What's Up with IBM's Watson?

CASE STUDY

In February, 2011 an IBM computer named Watson made history by handily defeating the two most decorated champions of the game show Jeopardy, Ken Jennings and Brad Rutter. Watson was named after IBM's founder, Thomas J. Watson, and its achievement marked a milestone in the ability of computers to process and interpret human language.

IBM had been working on Watson for years. The project's goal was to develop a more effective set of techniques that computers can use to process *natural language* – language that human beings instinctively use, not language specially formatted to be understood by computers. Watson had to be able to register the intent of a question, search through millions of lines of text and data, pick up nuances of meaning and context, and rank potential responses for a user to select, all in less than three seconds.

The hardware for Watson used in Jeopardy consisted of 10 racks of IBM POWER 750 servers running Linux, with 15 terabytes of RAM and 2,880 processor cores (equivalent to 6,000 top-end home computers), and operated at 80 teraflops. Watson needed this amount of power to quickly scan its enormous database of information, including information from the Internet. To prepare for Jeopardy, the IBM researchers downloaded over 10 million documents, including encyclopedias and Wikipedia, the Internet Movie Database (IMDB), and the entire archive of *The New York Times*. All of the data sat in Watson's primary memory, as opposed to a much slower hard drive, so that Watson could find the data it needed within three seconds.

Watson is able to learn from its mistakes as well as its successes. To solve a typical problem, Watson tries many of the thousands of algorithms that the team has programmed it to use. The algorithms evaluate the language used in each clue, gather information about the important people and places mentioned in the clue, and generate hundreds of solutions. Human beings don't need to take such a formal approach to generate the solutions that fit a question best, but Watson compensates for this with superior computing power and speed. If a certain algorithm works to solve a problem, Watson remembers what type of question it was and the algorithm it used to get the right answer. In this way, Watson improves at answering questions over time. Watson also learns another way — the team gave Watson thousands of old Jeopardy questions to process. Watson analyzed both questions and answers to determine patterns or similarities between clues, and using these patterns, it assigns varying degrees of confidence to the answers it gives.

Although Watson was only able to correctly answer a small fraction of the questions it was initially given, machine learning allowed the system to continue to improve until it reached Jeopardy champion level. IBM term cognitive computing to refer to Watson's ability to interpret speech and text, rapidly mine large volumes of data, answer questions, draw conclusions, and learn from its mistakes.

The Watson version used in Jeopardy took 20 IBM engineers three years to build at an $18 million labor cost, and an estimated $1 million in equipment. IBM saw the investment as a stepping stone to broader commercial uses of its AI technology, including applications for health care, financial services, or any industry where sifting through large amounts of data (including unstructured data) to answer questions is important. Watson is expected to become more useful and powerful by learning from new sets of experts in new fields of knowledge. In January 2014 the company created a new division, the Watson Business Group, which will have 2500 employees working largely in New York City's Silicon Alley. IBM has invested more than $1 billion in this Group, and has allocated one-third of its overall research efforts to Watson.

In September 2011, WellPoint Inc., the largest U.S. health care provider, with 34.2 million members, enlisted Watson for utilization management. The WellPoint Interactive Care Reviewer application is designed to determine if physicians' requested treatment meets the guidelines of the company and a patient's insurance policy. The Watson WellPoint application combines data from three sources: a patient's chart and electronic records maintained by a physician or hospital, the insurance company's history of medicines and treatments, and Watson's huge library of textbooks and medical journals. According to WellPoint vice president Elizabeth Bingham, Watson initially took too long to "learn" WellPoint's policies. IBM was able to improve the system by revising the Watson training routine for

WellPoint, and the Interactive Care Reviewer is being adopted by 1600 health care providers.

Cancer treatment appears to be an especially promising application for Watson. Current guidelines aren't precise enough to determine treatments that are most appropriate for a specific patient. For example, the recommended treatment may be chemotherapy, but how do you pick among ten or more possible chemotherapy options? How do you choose the dosage? What treatment frequency would work best? Oncologists also can't keep pace with the torrent of cancer research findings and therapies, genomic techniques, and patient record data. It is just too much for even a highly-trained scientist to manage.

In 2012 Memorial Sloan-Kettering Cancer Center began work on a Watson application to recommend cancer treatments, using data from Sloan-Kettering's clinical database of over one million patients along with treatment guidelines and published research to help Sloan-Kettering researchers to recommend personalized treatment options for lung cancer patients. The Watson application needs to pass a series of tests in order to be used on cancer patients, and actually being able to use Watson is more complex than originally envisioned. For instance, Sloan-Kettering oncologist Dr. Mark Kris displayed a screen from Watson that listed three potential treatments, but Watson was less than 32 percent confident that any of them were correct. Ari Caroline, director of Sloan Kettering's quantitative analysis and strategic initiatives group has tutored Watson and has said that system was still in pilot mode. But progress is genuine, and Caroline believes Watson will soon be able to guide oncologists in selecting treatment options and tackling new research. A final version of the system has not yet been released.

Researchers at the University of Texas MD Anderson Cancer Center worked with IBM for a year to build a version of Watson called Oncology Expert Advisor (OEA) to recommend cancer treatments by mining medical literature, with an initial focus on acute leukemia. Watson learned from a variety of data about which cancer treatments worked best and which should be avoided for specific patients. OEA "reads" the medical records of patients to generate case summaries. It then weighs the patient profile against its knowledge base to suggest treatment options relevant to that particular patient, based on literature, guidelines and expert recommendations. When asked by a doctor about a patient, Watson's algorithms search for possible treatments and rank them according to levels of confidence up to 100%, with each option linked to supporting evidence.

The project initially stumbled because IBM engineers and Anderson doctors couldn't understand each other. IBM developers worked elsewhere and only visited Anderson every few weeks to talk to doctors. When IBM developers and doctors started meeting several times a week, the application became much better and the leukemia advisor is nearly ready for use. However, it might take two more years before Watson could handle other types of cancer. And although Watson might help oncology specialists at M.D. Anderson identify leukemia treatment options, it can't substitute for the expertise of an experienced doctor, according to Lynda Chin, chairperson of the M.D. Anderson genomic department. The cancer experts have seen patients a thousand times, and sometimes their decisions are based on intuition that's difficult to explain. The Anderson project was valued at nearly $15 million, and IBM management is hoping it could grow to $100 million. The Anderson project plans to expand to other cancer types once the prototype becomes more developed.

In November 2013 IBM announced it would make Watson technology available via the Internet as a cloud service that could be used by many different industries. IBM will open parts of the system to outside developers to create businesses and mobile applications based on cognitive computing. A Watson Developer Cloud provides tools and methodologies for developers to work with a Watson system, a content store supplying both free and fee-based data for new applications, and about five hundred subject matter experts from IBM and third parties. Welltok used these tools to create a mobile Watson app called CareWell Concierge for Intelligent Health Itineraries for consumers. Users will be able to participate with Watson in conversations about their health. Fluid Retail is developing a personalized shopping assistant. MD Buyline is developing a Watson app to advise hospital managers about procurement of medical equipment and supplies.

IBM will deliver three new cloud-based products based on Watson's cognitive intelligence and capabilities. IBM Watson Discovery Advisor is aimed at the pharmaceutical, publishing, and education industries, and will wade through search results to deliver data faster and help researchers formulate conclusions. IBM Watson Analytics is a cloud-based service that provides insights, including visual representations, based on raw big data enterprises send to Watson. IBM Watson Explorer is a cloud service that will provide a unified view of a user's information, facilitating the revelation and sharing of data-driven insights .

IBM has also made Watson easier and less expensive to use. The latest version of Watson is 24 times faster than the version used in the 2011 Jeopardy contest, using only 10 percent of the hardware used in the Jeopardy version.

Nevertheless, Watson thus far has not produced much revenue for IBM—only about $100 million from commercialization efforts between 2011 and 2014. IBM CEO Virginia Rometty is hoping Watson will be able to produce $10 billion in annual revenue within a decade, and that Watson will bring in $1 billion in revenue per year by 2018.

In order to effectively commercialize the technology, IBM will need to expand Watson's knowledge domains, and this is its greatest challenge. Turning Watson into a useful business tool requires an enormous amount of work. Watson has to learn the terminology and master the domains of expertise in many different areas, including health care and scientific research, understand the context of how that language is used, and how to correlate questions with the correct answers. Watson doesn't work yet with data from audio, video, and animations and with languages other than English. It can't come up yet with its own ideas.

IBM will have to be careful not to oversell what Watson can do, so that Watson does not end up like other artificial intelligence systems where expectations were way overblown. Making machines that beat humans at chess or a TV game show is much easier than solving problems in the real world. According to Curt Monash, president of Monash Research, Watson hasn't yet overcome the hurdle that derailed AI in the 1980s, which was that AI was only able to capture small pieces of a limited knowledge domain for a single-purpose use. Watson is having more trouble solving real-life problems than Jeopardy questions. Watson's basic learning process requiring IBM engineers to master the technicalities of a customer's business and translate those requirements into usable software has been very arduous.

It remains to be seen whether the complexity of establishing a body of knowledge and training an intelligent system is repeatable and scalable for other types of work and whether it creates an opportunities for differentiation and competitive advantage. Watson is very much a work in progress.

Sources: Mohana Ravindranath, "How IBM Is Trying to Commercialize Watson," *Washington Post*, May 11, 2014; Spencer E. Ante, "IBM Struggles to Turn Watson Computer Into Big Business, Wall Street Journal, January 7, 2014; Lynda Chin, "IBM Watson: Providing a Second Opinion for Oncologists," www.ibm.com, accessed July 9, 2014; George Lawton, "IBM's Watson Supercomputer Gives Developers Access to Cognitive Cloud," SearchCloudApplications.com, March 28, 2014; Jack Vaughan, "For IBM Watson, No easy Answers on Commercial Cognitive Computing," Searchdatamanagement.com, January 10, 2014; Michael Goldberg, "Five Things to Know about IBM Watson, Where It Is and Where It's Going," DataInformed, January 14, 2014; Larry Dignan, "IBM Forms Watson Business Group: Will Commercialization Follow?," ZDNet, January 9, 2014; Quentin Hardy, "IBM Bets Watson Can Earn Its Keep," *New York Times*, January 8, 2014 and "IBM to Announce More Powerful Watson via the Internet," *New York Times*, November 13, 2013; Ian B Murphy, "Predictive Analytics in Development: IBM Watson at Memorial Sloan-Kettering, RPI Research Lab," DataInformed, February 20, 2013; Anna Wilde Mathews, John Markoff, "Computer Wins on 'Jeopardy!': Trivial, It's Not," *The New York Times*, February 16, 2011; Stanley Fish, "What Did Watson the Computer Do?" *The New York Times*, February 21, 2011; and Stephen Baker, "The Programmer's Dilemma: Building a Jeopardy! Champion," *McKinsey Quarterly*, February 2011.

CASE STUDY QUESTIONS

11-13 How powerful is Watson? Describe its technology. Why does it require so much powerful hardware?

11-14 How "intelligent" is Watson? What can it do? What can't it do?

11-15 What kinds of problems is Watson able to solve? How useful a tool is it for knowledge management and decision making?

11-16 Do you think Watson will be as useful in other industries and disciplines as IBM hopes? Will it be beneficial to everyone? Explain your answer.

MyMISLab

Go to **mymislab.com** for Auto-graded writing questions as well as the following Assisted-graded writing questions:

11-17 How do each of the following types of systems acquire and represent knowledge: Expert system, case-based reasoning, neural network?

11-18 How do enterprise content management systems help organizations manage structured and semistructured knowledge? What are two examples of each type of knowledge handled by these systems?