to as *Erika*, requiring a rule for extracting such instance-level information, as shown in Listing 3:

Listing 3: JAPE Rule Ex. 3

```
Rule: WeatherDisaster
(
 ({Lookup.majorType == disaster, Lookup.minorType ==
      weather}
 ):disastertype
 (({Lookup.majorType == person_first} | {Hashtag}
   | {Token.kind == "word", Token.category == "NNP"})?
 ):name
):tag
--> { /* create annotations */ }
```

Tweet 4, Semantically Annotated, Instance Extraction

BREAKING: **[Police]**$_{Organization}$ in **[Dominica]**$_{Location}$ say at least **[4 dead]**$_{DisasterEffect}$ as **[Tropical Storm [Erika]**$_{Name}$ **]**$_{WeatherDisaster}$ hits **[Caribbean island]**$_{Location}$... #PrayFor**[Dominica]**$_{Location}$

If we were to observe a considerable amount of such instance-level information for a specific time window, we can assume this corresponds to a notable real-world event, and populate our ontology with a corresponding instance, i.e., *TropicalStormErika*. To enable such dynamic, crowd-based ontology population, we thus need to provide a situational feedback loop from the aggregation level, which will be explained in the subsequent sections, in order to produce the following *WeatherDisaster*-instance-based annotations:

Tweet 3, Semantically & Situation-Adaptively Annotated

RT @nbc6: **[[Threat of]**$_{Forecast}$ **[floods]**$_{DisasterEffect}$**]**$_{Forecast}$ over for **[S. Fla.]**$_{Location}$, remnants of **[[#Erika]**$_{WeatherDisaster}$ **[dissipate]**$_{Lifecycle}$**]** as **[[Hurricane [#Fred]**$_{Name}$ **]**$_{WeatherDisaster}$ **[forms]**$_{Lifecycle}$**]**: http://t.co/Pv1Sd0bLSc http://t.co/8zyG...")

B. Spatial Grounding

Motivation. Furthermore, no matter what kinds of useful disaster-relevant information have been extracted, ultimately, emergency professionals need to know *where* critical events are happening to coordinate appropriate response actions. Thus, the system actually should not extract the location of the tweet's author, but the location of the event mentioned in the tweet's text (as the two locations may be disparate), i.e., perform *toponym recognition* (i.e., the detection of location names in texts), and *toponym resolution* (i.e., associating a toponym with its corresponding geographical coordinates) [16].

Realization. To perform toponym recognition and toponym resolution, we based upon CLAVIN-NERD[4] ("Cartographic Location And Vicinity INdexer"), a software package for toponym recognition and context-based toponym resolution. However, after inspecting the results, it seemed that CLAVIN's context-based toponym resolution is biased towards higher population numbers and populated places. Although results could be improved by increasing the maximum context window and the maximum number of considered, top-matching toponyms, which, however, decreases performance, we faced false resolutions on our data set particularly crucial w.r.t. the fine-grained location information required for CM, such as

- "Big Island" was resolved to Big Island in Virginia, even when co-occurring with the location "Hawaii", which should allow to disambiguate this toponym to the correct location (the Island of Hawaii, which is commonly referred to as the "Big Island")
- "Puna", a region on the Island of Hawaii, which was resolved to Pune in India
- "Wailea, Maui", was not resolved to Maui, but to the location of Wailea on the Island of Hawaii

To overcome this limitation, our *Spatial Grounding* implementation consists of two components (cf. Fig. 2): The first component, the *GeoAnnotator*, utilizes CLAVIN-NERD for initial toponym resolution. Subsequently, our custom solution, a so-called *GeoDisambiguator*, represents a dedicated toponym resolution component for determining the most specific geo-locations the tweet is assigned to (the so-called `AssignedGeoLocation`), and adapts the results obtained with CLAVIN-NERD, by incorporating the following types of *contextual information*:

- (C.1) "in-tweet-context": disambiguation within a single tweet based on the joint context of all location mentions occurring in this tweet
- (C.2) external context: "between-tweets-context" or "situative context": employs already established information from subsequent processing components (obtained from detected event hotspots)

For incorporating these two types of contextual information, we employ a multi-step approach consisting of an iterative refinement of toponym resolution, which uses ontology-based reasoning on the GeoNames ontology[5]. Regarding (C.1), we disambiguate the toponyms based on the "joint" context of all toponyms encountered in the tweet. The *GeoDisambiguator* queries the GeoNames web service to retrieve the toponyms' ancestry hierarchy, based upon which it constructs ancestry trees (a local tree per tweet is generated), and tries to find a better configuration which corresponds to more specific results. Finally, the leaves of the constructed ancestry tree are selected as *AssignedLocations*. Thus, we note that a tweet may also be assigned to multiple locations, which is also the case if it contains GPS meta-data, which are also added as *AssignedLocations*. Tweet 2 illustrates this principle, which is assigned to Pahoa, corresponding to its most specific location. If multiple locations are on the same hierarchical level, the tweet is assigned to all of these.

Tweet 2, Semantically Annotated

RT @reedtimmerTVN: 4 foot diameter Albezia **[Tree down]**$_{DisasterEffect}$ across **[road]**$_{Infrastructure}$ near **[Pahoa]**$_{Location}$, **[Hawaii]**$_{Location}$ from **[Tropical Storm [Iselle]**$_{Name}$ **]**$_{WeatherDisaster}$ http://[...]

Enabling Situation-Adaptivity. Whereas the *GeoDisambiguator* can correctly resolve the toponym "Puna", if given a tweet with multiple locations mentions (e.g., which also mentions "Hawaii"), the question remains of how to resolve ambiguous topoynms if such contextual information is omitted. How do we resolve a tweet that solely mentions "Puna", without providing a further indication whether it should be attributed to India or Hawaii? Similarly to the IE phase, we can aim at providing the component with a situational context based on using information obtained with subsequent components - thus, this limitation will be tackled with a situational feedback loop from the aggregation level.

C. Hotspot Detection

Motivation. After mapping the extracted information to real-world locations, the system can analyze which regions exhibit considerable coverage, thus corresponding to *hotspot* locations receiving substantial social media attention, which typically indicate large-scale events. Upon extracting these hotspot locations, it can aggregate information extracted from its assigned tweets to identify frequently corroborated information, and thus, infer the underlying events and provide a suitable summarization on these.

[4]https://clavin.bericotechnologies.com/clavin-core/

[5]http://www.geonames.org

Realization. $crowd^{SA}$ identifies these hotspot locations by employing a spatial kernel density estimation on the *AssignedLocations*. Since it operates on the assigned locations, and not only on the tweets' meta-data, it can determine the most-discussed regions, as opposed to approaches operating the locations of the tweets' authors, which are likely to be biased towards densely populated regions. The required spatial aggregation computations are implemented in R [17], which is accessed from the core Java implementation via *Rserve* [18], a binary R server providing R functionality via TCP/IP sockets. The Java-based *Hotspot Detector* component thus opens a connection to the R server, sends the data to the created R session, and invokes an evaluation function performing and returning the density computation, employing the `sp` [19], [20] and `spatstat` [21] packages to perform a fixed-bandwidth kernel estimation of the point pattern intensity of the assigned geolocations using an isotropic Gaussian smoothing kernel. Bandwidth is adjusted by setting the corresponding parameter `adjust` to 0.1, which is multiplied with the value of σ, the standard deviation of the isotropic Gaussian smoothing kernel. Following kernel density estimation, the contour lines of high-density regions are extracted and sent back to the Java-based *Hotspot Detector* component, which aggregates these to the different *Hotspot* locations. In our *Iselle data set*, two major hotspot locations could be determined, as shown in Fig. 3, which seems surprising in the light we only intended to monitor hurricane *Iselle*'s effects on the Hawaiian islands, as tracking the Twitter stream for the general keyword "hurricane" also recorded another hurricane event in a different part of the world, which, however, should be accordingly detected by the system. Determining the hotspot locations, however, only allows to detect locations where presumably some event is happening - to infer the further characteristics of this event (i.e., which type of disaster event is striking, and its effects), a so-called *Event Detector* subsequently aggregates the information extracted from the tweets assigned to each hotspot. Thus, for each hotspot, its extracted concepts and hashtags are accumulated — the most frequent concepts and hashtags per hotspot then are considered as being descriptive for the hotspot, i.e., form the general situational context. This spatially-driven aggregation thus delivers us the spatial extent and semantically-grounded information of the on-going events, the results obtained on our *Iselle data set* are shown in Fig. 3, thus, detects that the hurricanes *Iselle* and *Julio* are affecting the Hawaiian islands, and another hurricane event - which we did not intend to monitor, but which could be solely determined based on the tweets matching our general keywords - named *Bertha* is making landfall in the UK and causing flooding.

D. Incorporating Situative Feedback Loops

The established situative context furthermore can be used to fine-tune the sensing and extraction pipeline towards the detected hotspots. Based on the detected situation, we outline how the derived situative context can inform the processing pipeline. Currently, we have examined the following situational feedback loops:

Situation-adaptive Crowd-sensing. The system can immediately aim at collecting more detailed information on the detected hotspot events, by reconfiguring its adapter to incorporate the determined information: Therefore, we can include a feedback loop to the adapter, which appends keywords derived from the top descriptive concepts and most frequent hashtags to the utilized filter query (e.g., includes the keyword "Bertha"). Furthermore, for CM tasks, it would be highly desirable to receive more fine-grained location information on crucial on-goings (e.g., knowing what types of damage occurred in *Puna* is of more actionable value, than knowing that the *state of Hawaii* is affected). Thus, we seek to increase the fraction of tweets reporting more localized information, by incorporating a query expansion on presumably affected locations: $crowd^{SA}$ thus queries the GeoNames web service to receive all toponyms located within the identified hotspot regions, which can be included in the adapter query. Since this typically results in an extensive list exceeding the number of keywords allowed for Twitter monitoring, we seek to include strategies for selecting promising toponyms for future work.

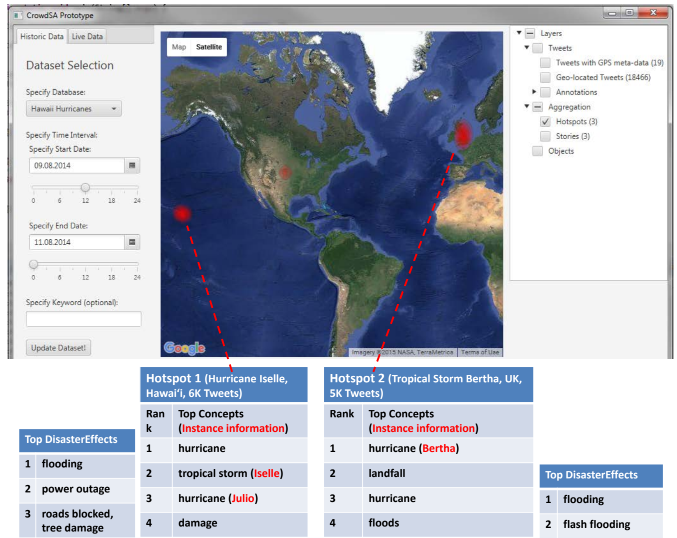

Figure 3: A screenshot of $crowd^{SA}$'s UI, showing the hotspot areas detected by the system. Their corresponding extracted event-level contexts, i.e., their top annotation aggregation results, are shown below the screenshot.

Situation-adaptive Entity Disambiguation. In $crowd^{SA}$'s default configuration, the text "Iselle" would match the concepts "Person" (corresponding to a name) or "Location" (corresponding to a place in Italy). However, in the temporal context of the disaster event, its occurrence in a tweet may likely correspond to a hurricane instance. Therefore, after crowd-based ontology population by extracting a considerable fraction of tweets mentioning a weather disaster instance named Iselle, $crowd^{SA}$ furthermore can utilize a feedback loop to the IE component, which prioritizes the disambiguation of the concept "Iselle" to the weather disaster instance in the corresponding time window, as long as a considerable mentioning of this concept is observed.

Situation-adaptive Toponym Resolution. Furthermore, the geo-fencing performed by hotspot detection informs the *GeoDisambiguator*, which thus can incorporate a geographic bias (derived from the hotspot area) as query expansion in the Geonames query issued during toponym resolution. Thus, for instance, it searches the toponym "Puna" under the spatial context of Hawaii, which thereupon is correctly resolved to a village on the Island of Hawai'i (instead of being wrongly mapped to Pune in India).

Whereas we examined the feasibility of these feedback loops in initial off-line experiments, a comprehensive evaluation of these is beyond the scope of the present work, as ground truth data sets are difficult to obtain. For future work, we plan to study the effectiveness of these feedback loops based on comprehensive long-term evaluations on online stream-monitoring, for which we need to elaborate a test bed allowing to log the behavior of the $crowd^{SA}$ system and correlate it with "ground truth" on the recorded events gathered from external reports (e.g., by analyzing news reports on areas severely hit by a specific disaster, and examining whether $crowd^{SA}$ has been able to fine-tune its crowd-sensing behavior towards these areas in the course of its monitoring process).

Addressing user refinement represents a further strain of planned future work. Naturally, the proposed feedback loops also may serve as a means for integrating the human operator's knowledge into the system, who can use these interfaces to manually supply required situational context information: The operator may enter novel keywords driving crowd-sensing, mark known hotspot areas on the user interface map to guide toponym resolution, correct or guide entity disambiguation, as well as manually populate the ontology (e.g., by adding a forecast weather disaster's name). Whereas the focus of the present work has been on examining the potential of automated processing, we plan to develop dedicated interfaces for human interventions for future work, to provide a unique interface for integrating both, knowledge derived from the system as well as the human operator's expertise.

IV. Related Work

In the present section, we contrast our devised implementation to related work, and pinpoint its differences w.r.t. existing solutions.

Information Extraction. Whereas several crowd-sensing systems apply NLP for IE of CM-relevant information from the tweets' textual content [3], most systems focus on NE extraction, such as extracting locations, persons and organizations, but do not further examine their interrelations and compositional semantics. We therefore concentrate our discussion on approaches aiming at the extraction of CM-relevant, high-level information types as classified in [22] (such as information regarding the categories *Caution and Advice*, *Casualties and damage* etc.). Imran et al. [23] proposed a machine-learning approach for automated extraction of such CM-relevant information from tweets. Consequently, their approach requires human-labeled, thus costly to obtain, training data, as a separate classifier for each type of information is needed. Since examinations of different crises revealed non-satisfactory domain adaptation - i.e., a classifier trained on one crisis event poorly generalizes towards another crisis event (even when corresponding to a similar type of crisis) [2], they developed AIDR, a system incorporating a crowd-sourcing platform for informativeness filtering of crisis tweets. Thus, human crowd-workers need to label incoming data streams to generate training data for the subsequent online learning suite in real-time [24]. ESA, the CM system proposed in [25], classifies incoming tweets whether they contain any information on infrastructural damage, but does neither extract the specific type of damage nor identify its location, deferring these tasks to the system's human operators.

Aggregation. AIDR and ESA produce lists of informative tweets, but do not investigate towards further extraction of their content and fusion to object-level information (i.e., inferring descriptions of the underlying real-world events and objects reported in these tweets). Although several crowd-sensing enhanced CM systems aggregate retrieved messages to summarize the overall event topic, these mainly perform text clustering, and do not aim towards a semantically grounded fusion of the clustered textual content to object-level descriptions [3]. Thus, our approach seeks to extend these valuable preparatory works by studying the integration of these approaches to synthesize an overall crowd-knowledge extraction pipeline.

Situation-Adaptivity. To the best of our knowledge, automatically evolving crowd-sensing along-side the monitored events in a knowledge-driven fashion has not been investigated up to now [3]. Whereas rule-based IE techniques for ontology population, as utilized in our approach, have been investigated for CM [8], [9] and in other domains (cf. [7]), to the best of our knowledge, current IE techniques for ontology population make use of the contextual information in terms of the surrounding textual content encountered *within* the analyzed document (i.e., tweet), but do not incorporate a dynamically determined, external situative context.

V. Conclusion

On the one hand, the basic information types of relevance for CM are known. On the other hand, each disaster's characteristics are unique - thus, machine learning-based IE techniques so far have shown poor generalization towards other types of disasters. We therefore have proposed an approach that seeks to interweave knowledge-based with data-driven strategies: Whereas the general domain of discourse (e.g., CM) needs to be modeled, our approach aims at dynamically learning new instance-information from the crowd, which we illustrated based on a case study on a historic crisis data set. Furthermore, we proposed three concrete situational feedback loops enabling a situation-adaptive processing, which thus incorporate established information on the detected situation to provide additional context refining its sensing and processing configuration.

VI. Acknowledgments

This work has been funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under grant FFG BRIDGE 838526, WTZ AR 18/2013 and WTZ AR 2015. We furthermore thank Caroline Salfinger for creating Fig. 2.

References

[1] A. Olteanu *et al.*, "What to Expect When the Unexpected Happens: Social Media Communications Across Crises," in *Proc. of the 18th ACM Conf. on Computer Supported Cooperative Work & Social Computing*, ser. CSCW '15. ACM, 2015, pp. 994–1009.

[2] M. Imran *et al.*, "Processing Social Media Messages in Mass Emergency: A Survey," *ACM Computing Surveys*, vol. 47, no. 4, 2015.

[3] A. Salfinger *et al.*, "Crowd-Sensing Meets Situation Awareness - A Research Roadmap for Crisis Management," in *Proc. of the 48th Annual Hawaii Intl. Conf. on System Sciences (HICSS-48)*, 2015.

[4] B. Pröll *et al.*, "crowdSA - Crowdsourced Situation Awareness for Crisis Management," in *Proc. of Social Media and Semantic Technologies in Emergency Response (SMERST)*, 2013.

[5] A. Salfinger *et al.*, "crowdSA - Towards Adaptive and Situation-Driven Crowd-Sensing for Disaster Situation Awareness," in *Proc. of IEEE Intl. Multi-Disciplinary Conf. on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, 2015.

[6] ——, "Towards a Crowd-Sensing Enhanced Situation Awareness System for Crisis Management," in *Fusion Methodologies in Crisis Management*, G. L. Rogova and P. Scott, Eds. Springer Intl. Publishing, 2016.

[7] D. Maynard *et al.*, "NLP Techniques for Term Extraction and Ontology Population," in *Proc. of the 2008 Conf. on Ontology Learning and Population: Bridging the Gap Between Text and Knowledge*. IOS Press, 2008, pp. 107–127.

[8] P. Sandeep *et al.*, "Structured Information Extraction from Natural Disaster Events on Twitter," 2014.

[9] V. Zavarella *et al.*, "An Ontology-Based Approach to Social Media Mining for Crisis Management," in *ESWC 2014*, 2014.

[10] M. Kokar *et al.*, "Situation tracking: The Concept and a Scenario," in *Military Communications Conf., 2008. MILCOM 2008. IEEE*.

[11] A. Salfinger *et al.*, "Staying Aware in an Evolving World — Specifying and Tracking Evolving Situations," in *IEEE Intl. Inter-Disciplinary Conf. on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE, 2014.

[12] K. Bontcheva *et al.*, "TwitIE: An Open-Source Information Extraction Pipeline for Microblog Text," in *Proc. of the Intl. Conf. on Recent Advances in Natural Language Processing*. ACL, 2013.

[13] M. Cunningham *et al.*, "GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications," in *Proc. of the 40th Anniversary Meeting of the Assoc. for Computational Linguistics (ACL'02)*, 2002.

[14] S. Vieweg *et al.*, "Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness," in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, 2010.

[15] H. Purohit *et al.*, "Identifying Seekers and Suppliers in Social Media Communities to Support Crisis Coordination," *Computer Supported Cooperative Work (CSCW)*, vol. 23, no. 4-6, 2014.

[16] J. L. Leidner, "Toponym resolution in text," Ph.D. dissertation, 2007.

[17] R Core Team, "R: A Language and Environment for Statistical Computing," Vienna, Austria, 2015.

[18] S. Urbanek, "Rserve: Binary R server," 2013.

[19] E. J. Pebesma and R. S. Bivand, "Classes and methods for spatial data in R," *R News*, vol. 5, no. 2, pp. 9–13, 2005.

[20] R. S. Bivand *et al.*, *Applied spatial data analysis with R, Second edition*. Springer, NY, 2013.

[21] A. Baddeley and R. Turner, "spatstat: An R Package for Analyzing Spatial Point Patterns," *J. of Statistical Software*, vol. 12, no. 6, pp. 1–42, 2005.

[22] S. E. Vieweg, "Situational Awareness in Mass Emergency: A Behavioral and Linguistic Analysis of Microblogged Communications," Ph.D. dissertation, University of Colorado at Boulder, Colorado, 2012.

[23] M. Imran *et al.*, "Extracting information nuggets from disaster-related messages in social media," in *ISCRAM 2013 Conf. Proc. - 10th Intl. Conf. on Information Systems for Crisis Response and Management*. Karlsruher Institut fur Technologie (KIT), 2013, pp. 791–801.

[24] ——, "AIDR: Artificial Intelligence for Disaster Response," ser. WWW Companion '14, 2014.

[25] J. Yin, *et al.*, "Using Social Media to Enhance Emergency Situation Awareness," *Intelligent Systems, IEEE*, vol. 27, no. 6, 2012.