

Web Fact Seeking For Business Intelligence: a Meta Engine Approach

Abstract

Inspired by our exploration of the applicability of automated question answering (QA) technology to the task of business intelligence and general design science principle, we advocate a meta approach to the QA (fact seeking) applications since it does not rely on a single system and may provide the necessary accuracy and responsiveness. We present our prototype of a fact seeking meta engine, the first known meta QA prototype, and report its empirical evaluation. Our results support the added value of the meta approach: the performance of the combined system surpasses the underlying performances of its components. The evaluation also indicates that, on average, our meta system answers every second answer perfectly, which we believe is very encouraging for its practical applications. We also discuss the challenges that meta approach faces, how they can be addressed and what are the design implications.

1. INTRODUCTION: BUSINESS INTELLIGENCE AND AUTOMATED QUESTION ANSWERING ON THE WEB

In order to succeed in today's competitive business world, organizations must constantly keep an eye on what is happening in the environment every day to make decisions and adjust quickly to the changes. Business intelligence (BI) is the process of monitoring the firm's external environment to obtain information relevant to its decision-making process (Gilad and Gilad, 1988). It consists of a series of activities that involve identifying, gathering, developing, analyzing and disseminating information. As the web has become an important information resource for individuals and organizations, the search has become indispensable for the second step of the business intelligence processing -- information and fact gathering on the web. Web search engines are commonly used to locate information for business analysis (Chung et al., 2004; Lyman & Varian, 2000). However, they typically retrieve a large number of Web pages when receiving a keyword query only to overload business analysts with irrelevant information (Chung et al., 2004). More fine grained technologies capable of understanding BI tasks and representing their results in comprehensible format are emerging, among them is the Automated Question answering (QA) technology (Roussinov & Robles, 2005b). The goal of Question Answering (QA) is to locate, extract, and represent a specific answer to a user question expressed in natural language. A QA system takes an input such as "How many Kurds live in Turkey?" and it provides an output such as "About 15 million Kurds live in Turkey", or simply "15 million".

Keyword based web search engines like Google or Yahoo have incredibly improved their ability to find the most popular and lexically related pages to a given query by performing link analysis and counting the number of query words. However, search engines are not designed to deal with natural language questions, and they treat these questions as "bags of words". When a user types a question such as "Who is the largest producer of software?" into Google, it will treat it the same way as if the user typed "software producer largest," which will lead to unexpected results. It displays pages about the largest producers of dairy products, trucks, and "catholic software", but not the answer that the user was expecting. Even if the correct answer is among the search results, it still takes some time to sift through all the returned results and locate the correct answer among all the other results.

QA systems are gaining increased popularity and they are playing a more important role in the current ubiquitous computing arena. It is more natural for people to type a question, such as "Who wrote King Lear?", rather than to perform queries such as "wrote OR written OR author AND

King Lear”. Precise, timely, and factual answers are especially important when the communication channel is limited. A growing number of Internet users are using mobile devices such as Internet enabled cell phones, which do not have the luxury of a large screen space. Military, first-responder, and security systems frequently put their users under such time constraints that every additional second that is spent browsing through search results may put human lives at risk. Finally, visually impaired computer users, who are currently under-served, simply cannot enjoy the vast quantity of information that is available on the Web, since they cannot glance through the pages of snippets that are returned by search engines. The available reader software and refreshable Braille screens do not offer enough bandwidth.

The enabling technologies (Automated Question Answering and Natural Language Processing), have made big strides in recent years. Many systems capable of answering questions expressed in a natural language have been introduced and made available online. For example, AskJeeves, a public company that positions itself as the pioneer of Web Question Answering, was recently acquired by IAC/InterActiveCorp for \$1.9 billion (Claburn, 2005), a price comparable even with the total stock value of Google. The interest in both companies is not surprising considering that the online advertising market is positioned to exceed the television industry (The Economist, 2005). Currently Google and Yahoo rival the combined prime-time ad revenues of America’s three big television networks, ABC, CBS and NBC.

In the next section, we review the state of the art technology involved in the online fact seeking (question answering) to demonstrate that it has matured enough to become an attractive investment opportunity and subject of interest to not only computational linguists but also information systems practitioners. We also present our observations from the surveys and interviews that we conducted with business IT professional with respect to their expectations and awareness of fact-seeking/question-answering technological capabilities. By following the example of meta search engines on the Web (Selberg & Etzioni, 1995), *we advocate combining several fact seeking engines into a single “Meta” approach*, which is elaborated in the section four. We discuss the benefits of the Meta approach, the challenges that it faces and the design implications. Finally, we present the prototype that we have developed to illustrate our suggested approach and its empirical evaluation.

2. AUTOMATED WEB QUESTION ANSWERING: TECHNOLOGY REVIEW

The National Institute of Standards and Technology (NIST) has been organizing the annual Text Retrieval Conference (TREC) since 1992, in which researchers and commercial companies compete in such tasks as document retrieval and filtering. The performance of each research team at the competition had significant impact on the government funding of their research efforts. For the last few years, the conference and the funding agencies’ priorities have shifted to novel applications, such as question answering, novelty and topic detection, summarization, and interactive Web searching. The participating systems have to identify exact answers to so-called *factual* questions (or *factoids*), such as *who*, *when*, *where*, *what*, etc., list questions (*What companies manufacture rod hockey games?*) and definitions (*What is bulimia?*). In order to answer such questions, a typical system would: (a) transform the user query into a form it can use to search for relevant documents, (b) identify the relevant passages within the retrieved documents that may provide the answer to the question, and (c) identify the most promising candidate answers from the relevant passages. Most of TREC QA systems are designed based on techniques from natural language processing (NLP), information retrieval (IR) and computational linguistics (CL). For example, Falcon (Harabagiu et al., 2000), one of the most successful systems, is based on a pre-built hierarchy of dozens of semantic types of expected answers (*person*, *place*, *profession*,

date, etc.), complete syntactic parsing of all potential answer sources, and automated theorem proving to validate the answers.

In contrast to the NLP-based approaches that rely on laboriously created linguistic resources, “shallow” approaches that use only simple pattern matching have been recently successfully tried, e.g. the system from InsightSoft (Soubotin & Soubotin, 2002) won the 1st place in 2002 and the 2nd place in 2001 TREC competitions. However, none of the best performing systems is publicly available for independent evaluation or for inclusion in a research prototype.

There are several important distinctions between QA from a fixed corpus (also called *close domain*, such as in a TREC competition) and QA from the entire Web, which is typically referred to as *open corpus* or *open domain QA*:

1) Typically, the Web has a much larger variation in answers that can be stated. This allows the Web QA fact seeking systems to look for the most explicit answers, making the task more robust at times.

2) The users of the Web fact seeking engines do not necessarily need the answers extracted precisely. In fact, we personally observed from the interaction with practitioners, that they prefer to read with the context to verify that the source is credible.

3) Web fact seeking engines need to be quick, while TREC competition does not impose any real time constraints. This places an emphasis on the simple and computationally efficient algorithms and implementations, such as simple pattern matching vs. “deep” linguistic analysis.

Those differences shape the design decision while porting and adapting existing QA techniques to the much larger context of the World Wide Web – the efforts of which there are numerous examples. AskJeeves (www.ask.com), a public company positions itself as the pioneer of Web QA. However their knowledge sources are limited to a small set of specially created databases (e.g. geographical locations). When answers are not found there, AskJeeves reroutes the question as a simple keyword query to a general purpose search engine (Teoma, <http://www.teoma.com/>). A relatively complete, general-purpose, web-based QA system, called NSIR, was presented in (Radev et al., 2002; Radev et al., 2005). Dumais et al. (2002) presented another open-domain Web QA system that applies simple combinatorial permutations of words (so called “re-writes”) to the snippets returned by Google and a set of 15 handcrafted semantic filters to verify seven possible categories to achieve striking accuracy. Their work followed the work by other researchers on using the inherent redundancy on the Web (Clarke et al., 2001).

The prototypes based on Web fact seeking technologies have been demonstrated to surpass human performance in answering trivia questions (Lam et al., 2003) and solving crossword puzzles (Castellani, 2004). Roussinov and Robles (2005b) first studied how and automated open domain (Web) question answering can help Business intelligence tasks and also locating malevolent online content for cyber security applications (Roussinov & Robles, 2005a). Their prototype has been featured in Information Week (Claburn, 2005) as one of the promising directions in the “Web Search of Tomorrow.”

3. QUESTION ANSWERING AND BUSINESS INTELLIGENCE NEEDS

This section reports our experience with business practitioners while exploring possible applications of Web QA technology. The experience has contributed to the motivation and the designed decisions behind our meta approach proposed in the subsequent sections.

We were specifically curious to find out what questions may be asked in a typical business environment, especially while performing business intelligence tasks. We distributed a set of 100 questions to each of 16 students in an MBA class. The students had at least 2 years of experience

of being employed in typically managerial or other leadership positions. The set of questions was obtained from random drawing from the TREC questions in years up to 2005 and from the Excite set. (Radev et al., 2002). Although the exact process of building TREC test sets is not specified, it is known that most of them came from real web search logs. All of the Excite set questions are from real search sessions as well and cover broad range of types and topics, e.g. from investment to entertainment. We asked the participants to select 5-10 that may be representative of business information needs based on their personal experience.

The survey was followed by several interviews with those students who volunteered. During the interviews we informally asked open ended questions about the subject experience with business intelligence and possible needs and the expectations from the QA technology. We also asked to comment on the question selections mostly to verify that the subjects interpreted their task correctly. We subsequently analyzed the set of selected questions (107 total) and our interview notes. So far, we can only informally report the following observations: 1) The proportion of factoid questions (43 out of 107) was lower than in the original distributed set, which indicates less interest to factoid questions than they are represented in TREC sets. 2) “How to” types of questions (e.g. *How do I get rid of spam mail?*) were also more popular than in the original set. The follow up interviews also indicated interest in those types of questions. However, the interviewees expressed doubts whether the technology can provide the answer in the format preferred by them: a step-by-step set of instructions on how to accomplish the task. Indeed, the only prototype that is known to support this type, although in a limited form, is the one from Roussinov and Robles (2005a). 3) List questions (E.g. *Rotary engines were manufactured by which company?*) were extremely popular (19 out of 108), which was also confirmed by the follow up interviews. This may be explained by the already known difficulty of humans to come up with comprehensive lists (perfect recall), that is why the technology is expected to come to the rescue. Although the list questions are used in TREC, they are not supported by online engines. 4) “How much” questions were also fairly represented (4 total), which is not surprising in a business “bottom-line” oriented environment. 5) Definition questions were less popular than in the original set. The interviewees explained it by the fact that the definitional questions can be easier answered by entering keywords to Google, rather than trying to solve the problem as a QA task. 6) The questions about When/Where somebody was born, When/Where/How some famous person died and how old He/She was at the time, widely represented in the recent TREC sets, were not selected at all, which indicates that they would not be asked in a typical business environment.

The important conclusion drawn from our interviews was that the trustworthiness of the source that provides the answer is very important. As one interviewee said “People tend to think that all that is coming from the web is true, but that is not the case.” The implication of this observation to QA designers is that the attributes of the source (e.g. URL) should be also visible in the interface along with the context (snippet) in which the answer was found.

4. META QA APPROACH: BENEFITS AND CHALLENGES

4.1. Building a Case for Meta Approach

The Natural Language Processing (NLP) task, which is behind QA technology, is known to be Artificial Intelligence (AI) complete: it requires the computers to be as intelligent as people, to understand the deep semantics of human communication, and to be capable of common sense reasoning. As a result, different systems have different capabilities. They vary in the range of tasks that they support, the types of questions they can handle, and the ways in which they present the answers. While looking for answers, users have to switch between several systems, and start their search all over again each time. The beginners can easily get disoriented. They do not have

adequate knowledge to realize with what system to begin for each specific task and where to go if the system fails.

By following the example of meta search engines on the Web (Selberg & Etzioni, 1995), *we advocate combining several fact seeking engines into a single “Meta” approach.* Meta search engines (sometimes called metacrawlers) can take one a query consisting of keywords (e.g. “Rotary engines”) and sends then to several portals (e.g. Google, MSN, etc.), then combine the results. This allows them to provide better coverage and specialization. The examples are MetaCrawler (Selberg & Etzioni, 1995), 37.com (www.37.com), and Dogpile (www.dogpile.com). Although, the keyword based meta search engines have been suggested and explored in the past, we are not aware of the similar approach tried for the task of question answering (fact seeking), which we pursue in this paper.

The practical benefits of the meta approach are justified by general considerations and the design science: eliminating “weakest link” dependency. *It does not rely on a single system which may fail or may simply not be designed for a specific type of tasks (questions).* The meta approach promises *higher coverage and recall of the correct answers* since different QA engines may cover different databases or different parts of the Web. In addition, the Meta approach *can reduce subjectivity* by querying several engines; like in the real-world, one can gather the views from several people in order to make the answers more accurate and objective. The speed provided by several systems queried in parallel can also significantly exceed those obtained by working with only one system, since their responsiveness may vary with the task and network traffic conditions. In addition, the meta approach fits nicely into a becoming-popular Web services model, where each service (QA engine) is independently developed and maintained and the meta engine integrates them together, while still being organizationally independent from them. Since each engine may be provided by a commercial company interested in increasing their advertising revenue or a research group showcasing their cutting edge technology, the *competition mechanism will also ensure quality and diversity* among the services. Finally, a meta engine can be *customized* for a particular portal such as those supporting business intelligence, education, serving visually impaired or mobile phone users.



Figure 1. Example of START output.

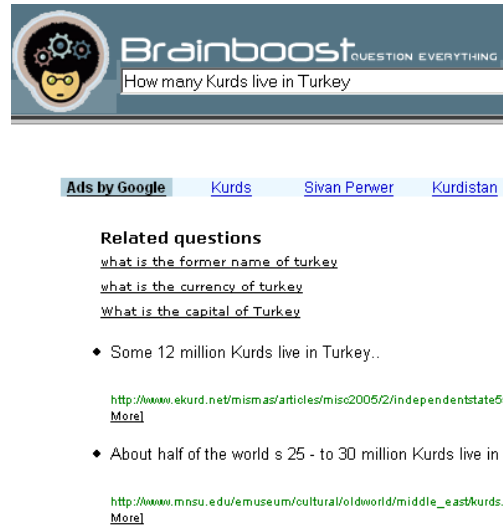


Figure 2. Example of Brainboost output.

4.2. Meta Approach Defined

We define a *fact seeking meta engine* as the system that can combine, analyze, and represent the answers that are obtained from several underlying systems (called *answer services* throughout our paper). At least some of these underlying services (systems) have to be capable of providing *candidate answers* to some types of questions asked in a natural language form, otherwise the overall architecture would not be any different from a single fact seeking engine which are typically based on a commercial keyword search engines, e.g. Google. The technology behind each of the answer services can be as complex as deep semantic NLP or as simple as shallow pattern matching.

4.3. Challenges Faced and Addressed

Combing multiple fact seeking engines also faces several challenges. First, *the output formats of them may differ*: some engines produce exact answer (START, NSIR), some other present a sentence or entire snippet (several sentences) similar to web search engines, as shown on figures 1-4. Table 1 (below) summarizes those differences and other capabilities for the popular fact seeking engines. Second, *the accuracy of responses may differ* overall and have even higher variability depending on a specific type of a question. And finally, we have to *deal with multiple answers*, thus removing duplicates, and resolving answer variations is necessary. The issues with merging search results from multiple engines have been already explored by MetaCrawler (Selberg & Etzioni, 1995) and fusion studies in information retrieval_[A2] (e.g. Vogt & Cottrell, 1999) but only in the context or merging lists of retrieved text documents. We argue that *the task of fusing multiple short answers, which may potentially conflict or confirm each other, is fundamentally different and poses a new challenge for the systems researchers*. For example, some answer services (components) may be very precise (e.g. START), but cover only a small proportion of questions. They need to be backed up by may be less precise services that have higher coverage (e.g. AskJeeves). However, backing up may easily result in diluting the answer set by spurious (wrong) answers. Thus, *there is a need for some kind of triangulation of the candidate answers provided by the different services or multiple candidate answers provided by the same service*.

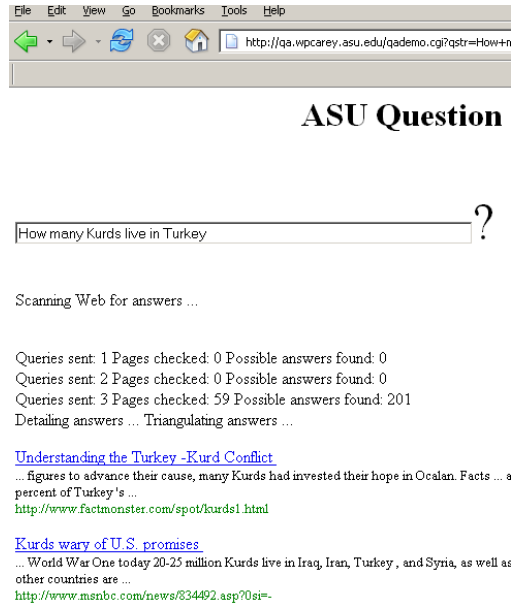



Figure 3. Example of Ask Jeeves output.

Triangulation, a term which is widely used in intelligence and journalism, stands for confirming or disconfirming facts, by using multiple sources. Roussinov and Robles (2004; 2005b) went one step further than using the frequency counts explored earlier by Dumais et al. (2002) and groups involved in TREC competitions. They explored a more fine-grained triangulation process which we also used in our prototype. Their algorithm can be demonstrated by the following intuitive example. Imagine that we have two candidate answers for the question “*What was the purpose of the Manhattan Project?*”: 1) “*To develop a nuclear bomb*” 2) “*To create an atomic weapon*”. These two answers support (triangulate) each other since they are semantically similar. However, a straightforward frequency count approach would not pick this similarity. The advantage of triangulation over simple frequency counting is that it is more powerful for less “factual” questions, such as those that may allow variations in the correct answers.

In order to enjoy the full power of triangulation with factoid questions (e.g. *Who is the CEO of IBM?*), the candidate answers have to be extracted from their sentences (e.g. *Samuel Palmesano*), so they can be more accurately compared with the other candidate answers (e.g. *Sam Palmesano*). That is why *the meta engine needs to possess answer understanding capabilities as well, including such crucial capability as question interpretation and semantic verification of the candidate answers* to check that they belong to a desired category (*person* in the example above).

Even if the answer is triangulated by the system as reliable (high probability of being correct), there is still a task of convincing the user that it is the case. The current systems typically approach the issue by representing the answer within its context (search engine snippet, sentence, or a longer passage). When dealing with meta engines and possibly large number of identical answers, one research question that still remains an open problem is *how to select the most convincing answer, among those positively triangulated?*

And finally, as has been discussed in “meta crawlers” applications, as the traffic through a meta prototype increases, there may be a need to obtain permission from the commercial portals used as answer services in order to ensure that the advertisement revenue is not “stolen”. This is typically accomplished by preserving all the advertising links.

Figure 4. Example of ASU QA

