

Social Radar Workflows, Dashboards, and Environments

Jennifer Mathieu^{1*}, Michael Fulk¹, Martha Lorber¹, Gary Klein¹,
Barry Costa¹, and Dylan Schmorow²

¹ MITRE Corporation
Bedford, MA 01730, U.S.A

² Office of the Assistant Secretary of Defense Research and Engineering
4800 Mark Center Drive, Suite 17E08
Alexandria, VA 22350-3600, USA

jmathieu@mitre.org
mfulk@mitre.org
mlorber@mitre.org
gklein@mitre.org
bac@mitre.org
dylan.schmorow@osd.mil

ABSTRACT

To build a global “Social Radar,” an integrated set of capabilities supporting strategic and operational level situation awareness, alerts, and option awareness, there is a need for an overarching enterprise approach. Social Radar’s objective is to demonstrate a useful, mission focused, end-to-end environment—a dashboard of socio-cultural indicators created from online news, blogs, and social media data processed at scale, as well as decision support tools. An enterprise focused testbed for early transition of sociocultural tools requires a data-to-decision support system. Such a system needs to provide tools to allow analysts to tailor and weight the fusion of indicators, to use online sources to update simulation model parameters to evaluate courses of action, and to use outcomes of course of action models to provide quantitative metrics for indicator integration strategies. This large scope requires an analysis environment that supports the development of common output measures, management of uncertainty analyses, and system evaluation and validation. Making these items requirements from the start ensures that Social Radar is addressing the most challenging aspects for use of sociocultural data and tools to support missions.

Based on ten months of work on the Social Radar prototype, the following capabilities were built: (1) federated search of all available widgets for a topic, (2) hotspots on a globe for instability monitoring, (3) ability to analyze datasets with timelines and topic cloud visualizations, (4) news search capability, and (5) use of various other analytic drilldown tools. This demonstrated the possibilities of a Social Radar when fully mature. Business process lessons learned and next steps are shared for rigorous use of analytic tools and series of tools (workflows); data, indicator, and model outcome visualization strategies (dashboards); and supporting architecture for data processed at scale and near-real time monitoring (environments).

1.0 CHALLENGES

Unanticipated social events have corollary and perhaps precursor indicators in open source news, blogs, and social media that if monitored could potentially provide warning prior to these events. Driving questions for

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012
4. TITLE AND SUBTITLE Social Radar Workflows, Dashboards, and Environments		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MITRE Corporation, Bedford, MA, 01730		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		
14. ABSTRACT To build a global ?Social Radar,? an integrated set of capabilities supporting strategic and operational level situation awareness, alerts, and option awareness, there is a need for an overarching enterprise approach. Social Radar?s objective is to demonstrate a useful, mission focused, end-to-end environment? a dashboard of socio-cultural indicators created from online news, blogs, and social media data processed at scale, as well as decision support tools. An enterprise focused testbed for early transition of sociocultural tools requires a data-to-decision support system. Such a system needs to provide tools to allow analysts to tailor and weight the fusion of indicators, to use online sources to update simulation model parameters to evaluate courses of action, and to use outcomes of course of action models to provide quantitative metrics for indicator integration strategies. This large scope requires an analysis environment that supports the development of common output measures, management of uncertainty analyses, and system evaluation and validation. Making these items requirements from the start ensures that Social Radar is addressing the most challenging aspects for use of sociocultural data and tools to support missions. Based on ten months of work on the Social Radar prototype, the following capabilities were built: (1) federated search of all available widgets for a topic, (2) hotspots on a globe for instability monitoring, (3) ability to analyze datasets with timelines and topic cloud visualizations, (4) news search capability, and (5) use of various other analytic drilldown tools. This demonstrated the possibilities of a Social Radar when fully mature. Business process lessons learned and next steps are shared for rigorous use of analytic tools and series of tools (workflows); data, indicator, and model outcome visualization strategies (dashboards); and supporting architecture for data processed at scale and near-real time monitoring (environments).		
15. SUBJECT TERMS		

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 26	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

a Social Radar, an integrated set of capabilities to support strategic and operational level situation awareness, alerts and option awareness, include:

- What are the measurable indicators from online news, blogs, and social media?
- When are the indicators observable in relation to events?
- How do the indicators change over time?
- How can these indicators be used to support alerts and analysis to aid in decision-making?
- How can we generalize the analytic procedures for different events across the globe?
- What organizations would likely use Social Radar and for what missions?

Workflows, dashboards, and environments to support the analyst in gaining situation awareness, using and tailoring alerts, and making use of near-real time information in decision support tools to obtain options awareness for unanticipated social events have the following challenges:

- What business process workflow strategies are appropriate for an analyst using sociocultural tools, and how to ensure the rigorous use of the tools and series of tools?
- What visualization strategies and dashboards are best for data, indicators, and model outcomes?
- What environments / frameworks / architectures should be used?
- What scalable data strategies can be used to populate the architecture?
- What sustainment strategies are needed for such environments?

2.0 GOAL

Social Radar's objective is to demonstrate a useful, mission-focused, end-to-end environment—a dashboard of sociocultural indicators created using online news, blogs, and social media data processed at scale, as well as tools that support decision-making. This will be done by creating an environment in which the analyst can access: representative data in near-real time; tools that support tailored alerts; and business process tools to create workflows using specific sociocultural tools. These capabilities (tools, models, and forecasts) would allow the analyst to explore data, perform diverse analyses, generate products for decision-makers, and to help communicate analyses through tailored dashboards that support drilldown and knowledge management.

Social Radar is being built as an enterprise-focused testbed for early transition of sociocultural tools. To this end, the Ozone Widget Framework and a supporting architecture is being used. This end-to-end, data-to-decision support system is being built to provide tools to allow analysts to tailor and weight the fusion of indicators, to use online sources to update simulation model parameters to evaluate courses of action, and to use outcomes of course of action models to provide quantitative metrics for indicator integration strategies. This large scope requires an analysis environment that supports the development of common output measures, management of uncertainty analyses, and system evaluation and validation. Making these items requirements from the start ensures that we are addressing the most challenging aspects for use of sociocultural data and tools to support missions.

Designing analyst workflows requires a deep understanding of the analytic task and mission. We are beginning to work with analysts to understand what tools would help them and the kinds of business processes they might use to analyse online data sources. Once the tasks and missions are understood we can

adapt Social Radar to support analyst workflows—especially rigorous, easy-to-use methodologies for advanced sociocultural capabilities. For Social Radar to be adopted and for it to be used appropriately, we need to make it useful, easy to use, and convey the rationale that underpins the sociocultural tools. We are focusing on identifying common tasks and planning more broadly for indicator integration. The Social Radar system can also be designed to collect data on system use (searches done, tools used, tools used together, products created). And of course, Social Radar needs to be continuously refined to improve the analyst experience, potentially through the use of an evaluation application or widget.

Based on ten months of prototype development, lessons learned will be shared in these categories:

- Business processes for rigorous use of analytic tools and series of tools (workflows)
- Data, indicator, and model outcome visualization strategies (dashboards)
- Supporting architecture for data processed at scale and near-real time monitoring (environments)


3.0 WORKFLOWS


Social Radar was used to conduct a thorough analysis of one historic events—2011 riots in the United Kingdom (UK)—and one post-election protest. We selected the events and then obtained relevant news, blog, and Twitter data. Retrospectively, based on the specific findings, we created a story as if the Social Radar tools were being used as events unfolded. We acknowledge that it is simpler to create a story in hindsight for any event than it is to tell the story as events are unfolding. Table 1 and Table 2 show the resulting tool workflows and specific results for the 2011 UK riots (Table 1) and election protest (Table 2).




3.1 2011 United Kingdom Riots Historical Case Study

Social Radar was applied to the August 2011 UK riots. This event was selected based on the fact that our access to Twitter began in mid-July and our proof-of-concept for global instability hotspotting was needed in October. Here are the details surrounding the event. Just after 6 p.m. on August 4th, 2011, Mark Duggan was shot by police in London (Tottenham). On the evening of the 6th, a peaceful protest was held outside the Tottenham police station. After a few hours violence broke out. Buildings and vehicles were set on fire. Sixty-one people were arrested and 26 officers were injured [1]. Table 1 is organized by: (1) the name and/or description for the tool(s) used, (2) orientation or results found using the tool(s), and (3) a representative screen shot. The ‘Specific Results Using Tools’ column in the table continues this paper’s narrative.

Table 1: 2011 United Kingdom (UK) Riots Historical Case Study

Tool(s) in Workflow	Specific Results Using Tool(s)	Social Radar Widget(s) Screen Shots
Social Radar Search Widget Global Instability Hotspotting Dashboard	The global instability hotspotting dashboard depicts colors in red (unstable), yellow (moderately stable), and green (stable). Red signifies current negative change, unrest, or instability at the nation state level. Yellow is a moderate level of instability, and green is no change or stable. The hotspotting layer (circled) gives us an aggregate view of instability.	
Social Radar Search Widget Global Instability Hotspotting Dashboard Search on 'UK' from July 16 to August 10, 2011	The UK is red indicating unrest based on an aggregate of 3 analyses (top table): 1. Political Instability Task Force (PITF) forecast model; based on annual data (green, stable); 2. Sentiment Analysis and forecast based on news (red, unstable); and 3. Emotion Analysis based on Twitter (yellow, moderately stable). Supporting results about the UK are shown in the bottom table.	
Social Rader Search and Data Layers Widgets Click on PITF 'Instability Model' Then Click on 'Instability Model' then 'PITF Model'	The UK is green based on the PITF forecast model [2] annual data for infant mortality rate, percentage of state led discrimination, number of countries in armed conflict in bordering states, and regime type. Each of the input indicators can be viewed on the map for the 20 countries. These indicators were updated as much as possible for 2011.	



Sentimeditr Widget Sentiment Analysis Drilldown Tool Click on 'UK Sentiment' in bottom table	The UK is red based on the rate and direction of change in sentiment on topics [3] relating to instability. UK-sourced Open Source Center news (processed at scale) was searched on the topics of leaders, domestic politics, terrorism, human rights, and dissent. Sentiment analysis shows steadily declining sentiment regarding economic issues starting at the end of July 2011—before the riots began.	
Social Rader Search and Data Layers Widgets Click on Twitter 'Instability-Sensitive Instability Detector' Then Click on 'Emotion Analysis' then 'Twitter'	The UK is yellow based on emotion analysis of profanity and instability-related words [4] expressed in a random 10% of the Twitter feed (processed at scale) [5]. The bulls eye (size indicates number of tweets) layer is just the profanity indicator, and it shows that the relative volume increased starting on August 4th (before the shooting), is very high on the 5th (before the riots), and remains relatively high through the 10th.	
Emotion Graph Widget Click on Emotion Graph 'UK Riot Emotion'	The data illustrated using the bulls eyes on the map can also be visualized in a timeline widget. The large increase in relative profanity words starts on the 6th, peaks between the 8th and 9th, and returns to the baseline on the 10th. The shaded and white areas on the graph are determined by an automated algorithm to detect change in specific indicators or a combination of weighted indicators [4].	
Comment Filter (CoFi) Widget Click on Topic Clustering 'UK Protests – Duggan Tweets'	To investigate the small peak on the 5th to see if there were early indications of unrest, the topic clustering tool [6] was used on tweets obtained using the name 'Duggan.' This dataset led to the first tweet (in a random 10% of the Twitter feed) related to the shooting. It occurred 24 hours after the shooting—on the 5th. The shooting was being discussed on Twitter before the riots.	

Topic Cloud Widget Click on Topic Cloud 'Twitter UK Topics'	The topic cloud was used to identify emerging hashtags after the shooting (processed at scale). On August 5th, 'jobs' is the most prominent hashtag. On the 7th, 'Enfield' and 'Tottenham' stand out. Riots occurred in those locations on that day. On the 10th, 'jobs' is the most prominent hashtag again. For the period 4-9 August, 'London Riots' is the most common hashtag—the discussion continued.	
Emotion Graph Widget Click on Emotion Graph 'UK Riot Emotion'	We know that the first tweet related to Duggan was on the 5th, however, on the 4th and 5th, the automated breakpoint analysis detected a rise in negative emotions based on a dataset related to terms denoting government. This analysis could be the basis of an alert [4]—it would have triggered the day Duggan was shot (grey to white shading on graph) as well as the day after. A few new breakpoints indicate a change in mood, and they occurred after the shooting but before the riots.	
COA Models Widget Select Course of Action (COA): Baseline (ground truth) Do Nothing Moderate Response Aggressive Response	On the 9th, the police presence in the UK went from 6000 to 16,000 officers [7] and the unrest ended. The State Stability System Dynamics Model (S3DM, [8]) was used to represent the UK Riots using news and reports [7, 9]. The model output is number of protesters (shown), rioters, and arrests. Course of action options included the baseline (ground truth), do nothing, moderate response (earlier response with less police), and aggressive response (deploy police in 12 hours)--illustrating 'what if' COA analysis.	

3.2 Election Protests Case Study

Social Radar was also applied to a protest following an election. Table 2 is organized by: (1) orientation or results found using the tool(s) and (2) a representative screen shot. The 'Specific Results Using Tools' column in the table continues this paper's narrative.

Table 2: Election Protests Case Study

Specific Results Using Tools	Social Radar Widget Screen Shots
Social Radar was used to analyze election protests. Emotion expressed in Twitter was monitored. The automated change detection algorithm was used and detected a change [4]. There was an increase in emotion in tweets as news spread about the death of a dissident. Emotion indicators in native language blog posts were also peaking right after the dissident's death [10]. Detection of topics in blogs indicated anger from two sides. Opposition bloggers were discussing the protest movement favorably and were criticizing the regime. On the other side, regime-supporting bloggers were strongly criticizing the opposition movement [10]. A week before the protests, these tools could have detected a sharp increase in population sadness based on Twitter and native language blogs.	 Twitter Blogs
Another change could have been automatically detected [4], indicating increased activity in sharing of information through forwarding tweets or retweets [12]. Sending the Twitter data to the clustering tool [6] to group and read some of the tweets revealed that three days before the protests started, a single tweet was rapidly and widely shared. It was a link to a map showing protest routes. So, an alert based on the indicator retweets using the automated detection algorithm could have helped analysts discover critical information three days before the protests began.	
With this understanding, US decision makers might have explored notional alternative diplomatic courses of action to impact unfolding events. The National Operational Environment Model [The NOEM, 11] was used to represent the election protests using available news and reports. One model output is number of activist (bottom graph). Options include four diplomatic approaches (top graph)—illustrating ‘what if’ COA analysis.	 US Diplomatic COA Activists Over Time

3.3 Workflow Lessons Learned

The most important lesson learned from the two historical case studies is the important integration work that is required to provide a sensible workflow uniting disparate projects. When tools are built by a team focused on one particular task, it is not always painless or simple to bring that work into the existing system. A coordinating layer of systems engineering should manage the widget and service interactions at a top level. Sufficient systems engineering is needed to avoid confusion between different projects. There is power in using widget frameworks that allow interoperability and communication between different tools, but the caveat is that this must be planned and system-engineered to be most effective.

It also is necessary to determine the potential analytic workflows that Social Radar will support. The broad scope of Social Radar from data to decisions—using online-based indicators and simulation models (Figure 1)—requires an analysis environment that supports rigorous procedures for using a series of tools together. For example, when the three models—PITF forecast model, sentiment analysis forecast, and emotion analysis—were used together to implement the global instability hotspotting layer, a common normalized output measure for the analyses had to be defined so they could be mapped to the red, yellow, and green ranges for change over time. The goal was to provide an analyst an aggregate visualization at the nation state level that represented the combination of available analyses. The tools had to support the analyst workflow of seeing the hotspot layer as well as drilling down into the details to discover why a particular country was a certain color. Research is needed to effectively develop common output metrics for specific cases and the results need to be analyzed to generalize the lessons learned. With common, normalized output metrics established, the uncertainty of the analytic results needs to be characterized and visualized appropriately. The visualization of indicators over time provides situation awareness by enabling the analyst to visually compare short-term with long-term trends in a region.

In addition, exploratory modeling [13] of courses of action with simulation models like S3DM and the NOEM allow the landscape of plausible options to be visualized. This visualization provides information describing the decision options and their desirability relative to one another, what Hall et al. [14] defined as a decision space. The extension to this decision space of perception, comprehension, and projection, which Endsley [15] defined for situation awareness, has been called option awareness [16, 17]¹. A visualization of the decision space provides the decision maker with something akin to night-vision goggles for the mind, allowing the decision maker to actually see otherwise obscured relationships between options rather than requiring them to mentally simulate each one. By making choice a perceptual comprehension process, it enables decision makers to apply their more powerful visual, pattern matching, recognition capabilities rather than the more limited mental simulation process. In addition, the decision space can be data mined to identify those factors that lead to better and to worse outcomes. With the level of option awareness this enables, decision makers can craft new courses of action that mitigate the bad factors and facilitate the good.

Bringing these approaches together in one tool suite, a workflow can be developed that links the assessment of online data-based indicators to course of action analysis through exploratory modeling, as shown in Figure 1. This was the case for the UK Riots, where a short term (days) assessment or forecast was done with sentiment analysis, which could be compared with a long term (two years) assessment or forecast done using the PITF forecast model. With the criticality of the problem identified (situation space), the S3DM simulation was used to evaluate four courses of action (baseline or ground truth, do nothing, moderate response, and aggressive response) with a forecast of the number of protesters, rioters, and arrests in each case over the next few days. In this way, the situation space (regarding the facts, their implication for the criticality of the current situation and the near future) was linked with the decision space (viable options and the plausible outcomes of each option). This linkage provides the analyst with not only situation awareness, but also option awareness—comprehension of the landscape of plausible futures linked to viable options—which will allow decision makers to identify the most robust option that will work across the broadest swath of those futures [17, 19].

When Social Radar indicators are used in simulation models to evaluate courses of action the uncertainty needs to be propagated in the model and visualized as part of the simulation model results. As indicators are developed and models adapted to use them to evaluate courses of action, it will be necessary to create

¹ Research has shown that displaying decision spaces does, indeed, provide option awareness [16, 18]. In addition, Pfaff et al. [18] demonstrated that under circumstances of deep uncertainty these decision space visualizations not only enabled decision makers more often to identify robust options, but to make decisions faster and with more confidence than unaided decision makers.

rigorous workflows. The necessary workflow must include, but is not limited to, the following steps: translate indicators with their uncertainty into model components; verify it is done correctly; design and run sensitivity analysis; determine output value scoring function; determine best visualization (proof of concept to this point, [20]); use ‘analyst in the loop’ experimentation to obtain feedback on tool transparency and workflow ease of use; and use subject matter experts in the loop experimentation to help with system validation (illustrated in Figure 1). Such a workflow will also provide metrics for indicator integration and indicator translation. For example, if varying how indicators are defined does not impact the decision space, it will tell us either that the models are not sensitive enough or the indicator variation does not make a difference to the decision under consideration. If the models are not sensitive to the social indicators that analysts and social scientists believe are significant to a particular area of interest, then the model may need to be localized and validated to the particular area. However, if the model has been localized and validated for the area, collecting these indicators is not useful to decision making. By indicating the factors that are critical to better and worse outcomes, the models also clarify the types of data that needs to be collected.




Figure 1: Social Radar Scope—Using Online Based-Indicators and Simulation Models to Support Situation Awareness, Alerts, and Option Awareness

4.0 DASHBOARDS

Social Radar has developed a Global Instability Hotspotting Dashboard (widget) that combines a map and two categories of results, as shown in Figure 2. All of the drilldown tools can be launched from this one widget and federated search (information retrieval technology that allows the simultaneous search of multiple searchable resources) is used to obtain all widgets with information for a particular country searched.




Figure 2: Social Radar Global Instability Hotspotting Dashboard

4.1 Dashboard Lessons Learned

Widgets generally focus on the data retrieval, data analysis, and data visualization. This dashboard has all three components. For example, monitoring tools are complemented by analytic tools that allow analysts to drill down in the data, visualization tools that aid in understanding, and other helpful widgets an analyst might want while performing their work. Common functionality identified for the workflow includes, but is not limited to:

- Custom and saved data searches
- Alerts, messages, push notifications
- Monitoring on near-real time data feeds
- Status on long running analytic tasks
- Modeling and simulation analysis
- Report generation

Creating a domain independent Social Radar requires an understanding of how analysts from various fields search data, use tools, and create products. The Global Instability Hotspotting Dashboard was a useful widget, but a more versatile interface is needed for Social Radar—one that enables business processes in support of many analytic tasks for multiple domains. Social Radar is also a set of analytic widgets that are deployed in the Ozone Widget Framework (Ozone), but it is possible that the Social Radar widgets will need to be deployed inside another widget marketplace or Ozone instance, and that these tools will need to co-

exist with other analytic tools for other domains. An analyst working in this environment may have issues understanding how Social Radar tools are meant to be used, and how to get started using them.

The dashboard for Social Radar must be designed so the analyst can get started quickly and find relevant data and tools. Analyst-centered design must be used to allow the capabilities from the common functionality list above to support novel business process workflows. This dashboard must be designed to clearly show the data sources, analytical services, options to compose workflows, and options to create visualizations that can be used to create analytic products and custom dashboards. Figure 3 shows a mockup of such an interface.




Figure 3: Social Radar Interface Mock-Up [21]

A major lesson learned is that each data source needs to be apparent and searchable. Social Radar has multiple data sources and should not attempt to exclusively search them in a federated manner (see section 5.2), but rather allow the analyst to explore the various data sources through their respective interfaces. Another mechanism for assisting the analysts is saved searches. Using the search tools for each data source, the analysts should have the ability to save their searches and return to them later. This was seen as a crucial capability for repeatability of the analyses. Using the saved searches also allows analysts to send the data from a saved search to analytic services that would transform or analyse the data.

Capabilities of the tools available must be made clear and the options for creating tool workflows must be easy to understand. This problem is not unique to the Social Radar widgets. It exists for all sets of coordinated widgets within environments because accessing the widgets from the toolbox has no apparent order. Choosing among many small applications to perform larger jobs requires some widget organization based on knowledge of the analytic workflow. To better understand this problem, Social Radar should use an analyst-centered design to identify problems that an analyst faces when using Ozone and to formulate a plan to alleviate those problems as much as possible. Finally, the dashboard must be readily customizable by the analyst to aid in the understanding of the environment.

5.0 ENVIRONMENT

The main goal of the Social Radar environment is to demonstrate a modular, enterprise architecture of small analytic tools working in conjunction with one another. The analytic tools currently work on online data sources that have been previously collected and stored, as well as tools that work on manually updated data that supports evaluating courses of action. The analytic tools allow access to and help with visualizing the data and support for evaluating courses of action. For example, techniques such as sentiment and emotion analysis are applied to the data sources in a way that helps monitor and detect changes that may indicate an analyst should investigate more deeply. The monitoring tools are complemented by analytic tools that allow analysts to drill down into the data and visualization tools that aid in understanding. Awareness of the architectural challenges associated with this goal is essential for making a robust, enterprise Social Radar prototype.

The user interface environment for Social Radar is the Ozone Widget Framework (Ozone). This is a lightweight framework that wraps web applications and exposes them to the analyst as small applications or widgets inside of a web browser. Additionally, Ozone allows the widgets to communicate with one another via provided publish and subscribe channels.

Based on 10 months of work on the Social Radar prototype, the following capabilities were built: (1) federated search of all available widgets for a topic, (2) hotspots on a globe for instability monitoring, (3) ability to analyze datasets with timelines and word cloud visualizations, (4) news search capability, and (5) use of various other analytic drilldown tools. This demonstrated the possibilities of a Social Radar when fully mature.

The following subsections describe the architectural lessons learned from building the prototype. Section 5.1 focuses on integration and development of the tools. Section 5.2 covers the Service Oriented Architecture (SOA) and specific components that proved valuable in the prototype. Section 5.3 discusses enabling access to Social Radar from web browsers and mobile devices. Section 5.4 includes recommendations for improving the agile development process. Finally, Section 5.5 addresses maintaining a Social Radar as tools and technologies evolve.

5.1 Analytic and Developer Tools Lessons Learned

The advantages of widget frameworks and Ozone in particular, include: (1) powerful capability by bringing a group of small analytic tools together to work in combination on specific problems; (2) rapid development by wrapping existing web applications in a widget that can be used inside of a web browser; and (3) flexible development by allowing widgets to communicate with one another via publish and subscribe communication channels using JavaScript libraries.

Because the widgets display as tools with icons in the Ozone launch menu (the widget palette), it is tempting to think that it is easy to share or move the individual widgets as individual units. This is not the case. The widget is a web application that has been wrapped in a frame to be displayed in the Ozone environment. It cannot be traded easily without having the underlying data and services that it relies on in place. The Social Radar widgets rely on the data and services from the Social Radar architecture to function properly. Social Radar has used the Ozone concept of a group to assemble all of the Social Radar widgets into one common group. The widgets should almost always be imported and used together as there are dependencies between them that would make copying subsets of the group problematic. They can, however, be used in conjunction with other groups of analytic widgets or even other frameworks.

To support a project like Social Radar that has many analytic widgets, developer tools are important for productivity. A normal developer workflow might include editing a web application, deploying it to a local web server, testing it in the local browser, adding this widget to a local ozone instance, and then testing it as a widget. To streamline the process, Social Radar automated much of this using Eclipse and custom build scripts (small programs to automate the compilation and other tasks in making an executable computer program). Using Eclipse, a developer is able to edit web application files, and deploy automatically to the Ozone Jetty server. Simply pressing the Jetty run button deploys Ozone and all the widgets that a developer specifies in a spreadsheet. This allows a developer to build and test many web applications quickly and efficiently for Ozone. Expanding upon this idea, Social Radar used tools to automatically build Eclipse projects for a developer if they already have either html or Java files. Finally, using GIT (distributed revision control system) as a repository allowed easy deployment to servers.

5.2 Service Oriented Architecture Lessons Learned

The high-level Social Radar architecture is shown in Figure 3. It consists of a presentation layer (green), service layer (blue), and data layer (orange). Dividing the system into these layers provided many benefits, two of which are highlighted here. First, because the presentation layer uses standard network communication mechanisms, only the presentation layer needs to be adapted when making the Social Radar system accessible via a new device (operating system, browser, mobile device, etc.). For example, a mobile application can use the same data access services as a thick client application to provide visualizations. Second, by adopting domain-appropriate formats and mechanisms for communication between layers, the likelihood of changes to a component in one layer requiring additional software development effort to make corresponding changes to components in other layers are reduced. For example, adding a new type of database to the data layer does not require changing components in the presentation or service layers as long as the new database conforms to the existing messaging format used by other components in the data layer. This design principle makes for a cleaner and more maintainable system design.

The presentation layer contains Ozone widgets responsible for exposing analysts to the services and data developed by technologists. The types of services (method of communication between two system components) include (1) web services that return a result over the network synchronously, (2) services that run at the location where the data is stored, and (3) services that run in near- real time. To support the set of services, the data layer must handle data sources that are, in some cases, very large, and it must make these data sources accessible to multiple services. Early in the prototyping effort, many projects maintained separate databases, resulting in a duplication of effort and increased difficulty in reusing code and data across services. Although each project must invest effort to convert to using a common data service solution in the data layer, steps taken to date to unify the data sources have already reduced redundant effort and provided opportunities for combining data that enables developing more powerful tools to support analysts. The lesson learned is that a common data service that provides multiple projects with a common access mechanism simplifies data management, improves the reusability of code, and enables leveraging heterogeneous data sources to develop more powerful tools.

The three categories of web services listed in the preceding paragraph vary in their response time depending on the complexity of the requests they receive. A web service may respond slowly to a request for statistical data derived from social media sources describing changing trends in sentiment across the entire globe and over an entire previous year, while the same request for a single country for the past month will return nearly instantaneously. This requires coordination between all three layers to provide the analyst with a user-friendly interface. When an analyst uses an Ozone widget to search for data and then requests that a service

process the search results, the analyst expects feedback on the status of the request. When the request can be fulfilled nearly instantaneously, the presentation layer can use a synchronous call. For example, the automated breakpoint algorithm's service, which is used to detect significant shifts in time series data, processes requests very quickly. However, the social media querying tools may require minutes or hours to complete a sufficiently complex request. As a consequence, the presentation layer must also support asynchronous communication with the services, so an analyst may receive notifications that a request will take a significant length of time to fulfill, receive periodic status updates on the request, and perform other actions while waiting for results.




Figure 3: Social Radar Architecture

Another lesson learned using the Ozone widgets for the presentation layer involves the use of common JSON data formats for visualizations and for passing messages in the Ozone communication channels. Services built by individual projects for accessing data benefit from a common data access mechanism that can return data in a JSON format, both in terms of reusability of this access code as well as minimizing the coupling of services to data or to visualization widgets. Using a common, data-agnostic JSON format also eases the passing of data between widgets. For example, Social Radar uses a common timeline widget that displays data received from any other widget, as long as the data obeys the common JSON format. This allows multiple widgets to send results to the timeline widget for display, and also enables results from different services to be displayed simultaneously on a single timeline for visual comparison. The common JSON format significantly helps decouple the presentation, service, and data layers.

A significant portion of the Social Radar data falls into the big data category, as it is composed of billions of small text messages. Technologists may desire access to subsets of this data that are prohibitively large for transferring over the network. For example, a researcher may perform a query that will produce a result set containing several million text messages. An analyst may want to perform the same query, but only see aggregate results produced by word counting algorithms that compute sentiment and emotion scores, rather

than viewing many millions of messages. In both cases, ideally, the service code representing the technologist's analysis, or the service which implements the analyst's word counting algorithm, would reside on the same servers where the data is stored. This eliminates the need to send the data across the network. One method for performing processing on the same servers that contain the data is using column-oriented databases (NoSQL) that allow for scalable data processing and storage, and allow the use of MapReduce (software framework to support distributed computing on large datasets on clusters of computers) to parallelize the computations performed by services. Social Radar would like to explore the use of these tools in the future since they may yield superior performance in the service and data layers.

Displaying hotspots on the globe is a prominent feature of the Social Radar presentation layer. For near-real time online data sources, an incoming data stream needs to be monitored continuously to update the hotspot visualization. The time between receiving new data and presenting it on the visualization is latency. The latency can be minimized by enhancing the existing architecture to allow certain services to process incoming social media data in parallel with storing the social media data. Social Radar is experimenting with the Yahoo S4 package to split the incoming data, allowing services to process it to produce updated, global hotspot visualizations based on emotion scores with minimal latency.

As the number of supported widgets in the Social Radar prototype grows, it has become clear that informing the analyst of each widget's availability and applicability to the user's task is important. Social Radar implemented a simple form of federated search in the early prototype that allowed the user to perform a search in a common widget, and then presented the user with a list of each individual widget that had results, as well as links to the results. For example, a user query for instability in a certain country might include the emotions indicator widget's results, the sentiment indicator's results, and results from a widget which searches for related news articles. However, this approach becomes cumbersome when queries are run on large data sets, since some widgets or services take a long time to return results leading to a user missing widgets that were simply taking a long time to respond to the query.

Perhaps most importantly, the federated search solution does not provide an analyst with an overview of the capabilities of each tool, or of the data sources available. This awareness is needed by the analyst in selecting analytical approaches most likely to yield operationally useful results. For example, the query described earlier produces interesting results and may eventually enable the analyst to understand the causes of the instability, but an analyst with greater awareness of the tools and data sources might more efficiently and robustness find information by restricting the search to social media data that is not controlled by the state, by excluding state-run news sources, and then sending the data to a service which performs automatic clustering to summarize the common themes. These themes may then be fruitful for further, more detailed searches to understand the causes of the instability. To address these concerns, the next iteration of Social Radar is building a widget that serves as starting point for the analysts, providing more detailed information and help. This will work in conjunction with the simple federated search approach.

5.3 Devices Lessons Learned

One of Social Radar's goals is to support analysts accessing data and services from multiple devices. As mentioned earlier, the separation of layers in the service oriented architecture makes it possible to support new devices by providing an adapted presentation layer. Using modular web applications inside a web browser as the presentation layer makes it more portable across platforms, and allows users to customize their Social Radar interfaces, picking and choosing which widgets should be displayed and organizing them on their desktops. However, this approach does not always take advantage of the strengths of specific

devices. For example, touch based mobile applications may provide a better alerting mechanism or even touch based data exploration in the drilldown stages. Because devices have different strong points, the plan is to leverage their strengths by making native applications for each platform and to allow these applications to work in conjunction with one another. The strengths of the various devices are listed below.

- **Native applications** are fast since they have direct access to hardware. Large applications with advanced capabilities that need big screen monitors are supported.
- **Web applications** are easy to manage and distribute. They have a common browser interface. Full featured software is supported.
- **Widget applications** are small and have a dedicated purpose. Multiple widgets can run in a browser.
- **Phone applications** are always in your pocket so collecting field data is easy. Touch interface is natural to use.
- **Tablets applications** have a bigger screen than a phone but still have a natural touch interface. Viewing imagery, maps, reading documents, and collaboration are strengths.

More generally, while Ozone provides a nice set of web-based analytic applications that can be used to explore data, Ozone is limited to the web browser. This limits usability, development, and performance on non-browser oriented platforms. With the explosion of mobile platforms, it is not difficult to imagine analysts will want to explore and use data on familiar mobile hardware in addition to web tools. Customized presentation layers could be developed for mobile devices while reusing the existing service and data layer capabilities for saved searches, alerts, and other tools. Taking advantage of the strengths of touch based and mobile devices by developing customized presentation layers is a possible avenue for further exploration. Mobile applications built to access data via the Social Radar services can leverage the infrastructure for saved searches, alerts, and other tools, as illustrated in the Figure 5 mock-ups.

5.4 Agile Software Process Lessons Learned

Using an agile software development process allowed software developers to build widgets quickly, and to respond to new or changing requirements while minimizing the cost and schedule impact. During the 10 months of prototyping Social Radar, one of the most useful agile tools proved to be developer scrum style meetings held twice a week. These meetings allow each developer to list his or her current tasks and voice any existing issues. It encourages participation and provides a way for those facing problems to get tips and advice from others who might already have seen similar issues. Social Radar had developers in widely separated geographical locations and used telecoms and screen sharing applications to support collaboration in the scrum meetings. In the initial, rapid prototyping effort, building development tools to streamline the process of integrating new capabilities with Ozone was essential to responding quickly to customer feedback. Because of the rapid build and deployment system, developers were able to focus on building widgets that operated together, listening to the feedback from users, and then iterating on the designs. This process, combined with frequent collaboration and input from technologists and researchers, helped evolve the Social Radar prototype and refine the team's vision for the desired final product much more quickly than would have been possible using a traditional software development method. This has continued to be true despite a gradual shift in emphasis from rapidly prototyping, altering and demonstrating capabilities to an emphasis on improving the robustness of capabilities whose design has matured beyond the prototyping stage.

One of the main lessons learned in the development process is that there should be an “agile analytic process” that goes hand in hand with the agile development process. The agile analytic process is an extension of the agile software development process to include analysts. Having analysts working with the developers in a tight feedback loop allows the developers to rapidly develop prototypes and update them based on feedback from the analysts who will eventually be using the tools. This tight coupling helps the developers understand the problems faced by the analysts and their work processes to tailor prototypes to the analysts’ needs. It also helps the analysts understand the strengths and limits of current technologies, leading to better feedback and new ideas for tools and services. For example, during the agile analytic process, an analyst may discover a useful or unexpected property of the data being studied. When analysts and developers discuss this property, they may be able to identify new tools to help analyze the property or improve existing tools by enabling them to exploit the pattern in the data. These new ideas reflect the technologist’s understanding of the potential of modern technology and the analyst’s understanding of the problem domain, ensuring that the new tool is both technologically feasible and useful to the analyst. The combination of the agile software development process and the agile analytic process produces excellent results in efficiently building more focused and effective analytic widgets.




Figure 5: Mobile iPad App Mock-up, Instability Toolkit, Showing the Analyst Can Set Alerts and Perform Near-Real Time Monitoring

5.5 Sustainment Strategies Lessons Learned

The landscape of online and social media sites is constantly shifting. New sites appear, others disappear, social media service capabilities change, and the popularity of sites can change quickly. To keep up-to-date with the ever changing landscape of online and social media sources and discussions, Social Radar must be kept current with the latest data sources and applications that allow for monitoring and analysis. Social Radar's modular, service oriented architecture helps it serve as a testbed for developers and technologists who are working to understand the social landscape and the information that can be gleaned from it for decision making. This provides valuable experience and design maturity before transitioning tools to use in the field.

6.0 EVALUATION AND VALIDATION

6.1 Evaluation

Social Radar is working with analysts from MITRE in our agile process, but we will also work with a new set of analysts from sponsor organizations to perform technology evaluation. Analytic technologies are often treated as discrete objects with static properties; however, a sociotechnical approach positions technology and users as coexisting in a more lateral system—human actions are to some extent guided by technological forces but humans also have the ability to resist technology, subvert its intended capabilities, or innovate in unforeseen ways. Social, cognitive, political, and cultural factors contribute to the innovative potential of tools in context. To this extent, our approach to evaluation is formative, seeking a broad base of information to understand how Social Radar widgets might be improved as technical and institutional capabilities evolve.

Social Radar technologies can be evaluated across several layers (see Figure 1 as a point of reference)—this discussion focuses on tools but can be leveraged for simulation models. First, each analytic widget purports to provide a specific capability, and the extent to which a widget performs the capability as claimed is one important task for evaluation. The sentiment and emotion analysis, for example, seeks to identify trends in news and Twitter data, respectively. For these widgets, one task for evaluation is to determine whether the analyst's expressed topic of interest is actually what the algorithms address in their analytic output. Ascertaining each widget's validity in meeting its analytic task is one aspect of evaluation. Another potential task for evaluation is: if analytic widgets are accurate in their algorithmic rationale, do they in fact address "ground truth" in the world? Or, if this cannot be satisfactorily ascertained, how might we think about these issues given the datasets and the analytic approaches employed? How might accuracy be assessed in retrospect, and can we make this process clear for users' to monitor the tools' performance?

Widget-specific usability and sensemaking, as well as cross-widget functionality, provide another layer for evaluation. Widgets are a new type of tool for analysts, and design and usability issues are still in the early stages. Widgets imply a degree of composability and flexibility, yet there exists no standard governance mechanism for organizing and representing them to users. Multiple widgets exist in a widget palette, where any number of widgets may do similar yet different things—how do analysts distinguish among widgets? How do they know what any one widget is supposed to accomplish? Which widgets can be used together to create effective workflows for any given task? An end goal for the design and implementation of Social Radar widgets is to enable users to create effective widget-task mappings without excessive in-depth training.

Widgets can be highly customized, yet they are also intended to facilitate collaborative analyses. Another layer to evaluation considers individual (cognitive) and group (social) interaction with Social Radar tools for transformative and synergistic value. How do individuals use Social Radar widgets to accomplish work? Is it better than existing capabilities? How do analysts collaborate with Social Radar widgets? What coordination or communication issues should be considered for distributed and asynchronous collaboration? How do Social Radar widgets work with other tools and processes that are commonly used to understand social phenomenon? In working with real analysts, we will seek to understand what key stakeholders have to say about the object of evaluation. The introduction of a new analytic technology has the ability to profoundly change (for better or worse) analysts' work as well as the organization of work throughout the enterprise. Changes that are likely to occur, analysts attitudes towards the technology, and other contextual atmospherics are important to assess in determining a potential technology's feasibility and impact.

A final key evaluation layer addresses whether, and to whom, Social Radar tools are relevant. While they bring to bear state-of-the art technology for understanding social phenomenon, do these tools provide granular or novel insight so that analysts will want to use them? Is Social Radar technological "state-of-the-art" good enough for, or relevant to, specific intelligence analysis tasks? If so, which ones? Will analysts trust these tools? Do the tools convey enough about their workings for analysts to feel confident in them? Analysts are notoriously stubborn about adopting automated approaches to make complex decisions. As compared to technologies that automate repetitive or logical operations, technologies seeking to accomplish increasingly complex tasks require broadly scoped evaluation approaches that result from harder-to-quantify, subjective activities such as "adding value" and "making better decisions" [22]. The layered approach to evaluation described seeks to provide rich data about the applications of this complex technology.

6.2 Validation

Social Radar technologies can be validated (see Figure 1 as a point of reference)—this discussion focuses on models but can be leveraged for tools as well. Harmon and Youngblood [23] write that ultimately the validation goal for a simulation model is establishing the conditions under which it is useful. By their criteria, a simulation is fit enough if it has sufficient fidelity to the target of the simulation (the "simuland") to meet an analyst's purpose. Bankes [13] notes that when a model is used for predictive purposes (as in scientific use), validation of its fidelity normally requires confirmation by experiment. However given the level of uncertainty in sociocultural models, large amounts of behavioral comparisons would need to be available to establish a statistically significant relationship between a model's prediction and a simuland's behavior. When we are trying to assess the impact of novel courses of action that have never before been applied to a particular situation, the needed data usually is unavailable (one goal of Social Radar is to improve this situation). So, although the underlying theories of these detail models can be, and have been, scientifically validated in the abstract, the specific applications of the models themselves often cannot be validated against existing data.

In theoretical models like S3DM, the reasonableness of the computational expression of each of the parameters needs to be made. In addition, the conceptual relationships between the elements of the model need to be examined. Do violent dissidents come from non-violent dissidents as modeled in S3DM? To address this, an unbiased, albeit subjective, process for using subject matter experts to validate these models is needed. The process acknowledges that we cannot simply compare a model's forecast with those of subject matter experts because agreement among the experts themselves is low. Instead, the process plays to the strength of subject matter experts to identify the critical causal factors in a situation that lead from initial conditions to plausible outcomes.

The process under development begins by training the subject matter experts in the particulars of the simulation model. This is necessary so they are not only conversant with elements in the model, but also what is not in the model. This is important so they do not look for causal factors that are not in the models. The model is executed for a given set of courses of action to generate a set of course of action-outcome pairs. Interesting examples can be identified from these results. For each course of action, examples of bad, good, and extreme outcomes would be chosen. In addition, outcomes near tipping points, and outcomes for apparently robust options would also make good candidates.

For each course of action-outcome pair, the subject matter experts are asked to identify the key factors and the causal connections that lead from initial conditions to the generated outcome. This is where the initial training is critical, to keep these identifications within the actual scope of the model. The model-identified key factors and causal connections are compared with those of the subject matter experts. Agreement between the independent explanations and those of the model will confirm the validity of the model for providing sociocultural understanding. When the model identifies factors and connections that were not originally identified by the subject matter experts, the model results must be traced and explained. If the subject matter experts accept the explanation, then this also confirms the validity of the model. However, if the model-explanation is not acceptable, the model's validity is brought into question.

7.0 MISSION SPACE

Social Radar is attempting to shed light on the question: what socio-cultural tools are needed at what echelon (analysts in what organization?) and across what mission spaces? Based on the example operations identified in this Joint Operations [24] publication the following missions were identified as being candidate domains where Social Radar would be applicable to at least part of the mission. Partial definitions are quoted from the publication [24] to help understand the missions.

- **Stability Operations** occur outside of the US in coordination with other “national powers to maintain or re-establish a safe and secure environment and to provide essential government services, emergency infrastructure reconstruction, and humanitarian relief.”
- **Counterinsurgency** “encompasses comprehensive civilian and military efforts taken to defeat an insurgency and to address core grievances.”
- **A Peace Operation** “contains conflict, redresses the peace, and shapes the environment to support reconciliation and rebuilding, and facilitates the transition to legitimate governance.”
- **Foreign Internal Defense** participates with civilian agencies in action taken by another government “to free and protect its society from subversion, lawlessness, insurgency, terrorism, and other threats to its security” (e.g., nation assistance).
- **Foreign Humanitarian Assistance** occurs outside of the US in coordination with the Department of State “to relieve or reduce human suffering, disease, hunger, or privation.”
- **Combating Weapons of Mass Destruction** includes “offensive operations against WMD, defensive operations, and managing the consequences of WMD attacks.”

Social Radar would likely be useful across all five mission planning phases, perhaps to varying degrees. Those phases include: Phase I Shape, Phase I Deter, Phase II Seize Initiative, Phase III Dominate, Phase IV Stabilize, and Phase V Enable Civil Authority [24]².

²**Phase 0, Shape** is the beginning phase designed to shape perception and influence behaviour to assure friendly nations and dissuade and/or deter adversaries. **Phase I, Deter** includes deterring adversaries from undesirable actions via use of preparatory actions to protect friendly forces and convey the intent to execute subsequent phases. **Phase II, Seize Initiative** is early and decisive use of joint force capabilities. **Phase III, Dominate** is to break the enemy's will to

Table 3 indicates potential Social Radar missions and users; with an ‘X’ marking which users might employ Social Radar for given missions. Stability Operations and Counterinsurgency would most likely occur with boots on the ground so Social Radar would also be used at that level. Social Radar would likely not be used for tactical support (below Battalion), but for operational and strategic support. For Foreign Internal Defense, Foreign Humanitarian Assistance, and Combating WMD, Social Radar would most likely be used at Combatant Commands, State Department, and Secretary of Defense levels (strategic).

Table 3: Potential Use of Social Radar by Mission

Mission Phase [24]	Stability Operations	Counter-insurgency	Peace Operations	Foreign Internal Defense	Foreign Humanitarian Assistance	Combating WMD
Notional Users*						
Secretary of Defense				X		X
State Department			X	X	X	
Combatant Command			X	X	X	X
In Country Headquarters	X	X				
Corps / Division	X	X				
Brigade/ Brigade Combat Team	X	X				
Battalion	X	X				
Company						

*A **Combatant Command** is a command that is composed of forces from at least two departments and has a broad and continuing mission. A Field Army is comprised of two to five Corps. In turn, a **Corps** is comprised of two to five divisions, and 20,000 to 40,000 troops and is typically commanded by a Lieutenant General assisted by a Command Sergeant Major and an extensive Corps staff. There are currently four in the Active Army. The Corps provides the basic “backbone” for multi-national operations. A **Division**, consisting of between 10,000 and 16,000 soldiers, usually consists of three brigade-sized (3,000-5,000 troops) elements commanded by a Major General assisted by two principal Brigadier General and a Command Sergeant Major. The division performs major tactical operations and can conduct sustained battles and engagements. A **Brigade** consists of 1,500 to 3,200 soldiers. Led by a Colonel, or in some instances a Brigadier General, a brigade headquarters can be employed on independent or semi-independent operations. A **Brigade Combat Team (BCT)** is the basic deployable unit of maneuver in the US Army. A **Battalion** is comprised of between 300 and 1,000 soldiers who form **Companies** therein. A battalion is typically commanded by a Lieutenant

resist or control the operational environment. Operations can range from large-scale combat to stability operations. **Phase IV, Stabilize** is characterized by a change in focus from sustained combat to stability operations. The intent of the phase is to help restore local political, economic, and infrastructure stability, and provide essential government, humanitarian, and infrastructure support. **Phase V, Enable Civil Authority** is characterized by joint force support to legitimate civil governance brought about through agreement with the appropriate civil authority [24].



Colonel and is used in independent operations of limited duration and scope. A **Company** consists of 100-200 soldiers who form three to five Platoons and is led by a Captain [25].

8.0 CONCLUSIONS AND NEXT STEPS

Social Radar has been in development for ten months. We have demonstrated a proof of concept using analytic tasks (2011 UK Riots, elections protests) and designed and are building a Service Oriented Architecture for integrating multiple internal and external capabilities. The following are the key lessons learned ordered in priority:

- Prototype architecture must separate the presentation, service, and data layers
- A layer of systems engineering for integrating tools from multiple projects is needed
- The agile software development process used was essential, and an agile analytic process is needed to sustain optimal Social Radar development starting with a domain-independent interface
- Data sources must be searchable by the analyst through their own interface, and research needs to be performed to process data at scale
- Social Radar tools need common, normalized output metrics
- The prototype must visually link the situation space with the decision space

Currently we are building the presentation, service, and data layers using the agile development and analytic processes. The new tools will be added to the system incrementally while getting feedback from analysts at every step in the process. The end vision will have a domain-independent interface focused on business process workflows and making it easier for the analyst to understand what is possible with the Social Radar tools. Research is being done on how to make Social Radar scalable to many domains (e.g., instability, violent extremism, WMD, healthcare, cyber, narcotics) at once. This requires a sophisticated domain-independent interface to search data, create workflows using analytic capabilities, and evaluate courses of action. Prototyping tools with real data, real analyst workflows, and integration with a Social Radar dashboard provides greater insight into the problems, bugs, and issues that will arise in the field, reducing the total monetary and schedule cost of developing new capabilities to address the ever-changing online data landscape. Social Radar must simultaneously be an environment to support research and a testbed for early transition of capabilities. Putting an instance of Social Radar on MITRE's external network will allow sponsors and collaborators to interact with the capabilities. Further discussion of the Social Radar vision can be found in related papers. [26, 27, 28].

9.0 ACKNOWLEDGEMENTS

Authors gratefully acknowledge help from the following individuals in preparing this paper: John James, Paula Mahoney, Allison Ounanian, Scot Lunsford, Tom Kiley, Rohit Kalkur, Holly Babin, Don Poor, Jared Mowery, and Lynette Hirschman.

10.0 REFERENCES

- [1] WN Network. "How the London Riots 'Allegedly' Started." 8 August 2011. URL: http://wn.com/How_London's_riots_'allegedly'_started_08Aug11 (last access 6 February 2012).
- [2] Goldstone, J. A. et al. (2010). "A Global Model for Forecasting Political Instability." American Journal of Political Science, Vol. 54:1 pp. 190-208. URL: <http://www.systemicpeace.org/Global%20Model%20for%20Forecasting%20AJPS%202010.pdf> (last access 6 February 2012).



[3] Day, D. and J. Boiney. "Indications and Warnings Using Natural Language Processing to Analyze Sentiment to Monitor Attitude and Behavior Trends in an Area of Operation." Sentiment Analysis for Strategic Communication Assessment (SASCA). URL: <http://www.mitre.org/tech/smartpower/research.html> (last access 6 February 2012).

[4] Elson, S. B. (2011) "Social-Global Mood Tracking Indicator." MITRE Research. URL: <http://www.mitre.org/work/areas/research/#Measuring> (last access 28 February 2012).

[5] Rasmussen, C. (2005) "Scalable Semantic Big Data Analytics in a Diverse Data Environment." URL: MITRE Research. <http://www.mitre.org/work/areas/research/#Enhancing> (last access 28 February 2012).

[6] Doran, C. (2010). "Forum and Blog Threaded Comment Analysis (FABTAC)." (Analysis of forum comment threads (e.g., YouTube comments, blogs) to provide interactive visualizations and thread summaries for intelligence and operations). URL: <http://www.mitre.org/tech/smartpower/research.html> (last access 6 February 2012).

[7] Metropolitan Police. (2011). "Strategic Review of MPS Response to Disorder: Early Learning and Initial Findings." New Scotland Yard. 21 October 2011. <http://tiny.mitre.org/4EB1> (last access 6 February 2012).

"2011 England Riots." Wikipedia. URL: http://en.wikipedia.org/wiki/2011_England_riots (last access 6 February 2012).

"Blair is Right: Very Few of the Unemployed Rioted." Daily Telegraph, 23 August 2011. URL: <http://www.telegraph.co.uk/comment/letters/8716577/Blair-is-right-very-few-of-the-unemployed-rioted.html> (last access 6 February 2012).

Johnson, M. "Society Inside-out: Tomorrow's Riots Can Be Prevented." The Guardian, 19 October 2011, p17. (through Lexis-Nexis – last access 6 February 2012).

"London Protests Spread to Manchester." PressTV. 10 August 2011. URL: <http://www.presstv.ir/detail/193228.html> (last access 6 February 2012).

"UK Riots: the Key Facts and Figures." The Guardian, 9 August 2011. URL: http://en.wikipedia.org/wiki/2011_England_riots (last access 6 February 2012).

"UK Riots: Nearly 2000 Arrested So Far, Say Police." The Guardian. 21 August 2011. URL: <http://www.guardian.co.uk/uk/2011/aug/21/uk-riots-nearly-2000-arrested> (last access 6 February 2012).

Urban Riots: Seven Days that Shook Britain." The Guardian, 12 August 2011. URL: <http://www.guardian.co.uk/commentisfree/2011/aug/12/urban-riots-seven-days-britain> (last access 6 February 2012).

"We Will Use Water Cannons on Them: At Last Cameron Orders Police to Come Down Hard on the Looters." London Daily Mail, 11 August 2011. URL: <http://www.dailymail.co.uk/news/article-2024203/UK-RIOTS-2011-David-Cameron-orders-police-come-hard-looters.html> (last access 6 February 2012).

[8] Choucri, N., et al. (2006). "Understanding and Modeling State Stability: Exploiting System Dynamics." MIT Sloan Working Paper 4574-06. CISL Working Paper 2006-02. URL: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=876962 (last access 6 February 2012).

[9] Mathieu, J. (2011) "Integration of Socio-Cultural Indicators for Global Situational and Decision Awareness (Social Radar)." MITRE Research. <http://www.mitre.org/work/areas/research/#Measuring>

- [10] Megerdooian, K. (2011). “Exploring Soft Power in Weblogistan.” (Foreign language blog analysis) URL: <http://www.mitre.org/tech/smartpower/research.html> (last access 6 February 2012).
- [11] Salerno, J. J., et al. (2011) “The National Operational Environment Model (NOEM)- A Focus on Understanding the Populace,” in Proceedings of the Modeling & Simulation Conference, MODSIM World 2011 Conference and Expo, Virginia Beach, VA, 11- 14 October 2011.
- [12] Zarrella, G. (2012). “Mapping Influence Via Online Postings.” MITRE Research. URL: <http://www.mitre.org/work/areas/research/#Measuring> (last access 28 February 2012).
- [13] Bankes, S. (1993). “Exploratory Modeling for Policy Analysis.” Operations Research, Vol. 41:3 (May – June), pp. 435-449.
- [14] Hall, D. L., et al. (2007). “Rethinking the Data Overload Problem: Closing the Gap Between Situation Assessment and Decision Making in Proceedings of the 2007 National Symposium on Sensor and Data Fusion (NSSDF) Military Sensing Symposia (MSS), McLean, VA.
- [15] Endsley, M. R. (1995). “Toward a Theory of Situation Awareness in Dynamic Systems.” Human Factors Vol. 37:1, pp. 32-64.
- [16] Drury, J. L., et al. (2009a). “Dynamic Decision Support for Emergency Responders” in Proceedings of the 2009 IEEE Conference on Technologies for Homeland Security (HST '09), pp. 537-544.
- [17] Klein, G. L., et al. (2010). “COAction: Enabling Collaborative Option Awareness” in Proceedings of the 15th International Command and Control Research and Technology Symposium (ICCRTS), Santa Monica, CA.
- [18] Pfaff, M. S., et al. (2010). Weighing Decisions: Aiding Emergency Response Decision Making via Option Awareness in Proceedings of the 2010 IEEE International Conference on Technologies for Homeland Security (HST), pp. 251-257.
- [19] Klein, G.L., et al. (2009). “Modeling as an Aid to Robust Tactical Decision-Making.” 14th International Command and Control Research and Technology Symposium, June 2009, Washington, DC.
- [20] Mathieu, J. et al. (2010). “Tactical Robust Decision-Making Methodology: Effect of Disease Spread Model Fidelity on Option Awareness” in Proceedings of the 7th International ISCRAM Conference – Seattle, USA, May 2010.
- [21] Troop, J. (2011) “Better Leveraging MITRE's Big Data Analytics Capability.” MITRE Research. URL: <http://www.mitre.org/work/areas/research/#Enhancing> (last access 28 February 2012).
- [22] Farbey, B., F. Land, et al. (1993). “How to Assess Your IT Investment: A Study of Methods and Practices.” Oxford: Butterworth Heinemann.
- [23] Harmon, S.Y., & Youngblood, S. M. (2005). A Proposed Model for Simulation Validation Process Maturity, *JDMS*, Volume 2, Issue 4, pp. 179-190.



- [24] Joint Publication 5-0. "Joint Operation Planning" 11 August 2012. URL: http://www.dtic.mil/doctrine/new_pubs/jp5_0.pdf (last access 2 February 2012)
"Combatant Command." DoD Dictionary of Military Terms URL: http://www.dtic.mil/doctrine/dod_dictionary/data/c/3206.html (last access 2 February 2012).
- [25] "Evolution of the US Army Divisional Structures" URL: <http://www.fas.org/man/dod-101/army/unit/docs/divhist.htm> (last access 2 February 2012)
"Military Structure: Army." GlobalSecurity.org URL: <http://www.globalsecurity.org/military/agency/army/> (last access 2 February 2012)
"United States Army Chain of Command." About.com: US Military. <http://usmilitary.about.com/od/army/l/blchaincommand.htm> (last access 2 February 2012).
"U.S. Army Organization." Olive-Drab. URL: http://olive-drab.com/od_milorgs_us_army.php (last access 2 February 2012).
- [26] Maybury, M. (2010) "Social Radar for Smart Power." The MITRE Corporation. http://www.mitre.org/work/tech_papers/2010/10_0745/10_0745.pdf (last access 28 February 2012).
- [27] Costa, B., et al. (2012) "Social Media Analysis with Social Radar." NATO HFM-201 Meeting on Social Media: Risks and Opportunities In Military Applications. Tallinn, Estonia, 16-18 April.
- [28] Servi, L.D. and S. B. Elson. (2012). "A Mathematical Approach to Identifying and Forecasting Shifts in the Mood of Social Media Users" NATO HFM-201 Meeting on Social Media: Risks and Opportunities in Military Applications. Tallinn, Estonia, 16-18 April.