CHAPTER 11 MANAGING KNOWLEDGE AND COLLABORATION

CASE 1

How IBM's Watson Became a Jeopardy Champion

VIDEO CASE

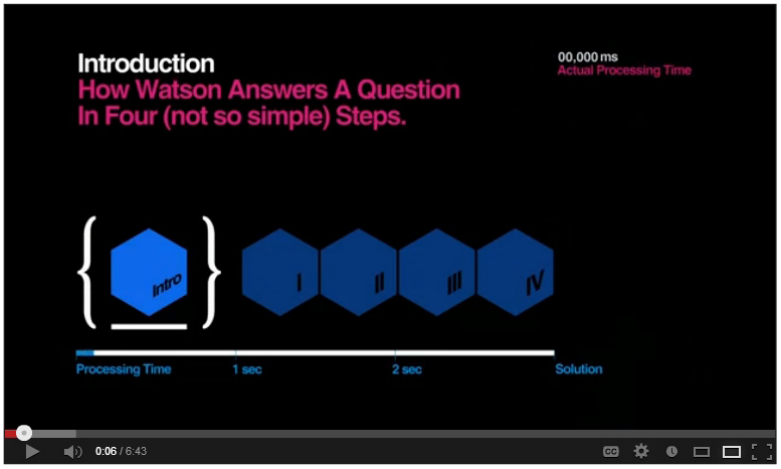

SUMMARY This video describes how IBM's Watson computer system is able to compete against humans in answering questions on the popular television game show *Jeopardy!* and how this technology will impact business and industry. L=6:42.

URL http://www.youtube.com/watch?v=DywO4zksfXw

CASE *Jeopardy!* is an American television quiz show created in 1964. The show asks contestants questions on a very wide range of topics, including history, language, literature, the arts, the sciences, popular culture, geography, and wordplay. The show had a unique question format: contestants are presented with answers containing clues, and must phrase their responses in the form of a question that would generate that answer. For instance, the clue "5280" is a response to the question, "How many feet in a mile?"

Here are some sample questions and appropriate answers in *Jeopordy!* style:

Q: Hard times, indeed! A giant quake struck New Madrid, Missouri, on February 7, 1812, the day this author struck England.
A: *Who is Charles Dickens?*

Q: According to C.S. Lewis, it was bordered on the east by the Eastern Ocean and on the north by the River Shribble.
A: *What is Narnia?*

continued

Q: Pseudonym of labor activist & magazine namesake Mary Harris Jones.
A: *Who is Mother Jones?*

Q: To marry Elizabeth, Prince Philip had to renounce claims to this southern European country's crown.
A: *What is Greece?*

The game of *Jeopardy!* makes great demands on its players—from the range of topical knowledge covered to the nuances in language employed in the clues. For instance, to answer the last question you would need to know something about Elizabeth the Queen, and her husband Prince Phillip, and what southern European country the Prince had claims on. Can the analytical power of a computer system—normally accustomed to executing precise requests—overcome these obstacles? Can the troves of knowledge written in human terms become easily searchable by a machine in order to deliver a single, precise answer? IBM's Watson took on that challenge.

Watson is a question answering (QA) computing system. IBM describes it as "an application of advanced Natural Language Processing, Information Retrieval, Knowledge Representation and Reasoning, and Machine Learning technologies to the field of open domain question answering" using IBM's DeepQA technology for hypothesis generation, massive evidence gathering, analysis, and scoring. Watson was named after IBM's first president, Thomas J. Watson.

In February 2011, Watson competed on *Jeopardy!* against two top human *Jeopardy!* players. Watson beat Brad Rutter, the biggest all-time money winner on *Jeopardy!*, and Ken Jennings, the record holder for the longest championship streak (74 wins). Watson received the first prize of $1 million, while Jennings and Rutter received $300,000 and $200,000, respectively. IBM divided Watson's winnings between two charities.

In order to win at *Jeopardy!* Watson had to be able to understand the language of a clue, register the intent of a question, scour millions of lines of human language, and return a single, precise answer—in less than three seconds. How did Watson do it? Watson had access to 200 million pages of structured and unstructured information consuming four terabytes of disk storage, including the full text of Wikipedia, but was not connected to the Internet during the game. That's because the Internet is far too slow to meet the response-time requirements of the game. Watson's sources of information included encyclopedias, dictionaries, thesauri, newswire articles, and literary works. Watson also used databases, taxonomies, and ontologies. All of this content was stored in Watson's RAM memory for the game because data stored on hard drives are too slow to access. Watson is able to process 500 gigabytes of data, the equivalent of a million books, per second.

continued

Watson received the *Jeopardy!* clues as electronic texts at the same moment they were made visible to the human players. It would then parse the clues into different keywords and sentence fragments in order to find statistically related phrases. The machine zeroes in on key words in a clue, then combs its knowledge databank for clusters of associations with those words. It rigorously checks the top hits against all the contextual information it can muster: the category name; the kind of answer being sought; the time, place, and gender hinted at in the clue; and so on.

Watson's main innovation was not in the creation of a new algorithm for this task but rather its ability to quickly execute thousands of proven language analysis algorithms simultaneously to find the correct answer. The more algorithms that find the same answer independently, the more likely Watson is to be correct. Once Watson has a small number of potential solutions it is able to check against its database to ascertain whether the solution makes sense. When Watson feels "sure" enough, it signals it is ready, and speaks with an electronic voice and to deliver the responses in *Jeopardy!*'s question format.

This is all an instant, intuitive process for a human *Jeopardy!* player. Watson has deficiencies in understanding the contexts of the clues. As a result, human players usually generate responses faster than Watson, especially to short clues. The video shows some of the questions Watson flubbed. Watson's nerve center is 10 racks of IBM Power750 servers running in Yorktown Heights, NY, with the processing power of 6,000 desktop computers. The servers and software are set to run massively parallel computations. The Watson computers consumed over 80,000 watts of electricity an hour when answering *Jeopardy!*'s questions. In contrast, the human brain operates with about 20 watts of power.

Just as Watson analyzed massive data in *Jeopardy!* to reach a set of hypotheses and list the most likely outcomes, it could help doctors in diagnosing patients. Watson could analyze the patient's specific symptoms, medical history, hereditary history (as provided to it), and synthesize that data along with the mass of all unstructured and structured medical information in the world including (but not limited to) medical records of similar situations available to it. Watson would use this information in addition to the wealth of fundamental medical knowledge available in the form of all published medical books and articles fed into Watson. IBM has made it clear that Watson is not intended to replace doctors, but can assist them to avoid medical errors and sharpen medical diagnosis with the help of its advanced analytics technology. IBM has projects underway with Nuance Communications, a speech and imaging technology company, and Wellpoint, a large managed health care company, for this purpose. IBM intends to use Watson in other information intensive fields as well, including legal research.

continued

VIDEO CASE QUESTIONS

1. Was playing Jeopardy a good way to test machine intelligence? Why or why not?
2. Is Watson a good example of a computer system demonstrating intelligence similar to that of a human?
3. Would Watson be useful in customer service systems? What would it take for businesses to use Watson for this purpose?
4. Suggest some other applications for Watson.