Towards a Complete System for Answering Generalized Subjective Questions using Internet-Based Information

Barry Webster, Maryam Abdirad

Abstract— Many people use generic search engines to find answers/results for their search terms. However, in many cases, there is no effective way to find appropriate answers to subjective questions.

The objective of this paper is to investigate the best way to find the best answer for subjective questions posed through online search, and put this process on the path to automation. In order to meet this challenge, it is necessary to understand the methodology for obtaining and selecting online search results. This paper presents a discussion of a suggested process for Question Answering (QA) that consists of a number of phases, including: (1) selection of a question, (2) selection of a search engine, (3) execution of a search query, (4) searching selected websites for evidence, (5) processing the discovered evidence, and (6) delivering a final answer.

As a preliminary case study, four questions were queried using Google. From the list of websites returned for each question, 10 related websites were selected for further examination, and conceptual data potentially relevant to answering each question were extracted from each of the selected websites. The collected data were analyzed, including the resolution of conflicting data, and a final answer for each question was returned. The results returned were very reasonable, demonstrating that additional development and automation of this process can improve the current QA Process.

***Index Terms*— Question Answering, Search Engines, Concept Mining, Website Selection**

I. Introduction

TODAY when people look for answers to questions, they often use the Internet to find those answers [1]. It is important for them to be able to obtain correct answers for their search interests. However, some of the questions they want to have answered are very subjective, making it extremely difficult to definitively determine what the "correct" answer is. This is where the process of Question Answering (QA) comes into play.

Manuscript received June 13, 2016; revised June 28, 2016.

Maryam Abdirad is a M.S student in the Department of Engineering Systems at the Florida Institute of Technology, Florida, USA (e-mail: Mabdirad2013@my.fit.edu).

Barry Webster is a faculty member in the Department of Engineering Systems, Florida Institute of Technology (e-mail: bwebster@fit.edu).

QA has been defined in various (but related) ways, including: "Question Answering (QA) is a multidiscipline field of computer science that involves information technology, artificial intelligence, natural language processing, knowledge and database management and cognitive science that automatically answer questions posed by humans in a natural language" [2]. "A Question-answering system searches a large text collection and finds a short phrase or sentence that precisely answers a user's question" [3]. In short, QA systems attempt to solve the problem of subjectivity and extract a final answer for these kinds of questions or search terms.

Research into QA has had a long history. As far back as the 1960s, a QA system called BASEBALL answered questions regarding baseball statistics, and another system called LUNAR answered questions about the analyses of rock samples returned by the Apollo moon missions [4]. There is a large extant volume of published studies describing the role of QA [2][3], and research into this area continues into the present. For example, AskMSR is a QA system that uses an architecture which includes the techniques of query reformulation, n-Gram mining, filtering, and n-Gram tiling [5][6]. Another work proposed a system called Mulder, which follows these steps to find an answer: question parsing in search-engines, question classification, query formulation, choosing a search engine, answer extraction, and selecting the answer by a voting procedure [7].

Predictive Annotation is another QA technique that attempts to locate the best answers by analyzing questions and ranking the selected answers [3]. In a similar work, a system was designed that uses the following mechanism as a framework: parse the query, and select web sentences via search engine. Another system uses feature generation and ranking between collections of answers [8].

Zhiping Zheng reported a new QA search engine called AnswerBus that "is an open-domain question answering system based on sentence level web information retrieval" [9]. In this paper, the suggested method has four steps: 1) selection of search engines and formation of engine-specific queries based on the question, 2) retrieval of documents found at the top of the search result lists, 3) extraction of sentences from the retrieved documents that may contain answers to the questions, and 4) ranking of answers and returning the highest ranked to the user [9].

ISBN: 978-988-14047-1-8
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

In 2014, Swami Chandrasekaran and his coworker commented regarding the Watson QA system that the system was capable of determining answers through acquired knowledge rather than by using prepared answers [10]. A seminal study in this area is the work of Mile Pavlić that discusses a QA system that forms part of a larger system based on what is called a "Node of Knowledge" conceptual framework for knowledge-based system development [11].

These research efforts developed functional QA systems, but there remain some gaps that need to be filled in. This is particularly true in cases where the question to be answered is highly subjective. That is, the answer to the question is not something that is grounded in hard fact, and cannot be verified or proven. Rather, the answer is going to be based on the opinion of one or more individuals. It may be true that the answer has some basis in fact, but ultimately it is the result of how available facts are interpreted. In such cases, value judgments are required, and such judgment calls are difficult enough for humans to effectively make, much less computers.

This paper presents foundational work that was done towards development of an effective automated QA system specifically designed to answer subjective questions. The work was intended to serve as a proof-of-concept, demonstrating that the proposed process could reliably produce reasonable answers to highly subjective questions. This foundation could then be used to develop a QA system by automating (and improving upon) each stage of the foundational process (each of which having already been demonstrated to be workable).

In the following sections, the methodology used in each phase of the research conducted is discussed. The results of the research work are then discussed, and finally conclusions drawn from the research are presented, along with discussion of future work that can be done.

II. Research Methodology

The general process for conducting QA is very straightforward. Regardless of implementation, certain operations need to be performed, which are: 1) formation of the question to be answered, 2) gathering of information from the Internet relevant to answering the question, and 3) processing of the information that was gathered in order to determine a final answer to the question. As evidenced by the systems discussed in the previous section, the ways in which these basic operations can be implemented and/or augmented are myriad. For the purposes of the research conducted for this paper, a slightly modified form of the basic operations was selected, which is shown in Fig 1.

As can be seen, the selected process was essentially the basic operations, with some additional details filled in. A question is chosen, followed by selection of a vehicle for conducting the search for information related to the question. A search is conducted, and information resulting from the search is collected. The total amount of information to be considered is reduced to a manageable size, then that remaining collection is searched more thoroughly for specific evidence related to producing an answer to the question. Any evidence found is then processed to determine a final answer, which is then returned to the questioner.

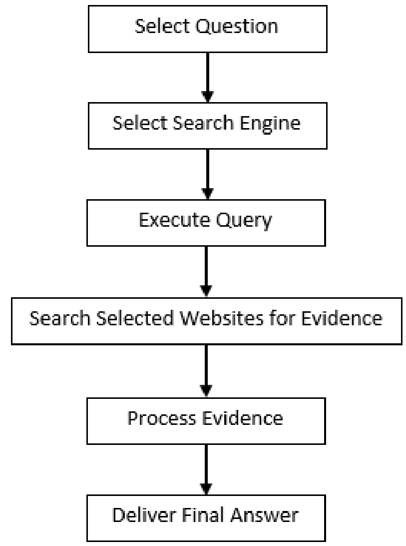

Fig 1 – QA Process Used for Current Research

In order to assess the selected process for its suitability as an automatable QA process, and to evaluate some ways in which such automation could potentially be implemented, a case study was formulated. A description of how the case study was used to assess the process follows.

A. *Selection of Questions*

The first step in the process was to develop a set of test questions which would constitute the case study. Whatever questions would be included in this set needed to be highly subjective, since that was precisely the type of question that this study was intended to answer. The questions would also need to be somewhat controversial, since the intent was to answer questions for which there would likely be a sizeable amount of available information on the Internet advocating different answers. However, even though the intent was to be able to answer complicated questions, it was not the intent to produce complicated answers, only to produce answers that were reasonable given the available information. Thus, the questions had to be answerable in a simplistic manner.

To accomplish these goals, an initial set of four questions was decided upon. Each question in the set was designed to represent a general topic of interest to one of the Colleges at the Florida Institute of Technology (Business, Engineering, Psychology and Liberal Arts, and Science). It was also the case that each question could be answered with a simple "yes" or "no". The questions in the set thus became:

1. "Should the federal government have bailed out companies like AIG in the financial crisis of 2008 – 2009?" (Business)
2. "Is solar energy better than wind energy as an alternative energy source?" (Engineering)
3. "Are arranged marriages more successful than marriages for love?" (Psychology and Liberal Arts)
4. "Is faster-than-light travel possible?" (Science)

ISBN: 978-988-14047-1-8
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

B. Selection of a Search Engine

The next phase of the QA process was to select a mechanism for searching the Internet to find information related to the selected questions. The most accessible way to do this was to use one of the existing general-purpose search engines, of which many are available. An initial search was conducted, using the top three most used search engines [1]. The questions from the case study set were put to each of the three search engines using the original wording, in order to get a base notion of the level of response that would be generated by each search engine. The results of these searches are given in Table 1. As can be seen, in general Google returned more results per question than the other two search engines, though it is interesting to note that the results are rather lopsided based on the question. For questions 1 and 3, Bing returned an order of magnitude more results than Google. However, for questions 2 and 4 exactly the opposite was true (and in the case of question 4, there was actually a two orders of magnitude difference in the number of results returned). The reason for this is unknown, and since it was not within the original scope of the study it was not investigated further.

Table 1 – Results from Major Search Engines (in thousands of hits)

No.	Question	Approximate Results		
		Google	Bing	Yahoo
Q1	Should the federal government have bailed out companies like AIG in the financial crisis of 2008 – 2009?	802	2,200	57.9
Q2	Is solar energy better than wind energy as an alternative energy source?	27,100	3,560	2,550
Q3	Are arranged marriages more successful than marriages for love?	1,370	10,400	663
Q4	Is faster-than-light travel possible?	37,400	745	745

The consideration then became why to use only one search engine rather than several. The rationale was primarily one of simplicity; even though different search engines produced different results, it was observed that there was a large amount of overlap, particularly in the highest-ranked results. Since this was the case, and since the evidence that would eventually be used to answer the case study questions would most likely come from the highest ranked sources, it was decided to just use one search engine.

Having made this decision, the obvious next decision was which search engine to use. The designation of a search engine as “best” is itself subjective, since the definition of what is “best” can vary substantially by searcher. Google had demonstrated, as a general rule, that it would return more results, but it was not at all definitive that more results equated to better results. However, according to studies performed in [1], Google as a rule delivers better results for answering general purpose questions. For this reason, Google was selected as the search engine for this study.

C. Execution of the Queries

Once a search engine had been selected, the next phase of the QA process was to formulate and execute the queries that would be used to attempt to locate information on the Internet that could be used to answer the questions in the case study set. The primary intent here was to feed the questions to the search engine in the same manner in which one person would ask the questions if they were talking to another person, which is how users would likely phrase their questions to an automated QA system. Since a standard search engine was being used, responses to the questions would come in the form of lists of websites.

D. Searching of Selected Websites for Evidence

The lists of websites returned from execution of the queries was, as can be seen in Table 1, extremely large. Conducting a search for evidence related to answering the case study questions in all of the returned websites would have been both infeasible and unnecessary. Such an exhaustive search would have been infeasible for the simple reason that it would have involved systematically searching for particular types of information through hundreds of thousands, if not millions, of websites. This would have rendered even a fully automatic QA system completely impractical. As an example, consider the smallest number of websites returned by the test queries using Google, which was roughly 802,000 for question 1. Even if it would have been possible to completely process one website per second, this would have resulted in the QA process taking over nine days to return an answer to just this one question. An exhaustive search of this kind would also have been unnecessary, since despite improvements in the precision and recall of general purpose search engines, it is still the case that not all websites returned by queries contain information that the searcher is actually interested in.

Knowing this, it was necessary to reduce the amount of material being searched for evidence towards answering the case study questions. This meant that the next phase in the QA process was to decide which of the websites returned by the queries would actually be searched for evidence, which in turn would necessitate a methodology for making that determination. According to a previous study conducted by the authors, it was found that a particular set of factors were preeminent in determining which websites human users would choose to look at from a list returned by a query [12]. This set of factors included: the websites appeared very early in the list, the websites were well-known to the users, and items of information shown in the list (e.g. website title, and description) included words and phrases similar to those included in the queries. Of these factors, the study found that, with few exceptions, by far the most common way that users decide which websites to look at was simply to examine the first several websites that appeared in the list [12]. Thus, it was decided that the QA process for this study would search a collection of websites found at the beginning of the list returned by the question queries, in an attempt to emulate the behavior of the human users from the previous study.

In order to do this, it next became necessary to determine how many websites in total would be searched. Some initial testing was done of the processing of individual websites to look for evidence related to answering the case study questions. This processing, which will be further discussed in later

ISBN: 978-988-14047-1-8
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

sections, took several hours to complete for lists with as few as 100 websites. In order to allow for repeated testing of the QA process in reasonable amounts of time, it became necessary to restrict the lists to 10 websites each. This kept processing times to under one hour per list.

E. Processing of Evidence from the Selected Websites

Given a list of selected websites, the next phase of the QA process called for each website on the list to be searched for evidence that could be used to produce an answer for the question at hand. In doing this, it would not be enough to just look for particular words or phrases. Those words/phrases would need to be placed within a proper context capable of being judged for applicability for answering the associated question. For example, if the question was, "Is the water in Lake Superior cold?", then it would be insufficient simply to look for the words "water", "superior", and "cold", since a query on these words could just as easily return websites containing information regarding high quality faucets. This information of course has nothing to do with the temperature of the water in Lake Superior, and would be of no use as evidence for answering the question.

Thus, more sophisticated techniques were necessary to be able to place words and phrases into a useful evidentiary context. This pointed in the direction of web mining, which can be loosely defined as a collection of techniques for finding knowledge in web pages not obvious from simple scans of the information in those web pages. Specifically, techniques would need to be used that could mine unstructured text within web pages (collectively known as web content mining). Of the currently available web mining/text mining operations, natural language processing and concept extraction best fit what was intended to be accomplished in this phase of the QA process. One technique in particular is used in both areas: sentiment analysis [13]. Sentiment analysis is a technique designed to determine how people feel about a topic, evaluating whether their opinions are positive, negative, or neutral towards that topic (and how strongly those opinions are felt) [14][15][16].

To use sentiment analysis in this study, a web service called AYLIEN was used. The AYLIEN service is "a package of Natural Language Processing and Machine Learning-powered tools for analyzing and extracting various kinds of information from text and images" [14]. It is available either as a callable application programming interface (API) or directly at AYLIEN's website. Locations from the lists of selected websites could be passed to AYLIEN, where the sentiment of the information within the websites would be determined in terms of polarity (positive, negative or neutral) and confidence level (0.0 to 1.0). For polarity, a positive value indicated that the content of the website had been determined to have a favorable tone, a negative value indicated an unfavorable tone to the website's contents, and a neutral polarity indicated that no overall tone could be determined within the website's contents [14]. The confidence level was a measure of the relative certainty of the polarity assessment, with 0.0 indicating no certainty at all, and 1.0 indicating absolute certainty [14].

Initially, the sentiment analyzer was tested by sending it controlled blocks of information, where the correct sentiment was known from the content being sent. After a parameter adjustment, the analyzer was consistently responding with the correct polarities, and was doing so at high confidence levels. This provided reason to believe that the analyzer was functioning correctly, and would provide reasonable assessments of the websites in the lists for the case study questions.

F. Delivery of a Final Answer

Since the questions selected for the case study were highly subjective and controversial, it was expected that for each question, some websites in the associated search list would generate positive sentiment analysis polarity results, while others would generate negative polarities. If and when such discrepancies occurred, it would be necessary to have a method for reconciling the conflicts such that a final answer could be produced. A variety of techniques for accomplishing this were examined, some of which were very simple, and some of which were quite complicated.

Unfortunately, it was also true that different techniques could arrive at different conclusions with regard to how the conflicting evidence should be reconciled. This meant that a determination had to be made as to which technique produced the best overall results. However, this determination could not be easily made. In the case of factual questions, a technique that produces the correct, or more correct, answer can safely be said to be better than a technique that does not. For example, if the question is, "What is 1 + 1?", then a technique that returns the answer "2" would be better than one that returns the answer "0". Even when a precisely correct answer cannot be given, it may be possible to make a determination of the best technique. If the question is, "What is the square root of 2?", then a technique that returns the answer "1.414" is better than one that returns the answer "1.4". Both answers are in some sense correct, but the former is more correct than the latter by virtue of its greater precision.

This study, though, was dealing with questions of opinion, not questions of fact. Thus, the issue was which of the examined techniques reconciled to the best opinion, and here was where the problem lay. If the answer to a question is based on evidence which is subjective, then there is really no way to say that one possible answer is better than another. For instance, suppose the question is, "Are blue flowers prettier than red ones?" This question is subjective, and a value judgment. One cannot prove that the correct answer is either yes or no using objective means such as a mathematical proof or even a set of empirical observations (which would not constitute a full proof but which would provide objective data favoring a particular answer). One cannot even rely on preponderance of the evidence in such cases. If there happen to be more people who prefer blue, then there is likely to be more evidence indicating that the answer to the question should be yes, and of course the reverse is also true.

Even if the answer to a subjective question is based in part on objective evidence, the fact that the question is subjective means that the objective evidence is still being interpreted differently by different evaluators. Provided the interpretations are logically consistent, then different interpretations can each be reasonable. Fundamentally, then, one cannot say definitively that any one possible answer to a subjective question is better than another.

Because of this, it was decided that there was no convincing rationale for using any of the more complicated methods for reconciling conflicting evidence for the questions in the case study. Rather, it was determined that the simplest technique should be used. Consequently, a basic "majority rules" was used; that is, the answer that had the highest count of websites with a particular polarity won. This way, if six websites from the list returned a positive polarity, and four returned a negative polarity, the final answer that the QA process would return, would be yes.

This left the issue of what to do if the count of polarities resulted in a tie. In such cases, again the choice was to use the simplest procedure, which in this instance was to sum the polarity confidence values. The polarity having the greatest sum of its respective confidence values would then be used as the final answer. It should be noted here that summing the confidence values could be used as a standalone technique. The reason it was not in this study was that work early on with the sentiment analysis consisted only of the polarities. It was only later on that the confidence level data were collected, and at that time it was decided not to throw out the early results but to augment the later results.

III. Results

The entirety of the QA process described in the previous section was performed for each of the questions in the case study, and then repeated multiple times to ensure that the results were being consistent. The final results are shown in Table 2.

Table 2 – Summary of Results According to Majority Rule

No.	Positive	Negative	Neutral	Majority Rule	Final Result
Q1	3	5	2	5	Negative
Q2	4	4	2	?	?
Q3	7	1	2	7	Positive
Q4	2	4	4	?	?

As can be seen in the table, the QA process was able to draw a conclusion for questions 1 and 3 using majority rule for reconciling conflicting evidence. However, a tie resulted for questions 2 and 4, meaning that summation of the polarity confidence levels was needed to break the tie. These results are shown in Table 3.

Table 3 – Confidence Level Sums for Questions 2 and 4

No.	Positive	Negative	Neutral
Q2	3.96	3.89	2
Q4	1.94	3.47	3.99

By comparing the sum of polarity confidence levels, it can be seen that for question 2, positive polarity has the highest confidence level sum, at 3.96. For question 4, the highest sum was for the neutral polarity, at 3.99. Given these results, it was now possible for the QA process to return final answers for all four questions. These answers are shown in Table 4

Table 4 – Final Answers Returned for Case Study Questions

No.	Question	Answer
Q1	Should the federal government have bailed out companies like AIG in the financial crisis of 2008 – 2009?	No
Q2	Is solar energy better than wind energy as an alternative energy source?	Yes
Q3	Are arranged marriages more successful than marriages for love?	Yes
Q4	Is faster-than-light travel possible?	Neutral

Once the answers were obtained, an examination was conducted of the source data in an effort to determine the reasonability of the manner in which the QA process obtained and analyzed its data, and thus whether the particular answers returned were reasonable given what the QA process had to work with.

For question 1, it was clear that there were a lot of strong opinions regarding the U.S. federal government's bailout of private companies following the financial crisis of 2008-2009. A majority of the opinions in the most popular websites (and thus the ones that were found at the top of the query list and included in the search) were very much against the bailout.

For question 2, the source data revealed a split in the opinions regarding which type of alternative energy is better. Both the solar power and wind power industries have large numbers of advocates, and this was reflected in the fact that there was a tie in the polarity counts.

For question 3, the majority of articles found in the websites on the search list seemed to have been written by persons composing apologetics for arranged marriage. Thus, the tone of those articles was very positive with respect to the opinion that arranged marriages are more successful than non-arranged.

Finally, for question 4 there was again a split, but this time between negative and neutral opinions rather than positive and negative. Many science fiction articles were of course in favor of the possibility of faster-than-light travel, but not many of these made it into the search list. Many scientific articles either took the position that faster-than-light travel was impossible or that they were not sure, and hence the tie in the polarity counts.

Another item of note was that the selection of questions for the case study turned out to be very successful. As mentioned, the intent was to choose questions where there would be a variety of opinions, and potential controversy. Somewhat ironically, the question on arranged marriages, which was expected to be very controversial, actually ended up being the most one-sided in terms of the polarity counts, with 70% of the websites on the search list coming back with positive polarities. Other than that, though, none of the other questions resulted the websites on their search lists having even a simple majority for a particular polarity. Even when the polarity sums were invoked, the difference between the two highest sums was no more than about one half of a single point across all websites in the search list. This demonstrated that there was indeed a great

ISBN: 978-988-14047-1-8
ISSN: 2078-0958 (Print); ISSN: 2078-0966 (Online)

deal of variety of opinion within the questions, which provided a good challenge for the QA process.

The examination of the source data showed that the QA process had correctly assessed the sentiment of the websites in the search lists with respect to the nature of the questions given. Thus, the answers returned for each question reasonably matched what the websites in the search lists were saying. Of course, as discussed previously these were subjective questions, which meant that there was no way to definitively state a "correct" answer. Nevertheless, the answers returned by the QA process accurately reflected the data given, so the answers returned were at the least reasonable, and consistent with the most popular opinions given by humans answering the same questions.

IV. Suggestions for Further research

As discussed at the beginning of this paper, the primary intent of this study was to provide a basis for automating the QA process as given. Since the completion of the study, an initial version of a completely automated system for following the QA process described herein has been developed and is functional. The next steps would be to make improvements to the individual elements within the process. Some suggestions for doing this are: 1) developing/using a metasearch engine integrating multiple single-source search engines which accounts for and filters out overlap of returned websites, 2) pre-processing of questions to maximize usefulness of returned websites, 3) expanding the types of questions that can be answered, being able to separate different types of material (e.g. article text vs. user comments), 4) using other methods of text analysis in lieu of or in addition to sentiment analysis, and 5) improvement in the ability to identify subtleties within text (e.g. humor, sarcasm, and double entendre).

V. Conclusion

This study demonstrated a complete QA system, using commercially available components, that could generate reasonable answers to subjective, controversial questions in a short amount of time. A selected set of case study questions was submitted to the system, and a query was formulated for each question using the Google search engine. The results of the queries were then narrowed down to a manageable number of selected websites. The selected websites were searched thoroughly and processed using algorithms for sentiment analysis provided by a web service. Any conflicts in the sentiment analysis were then resolved, and a final answer for each question was returned.

The results of the study were very promising, and though conducted on a limited basis, the process could be applied on a broader scope of questions using a variety of techniques. This is what was set out to be accomplished: not to develop a complete and finalized QA system, but to provide a functional framework that could be improved upon in each of its phases independently. This is the intent moving forward.

References

[1] K. Purcell, J. Brenner, and L. Rainie, "Search Engine Use 2012," *PEW Res. Cent.*, no. February, p. 42, 2012.

[2] P. Gupta and V. Gupta, "A Survey of Text Question Answering Techniques," *Int. J. Comput. Appl.*, vol. 53, no. 4, pp. 1–8, 2012.

[3] J. Prager, E. Brown, A. Coden, and D. Radev, "Question-Answering by Predictive Annotation," in *Proceedings of the 23rd annual international ACM SIGIR conference on research and development in information retrieval - SIGIR '00*, 2000, pp. 184–191.

[4] D. Jurafsky and J. H. Martin, "Question Answering and Summarization," in *Speech and Language Processing*, 2nd ed., Prentice Hall, 2008, p. 1024.

[5] E. Brill, S. Dumais, and M. Banko, "An Analysis of the AskMSR Question-Answering System," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing - (EMNLP)*, 2002, vol. 10, pp. 257–264.

[6] M. Banko, E. Brill, S. Dumais, and J. Lin, "AskMSR: Question Answering Using the Worldwide Web," 2002.

[7] C. Kwok, O. Etzioni, and D. S. Weld, "Scaling Question Answering to the Web," *ACM Trans. Inf. Syst.*, vol. 19, no. 3, pp. 242–262, 2001.

[8] H. Sun, H. Ma, W. Yih, C.-T. Tsai, J. Liu, and M.-W. Chang, "Open Domain Question Answering via Semantic Enrichment," in *Proceedings of the 24th International Conference on World Wide Web*, 2015, pp. 1045–1055.

[9] Z. Zheng, "AnswerBus Question Answering System," in *Proceedings of the second international conference on Human Language Technology Research*, 2002, pp. 399–404.

[10] S. Chandrasekaran and C. DiMascio, "Create a natural language question answering system with IBM Watson and Bluemix services," *IBM developerWorks*, 2014. [Online]. Available: http://www.ibm.com/developerworks/cloud/library/cl-watson-films-bluemix-app/. [Accessed: 01-Jan-2016].

[11] M. Pavlic, Z. D. Han, and A. Jakupovic, "Question answering with a conceptual framework for knowledge-based system development 'Node of Knowledge,'" *Expert Syst. Appl.*, vol. 42, no. 12, pp. 5264–5286, 2015.

[12] M. Abdirad and B. Webster, "Website Selection in Internet Search Engines," Melbourne, FL, Unpublished, 2014.

[13] G. Miner, D. Delen, J. I. Elder, A. Fast, T. Hill, and R. A. Nisbet, *Practical Text Mining and Statistical Analysis for Non-structured Text Data Applications*, 1st ed. Elsevier, 2012.

[14] AYLIEN, "AYLIEN | We Bring Intelligence To The Web" [Online]. Available: http://aylien.com. [Accessed: 01-Jan-2016].

[15] B. Snyder and R. Barzilay, "Multiple Aspect Ranking using the Good Grief Algorithm," in *Proceedings of NAACL HLT*, 2007, pp. 300–307.

[16] K. Khan, B. Baharudin, A. Khan, and A. Ullah, "Mining opinion components from unstructured reviews: A review," *J. King Saud Univ. - Comput. Inf. Sci.*, vol. 26, no. 3, pp. 258–275, 2014.