Approved for Public Release; Distribution Unlimited
Case #05-0140

Virtually Integrated
Visionary Intelligence Demonstration
(VIVID)

Mark Maybury, John Griffith, Roderick Holland, Laurie Damianos, Qian Hu, Randall Fish
Information Technology Division
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730, USA
(781) 271-7230
{maybury, rholland, griffith, laurie, qian, fishr}@mitre.org
http://itc.mitre.org

Abstract

Sophisticated information analysts are beginning to wield powerful tools to more effectively exploit massive, heterogeneous, and multisource global information. In this paper we present a set of innovative analytic capabilities within the context of an integrated scenario. Demonstrating a virtually integrated visionary intelligence demonstration (VIVID), we show how an analyst can quickly access thousands of documents and facts to make more effective decisions through the use of advanced capabilities such as event information extraction, geospatial and temporal visualization, audio hot spotting, and multilingual instant messaging.

Keywords: analysis, geospatial information processing, information extraction, question answering, geospatial and temporal visualization, audio hot spotting

1. Information Analysis Scenario

The sophisticated information analyst exploits tens to hundreds of sources daily, reviewing hundreds of documents, and extracting facts from thousands of statements, to establish relationships, discover evidence for or refute hypothesis, and/or support arguments in a highly iterative fashion. To support this analysis, we have created an integrated suite of information analysis tools (see Figure 1) that support multilingual information exploitation. As Figure 1 illustrates, our approach includes systems for intelligent dissemination/retrieval, extraction, summarization, translation, clustering, visualization, information use monitoring, and collaboration. To illustrate the application of specific implemented systems (indicated parenthetically in Figure 1), we exercise these tools in an integrated scenario.

Consider the fictitious scenario where the Columbian government has requested US assistance in protecting air space around several civilian and military airfields because of recent anti-government activities. To support operations, the Columbians suggest the use of commercial airports at Cali or Bogota as well as the military airfields at Tres esquinas, Apiay, and Palanquero. The job of the analyst is to determine which airfields to use.

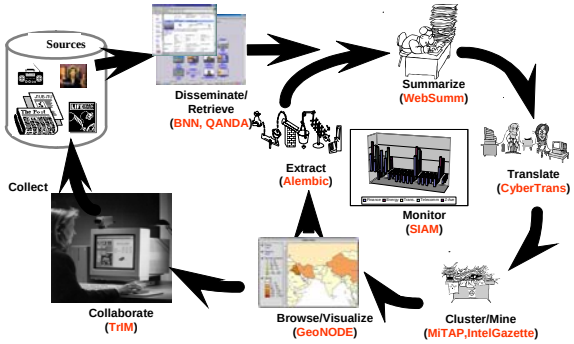

Figure 1. Analytic tool suite

2. Translingual Information Access

The analyst's first task is to determine threats at each of the suggested air fields. In addition to classical sources and methods, in this scenario the analyst exploits GeoNODE Americas to extract information from in country sources (Hyland et al. 1999 ; Vilain et al. 2000; Clifton et al. 2001). GeoNODE Americas has indexed a full year of leading Columbian newspapers from the web (e.g., El Tiempo, El Espectador, and Semana.com). As illustrated in Figure 2 (follow the arrows through each step in the analytic process), GeoNODE Americas enables the user to translate the English query +"Columbian Air Force" +("Tres esquinas" Apiay Palanquero) into its Spanish equivalent +"Fuerzas aereo Columbiana" +("Tres esquinas" Apiay Palanquero). Using this translation to search, GeoNODE finds relevant documents. The user selects from this list and invokes CyberTrans, a machine translation environment, to provide a semi- real-time, gist-quality translation as show at the right in Figure 2 with search terms highlighted. By retaining the structure of the original web page the results translated by CyberTrans are more understandable. This article describes how Columbian Air Force planes have been shot down at some of the air bases by ground fire. Related articles discuss how Palanquero was recently investigated by Department of Justice so may not be a good choice.

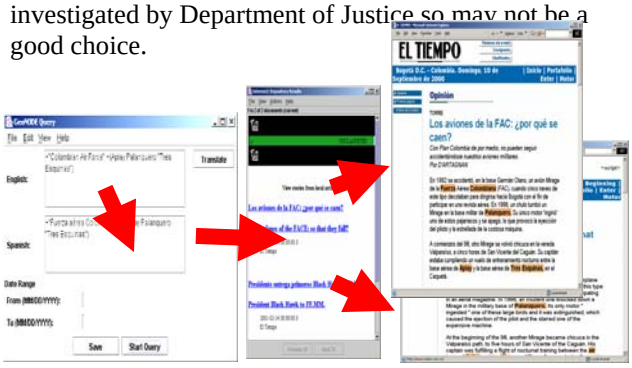

Figure 2. GeoNODE Americas: translingual information access

3. Event Extraction and Discovery: Airport Change

The analyst at this point decides she needs to take a look at the transportation status of Columbian Airfields. She takes a look at the German produced Momberger Report, a free text report which describes airfield changes world wide. Instead of searching the Momberger Report directly, however, she exploits GeoNODE's event extraction capability as shown in Figure 3. The first window to the left displays air field events and properties such as the country in which the change has occurred, type of modification, cost, and name of airport, all of which GeoNODE has extracted from the text report. Values are displayed in a kind of intelligent spreadsheet in which selecting cells (e.g., a country name) will select and display values in other columns that have occurred in events associated with that cell value (e.g., that country).

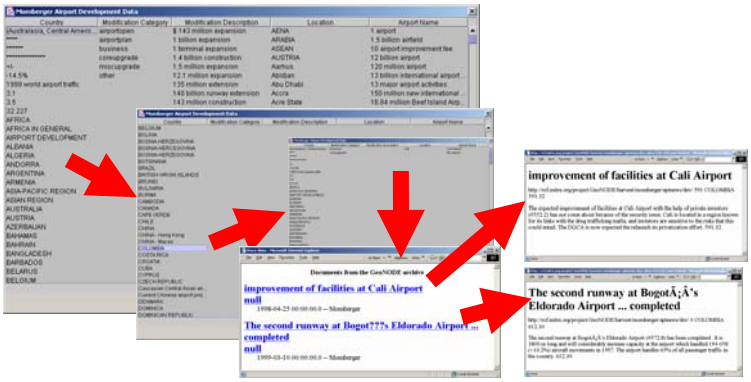

Figure 3. GeoNODE airfield modification tip-off prototype

Leftmost in Figure 3, the analyst first selects Columbia from the country column and all the values in the remaining rows are automatically constrained to those airport change events which occurred in Columbia. Two cities are displayed (Cali and Bogota) and when the articles associated with their changes are displayed the analyst learns that planned facility improvements have not occurred because: "Cali is located in a region known for its links with the drug trafficking mafia, and investors are sensitive to the risks that this could entail." In contrast, the other article states that a second runway at Bogota's El Dorado airport, which handles 65% of the country's passenger traffic, has been completed at 3800 m long considerably increasing airport capacity. The analyst was able to very quickly navigate to this critical information because airport change events were extracted by GeoNODE and presented in an intuitive, spreadsheet like form, linked back to the original event reports.

4. Event Extraction and Discovery: Violence

In addition to airfields for transportation, the analyst may be interested in violent events related to various groups in Columbia that could pose a risk. As illustrated in Figure 4, GeoNODE extracts so called mayhem events from open source newspaper reports. As with the airfield events above, the analyst first sees a display of the types of events (e.g., attack, bombing, murder), and names of locations, victims, human perpetrators, organizational perpetrators, and time. Thus when the analyst selects

AUC (Autodefensas Unidas de Columbia) under organizational perpetrators, only attacks are found. Expanding the selection to include ELN (Ejercito Liberacion National) and FARC (Fuerzas Armadas Revolucionarias de Columbia) adds murder and abduction to the event list. The analyst displays the list of articles containing these organizations and is able to quickly drill down to an article describing a compilation of recent guerilla and paramilitary activities (far bottom right in Figure 4).

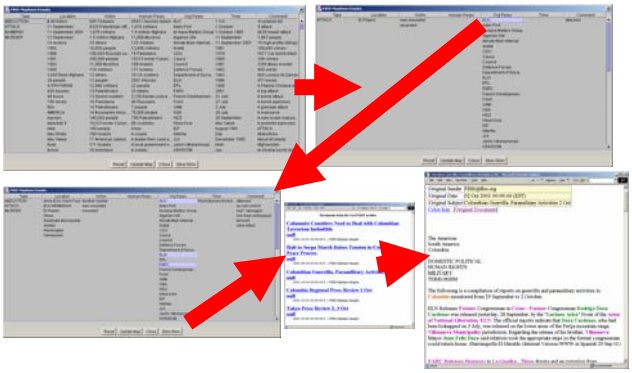

Figure 4. GeoNODE force protection prototype

5. Audio Hot Spotting: Tracking Terrorist Groups from Multimedia Source

Our analyst may also want to track terrorist group activities reported by a broadcast news source. As illustrated in Figure 5, the Audio Hot Spotting system allows the analyst to query the multimedia source either in English for *Basque Separatist Group* or in Spanish for *ETA.* The system retrieves all multi-media segments matching the query and presents both the speech transcripts from automatic speech recognition and the exact corresponding locations of these hits in the original broadcast news file. The analyst can either read the excerpts of the transcripts containing the query or listen to the segments of interest without wasting time listening to the entire broadcast. Since the speech transcript is likely to contain speech recognition errors, it is important that the analyst can quickly access the original audio segment and not rely wholly on the ASR output.

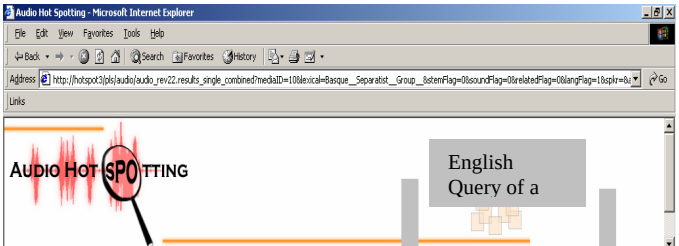

Figure 5. Audio Hot Spotting

When keyword queries do not yield the desired result, the analyst can also query by speaker identities in order to listen to the content spoken by a particular speaker. If the keyword queries return too many hits, speaker identities can also be used to refine the query. The analyst can retrieve all mentions of the *Basque Separatist Group* spoken by a particular speaker. Finally, the analyst can also search on and retrieve passages based upon audio cues in the background noise. In the current system, the analyst can retrieve laughter, applause, and speech rate. The system architecture is such that, depending on domain application requirements, other kinds of background noise can easily be added.

6. Event Visualization

As shown on the left in Figure 6, the GeoNODE geospatial visualization component, based on ESRI's ArcView, allows the user to display relevant documents as geospatial tags on a map. The user can drill-down to the articles that mention a particular place by selecting that place on the map. For GeoNODE Americas, there is also a web-based map display built on ESRI's ArcIMS (Internet Mapping Service). With this interface, as shown on the right in figure 6, the user can perform a geospatial search by a drawing a box, such as around Caqueta, the region containing Tres Esquinas. The GeoNODE geospatial displays allow analysts to assess the spatial signatures of queries and events of interest.

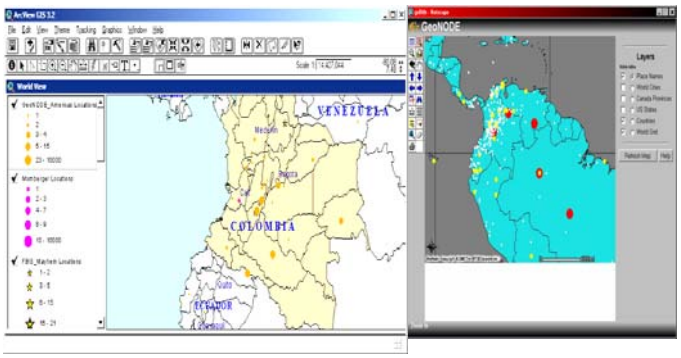

Figure 6. GeoNODE geospatial search and visualization

6. Event Monitoring: Global Infectious Diseases

Our analyst may also be concerned with health of personnel which may mean not only monitoring for violent events but also global infectious diseases. Accordingly, she invokes MITRE's MITAP system (Damianos et al. 2002) which monitors over 90 global sources. The analyst conveniently uses a conventional newsreader, as shown in Figure 7, wherein each news group contains articles about sources, countries, diseases, victims and so on. These are automatically populated and cross-posted based on disease event extraction. In our scenario, the analyst has come across the hemorrhagic disease Dengue Fever in the context of discussing force protection. By simply selecting the news group "Disease.Dengue" from the list on the left hand side of Figure 7, the relevant articles are listed in the top window pane. The analyst rapidly can browse the information color coding disease/symptoms, victims, locations, people, and organizations. One article catches the analyst's eye that mentions the FARC and after selecting and scrolling down he learns that the Andean village of Hobo has been terrorized by guerillas and is unable to obtain medicine for an epidemic of Dengue Fever.

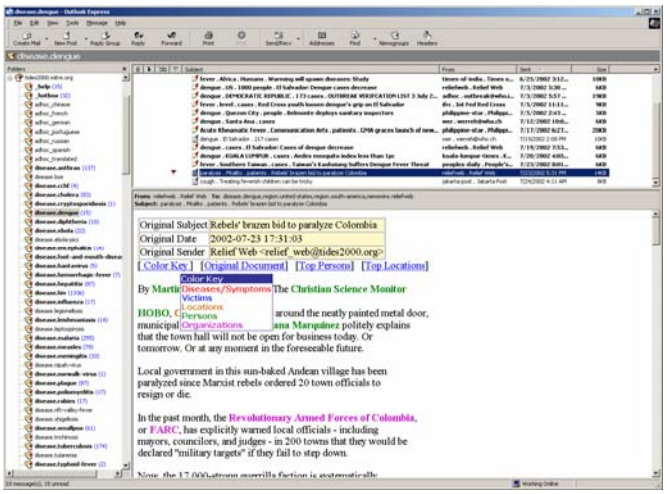

Figure 7. MiTAP

7. Question Answering

If the analyst has further questions about specific information displayed, a question answering system can provide answers (Maybury 2004). For example, in Figure 8 the analyst types a simple natural language query "What is Dengue Fever?" to Language Computer Corporation's question answering system and receives a direct response defining the disease. Question answering tools such as this can extract correct answers for factoid questions such as this example about three quarters of the time.

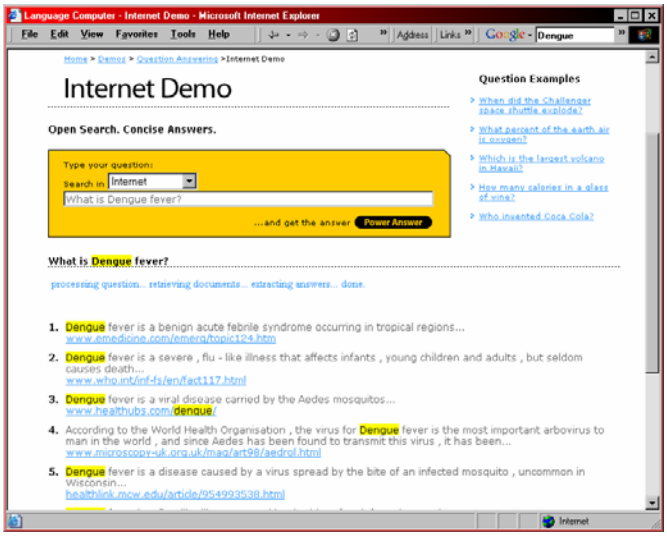

Figure 8. Question answering

8. Expert and Community Location

If an analyst wants access to experts in infectious disease they can query MITRE's Expert Locator (Maybury et al. 2001) software to find experts based on data mining of published documents, project activity, and public email lists. The analyst can also search an automatically generated social network of experts. In this case they search for groups of experts on "biological threats in Columbia" and discover the infectious disease analyst cell, displayed as a network, color-coded to indicate relevance. Further navigation reveals organizational relations and activities as shown in Figure 9.

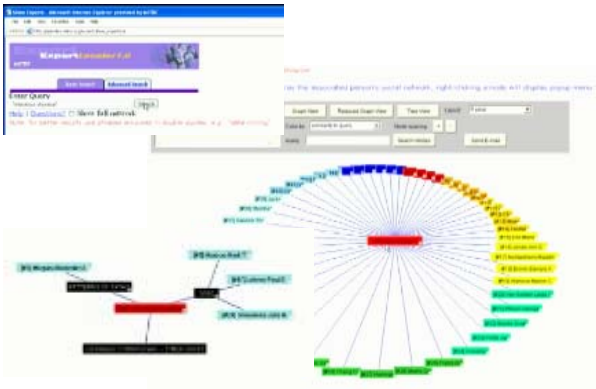

Figure 9. Expert Locator and ExperNet

9. Translingual Instant Messaging

Finally, the analyst decides to engage with some local in country experts to discuss the safety of the various airfields. While an English speaker, the analyst needs to engage with a Spanish speaker. She invokes an instant messaging system using MITRE's Translingual Instant Messaging (TrIM) system, which supports conversations among English, French, German, Italian, Korean, Japanese, and Thai speakers. In Figure 9, two users type as if in a conventional IM system, except that their input is typed in their native languages (left hand side of Figure 9) but sent to CyberTrans and translated to the target language of the other speaker (shown on right hand side of Figure 10). In this exchange, the analyst rapidly confirms concerns about Apiay as an airport and learns that Tres Esquinas has had minimal activity.

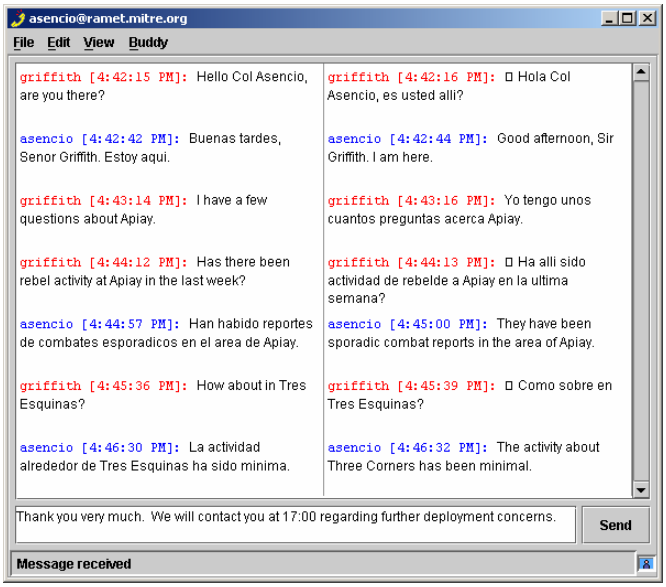

Figure 10. Translingual Instant Messaging (TrIM)

10. Summary and Future Research

In conclusion, we have illustrated how an analyst is able to rapidly formulate knowledge from a diverse array of multilingual sources using a suite of tools. We demonstrated novel capabilities in translingual retrieval, multilingual entity extraction, event extraction, machine translation, question answering, summarization, geospatial query and visualization, and multilingual instant messaging. All of these represent implemented prototypes operating on real data. Table 1 summarizes the technical performance of the modules described in this integrated scenario. For example, these systems are able to extract named entities in news with over 90% accuracy and relations among entities at 70-80% accuracy. The analyst can summarize documents to 20% of their source size without information loss, saving themselves 50% of task time. An analyst can retrieve answers to simple factual questions from relevant documents at 75% accuracy. Finally, systems can return documents across languages relevant to a particular subject with around 80% precision but low recall.

Table 1. Summary of advanced analytic tool performance

Capability	Performance	Example System/ Program
Cross-language Query	Return documents relevant to a particular subject with around 80% precision but low recall. Automated relevance feedback near human	GeoNODE Americas, Audio Hot Spotting

	performance.	
Event Extraction	~90+% precision and recall for entities (e.g., people, locations, organizations), ~70+% for event extraction, ~60% for multilingual event extraction	GeoNODE Force Protection, GeoNODE Airfields, MiTAP
Machine Translation	Gist quality translation (sufficient for relevancy assessment)	CyberTrans
Question Answering	~75% performance in fact based question answering (best system)	Qanda, LCC
Summarization	90%+ precision and recall for stories, 3 time faster retrieval from unstructured digital video	BNN
Multilingual Instant Messaging	Real time gist quality multilingual chat.	TrIM

Many outstanding research problems must be solved to further the analytic environment of the future. Important issues include:

1. *Instrumentation* of applications to log (e.g., Linton et al 1999) and support the inference of models of analyst and/or community interest (via group profiles).
2. *Tailoring.* Exploiting analyst models captured implicitly or explicitly to select content, order material to individuals preferences and needs.
3. *Extraction of Entities, Relations and Events.* Understanding sources to support extraction of named entities, key frames, or key sentences as well as relations and events.
4. *Multilinguality*. The ability to retrieve, extract, summarize, and translate foreign sources.
5. *Multimodality*. The ability to process and access sources including text, audio, and video.
6. *Evaluation*. Task oriented evaluations will be essential for progress that measure enhanced analyst productivity and effectiveness.

10. Acknowledgement

This visionary demonstration represents the integrated works of a broad range of individuals including but not limited to Rod Holland, John Griffith, John Aberdeen, John Burger, Laurie E. Damianos, Steven Wohlever, Marc Ubaldino, Michael (Andy) Chisholm, George Wilson, Janet Hitzeman, David Day, Lynette Hirschman, and Marc Vilain.

11. References

Aberdeen, J., Burger, J., Day, D., Hirschman, L., Robinson, P., and Vilain, M. 1995. Description of the Alembic System sused for MUC-6. In Proceedings of the Sixth Message Understanding Conference, 141-155. ARPA/ITO, 6-8. Columbia, MD.

Clifton, C., Griffith, J. and Holland, R. "GeoNODE: An End-to-End System", in *Proceedings of The 17th International Conference on Data Engineering*, IEEE, April 2-6, 2001, Heidelberg.

Damianos, L., Ponte, J., Wohlever, S., Reeder, F., Day, D., Wilson, G., Hirschman, L. 2002. MiTAP, Text and Audio Processing for Bio-Security: A Case Study. In Proceedings of IAAI-2002: The Fourteenth Innovative Applications of Artificial Intelligence Conference, Edmonton, Alberta, Canada, July 28 - August 1, 2002, pages 807-814.

Hyland, R., Clifton, C. and Holland, R."GeoNODE: Visualizing News in Geospatial Context", in *Proceedings of the Federal Data Mining Symposium*, AFCEA, March 9-10, 1999, Tysons Corner, VA

Light, M. and Maybury, M. 2002. Personalized Multimedia Information Access: Ask Questions, Get Personalized Answers. *Communications of the ACM* 45(5): 54-59. (www.acm.org/cacm/0502/0502toc.html).

Linton, F., Joy, D., & Schaefer, H-P. 1999. Building User and Expert Models by Long-Term Observation of Application Usage. In J. Kay (Ed.), UM99: User Modeling: Proceedings of the Seventh International Conference (pp. 129-138). New York: Springer Verlag. [Selected data are accessible from an archive on http://zeus.gmd.de/ml4um/]

Maybury, M., D'Amore, R, and House, D. December 2001. Expert Finding for Collaborative Virtual Environments. *Communications of the ACM* 14(12): 55-56. (www.acm.org/cacm/1201/1201toc.html).

Maybury, M. Feb. 2000. News on demand: Introduction. *Communications of the ACM*. Vol 43(2): 32-34.

Maybury, M. 2003. Broadcast News Understanding and Navigation. Innovative Applications of Artificial Intelligence. Acapulco, Mexico. 12-14 August, 2003.

Maybury, M. T. editor. 2004. *New Directions in Question Answering*. AAAI/MIT Press

Merlino, A. and Maybury, M. 1999. An Empirical Study of the Optimal Presentation of Multimedia Summaries of Broadcast News. Mani, I. and Maybury, M. (eds.) *Automated Text Summarization*, MIT Press.

Morey, D.; Maybury, M. and Thuraisingham, B. editors, Fall 2000. *Advances in Knowledge Management: Classic and Contemporary Works*. Cambridge: MIT Press.

Vilain, M. Hyland, R., and Holland, R. "Exploiting Semantic Extraction for Spatiotemporal Indexing in GeoNODE", in *Proceedings of Recherche d'Informations assistée par Ordinateur 2000*, RIAO, April 2-6, 2000, Paris.