19TH INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM
C2 AGILITY: LESSONS LEARNED FROM RESEARCH AND OPERATIONS

VISUALIZING HOST-NATION SENTIMENT AT THE TACTICAL EDGE

Suggested Tracks:
Topic 3: Data, Information, and Knowledge
Topic 8: Social Media

Sue E. Kase[1], PhD – sue.e.kase.civ@mail.mil
Heather E. Roy[1] – heather.e.roy2.ctr@mail.mil
Elizabeth K. Bowman[1], PhD – elizabeth.k.bowman.civ@mail.mil
Daniel N. Cassenti[2], PhD – daniel.n.cassenti.civ@mail.mil

1 - Computational and Information Sciences Directorate
Army Research Laboratory
321 Johnson Street
Aberdeen Proving Ground, MD 21005

2 - Computational and Information Sciences Directorate
Army Research Laboratory
Building 459
Aberdeen Proving Ground, MD 21005

Point of Contact: Sue E. Kase
sue.e.kase.civ@mail.mil
(410) 278-9762

Abstract

Measures of sentiment can help refine a Warfighter's knowledge and understanding of an unfamiliar operational environment, for example, the sentiment of civilians and insurgents to the Army, its operations, or its adversaries. In addition, observing changes in expressed sentiment about topics over time may provide baselines from which to detect if and when important shifts in attitude occur triggering further investigation. The paper discusses new visualization techniques for displaying sentiment offered by social media analysis platforms currently under development by the Army. The sentiment analysis capabilities of these platforms were tested within the realistic military conditions of the C4ISR OTM field exercise using an operational scenario composed of diversified data sources including social media content. A human subjects experiment is planned for more fine-grained testing to better understand the advantages offered by different visualization approaches to sentiment analysis.

.

1. Introduction

Social Media (SM) is a range of applications used for communal interaction via web-based communication platforms. As a largely unfettered means of sharing data, information, opinion, and for orchestrating collective action between and among individuals and groups, SM reaches broad audiences acting as an indicator, an interactive medium, and an archive of behavior in general and sentiment in particular.

We define *sentiment* as an attitude, thought, or judgment prompted by feeling. Sentiment can be relevant on both the group and individual levels. When applied to SM communications emanating from a particular community, indicators of sentiment can provide insights into how a particular group feels about a particular topic. On an individual level, detection of sentiment can indicate an author's suggestion for others to feel a certain way about that topic. Using techniques such as topic detection and sentiment analysis on SM can provide the Army with opportunities to leverage and influence communities in favorable ways.

After a topic or set of topics has been identified and the relevant content extracted, researchers may apply sentiment analysis to analyze and quantify the opinions expressed by SM users toward the topics. Typically, a sentiment analysis system searches for indicators of sentiment, such as lexicons of words or expressions that are in proximity to the topic of interest and aggregate the results into a score. Sentiment scoring can be examined over time by tracking how observed changes correlate with or precursor other external events. These measures of sentiment can help refine a Warfighter's knowledge and understanding of an unfamiliar operational environment, for example, the sentiment of SM civilians and insurgents to the Army, its operations, or its adversaries. In addition, observing topics and changes in expressed sentiment over time may provide baselines from which to detect if and when important shifts in attitude occur triggering further investigation.

Our paper presents new visualization techniques for displaying measures of sentiment from SM analysis platforms currently under development and sponsored by the Army. Each visualization technique is designed to highlight different aspects of sentiment analysis from SM content. We discuss sentiment visualization formats from the perspective of the user interface and a soldier at the tactical edge. Soldiers in the Army's lowest tactical formations need to pass information from the battlefield to a command post at critical decision points of an engagement. The paper concludes by describing our next step—a human subjects experiment investigating accuracy, response time, and preferences of participants performing intelligence analysis relevant tasks using the sentiment analysis capabilities offered by the SM analysis platforms discussed in the previous sections.

2. Social Media Analysis Platforms

Today, very large amounts of information are available through SM services. As part of the effort to better organize this information for analysts, researchers have been actively investigating the problem of automated text extraction and categorization. The bulk of such work has focused on topic categorization, attempting to sort text according to its subject matter. However, recent years have seen rapid growth in categorizing text by sentiment, or the overall opinion towards the subject matter. Labeling text with a sentiment metric can enable rapid and succinct summarization of SM communications for potential intelligence gathering.

Two SM analysis platforms offering topic trending and sentiment analysis are currently under development at the Army Research Laboratory (ARL): the Multi-Source Network Pulse Analyzer and Correlator (MiNPAC) and the Models of Information and Sentiment Transmission (MIST). The development of these technologies was facilitated by the collaboration of ARL and the Office of the Secretary of Defense (OSD) through the Small Business Innovation Research (SBIR) program. The SBIR program is a United States government program that awards competed contracts to small businesses for the development and advancement of future force technologies.

Both MiNPAC and MIST were awarded Phase I funding in 2010 and since then have advanced to Phase II development. In the summer of 2013, both platforms were successfully field tested by military intelligence personnel during the Command Control, Communications, and Computers, Intelligence, Surveillance, and Reconnaissance on the Move (C4ISR OTM) exercise conducted at Fort Dix, New Jersey.

During the C4ISR OTM field exercise MiNPAC and MIST were used by two intelligence analysts over the course of a week to analyze multi-source communications from a region within Kandahar, Afghanistan with little prior knowledge of the socio-cultural environment. This dataset, developed by ARL consisted of modified SM communications such as blog posts and Twitter messages, as well as, a variety of synthetic intelligence reports. The blog posts were collected from various Jihadist websites and modified to support the Kandahar scenario. Tweets were collected using a technology called Apollo (Khac Le et al., 2011) developed through an academic partnership within ARL's Network Science Collaborative Technology Alliance (NS CTA). The analysts used MiNPAC and MIST to extract prevalent topics and sentiment of the civilian population from the modified SM communications and intelligence reports.

The analysts identified the key leaders of several social groups and their primary topics of concern. The key leaders' sentiment overall and towards operationally important topics was monitored over time. The analysts' goal was to leverage this information combined with reported intelligence for planning a key leader engagement that would affect a positive influence on upcoming operational missions. The following subsections provide an overview of the MiNPAC and MIST analysis platforms.

2.1. MiNPAC Dashboard

MiNPAC's development was proposed by Knowledge Based Systems, Inc. (KBSI) as a solution to the 2010 OSD SBIR solicitation topic "Temporal and Conceptual Extractions of Complex Social Network Data." This topic called for the development of an information system utilizing knowledge extraction techniques suitable for large datasets of fused multi-source information that could identify and predict adaptations in a terrorist network.

The MiNPAC technology is a widget-based "Community Pulse Dashboard™" capable of sensing and extracting emerging themes, trends, and patterns of sentiment. It is intended to support an interactive analysis process through four inter-related quadrants: 1.) Theme Monitor, 2.) Sentiment Monitor, 3.) Network Analyzer, and 4.) Sentiment over Time (Figure 2.1).

The upper-left quadrant contains the Theme Monitor which allows the analyst to view themes over time. In order to identify key themes, MiNPAC uses a collection of keyword indicators that may be auto-discovered through the extraction algorithm using data mining techniques, and/or through a list provided by the analyst. The graph consists of a frequency-versus-time analysis of several extracted themes. The themes are legend entries which can hide or show their respective line in the graph to focus viewing.

The upper-right quadrant of MiNPAC contains the Sentiment Monitor which displays various sentiment score readings based on a selected time frame. The sentiment bar colored red (most negative sentiment, -1.0) to green (most positive sentiment, 1.0) displays the average sentiment of the selected metrics and time period. The first drop-down box selects a metric by which to view the data (All, Entities, or Messages). The second drop-down box allows further refinement if available.

In Figure 2.1, 'Entities' is selected in the first drop-down box, then 'Zabihullah Mujahid' in the second drop-down box from a list of entities extracted from the data sources. The sentiment score in the colored sentiment bar is an average of the sentiment from the data items appearing in the right list box. The buttons below this list box allow movement of data items. In Figure 2.1, the Sentiment Monitor shows an entity sentiment view with an average sentiment score of -0.39 calculated from Zabihullah Mujahid's communications. The right list box shows the ID numbers of these communications over the given time period.

The lower-right quadrant, Sentiment Over Time, is connected to the Sentiment Monitor. While the Monitor shows the average sentiment score over a period of time, the Sentiment Over Time graph displays the individual data points used to compute that average. Clicking on one of the data points provides a drill down capability to the actual text content.

The lower-left quadrant of MiNPAC contains the Network Analyzer which visualizes and monitors the social networks identified in the data sources. In a similar manner to the Sentiment Monitor, the drop-down box allows the analyst to select a particular metric by which to organize the data sources. The top five entities of the selected metric are displayed in the graph. The analyst may choose to view the entities based on ascending or descending order. The bar graph below the drop-down box shows the entities involved and their respective values based on the given metric.

Most of the MiNPAC dashboard quadrants contain a date slider that allows the analyst to indicate a date range over which to view the data analysis. In Figure 2.1, the date sliders (colored orange and located at the bottom of each quadrant) can be synchronized across quadrants.

As part of the C4ISR OTM scenario, the average sentiment for Zabihullah Mujahid displayed in the Sentiment Monitor is -0.39, which on a scale of -1.0 to 1.0, is between orange and red in color—the border of 'Negative' and 'Very Negative.' This is the average sentiment of Mujahid's SM communications as a prolific blogist followed by several anti-U.S. social groups.

The data points in the Sentiment Over Time quadrant represent individual blog postings by Mujahid positioned on the same color coded scale from green to red on the y-axis. Using the 'View Text' button, the text content of the blogs can be inspected in the scrolling text box below the graph. These data points show most of Mujahid's communications fall between neutral and negative sentiment with the most recent communications located in the red to orange area of very negative. For example, we see Mujahid referring to U.S presence in Afghanistan as the 'puppet army.'

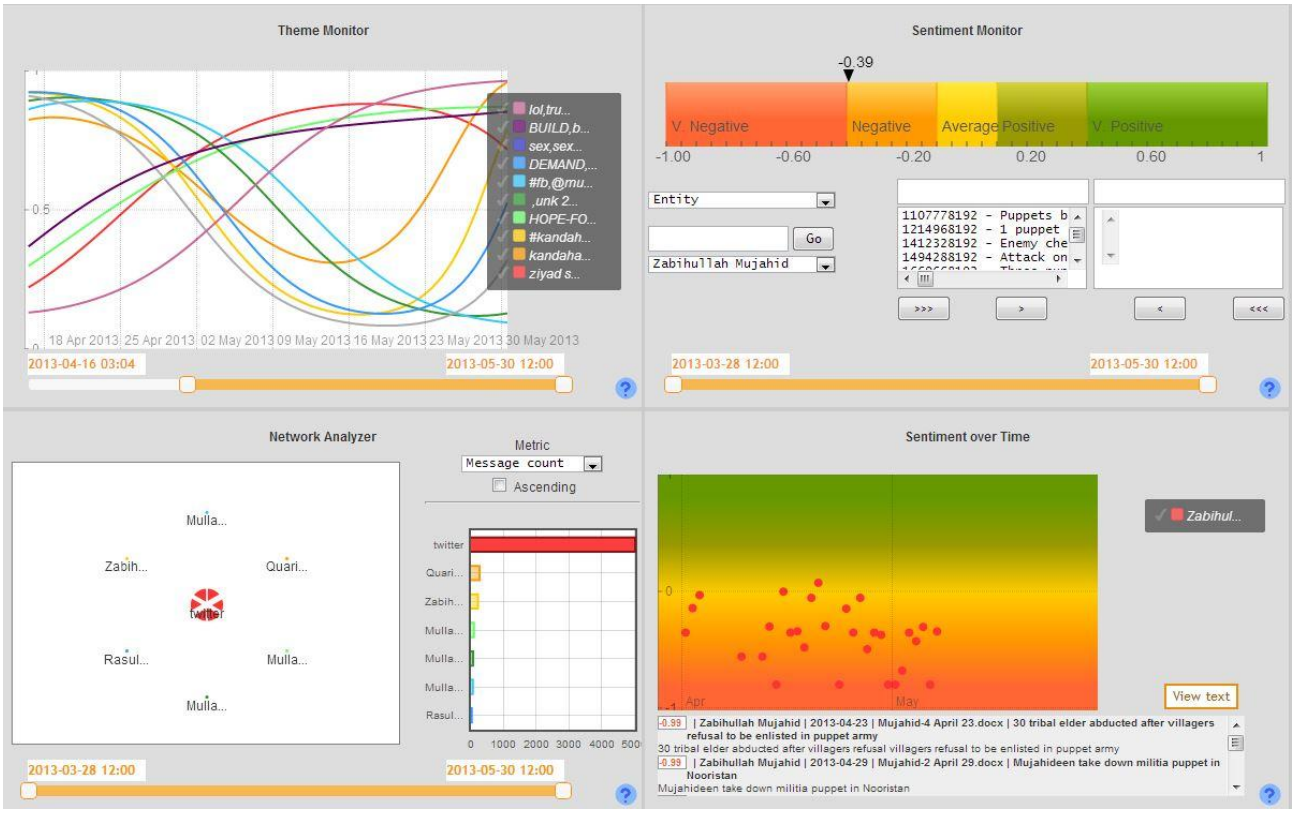

Figure 2.1 The MiNPAC "Community Pulse Dashboard™" develop by Knowledge Based Systems, Inc. (KBSI) consists of four integrated quadrants for extracting emerging themes and patterns of sentiment from multi-source data. The entity sentiment view shows an average sentiment score of -0.39 for the selected entity, Zabihullah Mujahid, who is an anti-U.S. blogist with connections to a drug distribution ring.

2.2. MIST Desktop

MIST's development was proposed by Aptima, Inc. as a solution to the 2010 OSD SBIR solicitation topic "Decision Superiority through Enhanced Cultural Intelligence Forecasting." This topic called for the development of a decision support tool capable of visually presenting complex cultural and social sentiment themes from large datasets for potential forecasting of attitudes and behaviors in a cultural context.

In response to the solicitation, Aptima developed MIST, a sentiment analysis desktop capable of calculating sentiment and topic prevalence of persons of interest by automatically monitoring and analyzing multi-source data streams. Different from MiNPAC, MIST portrays the local population's sentiment towards key individuals within the community in reference to topics of concern. This is an alternative perspective to MiNPAC's sentiment analysis where the sentiment of an author's communications is analyzed.

MIST's level of analysis also differs from MiNPAC in that it allows analysts to examine relationships between entities, topics, and sentiment at a pattern-level of abstraction. One of the benefits of this high-level visualization format is *discovery*—the ability to spontaneously discover relationships without knowing to look for them (Plaisant, 2004). In this way, MIST assists analysts in perceiving patterns in the data by exploring the data through comparative forms of topic and sentiment trend analysis. MIST capitalizes on analysts' perceptual abilities to identify visual patterns by measuring and visualizing comparisons at various levels of complexity (i.e., multiple entities for multiple topics and corresponding sentiment over time).

MIST leverages Aptima's Latent Variable Analysis (LaVA) automated language processing technology in order to extract entities and topics from multi-source data. Sentiment analysis is preformed and the entities' sentiment and topics are combined to derive 'memes' or units of cultural information about the attitudes and opinions observed in the data. The visualization format for the memes is inspired by Tufte's (2006) *sparklines*—a compact and uncluttered design to draw the analyst's attention to important trends and significant insights.

Left side of Figure 2.2 shows the MIST analysis desktop at the top level. The graphs to the right side of each entity image are sentiment and prevalence sparklines. Sentiment is on the y-axis where 1 represents positive sentiment, and -1 negative sentiment. Time is on the x-axis. Prevalence is similar except the y-axis scale is positive. In Figure 2.2, the seven entities pictured on the left side have been identified from the data as community leaders. Along the right side of each entity's image are the four top topics they have been associated with. For example, Habib Ala Ahmed is most often associated with the topics drug trade, attacks, patrols, and foreign invaders. Ahmed's sentiment and prevalence graphs are computed averages across his top topics. In the C4ISR OTM scenario, Habib Ala Ahmed is a police chief suspected of being corrupt. The second half of his sentiment graph shows variability with both positive and negative spikes. This pattern may signal to the analyst that a more in depth investigation is needed.

In the C4ISR OTM scenario, Zyad Guda Sultan is a known criminal associated with a drug distribution ring. The second half of Sultan's sentiment and prevalence graphs show some variability appearing about the same time as the spikes in Ahmed's graphs. The graph patterns of Ahmed and Sultan can be investigated further using an entity comparison level of analysis.

The screen capture on the right side of Figure 2.2 shows the Entity Comparison screen which results from requesting a comparison of two or three of the entities by checking the 'compare' boxes under each entity's image and then clicking the blue 'Compare' button in the upper right corner. A comparison between Habib Ala Ahmed and Zyad Guda Sultan has been requested by the analyst.

On the Entity Comparison screen, the analyst can select a specific topic to compare across entities. The topic 'Drug Trade' has been selected in the drop-down box at the top of the screen. The top portion of the Entity Comparison screen now displays a color coded comparison of sentiment and prevalence for Ahmed and Sultan associated with the topic Drug Trade. The bottom portion of the Entity Comparison screen shows Ahmed and Sultan's sentiment and prevalence averaged across all topics and each individual topic aligned horizontally side-by-side. In Figure 2.2, the Entity Comparison screen shows both Ahmed and Sultan with negative sentiment spikes about the same time, although Sultan's is less so, and corresponding spikes in prevalence. These patterns occurred as a result of SM communications expressing confusion and anger over Ahmed's release of two suspected drug traffickers under suspicious circumstances in the C4ISR OTM scenario.

The next section describes how SM analysis platforms can diversify intelligence sources and the role they could play in future Army operations. Two intelligence relevant examples are discussed in the context of MiNPAC and MIST sentiment analysis capabilities.

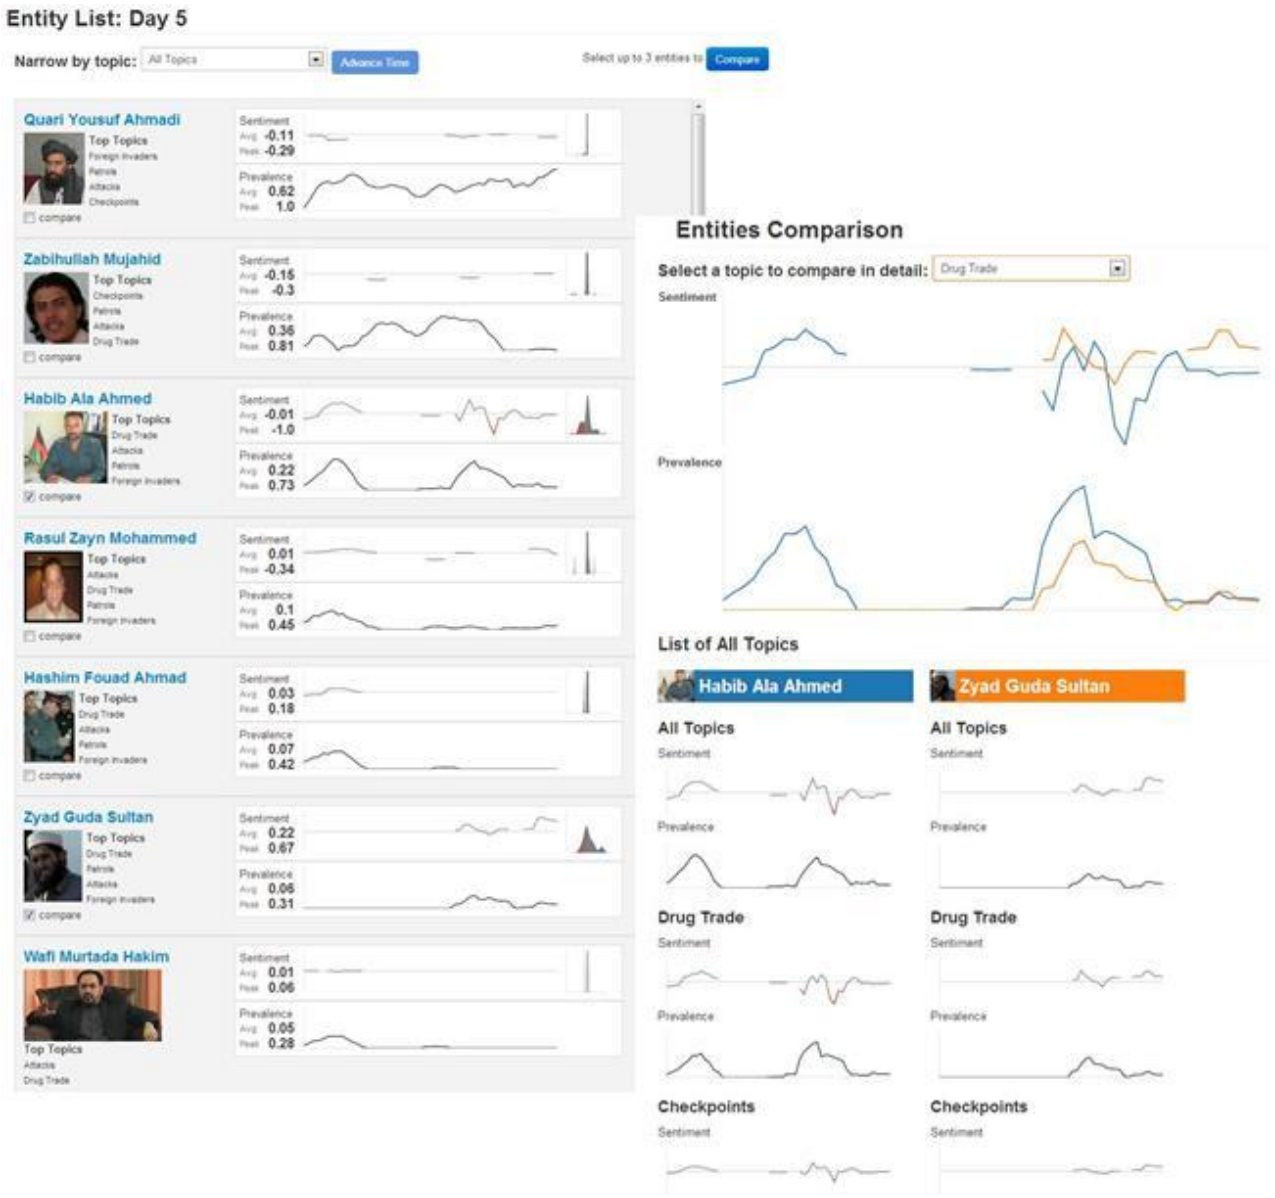

Figure 2.2 The MIST desktop developed by Aptima, Inc. examines relationships between entities, topics, and sentiment at a high-level of abstraction using Tufte-like sparklines (Tufte, 2006). The left screen capture shows seven of the most influential individuals extracted from the data. The right screen capture compares two of the entities "Habib Ala Ahmed" and "Zyad Guda Sultan" for sentiment and prevalence on the topic "Drug Trade."

3. Discussion

While MiNPAC and MIST provide a range of operational benefits, ARL views these systems within a larger context emphasized by Major General Michael Flynn in the need for a dramatic change in the organization and focus of intelligence operations (Flynn, Pottinger, & Batchelor, 2010). Flynn recognized a shift in recent military operations from standard, kinetic warfare to stability, security, transition and reconstruction (SSTR) operations, and argued there is a growing recognition that the lack of reliable information sources in the intelligence community may well be the major limiting constraint in developing an understanding of the socio-cultural landscape.

Today, there is increased ambiguity surrounding information sources as soldiers rely on more diverse intelligence sources during operations. Human intelligence (HUMINT) and open source intelligence (OSINT) become prioritized above more traditional intelligence based on signals (SIGINT) and electronic sources (ELINT). In addition to the change in information sources, there appears to be increased information management requirements associated with heavy reliance on HUMINT and OSINT. Soldiers as peacekeepers must manage multiple information assets and resources, often relying on local and international organizations, as well as, the local populace and SM content.

The above concerns set the stage for a series of OSD SBIR opportunities promoting technology advances in the processing and analysis of big data such as SM generated content. Currently, ARL is guiding the development of the MiNPAC and MIST systems. These SM analysis platforms are designed to organize multi-source data within an area of operations, and enable insightful analysis via visualizations of topic trending and sentiment analysis.

Here the discussion focuses on the sentiment analysis capabilities of MiNPAC and MIST. Both these systems categorize a document as positive, negative, or neutral based on a scale from 1.0 to -1.0. MiNPAC provides a measure of the author's sentiment—how positive/negative their writing is. In MiNPAC an author's documents are classified in two ways: by overall sentiment across documents (Sentiment Monitor quadrant), and by individual document (Sentiment Over Time quadrant). In both quadrants the color-coded sentiment scale from red to green supports agile leveraging and employment of sentiment in a time constrained operational tempo. This multi-level perspective is valuable for identifying inappropriate affect levels which might escalate to undesirable behavior detrimental to the mission.

For example, an author's blog postings that are consistently very negative and becoming more frequent would signal further investigation. Using MiNPAC an analyst can track how the author behaves over time, that is, what is the history of the author's sentiment. In Figure 2.1, Sentiment Over Time shows Zabihullan Mujahid's blog postings start off mostly neutral to slightly negative and then become increasing more negative over time. A trend of increased negativity and frequency of postings could indicate the author's intend to recruit like minded 'followers' and promote violence. The Sentiment Monitor presents an aggregate numeric sentiment score for Mujahid's communications of -0.39 positioned above the color scale. A numeric sentiment score is useful for command level reporting and tracking over time.

Whereas MiNPAC provides author generated measures of sentiment, MIST provides community measures of sentiment—how community members feel about leaders and their relationship to topics of concern. MIST's innovative pattern approach employs sparkline graph presentations of memes—sentiment towards an entity in the context of a topic over time. This type of visualization is easily interpretable by the analyst and helps suggest potential courses of action and interventions for command level decision makers.

The top level of the MIST desktop provides the analyst the ability to assess the changing attitudes of a population towards its community leaders. Then the entity comparison level enables the analyst to quickly and intuitively view comparisons among selected community leaders in the context of a particular topic.

For example, Habib Ala Ahmed's Drug Trade sentiment and prevalence graphs (right side of Figure 2.2, top graphs, blue lines) indicate MIST has detected the public's negative comments concerning the drug trade. The analyst sees Ahmed is gaining a negative reputation and investigates this further. The analyst learns that Ahmed released several suspected drug traffickers and the public is now accusing him of possible involvement in the drug distribution ring. Although his reputation appears to be on the rebound, the analyst can make the prediction that Ahmed will be increasingly susceptible to drug-related criticism. This information can be used by command level decision makers to plan a proactive engagement with Ahmed to discuss shortfalls in law enforcement and potential avenues of assistance.

Mujahid and Ahmed are characters in an operational scenario designed by ARL for the purpose of testing MiNPAC and MIST within the realistic military conditions of the C4ISR OTM field exercise. This operational scenario is composed of diversified data sources both SM and report-like documents (i.e., blog postings, tweets, news reports, and intelligence reports). The scenario begins shortly after the arrival of a unit responsible for an area within Kandahar, Afghanistan where little is known about the socio-cultural environment. The unit's mission is to build the area into an economically productive region supportive of the Afghan national government and U.S. presence. During the exercise, two intelligence analysts used MiNPAC and MIST to identify the key leaders of several pro- and anti-U.S groups, their primary topics of concern, and their corresponding sentiment. The exercise concluded with the analysts successfully planning a key leader engagement for ameliorating localized concerns, influencing the sentiment of the local populace, and gaining support for U.S. presence.

Despite successful field testing of both MiNPAC and MIST during the C4ISR OTM exercise, more fine-grained testing is needed to understand the advantages offered by the different types visual analytics for capturing sentiment. Our next step is to conduct a human subjects experiment involving the visualization capabilities of both MiNPAC and MIST.

4. Path Forward

A human subjects experiment will be conducted to discover if the visualization capabilities of the MiNPAC and MIST analytic platforms affect the accuracy, response time (RT), and preferences of participants performing intelligence analysis relevant tasks. As discussed in the previous sections these systems involve a variety of interface components for sensing emerging trends and patterns of sentiment from multi-source data. The primary utility of the systems is to understand how leaders in the community are influencing the sentiments of the populace. Although the MiNPAC and MIST systems analyze different measures (the attribution of the leader versus the attribution of the populace regarding the leader), the purpose to measure sentiment is common between the systems.

An area of interest is the nature of the visual analytics presented to the participants, which will answer questions on how to present network information extracted from SM communications. An example is how combining sentiments for different topics with one leader vs. comparison of different leaders on common topics affects the ability of the participant to determine a rank order of leaders to solicit for help. Another example is how information presented on one screen vs. coordinating between multiple screens affects performance.

4.1 Method

Participants will be instructed to analyze sentiments of leaders in online political discourse to determine rank orders of preferred leaders in recruiting to accomplish to-be-determined missions. They will make these judgments based on the information they receive from the systems they engage. In one condition, the participants will have only the blog entries from the leaders in text format (i.e., the control condition). A second condition will allow participants to use The MiNPAC dashboard. The third condition will allow participants to use the MIST desktop.

The conditions will be presented within subjects and counterbalanced across at least 15 participants. Dependent variables collected during trials include the following: proportion correct of answering questions regarding sentiment of entities, overall response time (i.e., from start of problem to completion), and sequence with time stamp of actions.

Following each trial, participants will be asked to rate their confidence in their choices. Following each block, participants will be asked to take an abbreviated form of the NASA-TLX to rate mental workload and a trust in software questionnaire to understand whether participants viewed the software positively or negatively, which will provide human factors feedback on the affective usability of the analysis platforms. The trust scale will be a

modification of Jian, Bisantz, and Drury (2000) trust-in-automation and used previously in Cassenti and Chan (2012).

4.2 Data Analysis

The data analysis will include ANOVAs and follow-up *t*-tests on the accuracy, overall RT, frequency of resource use, workload scale, and trust scale results to determine the relative effectiveness of the analysis platforms to convey useful information to participants in terms of performance and subjective assessment. In addition, a visual analysis of sequence of action data will be performed to determine characteristic patterns of action for possible future curve estimation analysis.

This will be the first of a series of human subjects experiments designed to explore the benefits of sentiment analysis through visualization techniques. The results of these experiments will inform ARL's continued development of a suite of tools to support today's Warfighters and intelligence personnel, and leverage a wide range of non-kinetic data sources to shape analysis and decision making from the company level on up.

5. Conclusion

A shift is under way in recent military operations from standard, kinetic warfare to stability, security, transition and reconstruction (SSTR) operations. There is growing recognition that new information sources and text processing techniques are needed by the intelligence community to understand the socio-cultural landscape of an area of operations. To meet this emerging need, OSD and ARL are making significant investments in purposefully funding and managing the development of SM analysis platforms with sentiment analysis and topic trending capabilities.

Two SM analysis platforms currently under development, MiNPAC and MIST, extract and visually present social sentiment themes from large datasets. MiNPAC provides measures of author based sentiment—how positive/negative are their communications to the public; while MIST provides community measures of sentiment—how community members feel about their leaders in reference to topics of concern.

The visualizations capabilities of MiNPAC and MIST were tested during the C4ISR OTM field exercise by two intelligence analysts over the course of a week to analyze multi-source communications from a region within Kandahar, Afghanistan with little prior knowledge of the socio-cultural environment. The analysts successfully identified the key leaders of several social groups and their primary topics of concern by monitoring sentiment and topic trending over time. As a next step, a human subjects experiment is planning for a more fine-grained investigation of the visualization capabilities offered by MiNPAC and MIST.

Acknowledgements

The authors acknowledge the developers of the MiNPAC technology, Knowledge Based Systems, Inc. (KBSI), and their supporting scientists: Dr. Timothy P. Darr and Dr. Perakath Benjamin. Additional, acknowledgements are due to the developers of the MIST technology, Aptima, Inc., and their supporting scientists: Drs. Bob McCormack, Dustin Burke, Mike Paley, and Ian Yohai. These technologies are currently under development with funding from the Office of the Secretary of Defense (OSD) through the Small Business Innovative Research (SBIR) program.

This research was supported in part by an appointment to the Student Research Participation Program at the U.S. Army Research Laboratory (ARL), administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and ARL.

References

Cassenti, D. N., & Chan, K. (2012). *Network science for decision making: Performance, trust, and workload with information from a network.* (ARL-TR-5890). Aberdeen Proving Ground, MD: U.S. Army Research Laboratory.

Flynn, M. T., Pottinger, M., & Batchelor, P. D. (2010). *Fixing Intel: A blueprint for making intelligence relevant in Afghanistan.* Washington DC: Center for a New American Security.

Jian, J. Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, *4*, 53-71.

Khac Le, H, Pasternack, J., Ahmadi, H., Gupta, M., Sun, Y., Abdelzaher, T., Adali, S. (2011). Apollo: Towards factfinding in participatory sensing. Proceedings from IPSN: *The 10th International Conference on Information Processing in Sensor Networks*. Chicago, IL.

Plaisant, C. (2004). The challenge of information visualization evaluation. Proceedings from AVI: *The working conference on Advanced Visual Interfaces*. New York, NY: ACM.

Tufte, E.R. (2006). *Beautiful Evidence.* Cheshire, CT: Graphics Press.