

Teaching With Watson

Michael Wollowski

Rose-Hulman Institute of Technology, 5500 Wabash Ave., Terre Haute, IN 47803, USA
wollowski@rose-hulman.edu

Abstract

In this paper, we describe how we integrated the materials from the 2013 IBM The Great Minds Challenge (TGMC) - Watson Technical Edition into our Introductory Artificial Intelligence course. We describe the variety of materials made available by IBM, as well as the nature of the competition and the datasets that are at the heart of it. We detail how, where and in what form we integrated the materials into our course. We describe assignments that are based on the materials from the competition as well as additional materials we incorporated into our course. We finish by evaluating our experience in teaching with the materials as well as summarize relevant student feedback. We make recommendations for those who wish to adopt the materials.

Introduction

IBM's WatsonTM incorporates a wide variety of Artificial Intelligence [AI] techniques. Many of them are at a very high level. As such, Watson^{*} provides a wonderful framework for introducing many key AI techniques. Additionally, Watson is a modern engineering marvel, centered in our field. It stands to reason that students should be familiar with it.

In this paper, we describe how we integrated the materials from the 2013 IBM The Great Minds Challenge [TGMC] - Watson Technical Edition into our Introductory AI course. In particular, we describe the materials, the nature of the competition, how we integrated the materials into our course, our experience in teaching with the materials as well as student feedback.

This was the first time the competition was held in the United State. and as such, we see this event as the beginnings of wonderful things to come. We fondly hope that we started down a path in which Watson becomes the center, the framework of an introductory AI course.

The Competition

IBM has been holding the TGMC for several years (Wikipedia 2013), however this is only the second time the competition focused on Watson, the software that won the Jeopardy!^{TM**} exhibition match against Ken Jennings and Brad Rutter. It is the first time the challenge is being held in the U.S.

For this challenge, teams of students were asked to write software which learns a model that is used for final scoring. Watson consists of hundreds of components that gather information from many sources with various confidence factors. The final scoring component distills all this data into a ranked short-list of answers and their confidence factors. The training data was quite large and some of the Machine Learning algorithms quite complex, so much so that some of our teams exhausted the most powerful computers on campus.

IBM made available three datasets: A training set, an evaluation set and a hold-back set. The *training set* contained a column indicating those answers that were correct. The *evaluation set* omitted that column of data and consisted of different data altogether. The *holdback set* contained yet different data and was used by IBM to determine the final standings. It was not made available until the last week of the competition. While students were permitted to submit their response to the evaluation set as many times as they wished, they only had one such opportunity for the holdback set. In essence, the holdback set controlled for software that over fitted the data.

The datasets in essence contained an answer id, a question id and then a little over 200 columns of data. The training set was about 378 Mbyte large and the evaluation set about 81 Mbyte. Figure 1 shows one row from the evaluation set, to give the reader a sense of the sort of data the participants had to process.

```

400001 844363 0 0 0 0 0 0 0 1 1 1 0 0
0.945096672 0 0 0.945096672 0 0.945096672
4.545120009 0.560993043 1 0.548314409 1 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 2.710147619 1 1
0 0 0 0 0 0 0 0 0 0 0 0 0 4 20.95855331 0 0 0 0
0 0 0 0 0 0 0 0 0 0 1 1.637038231 0 0 0 1 0 0 0
0 0 0 0.9225126 0.980986303 0.982530299 0.25 0 0
0 0 34.28676904 0 0 0 0 0 0 0 0 0 0 131 6 1893 2
0 0 0 0 0.095588235 0 0 0 0 0 0 0 0 0 0 0
3.835290015 0.945096672 0 0 0 0 0 0 0 8 0 0 0 0
1 0 71.71376184 0 0 0 19.81377727 6.447623113
45.05635569 71.31775607 4.64E-04 27.04109287
16.13377953 3.825283527 29.46219792 0 0 0 1.44E-
05 0 0 0 1 1 0 0 1 1 0 0 1 0 1 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 1 0 1 1 1 1 1 0 1 0 1 0 1 1
1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 1 1 1
1 1 1 1 1 1 0 0 1 1 1 1 0 0 0 0 0 0 0 0 0 1 1 0
1 1 1 1 0 0 1 0 1 0 0 1 1 0 0 0 0 0 0 0 1 0 1 0
0 0 0 0 0 0 1 0 0 0 0 0 1 0 0 0 0

```

Figure 1: Sample row from evaluation set.

The fall 2013 running of the TGMC – Watson Technical Edition saw four universities participate. IBM set up a website used to disseminate instruction materials and to aid in the running of the competition. The competition lasted a month and was aimed at graduate students; however, we are an undergraduate institution and were permitted to participate as well. We fielded about a third of the 37 teams. Our teams did very well; we tied for 1st and also came in third and fourth. Our teams ended up confirming the results published by IBM in (Gondek 2012). The teams that placed high in the competition used Logistic regression or non-linear Support Vector Machines. The teams that did not place high did, by and large, pursue other approaches.

Overview of the Learning Materials

While IBM used the format of a competition - not something that excites us too much as our students are already very competitive - their primary interest seemed to be for students to become familiar with Watson and to be able to apply Watson to interesting problems. To this extent, IBM went out of its way to provide learning materials so as to make participation in the competition a positive learning experience.

IBM provided four kinds of materials: (i) Those designed to help in the running of the competition, (ii) videos promoting IBM's definition and vision of "cognitive systems," (iii) promotional materials

highlighting current and future application of Watson and (iv) research papers about technical aspects of Watson.

We ended up using all but a few pieces of the materials made available on their competition site. We furthermore studied additional research papers about Watson. We now describe the materials in more detail.

Competition related materials. Right from the start, it was very obvious that IBM intended the competition to be a learning experience rather than a cutthroat competition. Among others, IBM hosted two question and answer sessions for interested faculty and eventually two Q&A sessions for the student teams. All of them were very cordial and were designed to help faculty and students to understand the nature and process of the competition. Additionally, there were several videos explaining the nature of the data sets and how to interact with the IBM hosted submission site. Their site contained several forums that enabled student teams to interact with IBM and other teams. However, they did not see a lot of use.

Cognitive Computing. The competition site hosted several videos in which IBM promoted its definition and vision of "cognitive computing." In essence cognitive computing is about developing Watson like systems that assist users to perform at a higher level. These videos were particularly illuminating when pitched against a mandatory review of Searle's "Minds, Brains and Programs." (Searle, 1980)

Promotional Materials about Cutting Edge use of Watson. There were about 10 short videos and leaflets that are aimed at getting people to think about the existing and potential impact of Watson in Healthcare. When covered as a whole and with a small amount of editing, they painted an impressive picture of the need and use of Watson derived technology in Healthcare.

Technical Papers about Watson. The website provided access to three technical papers about Watson, which originally appeared in the No. 3/4 issue of the 2012 IBM Journal of Research and Development (Ferrucci et al. 2012). All of them were useful and were assigned as individual readings. These technical papers are, in our opinion, the highlight of the materials made available. As a matter of fact, we asked teams of students to present additional articles from the above mentioned journal. While they did not aid in the performance of the competition, they did a wonderful job helping us understand further technical details of Watson.

Integration into Course

We made the decision to participate in the TGMC – Watson Technical Edition about three weeks before the beginning of our fall term. It was a bit of a leap of faith as the competition was aimed at graduate schools and our

course is the first (and only) course on Artificial Intelligence. Additionally, we are on the quarter system, as such, we meet four times a week for 50 minutes. The competition began on October 7, 2013. Our term began on September 5, as such we had a little over a month, or 18 sessions to give our students an introduction to some basic AI techniques and to get them ready for the competition. This meant that materials were not always covered in a linear fashion.

Since participation in the competition meant that we could not study all of the materials ahead of time, we decided that the best way to provide for a positive learning experience was to make the participation a research experience in which we study different machine learning approaches. This proved to be a good decision and we ended up doing a lot of science.

As a sign of the good faith efforts by IBM, there were negotiations on the team sizes. We assumed that the team size was unrestricted, only to learn that IBM intended the team size to be at most two students. In response to one of the Faculty Q&A sessions, IBM increased the team size to four.

We will now detail how and where we integrated the materials made available through the competition. They are listed in chronological order by the day the materials were covered.

Day 5. The first foray into Watson took place on day 5. We asked our students to review the paper “Introduction to *This is Watson*,” an overview of Watson (Ferrucci 2012a). This paper highlights the different phases of processing a question and gives an insight into the information to which Watson has access. To ensure a thorough reading of this paper, our students were asked to use the review guidelines as published by “computingreviews.com.” We spent an entire class session on introducing Watson. In addition to a class discussion of the assigned paper, we watched a brief YouTube^{TM**} video about Watson’s performance on Jeopardy! as well as a video made available through the competition site, entitled “IBM Watson: The Science Behind the Answer” (IBM 2011). Additionally, we studied selected slides from the talk “Building Watson - Beyond Jeopardy” by David Ferrucci in which he provides a good amount of detail about the process of developing Watson (Ferrucci 2012b).

Day 8. On this day, we discussed the paper entitled “Question analysis - How Watson reads a clue” (Lally et al. 2012). This is a paper made available through the competition site. Students were asked to review the paper in preparation for class.

Day 13. We spent day 13 discussing the task that the competition software was to accomplish. We discussed the key paper relevant for the competition, entitled “A framework for merging and ranking of answers in DeepQA” (Gondek et al. 2012). Additionally, we studied

figure 1 from (Ferrucci 2012a) to understand where the final merging and ranking takes place as well as its role in the overall performance of Watson. We watched and discussed the video “Understanding the Datasets” (IBMWatsonSolutions 2013e). In order to get a sense of the machine learning techniques that we would likely use, we had a quick look at the Wikipedia^{TM**} pages of all the machine learning techniques mentioned in (Gondek et al. 2012). We also spent a small amount of time watching the videos entitled “How Watson answers a Question” (IBMWatsonSolutions 2013c) and “Watson and Machine Learning” (IBMWatsonSolutions 2013d). Both of those videos were made available through the IBM Competition site. Finally, we pointed our students to the Weka Data Mining Software (Weka 2013).

The competition lasted from October 7 to November 8 corresponding to our days 19 to 36. By the time the competition began, we covered the following subjects:

- Agents
- The Turing Test
- Uninformed and heuristic search
- MinMax and alpha/beta pruning
- Various forms of knowledge representation: frames, scripts, production systems, expert systems, certainty factors, Bayesian reasoning.
- An introduction to machine learning

To ensure that our students were making good progress, we decided to split the competition into two rounds.

Day 14. Pre-competition exercise. By this time, the team memberships had been determined, partially by student request and partially based on our prior knowledge of our students’ abilities. Primary factors that influenced our decision included academic and organizational skills. To ensure that our students were ready for the competition, we asked each team to accomplish the following: (i) decide on a programming language with which to implement their software or select an existing package, (ii) implement a file reader to read csv files, (iii) decide on a preliminary Machine Learning approach, (iv) decide who calls in to the student Q&A session and (v) to ensure that the team is signed up on the IBM Competition site.

Days 19 to 27. Round 1. For the first round, we asked our teams to select three machine learning approaches and use them to train a model for the training set. Some teams implemented some of the approaches and some teams used existing packages that they found online. They were then to score their model on the evaluation set and fine-tune the software based on the feedback. The objective of this assignment was to cast a wide net and get a sense of the strengths and weaknesses of each approach. In addition to developing software, our students were asked to study and analyze each approach. As they experimented with an approach, they were asked to:

- Consider preprocessing the datasets so as to weed out poor data.
- Consider having several passes through the dataset.
- Consider that today's standout AI projects are an example of fine engineering.

Additionally, each team was asked to write a report in which they detailed their experiments. The report was to be about a page per approach, although most teams produced considerably longer reports. We expected that for each approach, the team ran several experiments, so as to fine-tune various learning parameters and to observe their effect on the accuracy of the model. For each experiment, we asked the teams to indicate the setup, any tuning parameters, the results, as well as a subjective evaluation of the experiment. We asked them to include graphs, figures and tables as appropriate as well as a conclusion in which they evaluated the strength of each approach and in which they speculated on potential other approaches.

The grading was specified as follows. “Your grade will depend on the number and quality of the experiments you conducted as well as the quality of your evaluations. The results of the experiments are not that important at this time. For now, we just want to get a sense of the lay of the land. We plan to assign a grade of an A to those papers in which each of the three approaches has been carefully evaluated by conducting about three well thought out experiments, such that subsequent experiments build on the insights gained from the prior experiments. We plan to assign a grade of a B for papers where the analysis is not as careful or where one of the three approaches is poorly evaluated. We plan to assign a grade of a C for those papers where two approaches are poorly evaluated. We plan to assign a grade of a D for those papers where all three approaches are poorly evaluated.”

During days 19 to 27, among others, we introduced the following Machine Learning techniques potentially relevant to the competition:

- Decision trees (1.5 days)
- Genetic Algorithms (1 day)
- Feed-forward neural networks (2 days)
- Case-based reasoning (1 day)
- K-means clustering (1 day)
- Logistic regression (1 day).

An integral part of the round 1 assignment was for the teams to further study their chosen approaches.

Day 21. One week into the competition, we studied and discussed the proposed and actual use of Watson in Healthcare by studying the following videos and leaflets:

- “Watson at Work” video. (Gold 2013)
- “IBM Watson & Wellpoint: A Progress Report” video. (IBMWatsonSolutions 2013a)
- “IBM Watson: The Voice of the Physician on the Future of Healthcare” video. (IBMWatsonSolutions 2012)

- “IBM and Memorial Sloan Kettering Cancer Center” video. (Connor 2012)
- “IBM Watson Demo: Oncology Diagnosis and Treatment” video. (IBMWatsonSolutions 2013b)
- “Watson in Healthcare - Oncology Diagnosis” case study. (IBMWatson 2013a)
- “Watson in Healthcare - Utilization Management” case study. (IBMWatson 2013b)

Day 28. Huddle. Halfway through the competition, we spend an entire class session discussing the results of the round 1 Machine Learning experiments. Each team was asked to give a 5 minute presentation of what worked and what did not. During the second part of class, we discussed general parameters on how to proceed during the second round. This proved to be a great day for science. Our teams did a wonderful job so far and had some excellent results. It became clear that certain approaches, such as Logistic Regression and non-linear Support Vector Machines would do well and some other approaches such as Neural Networks were not. This confirmed the results that IBM published in (Gondek 2012). However, we also had some promising results from approaches such as Genetic Algorithms and even Case-Based Reasoning.

Days 28 to 36. Round 2. For round 2, we wanted our students to continue to explore reasonable Machine Learning approaches. As such, we decided and announced that the standing in the competition will have no effect on their project grade whatsoever with one exception. Teams that placed top spots in the competition will get an A for their project. We felt comfortable with this decision, because our teams did a wonderful job conducting experiments and writing them up for round 1, so much so that it became obvious that teams that were going to do well on the competition were going to do produce excellent reports.

The round 2 assignment asked students to continue with the work they had done with the IBM Watson data. Based on our discussion of day 28, we encouraged them to consider combining approaches. This included the use of ensemble methods as well as hybrid approaches.

We implored our teams to continue conducting experiments that are based on what they have learned so far. We encouraged them to finely craft their software and experiments, and to submit the software to the IBM Competition site on a daily basis. Furthermore, we asked our students to meticulously document each experiment. Our teams were asked to provide a second write-up, similar to the first. This time, we were more precise, asking that they document their experiments in a table containing times, dates, tuning parameters and results as evidenced by a submission to the evaluation dataset. Additionally, they were asked to provide a justification for each experiment that they conducted as well as a subjective evaluation of each.

Days 29 to 32. *In-depth study of Watson.* While the videos, promotional materials, research papers and the competition itself provide wonderful insights into Watson, we were interested in learning more about the technical details of Watson. As such, we assigned additional papers about Watson, papers that were not made available on the competition site, but were independently obtained from the Watson team. Teams of students were asked to study and then give a 15 minute class presentation on the following papers:

- Deep parsing in Watson (McCord et al. 2012)
- Structured data and inference in DeepQA (Kalyanpur et al. 2012)
- Relation extraction and scoring in DeepQA (Wang et al. 2012)
- Textual evidence gathering and analysis (Murdock et al. 2012)
- Automatic knowledge extraction from documents (Fan et al. 2012)
- Finding needles in a haystack: Search and candidate generation (Chu-Carroll et al. 2012)
- Making Watson fast (Epstein et al. 2012)

This proved to be a great exercise. Universally, our students wished that the papers had even more detail.

Day 33. On this day, we took a last look at Watson, by pinning the information we learned from the paper presentations on the overview of the Watson architecture as shown in figure 1 of (Chu-Carroll et al. 2012).

Following those presentations, we spent four days on introducing planning. Among others, we studied another engineering marvel in our field: self-driving cars. Our students were asked to review the paper entitled “Autonomous Driving in Urban Environments: Boss and the Urban Challenge” (Urmson C. et al. 2008). We subsequently discussed this paper in class and watched some relevant videos.

Days 37 to 39. During these days, our teams were asked to present their work on the Machine Learning experiments. The teams were asked to spend about 15 minutes presenting their results. They were instructed to use slides for their presentations.

Day 40. Take home final. We felt that it would benefit our students to reflect on the potential impact of Watson. To this extent, we asked them to write an essay; this was the take-home final. The instructions were as follows: “We spent a good amount of time and effort to learn about IBM’s Watson. Watson-like technologies have the potential to change the world, just as the World Wide Web did. Imagine a world 15 years from now in which devices like Watson are in everyday use. While it is hard to predict the future, please speculate how a world like this might look like. Take into consideration what we learned about Watson, to discuss the feasibility of those devices. Justify your reasoning.”

Some of our students proposed some rather clever applications.

Instructor’s Evaluation

We have been meaning to integrate Watson into our course ever since its legendary performance on Jeopardy!. During the fall of 2012, we studied and read most of the papers from (Ferrucci et al. 2012). When we learned about the competition, we took a leap and decided to participate. Even though the Machine Learning techniques that were used in the competition were at a much lower level than we had hoped, we enjoyed integrating the competition into our course.

We were fortunate to be able to study in detail the nature and techniques of a modern engineering marvel in our field. We used a wide array of learning materials to do so: reviewing and discussing research papers, the development of software central to the workings of Watson, videos, promotional materials, class discussions presentations of the results of their experiments, and the conducting and documenting of experiments in the first place.

In this process, we were able to introduce our students to the research process and reinforce their research skills.

Student Feedback

We asked our students to give feedback about the participation in the competition through the existing instrument of our end-of-term anonymous course evaluations. About half of the students who stated something about the competition were very excited about it and about half were not.

Some of the students who were excited about the competition spent an extraordinary amount of time and effort on it, to the extent that they obtained accounts on the most powerful computers on campus and got them so busy that they received concerned emails from our system administrator.

Some of the students who did not like the competition stated that the dataset was just too uninteresting. They felt that the problem to be solved was too abstract. It appeared that some of our students literally wanted to see the outcome of the processing. We set up the participation in the competition as a research project with a lot of time built in for independent exploration. This did not go over well for some of our students; they wished more structure and guidance for the competition.

In the end of the day, IBM was correct in suggesting that graduate students participate in this competition. It appears that some of our students were just not ready for the research nature of the competition and since the competition was an integral part of the course, they were

by and large unhappy about the course. The saddest comment we read was the following: “Two months of Watson can get pretty boring.”

In response to the student course evaluations, we decided to remove the Watson competition from our introductory AI course and instead place it, as well as additional materials about Watson in a follow up course. We will still study Watson, however in a more traditional setting by studying and presenting the research papers, as was done on days 29 - 32.

Conclusions

Participating in the challenge was a bit of a risk, as by the time we had to make the decision to participate, we did not know much about the datasets nor about how to best incorporate the materials into our 10 week term. As described, we had mixed results.

Except for the video on the financial applications, we used all of the materials made available through the competition site. In addition, we used several of the technical papers as published in (Ferrucci et al 2012).

We wished that the AI techniques to be used were at a higher level. We would like to see a situation where students employ various inference techniques, higher levels of machine learning as well as knowledge representation.

We are hoping that in addition to the existing datasets, IBM makes available that data which shows the correspondence between the answer IDs and actual text of the answers. This would enable participants to explore higher level techniques that are geared towards answer merging.

Most importantly, we would like to go above the competition and develop an AI course built around Watson. Perhaps, this can be the result of a concerted effort by instructors across several institutions.

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of Jeopardy Productions, Inc., Wikimedia Foundation or Google in the United States, other countries, or both.

References

- Chu-Carroll J. et al. 2012. Finding Needles in a Haystack: Search and Candidate Generation. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 6;1 – 6;12.
- Connor 2012. IBM and Memorial Sloan Kettering Cancer Center. <https://www.youtube.com/watch?v=JKHG7PylZ8I&list=PL07C352F537A38A4C>
- Epstein E.A. et al. 2012. Making Watson Fast. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 15;1 – 15;12.
- Fan J. et al. 2012. Automatic Knowledge Extraction from Documents. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 5;1 – 5;10.
- Ferrucci D.A. 2012a. Introduction to *This is Watson*. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 1;1 – 1;15.
- Ferrucci D.A. 2012b. Building Watson - Beyond Jeopardy. <http://forum.stanford.edu/events/2012slides/plenary/2012plenaryDavidFerrucci.pdf>.
- Ferrucci D.A. et al. 2012. This is Watson. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4.
- Gold S. 2013. Watson at Work. <https://www.youtube.com/watch?v=yf1bZQRoU7Q&list=PLIGr-DeSOoFYbCId0tvP8cEmXWzKM89B4>
- Gondek D.C. et al. 2012. A Framework for Merging and Ranking of Answers in DeepQA. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 14;1 – 14;12.
- IBM 2011. IBM Watson: The Science Behind an Answer. <http://www.youtube.com/watch?v=DyW04zksfXw>.
- IBMWatson 2013a. IBM Watson Helps Fight Cancer with Evidence-Based Diagnosis and Treatment Suggestions. http://www-03.ibm.com/innovation/us/watson/pdf/MSK_Case_Study_IMC14794.pdf
- IBMWatson 2013b. IBM Watson Enables more Effective Healthcare Preapproval Decisions Using Evidence-Based Learning. http://www-03.ibm.com/innovation/us/watson/pdf/WellPoint_Case_Study_IMC14792.pdf
- IBMWatsonSolutions 2012. IBM Watson: The voice of the Physician on the Future of Healthcare. https://www.youtube.com/watch?v=_0i2P8atngg
- IBMWatsonSolutions 2013a. IBM Watson and WellPoint: a Progress Report. https://www.youtube.com/watch?v=_TNE5uMH_q4
- IBMWatsonSolutions 2013b. IBM Watson Demo: Oncology Diagnosis and Treatment. https://www.youtube.com/watch?v=8IGJ0h_jAp8&list=PLIGr-DeSOoFYbCId0tvP8cEmXWzKM89B4
- IBMWatsonSolutions 2013c. How Watson Answers a Question. <https://www.youtube.com/watch?v=-TpdT0wcM88&list=PLIGr-DeSOoFYbCId0tvP8cEmXWzKM89B4>
- IBMWatsonSolutions 2013d. Watson and Machine Learning. https://www.youtube.com/watch?v=Vgb2bzS_b2Y&list=PLIGr-DeSOoFYbCId0tvP8cEmXWzKM89B4
- IBMWatsonSolutions 2013e. Understanding the Datasets. <https://www.youtube.com/watch?v=mGZNGU5ju7I&list=PLIGr-DeSOoFYbCId0tvP8cEmXWzKM89B4>
- Kalyanpur A. et al. 2012. Structured Data and Inference in DeepQA. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 10;1 – 10;14.
- Lally A. et al. 2012. Question Analysis: How Watson Reads a Clue. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 2;1 – 2;14.
- McCord M.C. et al. 2012. Deep Parsing in Watson. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 3;1 – 3;15.
- Murdock J.W. et al. 2012. Textual Evidence Gathering and Analysis. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 8;1 – 8;14.
- Searle J. 1980. Minds, Brains and Programs, *Behavioral and Brain Sciences* 3 (3): pp 417–457,
- Urmson C. et al. 2008. Autonomous Driving in Urban Environments: Boss and the Urban Challenge. *Journal of Field Robotics* 25(8), pp 425–466.
- Wang C. et al. 2012. Relation Extraction and Scoring in DeepQA. *IBM Journal of Research and Development*, Vol.: 56, No: 3/4, pp 9;1 – 9;12.
- Weka 2013. Weka 3: Data Mining Software in Java. <http://www.cs.waikato.ac.nz/ml/weka/>
- Wikipedia 2013. IBM's The Great Mind Challenge. http://en.wikipedia.org/wiki/IBM's_The_Great_Mind_Challenge.