

MiTAP for Real Users, Real Data, Real Problems

Laurie Damianos, Steve Wohlever, Robyn Kozeriok, Jay Ponte

The MITRE Corporation

202 Burlington Road

Bedford, MA 01730 USA

{laurie, wohlever, robyn, ponte}@mitre.org

ABSTRACT

The MiTAP system was developed as an experimental prototype using human language technologies for monitoring disease outbreaks. The system provides timely, multi-lingual, global information access to analysts, medical experts and individuals involved in humanitarian assistance. Thousands of articles from electronic information sources spanning multiple languages are automatically captured, translated, tagged, summarized, and presented to users in a variety of ways. Real users access MiTAP daily to solve real problems. The successful adoption of MiTAP is attributed to its user-focused design that accommodates the imperfect component technologies and allows users to interact with the system in familiar ways. We will discuss the problem, design process, and implementation from the perspective of services provided and how these services support system capabilities that satisfy user requirements.

Keywords

Information detection, summarization, extraction

INTRODUCTION

Appropriate response to disease outbreaks and emerging threats depends on obtaining reliable and up-to-date information, which often means monitoring vast news sources in many languages worldwide - a task analysts cannot feasibly do. An effective solution requires automated support for global tracking of emerging threats. Analysts report that previous attempts at developing such tools have met with failure and frustration because the tools have often been designed with little consideration to the end user and poorly integrated into the work environment. As part of a research experiment on biological threats, we were tasked to integrate human language technologies. Our goal was to make a system that would be *useful* to analysts and also *used* by analysts. The resulting MiTAP (2001) system collects, annotates and categorizes documents from multiple open news sources. By automating these processes and by providing multiple views into the data, MiTAP enables analysts to spend less time on gathering and digesting raw information and more time on analysis tasks.

Acceptability through Accessibility

If users cannot access a system or cannot understand how to

use it the first time, it is unlikely they will try it a second time. For that reason, we chose familiar, intuitive, and reliable interfaces. Users can access data in MiTAP through a mail/news reader or via a web-based search engine. There are advantages to providing access through standard tools: there is no need to install custom software, the instant sense of familiarity with the interface is crucial in gaining user acceptance as little to no training is required, and data from the browsers are easily imported into other tools.

Information at a Glance

Given time constraints, analysts are not always able to read entire documents to determine their relevance or find important facts. As a way of providing overviews of document contents, MiTAP generates article summaries and draws attention to important text within each document. Pop-ups show lists of named entities (i.e., locations and people) found in the document (Vilain 1999). These summaries provide enough information to indicate relevance to the analyst and act as shortcuts to fact retrieval. Key words in articles are color-coded in the text so users can quickly scan for relevant information. Entities such as people, organizations, locations, diseases, and victim information, are highlighted (Aberdeen et al. 1996). Despite the errorful tagging, this feature helps the user quickly focus on key information in lengthy documents.

Information from Data

Thousands of articles are captured, processed, and posted to the system daily - too much to manage and read. To help organize, summarize, and navigate the data, the MiTAP news server hosts several hundred newsgroups organized by category (i.e., source, disease, region, person, and organization) to allow analysts, with specific information needs, to locate material quickly. Articles are cross-posted to various newsgroups, based on information extracted from their contents. Supplementing access to the data, articles are indexed by an information retrieval system, allowing full text, source-specific queries over the entire archive.

Knowledge from Information

Various MiTAP newsgroups have been created to provide high-level views of the news. The multi-document summarization feature (Columbia Newsblaster 2002) automatically tracks events and produces daily summaries or high-level views of the underlying data. Another summary is in the form of daily Top 10 watch lists of diseases or specific people in the news (Alias I, Inc. 2002). These views provide indications of top stories, extracts of relevant documents, and tables of associated entities.

Approved for Public Release; Distribution Unlimited ISO 03-0338.
©2003-The MITRE Corporation. All rights reserved.

CHI 2003, APRIL 5-10, 2003, FT. LAUDERDALE, FLORIDA.
ACM 1-58113-630-7/03/0004.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE MiTPA for Real Users, Real Data, Real Problems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MITRE Corporation, 202 Burlington Road, Bedford, MA, 01730-1420				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

User-Centered, Task-Driven Design Approach

The focus of integrating components into MiTAP was on providing value to end users. We wanted to create a tool that would require no installation, little to no training, and fit naturally into the analysts' workflow. We spent weeks in dialogue with analysts, understanding their requirements, and studying their work practices. This direct contact allowed us to understand the critical elements of the task, i.e., timely access to data, information extraction from that data, knowledge representation of that information, and seamless integration into the work environment.

During development, ongoing real exercises and evaluation proved to be invaluable for measuring utility, usability, and progress. Over a ten-day period, we used the system to gather information on potential biological threats and to monitor international coverage of specific events. We were able to produce relevant information that analysts, using other means, were not able to find. The exercise also helped improve the performance and robustness of the system; we were able to improve the data throughput from 4K to 8K articles a day. We learned that we needed to enhance search capabilities, integrate better tools for summarization and data visualization, and provide customizable mechanisms to allow users to track information.

The *Disease of the Month Experiment* was a series of minievaluations designed to measure progress on a monthly basis. We chose a scenario familiar to analysts (i.e., research a current disease outbreak and prepare a report) to help minimize dependent variables and reduce training. Test groups were compared monthly to control groups in order to measure system utility. Comparing MiTAP to the web and its vast amount of information, we hypothesized that 1) MiTAP users can produce better analytic reports in a shorter amount of time, where "better" means more up-to-date and more complete, and 2) MiTAP users spend less time reading documents and can digest more in a given period of time. Test groups were also compared across iterations to measure the progress of development.

Simultaneously, we performed independent usability studies. For purposes of contrasting and comparing test vs. control and test vs. test across months, we defined five categories of metrics: efficiency, task success, data quality, user satisfaction, and usability. In our experiments, MiTAP users provided more detail and more up-to-date information on disease outbreaks than just the web alone; however, they did not necessarily spend less time doing so. Our results also show that the test groups were able to find a larger number of relevant articles in fewer searches. In fact, the test groups, who were also permitted to use the web to find information, cited MiTAP articles in their reports an average of three times more than articles found on the web, and often the links to the relevant web information were found via MiTAP. During this experiment, feedback has guided development, provided a comprehensive

understanding of what real users do and how we can help them, and improved overall system performance (e.g., throughput increased by a factor of 2.5 while source integration time decreased by a factor of 4). As a result of improved performance, we were able to add many new sources, producing a significantly richer, broader, and larger data collection.

In addition to exercises and experiments, we used focus groups to help design analytical tools from combinations of integrated technologies. User surveys of early versions of the integrated tools, as well as unprovoked feedback, helped in ongoing improvements. We also examined the logs for usage patterns to help us understand how the system components were being used.

Real Users, Real Problems

MiTAP was originally designed for a group of medical analysts interested in monitoring infectious disease outbreaks. However, the dynamic and flexible nature of the system has allowed it to become larger in scope, encouraging a broad user base with a variety of interests. Currently, over 500 users have accounts, and the domain has expanded beyond diseases to include weapons of mass destruction, terrorism, and warfare.

Conclusion

MiTAP has demonstrated the utility of integrating multiple research technologies to address the requirements of a community of users. The end result is a set of in-demand tools that provide the capabilities to support individuals and organizations who must manage the information overload resulting from the need to keep current on numerous worldwide, multilingual events. We have learned that building a system is more than just integrating components. Focus on both user and task is critical, and ongoing evaluation aids in iterative development.

ACKNOWLEDGMENTS

This work was supported, in part, under DARPA contract number DAAB07-01-C-C201.

REFERENCES

1. Aberdeen, J., Burger, J., Day, D., Hirschman, L., Palmer, D., Robinson, P., Vilain, M. 1996. MITRE: Description of the Alembic System as Used in MET. In *Proceedings of the TIPSTER 24-Month Workshop*, 5 May, Tysons Corner, VA.
2. Alias I, Inc, 2002. New Technology for Information Access. <http://www.alias-i.com>
3. Columbia Newsblaster, 2002. <http://www.cs.columbia.edu/nlp/newsblaster>
4. MiTAP System 2001, <http://tides2000.mitre.org/>.
5. Vilain, M. 1999. Inferential Information Extraction in Pazienza, M., ed., *Information Extraction*, Lecture notes of the 1999 SCIE Summer School on Information Extraction, ed. M. Pazienza, 95-119. Springer Verlag.