

Question Answering from Frequently Asked Question Files

Experiences with the FAQ FINDER System

*Robin D. Burke, Kristian J. Hammond, Vladimir Kulyukin,
Steven L. Lytinen, Noriko Tomuro, and Scott Schoenberg*



■ This article describes FAQ FINDER, a natural language question-answering system that uses files of frequently asked questions as its knowledge base. Unlike AI question-answering systems that focus on the generation of new answers, FAQ FINDER retrieves existing ones found in frequently asked question files. Unlike information-retrieval approaches that rely on a purely lexical metric of similarity between query and document, FAQ FINDER uses a semantic knowledge base (WORDNET) to improve its ability to match question and answer. We include results from an evaluation of the system's performance and show that a combination of semantic and statistical techniques works better than any single approach.

In the vast information space of the internet, individuals and groups have created small pockets of order that are organized around their particular interests and hobbies. For the most part, those who build these information oases have been happy to make their work freely available to the general public. One of the most outstanding examples of this phenomenon can be found in the wide assortment of frequently asked question (FAQ) files, many associated with USENET news groups, which record the consensus of opinion among a group on some common question. So that newcomers do not ask the same questions again and again, most FAQs are periodically posted on the news groups to which they are relevant. This information-distribution mechanism works well for individuals who are sufficiently interested in a topic to subscribe to its news group but not necessarily to those with a more transient interest. What is needed is a centralized means of access to these answers.

We believe that the most natural kind of

interface to a database of answers is the question, stated in natural language. Although the general problem of understanding questions stated in natural language remains open, we believe that the simpler task of matching questions to corresponding question-answer (QA) pairs is feasible and practical. The aim of the FAQ FINDER Project is to construct a question-answering system that extends further the aim and intent of the FAQ-file phenomenon. The system is an information service, available on the World Wide Web at faqfinder.cs.uchicago.edu/. A user can pose a question and if it happens to be similar to one of the FAQs whose answer has been recorded in a FAQ file, the system will have a good chance of answering it.

FAQ FINDER is built on four assumptions about FAQ files: (1) QA format: all the information in a FAQ file is organized in QA format (Kulyukin, Hammond, and Burke 1996); (2) locality of information: all the information needed to determine the relevance of a QA pair can be found within the QA pair; (3) question relevance: the question half of the QA pair is the most relevant for determining the match to a user's question; (4) general knowledge: broad, shallow knowledge of language is sufficient for question matching.

Assumptions 1 to 3 lead to an essentially case-based (Kolodner 1993) view of the FAQ retrieval problem. A QA pair might loosely be considered a kind of case: It is a piece of knowledge that has been considered useful enough to be codified for reuse. The question serves as an index to the knowledge contained in the answer. These assumptions do not hold for all FAQ files, as we discuss later, but they hold often enough to form a good starting point for research.

FAQ FINDER

Please Enter a natural language question to be answered.
(for example, "Who is Lila Feng?")
FAQfinder is not a search engine:do not enter keywords

Is downshifting a good way to slow down my car?

Find Answer! Clear Question

Quick Match

Merge Related FAQs

Supported FAQs

System Status

Figure 1. Asking a Question of FAQ FINDER.

Sample Interaction

Figures 1 and 2 depict a typical interaction with FAQ FINDER. Figure 1 shows the World Wide Web (WWW) interface to the system. Suppose the user enters the following question: Is downshifting a good way to slow down my car?

FAQ FINDER compares the question to its set of FAQ files, returning a list of files ranked by their relevance to the question, including, in this case, the file `auto_consumer_FAQ` as the most relevant file. Some files of questionable relevance are also retrieved, such as the `car_audio_FAQ`, a typical artifact of a keyword-based retrieval system. If the user chooses Quick Match when entering a question, the system will skip this first step of manual file selection and automatically choose the top-ranked file to match against.

When a FAQ file is chosen, the system iterates through the QA pairs in the file, comparing each against the user's question and computing a score. The best five matches are returned to the user along with the first line of the answer; so, the validity of the answer

can easily be determined. Figure 2 shows the results of this matching step comparing our downshifting question with the entries in `auto_consumer_FAQ`. The correct answer, a question about downshifting and braking, is first, followed by two questions about brakes and two about tires. By selecting the appropriate link, the user can view the entire answer given in the FAQ file.

How the System Works

As the example in figures 1 and 2 shows, interaction with FAQ FINDER occurs in a series of stages. The first step is to narrow the search to a small set of FAQ files likely to contain an answer to the user's question. Second, each QA pair is matched against the user's question to find the ones that best match it.

For the first stage of processing, FAQ FINDER uses standard information-retrieval technology, the public-domain SMART information-retrieval system (Buckley 1985), to perform the initial step of narrowing the focus to a small subset of the FAQ files. The user's question is treated as a query to be matched



[FAQ Finder Help](#)
[About FAQ Finder](#)

Question: Is downshifting a good way to slow down my car?
 File: autos_consumer_FAQ

[View Entire FAQfile](#)

[Rephrase Question](#)

[Select Different FAQ](#)

[Start Over!](#)

- [They tell me I should downshift when braking to slow my car down. Is this really a good idea?](#)
 It used to be a very good idea, back in the days of medi...
- [What about DOT-5 brake fluids?](#)
 This breaks down in to two parts. The DOT-5 specificati...
- [How often should I replace my brake fluid?](#)
 Probably more often than you do. Traditional brake flui...
- [Can I rotate radials from side to side, or rotate them only on one side of my car?](#)
 Car and tire manufacturers have differing views on this ...
- [How many snow tires should I buy, and if I buy 2, which end of the car should I put them on?](#)
 In short, 4, and both ends. To explain, many drivers in...

Figure 2. Choosing an Answer.

against the library of FAQ files. SMART stems all the words in the query and removes those on its stop list of frequent words. It then forms a term vector from the query, which is matched against similar vectors already created for the FAQ files in an offline indexing step. The top-ranked files from this procedure are returned to the user for selection. We have not found it necessary to tinker with the default configuration of SMART. We treat this part of the system as a black box that returns relevant files.

The second stage of processing in FAQ FINDER is a question-matching process. Each question from the FAQ file is matched against the user's question and scored. We use three metrics in combination to arrive at a score for each QA pair: (1) a statistical term-vector similarity score t , (2) a semantic similarity score s , and (3) a coverage score c . These metrics are described in detail in the following discussion. Overall match similarity m is a weighted average

$$m = \frac{tT + sS + cC}{T + W + C}$$

where T , S , and C are constant weights that adjust the reliance of the system on each metric.

The statistical similarity score at the QA pair level is computed in a manner similar to SMART's document matching. A QA pair is represented by a *term vector*, a sparse vector that associates a significance value with each term in the QA pair. The significance value that we use is commonly known as *tfidf* (Salton and McGill 1983). If n is the *term frequency* (the number of times that a term appears in a QA pair), m is the number of QA pairs that the term appears in in the file, and M is the number of QA pairs in the file, then *tfidf* is equal to $n \times \log(M/m)$. The idea behind this measure is to evaluate the relative rarity of a term within a space of documents and use it as a factor to weight the frequency of the term in a particular document. A term that appears in every QA pair in a file is probably of little value, and its *idf*, or $\log(M/m)$ value, would correspondingly be zero. A term that appears in only a single QA pair would have the highest-possible *idf* value. Term vectors for user questions are computed similarly by using the *idf* values associated with terms in a given FAQ. Term vectors are then compared using another standard information-retrieval metric, the

cosine of the angle between the vector representing the user's question and the vector representing the QA pair.

The *term-vector metric* allows the system to judge the overall similarity of the user's question and the QA pair, taking into account the frequency of occurrence of different terms within the file as a whole. This metric does not require any understanding of the text—a good thing because the answers in FAQ files are free natural language text, often several paragraphs or more in length.

The *tfidf* measure has a reasonably long history in information retrieval and has generally been thought to work best only on relatively lengthy documents because only long documents have enough words for statistical comparisons to be considered meaningful. However, we have found that this metric contributes significantly to the overall performance of our matching algorithm (see evaluation discussion to follow), despite the extremely short length of the documents (QA pairs) involved in the matching.

The *semantic-similarity metric* enhances term-vector comparison by taking into account a shallow level of semantic analysis of lexical items that appear in user and FAQ questions. The semantic-matching algorithm in FAQ FINDER is designed to handle variations in lexical content between input and FAQ questions. For example, consider the following questions: How can I get my ex-spouse's debts off my credit report? Can I get credit if my ex-husband had a lot of debts?

Here, the difficulty is that there are many ways of expressing the same question, all using different (but semantically related) words and phrases. In large documents, these lexical variations might not affect term-vector comparison greatly because over the course of the document, a variety of synonymous terms might be used. However, in FAQ FINDER, because matching is being performed on a small number of terms, the system needs a means of matching such synonyms.

The need to match related words suggests the need for a level of semantic analysis of user and FAQ questions. However, in the FAQ FINDER system, it is important to balance the depth of analysis with the breadth of coverage. Deep causal reasoning about questions would not be feasible because it would require too much knowledge engineering to cover all the necessary areas of knowledge. For FAQ FINDER, we believe that a shallow lexical semantics provides an ideal level of analysis for the system. Such a semantics has three important advantages: (1) it provides critical

semantic relations between words, (2) it does not require expensive computation to compute relations, and (3) it is readily available.

For example, because the consumer credit FAQ file is full of questions about credit reports and debts, it is important that the system identify the relation between ex-spouse and ex-husband. Such a match can be performed at the level of the words themselves, hence our term *shallow lexical semantics*. As an example of deeper semantics, we can consider the following pair of questions: (1) How do I reboot my system? (2) What do I do when my computer crashes?

Here, there is a causal relation between the question variants: Rebooting is a causal consequence of having one's computer crash. To match these questions, the system would have to understand the causality of the computer domain. Because FAQ FINDER is intended to encompass the whole gamut of USENET topics, not just computers, it is impractical to expect even this simple level of domain-specific knowledge representation.

FAQ FINDER obtains its knowledge of shallow lexical semantics from WORDNET, a semantic network of English words (Miller 1995). WORDNET provides a system of relations between words and synonym sets and between synonym sets themselves. The level of knowledge representation does not go much deeper than the words themselves, but there is an impressive coverage of basic lexical relations. By using a marker-passing algorithm (Quillian 1968), the FAQ FINDER system uses the WORDNET database to accept variations such as ex-husband for ex-spouse. In particular, marker passing is performed over the network's synonym and hypernym (that is, is-a) links.¹

Marker passing is performed to compare each word in the user's question with each word in the FAQ file question. Let u_i be the i th word of the user's question. Let f_j be the j th word of the FAQ file question. The similarity scores of these two words is given by

$$s(u_i, f_j) = H - (p \frac{H-L}{D})$$

where p is the length of the path between u_i and f_j , D is the maximum path length permitted by the system, and H and L are constants. The score is in the range of H and L inclusively and linearly inverse to the number of links traversed between the two words. Marker passing is not used for words that are identical or related by morphology; they are given fixed scores.

The matrix s for a user question of length n and a FAQ file question of length m is an $n \times m$ matrix representing all possible comparisons of words in the two questions:

$$s_{u,f} = \left[\begin{array}{c} s(u_1, f_1) \cdots s(u_1, f_m) \\ \vdots \\ s(u_n, f_1) \cdots s(u_n, f_m) \end{array} \right]$$

S is reduced to a single value, w , by choosing the maximum match score for each user question and then averaging these maxima for all words. The value $w(u, f)$ for semantic relatedness of the two questions u and f is given by

$$w(u, f) = \frac{\sum_{i=1}^n \max(s(u_i, f_0), \dots, s(u_i, f_m))}{n}$$

WORDNET is not a single semantic network; separate networks exist for nouns, verbs, adjectives, and adverbs. Syntactically ambiguous lexical items, such as *run*, which could be either a noun or a verb, appear in more than one network. We found that unrestricted marker passing, using all networks in which a term appears, led to too many spurious matches, a common problem in marker-passing systems in general (Collins and Quillian 1972). We tried several approaches to disambiguate terms to a single WORDNET network, including using an existing part-of-speech tagger (Cutting et al. 1992) and context-free parsing of questions, but in the end, we found that simply relying on the default (most common) word sense for each word worked as well as any of the more sophisticated techniques. (The MOBY part of the speech dictionary [Ward 1993] was used to supply the most common sense of each word.)

The third metric measures the degree of coverage of user terms by the FAQ question. The intuition behind this measure is to penalize questions that are lacking corresponding words for each word in the user's question. In other words, we do not care if the FAQ file question answers many questions at once, but we want to make sure that the important concepts in the user's question are covered. The coverage value is the percent of words in the user question that have a nonzero s (computed in the semantic similarity metric) for some term in the FAQ question.

Evaluating FAQ FINDER

We evaluated the performance of FAQ FINDER on a corpus of questions drawn from the log files of the system's use during the period May to December 1996. A total of 241 test questions were used to perform the evaluation. We manually scanned each FAQ file for answers to each question and determined that 138 questions had answers in the FAQ file corpus,

and 103 questions were unanswered.

The most obvious precedents to FAQ FINDER are information-retrieval systems, and standard information-retrieval evaluation techniques are a starting point for the evaluation of the system. However, evaluation in FAQ FINDER is complicated by the fact that the task of the system is different than the information-retrieval problem as it is typically posed. Normally, the assumption is that there is a document collection in which there can be a number of documents relevant to the user's query. In contrast, FAQ FINDER works under the assumption that there is such a thing as a *correct answer*, a single FAQ QA pair that best addresses the user's question as it was posed. The system's job is to return the answer within the small fixed-size set of results that can be displayed on a single web page. Relevance is not that useful a measure to us because within a given FAQ, most of the answers are probably somewhat relevant to the user's query.

Because FAQ FINDER's task is different, the traditional infrared evaluation metrics of recall and precision must be modified somewhat. *Recall* typically is a measure of the percentage of relevant documents in the document set that are retrieved in response to a query, whereas *precision* is a measure of the percentage of retrieved documents that are relevant. In our case, however, because there is typically one correct answer to be retrieved from a FAQ, these are not independent measures of performance. Assuming that an answer to a user question exists in a FAQ file and that the system returns 5 QA pairs, FAQ FINDER will perform at either 100-percent recall and 20-percent precision (if the answer is retrieved), or 0-percent recall and 0-percent precision (if it is not). If no answer exists, then precision will be 0 percent, and recall is undefined.

To measure the quality of retrieval, we calculate our version of recall, which amounts to the percent of questions for which FAQ FINDER returns a correct answer when one exists. Our calculation is slightly different from traditional recall measures because it does not penalize the system if there is more than one correct answer in the file. If there are several answers within a file that respond to a user's question, it does not make sense to regard retrieval of only one of these answers as only partial success. If the user's question is answered, it is irrelevant that there was another QA pair that also answered it. Instead of precision, we calculate a value called *rejection*, the percentage of questions that FAQ FINDER correctly reports as being unanswered

FAQ FINDER works under the assumption that there is such a thing as a correct answer, a single FAQ QA pair that best addresses the user's question as it was posed.

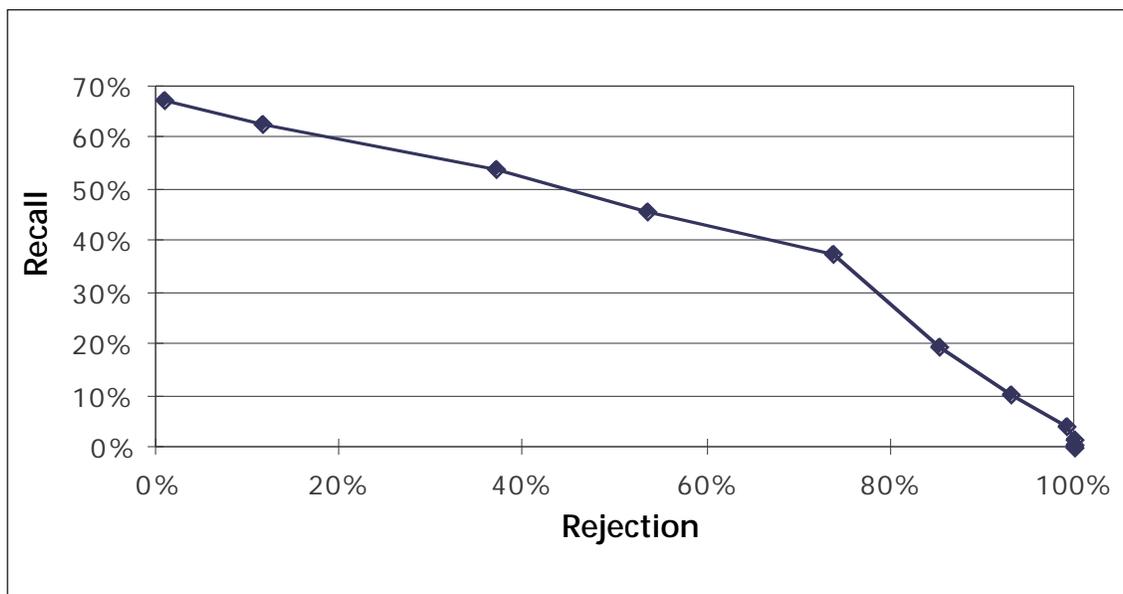


Figure 3. Recall versus Rejection for FAQ FINDER.

in the file. We feel that these metrics better reflect FAQ FINDER's real-world performance than traditional recall and precision.

Rejection is adjusted in FAQ FINDER by setting a cutoff point for the minimum-allowable match score. As with precision, there is a trade-off between recall and rejection. If the rejection threshold is set too high, some correct answers will be eliminated; however, if the threshold is too low, then incorrect responses will often be given to the user when no answer exists in the FAQ file.

Results

SMART is highly effective at the file-retrieval task. The correct file appears 88 percent of the time within the top 5 files returned to the user, and 48 percent of the time in the first position. These results translate to 88-percent recall and 23-percent precision.²

Figure 3 best shows the recall versus rejection results that we obtained for the second stage of the FAQ FINDER system. As the graph shows, rejection is somewhat low for reasonable values of recall, meaning that the system confidently returns garbage in most cases when there is no correct answer in the file. If the rejection threshold is lowered to make it easier to identify questions without good answers, recall drops dramatically. However, the top value for recall is encouraging: There is a better than two-thirds probability that the system will find the correct answer.

Ablation Study

Our next step was to evaluate the contribution of different components of the matching scheme through an ablation study. We selectively disabled different parts of the system and ran the same corpus of questions. There were four conditions: (1) a *random condition*, in which QA pairs were selected randomly from each FAQ file; (2) a *coverage-only condition*, in which the coverage score for each question was used by itself; (3) a *semantic score-only condition*, in which only the semantic scores derived from WORDNET were used in evaluating answers; and, (4) a *statistical score-only condition*, in which the term-vector comparison was used in isolation.

Figure 4 shows average recall results for these conditions. Interestingly, both WORDNET and our statistical technique are contributing strongly to system performance. These two methods had similar average recall but are clearly not equivalent measures because their combination yields results that are better than either individually. These results confirmed our earlier results with a small corpus of questions, which showed an even more dramatic benefit from the combination of methods.

Figure 5, which shows the recall versus rejection analysis for these conditions, provides even more evidence that the two measures differ. The curve for the semantic-scoring condition is particularly striking. Although recall in this condition is weaker than the system as a whole, this metric shows good rejection

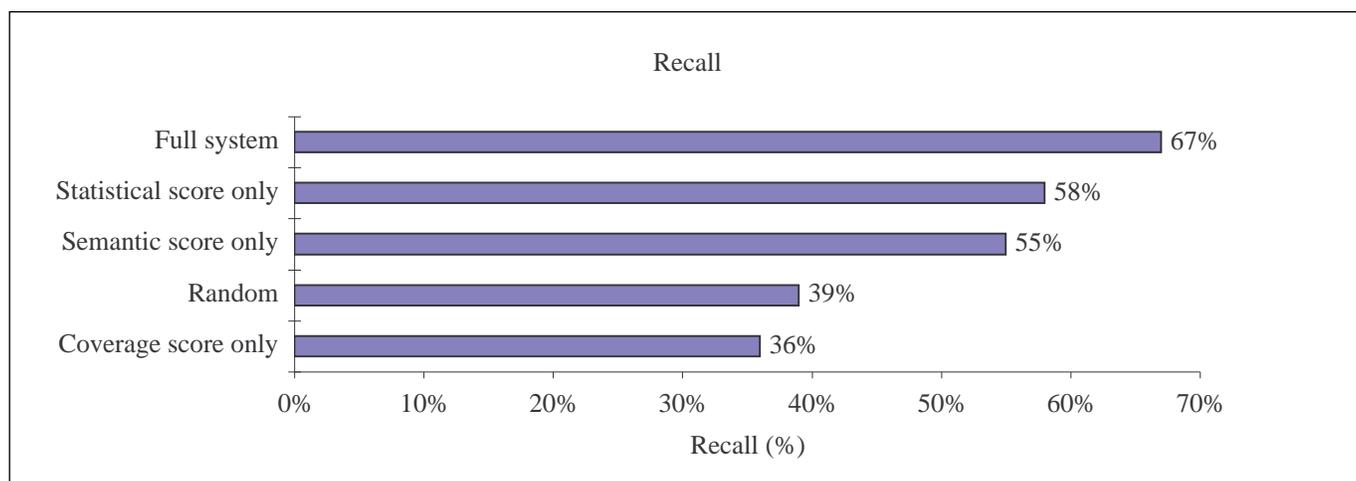


Figure 4. Recall for Ablation Study.

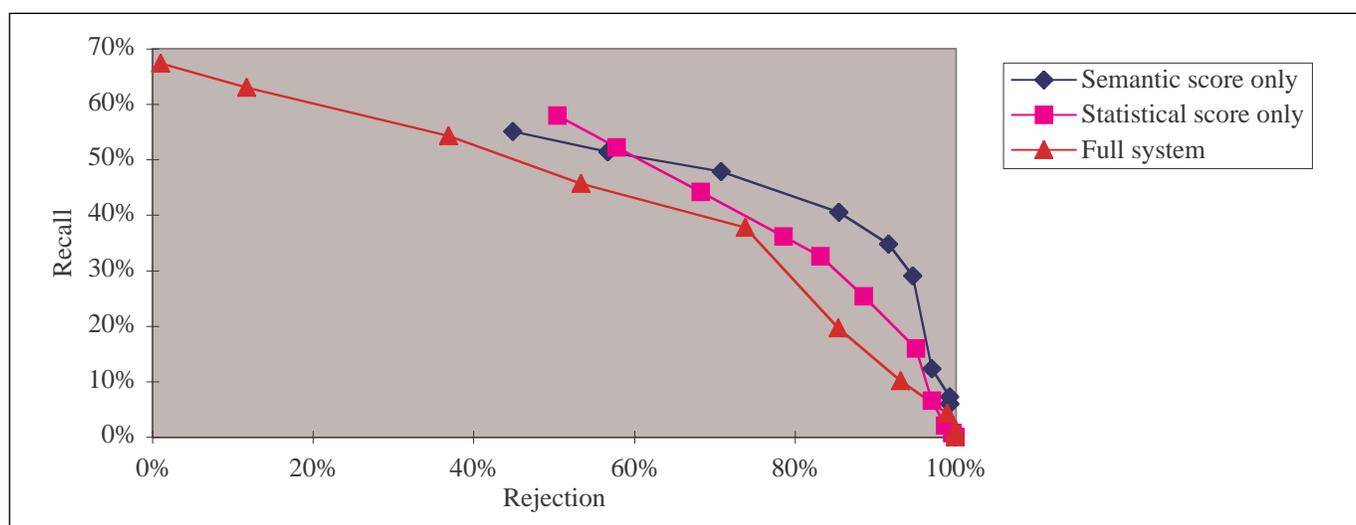


Figure 5. Recall versus Rejection for Ablation Study.

tion performance, suggesting that the application of semantic information might be used specifically to improve rejection.

Discussion

A preliminary analysis of the *failure cases*, questions for which the system failed to find answers, suggests that the biggest culprit is usually undue weight given to semantically useless words. For example, a question such as “where can I find woodworking plans for a futon?” retrieves questions that incorporate the word *woodworking* as strongly as those that contain *futon*, even though *futon* should be a much more informative term inside the woodworking FAQ than *woodworking*, which applies to everything. The problem is that the term *woodworking* does not appear that often in the FAQ despite its close semantic relation

to words that do appear.

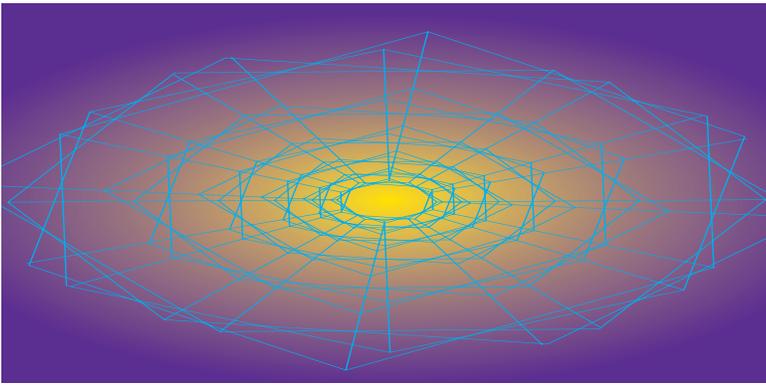
Another type of problem commonly encountered with FAQ FINDER is related to violations of the assumptions about FAQ files discussed at the beginning of this article: the assumptions of QA format, locality of information, question relevance, and sufficiency of general knowledge. We have found many instances in which these assumptions are violated. For example, FAQ writers frequently use headings to mark sections of their documents and rely on the reader’s interpretation of these headings in their question writing. In the Investment FAQ file, the following text can be found:

Subject: Analysis – Technical:

...

Q: Does it have any chance of working?

The *it* is, of course, intended to refer to tech-



nical analysis. However, FAQ FINDER is currently not capable of making use of this referent because it lies outside the QA pair, making it more difficult to match against a question such as “Does technical analysis work?”³ Part of our intent as we automate the tagging process is to make heading information available to the matcher.

There are other more difficult cases of ellipsis found in FAQ files. In the Wide-Area Information Server FAQ, the following passage can be found:

Q:What is Z39.50?
A: ...
Q:Do they interoperate?

The reference *they* refers to both Z39.50, an information-retrieval standard, and WAIS, the subject of the FAQ. We do not expect FAQ FINDER to be able to dissect references that are this oblique. It would, however, be useful to refer back to earlier questions if there is no heading information with which to resolve a referent.

One FAQ-specific phenomenon we have encountered is the use of *metasyntactic variables*, meaningless pieces of text that stand in for a filler, which can vary. For example, the Pool and Billiards FAQ contains the question

Q: What are the rules for XXX?
A: STRAIGHT POOL...
EQUAL OFFENSE...

Metasyntactic variables often have a distinct form and can easily be recognized. We anticipate that a mechanism similar to a heading recognizer could be used to recognize the subanswers within a multipart answer such as this. Not every variable can be so treated, however. The Woodworking FAQ contains the question

Q: Should I buy a Sears blurrf?

The answer does not enumerate the entire catalogue of Sears power tools: The same advice is intended to apply to all. The reader

is supposed be capable of matching the nonsense word against the name of any power tool. This knowledge is exactly the type of domain-specific knowledge that we have sought to avoid including in FAQ FINDER. FAQ FINDER can successfully match this question against such questions as “Are Sears power drills a good buy?” because the word *Sears* is sufficiently distinctive, but it would fail to match against a question such as “What kind of power drill should I buy?”

Future Work

The previous discussion suggests many areas that deserve future research attention. One of the most obvious open questions is how do we improve the system’s rejection of unanswerable questions. We have concentrated our tuning of the system to maximize recall, so that answerable questions will be answered. However, it is also useful to be informed that an answer does not exist within a FAQ file, possibly suggesting to the user that the question should be submitted to the FAQ’s related news group.⁴

One way of approaching this problem is to focus on the small retrieved set of QA pairs before they are returned to the user. We know from our evaluation that if an answer is present in the FAQ file, the system is likely to find it. Therefore, if none of the QA pairs returned by the system is, in fact, a good answer, the chances are that the system should report that the answer does not exist. We also know that semantic information seems to have better rejection characteristics than statistical information. We might be able to perform a more in-depth analysis, involving deeper natural language processing to accept or reject each returned set of questions. Because this set is, by definition, small, such intensive processing would not be as computationally prohibitive as performing deeper natural language processing throughout the entire matching process.

An important part of maintaining the performance of FAQ FINDER on a large set of FAQ files will be the incorporation of new vocabulary items into WORDNET. Because WORDNET was formed from a corpus of everyday English, its vocabulary does not include many technical terms or proper nouns. Unfortunately, because of the technical nature of many FAQ files, technical terms and proper nouns constitute a significant portion of the domain vocabulary of these files. In addition, these terms can be the most useful in retrieving relevant FAQ QA pairs because they are often

the most specific and discriminating terms. Thus, the fact that they are missing from WORDNET can significantly impair the performance of FAQ FINDER.

We are investigating ways in which information obtained from the parses of questions can be used to automatically acquire additional terms and build the appropriate synonym and hypernym links for these terms in one of the WORDNET hierarchies. We will rely on feedback from the user to tell the system when a good match has been found between a user question and a FAQ QA pair.⁵ If the user indicates the system retrieved a relevant answer, then any words in either the user or the FAQ question that are not contained in WORDNET have the potential to be acquired. The system will then attempt to match the unknown word with a synonym in the other question. Both questions will be parsed and position in the parse tree used to determine which words are candidate synonyms of the unknown word.

Because the matching process between question pairs is likely to incorrectly propose some synonyms of unknown words, our approach is to collect synonyms for unknown words over time and propose new WORDNET entries by analyzing collections of possible synonyms for each unknown term. Clustering algorithms are likely to be of use here in determining the likely best entry(ies).

Conclusion

FAQ FINDER is a web-accessible knowledge-based information-access system that relies on the knowledge engineering inherent in FAQ files distributed on the internet. The FAQ FINDER Project has shown that when there is an existing collection of questions and answers, question answering can be reduced to matching new questions against QA pairs, a considerably more tractable task than question understanding. The system combines statistical measures and shallow lexical semantics to match users' questions against QA pairs from FAQ files. Our evaluation, although conducted with a small corpus of questions, has demonstrated the effectiveness of the system.

The power of our approach rises from the fact that we are using highly organized knowledge sources that are designed to answer the commonly asked questions. We do not need our systems to actually comprehend the queries they receive (Lang et al. 1992) or generate new text that explains the answer (Souther et al. 1989). They only have

to identify the files that are relevant to the query and then match against the segments of text that are used to organize the files themselves.

Ultimately, FAQ files are a social phenomenon, created by people to record and make public their understanding of a field. In general, the FAQ FINDER Project is interesting in that it uses not just the existing archives on the internet but also the existing sociology. One of the more powerful aspects of the news groups is the collective desire on the part of the users to get it right. This drive has already resulted in the existence of FAQ files themselves. Our aim in FAQ FINDER is to further this goal by making the answers recorded in FAQ files more widely available.

Acknowledgments

In addition to the authors, the development team for FAQ FINDER 2.0 consisted of Jay Budzik and Benjamin Young of the University of Chicago and David Mortman of Commonwealth Edison. Previous versions of the system benefited from contributions by Edwin Cooper, Xiaobin Fu, Julia Kozlovsky, Charles Martin, and Tom McDougal. Research on FAQ FINDER has been supported with funds from the Office of Naval Research and hardware donations from Apple Computer.

Notes

1. Hypernym links are only provided in WORDNET for nouns and verbs; thus, only synonym links are used for adjectives and adverbs.
2. We do not use rejection in evaluating SMART's performance because it is a standard information-retrieval system designed to find relevant documents, not answer questions.
3. The statistical component of the matcher might enable the question to be retrieved anyway, but we would prefer that the semantic component also be able to contribute in this situation.
4. If rejection were sufficiently good, we could incorporate such an option into the system itself.
5. With our previous local version of FAQ FINDER, about 20 percent of the users gave this type of feedback. We expect that the improved interface in FAQ FINDER 2.0 will increase the frequency of user feedback on system performance.

References

- Buckley, C. 1985. Implementation of the SMART Information Retrieval [sic] System, Technical Report, 85-686, Computer Science Department, Cornell University.
- Collins, A. M., and Quillian, M. R. 1972. How to Make a Language User. In *Organization of Memory*, eds. E. Tulving and W. Donaldson. San Diego, Calif.: Academic.
- Cutting, D.; Kupiec, J.; Pederson, J.; and Sibun, P.

1992. A Practical Part-of-Speech Tagger. In Proceedings of the Third Conference on Applied Natural Language Processing. Trentano, Italy: Association of Computational Linguistics.

Ward, G. 1993. MOBY Part-of-Speech II. Computer file. Arcata, Calif.: Grady Ward. Available from G. Ward, 3449 Martha Court, Arcata, California.

Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, Calif.: Morgan Kaufmann.

Kulyukin, V.; Hammond, K.; and Burke, R. 1996. Automated Analysis of Structured On-Line Documents. Presented at the AAAI Workshop on Internet-Based Information Systems, 5 August, Portland, Oregon.

Lang, K. L.; Graesser, A. C.; Dumais, S. T.; and Kilman, D. 1992. Question Asking in Human-Computer Interfaces. In *Questions and Information Systems*, eds., 131-165. Hillsdale, N.J.: Lawrence Erlbaum.

Miller, G. A. 1995. WORDNET: A Lexical Database for English. *Communications of the ACM* 38(11): 39-41.

Quillian, M. R. 1968. Semantic Memory. In *Semantic Information Processing*, ed. Marvin Minsky, 216-270. Cambridge, Mass.: MIT Press.

Salton, G., and McGill, M. 1983. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill.

Souther, A.; Acker, L.; Lester, J.; and Porter, B. 1989. Using View Types to Generate Explanations in Intelligent Tutoring Systems. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, 123-130. Hillsdale, N.J.: Lawrence Erlbaum.



Robin Burke is a research scientist in the Department of Computer Science at the University of Chicago. He received his M.S. from Yale University and his Ph.D. from Northwestern University's Institute for the Learning Sciences. His research interests include intelligent information access on the World Wide Web, intelligent assistance for database browsing, and the role of strategic knowledge in retrieval. His e-mail address is burke@cs.uchicago.edu.



Kristian Hammond received his Ph.D. in computer science from Yale University in 1986. Since that time, he has been the director of the Artificial Intelligence Laboratory at the University of Chicago and has guided the development of intelligent agents in domains ranging from radiation treatment therapy to geometry problem solving and real-time game playing. Hammond has received support from the Defense Advanced Research Projects Agency (DARPA), the Air Force Office of Scientific Research, Rome Labs, the National Science Foundation, the Office of Naval Research, and the Whitaker Foundation. He is a fre-

quent member of the program committees of the American Association for Artificial Intelligence (AAAI) and the International Joint Conferences on Artificial Intelligence, has chaired the DARPA Case-Based Reasoning Workshop, and recently gave an invited talk on case-based reasoning at the national meeting of the AAAI. Hammond recently formed The Chicago Intelligent Information Laboratory. His e-mail address is hammond@cs.uchicago.edu.



Vladimir Kulyukin is pursuing a Ph.D. in computer science at the University of Chicago. He has an M.S. in computer science from the same university. His main research interest is intelligent information retrieval. He is particularly interested in anytime question-answering algorithms

that combine techniques of vector-space models and general-purpose semantic networks. Whenever possible, such algorithms take advantage of user-supplied hierarchies of topics and relevance feedback. His other interests include unicode, compilers for functional programming languages, and garbage collection algorithms. His e-mail address is kulyukin@cs.uchicago.edu.



Steven Lytinen is an associate professor of computer science at DePaul University. He received his Ph.D. in computer science from Yale University in 1984. His research is in natural language processing, particularly the extraction of semantic information from texts. He has also done

research in language acquisition.



Scott Schoenberg received his B.S. in computer science from DePaul University in 1997. He is currently employed as a programmer-analyst at the Institute for Learning Sciences at Northwestern University. His research interests include natural language and machine learning.



Noriko Tomuro is a Ph.D. student in the School of Computer Science at DePaul University. Her research interests include natural language processing, particularly in parsing and information-retrieval systems.