Semantic Visualization

Penny Chase, Ray D'Amore*, Nahum Gershon*, Rod Holland, Rob Hyland, Inderjeet Mani*, Mark Maybury, Andy Merlino, Jim Rayson

The MITRE Corporation
202 Burlington Road
Bedford, MA 01730, USA
*1820 Dolley Madison Blvd.
McLean, VA 22102, USA

{pc, rdamore, gershon, rholland, hyland, imani, maybury, andy, jrayson}@mitre.org

ACL-COLING Workshop on Content Visualization and Intermedia Representations

Abstract

This paper summarizes several initiatives at MITRE that are investigating the visualization of a range of content. We present results of our work in relevancy visualization, news visualization, world events visualization and sensor/battlefield visualization to enhance user interaction in information access and exploitation tasks. We summarize several initiatives we are currently pursuing and enumerate unsolved problems.

1. Visualizing Semantic Content

Visualization can support effective and efficient interaction with a range of information for a variety of tasks. As Figure 1 illustrates, information (data elements, attributes, relations, events) can be encoded in (possibly interactive) visual displays which users can exploit for a variety of cognitive tasks such as retrieval, analysis (e.g., of trends, anomalies, relations), summarization, and inference. In this paper we consider a range of semantic content, visual mechanisms, and cognitive tasks to deepen our understanding of the role of interactive visualization.

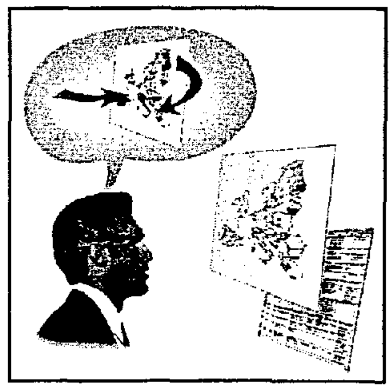

Figure 1. Information Visualization Process

2. Document Relevancy Visualization

Today's users are faced with a dizzying array of information sources. MITRE's Forager for Information on the SuperHighway (FISH) (Smotroff, Hirschman, and Bayer 1995) was developed to enable the rapid evaluation of information sources and servers. Figure 2a illustrates the application of FISH to three Wide Area Information Server (WAIS)™ databases containing information on joint ventures from the Message Understanding Conference (MUC). Figure 2b illustrates the application of FISH to visualize e-mail clustered by topic type for a moderator supporting a National Performance Review electronic town hall.

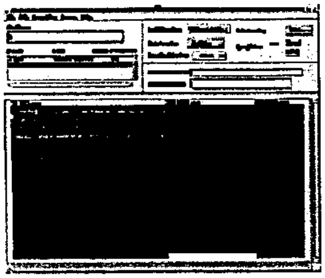

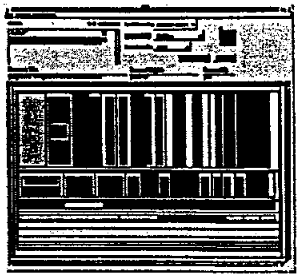

Figure 2a. WAIS FISH Figure 2b. NPR FISH

The traditional WAIS interface of a query box and a list of resulting hits is replaced by an interface which includes a query box, a historical list of queries, and a graphically encoded display of resulting hits (an example of which is shown in Figure 2a). In WAIS, the relevancy of a document to a given keyword query is measured on a scale from 1-1000 (where 1000 is the highest relevancy) by the frequency and location of (stems of) query keywords in documents. Motivated by the University of Maryland's TreeMap research for hierarchical information visualization, FISH encodes the relevance of each document to a given query (or set of compound queries) using both color saturation and size.

In the example presented in Figure 2a, each database is allocated screen size in proportion to the number of and degree with which documents are relevant to the given query. For example, the MEAD database on the left of the output window is given more space than the PROMT database in the middle because it has many more relevant documents. Similarly, individual documents that have higher relevancy measures for a given query are given proportionally more space and a higher color saturation. In this manner, a user can rapidly scan several large lists of documents to find relevant ones by focusing on those with higher color saturation and more space. Compound queries can be formulated via the "Document Restrictions" menu by selecting the union or intersection of previous queries, in effect an AND or OR Boolean operator across queries.

In Figure 2a, the user has selected the union of documents relevant to the query "japan" and the query "automobile", which will return all documents which contain the keywords "japan" or "automobile". Color coding can be varied on these documents, for example, to keep their color saturation distinct (e.g., blue vs. red) to enable rapid contrast of hits across queries within databases (e.g., hits on Japan vs. hits on automobile) or to mix their saturation so that intersecting keyword hits can be visualized (e.g., bright blue-reds could indicate highly relevant Japanese automobile documents, dark the opposite). In the example in Figure 2a, blue encodes Japan, red Automobile; the color coding is set for mixed saturation, the union of the relevant document sets for those two keywords is selected, and the order (from top to bottom in the display) is used to encode the WAIS relevancy ranking. One issue is just how effectively users can discriminate mixed colors.

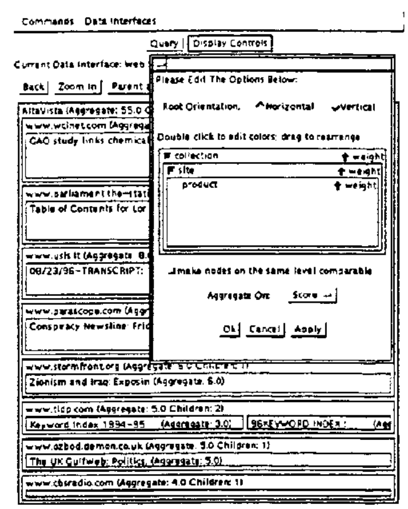

Figure 2c. J-FISH Multiserver Visualization

More recently, we have explored multiple server evaluation on popular World Wide Web search engines. For example, Figure 2c illustrates a query across multiple servers. Research issues include differences in relevancy ranking algorithms, encoding of multiple attributes beyond relevancy using color or size (e.g., length, quality, cost, source), and document collections which are heterogeneous in size, content, and format.

3. Document Structure/Content Visualization

Figure 3a (Gershon et al. 1995; Gershon 1996) illustrates another navigation mechanism in which the user is able to view a hierarchy of the browse space. The left hand of Figure 3a displays the traditional HTML layout of a web page whereas the right hand side illustrates a hierarchical, navigable view automatically generated from the underlying structure of the browsing space. The user can create a personal space by interactively and visually modify the structure of hyperspace or extracting segments of the documents.

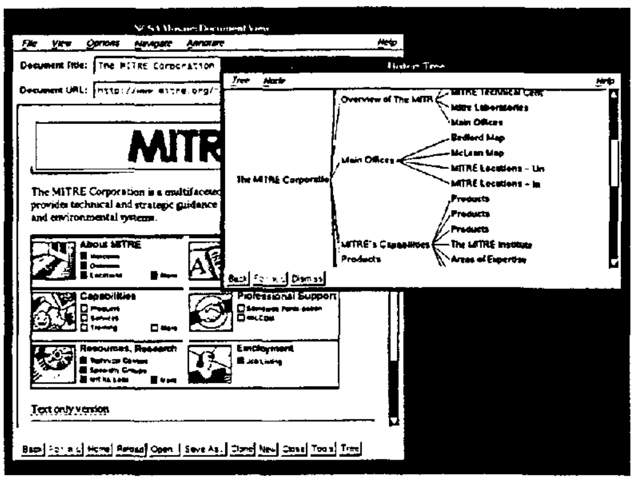

Figure 3a. Hyperspace Structure Visualization

For discovery and analysis of new information and relationships in retrieved documents, we have developed a method for aggregating relevant information and representing it visually (Gershon, et al, 1995). The method is based on representing correlations of words within a document in a table. These tables could be very large depending on the size of the document thus making it difficult for the user to perceive and make sense of all the highly relevant correlations. Since the order of the words is not usually based on contents, the order of the words is permuted until the highly relevant correlations are concentrated in one corner.

	Fat	Smoking	Snacks	sedentary
Fatigue	22	3	47	2
Aches	10	3	11	4
Alments	2	4	4	3
nausea	15	5	33	1

	fat	snacks	smoking	sedentary
Fatigue	22	47	3	2
Aches	10	11	3	4
Nausea	15	33	5	1
Ailments	2	4	4	3

Figure 3b. Example of Unaggregated (top) and Aggregated (bottom) Tables

Other research at MITRE has focused on automatic discovery and visualization of semantic relations among individual and groups of documents (Mani and Bloedorn 1997). Figure 3c illustrates the results of visualization of a set of documents using the NetMap visualization software after clustering these into related groups which appear around a circle. Outside of each cluster on the circle are displayed intracluster relations; in the center of the circle are intercluster relations (e.g., a shared named entity such as a person, place, or thing which appears in multiple documents). The user can zoom in any part of the graph. This is shown in Figure 3d, which shows individual people (green) and organizations (aquamarine).

Selecting an individual entity from a document returns a display such as that in Figure 3e. Figure 3e illustrates individual entities encoded with color and shapes (e.g., people in green stick figures, organizations in aquamarine diamonds, locations in purple jagged rectangles, documents in yellow circles, person-organization relations in white squares). Lines and their properties (e.g., color, dashed) can encode relations among these entities (e.g., co-occurrence in documents). This provides a richer mechanism for discovering interdocument and interentity relationships during analysis. Current research is investigating the role of automated text summarization, document retrieval and navigation and visualization.

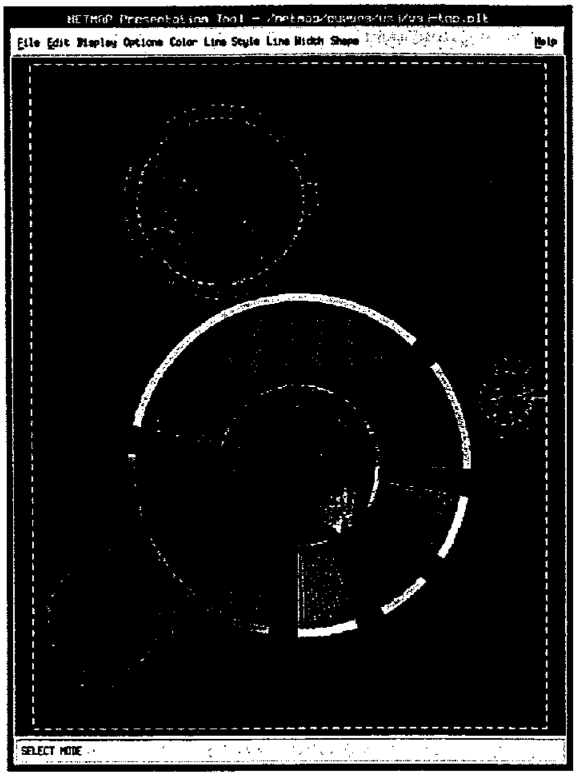

Figure 3c. Document Cluster Visualization

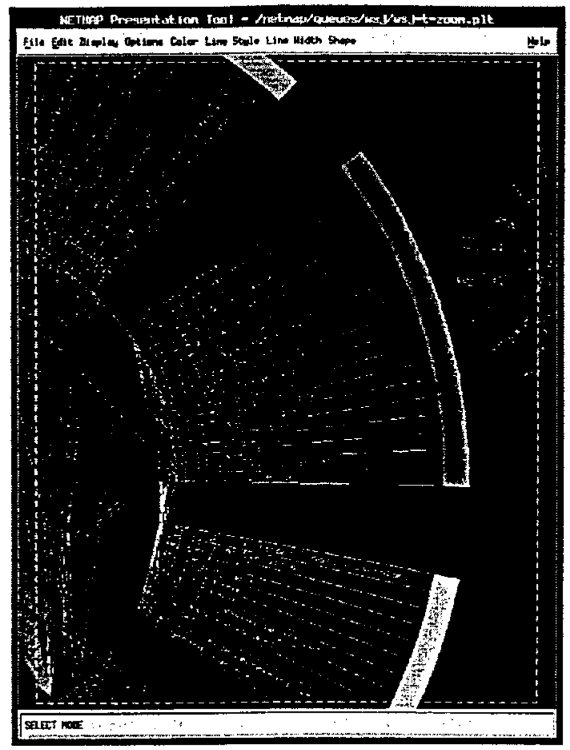

Figure 3d. Zooming in on Document Cluster

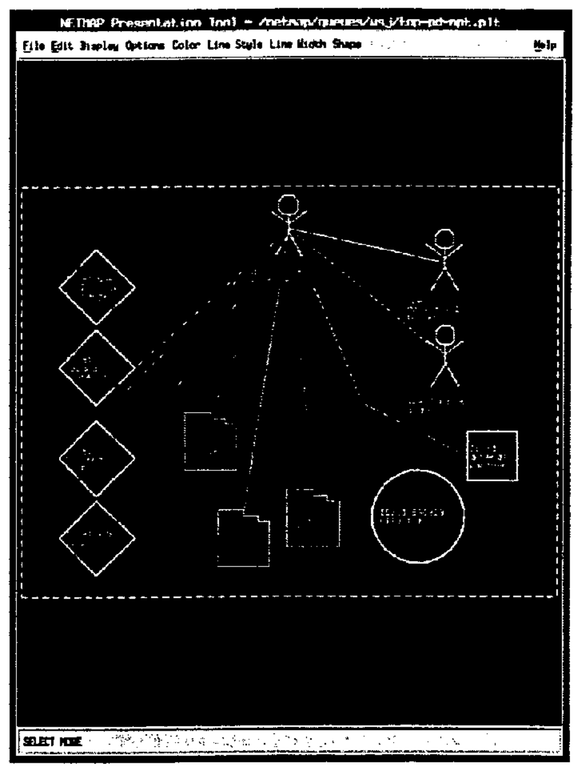

Figure 3e. Entity Relation Visualization

4. Named Entity/News Visualization

MITRE's Broadcast News Navigator (BNN) is a system that is investigating analysis of trends in news reporting. BNN performs multistream (audio, video, text) analysis to eliminate commercials, segment stories, extract named entities (i.e., people, organization, location) and keyframes, and classify and summarize stories (Merlino, Morey, and Maybury 1997). BNN's intuitive web-based interface gives the user the ability to browse, query, extract from and customize digitized broadcasts. Figure 4 illustrates a trend analysis display from BNN that shows the most frequently mentioned named entities reported on CNN Prime News™ from October to November of 1997. "China" spikes in the center of the graph, associated with a state visit to Washington. Later "Iraq" spikes which is correlated with news regarding UN site inspections. The user can click on any point on the line graphs and be brought to a list of stories that mention that named entity.

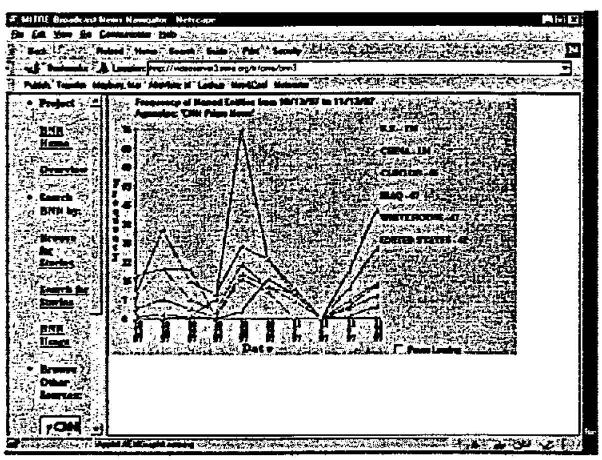

Figure 4. Broadcast News Visualization[1]

In contrast, the user can formulate a query specifying keywords, named entities or subjects. Figure 5a shows the results of executing the query: Find me stories which have a topic of "chemicals", the keywords "chemical weapons", person "Sadam Hussein", organization "Pentagon", and location "Iraq". Each story in this "Story Skim" view is represented by a keyframe and the three most frequent named entities. Selecting one of these stories yields a "Story Detail" display, which as shown in Figure 5b including a keyframe, named entities, subject classification and pointers to the closed caption and video source.

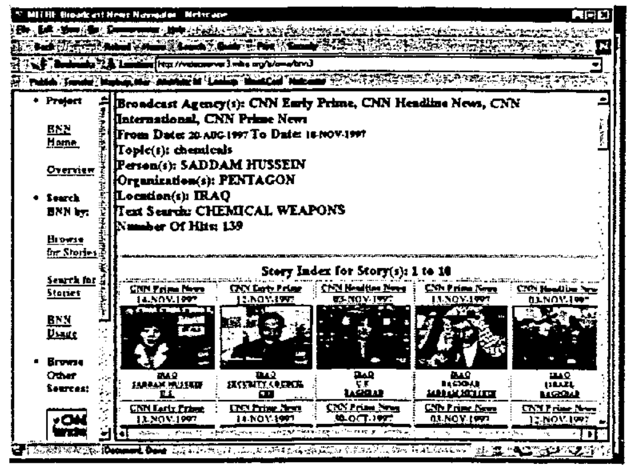

Figure 5a. BNN "Story Skim" Visualization

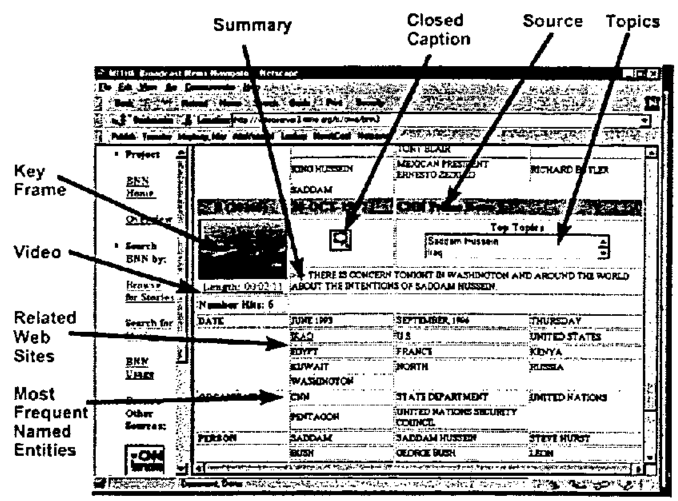

Figure 5b. "Story Detail" visualization

Current research is exploring connecting these broadcast news stories with internet stories, visualizing topic frequencies over time, and mechanisms for low quality spoken language transcriptions of broadcast stories. Other

[1] Note in the display the occurrence of the terms "U.S." and "United States". BNN performs no co-reference resolution, a topic of current research at MITRE.

investigations are focusing on which presentation mixes (e.g., keyframes, named entities, one line summary, full video source) are most effective for story retrieval and fact extraction from news (Merlino and Maybury 1998).

5. Geographic Event Visualization

The Geospatial News on Demand Environment (GeoNODE) initiative at MITRE is a new project investigating visualizing geographic aspects of news events. This program builds on MITRE's BNN, described in the previous section, and MSIIA, addressed in the subsequent section. GeoNODE is based on the research area of Geographic Visualization which investigates methods and tools that impact the way scientists and others conceptualize and explore georeferenced data, make decisions critical to society, and learn about the world (MacEachren and Ganter 1990, Taylor 1991). Since news reports are about events in the world, the reported events and trends can be assessed, queried, and reviewed effectively by leveraging a person's preexisting knowledge of the world's geography. The objective of GeoNODE is to understand the information integration of geospatial/temporal visualizations, information retrieval, multimedia, and other technologies to support browsing, analysis, and rapid inference from broadcast news.

As shown in Figure 6, GoeNODE will analyze global and local cooperation and conflict found in broadcast news, internet, newswire and radio sources as well as broadcast news. Processing will include the identification, extraction, and summarization of events from national and international sources. GeoNODE will consider event types (e.g., terrorist acts, narcotrafficking, peace accords), frequency, and severity in an interactive geo-spatial/temporal context that supports browsing, retrieval, analysis and inference.

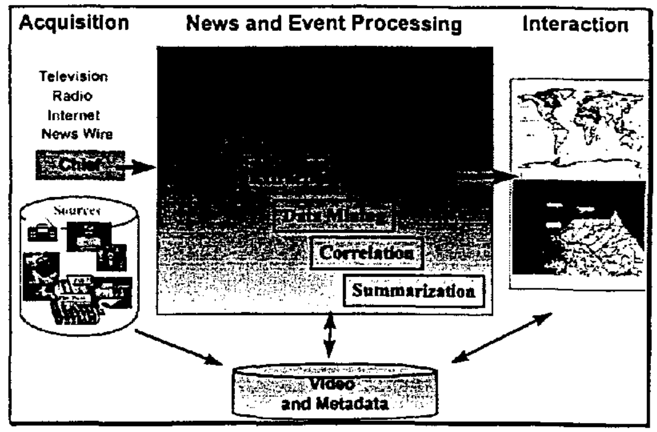

Figure 6. GeoNODE Architecture

Although a geographical context can enhance a person's understanding of reported events and therefore facilitate news retrieval and further queries, the same familiar visualization concerns apply to geographic presentation that are salient in visualizing any data rich multivariate information space. The GeoNODE user experience is derived from research, experience and standard practice in the visual search and retrieval domains: Overview first, zoom and filter, then details-on-demand (Shneiderman 1994). During each stage of the visualization process, cartographic methods and spatial analysis techniques are applied. These can be considered as a kind of grammar that allows for the optimal design, production and use of maps, depending on the application (Kraat 1997). Select cartographic generalization operators are applied to address key multi-scale and information overload problems (Buttenfield 1991).

GeoNODE addresses Knowledge Representation (KR) and information fusion issues that are important to the news event presentation. The KR activities specific to GeoNODE are concerned with discovering and manipulating geospatial and temporal information, specifically investigating the following:

- improved natural language processing of place names that are central to understanding a news report

- news event modeling
- cartographic generalization rules
- transformation of news events to visual metaphors

Spatial information management is currently growing in its utility to commercial applications, and several industries have already begun to explicitly rely on GIS systems, although most (53%) companies are evaluating while an average of only 7% are implementing or using a GIS (IDC 1997). Accompanying the growing interest in spatial information is a technology trend influencing the architecture of GeoNODE, mainly, a shift from single-purpose/standalone GIS applications to geospatial extensions and services for databases, component frameworks, data warehouses and data analysis applications. By supporting a component-based architecture, GeoNODE can more readily take advantage of future geospatial services and an expanding number of news sources (internet, newswire, radio, and other broadcast sources).

Further research will investigate incorporation of summarization, geospatial/temporal KR, and other traditional visualization techniques. For example, Figure 7 illustrates some of the kinds of visualizations that are being explored by other researchers, such as the use of color and geolocation to encode relations among geographic entries. Figure 7 is a geographic visualization of early WWW usage available at http://www.cybergeography.org/atlas/atlas.html. These and other research threads will shape GeoNODE into a visualization component for reasoning about news events in geographic space. As a long term objective, the system architecture should allow for navigation and retrieval from topic, conceptual, and web spaces where a user can access, update and annotate existing data with spatial information.

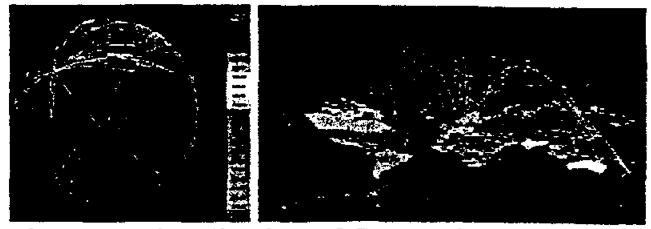

Figure 7. Visualization of Geospatial Relationships

6. Sensor Visualization

The Multisource Integrated Information Analysis (MSIIA) project, led by Steve Hansen at MITRE, is exploring effective mechanisms for sensor and battlefield visualization. For example, national and military intelligence analysts are charged with monitoring and exploiting dozens of sources of information in real time. These range from sensors which capture images (infrared, electro-optical, multispectral) to moving target indicators characterized by a point and some features (e.g., tracked vs. wheeled vehicle) to signals intelligence characterized by centroids and error elipses. Knowing which source to select and which sensors to task is paramount to successful situation assessment. An integrated view into what sensors are where when, as well as a fused picture of their disparate information types and outputs, would be invaluable. Figure 8 illustrates one such visualization. The x-y dimension of the upper display captures the coordinates of a geospatial area whereas the y coordinate displays time. This enables the user to view which areas are being sensed by which type of sensor (encoded by color or implicitly by the resultant characteristic shape). For example, a large purple cylinder represents the area over time imaged by a geosychronous satellite, the green cylinders are images taken over time of spots on the surface of the earth, whereas the wavy blue line is the ground track of a sensor flying across an area (e.g., characteristic of a unmanned air vehicle such as predator). If we take a slice at a particular time of the upper display in Figure 8 we get the coverage of particular areas from a specific time. If we project all sensor coverages over an area downward to the surface, we obtain the image shown in the lower display of Figure 8.

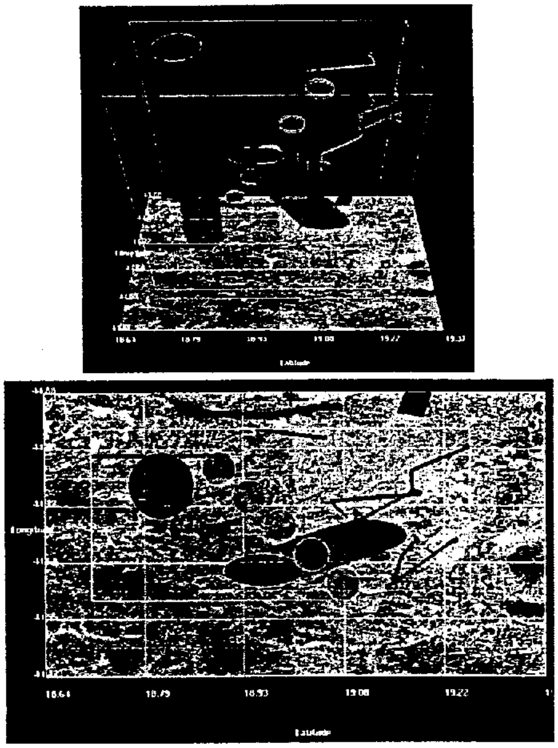

Figure 8. Sensor Coverage Visualization

A user can utilize this display to determine what material is available for a given time and space, analyze unattended coverage areas, and plan future collections. MSIIA is also investigating georegistration and display of the results of collections in an integrated, synthetic view of the world (e.g., fusing maps with images with radar returns). We consider next another example of synthetic views of the world.

7. Collaboration and Battlefield Visualization

Just as visualization plays an important role in information space visualization for MSIIA, MITRE's research on the Collaborative Omniscient Sandtable Metaphor (COSM) seeks to define a new generation of human-machine interfaces for military Command and Control (C2). The "sandtable" underlying COSM is a physical table whose top is rimmed with short walls and filled with sand. It is used in the military for planning and training. The sand can be sculpted to match the terrain in a specific geographic region. People standing around the table can place plastic or metal models of vehicles and other assets over this terrain model to indicate force deployment and move them around the terrain to indicate and/or rehearse force movements.

In defining COSM, we expanded the functionality of a sandtable and moved it into an electronic domain. It now taps into global gigabyte databases of C2 information which range from static data on airfield locations, to real-time feeds from hundreds ground, air, and space based sensors. This data is used to synthesize macroscopic or microscopic views of the world that form the foundation of a collaborative visualization system. Manipulating these views leads not only to modifying data, but also directing the actions of the underlying physical assets (e.g., moving an icon causes an aircraft to be redirected from point A to point B). A conceptual view of COSM is shown in Figure 9, where participants at air, land, and sea locations collaborate over an electronic sandtable. Some users are physically present, while others are represented by their avatars.

The key elements of COSM are geographic independence (transparent access to people, data, software, or assets regardless of location), a multimodal, direct manipulation interface with an initial emphasis is on the visual modality, heterogeneous platform support (enabling users to tailor data depictions to a range of platform capabilities), and data linkage (maintaining all parent, child, and peer relationships in the data).

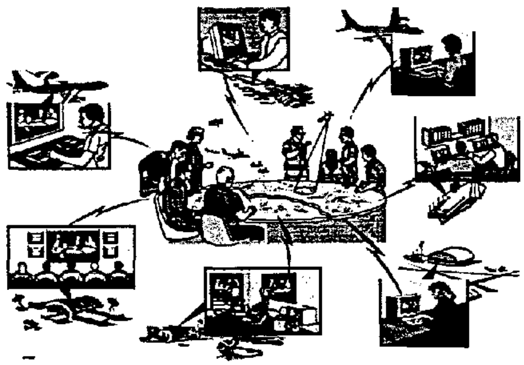

Figure 9. Conceptual View of COSM

Figure 10. Virtual Reality Instantiation

A first instantiation of COSM was implemented using Virtual Reality (VR) technology, as illustrated in Figure 10. The table is a stereoscopic projection system driven by a graphics workstation. It uses a horizontal display surface approximately 6 feet wide and 4 feet deep to display maps, imagery, and models of the terrain and objects upon or above the terrain. Since it is stereoscopic, objects above the terrain, such as airborne aircraft, appear to be above the surface of the table. The vertical screen behind the table is a rear-projection display used primarily used for collaboration support. At the top, we see a panel of faces representing all the remote users who have similar systems and are currently connected to this one with audio, video, and data links. The table serves as a shared whiteboard that is visible to all the users and can be manipulated by them. The larger faces at the bottom of the vertical screen are two users who have "stepped up to the podium" and currently have control of what it being seen on the table. The figure shows the user interacting with the table through the use of two magnetic position trackers. The first is attached to a pair of stereoscopic glasses, and as the user moves his head and walks around the table the computer determines his eyepoint location from the tracker and recomputes his view accordingly. The second tracker is attached to a glove that serves as an input device. The user's gloved hand becomes a cursor and he can use his fingers to touch an object to indicate selection or grab and move an object to indicate an action.

Several different kinds of information can be displayed on the table. Figure 11 illustrates a display of current air and ground information. There are several aircraft depicted as realistic models, with the relative scale of the models representing the relative sizes of the respective aircraft. They move in real-time, with the stereoscopic display making them appear to be flying above the table. Conceptually, the positions of the aircraft are provided in real-time by a radar system and the user has the option of displaying them as symbols or models. Remote users worldwide have real-time access to the data. The hemisphere in the upper left is a simple, unclassified representation of the threat dome of a Surface to Air Missile (SAM) emplacement. The large arrow is a cursor that is controlled by a remote user who is collaborating over this display. The amorphous blob in the lower left is a depiction of a small storm cell that is also moving through the region. This weather data is visually integrated in real-time with the current air picture data. The aircraft position, weather, and threat information are all provided by different sensor systems. However, they share a common spatiotemporal reference that allows them to be fused in this real-time synthetic view of the world. Every object in this synthetic view also serves as a visual index into the underlying global C2 database. Selecting an aircraft would let us determine its current status (airborne with a certain speed and heading) and plans (origin, destination, and mission), as well as associated information such as logistics at its base of origin.

Figure 11. Synthetic View of the World

Our current research is focused on the use of aggregation and deaggregation of data within visual depictions, in order to support a wide range of users. A weaponeer wants to study the details of a target (e.g., construction material, distance below ground) that is only a few hundred feet by a few hundred feet in size. A commander wants an overview of all airborne assets, targets, etc. for a region that is several hundred by several hundred miles in size. However, those examining an overview will frequently wish to "drill down" for maximum detail in certain areas, while those examining a detailed area may wish to examine a more global view to retain context. Allowing the visualization of data with this wide range of geographic scopes, as well as iterative travel between detail and overview, poses challenges in both data depiction, data simplification, and intuitive navigation techniques.

8. Conclusion and Research Areas

The above varied and rich application spaces – e.g., visualizing search results, topics, relations and events in news broadcasts, battlefield activities - provide a number of challenges for visualization research. Fundamental issues include:

1. What are effective *information encoding/ visualization techniques* for static and dynamic information visualization, including complex semantic objects such as properties, relations, and events?
2. What are the most effective methods for utilizing *geospatial, temporal, and other contexts* in synthetic displays of real world events that facilitate interface tasks (e.g., location, navigation), comprehension, analysis and inference?
3. What kinds of *interactive devices* (e.g., visual and spatial query) are most effective for which kinds of tasks (e.g., anomaly detection, trend analysis, comparative analysis).
4. What new *evaluation methods, metrics, and measures* are necessary for these new visualization methods?

In visualization, we tend to deal with complexity through methodologies involving abstraction, aggregation, filtering, and focusing. Insights from natural language processing promise to help extract semantic information from text channels, to provide a richer, task-relevant characterization of the information space. Visualization can certainly benefit from other aspects natural language processing in achieving economy of interaction such as notions of context in reference (e.g., "fast_forward <the next week>") or relation (e.g., move "<enemy_icon> behind <Bunker Hill_icon>" in the currently focused display). An investigation of many applications, tasks, and interaction methods will be required to make progress in better understanding and answering these and other fundamental questions.

References

MacEachren, A. M. Department of Geography, Pennsylvania State University, USA. Chair, ICA Commission on Visualization to be published in the Proceedings of the Polish Spatial Information Association Conference, May, 1998, Warsaw Poland (http://www.geog.psu.edu/ica/icavis/poland1.html)

Buttenfield, B. and McMaster, R. (1991). *Map Generalization: Making Rules for Knowledge Representation*, Longman Scientific Technical, England.

Exploratory Cartography: Maps as tools for discovery by Menno-Jan Kraat http://www.itc.nl/~carto/kraak/

Gershon, N., LeVasseur, Winstead, J., Croall, J., Pernick, A., and Rue, W. (1995). "Visualizing Internet Resources," In Gershon, N. & Eick, S.G. (eds), Proceedings for Information Visualization T95 Symposium , (pp. 122-128) IEEE Computer Society Press.

Gershon, N. (1996). "Moving Happily through the World Wide Web." *IEEE Computer Graphics and Applications*, March 1966 (pp. 72-75).

Gershon, N. and Eick, S. G. (eds.) 1997. Introduction: Information Visualization. *IEEE Computer Graphics and Applications*. 17 (4). July/August.

MacEachren, A. M. and Ganter, J. H. (1990). A pattern identification approach to cartographic visualization. *Cartographica*, 27(2): 64-81

Mani, I. and Bloedorn, E. (1997). Summarizing Similarities and Differences among Related Documents. In Proceedings of the Fourteenth National Conference on Artificial Intelligence, (pp. 622-628). AAAI Press, Menlo Park, CA.

Merlino, A., Morey, D. and Maybury, M. (1997). "Broadcast News Navigation using Story Segments", ACM International Multimedia Conference, Seattle, WA, November 8-14, (pp. 381-391). (http://www.mitre.org/resources/centers/advanced_info/g04f/bnn/mmhomeext.html)

Merlino, A. and Maybury, M. to appear. An Empirical Study of the Optimal Presentation of Multimedia Summaries of Broadcast News. In Mani, I. and Maybury, M. (eds.) *Automated Text Summarization*.

Mitchell, R., Day, D., and Hirschman, L. (1995). "Fishing for Information on the Internet." In Gershon, N. & Eick, S.G. (eds), Proceedings for Information Visualization T95 Symposium , (pp. 105-111) IEEE Computer Society Press. October.

Morris, H. and Sonnen, D. (1997). Spatial Information in Data Warehouses: Business Potential and Risks, Bulletin #14753, International data Corporation, November 1997. http://www.itresearch.com/

Shneiderman, B., (1994) Dynamic queries for visual information seeking, *IEEE Software* 11(6): 70-77.

Smotroff, I., Hirschman, L., and Bayer, S. (1995). "Integrating Natural Language with Large Dataspace Visualization," In Adam, N. and Bhargava, B. (eds), *Advances in Digital Libraries, Lecture Notes in Computer Science*, (pp. 209-224) Springer Verlag.

Taylor, D. R. F., Ed. (1991). *Geographic information systems: The microcomputer and modern cartography*. Oxford, UK: Pergamon.