


Welcome to CSC490:

Cognitive Computing Capstone

About the course

- Three months ago...




3.5 years ago...





Watson today

Too Much Information: The Doctor's Data Dilemma



Watson Today

- Watson Stats:
 - 2500 engaged innovators
 - \$100 million equity investments
 - 500+ subject matter experts
 - 10 employees → 2,000




Higher education interests

- Student engagement
 - Build awareness (advocate cognitive computing)
 - Provide contests & speakers.
- Skill development
 - Build skills pipeline, future leaders
 - Provide awards, competitions, curriculum development, faculty materials, internships.
- Research

Watson Competition

- Officially, the "*Cognitive Computer Curriculum Competition*" → aka "4C".
- 10 schools participating:
 - CMU, Northwestern, NYU, OSU, RPI, Stanford, University of California Berkeley, University of Michigan, University of Texas Austin
 - ...aaaaand the University of Toronto ☺

The Watson Competition



Cognitive Computing

- Systems that learn + unstructured data
- Processes:
 - Understand natural language
 - Generate and evaluate hypotheses
 - Adapt and learn (and repeat)
- Characteristics:
 - Understanding, engaging, learning/improving, discovery, trust, insight, timely.

How Watson works

- Training (Watson Experience Manager)
 - Document uploading
 - Supported files: .txt, .html, .doc, .pdf
 - Corpus creation
 - Domain expert training
- Question Answering
 - Submit question text to Watson's QAAP
 - (instructions in Getting Started Guide)

Things to know about Watson

- Watson is a black box.
- Passage-style answers.
- Single corpus for the class.
- Internally consistent corpus documents.
- No structured data.



Watson competition

- Final competition: January 9, 2015
 - One team from each school
 - Held in 51 Astor Place (IBM New York office)
- Each school selects a single team
 - Competition based on both the team's product & business plan.
 - Panel is made of IBM execs and VCs.
 - See rubric for details.
- \$100,000 award for winning team.

The Watson Course



Let's win us a contest.

- Course breakdown:
 - Crash course in Watson
 - Watson Hackathon (Experience Manager)
 - Corpus creation
 - Development milestones
 - Product Pitch
 - Project Design
 - Minimum Viable Product
 - Software Beta
 - Final Product
 - Project Showcase

How CSC490 works

- Meetings, not lectures.
- Industry mentors, advisors and speakers.
 - IBM special guests
- Agile production cycle
- Corpus domain?
 - Natural language data
 - Domain expertise needed
 - Keyword search insufficient
 - Easy (and free) access



Watson Hackathon Day



Corpus Training Boot Camp

- Saturday, Sept 13
- 10am – 2pm
- Learn how to:
 - Use the Watson Experience Manager to create and train a corpus.
 - Execute queries on the trained corpus.
- Meet at Digital Innovation Lab space, 2nd floor of Gerstein Library.
 - Pizza provided (let us know dietary restrictions)



What to prepare

- CanLII website: <https://www.canlii.org/en/>
- Corpus training:
 - Group 1: Criminal Code
 - Group 2: Income Tax Act
 - Group 3: Canada Labour Code
 - Group 4: Constitution Act
 - Group 5: Highway Traffic Act
 - Group 6: Family Law Act
- End goal: Use QAAPi to return answers to queries on your domain.