

# How to Effectively Train IBM Watson: Classroom Experience

<sup>1</sup>Syed Shariyar Murtaza, <sup>1</sup>Paris Lak, <sup>1</sup>Ayse Bener, <sup>2</sup>Armen Pischedotchian

<sup>1</sup>Data Science Lab, Ryerson University, Canada

<sup>2</sup>IBM IBM Watson, Littleton Lab, 550 King Street, Littleton, MA, USA

<sup>1</sup>{syed.shariyar, parisa.lak, ayse.bener}@ryerson.ca, <sup>2</sup>apischdo@us.ibm.com

**Abstract.** Watson is a question answering system that uses natural language processing, information retrieval, knowledge interpretation, automated reasoning and machine learning techniques. It can analyze millions of documents and answer most of the questions accurately with varying level of confidence. However, training IBM Watson may be tedious and may not be efficient if certain set of guidelines are not followed. In this paper, we discuss an effective strategy to train IBM Watson question answering system. We experienced this strategy during the classroom teaching of IBM Watson at Ryerson University in Big Data Analytics certification program. We have observed that if documents are well segmented, contain relevant titles and have consistent formatting, then the recall of the answers can be as high as 95%.

## I. INTRODUCTION

IBM Watson is a question answering system that won Jeopardy challenge in 2011 against former winners Brad Rutter and Ken Jennings [3]. Jeopardy is a famous American television game show that tests participants on their general knowledge and Watson turned out to be more knowledgeable than humans in that show.

Watson is built to imitate human brain in observing, interpreting, evaluating and making decisions. More specifically, it models human cognitive framework for decision-making and leverages cognitive computing. Cognitive computing systems learn through expressions that are more natural to human and they continue to learn as new scenarios evolve and new expression emerge [7].

Watson uses many innovations in natural language processing, machine learning and question answering domain [17]. From natural language processing, the system uses corpus processing and expansion [9], semantic evaluation and parallelism [6] as well as the techniques to find clues in questions. Watson also uses

search technologies, such as Lcuene [10] and Solr [14] to find and generate candidate answers [4]. According to IBM more than 100 different natural language processing techniques are used by Watson [8]

IBM Watson is ideal for the domains and applications in which large amount of information should be ingested, explored and decision-making is complex [2]. The system has been used in various domains, such as healthcare, banking, legal services, tourism and etc. IBM Watson was opened to the world in 2014. Its community now encompasses 280 commercial partners and thousands of students and developers, which are creating up to 3 billion monthly API requests on Watson [15]. For example, CafeWell is an app that provides heart patients with guidelines. Talkspace used Watson to best match the licensed therapists by providing real-time insights about patient's personality, thinking style and emotional stress. Go Moment is a mobile guest management app that provides front desk concierge staff in addressing the needs of guests [15].

We are also now part of the Watson community, and introduced Watson to students of Big Data Analytics certification program at Ryerson University in winter 2015. Our experience showed that teaching students IBM Watson could be tricky. It may result in frustration of students because Watson would not return correct answers if it is not trained properly. Initially, we experienced that the training done by students on questions and answers did not always result in correct answers during testing. However, we developed a set of guidelines with the help of an IBM engineer. This always resulted into accurate answers during testing.

In this paper, we present the best practices that we observed for training the IBM Watson during our class. These best practices are presented as guidelines which may be used to effectively train IBM Watson and successfully deploy it in the field. These may be helpful to anyone who wants to leverage the power of

cognitive computing by using IBM Watson and not just the classroom students. We hope that using these guidelines, the Watson community may build better applications.

The rest of the paper is organized as follows. In Section II, we provide the background knowledge about IBM Watson. In Section III, we provide the strategy that we have developed for its effective training and successful deployment. In Section I, we discuss the results that we achieved using our strategy during the classroom teaching. In Section II, we summarized the related work. In Section III, we conclude the paper.

## II. BACKGROUND: AN OVERVIEW OF WATSON

There are some key differences between question answering systems and search engines. In search engines a set of keywords are considered as an input and the output would be a list of documents ranked in the relevancy order of the keywords and their popularity on the World Wide Web. However in question answering systems, the input is a question in natural language. The first step in these systems is to understand the question and extract the key elements of the question. The second step is to provide a precise answer according to the relevancy to those key elements [12].

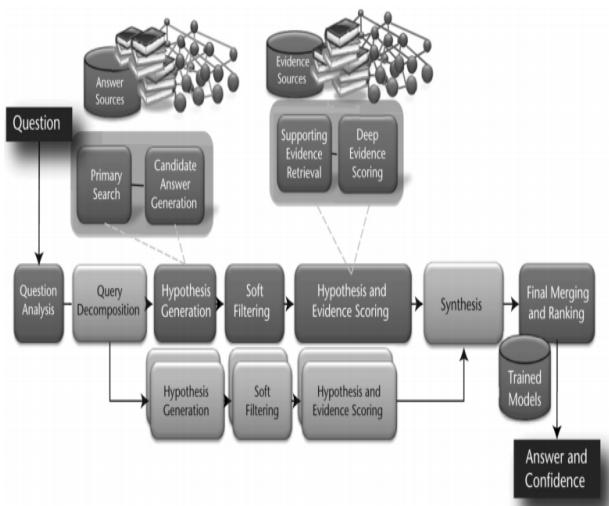


Figure 1. Watson Architecture [3]

Watson is a question answering system and uses natural language processing techniques and machine learning algorithms to comprehend questions and

return their answers. Watson is founded on an open source framework called Unstructured Information Management Architecture (UIMA) [16]. UIMA allows for different analysis engines to act as stages, where the analyzed, searched and scored outputs are saved in the form of a common data structure. Watson's architecture is also composed of stages and is illustrated in Figure 1 [3]. Below we explain this architecture with examples. The examples are adopted from Rob High's (IBM Fellow) presentation on IBM Watson [11].

First stage in Watson's architecture is Question Analysis. In this step, questions are decomposed into keywords and the answer type for a question is detected [1]. For example, consider the question: "*In which Michigan city in 1894, C.W. Post created his warm cereal drink?*" Watson will divide it into keywords as: "1894, C.W. Post, created, warm cereal, drink Potasum, Michigan, and city." It will also identify the answer type as "Michigan city" for this question because the question is asking about the city in Michigan. Answer types in a question answering system are usually detected base on a predefined taxonomy of name-entities. This taxonomy is extracted from a set of initial sample questions, or it is created separately. A grammar parser is used in Watson to extracts answer types along with the semantic and syntactic focus of the question. Similarly, a named entity detector identifies common and proper nouns and the entity type (person, place, etc.) and a predicate structure detector identifies the relationship between subjects and objects. All these become the input for the second stage.

The second stage is about the generation of hypothesis. In this stage Watson uses those keywords and search millions of documents for relevant passages. This primary search is done by using Lucene [10], Solr [14] and Prismatic [11]. If potential passages are not found at this time, they will not be discovered later. In the case of the above example, Watson found 5 relevant documents and 30 passages. In this stage candidate answers are identified in the search results by ranking the searched documents and the passages. Ranking is done by determining the keywords, answer types, etc. in the document titles (including a variety of title variants and expansions) and text in the passages [3].

Candidate Answers	Evidence Feature Scores				
	Doc Rank	Pass Rank	Ty Cor		
General Foods	0	1	0.1		
Post Foods	2	1	0.1		
Battle Creek	1	2	0.8		
Will Keith Kellogg	3		0.1		
Grand Rapids			0.9		
1895		0	0.0		

Figure 2. Evidence scores for candidate answers

The third stage is about answer scoring. In this stage, the candidate answers (selected passages from documents) are assigned scores using many answer scoring analytics. For example, type coercion is used to determine whether an answer is of a specific type or not. Figure 2 shows the scores assigned to the candidate answers using type coercion (Ty Cor), document ranking and passage ranking scoring mechanism (as mentioned in previous step). The first column in Figure 2 shows the titles of the documents and the remaining column shows the scores based on different analytics. In the case of type coercion for our example question, a score would be assigned by evaluating if the passage in the document is referring to a city or not, such as: (a) the score for the document “General Food” as a “city” would be 0.1; and (b) the score for the document “Grand Rapids” as a “city” would be 0.9. In addition to type coercion, hundreds of such scores are assigned to the candidate answers in Watson [5].

The fourth stage deals with final merging and ranking (FMR) of candidate answers. This step is done on IBM’s SPSS servers. Figure 2 shows that each row for a candidate answer represents a feature vector with columns as features. These feature vectors for each candidate answers are passed on to logistic regression based machine learning models. These models assign probabilities, called confidence, to each of the candidate answers. The candidate answers are then ranked in the decreasing order of their confidence and presented to users as the final answers. The answer with the highest confidence is considered as the best answer [5].

The logistic regression based machine learning models used in the fourth stage are actually developed from the sample of questions and answers provided to Watson during the training phase. During training,

users need to create well formatted documents with titles and passages as answers. Those documents are uploaded to Watson and then the users need to assign sample questions to different answers in the documents [7]. For each of the questions, Watson goes through the stages 1-3 as mentioned above but this time Watson knows the correct answers and labels each feature vector as “Yes” or “No” based on the answer-mapping done by the users [5]. Using this collection of feature vectors for candidate answers of sample questions, Watson trains the logistic regression based models which are then used to answer different questions of users not in sample questions (during testing) as mentioned in Stage 4 above [5].

### III. STRATEGY TO WORK WITH WATSON

During classroom teaching of IBM Watson, we experienced best practices that could be used to effectively train IBM Watson. In this section, we shall discuss them. We first discuss the best practices for creating documents that can be effectively ingested in Watson in Section 4.1. Second, we discuss the mapping of questions to answers—i.e., creation of labels for training of Watson in Section 4.2. Third, we discuss testing strategy for IBM Watson.

#### A. Creating Documents: The Best Practices

The most important step in successful training and deployment of IBM Watson is to create well-formed documents. Well-formed documents lead to better indexing at the time of ingestion, hence more efficiently analyzed by the search engines such as Lucene [10] and Solr [14]. Following are the seven rules that we identified for the creation of well-formed documents:

1. A good document should be divided into sections. A section should have a descriptive title, and the passage in the section should have an answer (answers) that matches the title. The answer must be descriptive and brief—too long answers will not get attention of users.
2. There should be some keywords in the passage in a section (i.e., the answer) that match the title of section too.
3. The title of the section should match a question or many questions. For example, if the keywords in the question coincide with the

keywords in the title or the passage that is following the title, then Watson will return relevant answers with high confidence.

4. Avoid documents with tables, especially nested tables. HTML pages with advertisement and Javascript running in the background or PDF files laden with images with no clear headings and corresponding text.
5. The document formatting should be consistent. For example, when using Microsoft Word, try to use “Heading 1” or “Heading 2” style for your section titles, and “Text”, “Body Text” or “Normal” style for your answers. It is important to consistently use the same heading style throughout the document.
6. A better practice is to have every student (or content creator) upload a separate document to have consistent formatting within each document. In this way, one can avoid inconsistencies in formatting a document. One student can upload as many documents as he/she likes.
7. Documents should not be too large or too small. Though there is no fixed criterion, a general rule that can be considered is to have 10-20 answers per document. The 10-20 answers (i.e., 10-20 pairs of titles and passages) should map to a minimum of 10-20 questions and a maximum of as many as possible. Answers for a related topic should be grouped and put in one document.

Once documents are formatted, they are uploaded to Watson using the Watson Corpus Manager [7]. All these questions should be entered into Watson using Watson’s Training Manager, and title-text pair (answers) should be uploaded separately using Watson’s Corpus Manager [7].

Below we provide examples of well-formed and poorly formed documents (titles and passages) based on the above rules. The examples are taken from the documents created by students at Ryerson University during the teaching of IBM Watson in the class.

#### 1) Example 1: Well-formed Document

##### “Cost of Living in Ontario

*Your cost of living in Ontario or monthly expenses in Ontario depend on many things, such as your city, housing type, where you shop, your lifestyle and more. Some household spending surveys estimate that the average family will spend C\$55,000 - C\$65,000 per year. Different stores charge different prices for food, clothing, and other things you need to live. You may want to ask your relatives, friends, or coworkers where they shop for good prices. There are also special services or stores in many communities where you can find items for free or a low price. When looking at the cost of living or monthly expenses, you'll want to budget carefully. Many community agencies can help you find information about living expenses and budgeting. To find help in your area, go to the following URL:*

*<http://settlement.org/ontario/housing/living-in-ontario/cost-of-living-in-ontario/what-are-the-monthly-living-expenses-for-a-family-of-3>.*

This example is a good example because the title is descriptive, and the passage matches with the title. Also, it can be mapped to multiple questions: “what is the cost of living in Ontario?” and “what are the monthly expenses in Ontario?”

#### 2) Example 2: Well-formed Document

##### “Replace a lost, stolen or damaged Driver’s License

*You can replace a lost, stolen or damaged license at a Service Ontario Centre. A fee of \$25 will be collected. You will need to bring a personal identification document (verifying your name, date of birth and signature). You will be given a temporary driver’s license and your new driver’s license will arrive in the mail within 4 to 6 weeks.”*

This is another good example because it is short, title is descriptive, and keywords in the title and the passage match each other. Also, it can be mapped to many questions, such as “how to replace a driving license?”, “how to replace the lost driver’s license?”, “how to replace a stolen driving license?”, “how to replace a damaged driving license?” and “how to get a new driving license when your driving license is lost, stolen or damaged?”

#### 3) Example 3: Poorly Formed Documents

## **Buying a Business**

*Find out what you need to know before buying a business: where to look, how to evaluate potential acquisitions, and what a fair price would be.*

### ***Leasing requirements***

*You may be inclined to begin by taking out a lease for your premises. However, leases can be one of your largest expenses. Make sure that your lease will be suitable to your business needs, in case you wish to break your lease or expand your business. A lawyer can give you advice on any pitfalls or costs that may be incurred, before you sign on the dotted line.”*

This example is a poor example because the title of the section is “Buying a Business” but the following text does not say anything about buying a business. In fact, the text itself is asking a question in the beginning. The text then starts to talk about leasing a business with an irrelevant sentence “leasing requirements” which looks like a subheading. It is very difficult to find out exactly what this text is talking about.

#### **4) Example 4: Poorly Formed Documents**

***“Check current operating dates below. Please note that some parks have multiple listings for campgrounds or seasonal operations.***

Park Name	Camping Dates	Open and Close Dates (day use)
<u>Aaron</u>	May 15, 2015 to September 20, 2015	May 15, 2015 to September 20, 2015
<u>Algonquin</u>	Day Use Only	April 1, 2015 to March 31, 2016
<u>Algonquin - Achray</u>	April 24, 2015 to October 12, 2015	April 24, 2015 to October 12, 2015
<u>Algonquin - Backcountry</u>	Interior - April 1, 2015 to March 31, 2016	April 24, 2015 to October 12, 2015

*Different parks have different timings.”*

This is one of the poorest examples. The title is not truly expressing the information in the text that follows it. It would be better to have separate titles and text for each of the parks. It also contains tables which are usually ingested after concatenating text by Watson. Tables should be avoided when creating documents.

#### **B. Training Watson: Mapping Questions to Answers and Deploying the Corpus**

The next step, after uploading documents, is to add questions in Watson related to the documents using the

How do you choose a ministry-approved driver education course?

Match a Similar Question   Match an Answer   Specify an Answer

Select a formatted answer for your question to create a new question cluster. If an answer has questions matched to it, select a question to join an existing cluster. [Learn more.](#)

Restrict by document: Enter a file name. Text search: To refine the list of documents, enter keywords and press Go. Or press Shift+Enter to start a new search. Go

Group10 DriverLic VehicleReg : Scheduling a knowledge/written test? : Choosing a ministry-approved driver education course.  Open Source Document  Selected

1 questions answered

Group10 DriverLic VehicleReg : Scheduling a knowledge/written test? : Government approved driving schools

1 questions answered

Group7 SettlementOrgAllScrapeXML

Choosing a ministry-approved driver education course:  
Not all driving schools offer ministry-approved beginner driver education courses. Look for schools that offer "MTO-approved beginner driver education courses", or "ministry-approved beginner driver education courses". Beginner driving schools in Ontario are regulated by MTO. Choosing a school is a personal choice. Choose carefully. You may want to check the list of driving schools that offer ministry-approved beginner driver education courses at the bottom of this page. You may wish to ask your friends or co-workers to recommend a driving school to you. You could consider asking to meet with the instructors who will be conducting the classroom instruction and the behind-the-wheel training. Not all driving schools offer ministry-approved beginner driver education courses. Look for schools that offer "MTO-approved beginner driver education courses", or "ministry-approved beginner driver education courses".

Figure 3. Selecting answers in the Watson Training Manager

Training Manager in Watson [7]. Once you enter questions in Watson, the relevant documents matching the answers will pop up. For example, Figure 3 shows a question: “how do you choose a ministry-approved driver education course?” When this question was entered in Watson, relevant answers (the titles and the corresponding passage) from documents appeared and the correct answer was selected.

If no passages are retrieved for your question, then perhaps the question is off topic or that a specific domain related document is missing from the corpus. You can mark the question as Off Topic or Needs Document and return to it at a later time.

IBM recommends a minimum of 400 questions per project [7]. It is common place to see projects with over 1000 questions. Users should also try to paraphrase a question, and enter it into Watson. Watson automatically discovers the similar questions, and asks the user to cluster similar questions for the same answer. The user can also add tags (or keywords) for questions related to one answer. According to our experience, adding tags increase the accuracy of answering correct questions during testing.

Furthermore, we have observed that answers should be short. It is better to add many one line answers rather than selecting individual lines for different questions in one large paragraph. IBM Watson provides this option of selecting individual lines from large paragraphs as answers but we found that the accuracy will be higher if the answers are provided separately as one line with titles.

Once all the answers in the documents are mapped to questions correctly, create and deploy the corpus [7]. IBM strongly recommends that you create a corpus not just after uploading the documents, but after pairing questions with answers and advancing their status to approved. This approach will index not only the uploaded documents but also the question and answer units. Indexed artifacts are that much more available to Lucene/Solr search engine, synonym expansion and the work that Prismatic does even this early in the process.

### C. Testing Watson

Once Watson is successfully deployed, it is the time to test Watson and the power of cognitive computing. During testing, Watson should be able to answer domain relevant questions, even if they are not in the sample questions of the training set. For example, you should not ask “who is president Obama?”, when there

is no such document containing that answer. However, the user can ask “who is president Obama?”, when the questions user entered were “who is the president of the United States?” and “who is the leader of the US?”

In addition, one may not see answers because of the spelling mistakes. For example, “what is the address of *foodbank* in Vancouver?” Watson could give a wrong answer of Stephenville Emergency food bank. This is because it may not pick up spelling mistakes and typos. However, on rewriting the question to “what is the address of food bank in Vancouver?”, it gives the right answer. In fact, Watson is extremely powerful in interpreting the questions. If the questions are not entered at all but the proper titles and passages are present in the well-formed documents, it can still provide the relevant answers.

A user can also provide feedback to Watson during testing and improve the accuracy of their answers.

## I. EVALUATION OF OUR STRATEGY

We validated our strategy to train and test Watson on students of the Big Data Analytics certificate program at Ryerson University. There were about 43 students in our class. We created five projects in Watson and divided the 43 students into 10 groups and each group was assigned to one of the five projects in IBM Watson. (Five projects are a maximum number of projects a class is usually allowed to create for free in Watson.) Students were asked to follow the guidelines described in Section III to train and test Watson. In particular high emphasis was placed on well-formed document guidelines. The five projects were created on the following topics:

- Cheap and healthy eating in Canada.
- Immigration, citizenship and real estate information for Canada.
- Legal services in Canada.
- Taxes, investment and startup information for Canada.
- Tourism and transportation information for Canada.

To evaluate our approach, we measured the recall, accuracy and precision of five projects. The measures are present in Watson and we used Watson to determine them. These measures are defined as follows. Recall is defined as the percentage of relevant answers that Watson returns for questions. Accuracy is defined as the percentage of questions for which the

top answer is correct. Precision is defined as the percentage of questions for which the top answer is correct and has a confidence greater than 70%.

Our five projects resulted into the recall of 85-100%, accuracy of 51-75% and the precision of 30-50%. The precision was low in few projects (less than 50%) but the low precision does not mean that the correct answer was not returned. It only meant that the confidence value returned by Watson was low. Infact, we observed that the top most answer is mostly correct but the confidence level generated by Watson could be low. Recall from Section II, confidence level is simply a probability generated by logistic regression model and it is used to rank the answers. In our five projects the top most answer was correct up to 75% of the times irrespective of the confidence value (i.e., accuracy). Also 100% recall means that the correct answer is always present in the top three to four answers. Nonetheless, the precision can be increased by giving feedback to Watson after deployment during testing.

On further investigation of the documents and question-answer mappings, we found out that some students did not exactly follow the guidelines. For example, some students took the options for selecting individual lines in the paragraphs for many questions and uploaded very large not well-formed documents. Many of these questions did not return proper answers during testing of Watson. Due to large number of students in our class, it was difficult to keep track of individual students.

In one project, students exactly followed the guidelines and it resulted into 100% recall, 75% accuracy and 50% precision. These students also developed a mobile application, and this project has been selected for commercialization by Ryerson University. It focuses on eating cheap and healthy foods while living in Canada. Users can type in questions about different products or foods, and can get the information about ingredients of them instantaneously.

In order to evaluate our strategy further, we provided these guidelines to the students of the Big data analytics program in the following two terms (in spring 2015 and summer 2015) after the first term (in Winter 2015). In the next two terms there were 30 and 40 students respectively. We observed that in all the three terms, students achieved similar recall, accuracy and precision. Thus, these results validate that the guidelines (or the strategy) developed by us are effective in training IBM Watson.

## II. RELATED WORK

Zadrozny et al. [17] taught students at the University of North Carolina at Charlotte question answering system and simulated Watson in the classroom. They created different groups of students and assign them tasks to develop different components of Watson from query decomposition to final candidate answer ranking using logistic regression. They couldn't get as high accuracy, precision and recall as Watson because of the difference in the level of expertise of students and professionals at IBM.

A series of research papers about the workings of IBM Watson have been published in IBM Journal of Research and Development (issue 3.4) in 2012 (e.g., [3],[1],[9],[6],[5]). These research papers describe the underlying technology of IBM Watson. Each component of the Watson, shown in Figure 1, is described in detail in different research papers. The innovations belonging to individual components are also described succinctly. In Section II, we already summarized how Watson works by using these research papers.

In this paper, we have introduced the effective training strategy of Watson based on our classroom experience. This can be useful to anyone, not only students, who wants to build applications using Watson. To the best of our knowledge, this work is novel and has not been done before.

## III. CONCLUSION

Watson is an exceptional question answering technology that leverages the power of cognitive computing. It encompasses many different innovations in the field of natural language processing and question answering system. However, training IBM Watson could be tedious and if it is not trained properly then the answers may not be accurate.

In this paper, we outline the strategy for effective training of IBM Watson question answering system that we learnt during the teaching of IBM Watson. We evaluated this strategy for three consecutive semesters with approximately 125 students. The results indicate that students following the guidelines to train Watson produced 100% recall, 75% accuracy and 50% precision.

During our class, students learned the fundamentals of natural language processing system, text analytics as well as the process of question answering systems. Students first learnt the theoretical concepts and then

received the hands-on experience in Watson question answering system. This allowed them to understand how a competent question answering system is trained, deployed and tested. They developed skills in the state of the art natural language processing system. The overall curriculum is designed to have experiential learning and hands-on approach in all of the courses as well as in the capstone project course. Watson fits very well in our program as students learn through experience, and gain the skills to apply this knowledge in other courses throughout the program. This will help them in enhancing their skills easily in future as the technology evolves.

Watson is a new technology, approximately a year old, and has an exceptional potential. It is maturing and will become an excellent product with the passage of time.

#### IV. REFERENCES

- [1] A. Lally, J. M. Prager, M.C. McCord, B. K. Boguraev, S. Patwardhan, J. Fan, P. Fodor, and J. Chu-Carroll, "Question Analysis: How Watson Reads a Clue." *IBM Journal of Research and Development*, Vol. 56, No. 3.4, 2012, pp. 2:1—2:14.
- [2] B. Anbarasan, "Using the Computer Brain Cognitive Computing" *Special Issue (NCRICA-14) of International Journal of Emerging Technologies in Computational and Applied Sciences*, Vol. 7, No. 1, 2014, pp. 22-26.
- [3] D. A. Ferrucci, "Introduction to "This is Watson"." *IBM Journal of Research and Development*, Vol. 56, No. 3.4, 2012, pp. 1:1—1:15.
- [4] D. Ferrucci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, J. Prager, N. Schlaefer, C. Welty, "Building Watson: An Overview of the DeepQA Project." *AI Magazine*, Vol. 31, No. 3, 2010, pp. 59-79.
- [5] D.C. Gondek, A. Lally, A. Kalyanpur; J. W. Murdock, P.A. Duboue, L. Zhang, Y. Pan, Z.M. Qiu, C. Welty, "A framework for merging and ranking of answers in DeepQA." *IBM Journal of Research and Development*, Vol. 56, No. 3.4, 2012, pp. 14:1—4:12.
- [6] E. A. Epstein, M. I. Schor, B. S. Iyer, A. Lally, E. W. Brown, and J. Cwiklik, "Making Watson Fast." *IBM Journal of Research and Development*, Vol. 56, No. 3.4, 2012, pp. 15:1—15:12.
- [7] IBM Guide, "IBM Watson Ecosystem Getting Started Guide", V 1.7, July 2014.
- [8] IBM White Paper, "IBM Systems and Technology. Watson – A System Designed for Answers: the Future of Workload Optimized Systems Design." February 2011. [Online]: <http://public.dhe.ibm.com/common/ssi/ecm/en/pow03061usen/POW03061USEN.PDF>
- [9] J. Chu-Carroll, J. Fan, N. Schlaefer, and W. Zadrozny, "Textual Resource Acquisition and Engineering." *IBM Journal of Research and Development*, Vol. 56, No. 3.4, 2012, pp. 4:1—4:11.
- [10] Lucene: <https://lucene.apache.org/>
- [11] Prismatic: <https://github.com/Prismatic>
- [12] R. Deepak, and E. Hovy. "Learning Surface Text Patterns for a Question Answering System." in *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, 2002, pp. 41-47.
- [13] R. High, Jr., "IBM Watson, How Watson Works", power point presentation, March 21, 2015.
- [14] Solr: <http://lucene.apache.org/solr/>
- [15] T. Olavsrud, "How IBM Watson Apps are Changing 7 Industries", Chief Information Officers. [Online]: <http://www.cio.com/>
- [16] UIMA: <https://uima.apache.org/>
- [17] W. W. Zadrozny, S. Gallagher, W. Shalaby, and A. Avadhani, "Simulating IBM Watson in the Classroom." in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ACM, 2015, pp. 72-77.