# Sympathetic Decisions: Incorporating Impacts on Others into Emergency Response Decision Spaces

**Jill L. Drury, Gary L. Klein, Jennifer Mathieu**
The MITRE Corporation
{jldrury, gklein, jmathieu}@mitre.org

**Yikun Liu, Mark S. Pfaff**
Indiana University Indianapolis
{yikliu, mpfaff}@iupui.edu

### ABSTRACT

We designed two decision support tools and employed them during a one-week, simulation-driven experiment that included emergency responders acting in their real-life roles. Each tool visualized a "decision space": a diagrammatic depiction of the relative desirability of one option versus another, including the inherent uncertainty in the potential outcomes. One requirement was to develop a tool accounting for the impacts of decisions on others, so that emergency responders can make "sympathetic decisions." For example, one decision space enabled responders to request resources from surrounding jurisdictions while also considering the potential negative effects on the lending organizations. Another decision space enabled responders to engage in a strategic dialogue with the public: "listening" to the public's greatest concerns by mining social media to measure emotion, and thereby suggesting strategic communications addressing those concerns. We report how we designed the decision spaces and the qualitative results of using these spaces during the experiment.

### Keywords

Decision making, robust decision making, decision space, decision support, situation awareness, instant message, emergency response, public communication

### INTRODUCTION

Decision makers in emergency management situations must evaluate multiple options quickly under circumstances that are often uncertain and complex. They have many tools that help them to understand the state of the situation, such as WebEOC (ESi, 2012); these tools provide a view of the *situation space* (Hall et al., 2007) and help users to attain *situation awareness* (Endsley, 1988). In contrast, few tools help them to understand their options and the relative desirability of those options under uncertain conditions: what Hall et al. (2007) term the *decision space,* and which we believe can lead to *option awareness* (OA; Drury, Klein, Pfaff, & More., 2009). In the decision space, *option awareness* is the understanding of how options compare to one another and the underlying factors that contribute to the plausible outcomes of choosing any particular option. While others have reported research on visualizations to aid decision making, in general these efforts have either focused on the situation space (e.g., Dong & Hayes, 2012), or have less comprehensively addressed how to most effectively convey the decision space (e.g., Hoffman, 2006). The value of the OA framework is its applicability across disparate decision making models, from rationalistic to naturalistic (Klein, Drury, & Pfaff, 2012). Indeed, decision makers rely on both strategies in complex high-technology work, referred to as "hybrid" decision making environments (Mosier, 1997), in which expert intuition and methodical analysis must coexist.

We have been working to close this "situation-space-to-decision-space gap" (Hall, Hellar, & McNeese, 2007) by developing decision space visualizations (DSVs) that enable decision makers to see a range of outcomes for each option, such that emergency managers can use their judgment to choose an acceptable option (Drury et al., 2009; Liu, Moon, Pfaff, Drury, & Klein, 2011; Pfaff et al., 2012). The DSVs are based on models that run tens, hundreds, or thousands of "what if?" cases for each option, with each case being constructed using different combinations of slight variations of the conditions that are beyond decision makers' control. The result is a landscape of plausible outcomes for each option that takes into account the inherent uncertainty of the situation, made explicit through an interactive visualization. This offloads complex mental simulation of the combinatorial explosion of potential outcomes, freeing cognitive resources for the decision maker to be more creative and agile under complex, uncertain, and stressful conditions.

Our previous work has shown empirically that decision makers in emergency response scenarios who had DSVs made decisions more quickly and accurately and with more confidence than those who did not have DSVs (Pfaff et al., 2012). These DSVs, however, focused on comparing the impacts of the options on the target mission, as opposed to how direct response organizations, supporting response organizations, or the population might be affected. We have defined *sympathetic decisions* as choosing options in a manner that is sensitive to the impact on the associated community.

In addition, previous decision spaces have been generated based on physical models of the environment. US Presidential Policy Directive 8 (PPD8), signed 30 March 2011, states that "Our national preparedness is the shared responsibility of all levels of government, the private and nonprofit sectors, and individual citizens. Everyone can contribute to safeguarding the Nation from harm." In other words, PPD8 sets the policy of incorporating citizen input into crisis response. New technology has provided us with the opportunity of actually incorporating that input into a decision space. In the last decade, concurrent with the rise of reader-contributed web content and other social media, citizens have been providing their information, observations, questions, and concerns about crisis situations online while the crisis is ongoing. Since then, the emergency management community has struggled with how to best incorporate social media information. The current work shows how DSVs can be used to meet this objective.

This paper describes these new sympathetic decision spaces and reports on their first use in the hands of practicing emergency responders.

## RELATED WORK

There are a few emergency response tools that incorporate a decision space for use in emergency response resource allocation decisions. For example, the ARGOS Accident Reporting and Operational Guidance System for chemical, biological, radiological, and nuclear incidents provides predictions for the consequences of the events and possible courses of action (Prolog Development Center 2010). Bonazountas et al. (2007) developed a fire fighting decision support tool that shows the costs for undertaking different firefighting strategies. Neither of these tools employs exploratory modeling for analyzing uncertainties (Chandresekaran, 2005; Chandresekaran & Goldman, 2007) nor a frequentist display of the results for effective depiction of the relative desirability of one option versus another in consideration of such analyses (Gigerenzer and Hoffrage, 1995). Also, neither of these decision spaces takes into account the effects of the potential decision choices on the associated community.

A number of researchers have examined the use of social media in crisis response. Some, such as Kavanaugh et al. (2012), analyze social media after the emergency has taken place to better understand what happened. Others, such as St. Denis et al. (2012), Starbird et al. (2012), and Schulz et al. (2012), investigate novel methods of helping responders sift through massive amounts of microblogging data ("tweets") in real- or near-real time during the crisis. Additional studies have examined how citizens impacted by catastrophic emergencies have used social media to provide status information to loved ones, help others, and find information (e.g., Heverin and Zach, 2010; Palen et al., 2009; Qu et al., 2009; Shklovski et al., 2008; and Starbird and Palen, 2011). In contrast, we are interested in seeing whether using social media to incorporate community emotion and sentiment into a dynamic decision space during a crisis can help responders to make more sympathetic decisions.

## DECISION SPACE VISUALIZATIONS

Before describing the new DSVs that support sympathetic decision making, it may be helpful to describe the techniques and approaches behind our DSVs. Specifically, our DSVs use the exploratory modeling approaches of Bankes (1993) to support *robust decision making*: decisions that may not be optimal under some conditions, but will be acceptable under the widest range of plausible conditions that may occur (Chandrasekaran & Goldman, 2007). By taking a broad range of uncertainty into account, decision makers can choose an option that will be robust across that range.

One of the ways that DSVs can be generated is via computer-based forecasting models that assess dozens of options with hundreds or thousands of variations due to uncertainty. Uncertainty arises when there are variables outside of the decision makers' control. For example, the chances for successfully dousing a fire by sending two fire trucks will be different if high winds arise to fan the fire's flames, or a drenching downpour occurs, or one of the trucks breaks down, or any combination of these conditions. Simulation models can run many "what if" variations of such conditions to determine a landscape of plausible outcomes from choosing a particular option. The goal is to significantly reduce the cognitive load of conceiving and evaluating this large array of contingencies by off-loading such computation to automation.

The cost of the outcome of each "what-if" case is computed by summing the execution-cost of enacting the option and the outcome-cost of consequences resulting from enacting the option. Thus there is a distribution of possible consequences for each option. Each distribution is a function of the uncertainty of the situation space (e.g., how serious the emergency actually is) and the uncertainty inherent in the decision option (e.g., what percent of emergency vehicles will get to the scene and when). Although an optimal plan would generate the highest expected return on investment under a given set of assumptions, under deep uncertainty (Lempert et. al. 2003), where uncertainty is irreducible, optimal strategies lose their prescriptive value if they are sensitive to changes in these assumptions due to these uncertainties. In other words, selecting an optimal strategy is problematic when there are multiple plausible futures for each option. Instead, Chandresekaran (2005) and Chandresekaran & Goldman (2007) suggest shifting from seeking optimality to seeking *robustness*: acceptable outcomes across the broadest swath of plausible futures.

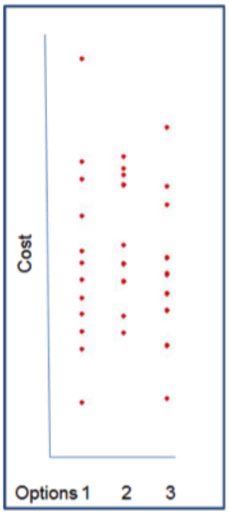


**Figure 1. This decision space shows the ranges of potential costs of three different options. Option 1 has the highest worst-case cost but the lowest best-**

Identifying robust options involves displays that enable decision makers to compare the distributions of results for each option. Hoffrage and Gigerenzer (1998) showed empirically that people are able to make better choices when they are given uncertainty information in terms of natural frequencies rather than probabilities. Natural frequencies are absolute (non-normalized) frequencies resulting from observing cases that have been representatively sampled from a population (Gigerenzer and Hoffrage, 1995). As an example of the efficacy of frequency displays, Hoffrage and Gigerenzer (1998) found that physicians who were provided information in the form of natural frequencies correctly estimated the positive predictive values of diagnostic tests more than four times as often than when they were given the same information in the form of probabilities.

This frequentist approach can also be applied to exploratory modeling results, as can be seen in Figure 1, which shows a scatter plot indicating the range of costs for each option in an example modeled scenario (in this case, lower costs are better). Each dot represents the cost of a case (or multiple cases, when they are the same for several cases and so the dots are plotted on top of one another). All three options have the same means (which are not shown in this illustration), so making a decision based on choosing the option with the lowest average cost is not a helpful discriminator. Decision makers who are most concerned about the possibility of a high-cost case occurring might choose option 2, which has lower worst-case costs than the other two options. Alternatively, if the decision maker closely examines the highest cost case for option 1 and determines that he or she can take steps to decrease the likelihood that this case will occur, or if the decision maker otherwise believes that this highest cost case will not happen, then option 1 may be a more robust choice than option 2. The decision space visualization facilitates *parallel* visual comparison of options, affording an evaluation that is difficult or impossible through *serial* mental simulation.

## COMMUNICATIONS DECISION SPACE

The first generic emergency decision space deals with a strategic dialogue between an emergency operations center and the public. The objective is similar to what Laterno et al. (2010) described for the Los Angeles Fire Department: using social media for both listening to the public as well as for informing the public. The purpose of the communications decision space is to help emergency responders in Public Affairs roles to decide what topic(s) to address next in communications to citizens in order to calm the public both about the emergency situation (e.g. fear for one's safety) and the emergency response (e.g. concern for the sufficiency of response). As such, this space is purely a sympathetic decision space. For example, even though emergency responders are indeed sending resources as quickly as possible, without adequate communication, members of the public may still worry that responders are not sending resources to trouble spots quickly enough. Communications from responders that include their plans for addressing the disaster can have a calming effect on the associated community of citizens (Sandman, 2003). So, a communications decision space should visualize for the decision makers the distribution of impacts that will come from applying various classes of communications to various topics.

As step toward implementing this communications decision space concept, we developed a simplified prototype suitable for experimentation. The prototype system generates an approximation to the communications decision space concept by automatically assessing the level of concern associated with a given topic. The prototype applies the Linguistic Inquiry and Word Count approach (LIWC; Pennebaker et al., 2007) to social media postings such as the140-character or less "tweets" used in Twitter. People who use Twitter tend to employ "hashtags" (the "#" sign followed by a topic name) to identify the keywords or principle subject(s) of a posting. Topics manually were identified using both hashtag and keyword searches. By comparing the content of tweets to LIWC's dictionary of words that have known negative emotional content, the system scored the level of negativity as an indicator of public concern regarding each topic in a twitter-like dataset (Elson et al., 2012).

Technically, instead of a forecast of outcomes for a communication action (which would yield option awareness), the experimental implementation provides the decision maker with situation awareness of the range of negativity for a topic. However, for experimental purposes, this estimate of current state was a feasible stand-in for a forecasting model whose development would have required more time than was available. See the Future Work section for how we could further develop this prototype toward a more complete implementation of the communications decision space concept.

In our initial version of the communications decision space, a decision-space administrator works with emergency response decision makers to enter candidate topics for assessment based on manually reviewing tweets as they are posted. Each candidate topic, which can be associated with a hashtag or a keyword, is at the core of an option. In Figure 2, the options are laid out on the horizontal axis and the level of negative emotion is indicated on the vertical axis. In this figure, four topics are being considered, with each being indicated by a set of three dots. The top and bottom dots of each set indicate the likely upper and lower bounds of negative emotion for that topic, with the middle dot indicating an average emotion level and the size of that dot indicating the number of people currently tweeting about that topic. The display is updated periodically to reflect the results of new communications and the occurrence of new topics.

The horizontal lines in the middle of the graph represent a 20-minute moving average of the emotion level of all of the tweets, regardless of hashtag or keyword, to be used as a standard of comparison for each option. The middle line indicates the computed average emotion level with the upper and lower lines showing upper and lower bounds for expected variance. These lines provide a reference so that decision makers can see at a glance whether a particular topic is garnering more or less negative emotion than the set of postings as a whole. The topic to the far left of Figure 2 is more negative than the postings as a whole, and so appears to be a robust option for immediate communications. But also notice that the next option to its right, although lower in predicted level, is greater in the number of people involved: this may make it a robust option as well. These options are not exclusive, and the number of communications is limited only by the resources (time, people) available to craft communications.

Once a topic-option is chosen for communication, there were three communication alternatives: providing situation awareness about the topic, providing information regarding the status of emergency response regarding the topic, or polling the citizens for further information about the topic. In this prototype, choosing among these alternative communication-actions was left for the users to evaluate, but in a more complete communications decision space, the distribution of impacts for each of these alternatives also would be depicted.

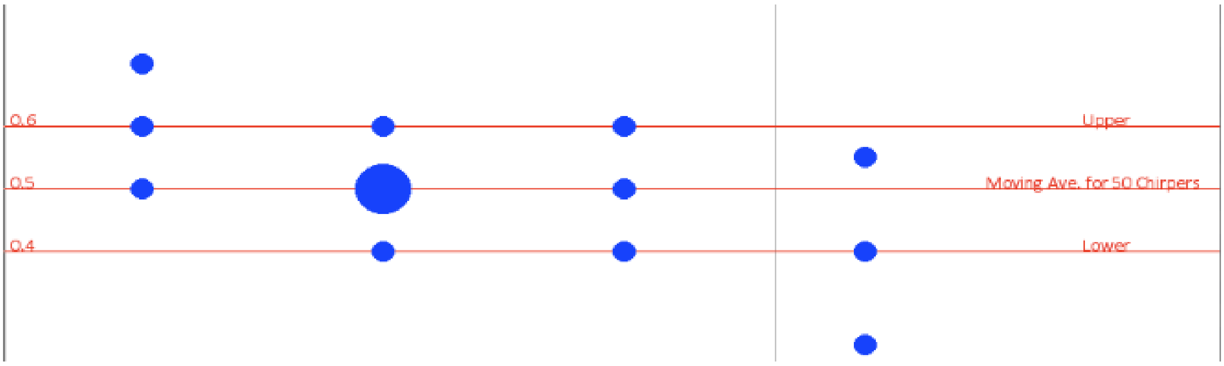


**Figure 2. Example Communications Decision Space with the options (topics to address via communications with the public) on the horizontal axis and the level of negative emotion exhibited by the public via their social media posts on the vertical axis. Each set of three dots indicates one topic, with the three lines indicating a 20-minute moving average for all people posting tweets (from top to bottom, 0.6 is the upper bound, 0.5 is the average, and 0.4 is the lower bound).**

## RESOURCE DECISION SPACE

The purpose of the Resource Decision Space is to enable decision makers to choose a neighboring city or county (or combinations of these jurisdictions) from which to request additional emergency vehicles: fire trucks, police cars, and/or ambulances. To drive that decision, the visualization of the decision space depicts both mission-oriented and sympathetic decision information. First, it displays arrival time (mission-oriented) information for different mixes of resources from the lending jurisdictions. Second, for each resource mix, it displays the forecasted impact (sympathetic information) that fulfilling the request will have on the jurisdictions. Using color-coded symbols to separate these two pieces of information, this visualization is able to display both types of information side-by-side, thereby enabling decision makers to quickly see the trade-off between mission-impacts and jurisdiction-impacts.

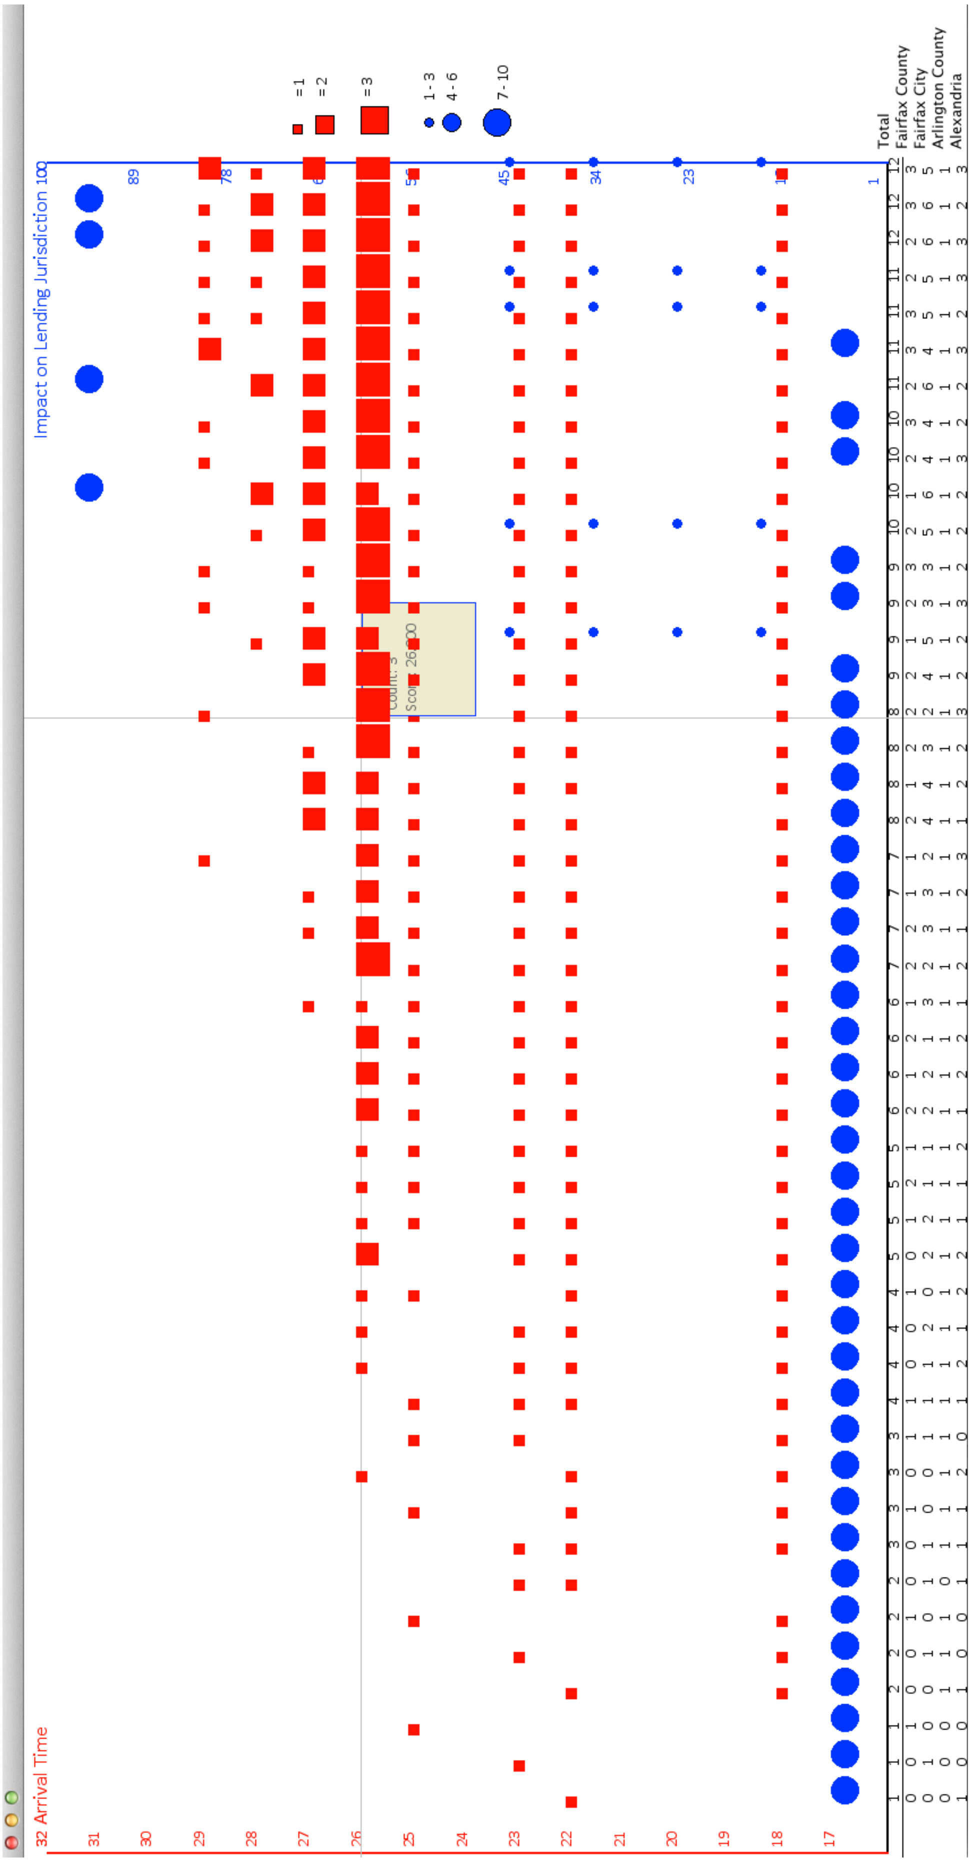


**Figure 3. Example Resource Decision Space showing the arrival times (red squares) and impacts on lending jurisdiction (blue dots) for the options of sending different combinations of between 1 and 12 fire vehicles to an emergency situation.**

The decision space in Figure 3 displays combinations of resource dispatches along the horizontal axis. The first line of the legend indicates the total number of units sent and remaining lines indicate quantities dispatched from each location. There are vertical axes on left and right sides, respectively, representing estimated "Arrival Time" (red squares) and "Impact on Lending Jurisdiction" (blue dots). "Arrival Time" forecasts for how long it will take for the individual units in a given resource mix to arrive at the location of interest. "Impact on Lending Jurisdiction" indicates how much a resource mix may impact the lending jurisdictions' capabilities for dealing with future events, based on historical trends (e.g. how much will dispatching three fire trucks now affect the ability to respond to other fires in the near future).

If multiple data points fall into close proximity, they are combined to form a larger square or dot. The legend on the left indicates how many data points the different sizes represent. Hovering a cursor over any of the data points reveals the actual value of the data as well how many data points that symbol represents. Figure 4 shows a close-up of the Resource Decision Space.

When designing the Resource Decision Space, we anticipated that decision makers would first examine the expected arrival times. Obviously, if one city is positioned to send resources faster than another city, the first city would tend to be the one chosen when asking for help. But the interesting cases happen when several jurisdictions have roughly the same response time, yet the burden on one of them is much larger than the expected impact on the others. As illustrated in Figure 4, there are four viable ways of requesting eleven (11) fire trucks from the jurisdictions. They all result in the same arrival times, but they differ greatly in their impact on the lenders. The second option (sending 3,4,1, and 3 vehicles) has the least impact (as indicated by the blue dot located lower in the graph) and so would be robust. Unaided naturalistic decision making (NDM) would not only be unable to consider the number of combinations displayed by the DSV, but because NDM considers them sequentially (Phillips et al., 2004) there would be no guarantee that the second option would even be considered. In this way, the DSV provides an extension of the NDM process (Klein et al., 2011) and we therefore expect that the Resource Decision Space could add value when making these decisions.

## EXPERIMENTAL EVALUATION OF THE DECISION SPACES

Both decision spaces were put to use in the Citizens' Emergency Response Portal System (CERPS) Simulation Experiment (SIMEX), which was held from 1 – 5 October 2012. The primary purpose of the CERPS-SIMEX was to investigate whether information communicated through social media from the general public during an emergency incident would result in a better overall response on the part of emergency operators and public entities. The scenario took place in a *completely* virtual environment in which a simulated radiological dispersal device was detonated on the campus of a simulated university. A group of students at a real university represented the public, with a response team being located in a (physical) Emergency Operations Center (EOC) created for this experiment. Operations Center personnel consisted of actual emergency managers, operators, and public relations staff from the university and government agencies, including national level, county and city emergency response managers. The CERPS-SIMEX was part of a series of SIMEX experiments conducted by The MITRE Corporation in which actual military and civilian operators use real command and control systems linked to simulated reporting and sensor systems in various crisis-based scenarios.

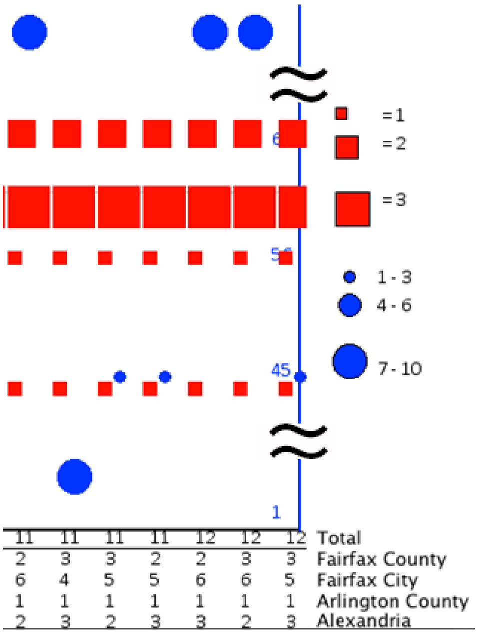


**Figure 4. Detail of the resource decision space visualization, edited to highlight the bottom section showing the available courses of action, and middle section showing the legend for the different symbols representing Arrival Time (red squares) and Impact on Lending Jurisdiction (blue dots).**

While the complete results of the CERPS-SIMEX experiment are still being analyzed and will be reported in the near future, we can share some observations based on EOC personnel's opportunities to use and comment on the decision spaces.

### Using the Communications Decision Space

The decision spaces were relevant to a subset of the EOC operators, based on their roles. Operators whose responsibilities included Public Affairs were among those who expressed the most interest in the communications decision space because they said, "this tells me when to communicate." An operator commented that the communications decision space was "used to bring down negative emotion" on the part of the citizens, via using a twitter-like application to post information about the simulated disaster.

One operator who was particularly engaged with the communications decision space likened it to a video game called "whack-a-mole" in which virtual "emotional" moles pop up and must be pushed down into their holes. He saw great potential for the communications decision space to help identify topics that generate high levels of negative emotion from a lot of people, which would motivate him to communicate immediately on those topics (thus he "whacked the mole" to lower the negative emotion). The decision space could then be re-examined to see the effect of the communications.

The pace of his interaction with the communications decision space also reminded the operator of the whack-a-mole game because the decision space's implied recommendations to communicate about topics were replaced quickly by suggestions to communicate on other topics after the initial topic(s) had been addressed or something additional happened to divert the public's attention and shift their emotions. We observed that the workflow pace made it challenging for operators to spend much time applying their judgment to weigh the merits of communicating on one topic that evidenced a *high* degree of negative emotion by a *few* people, versus another topic that showed a somewhat *lower* degree of negative emotion on the part of *many* people.

On a positive note, the communications decision space tool did enable us to find topics that had not previously been addressed by the emergency responders, yet merited attention based on the negative emotion level. However, it seemed that the emotion content was lower than would be seen in a real-life disaster situation, most likely because of the simulated nature of the experiment. Further, we sometimes had difficulty in choosing a hashtag or keyword to assess, because several similar hashtags or keywords were being used for essentially the same topic; therefore, it was difficult to isolate the topics.

### Using the Resource Decision Space

As expected, it was the operators responsible for resource allocation who explored the resource decision space. In our conversations with these operators during the CERPS-SIMEX, we determined that, although they had real-world systems that provided some arrival time information, they had not previously had access to such comprehensive information about the likely impacts of their requests on lending organizations. These operators were used to learning about the potential impacts of borrowing requests by talking with their counterparts in other jurisdictions, but they would not always call *all* of their counterparts before each borrowing decision and so would not have a comprehensive picture of which organization would be impacted the least by a request. Of course, in time-sensitive emergencies, the arrival time for each asset is the primary consideration, but if two or more jurisdictions could respond with assets in roughly equal times, ideally the jurisdiction that could most easily spare the asset should be the one to loan it.

While we did provide training for both decision spaces, one purpose of the experiment was to allow users to explore how new tools could yield new concepts of operation. We believe that performance would have been enhanced with experience, and that this experience could inform the composition of an explicit, written concept of operations for incorporating resource decision space information into their decision making process.

The resource decision space was designed to complement other applications or portals that contained a corresponding situation space (facts about what resources were available and their characteristics). Not all of the systems that are currently used for situation awareness regarding resources were provided in the CERPS-SIMEX, however, so the resource decision space was designed to show some situation space information (i.e., the response time of the candidate resources) as well as decision space information (i.e., the likely impact to the lending organizations of sending those resources to the disaster). Nevertheless, some operators noted that the absent resource status information constituted "missing pieces" of the puzzle. This experience reinforced our view regarding the necessity of having both situation and decision space information.

## FUTURE WORK

We have four areas of focus for future work: 1) further develop the communications decision space prototype toward a more complete implementation of the concept, 2) testing in new task environments, 3) refining the methods for mining social media, and 4) extending the visualization platform to support deeper levels of option awareness.

For the first focus, the prototype can be incrementally moved toward a more complete implementation of actually forecasting the effect of communications. As an initial step, the current state of negative emotion about a topic could be replaced with a forecast of future negative emotion when a given communication alternative is not executed. For example, what would be the range of emotion if the public was not informed of how emergency responders plan to address evacuating students from the blast area. Such forecasts could be calculated based on extending the trend of negative emotions regarding the evacuation topic and how that trend would be modified based on communication history (both what actions were taken and their timing). A more ambitious forecast model could be based on social psychological theory to actually estimate what the range of emotion would be in response to an action being executed. This would require forecasting how the communication action would be perceived by individuals given their current emotional state.

For the second focus, there are many analogous task domains where decision makers need to respond to uncertain and emerging situations by employing multiple streams of situation data, which increasingly includes social media. One project currently under development is applying techniques to those used in the study reported here, for the purpose of epidemiological decision making for the Indiana State Department of Health. In this task environment, decision makers are trying to detect potential outbreaks of infectious disease and make decisions about strategic communication and recommending countermeasures. We have developed a preliminary data mining and visualization tool for mining health-related tweets for a specified geographic area, which is currently under evaluation with healthcare experts. Lessons learned by applying decision space visualization techniques in diverse task environments will lead to more robust and generalizable knowledge about supporting complex decision making using social media.

For the third focus, there is widespread attention to the complex socio-technical problem of making sense of unstructured social media data, such as tweets. Further research is necessary to determine the relative advantages and disadvantages of automated hashtag and keyword disambiguation via natural language processing and other computational techniques (Han & Baldwin, 2011; Michelson & Macskassy, 2010) compared to volunteer social media curators such as those who picked out the top hashtags and keywords in our experiment. In addition to the hashtags, keywords, emotion and sentiment in the twitter-like messages, geographic information is of particular importance for contextualizing pieces of information about complex emerging scenarios. Less than 1% of tweets have geolocation information (specific longitude and latitude data identifying the geographic origin of the tweet), and only one-quarter of Twitter users provide a specific city location in their profile (Cheng, Caverlee, and Lee, 2010). While techniques are under development to infer the location of a tweet from its content and metadata, such as keywords and time stamps (Cheng et al, 2010; Li et al., 2011; Burger et al., 2012; Topsy Labs, Inc., 2012), more work is necessary to understand the risks and rewards of approximated geographic data, as well as whether there might be important individual differences in motivation or reliability between Twitter users who do or do not have the geolocation feature enabled.

Finally, for the fourth focus, we have developed but not yet tested a tree-based display for visualizing the factors and relationships mediating the consequences of choosing one option versus another. This display employs strategies of “fast-and-frugal” trees, which often perform as well as complex benchmark strategies, but with less information. The tree display represents a hierarchy of underlying interacting characteristics of the scenario (including present and future concerns), ordered in descending level of influence on the outcomes under inspection. An important feature of the fast-and-frugal tree visualization is that users can readily memorize and rapidly recognize recurrent tree patterns under stressful circumstances, enhancing the efficient development of the decision maker’s option awareness.

## ACKNOWLEDGMENTS

We would like to thank John James, Paula Mahoney, Dan Potter, and Laura Strickhart for their help with the decision space; and Scot Lunsford, Mike Fulk, Tom Kiley, and Eric Levine for their help with integrating the decision space with the emotion indicators. This work was partially funded by MITRE Corporation projects 43MSR003-KA, 0712M810-AA, 51SPI010-AZ, 0712M840-AA, and 05MSR003-HA. Approved for public release, case 12-4860. 

**REFERENCES**

1. Bankes, S. (1993). Exploratory modeling for policy analysis. *Operations Research*, 41, 3 (May – June), 435 – 449.
2. Bonazountas, M., Kallidromitou, D., Kassomenos, P., and Passas, N. (2007). A decision support system for managing forest fire casualties. *J. of Environmental Management*, 84 (2007), 412 – 418.
3. Burger, J., Henderson, J., Kim, G., and Zarrella, G. (2011). Discriminating gender on Twitter. In *Proc. of the 2011 Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, UK, July 2011, 1301 - 1309. Available at: http://www.aclweb.org/anthology-new/D/D11/D11-1120.pdf.
4. Chandrasekaran, B. (2005). *From optimal to robust COAs: Challenges in providing integrated decision support for simulation-based COA planning*. Columbus, OH: Laboratory for AI Research, Ohio State U.
5. Chandrasekaran, B. and Goldman, M. (2007). Exploring robustness of plans for simulation-based course of action planning. In *Proc. of the 2007 IEEE Symposium on Computational Intelligence in Multicriteria Decision Making,* Honolulu, HI, April 2007, 185 – 192.
6. Campbell, B., Mete, H. O., Furness, T., Weghorst, S., and Zabinsky, Z. (2008). Emergency response planning and training through interactive simulation and visualization with decision support. In *Proc. of the IEEE 2008 International Conference on Technologies for Homeland Security*. IEEE, Waltham, MA.
7. Cheng, Z., Caverlee, J., and Lee, K. (2010). You are where you tweet: A content-based approach to geo-locating twitter users. In *Proc. of the 19th ACM International Conference on Information and Knowledge Management* (CIKM '10). ACM, New York, NY, USA, 759 – 768.
8. Defense Group, Inc., undated. CoBRA: For the new breed of first responders. Available at: http://www.defensegroupinc.com/cobra/, accessed 25 November 2012.
9. Delongueville, B., Smith, R. S., and Luraschi, G. (2009). "OMG, from here, I can see the flames!": A use case of mining location based social networks to acquire spatio-temporal data on forest fires. In *Proc. of the 2009 International Workshop on Location Based Social Networks (LSBN '09),* November 2009.
10. Dong, X. and Hayes, C. C. (2012). Uncertainty visualizations: Helping decision makers to become more aware of uncertainty and its implications. *J. of Cognitive Engineering and Decision Making*, 6(1).
11. Drury, J. L., Klein, G. L., Pfaff, M. S., and More, L. D. (2009). Dynamic decision support for emergency responders. In *Proc. of the 2009 IEEE Technologies for Homeland Security Conference*, Waltham, MA.
12. Elson, S.B., Yeung, D., Roshan, R., Bohandy, S.R., Nader, A. (2012). "Using social media to gauge Iranian public opinion and mood after the 2009 election." Rand Technical Report.
13. Endsley, M. R. (1988). Design and evaluation for situation awareness enhancement, In *Proc. of the Human Factors Society 32nd Annual Meetin*g, Santa Monica, CA, Human Factors Society.
14. ESi (2012). ESi: Boundless Collaboration. http://esi911.com/esi.
15. Gigerenzer, G. and Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats, *Psychological Review, 102(4),* 684 – 704.
16. Hall, D. L., Hellar, B. and McNeese, M. (2007). Rethinking the data overload problem: Closing the gap between situation assessment and decision making. In *Proc. of the 2007 National Symposium on Sensor and Data Fusion (NSSDF) Military Sensing Symposia (MSS),* McLean, VA.
17. Han, B., and Baldwin, T. (2011). Lexical normalisation of short text messages: Makn sens a #twitter. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1* (HLT '11), Vol. 1. Association for Computational Linguistics, Stroudsburg, PA, USA, 368-378.
18. Heverin, T. and L. Zach (2010). Microblogging for Crisis Communication: Examination of Twitter Use in Response to a 2009 Violent Crisis in Seattle-Tacoma, Washington Area. In *Proc. of the Information Systems for Crisis Response and Management Conference (ISCRAM 2010).* Seattle, WA.
19. Hoffrage, U., and Gigerenzer, G. (1998). Using natural frequencies to improve diagnostic inferences. *Academic Medicine, 73(5),* 538- 540.
20. Hoffman, J. R., Wilkes, M. S., Day, F. C., Bell, D. S., Higa, J. K. (2006). The roulette wheel: An aid to informed decision making. *PLoS Medicine*, 3(6): e137.

21. Kavanaugh, A., Hassan, R., Elmongui, H. G., Magdy, M., Sheetz, S. D., Yang, S., Fox, E. A., and Shoemaker, D. J. (2012). Between a Rock and a Cell Phone: Communication and Information Technology Use during the 2011 Egyptian Uprising. In *Proc. of the 2012 Information Systems for Crisis Response and Management Conference*, Vancouver, BC, Canada, April 2012. Available at: http://www.iscramlive.org/ISCRAM2012/proceedings/150.pdf
22. Klein, G. L., Drury, J. L., and Pfaff, M. S. (2012). Providing an option awareness basis for naturalistic decision making. *The International Journal of Cognitive Technology, 16*(2), 10 – 19.
23. Latonero, M. and Shklovski, I. (2010). "Respectfully Yours in Safety and Service": Emergency Management & Social Media Evangelism. *Proceedings of the 7th International ISCRAM Conference,* Seattle, USA,
24. Lempert, R. J., Popper, S. W., and Bankes, S. C. (2003). *Shaping the next one hundred years: New methods for quantitative, long-term policy analysis*. Santa Monica, CA: RAND.
25. Li, W., Serdyukov, P., de Vries, A. P., Eickhoff, C., and Larson, M. (2011). The where in the tweet. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management* (CIKM '11), ACM, New York, NY, USA, 2473-2476.
26. Liu, Y., Moon, S. P., Pfaff, M. S., Drury, J. L., & Klein, G. L. (2011). *Collaborative option awareness for emergency response decision making*. Paper presented at the 8th Annual International Conference on Information Systems for Crisis Response and Management (ISCRAM), Lisbon, Portugal, May 2011.
27. Martignon, L., Katsikopoulos, K. V., and Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology, 52*(6), 352-361.
28. Michelson, M. and Macskassy, S. A. (2010). Discovering users' topics of interest on twitter: a first look. In *Proceedings of the fourth workshop on Analytics for noisy unstructured text data* (AND '10). ACM, New York, NY, USA, 73-80. DOI=10.1145/1871840.1871852 http://doi.acm.org/10.1145/1871840.1871852
29. Mosier, K. L. (1997). Technology and "naturalistic" decision making: Myths and realities. In C. E. Zsambok & G. A. Klein (Eds.), *Naturalistic decision making* (pp. 41-54). Mahwah, NJ: Lawrence Erlbaum Associates.
30. Palen, L., S. Vieweg, S.B. Liu and A.L. Hughes (2009). Crisis in a Networked World. *Social Science Computing Review* 27, no. 4: 467–480.
31. Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). *The development and psychometric properties of LIWC2007* [LIWC manual]. Austin, TX: LIWC.net.
32. Pfaff, M. S., Klein, G. L., Drury, J. L, Moon, S. P., Liu, Y., and Entezari, S. O. (2012). Supporting complex decision making through option awareness**.** *Journal of Cognitive Engineering and Decision Making*, first published online on September 10, 2012 as doi:10.1177/1555343412455799, available at http://edm.sagepub.com/content/early/recent.
33. Phillips, J. K., Klein, G. A., & Sieck, W. R. (2004). Expertise in judgment and decision making: A case for training intuitive decision skills. In D. Koehler & N. Harvey (Eds.), *Blackwell Handbook of Judgment and Decision Making*, Malden, MA: Blackwell Publishing Ltd, 297 - 315.
34. Prolog Development Center (2010). ARGOS CBRN information system for emergency management. White paper version 0.8, June 2010. http://www.pdc.dk/Argos/downloads/ ARGOS_whitepaper.pdf, accessed 22 February 2011.
35. Qu, Y., P.F. Wu and X. Wang (2009). Online Community Response to Major Disaster: A Study of Tianya Forum in the 2008 Sichuan Earthquake. In *Proceedings of the 2009 Hawaii International Conference on System Sciences (HICSS 2009),* 1–11. IEEE Computer Society.
36. Sandman, P. (2003). Beyond panic prevention: addressing emotion in emergency communication. Emergency Risk Communication CDCynergy (CD-ROM). Centers for Disease Control and Prevention, U.S. Department of Health and Human Services. February 2003. Available online at http://www.psandman.com/articles/beyond.pdf.
37. Schulz, A., Paulheim, H., and Probst, F. (2012). Crisis Information Management in the Web 3.0 Age. In *Proc. of the 2012 Conference on International Systems for Crisis Response and Management*. Available at: http://www.iscramlive.org/ISCRAM2012/proceedings/160.pdf.

38. Shklovski, I., L. Palen and J. Sutton (2008). Finding Community Through Information and Communication Technology in Disaster Response. In *Proc. of the 2008 Conference on Computer Supported Cooperative Work (CSCW 2008),* 127–136. San Diego, CA: ACM.
39. Starbird, K., Muzny, G., and Palen, L. (2012). Learning from the Crowd: Collaborative Filtering Techniques for Identifying On-the-Ground Twitterers during Mass Disruptions. In *Proc. of the 2012 Conference on International Systems for Crisis Response and Management*. Available at: http://www.iscramlive.org/ISCRAM2012/proceedings/148.pdf.
40. Starbird, K. and L. Palen (2011). "Voluntweeters:" Self-Organizing by Digital Volunteers in Times of Crisis. In *Proc. of the 2011 Conference on Human Factors in Computing Systems (CHI 2011),* 1071–1080. Vancouver, BC: ACM.
41. Starbird, K., Palen, L., Hughes, A. L., and Vieweg, S. (2010). Chatter on The Red: What hazards threat reveals about the social life of microblogged information. In *Proc. of the ACM Computer Supported Cooperative Work (CSCW) 2010 Conference,* Savannah, GA, February 2010, 241 – 250.
42. St. Denis, L. A., Hughes, A. L. and Palen, L. (2012). Trial by Fire: The Deployment of Trusted Digital Volunteers in the 2011 Shadow Lake Fire. In *Proc. of the 2012 Information Systems for Crisis Response and Management Conference*, Vancouver, BC, Canada, April 2012. Available at: http://www.iscramlive.org/ISCRAM2012/proceedings/150.pdf.
43. Topsy Labs, Inc. (2012). Topsy home page. Available at: http://topsy.com.