Tools for Cultural Intelligence Development: A Cognitive Engineering Approach

Dr. Michael Woodman, Dr. William Fitts, Thomas Santarelli, Andrew Rosoff, & Jennifer Engimann
CHI Systems, Inc.
Plymouth Meeting, PA
{mwoodman, bfitts, tsantarelli, arosoff, jengimann} @chisystems.com

ABSTRACT

Military Operations Other Than War, from humanitarian assistance to counterinsurgency support, require actionable cultural intelligence – information that provides understanding of the behavior of a general population and its friendly, neutral, and hostile sub-groups, plus ways to recognize and predict changes in that behavior over time. Cultural intelligence is made actionable through integration with tactical data via the Intelligence Preparation of the Battlespace (IPB) process. Cultural intelligence development and application are poorly supported by existing intelligence systems.

This paper explains how a cognitive engineering methodology was used to understand this problem from the analyst's perspective and to develop human-centered technology solutions. Domain and semantic analyses were used to identify the concepts and relationships central to creating and applying cultural intelligence. Cognitive work analysis was then applied to the processes and workflows used in cultural intelligence and its integration into IPB at the infantry battalion (and below), where the problem is most acute. User-centered design methods were employed to create and refine user interface, automation, and novel workflow solutions.

The paper also describes how the analysis demonstrated requirements for:

- flexible tools and interfaces allowing ad hoc creation and communication of intermediate products for rapid, collaborative evolution
- multiple combinable visualizations that support temporal, geo-spatial, and social network-based analysis

Lastly, the paper delineates how the requirements translated into the goal of creating shareable, user-constructed products termed "mashups," which enable users to tailor and annotate visualizations for customized, mission-oriented intelligence products. Request for Information (RFI) workflows allow saved mashups and underlying data sets to be restored for live processing, ongoing analyses, and exploration of multiple hypotheses. Automated analyses of data-intensive social media sources were created to meet needs for instantaneous assessment of "area atmospherics." These features are integrated into the CultureMap system, which was built upon CHI Systems' 4D-Viz framework for creating Command and Control (C2) applications, and has proceeded through establishment of face and construct validity with domain experts.

ABOUT THE AUTHORS

Dr. Michael Woodman is a Senior Scientist for CHI Systems. Dr. Woodman has over thirty years of experience in training and simulation and has been the Principal Investigator for the CultureMap effort and was the Modeling and Simulation Officer for Marine Air-Ground Task Force (MAGTF) Training Simulations Division at the USMC Training and Education Command.

Dr. William Fitts is a Software Engineer and Analyst for CHI Systems. He has been combining his background in the four sub-fields of anthropology (archaeology, physical anthropology, socio-cultural anthropology, and linguistics) with computer applications for the past eighteen years.

Thomas Santarelli is Director of Training at CHI Systems, Inc. Mr. Santarelli has more than twenty years' experience in research into cognitive modeling and has applied cognitive engineering techniques to intelligence analysis domains by developing automated decision-support capabilities to provide key accelerators of the decision-to-action cycle at the center of intelligence analysis.

Andrew Rosoff is CEO and Principal Software Engineer at CHI Systems, Inc. Mr. Rosoff has a background in the design and implementation of AI-based systems and engineering applications of cognitive science, with nearly 15 years of relevant experience in these areas.

Jennifer Engimann is a Software Developer at CHI Systems, Inc. Ms. Engimann has a background in the design and implementation of GPS/GIS systems and Big Data applications. She has applied her expertise in artificial intelligence and computational linguistics, including sentiment analysis and natural language processing,

Tools for Cultural Intelligence Development: A Cognitive Engineering Approach

Dr. Michael Woodman, Dr. William Fitts, Thomas Santarelli, Andrew Rosoff, & Jennifer Engimann
CHI Systems, Inc.
Plymouth Meeting, PA
{mwoodman, bfitts, tsantarelli, arosoff, jengimann} @chisystems.com

BACKGROUND AND PROBLEM DEFINITION

Soldiers and Marines in a tactical environment can benefit from the dissemination of cultural intelligence within the context of existing battlefield information systems. Cultural intelligence provides valuable context that can support the understanding of the behavior of a general population and enemy combatant groups, recognizing changes in behavior over time, and the ability to predict future behavior. To address this operational need, CHI Systems conducted Research and Development (R&D) that leveraged our work on advanced net-centric information systems. Our cognitive engineering approach to this R&D involved four layers of functionality designed to assist intelligence officers and the individual warfighter:

- A cultural information database populated with regionally-specific cultural information that can be browsed or searched semantically or geospatially;
- Graphically-defined, user-modifiable cultural intelligence templates that can be used to designate key cultural infrastructure;
- A cultural behavior and history database that contains multimedia data that document the operational history of the region; and
- Cultural analysis and alerting tools that allow an intelligence officer to identify trends, predict the cultural influence on tactical operations, and share this information with operational units he/she services.

RESEARCH AND FINDINGS

A cognitive engineering methodology, based on commonly applied methodologies (Norman, 1981; Vicente, 1999), was used to understand this problem from the intelligence analyst's perspective and to develop human-centered technology solutions. Domain and semantic analyses were used to identify the concepts and relationships central to creating and applying cultural intelligence. Cognitive work analysis was then applied to decompose and analyze the processes and workflows used in cultural intelligence and its integration into IPB at the infantry battalion (and below) where the problem is most acute.

Our first step in this project was to define the functionality required in the development of tools for Cultural Intelligence. We collected operational needs and then conducted and documented a requirements analysis. This allowed us to define the functionality of the solution based on the requirements analysis and to develop a Concept of Operation. Throughout the project, we revised the functional definition and Concept of Operation. Our research then focused on four areas related to dissemination of Cultural Intelligence:

- Cultural ontology (how to organize cultural information)
- Atmospherics (a framework for key cultural trend information)
- Visual analytics (integrated visualizations for exploratory and detailed analysis)
- Communication and collaborative sensemaking (how to share/collaborate on analysis results)

Cultural Ontology

In order to understand Cultural Intelligence from the analyst's perspective, we first had to develop a taxonomic representation of cultural HUMINT information and use ontology-building tools to convert this into a formal ontological structure. We defined six example geographical areas from which example/demonstration cultural backdrop datasets could be constructed. We analyzed and reviewed existing cultural data resources and developed a plan to populate the ontology with appropriate cultural backdrop data.

Because the concept of culture is extremely difficult to define (Geertz, 2003), the exact definitions of the groups and elements that make up culture in general and a culture specifically are extremely diverse. We have drawn heavily from the MCIA's 2008 Cultural Generic Information Requirements Handbook (C-GIRH) (MCIA, 2008) as a common reference to provide unified cultural terms and concepts. Drawing from this document's practical hands-on approach, we have focused our ontology on the observable expressions of culture within the area of operation. While not as nuanced as a modern socio-cultural analysis by an anthropologist, this approach provides relevant and useful data that is easily quantifiable and straightforward to collect. To relieve some of the confusion around similar terminologies, we use the term "population" to refer to a group that is the subject of cultural analysis.

<u>Population versus Network</u>
Borrowed from demography, the term "population" refers to groups of people living and interacting in a shared area. Even though many populations are in fact made up of individuals from multiple cultures and ethnicities, they often have shared or common cultural frameworks that govern daily life within and between their communities (such as food choices, language, behavioral patterns, and belief systems). By removing common cultural terms in favor of the more abstract "population," we help control the data collection process by removing loaded terms that can lead to unintentional interpretations during collection.

"Networks" are often also referred to as groups in many analysis systems. In our usage, "networks" consist of individuals, who may or may not share a cultural background, but are held together by motives for specific action. Relationships do exist between populations (or subpopulations) and these networks of highly-motivated and action-oriented individuals are related through individual affiliation (e.g., ethnic group membership) or through individual relationships (e.g., two persons attended same school).

Terrorist networks are often organized orthogonally to cultural divisions (i.e., networks may cross ethnic boundaries), while other networks of individuals (e.g., the power elites, kinship elites, economic elites, religious elites) may form the active core of specific population segments (e.g., Shia clerics). Population/network relationships are critical to understanding conflicts within single societies or nations and therefore are critical to understanding stability as well as pathways to influence stability.

Atmospherics

We have determined that there exist two different but equally important perspectives to understanding cultural intelligence: situational and sentiment. Situational data describes what is happening in absolute and quantifiable terms, for example:

> *Are people receiving enough food aid to prevent malnutrition?*
> *Yes, X tons of wheat and maize products are being distributed in affected areas, thus ensuring that each person receives 1500 calories of food per day.*

Sentiment data describes how the population within the area of interest perceives the situation through their own cultural perspective. For example:

> *Is there an adequate food supply?*
> *No, people say they are not receiving enough rice.*

For the analyst to apply atmospherics to the IPB process, he must understand the nature and sources of each indicator type.

<u>Situational Atmospherics (Facts on the ground)</u>. These indicators can be gleaned from national and international referential sources such as the World Bank Data, United Nations Agencies, Central Intelligence Agency Factbook, military personnel in the area, and even (with caution) news reports. Data on Situational Atmospherics are time consuming to develop because while relatively course data at yearly increments are straightforward to obtain, detailed information needs to be carefully assembled. These indicators include items such as:

- annual rice production/imports nationwide
- average price of rice on a national level
- current price of rice in a specific village market

- amount of rice available (known imported quantities)

Sentiment Atmospherics (Perceptions). These indicators are not as clear-cut as facts on the ground and must be accumulated, analyzed, and interpreted from sources such as social media (Twitter, Facebook, web logs or “blogs”), online news comments, and even military personnel in the area. Sentiment Atmospherics data can be acquired and processed quickly, course evaluations can happen in near real time, and more detailed analysis can be provided only a few hours after data collection begins. These indicators include items such as:

- the population's belief that the price of rice is too high (by comparison to neighboring countries it might actually be low)
- the population's belief that there is a shortage of rice (whether true or not, it might cause hoarding)
- the population's belief that the government is to blame for a real or perceived shortage (may tie into other, unrelated anti-government sentiments)
- the population’s believe that rice availability is being controlled by corrupt officials

These atmospherics are then combined into aggregate situational and sentiment indicators describing the most recent trends in the data and if they represent a positive impact on the operational environment.

Data Analysis: Mapping Atmospherics to Social Media Data

In order to move forward in the development of our sentiment analysis approach, we decided to focus exclusively on the analysis of Twitter data. To do this we needed to establish the most efficient and effective methods for categorizing “tweets” (individual messages or status updates posted on Twitter) across different cultural contexts. In order to determine the degree to which classification can be based on training from a single cultural context, it was necessary to collect substantial numbers of tweets in order to compare the various solutions (e.g., training on country A versus training on countries A and B, and then applying to country C). Expanding the process to a larger set of countries was also required in order to account for the additional classes of atmospheric-related sentiments remaining to be covered. Collection, tagging, and analysis are therefore ongoing for our six target countries.

While our set of analyzed and tagged data is still expanding, we have accumulated a large sample (over 105,000,000 tweets as of June 2014) which allows us to incorporate the new data, along with additional classes of atmospheric-related sentiments, into the automatic classification process. We successfully constructed automatic classifiers capable of discriminating between sentiments expressed toward the six atmospheric classes for which sufficiently large samples were collected. These include sentiments toward (1) Opposing Forces, (2) Blue Forces, (3) Government, (4) Stability, (5) Economy, and (6) Food Availability. While present in the data, the sets of exemplars found for the classes of sentiment toward Media Saturation, Medical Availability, and Public Works Status are currently still too small to be significant, and therefore were temporarily excluded. While more data for these classes may be found during ongoing analysis, it is likely that calculation of these atmospherics will rely more heavily on situational data sources.

Automatic classification across six classes is an inherently difficult problem. Inspecting individual elements of content reveals that instances of misclassification are often caused by a natural, true overlap existing between one class (e.g., such as Stability) and a small number of related classes. For example, the following excerpted tweet contains an indication of negative sentiment toward both the Government and Stability:

a) *“RT @#####: The poor in the Philippines are not only income poor but also suffer from deprivations in health and education.*

Given the ultimate goal of the current effort, the exact classification of any individual piece of content (e.g., tweet) is much less significant than the larger, summative results aggregated over different granularities of time and space. Therefore, accuracy in the overall, aggregated results can be best improved by allowing each content element to partially contribute simultaneously to multiple atmospheric categories in order to characterize mixed content such as that seen in the example provided above. The most accurate classification methods we have used, however, do not lend themselves easily to this as they involve employing a series of multiple, pairwise classifiers, each trained to discriminate between a single pair of atmospheric categories. We therefore implemented an additional post-classification algorithm which, for each content element, transforms the set of confidence values produced by the many pairwise classifiers into a single continuous probability distribution across all six atmospherics (Hastie & Tibsirani 1998, Wu, Lin, & Weng 2004).

With this post-classification algorithm implemented, we then began a process of manually inspecting the final results, focusing on individual cases that either clearly fell into one distinct category or that clearly straddled multiple categories simultaneously.

The large bulk of ingested social media content, however, will often be irrelevant to the current effort, and therefore should not be allowed to contribute to any metric illustrated in the dashboard. This is true even when location and keyword constraints are applied, such as in example (c) provided below:

b) "RT @#####: @##### Number 3 in the Philippines :) http://t.co/OkbRLZ3ZSj"

Tweets such as above are not relevant to atmospherics, but nevertheless will still be assigned some probability distribution across the atmospheric classes.

With the initial classification and post-classification algorithms implemented, we undertook the task of validating our methodology with real-world data. By this point, we had already collected nine continuous months of Twitter data for our target countries, including the Philippines. The timing of our validation phase coincided with the aftermath of Typhoon Haiyan; using the Twitter data preceding, during, and following the typhoon landfall represented a timely and useful case study. Samples from the data set were therefore processed through the entire analysis pipeline and the results inspected.

A closer look (Figure 1 below) at the classification of the tweets collected for the Philippines over a period of two months shows that the perturbations in the data coincide with events in the Philippines: the climax of a weeks-long standoff with Muslim rebels in the southern city of Zamboanga, the recent earthquake in the province of Bohol, and the landfall of Typhoon Haiyan. As seen in the chart, events surrounding Haiyan generated an enormous amount of Twitter traffic when compared to preceding months and even when specifically compared to other events of note. It should be noted that this increase in traffic corresponded with an increase in tweets about the Philippines from outside of the region (e.g., raising awareness and funds for relief supplies such as clean drinking water).

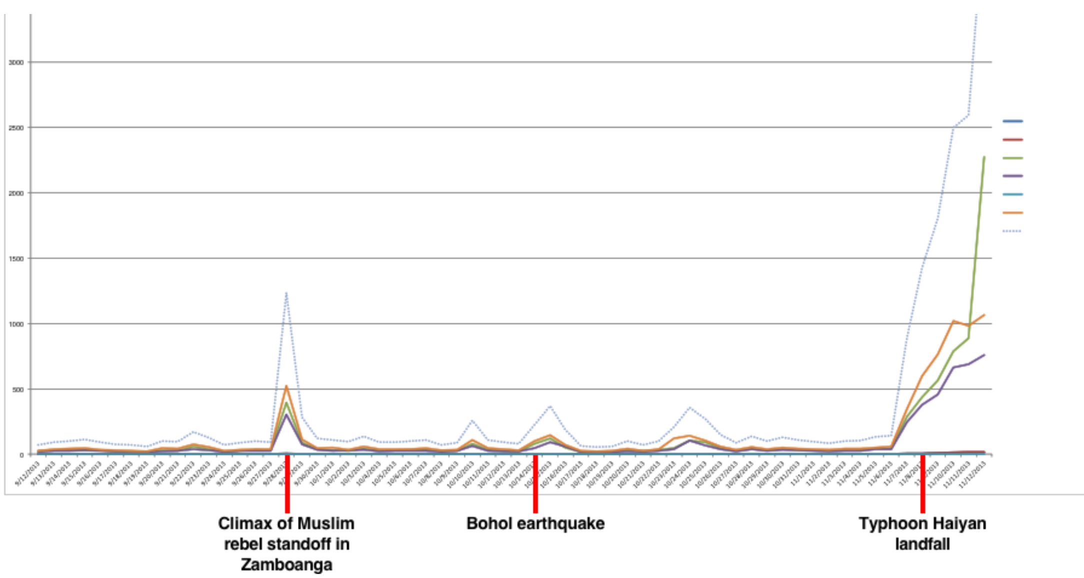

Figure 1. Atmospherics for landfall of Typhoon Haiyan Analysis

Work is on-going to fine tune our training and processing methodologies so that we will be able to remove the need for validating classifications after the initial training of the software.

Visual Analytics

Visual Analytics is "the science of analytical reasoning facilitated by interactive visual interfaces," (Thomas and Cook, 2005) or, in other words, how to put the atmospherics data together so that an analyst can use it. Analysts typically:

- Must merge access to cultural intelligence, HSCB modeling results, and tactical events into single suite of visualizations providing <u>geographical and temporal</u> context
- Must be possible to view a single data set through the lenses of various visualizations and readily move between them in order to allow analysts to integrate these multiple perspectives
- Rapid navigation along <u>threads of interest</u> through the data, moving between different visualizations and related data subsets of interest
- Drill-down to understand and evaluate summary visualizations
- Need multiple views of data sets, with cross-view interaction (e.g., highlighting of same elements in multiple views), and the ability to move from aggregated data to individual data elements (e.g., tweets).

In the development of a visually-oriented analysis system requiring user activities that involve moving across many displays, advanced human-computer interaction techniques are needed in order to support the user and to reduce cognitive load and complexity. To meet this complex user-interface challenge, we leveraged a concept developed by Woods (1984) termed visual momentum. Woods describes visual momentum as the "…observer's ability to extract task-relevant information." (Woods & Watts, 1997). More specifically, "…the amount of visual momentum supported by a display system is inversely proportional to the mental effort required to place a new display into the context of the total data base and the user's information needs. When visual momentum is high, there is an impetus or continuity across successive views that supports the rapid comprehension of data following the transition to a new display" (Woods, 1984). We have leveraged a number of techniques developed by Woods and colleagues to maximize visual momentum in order to reduce cognitive complexity of the user interface and workload of the user. Taken in aggregate, Figure 2 below depicts our user interface based on a hypothetical visual analysis.

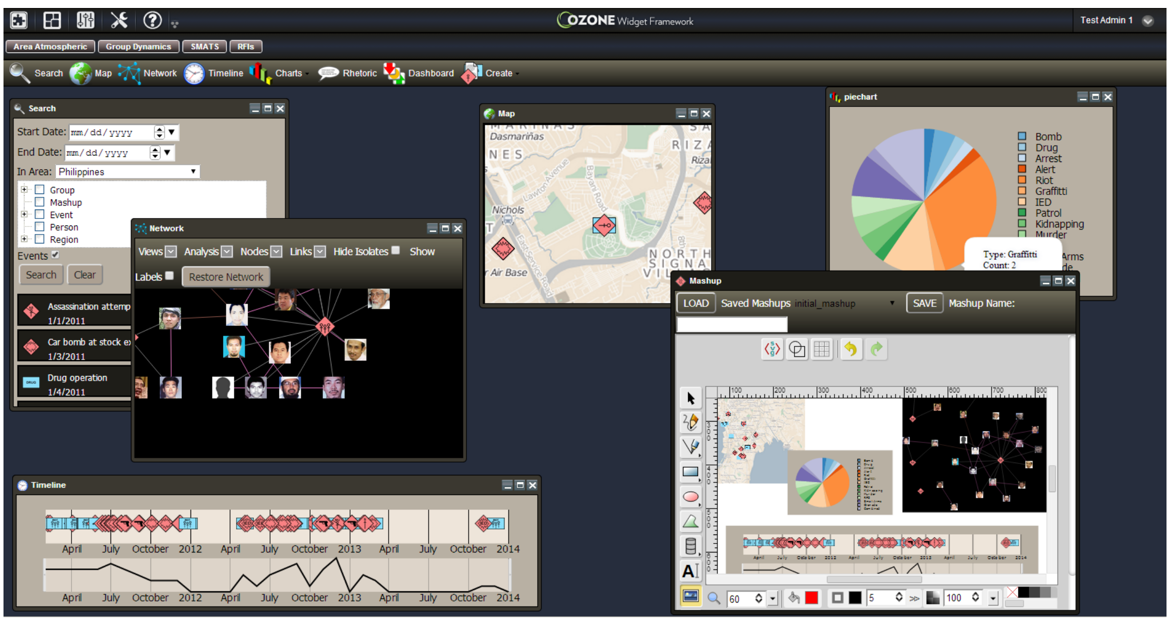

Figure 2. CultureMap Analytics – User Interface Example

The user interface requirements derived to meet the needs of Intelligence Analyst end-users, as well the design principles used to guide the design of the user interface constructs developed around these techniques, are both briefly described below.

<u>Principles Leveraged in the Design</u>

Woods and Watts (1997) describe a number of design techniques to maximize visual momentum, of which a subset was included within our design. The design therefore includes specific capabilities to support interactions with the user interface based on these design techniques, namely: (1) "longshot" displays; (2) related/parallel views; (3) center surround; (4) cues to status; and (5) conceptual spaces, all briefly explained below.

(1) <u>Longshot</u> displays allow a user to see "the big picture" or overall status of the work context (e.g., battlespace) without having to view every detail shown at one time. Where possible, information will be put in context in a meaningful manner given the informational needs of the observer, and in a form that helps guide the user's attention to that which deserves a more immediate response. From such a display the user is able to pursue more detail in a natural manner from the overview, and it is possible, where appropriate, to reorganize the information to efficiently provide alternate views that reveal further facts or relationships about the data shown. We integrated four longshot characteristics described by Woods and Watts (1997) into our design, including ensuring that data is distilled and abstracted, conveys change/sequence information, and conveys information relevant to the user's context.

(2) <u>Related views</u> enable users to view information with intrinsic relationships and to do so in parallel. While a simple concept, the idea of presenting related information in parallel is an important one in user interface design. It may be important, for instance, to show overview (i.e., longshot) information with the related detailed information to help put the detail in perspective. Other relationships can be captured by showing statistical charts in parallel or a table of significant events in parallel with data that represents effects they may have caused.

(3) <u>Center surround</u> draws upon concepts related to how human focal attention and vision supports the process of knowing where to look when. Within this set of concepts, design questions arise as to how to encapsulate, aggregate, and disaggregate information to highlight relationships within data relative to how that information is salient for the user depending on context. When designers effectively use center surround, it helps the user see overall relationships and where to focus further attention. One way to focus and obtain more information is to 'zoom', where more information is obtained about a subset of the data by changing the range of interest on a scalar metric. For instance, a user can change the date/time range or grid coordinates range within a search/filter component. While looking at the map display of events for a selected date range, the user could notice a large concentration within that narrow time-frame in close proximity to a specific landmark (e.g., mosque). The user could therefore zoom in on that part of the city and examine specific clusters of IED events near the mosque.

(4) <u>Cues to status</u> enable users to quickly 'size up' status or activities within the task space or to quickly understand what has changed, and the recency of that change, within a data-space. In particular, longshot displays benefit from well-designed cues to status as they should include a distillation of the most salient aspects of interest to the user, include abstracts of the data beyond simple drill-down, and enable users to maintain their workflow with limited interruptions.

<u>User Requirements for Integrated Visual Analytics</u>

Based on interviews with domain subject matter experts (Intelligence Analysts), we developed the requirements for visual analytic metaphors for both *view* and *control* constructs within the CultureMap design and implemented a set of user interface widgets based on those constructs and the accompanying techniques posited by Woods and Watts (1997). These constructs include:

- <u>Geospatial Views</u>: A predominant portion of information relevant to the CultureMap users is often highly geospatial in nature. Where data of any type has such information associated with it (e.g., a latitude and longitude), the user must be able to view that information in the context of a 2D display.
- <u>Network Views</u>: An interactive view of the information is presented in graph form displaying the data attributes and their relationships. Advanced versions could include graph-layout algorithms that impart meaning to the layout based on relationship data specified by the user.
- <u>Temporal Views</u>: The information of interest to the end user commonly has a temporal aspect associated with it as well (e.g., the observed time of an event or the time span of a religious holiday). Where this is the case, the user must have a temporal perspective, or a selection of temporal perspectives, available for visualizing the information in this context. This may include displays such as timelines, time-wheels, and time series charts.
- <u>Temporal Controls</u>: An equally common starting point for navigating through the available information is time. A traditional calendar, with the ability for one to select individual points as well as distinct ranges of time, is a familiar mechanism for achieving this functionality. However, as with the geographic control mechanism described above, wherever possible temporal views should 'double' as temporal controls. This

would imply, for example, the ability to highlight a particular region of time on a time series chart or time-wheel, consequently causing the data displayed in other views to be updated appropriately.

- Filtering Controls: The CultureMap design includes a filter control on top of the search function whereby the user may construct criteria for exactly which data, from the full set available, is presented. This functionality becomes essential once the initial query has been performed and the user then wants to drill-down or drill-up alternatively. Examples of this are data type (e.g. events, groups, places) and data subtype (e.g. IEDs, religious sects, mosques).
- Classic Search: While the above methods for querying and filtering the available data are visually-intuitive, often a user needs to search for particular words or phrases in the data. To this end, the CultureMap design includes a more traditional search mechanism.
- "Mashup": Individual visual analytics can be stored in private workspaces and can be annotated, saved, modified, and maintained over time, thus enabling the user to build and maintain an understanding of a region/culture of interest.
- Collaboration: All annotations and other changes to the workspace are pedigreed and tracked. In this way, a full history of which users have contributed to changes, and when, may be reproduced and shared with other CultureMap users. This is important as it promotes a more efficient RFI workflow capability, thus enabling users to draw upon knowledge of cultural or tactical information that is often spread across geospatially distributed users (e.g., reach back context).

COMMUNICATION AND COLLABORATIVE SENSEMAKING

Our cognitive engineering approach led to development based on two important architectural considerations, given the need for future iterations of the toolset to support advanced collaboration techniques, likely involving users in multiple geographic and temporal contexts. These considerations included the need to:

- Create an Ozone Widget Framework (OWF)-based/web-based implementation available as analytics services, Ozone widgets, or a stand-alone, integrated application, so that CultureMap could act as a framework that can be tailored to user roles/needs while at the same time supporting the ability to enable efficient sharing and dissemination of intelligence products derived via those widgets, and
- Design a computational Service Oriented Architecture (SOA) that is compatible with both the developmental testbed as it matures and to support ultimate integration with the Distributed Common Ground System (DCGS) so that CultureMap could be deployed to users based on their specific mission and usage needs.

An additional constraint was the need to be able to rapidly reconfigure variations of the application to accommodate differences in deployed users' needs based on mission differences.

To support our ultimate vision for a flexible application environment that can provide intelligence analysts with tools that support the analysis, visualization, reification, and dissemination of cultural and social information, we used a widget-based, services-oriented development approach as described above. The toolset supports shared battlespace understanding, as it is architected and designed around the recognition that sensemaking in a distributed environment relies upon a mix of individuals, each possessing unique knowledge and incremental insights in an area of responsibility based on cultural and social information requirements. To that end, the system was engineered to support collaboration – based on the concept of accommodating RFI workflows, whereby an RFI may start with a single user and require multiple additional consumers and producers to add to that RFI before it satisfies the intent of the originator of the initial RFI, as notionally depicted in Figure 3 below.

The core conceptual and functional capability that collaboration within the system is premised upon is that of the "mashup." Mashups are collections of individual visual analytics that individual users can create within their own workspace and can be annotated, saved, modified, and curated over time, thus enabling the user to build and maintain an understanding of a region/culture of interest. More importantly, they are the basis upon which visual analytics can be shared between system users as they include the entire context used to generate each intermediate product as well as all enhancements by each user as it is shared and distributed over time.

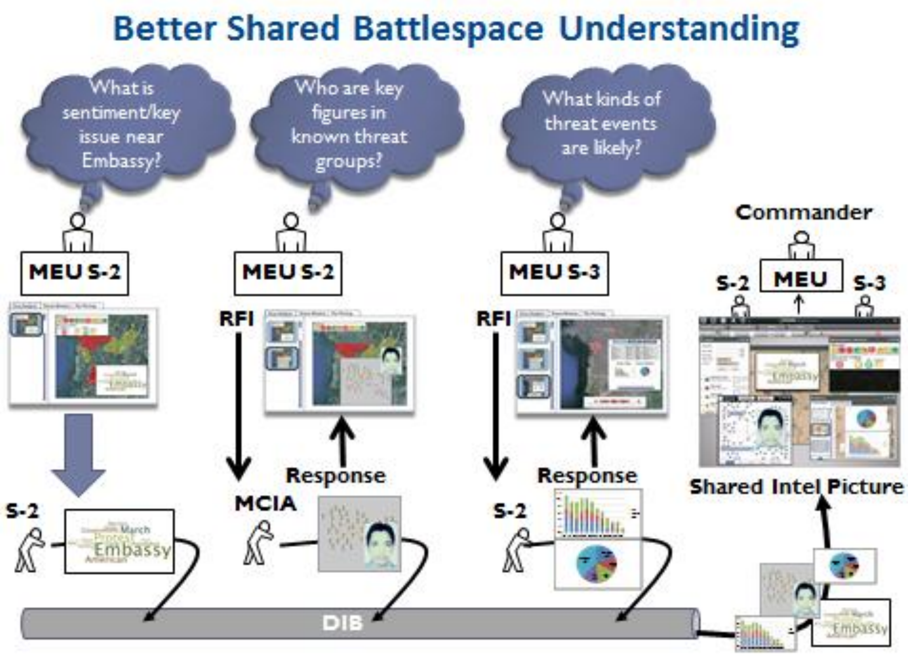

Figure 3. Example Producer-Consumer Workflows

Thus, as RFI workflow products iteratively progress across team-members, users can leverage these saved mashups to restore underlying data sets for live processing, allowing resumption of ongoing analyses, and exploration of multiple hypotheses of their own. Mashups are based on four key design considerations:

- Support the analysis capture, collaboration, and dissemination of individual user analytics,
- Composed as collection of annotatable visualization snapshots,
- Restorable for live manipulation as original visualizations & underlying data set/criteria, and
- Support emergent RFI workflows.

The ability for any user to designate exactly which pieces of information are of interest, based on combining visualizations and information of interest, and creating reified versions of information directly provided by the system, is a powerful capability and is consistent with the Woods and Watts (1997) Conceptual Spaces concept. Since mashups can be annotated, saved, modified, and maintained over time, users can build and maintain an understanding of events, entities, and information of interest, as seen in the example in Figure 4.

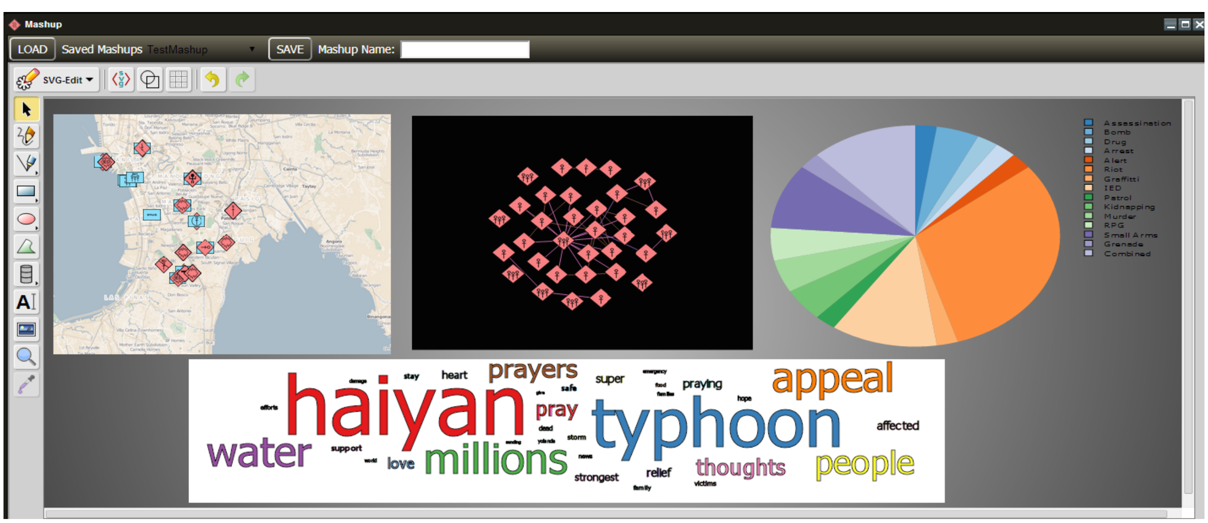

Figure 4. CultureMap Mashup Concept

CONCLUSIONS AND FUTURE RESEARCH

As the military continues to transform to more net-centric operations and incorporates systems to support collaborative sensemaking for a distributed command, there will be an increasing need for tools to help users across all echelons aggregate, interpret, and share resulting data and analytics. The framework and sample application that we have developed provide a flexible architecture that provides for the rapid development of visualization capabilities and graphical user interfaces that afford battlespace understanding using proven information visualization principles and techniques. Beyond the direct utility that the framework provides to prospective end-users, a number of lessons-learned have been garnered from a more general perspective regarding the development of sensemaking applications. These lessons learned can be applied to broader applications beyond CultureMap for systems that involve complex visual analysis, including the need to:

- design towards a flexible infrastructure to facilitate the evaluation of new visual analytics technologies and approaches based on principled human factors guidelines;
- create and use a common standards-based widget framework that includes security and privacy infrastructure;
- use a component-based software development approach for visual analytics software to facilitate evaluation of research results in integrated prototypes and deployment of promising components in diverse operational environments; and
- create a community of practice and develop guidelines for best-practices to guide the use of the visual analytic framework so that they can be effectively integrated within operational systems.

We believe that providing cultural data analysis tools, as developed within the CultureMap framework, coupled with visualizations that afford easy interpretation of the key features of complex, multifaceted data is a powerful approach. Future evolutions of CultureMap will focus on the expansion of the underlying RFI workflows, collaboration and sensemaking tools, and integration with other sentiment analysis tools and frameworks.

ACKNOWLEDGEMENTS

The research and development discussed in this paper has been sponsored by the Department of the Navy, Office of Naval Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Office of Naval Research.

REFERENCES

Geertz, C. (2003). Thick Description: Toward an Interpretive Theory of Culture, in Turning Points in Qualitative Research. Yvonna S. Lincoln and Norman K. Denzin eds., Altamira, New York.

Greenberg, T.L. (1993). The Handbook for Focus Group Research. Lexington, MA: Lexington Books.

Hastie, T. & Tibshirani, R. (1998). Classification by Pairwise Coupling. The Annals of Statistics 26 (1998), no. 2, 451--471. doi:10.1214/aos/1028144844. http://projecteuclid.org/euclid.aos/1028144844.

Marine Corps Intelligence Activity (2008). C-GIRH Cultural generic information requirements handbook. Quantico, VA

Murdock, G.P. (1938). *Outline of Cultural Materials*. New Haven, CT: Institute of Human Relations.

Norman, D.A. (1981). Steps Toward A Cognitive Engineering. Technical Report, Program in Cognitive Science, University of California at San Diego.

Thomas, J.J. & Cook, K.A., eds. (2005) Illuminating the Path—the Research and Development Agenda for Visual Analytics, IEEE CS, 2005.

Vicente, K. (1999). *Cognitive Work Analysis*. Lawrence Erlbaum Associates, Mahwah, NJ.

Wong, P.C. & Thomas, J. (2004). Visual Analytics, *IEEE Computer Graphics and Applications*, Volume 24, Issue 5, Sept.-Oct. 2004 pgs. 20–21.

Woods, D. D. (1984). Visual momentum: A concept to improve the cognitive coupling of per- son and computer. *International Journal of Man–Machine Studies, 21*, 229–244.

Woods, D. D. & Watts, J. C. (1997). How not to have to navigate through too many displays. In M. Helander, T. K. Landauer, & P. Prabhu (Eds.), *Handbook of human–computer cooperation* (2nd rev. ed., pp. 617–650). New York: Elsevier.