Towards Real-time Emergency Response using Crowd Supported Analysis of Social Media

Jakob Rogstadius, Vassilis Kostakos
Madeira Interactive Technologies Institute
University of Madeira
9000-390 Funchal, Portugal
{jakob,vk}@m-iti.org

Jim Laredo, Maja Vukovic
IBM T.J. Watson Research Center
Hawthorne NY 10532, USA
{laredoj,maja}@us.ibm.com

ABSTRACT
This position paper outlines an ongoing research project that aims to incorporate crowdsourcing as part of an emergency response system. The proposed system's novelty is that it integrates crowdsourcing into its architecture to analyze and structure social media content posted by microbloggers and service users, including emergency response coordinators and victims, during the event or disaster. An important challenge in this approach is identifying appropriate tasks to crowdsource, and adopting effective motivation strategies.

Author Keywords
Emergency response, social media, crowdsourcing, text mining.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms
Design, Experimentation.

INTRODUCTION
The period of time following a natural disaster or other large scale emergency is traditionally characterized by individuals having limited situational awareness bound to their immediate surroundings, combined with sparse high level summaries provided by traditional media. More recently, during events such as earthquakes, elections, bushfires and terrorist attacks, people have systematically chosen to share their knowledge on a micro level with others through online social media such as Twitter [3,9,10].

In fact, it is often the case that reports of incidents get published through social media before they reach regular media. However, despite the timeliness and volume of this new information source, it is highly challenging for users to overview and navigate the torrent of information which can result from such large scale events. In addition, the absence of summaries and validity checks of claims made by posters add further complexity to the already challenging task. In the near future we are likely to see an increase in volume of produced social media content, thus further increasing the need for improved structure and overview.

CHI 2011, May 7–12, 2011, Vancouver, BC, Canada.

ENVISIONED SYSTEM

Architecture
We envision a system design (Figure 1) in which machine learning and automated tools work hand in hand with a crowdsourcing community to quickly and efficiently organize and analyze information on microblogging websites during crisis and emergency situations. The system we envision has six main responsibilities or components.

1. Collect crowd generated data (e.g. by tracking keywords on Twitter).
2. Make a "best attempt" at structuring the data using NLP and other text mining techniques, as well as extracting named entities, locations and important points in time.
3. Identify shortcomings in the collected data and the inferred structure, formulate tasks and seek answers via a crowdsourcing platform.
4. Integrate the new knowledge provided by the crowd into the existing knowledge base.
5. Present aggregated and structured data back to the community, i.e. emergency response professionals, affected community members and others.
6. Wherever possible, support direct interaction between users of the presentation layer and the original information providers.

Crowd in the loop
Two vital feedback loops exist in this design. The *analysis loop* is one where the system gives users an improved understanding of the event; enabling improved actions and communication (for instance by directly addressing messages to service users whose reports have been collected by the system). This in turn changes the state of the event, which is reflected by a change in the inflow of information to the system.

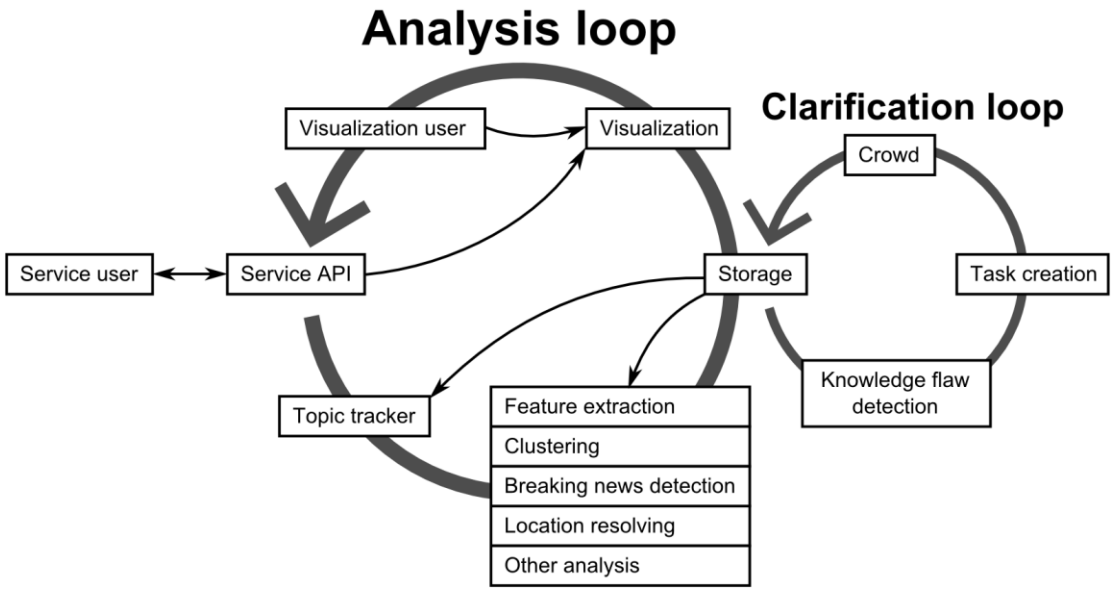

Figure 1. The proposed system architecture that integrates crowdsourcing into the analysis of social media content. Two feedback loops are present in the flow of information; the analysis loop and the clarification loop.

In the *clarification loop*, the system identifies information gaps, contradictions, weaknesses or uncertainties in the current information coverage of the situation. It then turns these flaws in the knowledge base into tasks suitable for crowdsourcing, sends them off to a crowdsourcing engine and integrates the results back into the knowledge base. We argue that by merging automatically aggregated information from social media with the output of crowdsourced work, the system can have the short processing times and scalability of an algorithmic approach, combined with the adaptability of humans. By integrating the crowd into the analytic process, the system will be able to infer structure in ways that closely match human cognitive models, even for topic domains where the training corpora is scarce.

Information generated through crowdsourced tasks should ideally be fed back into the medium which formed the original input to the system, to decouple the analysis and knowledge representation from the presentation layer and thereby simplifying development of clients for different technical platforms. In fact, if gathered data and drawn conclusions can be made publicly available (e.g. as social media updates or RSS feeds), this knowledge can be accessed in its raw form through any existing client. For instance, members of the crowd can be asked to track down images depicting an event, which the system then automatically shares through a designated Twitter account.

A significant strength of our proposed system over existing disaster tracking systems such as Ushahidi (www.ushahidi.com) is that it listens to communication channels that people already use in their pre-event lives, rather than attempting to rally information providers for a new channel once the event has already taken place. This social media content is available regardless of the success and popularity of the system itself and by merely acting as a crowd-algorithm-powered information catalyst, it becomes easier to deploy the system in particular during early stages of an event.

RESEARCH AGENDA

Related work

The proposed system builds on existing research in text mining methods, such as clustering, named entity extraction and relevance classification, and in particular methods adapted for social media content [2,5,6]. Furthermore, crowdsourcing platform design, e.g. Amazon's Mechanical Turk (mturk.com) and CrowdFlower (crowdflower.com) are directly relevant to this work, as are media aggregation systems such as Twitrix+ [8] and the Europe Media Monitor [7]. Finally, motivational factors governing the quality and quantity of crowdsourced work, both of extrinsic and intrinsic nature [1,4], are directly relevant.

Ongoing research

The ongoing research efforts in this project are currently focused on measuring the interaction effects of intrinsic and extrinsic motivation on crowdsourced work. In addition, we are currently adapting text mining methods to streaming social media, in ways that permit integration of a crowd-in-the-loop at different stages of the analysis. Finally, we are in the process of identifying types of crowdsourcing tasks suitable for being generated by the system.

Research challenges

There exists a series of research challenges that need to be addressed. In terms of *knowledge mining*, we believe there is a need for knowledge representations that support both the identification of missing information and turning the gaps into crowdsourced tasks. Additionally, we require suitable techniques for keyword extraction and pruning for

high quality topic tracking in real-time, as well as techniques for capturing the location (and possibly context) of people contributing information to the system.

In terms of *representation*, we expect to develop visualization techniques for the collected data, and suitable UI's for commenting and responding to content generated by others or the system. A further challenge will be the design of compact and accurate summarizations of large social media content clusters of similar topic.

Finally, in terms of *crowdsourcing*, there exists an open issue of managing the tradeoffs between quality, cost and time needed to complete tasks, in the context of the varying priorities that are applicable during disasters. In addition, the critical information flow in the system must be algorithmic, as processing bottlenecks are otherwise likely to appear due to lack of people willing to work, or lack of incentives to offer the workers. Part of this research must clearly identify the functionality that belongs to this critical path and define support tasks that can be delegated to a crowd. Even if human intelligence is necessary for high quality output, the system as a whole must still be functional without workers.

CONCLUSION

This paper has outlined our ongoing efforts at building a crowd-powered system for real-time response to emergency events by analysing information available on social media. We have outlined a proposed system architecture for involving crowds in the real-time analysis, and have designed two important feedback loops. These will enable a) the system to give users feedback about the status of the event, and b) the users to give feedback to the system to improve its analysis. Finally, the paper summarizes our initial findings, our ongoing efforts, as well as a set of challenges that we expect to tackle in the future.

ACKNOWLEDGMENTS

This work is funded by an IBM Open Collaboration Award and by the Portuguese Foundation for Science and Technology (FCT) grant CMU-PT/SE/0028/2008 (Web Security and Privacy).

AUTHOR BIOGRAPHIES

Jakob Rogstadius is pursuing a PhD in HCI at the University of Madeira, where he conducts research on how community generated content (such as Twitter) can be leveraged to create situational awareness in crisis situations through both crowdsourcing and algorithmic approaches. Before joining the University of Madeira, he designed visualization tools and data mining algorithms for a company in the fuel and energy sector; and he has been employed as research engineer at Sweden's National Center for Visual Analytics. He holds a MSc in Media Technology and Engineering from the University of Linköping, Sweden, where his final thesis project dealt with applied information visualization.

Vassilis Kostakos is an Assistant Professor in the Madeira Interactive Technologies Institute at the University of Madeira, and an Adjunct Assistant Professor at the Human Computer Interaction Institute at Carnegie Mellon University. He received an IBM Open Collaboration Award in 2010, and is a Fellow of the Finland Distinguished Professor Program. His research interests include: mobile and pervasive computing, human-computer interaction, social networks, crowdsourcing.

REFERENCES

1. von Ahn, L. Games with a purpose. *Computer*, IEEE Computer Society (2006), vol. 39, issue 6, 92-94.
2. Berry, M. *Survey of text mining: clustering, classification and retrieval, Second Edition*. Springer-Verlag, New York, USA, 2004. ISBN 0-387-95563-1.
3. Burns, A., Eltham, B. Twitter free Iran: An evaluation of Twitter's role in public diplomacy and information operations in Iran's 2009 election crisis. In *Record of the Communications Policy & Research Forum 2009*, Network Insight Institute (2009), 298-310.
4. Gneezy, U., Rustichini, A. Pay enough or don't pay at all. *The Quarterly Journal of Economics*, MIT Press (2000), 791-810.
5. Kotsiantis, S. B. Supervised Machine Learning: A Review of Classification Techniques. *Informatica 31* (2007), 249-268.
6. Petrović, S., Osborne, M., Lavrenko, V. Streaming First Story Detection with application to Twitter. In *Proc. HLT 2010*, Association for Computational Linguistics (2010), 181-189.
7. Piskorski, J., Tanev, H., Atkinson, M., van der Goot, E. Cluster-Centric Approach to News Event Extraction. In *Proceeding of the 2008 conference on New Trends in Multimedia and Network Information Systems*, IOS Press (2008), 276-290.
8. Sheth, A., Purohit, H., Jadhav, A., Kapanipathi, P., Chen, L. Understanding Events Through Analysis Of Social Media. In *Proc. WWW 2011*, ACM Press (2010).
9. Slagh, C. L. *Managing chaos, 140 characters at a time: How the usage of social media in the 2010 Haiti crisis enhanced disaster relief.* Georgetown University, USA, 2010.
10. Vieweg, S., Hughes, A., Starbird, K., Palen, L. Microblogging during two natural hazards events: what Twitter may contribute to situational awareness. In *Proc. CHI 2010*, ACM Press (2010), 1079-1088.