

# geek



**FREE**  
COLLECTIBLE  
POSTER

**EXCLUSIVE INSIDE**

Meet the  
Warrior  
Women of  
**The Red  
Star**



# Star Wars

## YES, IT'S COOL AGAIN!

Experts argue the return of **Han, Luke and Leia**, the prequels, the end of Clone Wars and J.J. Abrams' epic challenge.

"I'D LIKE TO GO BACK TO THE HEART AND THE FUN THAT WE HAD IN OUR MOVIES."  
— MARK HAMILL PG.80

### World War Z

Zombies Gone Global

### Geek Guide: Home Theater

The Good, The Better and The Best

### LAST CALL

Director Todd Phillips' final Hangover

### REVENGE OF THE NERDIST

Chris Hardwick Sounds Off



**GEEKEXCHANGE.COM**  
VOL. 2 / NO. 1  
\$6.99  
DISPLAY UNTIL 8/13/2013

## THE ULTIMATE COMIC-CON SURVIVAL GUIDE

Insider Tips for Geektastic Fun



■ HELLO, DAVE: Thanks to deep learning, our dreams of being dominated by a terrifying computer intelligence may finally come true, with help from the likes of HAL and the evil A.I. from such movies as *WarGames*; *TRON*; *Colossus: The Forbin Project*; *The Matrix*; *Resident Evil*; *I, Robot* and *The Terminator*.

HAL illustration by Daren Dochterman

# Elementary, My Dear Watson

## WILL IBM'S QUIZ SHOW CHAMPION OUTGROW HUMANKIND?

By Daniel DeFabio



Dr. James A. Hendler, the head of RPI's Computer Science Department.

// IBM's Watson supercomputer has been called the smartest machine on earth. Now researchers and students at Rensselaer Polytechnic Institute (RPI) have been given the chance to see what Watson can do beyond trouncing *Jeopardy* champions. That in mind, *Geek* sat down with the head of RPI's Computer Science Department, Dr. James A. Hendler, whose research has included robotics, A.I., the semantic Web and Big Data. Hendler offered us a glimpse into the future of Watson, the coming of "memory prosthetics" and revealed whether Watson would be a good *Dungeons & Dragons* player.

**GEEK:** *What's Watson doing at RPI? And what's RPI doing with Watson?*

**Hendler:** We are the first university to receive both the Watson hardware and the software. The relationship between RPI and IBM on Watson has been a long one. Many of the people on the team were either RPI alums or had worked with our professors. We've been talking for a long time about new things to do with Watson. [Such as] what can Watson teach us about how we think? About how that kind of cognitive computing can be embedded in a larger computing environment. One of the things we're excited by is rapid transition. We can do some research and it can go right to IBM and then out into the world. So our goal is give our students access to something new, exciting, different and see what these amazingly creative young geeks can do. This really is something we feel is the right way to bring a new kind of computing out into the world.

**There are now several Watsons. Are they self-contained? Are they networked together? Will there be a debrief where things learned here get shared?**

They are separate. My assumption is that the people who are doing research using Watson will start talking together. We definitely want to have a community.

**Watson can teach itself, or learn, correct?**

There are different kinds of ways Watson can learn, some of it is new information coming in. So I just give it some new documents to read. *Zap!* It's in there. Essentially Watson is almost like having a lot of little boxes that are each trying to vote on an answer. What the Watson team found in playing *Jeopardy* is that a

the coming decades. Useful digital assistants, like Apple's Siri, are already using deep learning techniques. They are still limited, but what they can do was unthinkable just a few years ago, and researchers say the technology is now advancing at an unprecedented pace. Deep learning is based on neural networks, simplified models of the way clusters of neurons act within the brain that were first proposed in the 1950s. The difference now is that new programming techniques combined with the incredible computing power we have today are allowing these neural networks to learn on their own, just as humans do. The computer is given a huge

pile of data and asked to sort the information into categories on its own, with no specific instruction. This is in contrast to previous systems that had to be programmed by hand. By learning incrementally, the machine can grasp the low-level stuff before the high-level stuff. For example, sorting through 10,000 handwritten letters and grouping them into like categories, the machine can then move on to entire words, sentences, signage, etc. This is called "unsupervised learning," and deep learning systems are very good at it.

Impressive, but we still have a long way to go. This is only a small step in building a truly

intelligent machine. Deep learning has no way of establishing causal relationships or making logical inferences. IBM's amazing, *Jeopardy*-winning Watson system uses deep learning only as part of a much larger ensemble of techniques, including statistical analysis, Bayesian inference and deductive reasoning. We still aren't even close to creating HAL 9000 or anything truly intelligent. However, if the pace of progress continues, self-driving cars may become commonplace and machines just might start to interact with humans in a way that doesn't result in one's iPhone being thrown against the wall.

much better way to do it is: For different kinds of questions, different entities have different votes. So for this kind of question we found these two things really well, so we'll give their votes a little higher value. You give it a whole lot of training examples. And then Watson learns from that which of the different components best answer those kinds of questions. So when you take Watson to a new domain, you use this learning algorithm to tune, essentially.

*Given a different experience of learning from different data is it possible that the RPI Watson one year out would function differently than any other Watson?*

That's our hope. Part of what we'll be doing is building some new software modules that extend its capabilities. Watson doesn't have very much numeric reasoning. As strange as it sounds, computers being very good at numbers, Watson is really good at text. We'd like to see what Watson does in those kind of capabilities. Suppose you wanted to build a game-playing system. Watson can know all the rules of the game, but the strategic component would have to live somewhere else. If we wanted a Watson Dungeons & Dragons player, you'd still need to build a little Dungeons & Dragons playing logic somewhere. But it could use Watson to know what all the different kinds of threats were and what they respond to and what happened last time we were in this situation. I'm talking about this as a conceptual idea. I'm not saying we're going to build it. But imagine you wanted to build a Watson D&D player. You would definitely want to have information about what a *wight* is. The word *mage* is something you'd want in your D&D player; it's a lot less likely to come up in the *Jeopardy* world. You'd have to know what's real and what's imaginary. But it wouldn't be a good decision maker. You probably wouldn't want it to be the dungeon master because that takes a lot of thinking about game situations and what would make it exciting.

*Could Watson sort of be let loose in a library and see how much it learns?*

Think of Watson as a very big memory. If you put a big memory out there in the library and it reads all the books... So now you have something that knows a lot of stuff from the books but it hasn't actually *done* anything yet. It's when you start asking it questions, you start to say, "This is the information I want." Watson is inherently a question-answering device, but, when you think about it, so is your memory. What Watson is doing is giving you access to a tremendous memory that can do a certain amount of this processing. Then if we want to use that memory, we build some other stuff that uses Watson. If you wanted Watson to have a conversation, you have to build

the conversation box on top of Watson. You could build a module around that.

*Would Watson pass the Turing test?*

The answer is no because it doesn't have any conversational capability. On the other hand, it would do an amazing job at the first few questions of the Turing test. You'd have trouble stumping it. What's exciting about Watson is that we've never had a computer system before that showed the capability to just answer complicated questions about almost anything. The reason we all think [Human *Jeopardy* champion] Ken Jennings is so amazing is it doesn't seem to matter what you ask him, he knows the answer. And Watson does too.

*You are one of the leaders of the "Semantic Web" movement, will that influence your work with Watson?*

Some of what we're going to be doing is taking government data that is now available around the world and making it a little more machine-readable with this Semantic Web stuff. You'll be able to ask Watson "Where can I find data about obesity in Europe?" Tools for gathering a lot of that information — tying that information to things you find on the Web — is hard to do. And Watson may be a thing that will really help us get a leg up on that. I have a student looking at how would we tie this thing to Google Glass.

*Is there a development on the near horizon that you're excited about?*


What I think will be the next game changer in the next 5 to 10 years is memory prosthesis. If I lose my leg, we have good technologies to replace it. But if I lose my memory... That kind of stuff could be living in a thing that is Watson's great grandchild. Even Watson already would help me with some of that if I could put my personal memories in it. My phone could be sitting in my pocket listening to this conversation so next time we got together it could remind me about what we talked about and where we met.

*What are some of your favorites in sci-fi?*

*Doctor Who, Red Dwarf and Star Trek.* The first time I ever saw a talking computer on *Star Trek*, and I was like, "I want one of those." And now with Watson we're getting to the early versions. Watson has reminded me of what I fell in love with in the early days. There's a whole generation of A.I. scientists who were inspired by the *Star Trek* computer.

FOR MORE:

- ▶ [watson.rpi.edu](http://watson.rpi.edu)
- ▶ [ibm.com/innovation/us/Watson](http://ibm.com/innovation/us/Watson)



## SHALL WE PLAY A GAME?

AUTHOR, FORMER SOFTWARE ENGINEER AND JEOPARDY CHAMPION **KEN JENNINGS** HAD A FEW WORDS ABOUT WATSON — THE MACHINE THAT BEAT HIM ON *JEOPARDY*.

*GEEK: What fictional computer is Watson most like?*

**Jennings:** We don't need to worry about it turning sentient any time soon. It's more like "Joshua" in the movie *WarGames*... it's a fast learner, but if it ends up nuking us, it'll only be because someone accidentally *told* it to.


*How would you improve Watson?*

I'd improve Watson's voice. It *totally* should have been Darrell Hammond doing Sean Connery from the *SNL Celebrity Jeopardy* sketches. Or, hey, get Alex Trebek to do the voice. People expect smug, dry superiority to have a Canadian accent.

*Are you up for a re-match?*

I'd play again any time. Brad Rutter and I both beat Watson in practice rounds. We'd love for Watson to get its comeuppance in an actual televised match. That said, Watson's reflex advantage on the *Jeopardy* buzzer means that the human players would need to get lucky with a better spread of categories. Watson isn't so hot at "high-concept" categories that require a lot of lateral thinking or some creative take on the relevant data. Straight factual retrieval is no problem for it, but ask it to intuit the logic behind, say, a category about hypothetical police reports for fictional literary villains, and it might choke. — *Daniel DeFazio*

Jennings latest book is "*Because I Said So! The Truth Behind the Myths, Tales, and Warnings Every Generation Passes Down to Its Kids.*" [ken-jennings.com](http://ken-jennings.com)



## What is 42?

- (A) The number of orbiting streaks around the Watson avatar logo.
- (B) The ultimate answer to life, the universe and everything, according to Douglas Adams.
- (C) Both are true. (Adjust tinfoil hat now.)