MANNING PUBLICATIONS

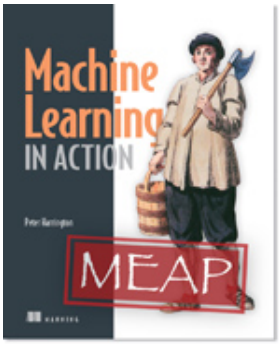

Machine Learning in Action
By Peter Harrington

This article, related to Machine Learning in Action, discusses the inner workings of Watson, the computer built by IBM to compete in Jeopardy! and beat the two greatest human players the game has ever known. Watson's achievement was building a massive system that could quickly generate an answer or choose not to answer. Watson used the computing power to run multiple natural language processing algorithms on a question and combine the answers from these algorithms.

You may also be interested in…

Machine Learning with Watson: The Greatest Achievement of Artificial Intelligence to Date

The most salient feat of artificial intelligence in 2011 has been IBM's Watson computer playing Jeopardy! Not only did the computer play Jeopardy, but it beat the two best players the game has ever known. IBM is not new to developing great artificial intelligence systems—in 1997, IBM's Deep Blue beat world chess champion Gary Kasparov. In 2004, a research manager at IBM went out to eat when he noticed everyone in the restaurant fixed to the TV. They were paying attention to Ken Jennings' 74-game run on Jeopardy! Within the next year, IBM launched a program to build a computer system to compete in Jeopardy! Initial versions of the system took minutes to respond and had low accuracy rates. By 2008, the IBM system could compete with Jeopardy! champions and, by 2010, it would beat humans on a regular basis.

Inside Watson

Let's take a look at how Watson does so well on Jeopardy. The first thing is its use of massive computing power. Watson has a data set of over four terabytes, with over 200 million pages of structured and unstructured content. This whole data set is kept in RAM because waiting for a hard drive takes too long. Secondly, Watson is made up of 90 Power 750 servers with a total of 2880 POWER7 processor cores. This machine can do 80 teraflops, which makes it the 94th most powerful computer in the world at the time of this writing.

Now, let's look at the software. C++ and Java were used to write the algorithms powering Watson. The algorithms were executed on SUSE Linux Enterprise Server 11. This is a big feat in natural language processing (NLP), and you must be asking yourself, "What was the special breakthrough?" You may be a little disappointed to find that Watson's main innovation was not a new algorithm. Instead, it was the ability to run hundreds of NLP algorithms at the same time, collect the results of these algorithms, and check them against a database to see if they are reasonable answers. IBM's literature reads, "More than 100 different techniques are used to analyze natural language, identify sources, find and generate hypotheses, find and score evidence, and merge and rank hypotheses."[1] This combining of multiple algorithms to see an answer is known as an ensemble. In Machine Learning, there are a number of different algorithms and each one has its own strengths and weaknesses

[1]"Watson—A System Designed for Answers" IBM Systems and Technology, February 2011
ftp://public.dhe.ibm.com/common/ssi/ecm/en/pow03061usen/POW03061USEN.PDF

depending on the problem and the data. You could spend time trying to figure out which is the best algorithm or you could use many algorithms at once. Ensemble methods use many algorithms at one time and combine their results in some way to get a single final answer. In most cases ensemble methods outperform single algorithms, often with much better performance. Ensemble methods are covered in more detail in chapter 7 of my book *Machine Learning in Action* published by Manning Publications.

IBM's history of AI successes

When IBM's Deep Blue defeated world chess champion Gary Kasparov in 1997, the game at hand was easier for a computer to understand than playing Jeopardy! Let me explain what I mean by that. In a game of chess, you have 64 squares and 16 pieces for each player. There are six different types of pieces with predefined patterns. Now take the game of Jeopardy!. The Oxford English Dictionary lists over 250,000 distinct words, not counting slang and people's names. Put together, combinations of words and the amount of inputs gets very large very quickly. You can see how NLP can be difficult with common language. Jeopardy! makes NLP even more difficult because the language used is full of puns, misleading statements, extra details, and cultural references.

The way Watson's software approached this problem was by having the full text of Wikipedia in its memory, along with additional data. The NLP algorithms generate a series of answers. For each answer a confidence score is assigned. If any of these answers has a confidence score over a certain threshold, Watson attempts to answer—the threshold appears to be set at 50%. Calculating a probability of a number of possible answers and then selecting the answer with the highest probability is the approach taken with the Bayesian classifier. Bayesian classification is addressed in chapter 4 of my book. Bayesian decision theory allows us to draw conclusions from the statistical nature of our inputs. From that point, we can assign a probability to each possible answer and select the answer with the highest probability.

IBM's victory at Jeopardy! was a great accomplishment, but was it just a an expensive publicity stunt or will IBM be able to cash in on this investment? The benefits of this project to humankind as a whole far outweigh its costs. This is arguably the greatest achievement in the field of artificial intelligence to date. Early approaches to artificial intelligence promised philosophizing computers yet failed to deliver machines that could perform humble tasks. For a long time, the idea of a machine answering questions in human terms has been the stuff of science fiction. Now it's no longer fiction. I hope Watson inspires students to pursue a career in the sciences.

It is reported that IBM does envision using the Deep Thought technology in the fields of legal research and medicine. There are a number of companies using Machine Learning extensively in legal research. The number of hours needed to read documents for a case often is greater than the number of hours a client can afford. A machine, on the other hand, can digest these documents. A machine that answers a question posed in the human format could prove to be an indispensible tool in the court room.

Physicians can often misdiagnose illnesses due to limited experience, insufficient information, or arrogance. For example, a physician that spent the majority of his time in one region of the country may be poorly prepared to deal with a sick patient from another region. A machine with encyclopedic information gathered from physicians all over the world would do a better job. Watson's software, or software similar to it, could use the encyclopedic information to make a proper diagnosis with a few human language questions.

Summary

IBM's Watson was a significant demonstration of machine learning, natural language processing, and information retrieval. The world will benefit from the science and engineering put into creating it. It has renewed my interest in some of these areas and their applications. I hope it has renewed yours too.

Here are some other Manning titles you might be interested in:

The Quick Python Book, Second Edition
Vernon L. Ceder

Real-World Functional Programming
Tomas Petricek with Jon Skeet

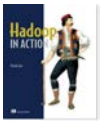

Hadoop in Action
Chuck Lam

Last updated: April 1, 2011