

Beyond Broadcast

Kevin Livingston, Mark Dredze, Kristian Hammond, and Larry Birnbaum
Intelligent Information Laboratory, Northwestern University
1890 Maple Ave, Evanston IL 60201 USA
+1 847 491 3500

{livingston, m-dredze}@northwestern.edu
{hammond, birnbaum}@infolab.northwestern.edu

ABSTRACT

The work presented in this paper takes a novel approach to the task of providing information to viewers of broadcast news. Instead of considering the broadcast news as the end product, this work uses it as a starting point to dynamically build an information space for the user to explore. This information space is designed to satisfy the user's information needs, by containing more breadth, depth, and points of view than the original broadcast story. The architecture and current implementation are discussed, and preliminary results from the analysis of some its components are presented.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval. H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing.

General Terms

Design, Human Factors.

Keywords

Information access, information retrieval, user information needs, broadcast news, proactive system

1. INTRODUCTION

The broadcast model of television news is insufficient for conveying complete and in-depth information to all viewers due to time and attention constraints. This research discusses a method for solving this problem by delivering just-in-time information to the viewer to supplement television broadcasts. We present an implementation of this method, Cronkite (see Figure 1).

A novel aspect of this research is that it uses broadcast news as a starting point for gathering information regarding specific stories, as opposed to considering the broadcast version to be the terminus of the viewer's exploration. Cronkite provides a gateway from broadcast news into expanded and useful information [2] gathered from both general and domain specific information repositories. Cronkite expands both the breadth and depth of information available to the viewers, and allows them to browse through this expanded information space as they see fit. Broadcast news is a good starting point for this exploration since it is fairly well structured, and viewer information needs are easy to identify.

Imagine a viewer is watching a news story about the latest negotiations in the Middle East peace process. He might want to learn about the Palestinian and Israeli reactions to these developments, so while viewing the story he presses the 'interest' button on his remote control. Within seconds, Cronkite provides him with expanded information, on a nearby computer display, related to the story he was watching. The system presents the viewer with maps and a history of the Palestinian-Israeli conflict, as well as information about political leaders, such as Yasser Arafat and Ariel Sharon. Cronkite also provides pointers to related stories from sources such as The New York Times and Google News. If the user wants to learn about the Palestinian and Israeli viewpoints he can navigate to the opinion section, where Cronkite displays summaries and links to stories from the Jerusalem Post and the Al-Ahram. These stories may include coverage of a press conference with the Israeli foreign minister as well as an interview with the Palestinian chief negotiator.



Figure 1. A television viewer receives information on a computer (left) after using the remote control to signal interest in the story currently on the television.

2. ARCHITECTURE

The basis for Cronkite's information retrieval system is a model of broadcast news that contains categories of potential viewer questions. We generated this model by looking at the questions

several viewers had while watching news broadcasts, and then collapsing them into five categories:

- a. Related coverage
- b. Opinions
- c. In-depth analysis of the subject
- d. Explanation of key entities
- e. General overview or background of the subject.

As it processes a story, Cronkite classifies it into a hierarchy of topics, such as general news, medical, or Middle East. Attached to each node of this hierarchy are topic-specific questions that stem from the categories of the broadcast news model described above, such as “who is this politician?”, or “what does this drug do?”. Each question is in the form of parameterized queries associated with different information repositories. Cronkite uses information from its representation of the story to instantiate these queries and then submit them to their respective repositories. The returned information is organized and formatted using templates based on the category of the question that the information answers.

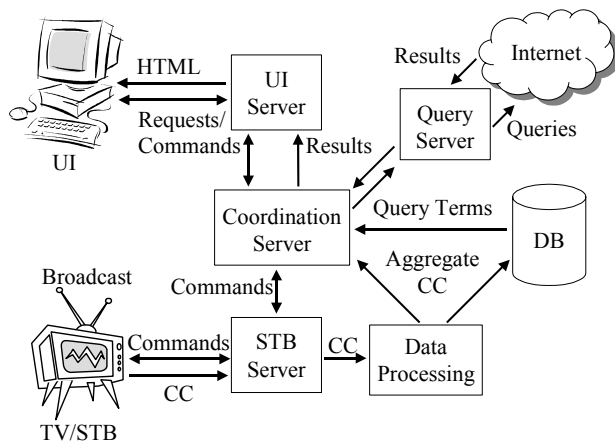


Figure 2: The system is composed of 5 servers, the set top box (STB) itself, and the user interface which is provided through a networked device that can render HTML. Cronkite also has access to offline storage and the Internet.

Cronkite is composed of several autonomous modules (see Figure 2). As a viewer watches the story on the Middle East, the set top box (STB) server streams closed captioning text (CC) to the data processing server. The data processing server contains four primary components: a segmentation engine, a classification engine, an entity detection engine, and a query formation engine.

The segmentation engine determines when a new story has begun by detecting either the CC new topic marker “>>>” or the occurrence of one of several cue phrases. For example, as the anchor says, “and now to the latest developments...”, the segmentation engine triggers a new story. Information about the previous story is archived to storage, and the classification, entity detection, and query formation engines are all reset for the new story.

The classification system searches for words related to several topics, for example, the Middle East, and after identifying sufficient evidence it flags the story as being ‘about the Middle East’. At the same time, the entity detection engine identifies important key players in the story, such as, Israel and the Palestinian Authority. The query formation engine dynamically maintains a term vector describing the story that is used in searches against more general information repositories.

If a user presses the ‘interest’ button on his remote control, the STB server sends a request to the coordination server. The coordination server polls the data processing server for the current story representation, which it uses to identify what information it needs by locating it in the topic hierarchy. Queries are selected, instantiated, and sent to the query server, a wrapper around the InfoLab Information Source Adapter (ISA) system [5]. Given a query, the ISA system can identify databases or web pages that contain information relevant to that query, request that information, parse it, and return an XML-based representation of the specific information requested. The responses from the ISA system are then aggregated and transmitted to the user interface (UI) server. The UI sever selects a template based on the type of information returned, which it uses to format the information as HTML for the user.

The UI is comprised of several templates used to construct different views. Each template specifies not only formatting information, but also what questions the user is currently exploring. The user interface is designed to aid the user in exploring the breadth and depth of the information space surrounding a given story. It allows the user freedom to view the information she is interested in by affording easy navigation of the information space.

3. STATUS

We have produced four topic classifiers: medical, Middle East, weather, and sports. The most developed of these, medical and Middle East, have been tested for precision and recall, and performed well. For both classifiers combined there were 582 stories in the test corpus with 91 being either medical or Middle East stories. The average precision was 96.3% and the recall was 86.8%. Early analysis of the data has shown that classification scores are all very polar compared to the threshold, and that stories can generally be classified after seeing only 30-50% of the closed captioning. These results are very promising, although more rigorous test should be conducted. We also have entity detectors that perform well for businesses, Congressmen, and medical conditions and treatments.

We have tested and demonstrated Cronkite’s functionality using real samples of broadcast news, and initial user reaction has been positive.

4. RELATED WORK

Research in retrieval of information useful to viewers of broadcast news has primarily focused on news summarization, the generation of personalized news broadcasts (personalcasts), and providing news on demand (NOD) [1] [11] [12]. In these systems, the viewer must explicitly convey her context or have a good understanding of what she is looking for, which is difficult since people cannot easily communicate their needs with Information

Retrieval (IR) engines [3]. Additionally, these systems do not fulfill viewer needs since broadcast news is the end product of the interaction and news stories usually do not go into sufficient detail. In contrast, Cronkite provides detailed information from both general and expert quality sources.

Our work is also related to automated question answering systems, which attempt to provide synthesized answers to user queries [9]. Instead of requiring the users to ask explicit questions, Cronkite anticipates them, and provides answers by querying specific information repositories.

5. FUTURE RESEARCH

We are currently designing user studies to assess the quality and usefulness of information retrieved by Cronkite. We are also expanding the system by adding more information templates and sources. Since we want to provide viewers with multiple directions for information exploration but not flood them with options, we are also trying to determine the quantities and extent of information that is appropriate. Segmentation independent of the CC topic change marker must be improved. Several methods, such as discourse cues studied by Hirschberg and Litman [7] or Hearst's Text Tiling algorithm [6] will likely increase accuracy.

More general goals include providing information via different forms of media, as well as integrating the user interface with the television itself. Finally, the long-term focus is to continue to develop an overall theory of viewer interaction with television. For example, sitcom viewers may be interested in fan sites, show listings, plot summaries, and actors. There are many different genres of television to explore, including: do-it-yourself and cooking shows, documentaries, kids' television, sports, and even commercials. Each genre has its own context and will require different models of interaction with the viewer.

6. CONCLUSION

Cronkite demonstrates the ability to segment broadcast news at the story level, identify the content of the story, and use this content to dynamically build a set of queries to generate an information space for the viewer to explore. This space provides both more depth and more breadth than the original broadcast version. This work takes the first steps at providing a theory for understanding viewer interaction and information needs while watching television.

7. ACKNOWLEDGMENTS

We thank Andy Crossen and Jay Budzik for the ISA system, Ayman Shamma for the photograph, and Mike Baker and the other members of the TiVo hacker community for their assistance.

8. REFERENCES

- [1] Brown, M.G., Foote, J.T., Jones, G.J.F., Sparck Jones, K., and Young, S.J. Automatic Content-based Retrieval of Broadcast News. *Proceedings of ACM Multimedia*. (San Francisco, CA, 1995), ACM Press, 35 - 43.
- [2] Budzik, J., Hammond, K., Birnbaum, L., and Krema, M. Beyond similarity. *Working Notes of the AAAI-2000 Workshop on AI for Web Search*. (Menlo Park, CA, 2000), AAAI Press.
- [3] Budzik, J., Hammond K., and Birnbaum, L., Information access in context. *Knowledge based systems 14* (1-2), (2001), Elsevier Science 37-53.
- [4] Budzik, J., Ruberry, M., Stein, H., and Hammond, K. J. Proactively searching the Web: Can we agree on what's relevant?. *Demo and Poster Proceedings of ACM HyperText* (2002).
- [5] Crossen, A., Budzik, J., Warner, M., Birnbaum, L., and Hammond, K. XLibris: An automated library research assistant. *Fifth International Conference on Intelligent User Interfaces*, (2001).
- [6] Hearst, M.A. Multi-paragraph Segmentation of Expository Text. *Proceedings of the ACL*, (1994).
- [7] Hirschberg, J., and Litman, D. Empirical Studies on the Disambiguation of Cue Phrases. *Computational Linguistics*, 19(3) (1993), 501-529.
- [8] Leake, D., Scherle, R., Budzik, J., and Hammond, K. Selecting task-relevant sources for just-in-time retrieval. *Proceedings of the AAAI-99 Workshop on Intelligent Information Systems*, (Menlo Park, CA, 1999) AAAI Press.
- [9] Light, M., Maybury, M. Personalized Multimedia Information Access. *Communications of the ACM*, 45(5) (May 2002), 54-59.
- [10] Mani, I., House, D., Maybury, M., Green, M., Towards Content-Based Browsing of Broadcast News Video. *Intelligent Multimedia Information Retrieval*, (1997), 241-258.
- [11] Maybury, M. (ed.). News on Demand. *Communications of the ACM*, 43(2) (February 2000), 33-34.
- [12] Wilcox, L., and Boreczky, J. S. Annotation and Segmentation for Multimedia Indexing and Retrieval. *HICSS* (1998) 259-266