

Intelligent Systems (AI-2)

Computer Science cpsc422, Lecture 34

Apr, 14, 2021

Teaching Evaluations

- **Please complete the teaching evaluations on Canvas**
- **We will start the actual class at 4:10**
- **Thank you!**



Lecture Overview

- **→ IBM Watson... And Debater**
- **After 422...**
- **(422) Highlights from conferences** ~~UI~~
- **Final Exam: how to prepare...**

JEEPARDY!



The IBM Challenge

→ 2014 →

Watson : analyzes natural language questions and content well enough and fast enough to compete and win against champion players at Jeopardy!

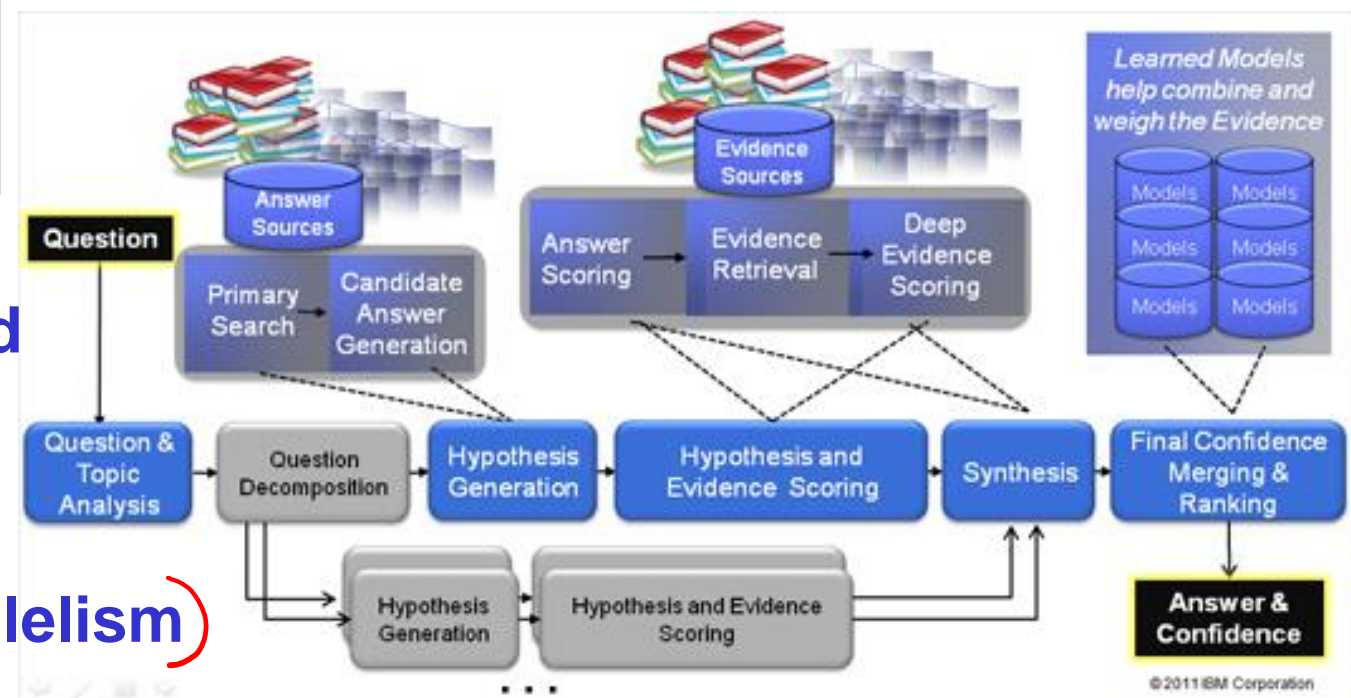
Clue: "This Drug has been shown to relieve the symptoms of ADD with relatively few side effects."

Answer: "What is Retail?"

- 1000s of algorithms and KBs

- 3 secs

- Massive parallelism



Source: IBM

AI techniques in 422 / Watson

- Parsing (**PCFGs**)
- Shallow parsing (NP segmentation with **CRFs**)
- Entity and relation Detection (NER with **CRFs**)
- **Logical** Form Generation and Matching
- **Logical** Temporal and Spatial Reasoning
- Leveraging many databases, taxonomies, and **ontologies** (help only 25% of questions)
- Confidence... **Probabilities** (**Bnets** to rank)
- Strategy for playing *Jeopardy*... statistical models of players and games, game-theoretic analyses and application of **reinforcement-learning** (**Buzz-in - Bets**)

From silly project to \$1 billion investment

2005-6 “IT’S a silly project to work on, it’s too gimmicky, it’s not a real computer-science test, and we probably can’t do it anyway.” These were reportedly the first reactions of the team of IBM researchers challenged to build a computer system capable of winning “Jeopardy”!

.....after 8-9 years...

On January 9th 2014, with much fanfare, the computing giant announced plans to invest \$1 billion in a new division, IBM Watson Group. By the end of the year, the division expects to have a staff of 2,000 plus an army of external app developersMike Rhodin, who will run the new division, calls it “one of the most significant innovations in the history of our company.” Ginni Rometty, IBM’s boss since early 2012, has reportedly predicted that it will be a \$10 billion a year business within a decade.

More complex questions in the future....



NLG


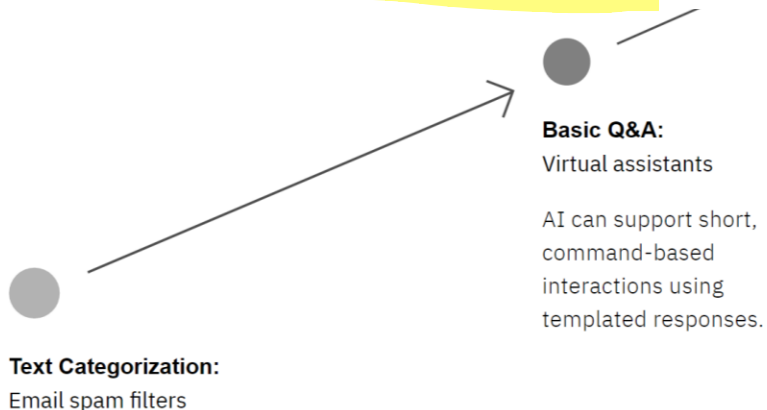
Or something like: “Should Europe reduce its energy dependency from Russia and what would it take?”

And interactive, collaborative question-answering / problem solving

Why is Project Debater important?

The rise of one-sided and doctored narratives is challenging society and our platforms. Too often, we talk past one another. We need a smarter way. New developments in language and reasoning in AI can help shine a light in the darkness of distorted facts to provide diverse, well-informed viewpoints—both the pro and the con.

The world is awash with information, misinformation, and superficial thinking. Project Debater pushes the frontiers of AI to facilitate intelligent debate so we can build well-informed arguments and make better decisions.



IBM Debater Architecture

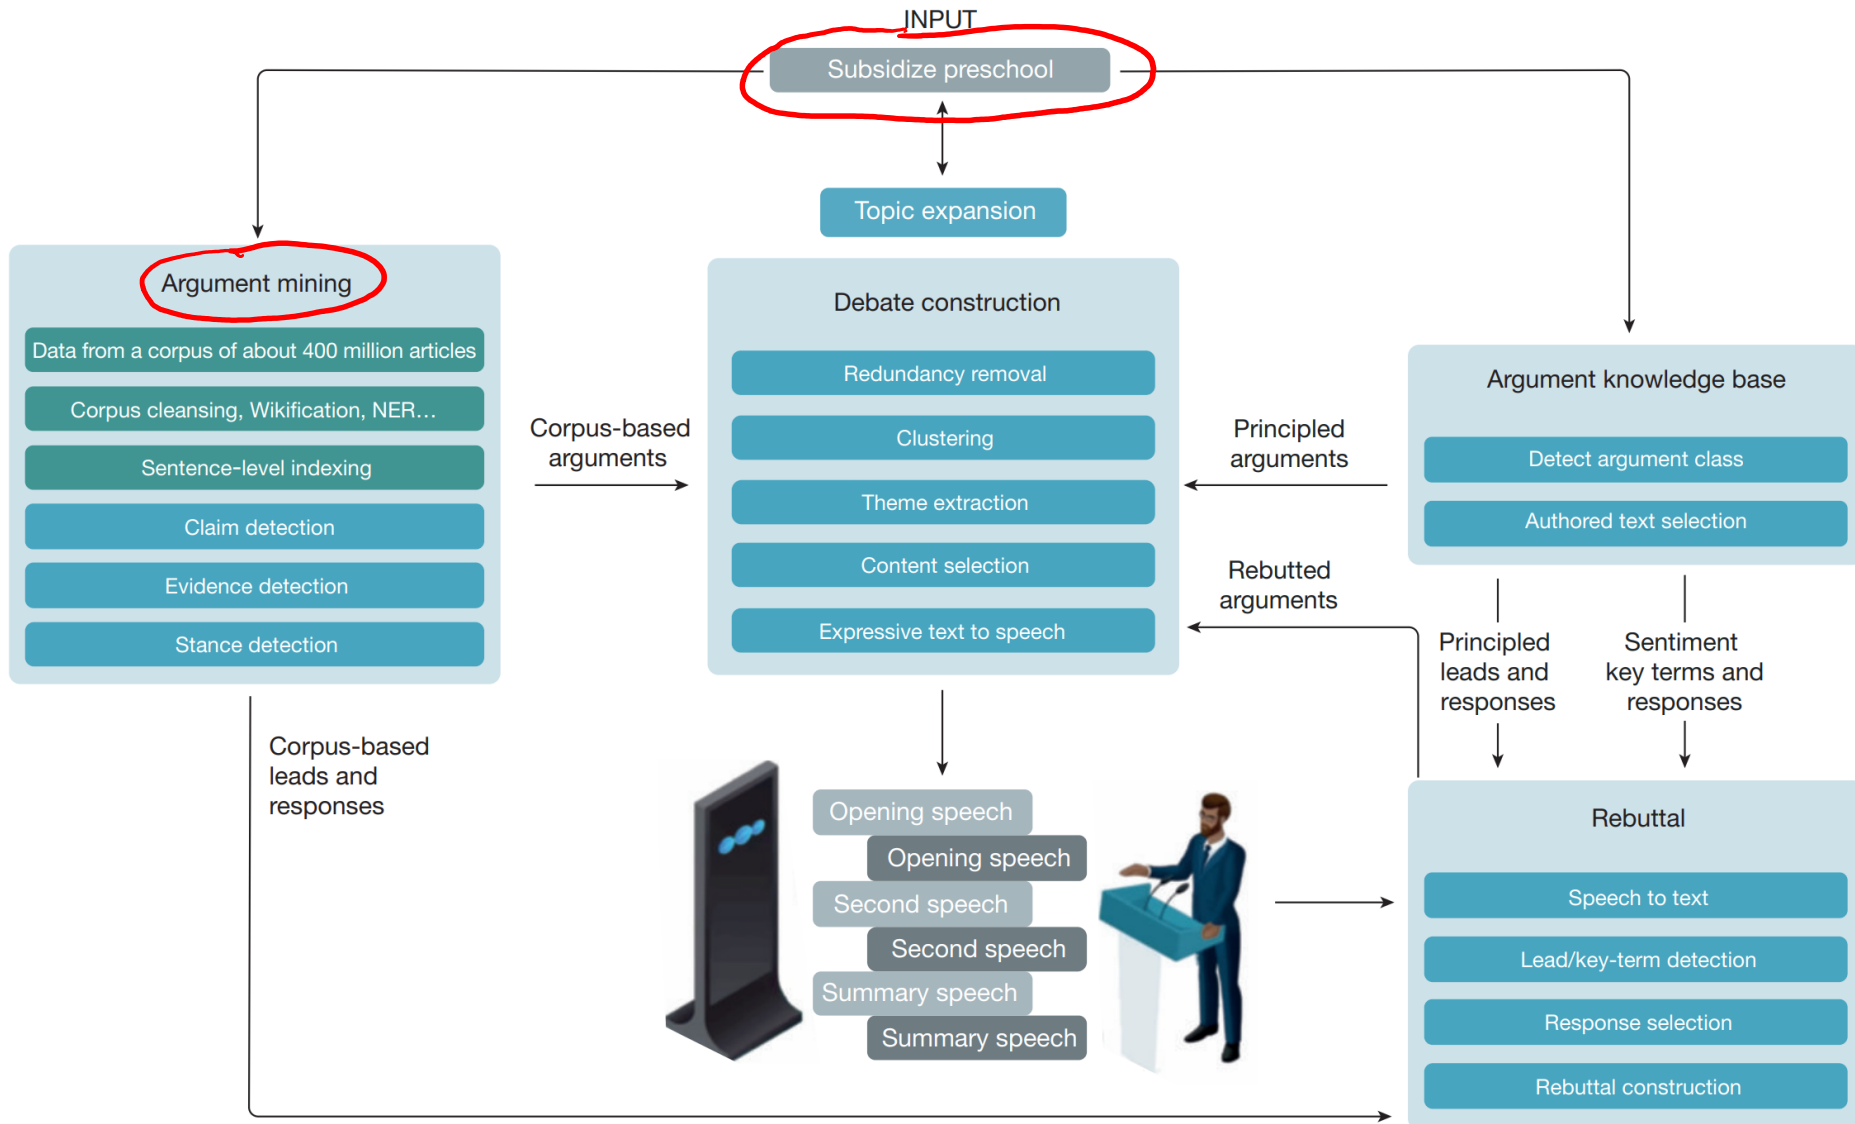


Fig. 2 | System architecture. Description of Project Debater components. Offline analysis is shaded in green; online analysis is shaded in blue. NER stands for named entity recognition.

AI applications.....

- DeepQA
- Personal Assistants
- Robotics
- Search Engines
- Games
- Tutoring Systems
- Medicine / Finance / ...
-

Most companies are investing in AI and/or developing/adopting AI technologies

Lecture Overview

- TA Evaluations / Teaching Evaluations
- IBM Watson...
- **After 422...**
- **(422) Highlights from conferences**
- **Final Exam: how to prepare...**

422 big picture

StarAI (statistical relational AI)
 Hybrid: Det +Sto
 Prob CFG
 Prob Relational Models
 Markov Logics

Deterministic

Stochastic

Query

Planning

Logics
Full Propositional
First Order Logics

Ontologies –
Knowledge Graphs

- Full Resolution
- SAT
- Concept Similarity

Belief Nets

Approx. : Gibbs

Markov Chains and HMMs

Forward, Viterbi....

Approx. : Particle Filtering

Markov Networks
Conditional Random Fields

Markov Decision Processes
and
Partially Observable MDP

- Value Iteration
- Approx. Inference

Reinforcement Learning
Q-learning, SARSA

Representation

Reasoning
 Technique

Applications of AI

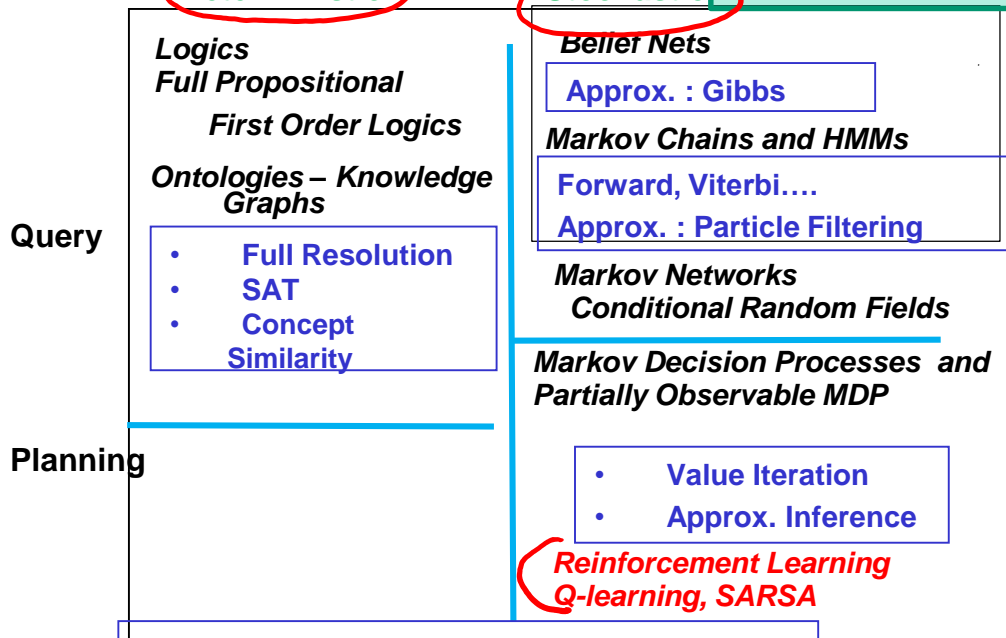
After 422.....

Machine Learning + Neural Models Knowledge Acquisition Preference Elicitation

StarAI (statistical relational AI)
Hybrid: Det +Sto
Prob CFG
Prob Relational Models
Markov Logics

~~Deterministic~~

Stochastic



Applications of AI

Representation
Reasoning
Technique

Where are the components of our representations coming from?

The probabilities?
The utilities?
The logical formulas?

From people and from data!

Some of our Grad Courses

522: Artificial Intelligence II : Reasoning and Acting Under Uncertainty

Sample Advanced Topics.....

Relational Reinforcement Learning for Agents in Worlds with Objects, relational learning.

- Probabilistic Relational Learning and Inductive Logic Programming at a Global Scale,

Some of our Grad Courses

503: Computational Linguistics I / Natural Language Processing

Sample Advanced Topics.....

- Topic Modeling (LDA) – Large Scale Graphical Models – Gibbs Sampling
- Discourse Parsing by Distant Supervision (CKY + Neural Nets) (*third 422 reading*) ←
- Abstractive Summarization (Neural Nets + Concept similarity)

Other AI Grad Courses: check them out

532: Topics in Artificial Intelligence (different courses)

- Human-Centered AI ←
- Foundations of Multiagent Systems
- Deep Learning for Computer Vision and Graphics
- Multimodal Learning with Vision, Language and Sound
- Probabilistic Programming

540: Machine Learning

349

505: Image Understanding I: Image Analysis


525: Image Understanding II: Scene Analysis

515: Computational Robotics

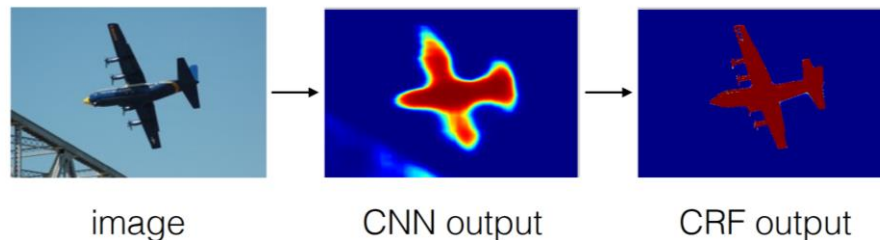
Connection with Neural Models / Deep-Learning

NN + Reinforcement Learning: e.g. Alpha-Go (Deep Mind)


<http://karpathy.github.io/2016/05/31/rl/>



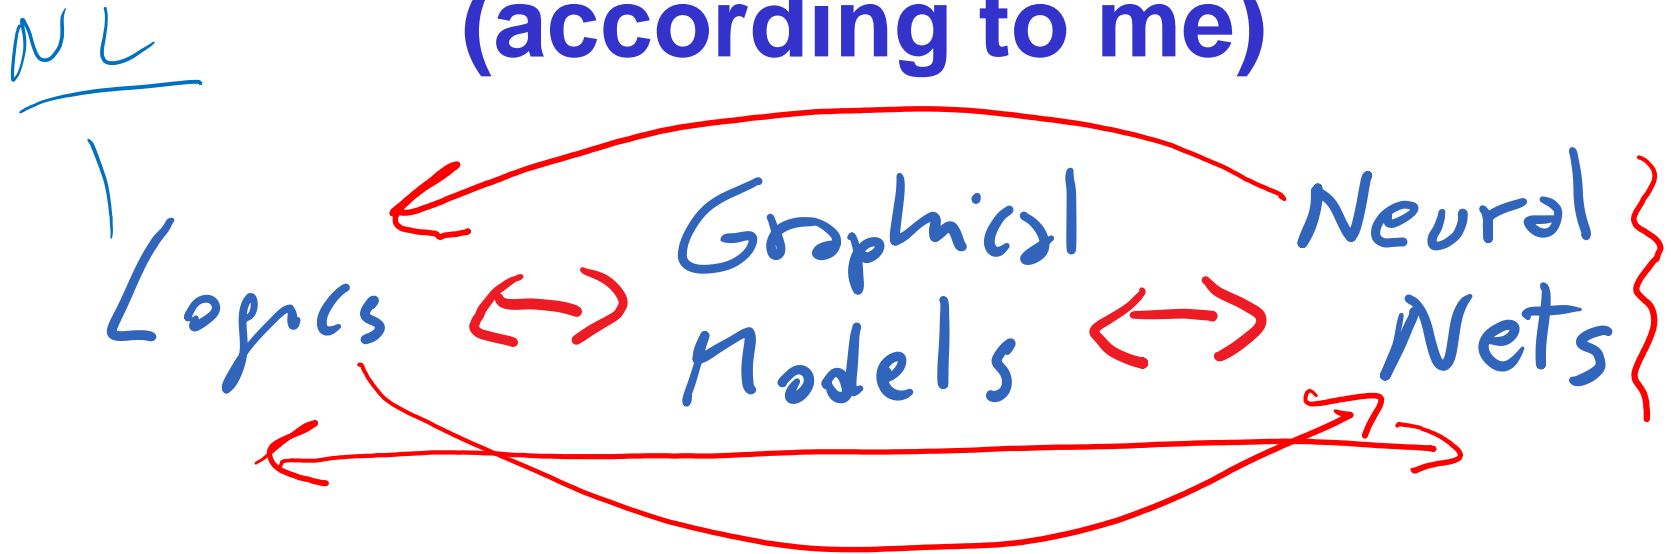
NN (rough prediction) + CRF (refine prediction): Vision (Google paper)



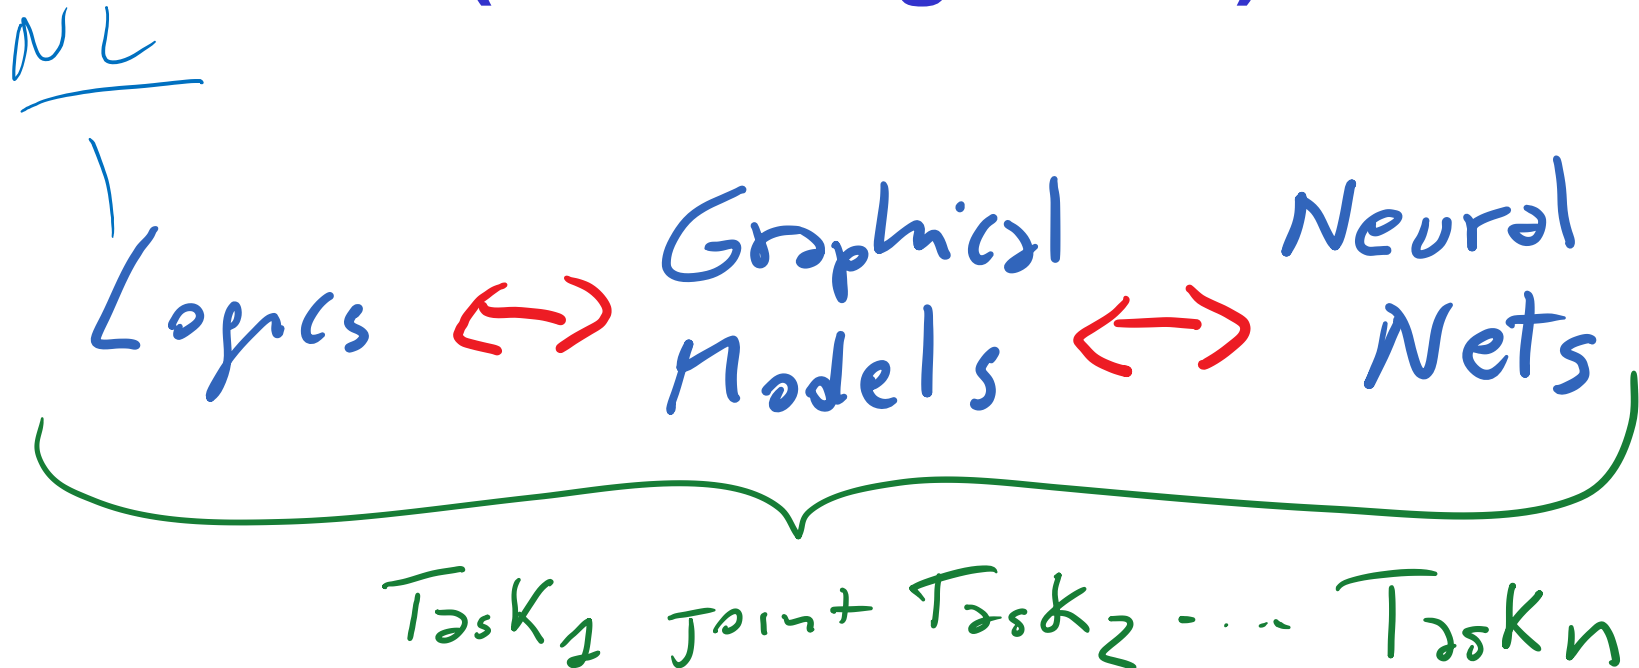
NN + CRF + CKY for NLP (Goldberg book. 2017, Pag 224)



AI Representation & Reasoning & Learning future in a nutshell (according to me)



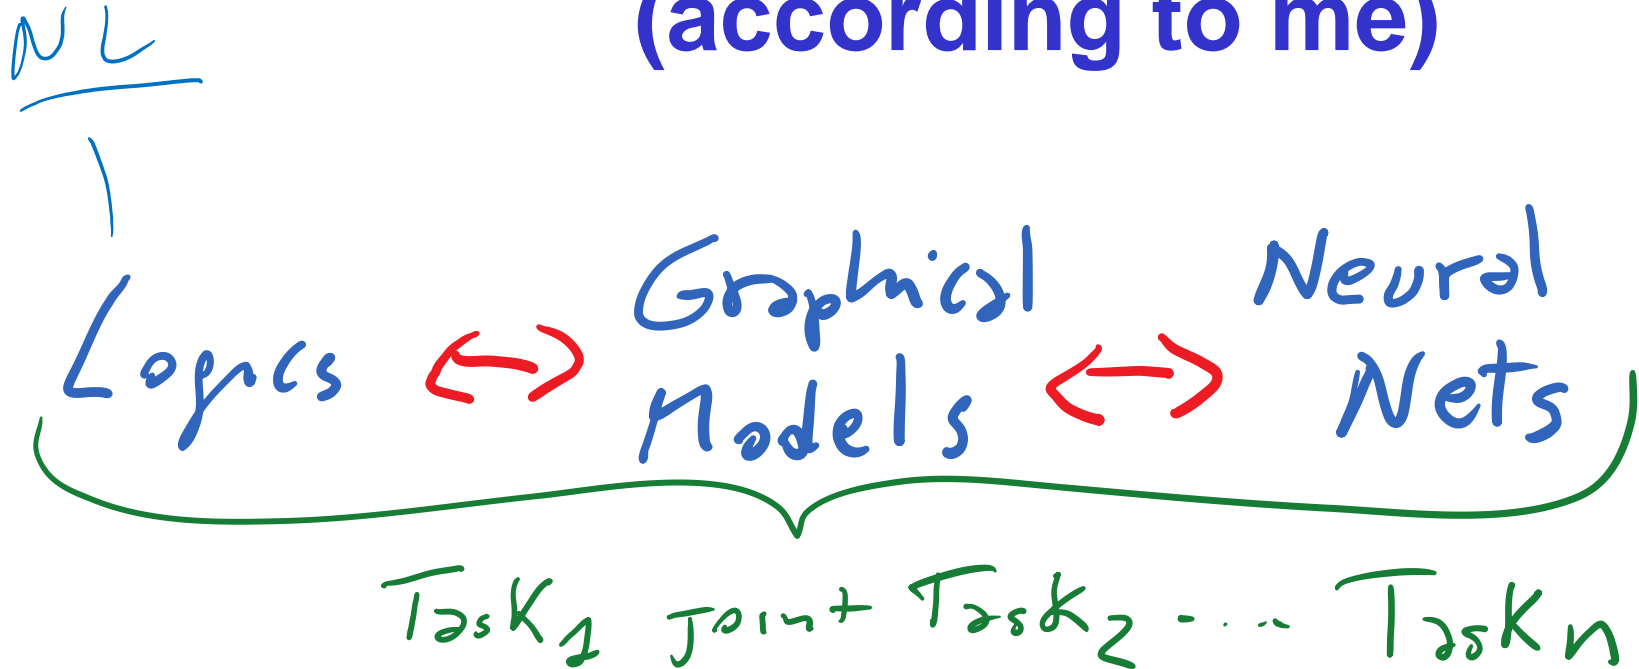
AI R&R&L future in a nutshell (according to me)



How to manage multi-tasking for an increasing number of tasks? **Multi-task**

Machine Translation, Text Summarization, Event recognition in videos

AI R&R&L future in a nutshell (according to me)



How to manage multi-tasking for
an increasing number of tasks?

Multi-task R&R&L

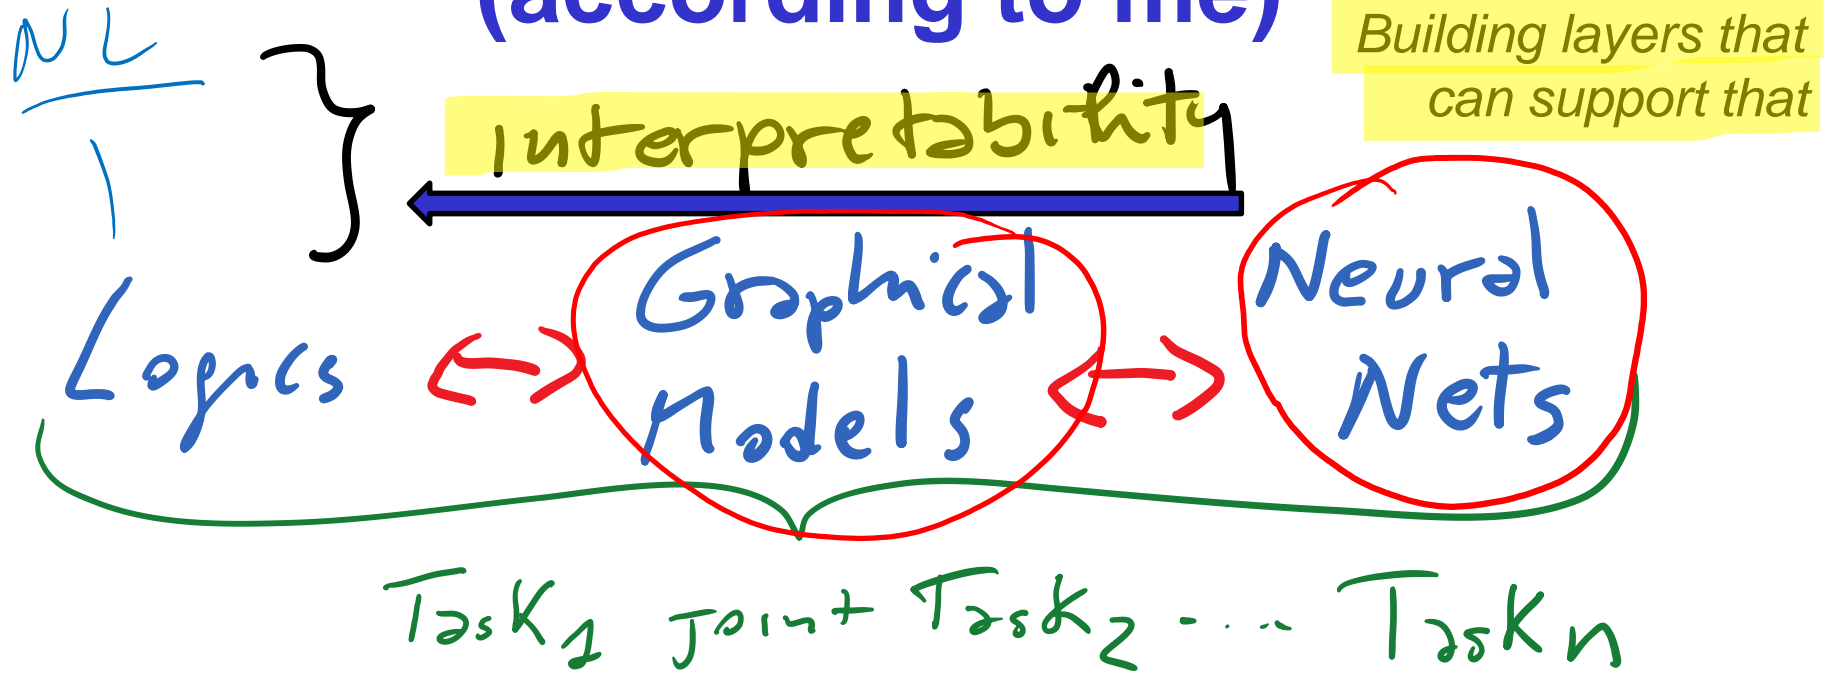
In different domains

*Medicine, Transportation,
Politics, Law...*

*Machine Translation, Text
Summarization, Event
recognition in videos*

Transfer learning

AI R&R&L future in a nutshell (according to me)



How to manage multi-tasking for
an increasing number of tasks?

Multi-task

In different domains

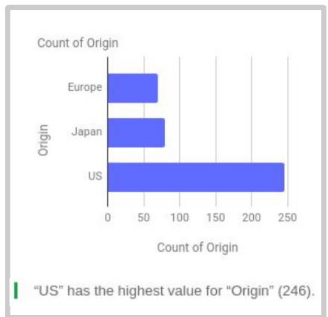
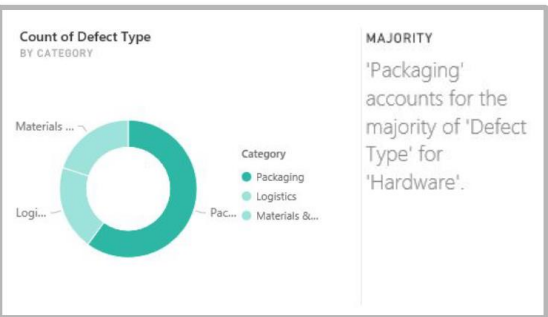
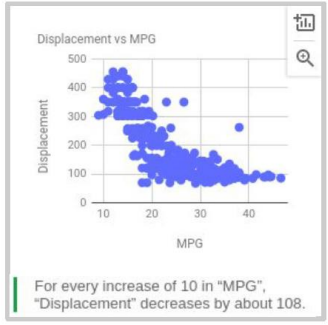
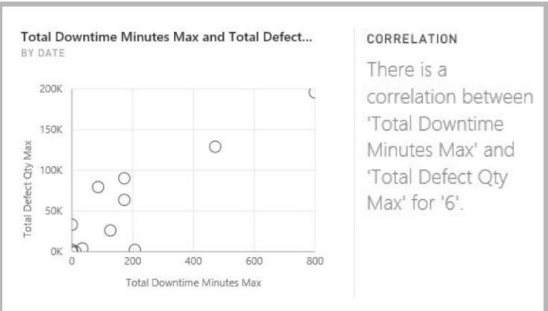
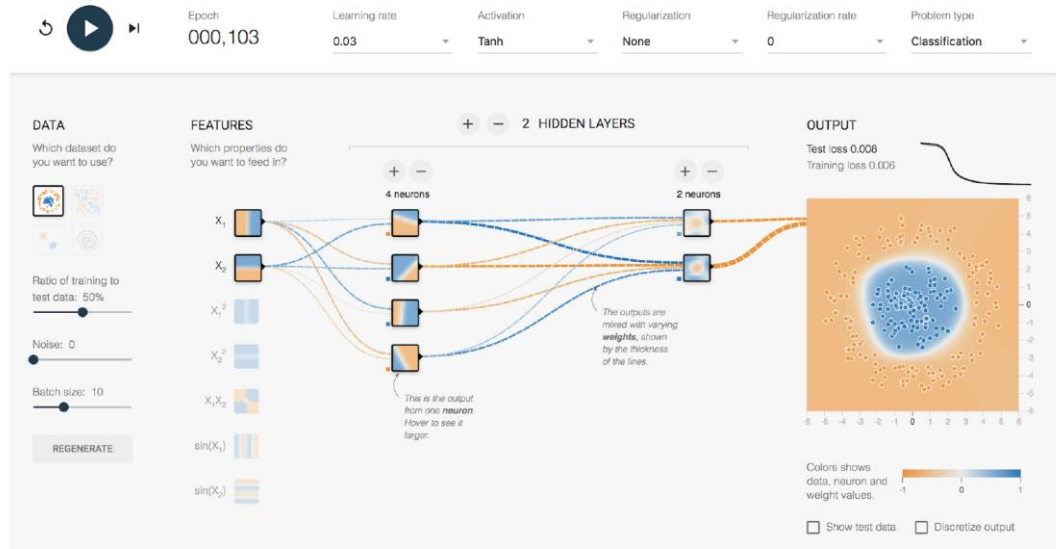
Medicine, Transportation,
Politics, Law...

Machine Translation, Text
Summarization, Event
recognition in videos

Transfer learning

Key to interpretability

Language +
mathematical
abstractions
+ Visualization
(we have a ugrad and a grad course)



Lecture Overview

- TA Evaluations / Teaching Evaluations
- IBM Watson....
- After 422....
- **(422) Highlights from IUI conferences and more....**
- **Final Exam: how to prepare.....**



AI and HCI meet

Keynote Speaker:

Prof. Dan Weld, University of Washington

Intelligent Control of Crowdsourcing

Crowd-sourcing labor markets (e.g., Amazon Mechanical Turk) are booming, use of **Partially-Observable Markov Decision Processes (POMDPs)** to control voting on binary-choice questions and iterative improvement workflows.

... applications range from photo tagging to audio-visual transcription and interlingual translation

Some papers from IUI

Unsupervised Modeling of Users' Interests from their Facebook Profiles and Activities

Preeti Bhargava (*University of Maryland*)

Oliver Brdiczka (*Vectra Networks, Inc.*)

Michael Roberts (*Palo Alto Research Center*)

named entity recognition, document categorization, sentiment analysis, semantic relatedness and social tagging

Semantic Textual Similarity (STS) system [13] for computing the SR scores. STS is based on LSA along with WordNet knowledge and is trained on LDC Gigawords and Stanford Webbase corpora

Some papers from IUI-15

BayesHeart: A Probabilistic Approach for Robust, Low-Latency Heart Rate Monitoring on Camera Phones

Xiangmin Fan (*University of Pittsburgh*)

Jingtao Wang (*University of Pittsburgh*)

BayesHeart is based on an adaptive hidden Markov model, requires minimal training data and is user-independent.

Two models, one with 2 states and one with 4 states, work in combination....

.... measuring people's heart rates through commodity cameras by capturing users' skin transparency changes, color changes, or involuntary motion

Applications: gaming, learning, and fitness training

Paper for IUI 2017

Analyza: Exploring Data with Conversation (Google research)

Applied in two systems

- (a) Question answering for a spreadsheet product.
- (b) provides convenient access to a revenue/inventory database for a large sales force.....

Support users who do not necessarily have coding skills....

We also derive an additional lexicon for entities in our

knowledge. We use a **context-free grammar** to parse with a the annotated query. The grammar

knowledge rules are written in different types.

We used multiple ways of establishing **semantic similarity** between a question term and the lexicon.

Another UI 2017 paper

"How May I Help You?": Modeling Twitter Customer Service Conversations Using Fine-Grained Dialogue Acts
(IBM, U of California..)

| Tag | Example |
|---------------------------|--|
| Statement Informative | The signal came back last night [...] |
| Request Information | Can you send us [...]? |
| Statement Complaint | Staff didn't honor online info, was dismissive [...] |
| Question Yes-No | Have you tried for availability at [...] |
| Statement Expressive Neg. | I don't trust places that do bad installations [...] |
| Statement Suggestion | Let's try clearing the cache <link> [...] |
| Answer (Other) | Depends on the responder [...] |
| Social Act Apology | I'm sorry for the trouble [...] |
| Social Act Thanks | Thanks for the help [...] |
| Question Wh- | Why was that? |
| Statement Offer | We can always double check the account [...] |
| Question Open | How come I can't get a [...] quote online? |
| (All Other Acts) | |

We develop an **SVM-HMM** model to identify dialogue acts in a conversation, in a real-time setting, and using a novel **multi-label approach** to capture different dialogic intents contained in a single turn.

RL is booming... (2020)

Applied in CS itself... eg in Software Engineering
(she was a ugrad here at UBC and finishing her
PhD at Berkeley ;-)

Sameer Reddy, Caroline Lemieux, Rohan Padhye,
Koushik Sen. **Quickly Generating Diverse Valid
Test Inputs with Reinforcement Learning**. *In
Proceedings of the 42nd International Conference
on Software Engineering, ICSE
2020.* pdf, slides, video.

Understanding basic RL and underlying
representations MDPs POMDPs is critical !

Last Clicker Question

I would like to learn more about AI....

- A. Yes
- B. Maybe
- C. No



Lecture Overview

- TA Evaluations / Teaching Evaluations
- IBM Watson....
- After 422....
- (422) Highlights from IUI conferences
- **Final Exam: how to prepare.....**

Final Exam

Time: Fri Apr 23 12:00 PM 2:30

Location: zoom

How to prepare....

- **Learning Goals** (posted on Connect): Each LG corresponds to one or more possible questions
- Revise all the **clicker questions**, practice exercises, **assignments** and **midterm** !
- Will post more practice material ✓
- Office Hours – My and TA office hours usual schedule. May add a few closer to exam
- Post Questions on Piazza

Can use a letter sized sheet of paper with anything written on it (double sided)