

## **Q&A: A system for the Capture, Organization and Reuse of Expertise**

**Jay Budzik, Kristian Hammond**  
Northwestern University, Evanston, Illinois

### **Abstract**

It is a time-consuming and difficult task for an individual, a group, or an organization to systematically express and organize their expertise so it can be captured and reused. Yet the expertise of individuals within an organization is perhaps its most valuable resource. Q&A attempts to address this tension by providing an environment in which textual representations of expertise are captured as a byproduct of using the system as a semiautomatic question-answering intermediary. Q&A mediates interactions between an expert and a question-asking user. It uses its experience referring questions to expert users to answer new questions by retrieving previously answered ones. If a user's question is not found within the collection of previously answered questions, Q&A suggests the set of experts who are most likely to be able to answer the question. The system then gives the user the option of passing a question along to one or more of these experts. When an expert answers a user's question, the resulting question-answer pair is captured and indexed under a topic of the expert's choice for later use, and the answer is sent to the user. Unlike previous work on question-answering systems of this sort, Q&A does not assume a fixed hierarchy of topics. Rather, experts build the hierarchy themselves, as their corpus of questions grows. One of the main contributions of this work is a set of techniques for managing the emerging organization of textual representations of expertise over time by mediating the negotiation of shared representations among multiple experts.

### **MOTIVATION: FOUR PROBLEMS IN KNOWLEDGE MANAGEMENT**

The "war stories" that represent the experience and expertise of individuals are invaluable resources to novices in similar situations. Access to expert advice might allow a novice to adopt past solutions or avoid common pitfalls. Expert advice on how to solve a problem can mean the difference between arriving at a solution quickly and failing altogether.

At the same time, once the availability of expert advice is made known to a large group of interested individuals and the expert receives more and more requests for the same kind of information, providing expertise becomes burdensome at best, and intolerable at worst. In such a situation, there is a direct conflict between the goals of the two parties involved. Specifically, the knowledge-seeking individual's goal to have constant and uncompromised access to expert advice is in conflict with the expert's goal to use his time for other things. We call this the problem of *expert/novice goal conflict*. In light of this, we would like to provide an environment that attempts resolve the above conflict by reaching a compromise in which both the expert and the information-seeking user are satisfied most of the time. Such a solution would maximize the availability of expert knowledge while minimizing the amount of effort and time required by the expert to provide advice to his colleagues.

Unfortunately, when an expert is unavailable, so is access to their experience. Moreover, having the expert simply "write it all down" so that other people could benefit from the expert's knowledge in their absence could take years and would require considerable motivation. We call this the *knowledge acquisition bottleneck* [1]. It is a time-consuming and difficult task for an individual, a group, or an organization to systematically express and organize their expertise. A more realistic solution would entail incrementally capturing knowledge over time, as a byproduct of some other activity.

Moreover, although knowledge may be available to an information-seeking individual, it may not be accessible. For example, consider an individual with a question to ask, and assume he does not know of anyone capable of answering it. Even though the answer to his question may be available to him (from his willing colleague in the office down the hall), the question will remain unanswered. We call this the *resource discovery problem*. In light of this second problem, we would like to provide an environment in which resources that could satisfy a request would be provided automatically. Such an environment would need to have an up-to-date model of the kinds of requests a resource might satisfy. Ideally, this model would not require manual construction.

Finally, any system that attempts to capture a large collection of knowledge and make it available to a wide audience must take seriously the modes of access it provides. Solving this *knowledge access problem* is essential to the design of any useful knowledge management system. In this arena, there are essentially two common approaches: (1) *search*: provide an interface which allows users to explicitly express their request explicitly and return the best results or (2) *browse*: provide an organized collection which can be navigated and explored by the knowledge-seeking user. Each have their benefits and pitfalls [2]. Grossly, searching is good because it can require less interaction between the user and the system, and requires no organization, while browsing allows for the user to opportunistically discover relevant material that might be excluded from the results of a search. Searching, on the other hand, tends to be problematic because of its unpredictable results. Browsing, too, can be problematic because the organization of information is a particularly costly overhead. Moreover, the number of pages a user must traverse can make systems based solely on browsing difficult to use. An ideal solution to the knowledge access problem would provide both browsing and searching, and minimize the aforementioned difficulties.

### Q&A: AN OVERVIEW

Q&A embodies a solution that attempts to satisfy all of the above requirements. Q&A is a Web-based question-answering system that mediates interactions between an expert and a question-asking user. It uses its experience referring questions to expert users to answer new natural language questions by retrieving previously answered ones. If a user's question is not found within the collection of previously answered questions, Q&A suggests a set of experts who are most likely to be able to answer the question. The system then gives the user the option of passing a question along to one or more of these experts. When an expert answers a user's question, the resulting question-answer pair is captured and indexed under a topic of the expert's choice for later use, and notification is sent to the user.

There are essentially two main parts of the system's operation: knowledge acquisition and organization, and knowledge retrieval. Knowledge is acquired and organized as a byproduct of experts interacting with the system, while information-seeking users interact with the system to initiate retrieval. Upon failure they initiate knowledge capture by evoking an actual expert. This kind of loop, or *learning cycle*, in which knowledge is acquired as a result of the system's failure is typical of case-based reasoning systems (Kolodner, 1993), of which Q&A is a part.

#### Supporting the Expert

An expert opens an account with Q&A by registering with a central Web server. Initially, the expert supplies contact information, a password, and a brief description of their areas of interest. This description is used by the system as bootstrapping information, until it has gathered enough information about the expert in the form of questions that he has answered. As a result of registering, the expert is given a fragment of HTML source code they can add to their home page. The source code provides a button for other users to browse their questions and answers, as well as a button for them to access (with password) an expert view of their account. When a new question arrives, the second button changes and e-mail is sent to notify the expert. The expert can then view his pending questions and answer them, forward them to another expert, or notify the system when they are irrelevant to his area of expertise.

Questions and answers are stored in a directed graph of topics, subtopics, and related topics. This organization is most easily thought of as an abstraction hierarchy of topics (this is what we mean when we refer to the expert's hierarchy). When the expert answers a question, notification is sent to the user who originally asked the question. The resulting question-answer pair is captured and indexed under one or many topics of the expert's choice. The system recommends the most appropriate categorization to aid the expert user in organizing his space of expertise. It also allows the expert to create new topics (in an arbitrarily deep hierarchy), move questions among the topics, as well as associate subtopics with multiple parent topics. The expert can also decide whether or not to make an answered question *public*, in which case it is indexed for

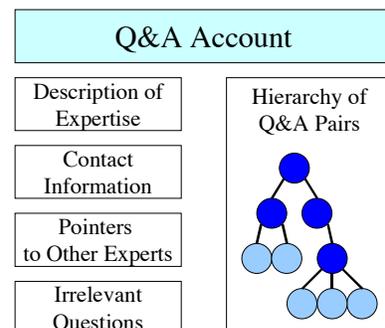


Figure 1: Summary of the contents of an expert's account.



Figure 2: The Q&A browsing interface.

automatic retrieval, or *private*, in which case it is only accessible by browsing his hierarchy. When the expert defers to another expert to answer the question, the system adds it to the other expert's list of pending questions, and sends notification. Should an expert decide the question is irrelevant to his area of expertise, the system indexes the question in a collection of irrelevant questions associated with his account. This collection is used when deciding whether or not an expert should be recommended as a resource for answering a question. This technique is called *negative evidence acquisition* (Kulyukin, Hammond, & Burke, 1998). In addition, the expert can also add pointers to nodes in other expert's hierarchies from within his own hierarchy of topics. These links represent topic-to-topic similarity across multiple collections of expertise. Figure 1 summarizes the contents of an expert's account.

The expert's view of his hierarchy of questions and answers provides him with an interface for editing question-answer pairs and the topics that organize them. Each topic contains a list of questions, subtopics, links to nodes in other expert's hierarchies, and a descriptive piece of summary text, which is displayed at the heading of the topic. The summary text is used in bootstrapping the system in the absence of questions indexed beneath it that would normally be used as description of the content of the topic. This view also provides the expert with a list of related topics in other expert's hierarchies that the system discovers automatically. This is meant to facilitate communication among communities of experts, as well as to encourage experts to link their spaces together. The links to other expert's hierarchies are used to guide expert referral and question retrieval, as well as to provide navigational facilities to information-seeking users who are browsing through the system's database of expertise.

### Supporting the User

Information-seeking users can interact with Q&A by browsing an expert's hierarchy of questions and answers, or by posing a question of the system directly. Users can browse an expert's hierarchy by entering it either from the expert's own home page, by selecting a retrieved expert (described in the next paragraph), or by selecting a retrieved question and navigating through the associated space. As stated earlier, an expert's knowledge is organized as lists of questions and answers indexed under topics. Users can navigate through this space by selecting topics, and eventually questions. Each page contains a topic heading and a list of questions and subtopics. The topic heading consists of a topic title, its path in the expert's hierarchy, a short description, and an optional list of related topics. Each question is linked to its answer. This interface is displayed in Figure 2 [4].

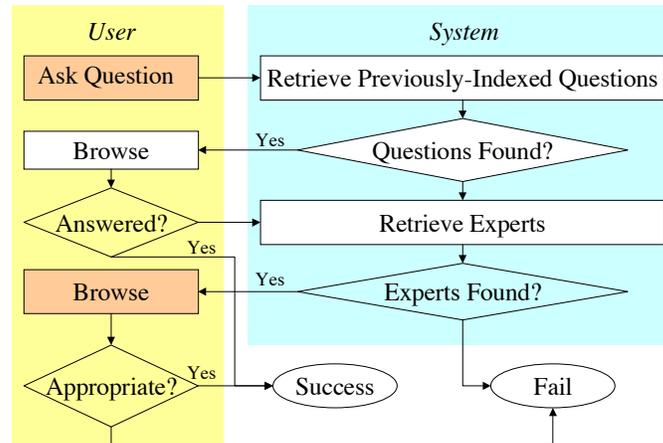


Figure 3: Flow of control between the user and the system during question answering.

When a user enters a question, the system searches for relevant experts. If the incoming question matches a question an expert has previously tagged as irrelevant, that expert is removed from consideration. Otherwise, the system traverses each expert's hierarchy of topics, collecting a list of matching questions. The retrieved questions are displayed as a list. Each question is linked to its answers, as well as to the topic in which it occurs in expert's collection of expertise. The questions are also annotated with the name of the expert, and a link to his contact information and the description of his expertise provided when he opened his account.

Should the user decide that none of the information presented satisfies his information need, he can view the list of experts retrieved, and forward his question along to any of the experts listed. The experts are listed in decreasing order of relevance as determined by the previous search for relevant questions. Each entry is linked to the expert's topic hierarchy, and is annotated with his contact information as well as a short description of the expert's areas of expertise. Should the user choose to browse an expert's hierarchy either before or after he requests for experts to be retrieved, the system remembers the user's question and allows the user to submit his question to the expert at any time. A summary of the flow of control in this situation is displayed in Figure 3.

### Discussion: Q&A is a Solution to the Above Problems

Q&A addresses the problem of expert/novice goal conflict by automatically capturing textual representations of expert knowledge in the form of answers to questions and indexing them by the questions they answer. If an incoming question is sufficiently similar to one that has been previously captured, the expert will never see the new request. The expert is only called upon when the user has determined no answer is available from the system. Moreover, through the use of negative evidence acquisition, the system avoids making the same referral mistakes twice.

In addition, Q&A attempts to address the knowledge acquisition bottleneck by making knowledge capture a byproduct of answering explicit requests. Furthermore, Q&A pragmatically constrains the knowledge captured to answers to questions that users have actually posed. Moreover, because questions originate from specific individuals, experts may tend to be more motivated to answer them.

Q&A addresses the resource discovery problem by suggesting relevant experts when the system has no direct answer to an incoming question. Q&A's model of the kinds of requests an expert can satisfy is determined by the questions the expert has already answered.

Finally, Q&A provides both browsing and searching as modes of interaction for access to its database. Experts organize questions and their answers in a hierarchy of their own creation when they answer a question. Information-seeking users can search for an expert and browse that expert's hierarchy of questions and answers. Experts' hierarchies are interconnected at the topic level to allow users to easily navigate among information stored in multiple accounts. To ease in the organization of expertise, Q&A suggests both topic-to-topic associations as well

as question and answer classifications. Users can also search for the answer to their question by entering their question in natural language and browsing the retrieved results.

### **Related Work**

Q&A is most directly related to work on the Chicago Information Exchange (Kulyukin, 1998a, 1998b) (Kulyukin, Hammond, & Burke, 1998), which was developed at the same time as Q&A [3]. Both Q&A and CIE are descendents of the FAQ Finder system (Burke, *et. al.*, 1997) (Burke, Hammond, & Koslovsky, 1995) which acts as a gateway to frequently asked questions files gathered from newsgroups. Like Q&A, CIE answers incoming questions by retrieving questions it has previously answered or by referring users to experts.

In terms of indexing and retrieval, the main difference between Q&A and the Information Exchange system is that Q&A has no *a priori* notion of organizational structure. In CIE, expertise in the form of question and answer pairs is organized in a fixed hierarchy that is built before the system can be deployed. The assumption made in CIE is that the organization of expertise should reflect the organization's internal structure. However, in practice, this simply is not the case. Many organizations have a fairly flat, non-hierarchical organization. Others have multiple locations for similar kinds of expertise. For example, in a university, expertise in statistical analysis of numerical data can not only be found in the Statistics Department, but in virtually any of the other science departments. Unfortunately, in the absence of organizational structure to guide the creation of the system's taxonomy of topics, the system designers must anticipate the contents of the knowledge the system will contain. In other words, the system designers must either be experts in all of the areas of knowledge contained within the system, or they must coordinate an effort within an organization to discover the most appropriate organizational structure. Moreover, should this structure change during the lifetime of the system, the designers must step in again to reorganize the information in the system. Work on Q&A was motivated by the fact that such assumptions and requirements are unrealistic. Instead of relying on system designers to provide a fixed organizational structure, Q&A relies on the experts to structure to their own areas of expertise.

The second important difference is in research focus. Research on CIE was focused on new techniques for the indexing and retrieval of semi-structured text. As a result, the user interface was more of an afterthought. Work on Q&A was also motivated by the usability problems associated with CIE's interface.

Q&A is also closely related to work on the Answer Garden system (Ackerman & Malone, 1990) (Ackerman, 1998). Answer Garden is a tool for capturing and organizing a memory of expertise. The primary interface to its database of expertise is a branching network of diagnostic questions that users can browse to find the answers they need. Answer Garden routes new questions to appropriate experts by associating experts with parts of this branching network and allowing users to ask a question as they are navigating the hierarchy. Answer Garden also allows experts to modify the network in response to user's questions.

Answer Garden's strength is in the flexibility it allows experts in their task of designing networks of diagnostic questions. However, the system has serious scalability issues. As the number of experts increases, the more burdensome the task of collaboratively organizing networks of expertise becomes. Moreover, having to navigate through a "20 questions" interface is cumbersome for users. Q&A addresses this issue by providing automatic indexing and retrieval of questions and answers. Instead of having to navigate a diagnostic network of questions, users simply enter their question in natural language and browse the retrieved answers. With Q&A, experts need only associate their expertise with topic areas of their own specification; they do not have to construct complicated discrimination networks aimed at leading the user down the right path.

In addition, Q&A is related to work on the design of explanatory hypermedia systems. The most prominent example of these are ASK systems (Ferguson, *et al.*, 1992) (Bareiss & Osgood, 1993). ASK systems organize expertise in the form of videotaped interviews in a network of questions. Unlike most other systems of this type, once the user has viewed a video, the system presents the user with follow-up questions that lead to other videos. This gives the system a conversational feel that is unique to the ASK approach.

While the strengths of ASK systems are their extremely natural conversational feel and the high quality of results retrieved, their main weakness is that they can take months to construct. Moreover, system designers must gather and organize all of the information that will be contained within the system before it can be deployed. If the

collected material becomes inaccurate or new information needs to be added, the system must be augmented by the designers. In other words, the main problem with ASK systems is that they fall prey to the knowledge acquisition bottleneck described above. Q&A (and Answer Garden) attempt to solve this problem by making knowledge acquisition a byproduct of using the system. Unfortunately, although Q&A has several benefits over the ASK approach, the system does not typically achieve the conversational feel the ASK systems provide. Future work on Q&A will attempt to address these shortcomings.

## DETAILS OF THE IMPLEMENTATION

Q&A contains a database of experts and their expertise in the form of question and answer pairs organized under topics. This database is accessible from the Web. In addition, the vector space model (Salton & McGill, 1983) and the cosine metric are used for the automatic indexing retrieval of questions, the indexing and prediction of referral failures, and the discovery of similar topics.

### Indexing and Retrieval of Textual Representations of Expertise

Q&A organizes the representations of the content of its expertise hierarchically. The hierarchy consists of experts at the top, then topics, then questions and their associated answers. The retrieval of expertise proceeds from the top of this hierarchy downward, whereas indexing proceeds from the bottom up. The hierarchical organization is used in anticipation of the need to reduce the cost of retrieval, and is motivated by previous approaches (Burke, et al., 1997) (Burke, Hammond, & Koslovskysy, 1996). An example of this hierarchy is shown in Figure 4. Each level of the hierarchy can be thought of as a separate collection of texts, where each text is the concatenation of the texts below it in the hierarchy. Each level will result in a separate vector space representation, which will be used in the indexing and retrieval of expertise.

Indexing occurs when the expert modifies his account. Specifically, automatic indexing occurs when an expert adds, modifies or removes a topic description, question and answer pair, description of his area of expertise, or an irrelevant question. The indexing occurs locally, and then moves up the hierarchy. For example, if an expert inserts a new question in a topic, a vector representing that question will be added to the vector space representation of that topic. The vector representing that topic in the parent topic's vector space will also be recomputed, and so on, up the hierarchy, until the vector representing the expert is also updated.

Retrieval occurs in the opposite direction. The first phase of retrieval entails selecting candidate experts. For retrieval purposes, experts are represented as term vectors. The vectors are derived from an analysis of their collections of questions and answers, as well as the textual descriptions they provide of the topics in their collection, and their expertise, in general. Associated with each expert is a collection of irrelevant questions, gathered from previous referral failures. An expert is selected if the cosine of the angle between the vector representing the incoming question and the term vector representing the expert is below a threshold, and none of the questions in his collection of referral failures match the incoming question.

Once experts are selected, retrieval proceeds downward through the hierarchy. For each expert, the top level of their topic hierarchy is searched for matching questions and matching subtopics. Topics at each level of the expert's

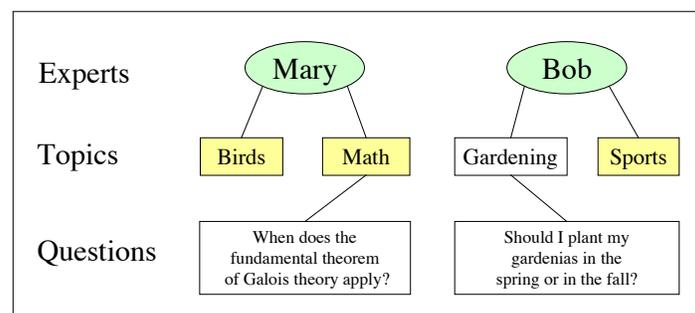


Figure 4: Hierarchical organization for indexing and retrieval.

hierarchy are represented as term vectors derived from an analysis of the questions and subtopics contained within them. If a matching subtopic is found, the retrieval process continues in that topic. If a topic contains a link to a node in another expert's hierarchy, the system ensures that subgraph will be searched, as well. If a matching question is found, it is added to a heap data structure which is responsible for sorting retrieved questions by decreasing degree of similarity to the incoming question. The search terminates when no matching questions or topics remain. The top 10 questions are removed from the heap and presented to the user.

Comparison of textual representations for retrieval is conducted using techniques similar to the ones used by Kulyukin (1998a, 1998b) in CIE. These algorithms have been shown effective in similar contexts of use, and are appropriate because they support hierarchical and predictive indexing, abstraction-based retrieval and negative evidence acquisition.

Indexing and retrieval in Q&A differs from CIE in its organization, the search strategy it uses, the kind of relevance feedback it employs, and its level of autonomy. Q&A retrieves questions fully autonomously, without requiring feedback from the user. Because the goal is to retrieve a few good matches, and provide a list of results, ambiguities in the direction search should proceed are accommodated by following both directions simultaneously. Furthermore, unlike CIE, relevance feedback from the information-seeking user is not used, as it is cumbersome for users, and unnecessary given the choice to display a list of questions, instead of each question in sequence.

Because the techniques we use are very similar to those used by Kulyukin, we provide only a sketch of how they work. Readers interested in the details of these methods are encouraged to refer to Kulyukin's (1998a) ASIS-98 paper.

### **Computation of Indices**

Indices are computed by analyzing the text to be indexed and removing words that occur in an extended version of the stop list derived by Francis and Kucera's (1982). Next, a greedy morphological analysis based on the morphological component of WordNet (Miller, 1995), a lexical database of English words, is performed on the resulting words. This analysis attempts to convert terms into their base forms. For example, "running" is transformed to "run". As a byproduct of applying the morphological conversion rules, words are tagged with their part of speech. Words are tagged as nouns, verbs, adjectives, or adverbs. The default is that the word is tagged as a noun. The resulting words are annotated with their part of speech and form the vector to be indexed.

Questions are treated differently, because of their centrality to the retrieval process. In order to account for lexical variance in the phrasing of questions, the indexer uses a spreading activation technique (Quillian, 1968) (Cohen and Kjeldsen, 1987) backed by WordNet to expand the original vector of indices. WordNet (in addition to its morphological component) provides a semantic network of English words. Words are organized in lists of synonyms, called *synsets*. Each synset corresponds to a particular context in which its words are interchangeable. For example, one synset might be {bank, financial\_institution, financial\_organization} while another might be {bank, slope, incline, side}. Synsets are organized into orthogonal networks of nouns, verbs, adverbs and adjectives. Each network has its own set of relations. For example, verbs are related by synonymy, troponymy, and entailment (e.g., to turn entails to move). Activation is spread from each term along relations in the semantic network for a given depth. The terms activated are collected and annotated by their part of speech and their depth from the origin term. These terms are stored in the vector associated with a given question. This technique is called *predictive indexing* (Kulyukin, Hammond, & Burke, 1998).

### **Computation of Weights**

The overall weight of a word combines three metrics: two statistical, and one semantic. The statistical metrics are *tfidf* (Salton & McGill, 1983) and *condensation clustering* (Bookstein, Klein, and Raita, 1998) or CC. The *tfidf* metric measures the relative importance of a word based on its local occurrence frequency, as well as its role as a predictor across the collection. The CC weight of a term is based on the intuition that terms that follow a random distribution across texts in a collection are not content-bearing, while terms that deviate from randomness indicate content. The semantic weight of a word takes into account the word's part of speech, the depth of activation, and its polysemy (number of senses) in the semantic network. A word's semantic weight is inversely proportional to its polysemy and its distance from the origin of activation.

## Discovery of Similar Topics

In addition to retrieving questions and referring information-seeking users to experts, Q&A proactively encourages connections among expert hierarchies by attempting to discover similar topics. Similar topics in other expert's hierarchies are presented to an expert when he views his hierarchy of topics. In addition, when topics are sufficiently similar, experts are sent notification via e-mail. It is the expert's responsibility to actually add the link to his hierarchy. Similar topics are discovered by computing topic-to-topic similarity in the same way as question-to-question similarity is computed during question retrieval. The only difference is that the input to the search process is the vector representing the topic, and the search terminates at the topic level. The discovery happens off-line every night. When the cosine measure of similarity of two topics is found to be above a threshold  $t_1$ , the topics are added to each other's list of recommended links. When it is found to be above another threshold  $t_2 > t_1$ , the corresponding experts are sent e-mail notification.

## EVALUATION

A small-scale evaluation of Q&A's retrieval algorithms is currently under way. Initial results suggest the system performs well on moderately-sized collections of questions and answers distributed among several experts. However, the importance of an evaluation of the automatic indexing and retrieval algorithms is diminished because they have been previously evaluated in similar contexts, and shown to be effective (Kulyukin, 1998a) (Kulyukin, Hammond, & Burke, 1998). However, an evaluation of the system's usability is clearly needed. The success or failure of knowledge management systems hinges upon their ability to be quickly deployed and easily used by both experts and by information-seeking users. Our position is that while indexing and retrieval algorithms are important to the success of a system, user interface design is equally important.

### Heuristic Usability Evaluation

To the end of demonstrating the usability of Q&A, we performed an informal heuristic evaluation of the usability of Q&A by inspecting the interface (Nielsen & Molich, 1990) (Nielsen, 1992). This kind of inspection is appropriate for laboratories to perform when more detailed ethnographic usability studies (Blomberg & Giacomi, 1993) cannot be completed [6]. In this case, we do not have a sufficient user population to perform such studies. In the following paragraphs, we evaluate the facilities provided by Q&A. Our intention here is simply to demonstrate the system's usability and argue that an overall evaluation of a system should include both an evaluation of retrieval algorithms, and an evaluation of the system's usability.

First, Q&A consistently provides shortcuts to the user. In response to an incoming question, Q&A produces a list retrieved questions. In contrast, systems that require incremental relevance feedback during retrieval (such as CIE) tend to violate this principle of usability, because the user must enter in a dialogue in order to view the next result. Moreover, should the user decide to browse an expert's hierarchy of expertise, the system remembers the user's last question.

Second, Q&A consistently attempts to discourage errors on the part of the user by providing topic summaries at any point in the hierarchy, as well as a standard hyperlinked interface, familiar to most users. In addition, the user's position within the hierarchy of topics is consistently displayed at the heading of the topic page, and search is always conducted across the entire hierarchy. In contrast, the user interface for navigating CIE's hierarchy of expertise is cumbersome to use and often results in error messages. Moreover, at any point in navigation, the user can only see one level of abstraction, and only the title of the topic, which can be daunting to novice users who do not understand the terms used to describe the information therein.

Third, Q&A attempts to speak the language of the user by providing a natural language interface to its database of expertise. A point of concern within the system is the fact that the expert's answer to a question might not as intelligible to the user as it could be. This issue was uncovered in field studies of AnswerGarden (Ackerman, 1998), as well. Future efforts on improving Q&A will be aimed at addressing this issue.

Finally, Q&A provides a combination of two prevalent information access paradigms: browsing and searching. Similar systems focus on browsing (e.g., AnswerGarden) or searching (e.g., CIE) alone. Q&A augments the information access modalities exemplified by these systems by providing both search facilities, as well as easy-to-use navigational facilities, in an attempt to garner the benefits of both approaches.

## FUTURE WORK

In the near future, work on Q&A will focus on including new techniques for text classification that are appearing in the digital libraries literature (Dumais, *et al.*, 1998). It is our hope that adopting techniques such as the use of support vector machines (SVMs) for classification will improve the classification accuracy, and therefore improve precision in retrieval. In addition, we are considering investigating the problem of building uniform abstraction hierarchies across multiple accounts. As it is, the system provides the facilities for sharing representations of expertise between experts by the formation of topic-to-topic links. Extending this notion to allow users of the system to collaboratively build a uniform abstraction hierarchy is a significant challenge.

Because Q&A is intended to be a wide-scale system for large groups of experts, issues of validity and quality arise. In order to accommodate for these situations, feedback from information-seeking users will need to be incorporated. Moreover, the problem of malicious experts will need to be addressed.

In addition, we would like to extend Q&A's organization to support follow-up questions, as was done in the ASK systems (Ferguson, *et al.*, 1992) (Bareiss & Osgood, 1993). One of the major contributions of efforts on building ASK systems was the taxonomy of follow-up questions. Using this taxonomy to structure follow up questions to answers in Q&A may encourage experts to provide useful, audience-appropriate information, and hence address some of the issues elucidated by Ackerman's (1998) field study of Answer Garden.

Finally, we are examining the role of context in information retrieval. It is a significant challenge for experts to determine the appropriateness of a response to a question without any knowledge of the user's background, their current task, the seriousness of their request, and many other variables. The elision of contextual information poses a significant limitation to any information retrieval system, but it becomes even more apparent when humans are placed in the loop. Work on a system called Watson (Budzik & Hammond, 1999), attempts to address these issues by observing users as they work in everyday applications such as Web browsers and word processors. Nonetheless, more research is needed in order to determine the best way to communicate contextual information to human experts.

## CONCLUSION

We presented Q&A, a system for the capture, organization and reuse of expertise. Q&A answers questions on behalf of an expert by retrieving questions that expert has previously answered. If a user's question is not found within the collection of previously answered questions, Q&A suggests the set of experts who are most likely to be able to answer the question. The system then gives the user the option of passing a question along to one or more of these experts. When an expert answers a user's question, the resulting question-answer pair is captured and indexed under a topic of the expert's choice for later use, and the answer is sent to the user. In addition to facilitating knowledge capture and reuse, Q&A fosters communication between experts by discovering experts with like interests. Q&A builds on previous techniques for building question-driven information retrieval systems, but represents a significant step forward in the flexibility it provides for the organization of knowledge, as well as its user interface.

## NOTES

1. This phrase was first used by (Hayes-Roth *et al.*, 1983) in reference to the time-intensive effort of constructing rule-based expert systems. We think it is apt here as well, even though this research is primarily concerned with the capture of textual representations.

2. It is well beyond the scope of this paper to discuss them all.
3. Q&A was first described in (Budzik & Hammond, 1998). Q&A uses many of the same indexing and retrieval algorithms used in CIE. We are indebted to Valadimir Kulyukin for providing an implementation of these algorithms for use in this system.
4. These questions are from an early version of Q&A that was used by visitors during the 1996 Democratic National Convention in Chicago.
5. WordNet is a trademark of Princeton University.
6. For an excellent overview of usability evaluation techniques, see (Mark & Nielsen, 1994).

## REFERENCES

- Ackerman, M., and Malone, T. (1990). "Answer Garden: A Tool for Growing Organizational Memory." In *Proceedings of CSCW-90*.
- Ackerman M. (1998). Augmenting Organizational Memory: A Field Study of Answer Garden. *ACM Transactions on Information Systems* 16, 3, 203-224.
- Bareiss, R., and Osgood, R. (1993). Applying AI Models to the Design of Exploratory Hypermedia Systems, in *Hypertext '93*, ACM Press, 94-105.
- Blomberg, J., and Giacomi, J. (1993). "Ethnographic Field Methods and Their Relation to Design." In Schuler, D. and Namioka, A., (Eds.) *Participatory Design: Principles and Practices*. Lawrence Erlbaum & Assoc., Hillsdale, NJ.
- Budzik, J., and Hammond K. (1998). "Learning for Question-Answering and Text Classification: Combining Knowledge-Based and Statistical Techniques." In *Proceedings of the AAAI-98 Workshop on Learning For Text Classification*. Menlo Park, CA: AAAI Press, 1998.
- Budzik, J., and Hammond K. (1999). "Watson: Anticipating and Contextualizing Information Needs." In *Proceedings of the ASIS 1999 Annual Conference*.
- Burke, R., Hammond, K., Kulyukin, V., Lytinen, S., and Tomuro, N. (1997). Question Answering from Frequently-Asked Question Files: Experiences with the FAQ Finder System. Technical Report TR-97-05, University of Chicago, Department of Computer Science.
- Burke, R., Hammond, K., Kozlovsky, J. (1995). "Knowledge-based Information Retrieval for Semi-Structured Text." In *Working Notes from the AAAI-95 Fall Symposium on AI Applications in Knowledge Navigation and Retrieval*. Menlo Park, CA: AAAI Press, 19-24.
- Dumais, S., Platt, J., Heckerman, D., and Sahami, M. (1998). "Inductive Learning Algorithms and Representations for Text Categorization." In *Proceedings of CIKM-98*. , New York: ACM Press.
- Ferguson, W., Bareiss, R., Birnbaum, L., and Osgood, R. (1992). ASK Systems: An Approach to the Realization of Story-Based Teachers. *The Journal of the Learning Sciences* 1, 2, 95-134.
- Francis, W., & Kucera, H. (1982). *Frequency Analysis of English Usage*. New York: Houghton Mufflin.
- Hayes-Roth, F., Waterman, D., & Lenat, D. (Eds.). (1983) *Building Expert Systems*. Reading, MA: Addison-Wesley.
- Kolodner, J. (1993). *Case-Based Reasoning*. San Francisco: Morgan Kaufmann.

Kulyukin, V. (1998a). "An Interactive and Collaborative Approach to Answering Questions for an Organization." In *Proceedings of the ASIS-98 Midyear Conference*.

Kulyukin, V. (1998b). *Question-Driven Information Retrieval Systems*. Ph.D. Thesis, Department of Computer Science, The University of Chicago.

Kulyukin, V., Hammond, K., and Burke, R. (1998). "Answering Questions for an Organization Online." In *Proceedings of AAAI-98*. AAAI Press, Menlo Park, CA.

Mark, R. L., and Nielsen, J. (1994). "Usability Inspection Methods: Executive Summary." In Baecker R., Gruden, J., Buzton, W., and Greenberg, S. (Eds.), *Readings in Human-Computer Interaction: Toward the Year 2000*. Morgan Kaufmann, San Francisco, CA.

Miller, G. A. (1995). WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11).

Nielsen, J. (1992). "Finding Usability Problems Through Heuristic Evaluation." In *Proceedings of CHI-92*. ACM Press, New York, NY.

Quillian, M. R. (1968). Semantic Memory. In Minsky, M. (Ed.), *Semantic Information Processing* (pp. 216-270). Cambridge, MA: MIT Press.

Salton, G., & McGill, M. (1983). *Introduction to modern information retrieval*. New York: McGraw-Hill.