

Advanced Question and Answering Systems for Community Development

Queen Booker, Scott Farrar, Mick McQuaid, Ana Lopez,
University of Arizona

Center for the Management of Information
Tucson, AZ

qbooker@cmi.arizona.edu; farrar@email.arizona.edu; mmcquaid@bpa.arizona.edu;
alopez@cmi.arizona.edu;

Abstract

The Internet has made transfer of knowledge of successful community development programs and processes possible by allowing practitioners to connect through computer-mediated-communications such as web pages, email, chat rooms and groupware. However, because community development needs are specific to each community, it becomes difficult to design programs to solve the problem despite availability of funds, hence one of the problems of enduring poverty and distress. We propose that one of the ways to alleviate the burden of knowledge transfer among practitioners is to build a large repository of community development practices and practitioner knowledge using video or automated slide shows and an intelligent agent question/answering system to simulate a conversation between a practitioner seeking information and a practitioner sharing information. Such a system could provide the information seeking practitioner with the perception of dealing with one or more knowledgeable experts without having to actually interact, expediting the transfer of knowledge when the practitioner actually needs it.

Introduction

Community development is comprised of programs, processes and policies designed to aid community residents in building and sustaining a strong and healthy environment for living. Characteristics of a healthy community environment include low unemployment, living wages for workers, a strong educational system, homeownership with property values sufficient to support most public infrastructure (e.g., water/sewer systems, roads, schools, other public facilities), clean air, strong local business ownership, strong retail sector, low mortality and morbidity rates, a stratified population, a strong transportation system and strong current and future leadership. In addition to these static indicators, a healthy community also has indications of growth in job creation/transition and

housing development. The community development workforce is comprised of program officers who make funding available for community development activities, policy makers who establish policy that affect and impact those funding decision and sets restrictions or reduce barriers on activities, and community development practitioners who put those programs and policies into action. A community development practitioner is a person who engages in one or more activities to create, maintain or adjust programs and policies to support achieving or maintaining a healthy community. While some practitioners focus on only one of the activities (e.g., housing development, job development, business development), many engage in more than one, oftentimes without receiving the appropriate training to perform any one proficiently. While training is available to practitioners through both federal and state agencies, and through non profit organizations such as the National Congress of Community and Economic Development (NCCED), Neighborhood Reinvestment Institute, as well as a host of colleges and universities that offer community development coursework, one of the challenges is that the training is usually general and not available when the practitioner needs it and is able to employ the techniques. While this may be viewed as a reactionary rather than proactive approach to community development, it may not necessarily be. For example, through an annual analysis a practitioner may recognize a trend in low wage job creation and wish to counteract the trend but is a housing specialist without the knowledge to do so. In most circumstances, the trend goes unchallenged until too late. Also, attending training usually requires the person to travel to some location other than where they work, and the trainee generally finds that the information provided in the training is not relevant to their specific needs. This is particularly true for people in rural communities.

Many funders make technical assistance providers available who offer individualized, hands-on support to

its program participants. Unfortunately, this type of assistance is limited by the knowledge of the provider as well as the cost to maintain. To support training or information sharing, much has been developed for the field. Among these are books, with a popular one being *Building Communities from the Inside Out* by Kretzmann and McKnight; email where practitioners can share information with each other on a one on one basis; chat-rooms; and even on-line discussion groups. Community groups have created web-pages, to the tune of over 2 million as of June 2001 based on a search by google.com.

We propose a knowledge management system (KMS) that is designed to aid the community development practitioner in obtaining knowledge from other community development practitioners via Internet through the use of video taped conversations with an expert describing aspects of a particular community development program or activities of a process, or by locating relevant information in online text. By documenting the knowledge of existing practitioners or creating a mechanism for locating relevant text, we are able to create a growing and diverse repository of practices and processes for solving community problems that could be available to new and enduring practitioners, creating a collaborative system allowing a practitioner to obtain information from his peers without having to disrupt the colleague. Because we are using the Internet as our primary means of knowledge transfer, these colleagues can be and learn anywhere without really ever having to meet the colleague, reducing the need to travel to learn how to accomplish a particular task.

The research questions we are trying to answer are:

- What constitutes an effective Q&A workflow that can provide information to an analyst as he needs it?
- Can a Q&A system improve the quality of services received from the workforce?
- Can current NLP processes be improved enough such that a viable dialogue and intelligent agents can retrieve precise responses to support just-in-time knowledge transfer?
- Does streamed video improve workflow and task completion in the absence of an actual expert?
- Can segmented video and improved web page information retrieval reduce the need for extensive training for community development practitioners?
- Can video segments and web pages be organized in such a manner as to provide

immediate workflow support and strengthen the ability of the practitioner to perform his job more efficiently?

- Can poverty be reduced more readily with such a system
-

While the ultimate success of the system will be based on the quality of the information collected, the architecture that enables the user to obtain the information he requires is equally important. The rest of the paper is organized as follows. Section 2 focuses on system design considerations, Section 3 focuses on the research approach, Section 4 describes content development and finally we summarize with next steps.

System Design Considerations

Our proposed system has as its several subsystems: video processing, web page processing, ontology construction, question understanding and refinement, and answer formulation. The main new technology we offer is the dialogue manager, which combines the question understanding and answer formulation components. Existing work has adequately shown how to effectively develop video processing, web page processing and ontology construction. Thus, this paper focuses on the integration of the dialogue manager and deep content analysis as the critical new components to the system.

The Dialogue Manager

A major obstacle to retrieving information from any knowledge system is the ability to determine exactly what the user wants to know. One way to address this is to implement commonsense reasoning in identifying the correct context or domain theory in which to operate. The approach taken in the present Q&A system assumes that an artificial intelligence complete solution, one in which the problem of artificial intelligence is solved, is not possible at the present stage of research. Instead, certain compromises must be accommodated. One promising area of research that may yield fruitful results is dialogue processing. A dialogue manager would provide a means of interaction between the user and the system to identify the user's intentions in asking a question, thereby narrowing the topic domain and providing a means for question disambiguation.

In the past, question understanding systems have suffered from a lack of knowledge [11]. In this work it is shown that utilizing knowledge and reasoning techniques can greatly enhance the precision of question understanding and the eventual answer. Even lightweight abduction of queries based on [12] can

help. But given recent advances in knowledge base technology, knowledge can be used even more effectively through *deep content analysis*. If precise answers are needed, especially at a high level of complexity, then it is necessary to perform a deep content analysis. Content analysis is too computationally expensive to perform on every document. The rationale is to first narrow the document pool to a few highly relevant documents by using indexing techniques, and only then perform a deeper analysis. Since the system is already designed with knowledge sources and an NLP system, expanding the use of these resources is an attainable task.

Knowledge sources

Two basic knowledge sources are used in the proposed question understanding system. One is *a priori* knowledge, including WordNet and a general knowledge base, modeled after the efforts of CYC. The *a priori* component will be used in NLP, both in processing user input and deep analysis of the data (though this connection is not shown in the diagram). The other knowledge source represents data that has been processed from the input data, including a traditional index and a knowledge base, which is updated as needed by deep content analysis.

NLP

Central to the overall system design is a fully integrated dialogue manager. The dialogue manager tracks user input based on system response. It is connected to the front-end NLP components (parser, Q-analyzer) and the answer generation portion. Since the primary reason for the dialogue is to narrow the context and to refine the user's input, the dialogue manager also communicates with the GUI and the context module.

The component responsible for the query formulation is the question analyzer (Q-analyzer). The question-analyzer determines the question type and focus as well as projected answer type.

Document Processing

Initially, multi-media data (documents, video, audio, etc.) are processed by traditional indexing techniques and key words and major phrases are stored in the document index. If the index does not contain an adequate answer, further analysis on a subset of the data is required. The particular context emerges from the context manager and the relevant subset of data is chosen to be analyzed further. Further analysis consists of deep NLP and conceptual processing with the goal of updating the document knowledge base. The

document knowledge base grows as more user input is collected. So at any given time the knowledge base consists of a portion of the overall data corpus. The knowledge base is subject to general reasoning rules and may be incorporated with the general knowledge base.

NLP Component

Dialogue Processing

The open-domain requirement for a Q&A system poses a major problem to the NLP component. Essentially, treating user input as domain independent requires a solution to the AI problem of language interpretation. As of yet, no robust solution has been found, but any solution to the interpretation of natural language must at least accommodate commonsense reasoning and apply an understanding of the bridge between natural language and symbolic thought. The one promising area of research that may yield fruitful results is *dialogue processing*.

Dialogue processing is a sub-domain of NLP that attempts to model and implement language behavior at the level of *discourse*. Discourse involves the combination of utterances to form coherent speech events, either in the form of whole paragraphs or single sentences. The current work concerns endowing Q&A systems with the discourse of conversation. The study of discourse in modern linguistics and philosophy begins with [4] and [10]. [16,17] extended the early ideas and created taxonomies of everyday *speech acts*. For instance, *requesting*, *thanking*, *declaring*, and *asking* are some common speech acts used in dialogues. As far as AI is concerned, a major step in the ability to process dialogue was the incorporation of planning techniques into language systems [7]. The planning approach remains the most successful method of processing conversation, as exemplified by such systems as TRAINS [1]. The success of the TRAINS system is partly based on a system of dialogue mark-up known as DAMSL [2]. With the envisioned Q&A system a specific dialogue mark-up for Q&A will be necessary. A recent attempt to incorporate dialogue processing and planning is the ARTEMIS system [15]. Furthermore, [13] presents a general applied framework for implementing such a dialogue processing system.

Dialogue interaction in a Q&A system is key, because such behavior produces the illusion of domain dependence. For instance, consider the following dialogue:

() System: How can I help you?

User: How do I create a housing program?

System: Do you want to build new housing or renovate housing?

User: Renovate housing.

At this point, the search for an answer can begin with focus on those documents containing terms relevant to housing renovations. So in a sense only focusing on a certain subset of the document corpus narrows the domain. Relying on a cooperative dialogue model does not eliminate the need for commonsense reasoning altogether but it aids the system by allowing the user to guide it when ambiguities arise. A dialogue processing approach, then, allows the system to deal with non-specific questions in a natural manner. For a detailed discussion of interaction in an information retrieval context, see [8].

Cooperation is key to human-computer interaction in a Q&A system especially in building up the query. Q&A systems should be a joint activity, not one-sided and user-focused. Non-expert users are assumed to be naïve with respect to their ability to pose ideal questions, *ideal* as it relates to the question model being used. For a Q&A system to be precise some mechanism must map user input to a query [5]. At this point the negotiation aspect of dialogue becomes important.

Other than the obvious NLP challenges (prepositional phrase attachment, named entity extraction, lexical ambiguity), this complex question is difficult to answer precisely because of the difficulty in breaking the input down into individual queries. A dialogue approach to question and answer can help.

A cooperative dialogue demonstrates interactive lexical disambiguation and a narrowing of the topic domain. The word ‘programming’ is ambiguous, so the system asks the user to specify. Once the domain is established, the user helps the system identify a particular context prior to being a specific question. Furthermore, narrowing the topic domain may aid in extracting the correct type of media desired, in this case text. So, a precise query is built up through cooperative dialogue, even when a complete parse of the input is not achieved as seen with the clarification in the last interaction pair.

Work in the area of intention, in particular [6] indicates that speakers’ intentions drive dialogue. In a Q&A application, user intention must be recognized by the system and also be mediated by system’s capabilities.

The system uses plan recognition to identify the user’s intention behind the request. Intention recognition forms the backbone of the successful TRAINS system described in [1].

Linguistic Knowledge Bases

Linguistics analysis is impossible without a database of lexical knowledge. The primary lexical database in use in many Q&A systems is WordNet [9]. Some functions of WordNet in a Q&A framework are: lexical disambiguation, parsing, and query expansion. In query expansion key words related to the actual user input are added to the knowledge base query, thereby increasing the coverage of documents returned. WordNet, however, has a few undesirable characteristics. The primary one being its unsystematic organization. Often when WordNet is used in conjunction with a knowledge base, for example CYC [14], the ontologies of the two systems do not mesh. To overcome this difficulty it is proposed to collapse WordNet and a portion of the CYC ontology into an object-oriented, combined version. In this way, greater uniformity will be achieved in keeping with the overall goal of producing a tightly organized system.

Tagging and Parsing the Input

A finite-state parsing scheme similar to that used in the FASTUS system [3] will be employed. The output of such a parser is not the complete parse, but rather a list of the basic phrases, including most importantly the main noun groups and the main verb phrase. A complete parse is not necessary, mainly because of the possibility of correction through a cooperative dialogue, as was illustrated previously.

Mapping the Input onto a Knowledge Base Query

The goal of NLP is to produce a complex query to send to the knowledge base. Crucial to this process is converting the partially parsed input into some kind of semantic representation, referred to as the logical form. A scheme similar to that of [12] is proposed.

References:

- [1]Allen, J. (1995) *Natural Language Understanding*. Menlo Park, CA: Benjamin Cummings.
- [2]Allen, J. and Core, M. (1997) “DAMSL: Dialogue act markup in Several Layers.” *AAAI 97*.
- [3]Appelt, D. and Israel, D. (1997) “ANLP-97 Tutorial: Building Information Extraction Systems.” Available at www.ai.sri.com/~appelt/ie-tutorial
- [4]Austin, J.L. (1962) *How to Do Things with Words* Cambridge, MA: Harvard Univ. Press.

- [5]Belkin, N. J., Oddy, R. N. and Brooks (1982) "ASK for Information Retrieval. Part I: Background and Theory." *Journal of Documentation*, 38, 2, 61-71.
- [6]Cohen, P.R., Morgan, J., and Pollack, M. (1990) *Intentions in Communication*. Cambridge, MA: The MIT Press.
- [7]Cohen, P.R. and Perrault, C.R. (1979). "Elements of a Plan-Based Theory of Speech Acts." *Cognitive Science*, 3(3), pp. 177-212.
- [8]Daniels, P. J., and Brooks, H. M., and Belkin, N. J. (1997) In K. Sparck Jones and P. Willett (Eds.) *Readings in Information Retrieval*. San Francisco: Morgan Kaufmann.
- [9]Felbaum, C. (Ed.) (1998) *WordNet: An Electronic Database*. Cambridge, MA: MIT Press.
- [10]Grice, H. P. (1969) "Utterer's Meaning and Intentions." *Philosophical Review*, 68, pp. 147-177.
- [11]Harabagiu, S.M., Moldovan D., Pasca, M.A., Mihalcea, R., Surdeanu, M., Bunescu, R., Girju R., Rus, V. and Morarescu, P. (2001) FALCON: Boosting Knowledge for Answer Engines, in *TREC-9 Proceedings*.
- [12]Hobbs, J. (1986) "Overview of the TACITUS Project." *Computational Linguistics*, 12, (3).
- [13]Houghton, G. (1997) *The Production of Spoken Discourse*. Greenwich, CT: Ablex Publishing Corp.
- [14]Lenat, D. and Guha, R. (1991) *Building Large Knowledge-Based Systems: Representation and Inference in CYC*. Reading, MA: Addison-Wesley.
 Proceedings of AAAI 2000 workshop on Representational Issues for Real-world Planning Systems (July 2000).
- [15]Sadek, D. and de Mori, R. (1998) "Dialogue Systems" In R. de Mori (ed.) *Spoken Dialogues with Computers* pp. 523-561. London: Academic Press.
- [16]Searle, J.R. (1970) *Speech Acts: An Essay in the Philosophy of Language*. Cambridge England: Cambridge University Press.
- [17]Searle, J.R. (1975) "Indirect speech acts." In P. Cole and J. Morgan (eds.) *Syntax and Semantics*. Vol. 3: *Speech Acts*. New York: Academic Press, 59-82.