
CrowdHelp: Application for Improved Emergency Response through Crowdsourced Information

Liliya I. Besaleva

Department of Computer Science
University of Virginia
Charlottesville, VA, USA
besaleva@virginia.edu

Alfred C. Weaver

Department of Computer Science
University of Virginia
Charlottesville, VA, USA
weaver@virginia.edu

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
UbiComp'13 Adjunct, September 8–12, 2013, Zurich, Switzerland.
Copyright © 2013 978-1-4503-2215-7/13/09...\$15.00.

Abstract

Emergency resources are often insufficient to satisfy fully the demands for professional help and supplies after a public disaster. Furthermore, in a mass casualty situation, the emphasis shifts from ensuring the best possible outcome for each individual patient to ensuring the best possible outcome for the greatest number of patients. In the past several years, an ongoing movement among crisis management organizations is the incorporation of ubiquitous Web 2.0 tools into their practices for the improvement of their critical situations response. In unison with this trend and the latest discoveries in crowdsourcing, we have developed a system, called CrowdHelp, for real time patient assessment which uses mobile electronic triaging accomplished via crowdsourced and sensor-detected information. With the use of our system, emergency management professionals receive most of the information they need for preparing themselves to perform a timely and accurate treatment of their patients even before dispatching a response team to the event.

Author Keywords

Crowdsourcing; Triage; Mobile Devices; Clusterization, Feedback

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous. J.3 [Computer Applications] --- Life and Medical Sciences, Health

Introduction

Preventing natural disasters is beyond our capabilities, but providing better information to disaster management professionals and affected persons (citizens) is not. Earthquakes, tsunamis, tornadoes, violent storms, and influenza epidemics have proven difficult to manage despite all of our technological advancements. When a disaster occurs, sometimes the best (or indeed only) way to save numerous lives is to generate a timely reaction to its outcomes. That puts extra emphasis on the importance of high-quality information management for emergency communications. This is particularly visible in critical situations because internal and cluster communication within relief groups, non-governmental organizations, civil societies and other agencies is crucial to getting their work accomplished in a timely manner.

One of the most important steps to be taken during mass-casualty emergencies is to quickly and accurately triage (assess, sort and count) those who are in need of help. A standard medical triage scheme typically helps in the process of determining the priority of patients' treatments based on the severity of their condition (see Figure 1).

As a result of this triaging, numerous vital decisions about the best resource division and accommodation are made within the emergency response organizations.

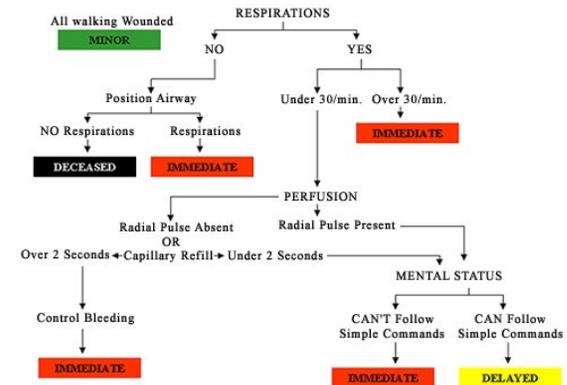


Figure 1: Simplified schematic representation of a medical triage system.

Historically, various manual and electronic medical triage systems have been developed and used both under civil and military conditions to determine the order and priority of emergency treatment, transport, and best possible destination for the patients [4][5][8][9]. These solutions, however, have proven labor intensive, slow, and prone to error when done by hand, or simply too obtrusive and unavailable for mass use when in the form of specialized systems [1].

Consequently, many of the leading rescue and disaster management organizations have been looking for other strategies for acquiring relevant information. Most recently, they have been turning towards social networking websites to take advantage of their huge user base, accessibility, and fast response. [3][7] Hence, one of the positive characteristics of social media, which is its capacity to harness collective knowledge for learning and problem-solving, is no longer being overlooked.

The reason behind the rising usefulness of Web 2.0 technologies to the healthcare systems is that through their highly crowdsourced information, constant availability, and ease of use they greatly improve information accessibility. While the local widespread communication and information services (e.g., telephone, television, radio) often crash during large natural disasters, the fabric of the social networking websites can remain intact due to its distributed and agile nature. Social networking can provide a reliable means of communications for both those involved and those witnessing an incident.

During the past several years, considerable effort has been expended to utilize the power of online social networks for crisis management. Both commercial and research organizations have developed a large number of applications that attempt to gather, sort and analyze Internet-available information. They frequently revolve around identifying social phenomena through tweets, analyzing Facebook posts, and recognizing images available on online community websites.

Crisis management using crowdsourcing has been attempted by numerous organizations from all around the world and of all sizes. Following that trend, and taking into consideration the latest movements toward mobility in the technology world, we now present our design for a practical and highly capable application called CrowdHelp. CrowdHelp is a versatile application (accessible on smartphones, tablets and computers alike) which collects direct feedback from its users about their current medical state, in combination with data coming from sensors in smart devices

(smartphones, tablets, laptops, etc.) and is to be used for enabling fast response to devastating events. CrowdHelp is based squarely on the idea of social networks, and its functionality depends upon gathering the “wisdom of the crowd.” It is an advance beyond currently existing technologies because of its centralized approach for data availability, fast data analysis, innovative combination of techniques (machine learning, cyber-physical systems, and crowdsourcing) and its focus on public safety operations. CrowdHelp was developed in collaboration with emergency response and medical specialists with experience from large natural disasters such as the 2010 Haiti earthquake [11].

System design

CrowdHelp is a software platform which combines two major components into its workflow (see Figure 2)—an application serving as a gateway to CrowdHelp’s functionalities for ordinary users and victims, and a server layer providing the backend features of the system to be used by disaster management professionals (DMPs). From a user’s point of view CrowdHelp is a downloadable application which is also accessible as a website online. CrowdHelp could be used equally well by numerous people reporting an event, as well as by a single individual who wants to determine what his/her physical condition might be and what are the possible symptoms and treatments of that condition. Additionally, CrowdHelp allows its users to submit reports both from the perspective of an injured victim, and of a bystander who is reporting someone else.

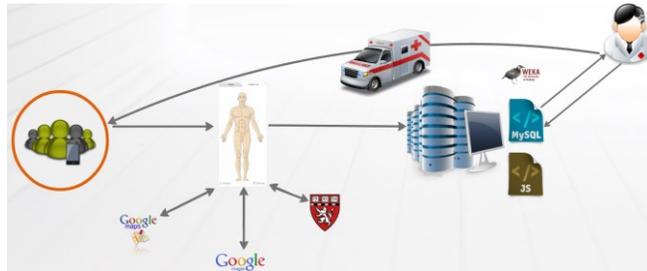


Figure 2: Schematic presentation of the workflow of the CrowdHelp platform

When used for emergency reporting our application has several main features: it allows its users to submit information relevant to the event; provides the users with information about their possible conditions, as well as the possible causes for their symptoms; dynamically populates a list built via the official Google Maps Javascript API, Version 3, which displays all places capable of treating the victims; and maintains a user profile with information about the user's location and reported symptoms. The data submission is done through an intuitive user interface which uses a simplistic visual reporting technique (see Figure 3).

Information about the possible symptoms and causes is generated based on the user's report of their (or the person they are reporting) injured body area(s) and gender. As a result, a number of related imagery and videos which describe the symptoms, as well as a list of emergency medical treatment centers (which could be sorted based on their features or location) is populated dynamically (see Figure 4).

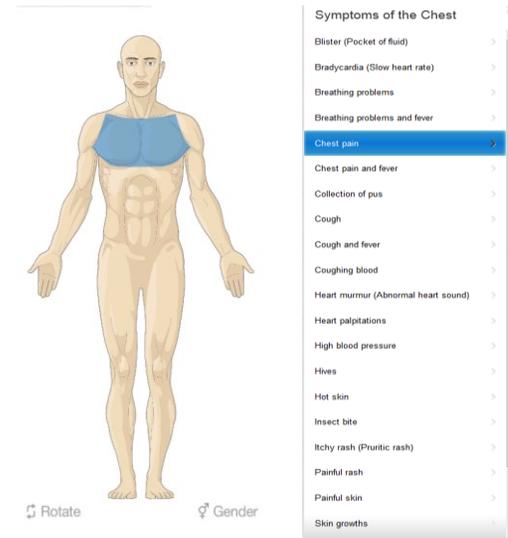


Figure 3: CrowdHelp's interface with body part selection indicating injured area and the most likely symptoms

From the point of view of emergency professionals who have server access, CrowdHelp is a service for data analysis (see Figure 5) which works together with the machine learning software WEKA [6][10]. It provides a number of machine learning algorithms and computational formulas for clarifying, deciphering, summarizing and clustering all inputs into easily comprehensible visual images representing the geographical location, urgency and association of each entry to a specific cluster (see Figure 6).

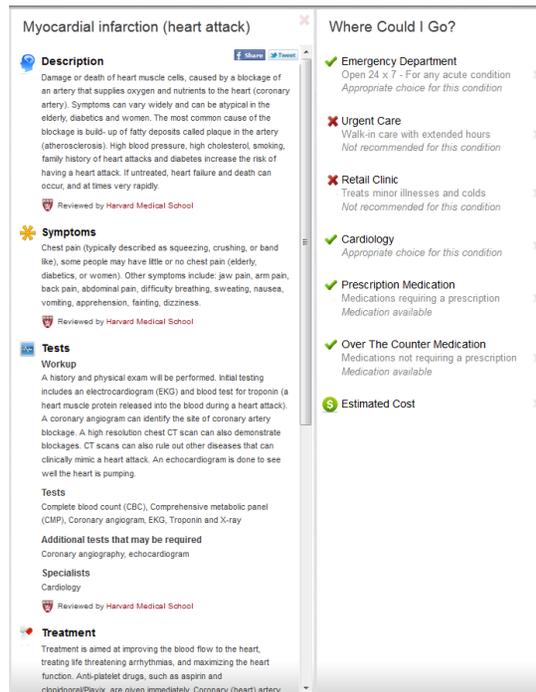


Figure 4: Symptoms and medical centers list

CrowdHelp provides its backend users with the ability to customize all testing, analytical and clusterization methods by choosing their preferred machine learning algorithms, visual characteristics of the representation and data attributes. The data that is collected through that application could be used by the responsible organizations both for immediate response and for future performance improvement analysis and event model building.

The server layer is designed with a user management module which provides its users with a hierarchical system of access rights, and an easy to operate

interface for performing tasks such as data manipulation and analysis configuration customization (selecting the clusterization parameters).

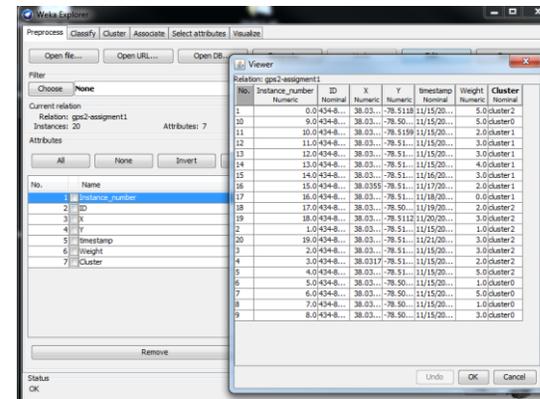


Figure 5: CrowdHelp's backend

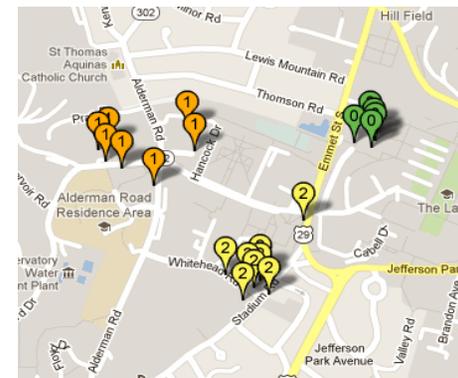


Figure 6: Example clustering based on medical urgency

For security, as well as for performance optimization, the user management system comes with user roles—guests, operators, and administrators. A guest user can only view the server and cluster configurations and

monitor the selected cluster. An operator can view the server and cluster configurations and perform basic administrative tasks. An administrator can view the server and cluster configurations, perform administrative tasks (e.g., switch to operations, freeze operations, offline operations), modify configuration (e.g., create new groups, add resources, delete resources and or groups), perform advanced cluster manager functions. Initially everyone begins as a guest user with the option to be promoted to an operator or administrator user by another administrator.

In CrowdHelp's workflow the following steps take place:

(a) Download and install the *CrowdHelp* application onto a smartphone/tablet (or open in a web browser).

(b) Check the user's status (e-mail login required, optional social networks login is available). If this is an established user skip to step (d).

(c) Create a new user. Confirm the newly created registration with a confirmation email (not necessary when logged in with an existing social networking registration). Ask for permission to access the user's GPS coordinates, network location and time, compass and accelerometer/gyroscope readings.

(d) Present the user with a simple interface for symptoms selection, and then send the combined sensor readings and answers to the server.

(e) In the server: store the GPS coordinates in a spatial database. Analyze the data with machine learning algorithms for clusterization (e.g., XMeans, Density Based Clustering, Expectations-maximization).

(f) In the emergency management organization: Access the stored data, choose the preferred clusterization

settings and manage the stored data by marking 'jobs' (entries) as open, pending, or finished. Read automatically generated suggestions which are based on the chosen answers in the questionnaire.

(g) In the user device: Receive warning push messages for forthcoming dangerous events (e.g., blocked roads, floods, storms).

What follows next is information distribution, which is an important step of the process of crisis management with its potentially vital social importance. Correct distribution of the newly gained knowledge to specialized organizations (e.g., governmental organizations, local rescue teams, volunteers) or the general public could result in saving lives and minimizing the devastation caused by disasters.

Testing

Testing objective

CrowdHelp is a support tool for both disaster management professionals and citizens. It has three main contributions. On one side, two of those contributions lie in providing a previously unavailable service to emergency professionals by assisting them in processing large amounts of data more efficiently and by providing them with helpful priority suggestions. This is done through a suite of techniques for accurate data clusterization and visualization for easier data comprehension. And on the other side, it empowers the victims to help themselves, and each other, by giving them a tool for self-assessment (self-triaging) and by providing them with a useful package of information about their condition.

As this is a new service, not comparable to existing methods, our leading testing objectives focused on the technical feasibility of this project. Our service is

designed to be highly flexible and inclusive in the types of information it provides. Clusterization of the data is done based on three different criteria – medical urgency, physical proximity to dangerous events (especially useful in the case of pollution), and geographical location with regards to neighboring entries (victims). All three clusterization maps could be visualized on separate system views (see Figures 6,7,8).

For the visualization we used the standard medical triaging conventions where all entries have levels of urgency ranging from 0-5 with zero being the lowest, therefore least urgent, and five being of the highest possible urgency. Each entry is given a number representing to which cluster it belongs, as well as a color representing the overall urgency of the cluster. The colors are selected as follows: green for urgency level of <1, yellow for urgency level of 1–2.5, orange for urgency level of >2.5 and <=3.5, red for urgency levels of >3.5.

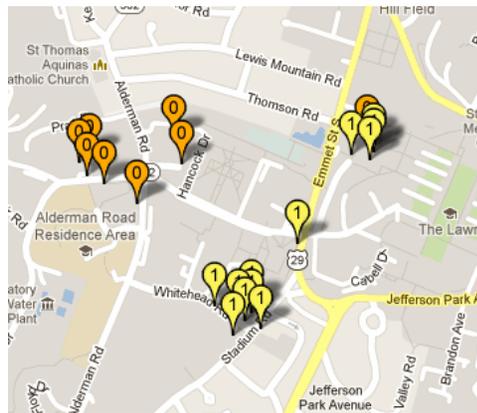


Figure 7: Clustering based on proximity to neighboring entries

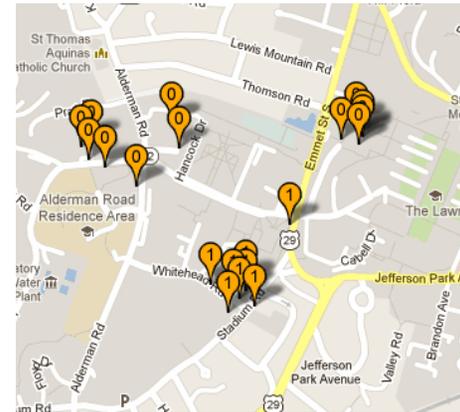


Figure 8: Clustering based on physical proximity to dangerous events

Testing Setup

Priorities were assigned on two main levels – entry-wise and cluster-wise.

- Entries level

In the case of doing clusterization according to **medical urgency**, because *CrowdHelp* provides information about concrete medical symptoms (e.g., bleeding, numbness), we assigned priorities based on the standard medical triage system used in the United States (see Figure 1). Our testing schema consisted of the following. We designed 100+ test cases with different symptoms (all combinations were made based on the likelihood of their combined occurrence according to the Harvard Medical DB) and manually assign each an urgency level based on the triage system. Then we randomly generated 50000 combinations of those test cases where each combination takes the highest urgency of its members. For instance, if we categorize a bleeding head trauma

as an urgency level 4, and a minor leg injury as 1, a case with both a bleeding head trauma and a minor leg injury will be considered as urgency 5. After that, we customized an XMeans algorithm to use our specific data, and used the previously assigned urgency levels and color coordinated the entries according to the triage scheme.

Over 5000 tests of different sizes of the test data were performed and all of them showed uniform results representing a neat visualization of the entries on a map. In the cases of entry overlaying some of the entries are automatically merged into a group shown through a single entry point on the map which contains meta-data about the elements of which it is made.

When doing clusterization based on **physical proximity to dangerous events**, we assigned priority levels based on the location of the victim and the anticipated size of the area to be impacted by the event. The nearest 15 % are assigned as high priority, following by medium priority for the next 35%, and green for the remainder. For instance, if we have an event which is expected to affect an area with diameter three miles (15840 feet) from its epicenter, all people within 2376 feet from it will be priority red, within 2377 – 7920 feet will be yellow, and the rest who are further away will be green.

To confirm the integrity of the resulting output from our system we performed over 350 tests and compared the actual physical location of the entries to a fixed point on the map and each time they were properly color prioritized according to our clustering rules.

Lastly, when we cluster based on **proximity to neighboring entries**, we do density based clustering which is used for finding non-linear shapes structures

based on the entries' density. When doing density clustering we took into consideration two factors – density reachability and density connectivity.

For reachability we assumed the following rule: *a point p is said to be density reachable from a point q if point p is within distance ϵ from point q and q has a sufficient number of points in its neighbors which are within distance ϵ .*

For connectivity we used the following rule: *points p and q are said to be density connected if there exists a point r which has a sufficient number of points in its neighbors and both the points p and q are within the distance ϵ . This is a chaining process. So, if q is a neighbor of r , r is a neighbor of s , and s is a neighbor of t which in turn is a neighbor of p , this implies that q is a neighbor of p*

- Clusters level

All formed clusters are assigned overall priority levels which represent the median of all priorities of the entries in that cluster. This was a straightforward task which required going through the priorities of all entries and sorting them out to find the median. This is an additional feature of the system which is added for richer functionality and higher customization capabilities.

Application training

Modern mobile devices are highly capable of providing various types of useful information about their users through the rich suite of sensors available within them. Our application is exploiting smartphone sensors to infer user behavior, activity, or context to refine its preliminary results coming from crowdsourced data. Automated decision making requires training using labeled data as ground truth. Obtaining such data is not

at all easy, especially when the recreation of the appropriate circumstances on a large enough scale is very difficult (e.g., earthquakes, massive riots). Therefore, in our system we are adopting a technique from Bao, Kansal, and Chouhury's paper 'Helping Mobile Apps Bootstrap with Fewer Users' [2] for automatically adapting the inference model required by the application to multiple users. This method is based on an insight that applying well-known semi-supervised learning methods would display certain data points that occur with higher frequency and correspond to higher inference. Using only the high-confidence points as seeds for semi-supervised learning, the unlabeled data is incorporated for each user to train a new model specific to that user. As a result, with only small numbers of trained data we are able to 'teach' our application how to behave appropriately in the case of a mass-casualty emergency.

References

- [1] Asaeda G. MD. The Day That the START Triage System Came to a STOP: Observations from the World Trade Center Disaster. *Academic Emergency Medicine*, vol. 9, Issue 3, pp. 255-256, March 2002. Article first published online: June 28, 2008. DOI: 10.1197/aemj.9.3.255.
- [2] Bao X., Bahl P., Kansal A., Chu D., Choudhury R. R., and Wolman A. Helping mobile apps bootstrap with fewer users. *In proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12)*. ACM, New York, NY, USA, pp. 491-500. <http://doi.acm.org/10.1145/2370216.2370289>
- [3] Briones R. L., Kuch B., Liu B.F., and Jin Y. Keeping up with the digital age: How the American Red Cross uses social media to build relationships. *Public Relations Review*, vol. 37, Issue 1, pp. 37-43, March 2011.
- [4] Gao T., Massey, T., Selavo, L., Crawford, D., Chen B.; Lorincz, K., Shnayder, V., Hauenstein, L., Dabiri, F., Jeng, J., Chanmugam, A., White, D., Sarrafzadeh, M., and Welsh, M., The Advanced Health and Disaster Aid Network: A Light-Weight Wireless Medical System for Triage, *Biomedical Circuits and Systems, IEEE Transactions on* , vol.1, Issue 3, pp.203-216, September 2007. DOI: 10.1109/TBCAS.2007.910901.
- [5] Gao, T., and White, D. A next generation electronic triage to aid mass casualty emergency medical response. *In proceedings of the 28th IEEE EMBS Annual International Conference*, 2006.
- [6] Hall M., Frank E., Holmes G., Pfahringer B., Reutemann P., and Witten I. H. The WEKA data mining software: an update. *Published in SIGKDD Explorations Newsletter*, vol. 11, Issue 1, pp. 10-18, June 2009. DOI: 10.1145/1656274.1656278. <http://doi.acm.org/10.1145/1656274.1656278>
- [7] Magro M. J. A Review of Social Media Use in E-Government. *Administrative Sciences*. 2012; 2(2), pp.148-161.
- [8] Massey T., Gao T., Welsh M., Sharp J. H., and Sarrafzadeh M. The Design of a Decentralized Electronic Triage System. *In proceedings of AMIA Annual Symposium 2006*, pp. 544-548.
- [9] Weinerman, R. E., Rutzen, S. R., and Pearson, D.A. 1965. Effects of medical 'triage' in hospital emergency service. *Public Health Rep.*, May 1965, 80(5), pp. 389-399.
- [10] Weka. [Online] Available at: <http://www.cs.waikato.ac.nz/ml/weka/>
- [11] Wikipedia, Haiti earthquake (2010) [Online]. Available: http://en.wikipedia.org/wiki/2010_Haiti_earthquake