

Training Watson - A Cognitive Systems Course

Michael Wollowski

Rose-Hulman Institute of Technology
wollowski@rose-hulman.edu

Abstract

We developed 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 poster, we justify this course, characterize major assessment items and provide advice on choosing a domain.

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 develop an application for IBM's Watson, thereby studying Q/A techniques from an applied perspective. The main 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.

A second significant portion of the course is concerned with developing students' entrepreneurial skills. 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.

A third component of this course centers around studying technical details underlying Watson. Through technical

papers and an IBM produced MOOC about Watson [IBM], there is an opportunity to study key technologies underlying advanced Q/A systems.

There is a wild-held belief that cognitive assistants will assume a significant role in our future. Studying Q/A systems like Watson enables us to critically examine some of the claims made about them.

Best Practices

Perhaps the biggest challenge in teaching this course is to manage expectations. 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, related challenge is the choice of a good domain. Watson requires a good amount of training. IBM recommends the development of about 1200-1600 question-answer pairs. In essence, one trains Watson to learn the way a user would ask a question. In a sense, Watson enables the user to pose a question in 1,001 ways. Considering the relatively small number of answers, this suggests that the domain ought to 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.

The precision and recall of a Watson application is highly correlated to quality of the documents to be ingested. As such, a 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.

Prerequisites

For the first offering of this course, we required our upper level Artificial Intelligence course as a prerequisite. As such, we were able to dive fairly deep into the technology of Watson.

For the next offering, we lowered the prerequisites to our data structures course. There is a push towards entrepreneurship at our institution and such a move would enable our students to get exposed to entrepreneurship at a much earlier point in their academic careers.

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 (several rounds, recruitment).
- 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

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.

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 project proposal phase consists of a four week process in which students will 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.

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 the most advanced Q/A system. Additionally, students will gain practice advertising their

project through an oral presentation, a technical write-up and a video demonstration.

The take-home final is designed to think about the future of Cognitive Computing, 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.

Conclusions

We taught this course for the first time during the spring of 2015. We since improved the course significantly and the course documents can be found on our course site [Cognitive Systems Course]. The improvements were based on student feedback and our own experience teaching the course.

Our students very much enjoyed the course; in particular, they enjoyed learning about Watson and to work on something special. We too enjoyed working with our students on a cutting edge project. Our 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.

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 reasonably 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

- Cognitive Systems Course. www.rose-hulman.edu/class/csse/csse453/watson
- Ferrucci, David et al. 2012. This is Watson. *IBM Journal of Research and Development*, Issue 3.4
- IBM. *Cognitive Systems MOOC*. <http://www-304.ibm.com/services/weblectures/watsonacademy/#discover>
- IBM Cognitive Systems Institute. *Point of View documents*. <http://cognitive-science.info/learning-resources/point-of-view-documents/>
- IBM Watson. *Watson University Competition*. <https://www.youtube.com/playlist?list=PLZDyxLINKRY-2oGTHzrKILIDz7u4-p3dR>