Proceedings of the Twenty-Ninth International
Florida Artificial Intelligence Research Society Conference

Studying Watson Inside Out — A Cognitive Systems Course

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

Abstract

We describe a course in which students train an instance of Watson and develop an application that interacts with the trained instance. Additionally, students learn technical information about the Jeopardy! version of Watson and they discuss a future infused with cognitive assistants. In this paper, we provide learning outcomes and course assessment items. We provide detailed course materials and advice for instructors interested in teaching such a course. The advice is in the form of best practices, a description of a successful use case and an evaluation of our experience teaching this course.

Introduction

IBM's Watson captured the imagination of people during the memorable match on Valentine's day of 2011. Personal assistants such as Siri, Google Now and Cortana are in widespread use. The technologies and prospects of those technologies are of interest to academia and industry alike. We developed a course that enables students to study Q/A techniques in particular those underlying IBM's Watson from an applied perspective. The primary objectives of this course are to learn about the

- Power and limitations of IBM's Watson
- Technology behind Q/A in general and Watson in particular

The major activity in this course is centered on training an instance of Watson. IBM recommends a training set consisting of 300 – 400 answers and 4 - 5 questions for each answer, amounting to 1200 – 1600 question answer pairs. Going through this exercise is a valuable experience in determining the power and limitations of one of the most powerful question answering system. We were told that the Jeopardy! version of Watson was trained over a period of one year.

Since Watson is relatively new, it is not uncommon for students to develop an application that can be brought to market. In addition to the excitement that exposure to the process of taking a product to market brings, there is an opportunity for students to learn about and practice the creative process of developing new or improved products. This course is designed for students who want to learn about:

- the technology of modern Q/A systems
- entrepreneurship

Through technical papers and an IBM produced MOOC about Watson (IBM 2016), there is an opportunity to study key technologies underlying advanced Q/A systems.

There is a broadly held belief that cognitive assistants are in our future. Studying Q/A systems like Watson enables us to critically examine some of the claims made about them.

In this paper, we motivate, describe and evaluate a course in which we study the technologies driving Watson, develop an application for Watson and train an instance of it.

Learning Outcomes

The learning outcomes for our course are as follows. Students who successfully complete this course should be able to:

I. Develop applications for Watson.
II. Develop potentially successful applications in Cognitive Systems.
III. Explain the architecture of Watson.
IV. Evaluate future directions of Cognitive Systems.
V. Use the Innovation Canvas to justify potentially successful products.
VI. Explain various ways in which to develop a product idea.

Prerequisites

The prerequisite for the first offering of this course was our upper level artificial intelligence course. There are two benefits to this choice: (a) One can dive much deeper into the technology of Watson and (b) the technical skills as well as the maturity of the students are higher, enabling a much more focused discussion of Watson and cognitive systems.

Most of the students of the first offering were advanced CS undergraduate students. Two students were advanced CS minors.

We are scheduled to teach this course again during the spring of 2016. For this offering, we lowered the prerequisite to our data structures course and junior standing. There were several reasons for this decision. (i) To increase the audience for this course. Among others, it would place a course with a significant entrepreneurial component earlier in the curriculum. (ii) The software development component is small. The development of an iPhone app does not require writing a lot of code; as such students with reasonable software development skills should be able to do well in the course. (iii) The primary activity in this course is to develop training data and train an instance of Watson. This is something that a mature college student should be able to achieve. (iv) The videos about the inner workings of Watson are designed for a general technical audience and do not require the knowledge of an AI course.

Assessment

The following types of assignments will be given in this course and will be used to assess the course objectives.

A. Slides, presentation and write-up justifying a project idea.
B. Slides, presentation and write-up of a project proposal.
C. Presentation of ways in which innovators/artists develop an idea.
D. Training data and training of our instance of Watson.
E. Project software.
F. Documentation of project through a technical paper, a slide presentation and a videotaped demo.
G. Reviews of papers, videos and presentations about Watson and cognitive systems.
H. Take-home final justifying a future of cognitive systems.
I. Participation in class discussions.

Assignments and Activities

We will now provide insight into the assessment items and how they are used to satisfy the learning outcomes. To give the reader as sense of the timing, we will provide scheduling information as well.

In order to understand the scheduling of the activities, the reader should realize that we operate on the quarter system; we meet four times a week for ten weeks. Each class session lasts 50 minutes.

There are four kinds of assignments: reviews and presentations (20%), the project proposal phase (20%), the project phase (50%) and the take-home final (5%), contributing to the course grade as indicated. The remaining 5% of the course grade are determined by participation in the course, largely through participation in class discussions.

Reviews and Presentations

The reviews are designed to learn about (i) the inner workings of Watson as well as (ii) the future impact of technologies like Watson. The learning materials for the 1st objective are select lectures from a MOOC about Watson (IBM 2016). To ensure productive classroom discussions, students are asked to provide a 3/4 page review of each of the lectures, using the review guidelines from computingreviews.com (Computing Reviews 2016).

The materials for the second objective will be assigned from a website maintained by IBM (IBM Cognitive Systems Institute 2016). This site is designed to discuss and learn about a future with cognitive assistants. The materials assigned will depend on student interest as well as availability of timely materials. As such, we do not specify concrete resources in the schedule below.

Below are the assignments and their due dates.

- Day 3: IBM Lecture 1: Introduction to IBM Watson
- Day 7: IBM Lecture 2: Deep Question Answering Architecture
- Day 11: IBM Lecture 3: Semantic Integration and Machine Learning
- Day 15: IBM Lecture 4: Natural Language Processing
- Day 19: IBM Lecture 5: UIMA at IBM Watson
- Day 23: IBM Lecture 6: Structured Knowledge in IBM Watson
- Day 27: Video or readings about cognitive systems
- Day 31: Video or readings about cognitive systems
- Day 35: Video or readings about cognitive systems

To expose our students to an entrepreneurial mindset, we assign several presentations in which our students present the results of their investigation of how entrepreneurs develop new product ideas. They are interspersed with regular class sessions.

Project Proposal

The project proposal phase consists of a four week process in which students develop an idea for an application, research it, obtain feedback from classmates, and attempt to recruit classmates to join their proposal. If necessary, we

conduct a class vote on the domain to be used for training Watson.

It should be noted that due to the good amount of training, we will train only one instance of Watson. For more information about these constraints, see the sub-section about the project itself. We make specific project constraints available at this time to assist students with proposing a good domain and application.

The proposal phase is designed to learn about ways in which people creatively think about product ideas, how to obtain fail-fast feedback/validation of an idea, how to research a market opportunity, how to make a good argument for a project proposal and to engage in the back and forth of negotiating a commonly accepted project proposal.

Timing of Project Proposal Phase

We established the following deadlines for the proposal phase.

- Day 2: Review of various IBM Watson projects
- Day 4: Idea for an application
- Day 8: Revised application idea
- Day 10: Innovation Canvas for Amazon.com
- Day 12: Project proposal

Most of the assignments are probably fairly obvious. We will explain those that may need some background.

The purpose of the reviews of various IBM Watson projects is to study the current state of the art of Watson infused Q/A applications. The projects are assigned from the Watson University Competition YouTube video collection (IBM Watson 2016) and other sources.

The *innovation canvas* is a tool designed by several colleagues here at Rose-Hulman (Kline 2016, Ahmed 2014). It is a tool designed to holistically think about several aspects of developing a product while paying attention to a central value proposition. To get to know the canvas, our students are asked to fill in the innovation canvas the way Jeff Bezos might have filled it in a long time ago. My students are asked to read "The Inner Bezos," (Bayers 1999) which provides sufficient information to make this a good exercise.

As part of their project proposal, our students were asked to fill in a blank innovation canvas. In addition to the specification of their proposed application, students were asked to pay attention to the value that might be generated by their proposed app as well as propose possible revenue streams and estimate costs. This aspect of the project proposal is designed to hone our students' entrepreneurial skills.

Project

For the term project, students are asked to develop an application in which Watson assumes a central role. We will train Watson with about 1200-1500 question-answer pairs. This is a time consuming task and as such, we will select just one domain for the entire class. Groups of students may decide to develop different applications based on the domain, but we will train Watson on just one domain.

The aim of the project is for our students to gain first-hand experience with training Watson and to learn about the power and limitation of Q/A system. Additionally, students will gain practice advertising their project through an oral presentation, a technical write-up and a video demonstration.

Timing of project phase

We established the following deadlines for the project phase.

- Day 18: Revised innovation canvas for the application
- Day 20: Initial architecture of application and set of documents to be curated
- Day 24: Curated documents and initial set of questions
- Day 28: Questions developed, training completed and progress report of application development
- Day 32: Questions developed, training completed and progress report of application development
- Day 36: PPT presentation of application, training completed and progress report of application development
- Day 38: Video demo of application
- Day 40: Final version of application, training completed and technical write-up of application

Here too, we provide additional background on those assignments that may not be intuitively obvious.

It is very likely that the chosen domain and application has not been favored by all students. To familiarize all students to the chosen subset of applications and to incorporate their input, a revised innovation canvas was to be developed, to capture the most recent thinking about the app.

While a suitable set of source documents has to be identified as part of the project proposal, we ask our teams to conduct a more thorough search. This search was split into two very short phases: locating documents and curating them. The curation process consists of keeping actual information and removing navigation and non-textual information from the documents. Additionally, we ensured that title and section headings were meaningful and relevant.

For the training phase, small teams of students were asked to develop 33 questions per milestone and train Watson with those questions.

Final

The take-home final is designed to think about the future of cognitive systems, i.e. a future in which there are personal assistants in the form of software applications, in particular applications that assist people who operate at a very high level.

Our students are asked to write a three page essay in which they envision a not too distant future in which software personal assistants are prevalent. Additionally, they are asked to evaluate this future by addressing concrete benefits and drawbacks.

Best Practices

Perhaps the biggest challenge in teaching this course is to manage expectations. Watson has been portrayed as considerable more powerful than it actually is. To be precise, Watson is a question-answering system. In particular, it is not a reasoning system. It cannot evaluate choices. This may change in the future, but for now, the version of Watson to which students get access is the Q/A version.

A second challenge is the choice of a good domain. Watson requires a good amount of training. IBM recommends the development of about 1200-1500 question-answer pairs. There will only be about 300-400 answers, as such, there will be about three to five questions per answer. In essence, one trains Watson to learn the way users ask a question. Considering the relatively small number of answers, this suggests that the domain will be fairly small.

The above exposition suggests another criterion for selecting a good domain: One in which people may not know how to phrase a question in the first place. A good example and one for which Watson is used in industry is customer support.

Finally, the precision and recall of a Watson application is highly correlated to quality of the documents that are ingested. As such, as good amount of time should be spent curating the documents. In particular, the documents should have a good amount of structure to them in the form of a title, section and sub+-section headings. The information in the headings should be relevant to the section body. Ideally, the section heading is repeated in the section body. It is perfectly fine to edit out information that is not pertinent to the answers.

As a final thought, it has been said that the answer to about 95% of Jeopardy! questions occur in Wikipedia titles.

A Use Case

To give the reader a sense of domains that might work well, we now present the application that we developed in our course. We developed a First-Aid iOS app. There are several benefits to this choice, chiefly, the source documents are already highly vetted and edited. Secondly, a smart phone adds many other features that greatly embellish the application.

According to personal communication by a technical advisor to the IBM Academic Initiative, our application achieved 63% precision and 83% recall. This was accomplished with only 30% of the training completed.

The reader may be interested in learning that this domain was chosen by a class vote. We narrowed down project proposals to two choices, partly through feedback from contacts at IBM. The First-Aid app won by a 4-3 vote.

We felt that the primary benefit of using Watson in a first aid situation is that in a panic situation people do not necessarily phrase their questions properly.

Additional benefits offered by a smart phone include the option of hands-free access through speech-to-text and text-to-speech components as well as speech activated, location based emergency dialing and directions to nearby health care facilities.

We will provide details of our solution in the next few sections.

Functionality

There are several groups of components that we implemented for our application. Please refer to figure 1 for details. Starting at the bottom, we provide one-handed access to the Watson First-Aid Q/A application.

Above the first aid button is a row of buttons which will call various emergency services. At this time, we are able to call 911, the local police department and the national number for poison control. While emergency services may be called at any time during a Q/A session, undoubtedly there are situations when it is obvious that emergency help is needed.

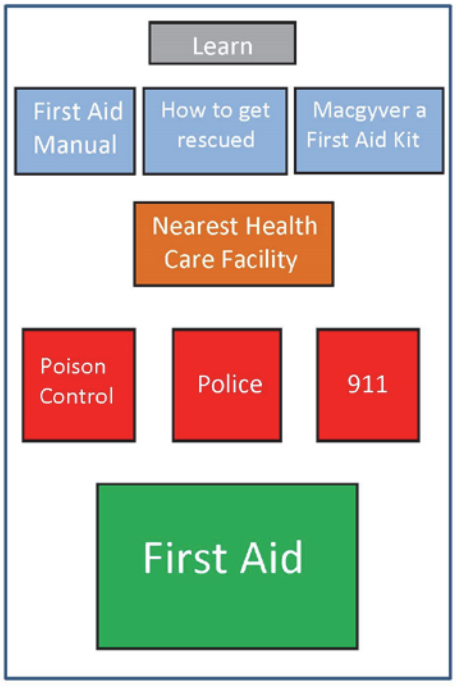

Figure 1: UI of First-Aid App

Similarly, there are cases when one wishes to locate the nearest hospital or, in case of remote areas, simply a nearby health care facility of any kind. The "Nearest Health Care Facility" button will use the iPhone's mapping software to indicate and if selected, provide directions to nearby health-care facilities, including hospitals.

A number of first aid cases occur in locations without cell or data reception. Our application has a row of buttons that provide various pieces of information that is packaged with our application. They include a first aid manual, a guide on how to get rescued in case one is lost and finally, instructions on how to fashion first aid materials from everyday items, in case a first aid kit is not available.

Lastly, we propose to add a repository of materials designed for the user to learn or brush up on first aid knowledge as well as information on how to put together an appropriate first aid kit. The "Learn" button would serve as a gateway to such materials.

Architecture

In this section, we describe the architecture of the Watson component proper. Please refer to figure 2 for the components and the control flow between them. When the "First aid" button is clicked, a dictation field opens up. This dictation field enables the user to use the built-in speech-to-text software to enter a question or to type it in. The benefit of speech-to-text is that it enables hands-free operations, a critical feature in a first aid situation. We decided to add the typing functionality for situations when there is a significant amount of ambient noise, something that is to be expected in some first-aid situations.

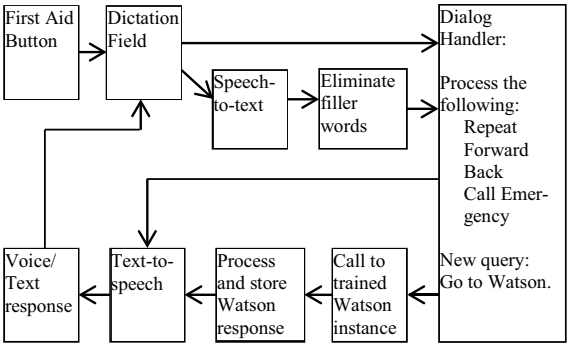

Figure 2: Architecture of First-Aid app

The speech-to-text component will produce only correctly spelled words. The auto correction feature on the iPhone is of a quality that no additional spell-checking is necessary.

For speech input, we plan to eliminate filler words before sending the query to Watson. Based on an informal study of recordings of emergency calls found on the web, people do not seem to add a lot of filler words when calling 911. Additionally, the very nature of filler words is that they do not change the grammar of a sentence. As such, they can simply be eliminated without much effort.

At this point of the development, our system returns the entire response as returned by Watson. It is displayed on the screen as well as read to the user.

Training Data

The benefit of choosing first-aid information as a domain is that there is a good amount of highly edited and highly vetted information available. We used information from the following sites: Mayo Clinic (Mayo Clinic 2016), WebMD (WebMD 2016), and St John Ambulance (St John Ambulance 2016). We combined information from all three sites and generated about 28 documents for the 23 groupings of topics that we identified.

Due to the high quality of the information, the curation process consisted primarily of ensuring that the documents have the sort of structure that Watson likes to have: titles, a fair number of headings and the repetition of the heading in the body of the text. All of our documents are in HTML format.

It should be pointed out that due to the copyrighted nature of the information we obtained, we are not able to make this application publicly available at this time.

Evaluation of Our Course Offering

We taught this course for the first time during the spring of 2015 with a good amount of support from IBM. This support came in the form of course materials (IBM Course Materials 2016) and advice from an IBM employee who is knowledgeable about Watson and volunteered their time to offer advice. We were very lucky because our contact is an alumnus. We spent about an hour a week discussing characteristics of a good domain and application. This was in addition to reviewing student project proposals. The assistance offered by our contact was invaluable. It would have been useful to start the consultation process several weeks before the course began so that we could have given our students good advice from the beginning of the term, however, our workload prior to the spring term made this not feasible.

Consonant with our learning "on the job," we began the training process later than would have been desirable. We did not fully train our instance of Watson. This did not seem to distract from the learning experience as we were able to fully train our instance on some sub-areas of First-Aid knowledge and as such were able to appreciate the power of Watson.

We did not integrate the videos about Watson as well as we could have. They were assigned throughout the course, however, we should have dedicated class time towards

discussing the contents of the videos and how they relate to the work we performed with Watson.

The phrasing of the questions, according to the way it was related to us, is important. As such, we plan to insure more consistency of the way the questions are phrased across groups.

Something we had never done before turned out to be a refreshing component of the course. To receive credit for the IBM MOOC, we agreed with IBM to make a video advertising key features and benefits of our iPhone app. We went as far as writing a script for our skid and practiced it before we videotaped it.

Seven students were enrolled in our course. I knew all of them from prior courses and they were some of the best students we have had. This small setting and the quality of our students made for a very enjoyable course for all of us. Among others, my students enjoyed learning about Watson and to have the privilege to get access to it. We too enjoyed working with our students on a cutting edge project. My students suggested bettering the pace of the course by pushing the workload towards the front of the course. In particular, they suggested starting the project earlier. The course as described here implements those suggestions.

Conclusions

The course as presented in this paper is an improved version of what we taught. The improvements were based on student feedback and our own experience teaching the course. For the complete set of course materials, please visit our course site (Wollowski 2016).

Absent from this course is a module in which students present a business proposal to some entity external to this course, commonly known as "Shark Tanks." We originally planned on such a learning component but ran out of time.

Based on our experience with attempting to commercialize our application and discussions with experts in the field of taking a product to market, we came to the conclusion that taking a product to market requires skills and experience of a nature that we cannot reasonable teach in two weeks. As a matter of fact, we believe that the best advice we can give our students is to network and connect with people who have business skills that are on par with our students' exceptional technical skills.

References

Ahmed, Jameel et. al. 2014. *The Innovation Canvas: An Instructor's Guide.* Proceedings of the 2014 ASEE conference.

Bayers, Chip. 1999. *The Inner Bezos.* http://archive.wired.com/wired/archive/7.03/bezos.html

Computing Reviews. 2016. CR Writing Guidelines. http://www.computingreviews.com/Reviewer/Rev-info.html

Ferrucci, David et.al. 2012. *This is Watson*. IBM Journal of Research and Development, Issue 3.4

IBM. *Cognitive Systems MOOC*. 2016. http://www-304.ibm.com/services/weblectures/watsonacademy/#discover

IBM Cognitive Systems Institute. 2016. *Point of View documents.* http://cognitive-science.info/learning-resources/point-of-view-documents/

IBM Course Materials. 2016. Watson Course Outline and Resources. http://www.rose-hulman.edu/class/csse/csse290-Cognitive Systems/WatsonCourseOutlineAndResources02Sept2014.xlsx

IBM Watson. 2015. *Watson University Competition.* https://www.youtube.com/playlist?list=PLZDyxLlNKRY-2oGTHZrKILIDz7u4-p3dR

Kline, Bill. 2016. *The Innovation Canvas*. http://www.rose-hulman.edu/media/882769/blank_innovation_canvas_a3v1.pdf

Mayo Clinic. 2016. *First Aid - Mayo Clinic*. http://www.mayoclinic.org/first-aid

St John Ambulance. 2016. *First aid tips and information.* http://www.sja.org.uk/sja/first-aid-advice.aspx

WebMD. 2016. *First Aid Guide and Emergency Treatment Instructions*. http://www.webmd.com/first-aid/

Wikipedia. 2016. *Watson*. https://en.wikipedia.org/wiki/Watson_(computer)

Wollowski, Michael. 2016. *Cognitive Systems Course Materials*. http://www.rose-hulman.edu/class/csse/csse290-CognitiveSystems/