Beyond Deep Learning : Learning+Reasoning

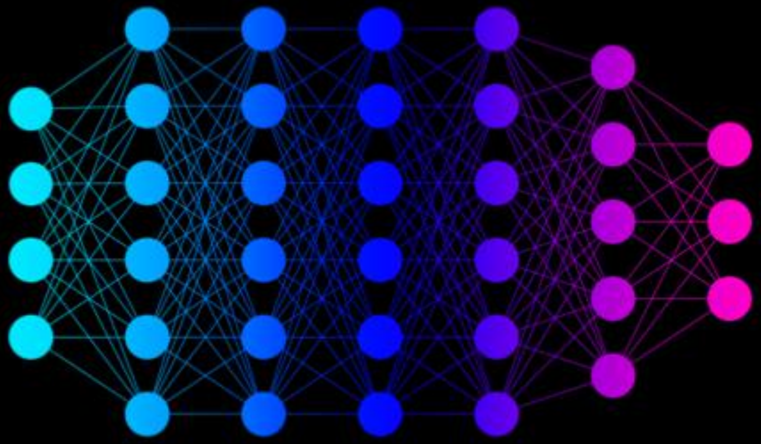

Lisa Amini

Director, IBM Research Cambridge,

Acting Director, MIT-IBM Watson AI Lab

MIT 6.S191 Intro to Deep Learning

2011, IBM Watson computer wins human champions at Jeopardy!

2017, World’s first 50 qubit quantum computer

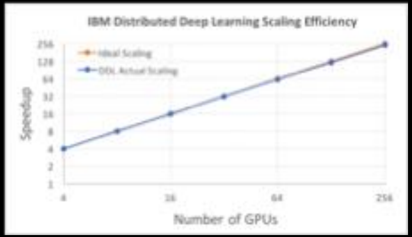

2017, IBM demonstrates 95% scaling efficiency on Caffe deep learning framework

2017, quantum algo efficiently computes lowest energy state of small molecules.

Leading corporate institution for high-quality science

IBM Research

3000 creative, scientific and technical minds worldwide

6 Nobel Laureates

10 National Medals of Technology

5 National Medals of Science

6 Turing Awards

The MIT-IBM Watson AI Lab

$240M 10 year commitment to jointly create the future of artificial intelligence

- Fundamental advances in AI algorithms
- Physics of AI
- AI Transforming Industries:
  Healthcare, Life Sciences & Cybersecurity
- Advancing Shared prosperity through AI

http://mitibmwatsonailab.mit.edu/

Moments in Time

Landmark 1 Million video dataset to transform AI Vision

Pushing

Carrying

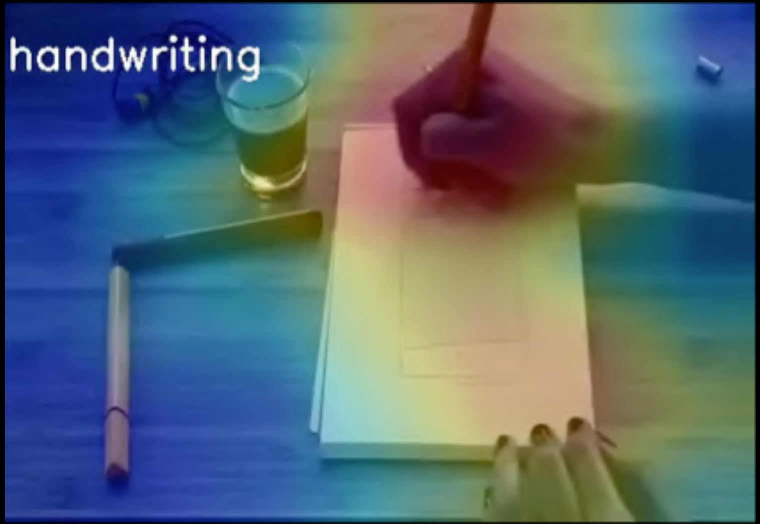

- Three seconds events
- Open access
- Goal: Recognizing and understanding actions in video

http://moments.csail.mit.edu/

Recent successes in Deep Learning are awe-inspiring, but epic breakthroughs are still needed for Machine Intelligence

- Humans learn without a lot of labeled data per task
  - *Why can't machines?*

- People learn continuously throughout their lives, remembering what they've learned and leveraging it for new tasks
  - *Current algorithms suffer from catastrophic forgetting and are unable to recognize and generalize to analogous situations or tasks*

- To interact sensibly with humans, machines must be able to remember, reason, explain, and seek to fill knowledge gaps

- **Learning+reasoning**

Making Language Computational

- *Word* Embeddings
  - Represent words as a real-valued vector in some abstract space
  - Goal: representations that capture multiple degrees of similarity

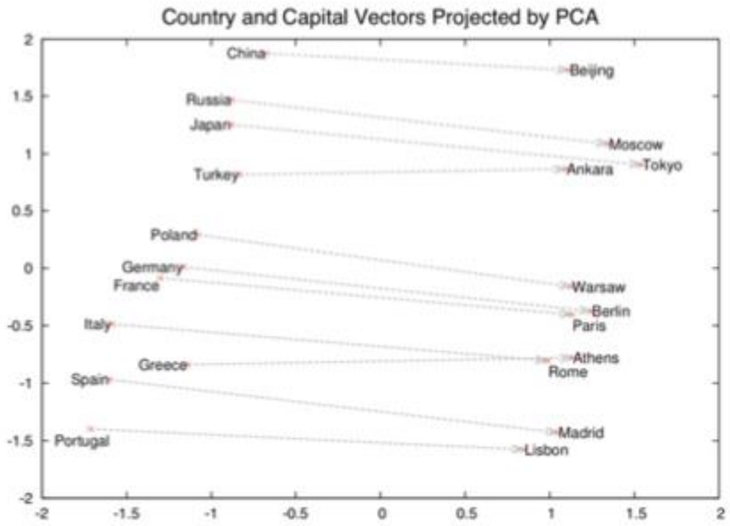

Skip-gram model

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{-c\leq j\leq c, j\neq 0} \log p(w_{t+j}|w_t)$$

Maximize the average log probability of predicting surrounding words

Distributed representations of words and phrases and their compositionality, Mikolov, et al, 2013

Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors, Baroni, Dinu, Kruszewski, 2014

FastText – Open library for unsupervised learning of word embeddings. http://fasttext.cc

Embeddings Impact on Automated Knowledgebase Construction (AKBC)

Formal knowledge representation

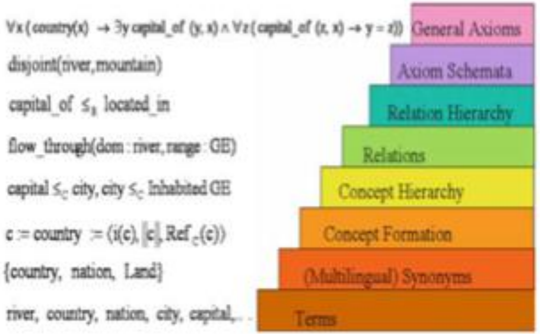

Freebase relation types	Incompleteness
/people/person/education	0.792
/people/person/employment_history	0.923
/people/person/nationality*	0.785
/people/person/parents*	0.988
/people/person/place_of_birth*	0.938
/people/person/places_lived*	0.966

Distant Supervision for Relation Extraction with an Incomplete KB, Min, Grishman, Wan, Wang, Gondek, 2013

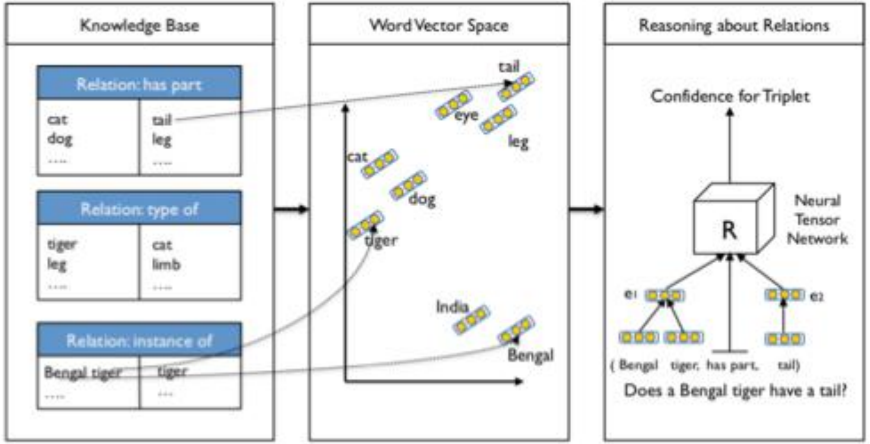

Reasoning with Neural Tensor Networks for Knowledge Base Completion, Socher, et al, 2013

Compositional Vector Space Models for Knowledge Base Completion, Neelakantan, Roth, McCallum, 2015

Embeddings Impact on Automated Knowledgebase Construction (AKBC)

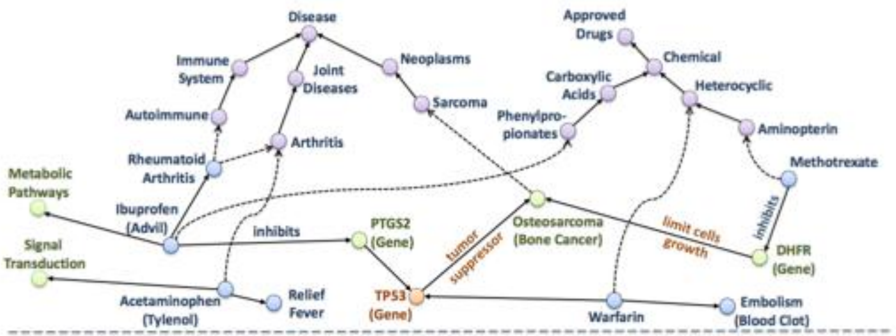

DrugBank: Bioinformatics & Cheminformatics Resource

Drug Name	Drug Targets (Genes)	Symptomatic Treatment
Ibuprofen	PTGS2	Rheumatoid Arthritis
Acetaminophen	PTGS2	Relief Fever
Methotrexate	DHFR	Antineoplastic Anti-metabolite
Warfarin	TP53	Embolism (Blood Clot)

CTD: Comparative Toxicogenomics Database

Gene	Interaction
PTGS2	TP53 (Gene)

Gene	Disease
TP53	Osteosarcoma

Chemical	Pathways	Linked Data Source
Ibuprofen	Metabolic Pathways	KEGG
Acetaminophen	Signal Transduction	Reactome
Methotrexate	Immune System	Reactome

Uniprot: Universal Protein Resource

Gene	Function
TP53	Tumor Suppressor
DHFR	Limits Cell Growth

Predicting Drug-Drug Interactions Through Large-Scale Similarity-based Link Prediction, Fokoue, et al 2016

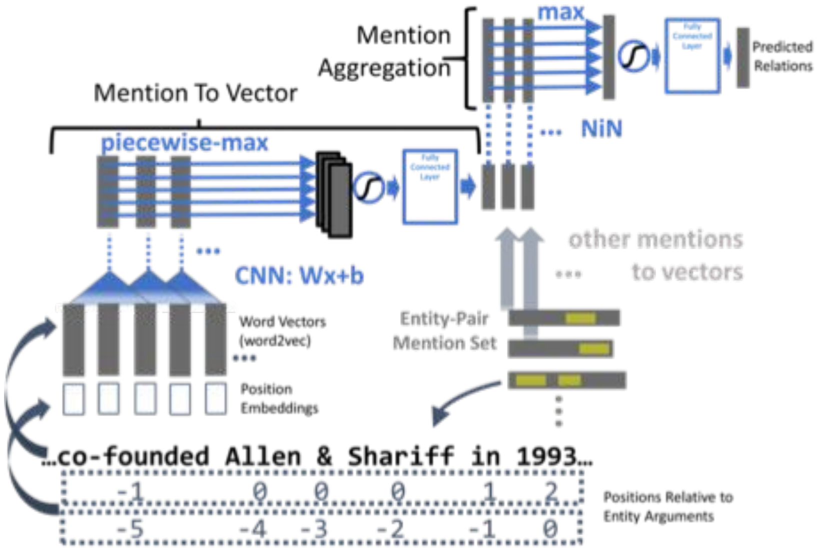

Relation prediction with confidence, leveraging disparate structured and unstructured data

Socrates: Deep Relational Knowledge Induction, Glass, et al, 2017

!st place winner: ISWC Semantic Web Challenge on AKBC

How to create differentiable machines to reason leveraging learned external knowledge bases?

Example Task: Question Answering with Long-term Memories

Sam walks into the kitchen.
Sam picks up an apple.
Sam walks into the bedroom.
Sam drops the apple.
Q: Where is the apple?
A. Bedroom

Brian is a lion.
Julius is a lion.
Julius is white.
Bernhard is green.
Q: What color is Brian?
A. White

Mary journeyed to the den.
Mary went back to the kitchen.
John journeyed to the bedroom.
Mary discarded the milk.
Q: Where was the milk before the den?
A. Hallway

Towards AI-Complete Question Answering: A set of pre-requisite toy tasks, Weston, et al , 2015

Question Answering with External Memories

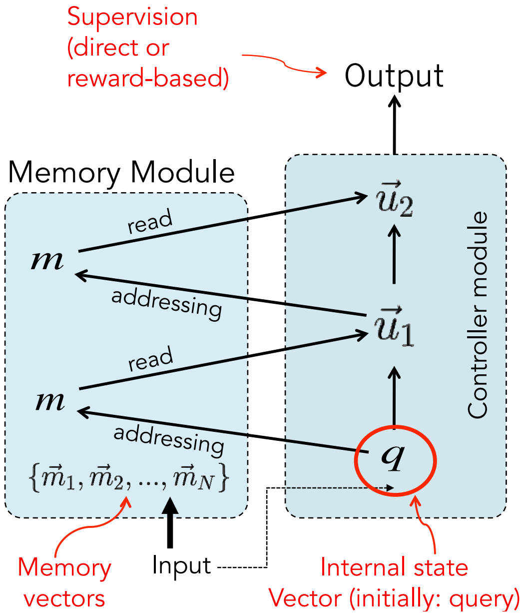

Example Inputs

```
Mary journeyed to the den.
Mary went back to the kitchen.
John journeyed to the bedroom.
Mary discarded the milk.
Q: Where was the milk before the den?
A. Hallway
```

Jointly trained with Inputs (I ➔ *x* ➔ *m*), Questions (Q➔*q*), Answer (*u* ➔ *o*)

Memory Networks, Weston, et al , 2015

Memory Networks for Language Understanding, ICML Tutorial, Weston, et al , 2016

Question Answering with External Memories

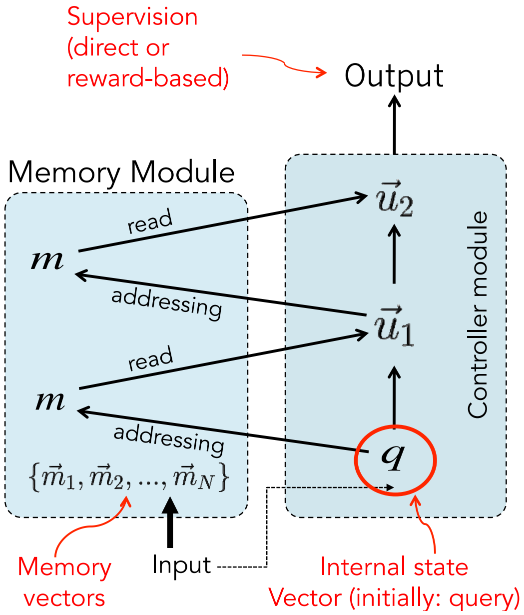

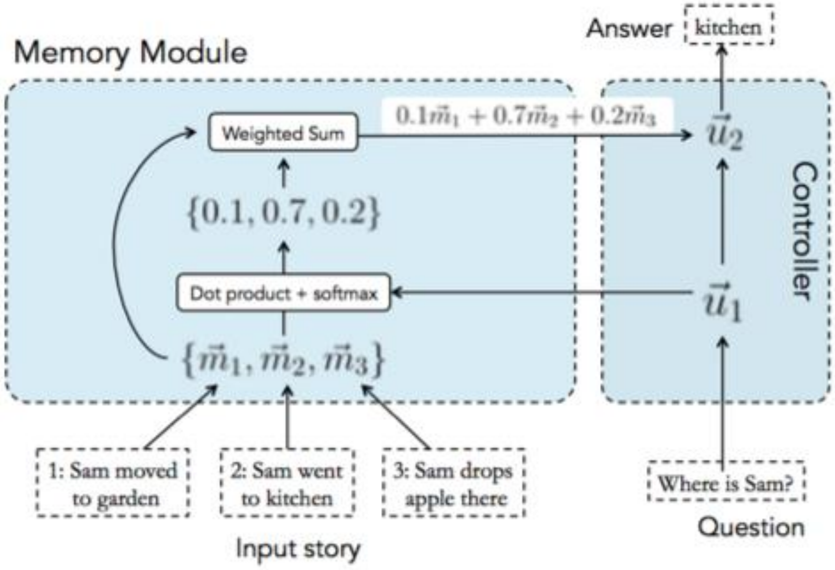

Memory Networks, Weston, et al , 2015

End-to-End Memory Networks, Sukhbaatar, et al , 2015

Memory Networks for Language Understanding, ICML Tutorial, Weston, et al , 2016

Want to Learn More?

- Improved detection of key relations KBQA

  Improved Neural Relation Detection for Knowledge Base Question Answering, Yu, et al 2016

- Simulator to generate challenge questions from ambiguous texts

  Learning to Query, Reason, and Answer Questions on Ambiguous Texts, Guo et al, 2017

- Bringing commonsense knowledge into vector space

  Lifted Rule Injection for Relation Embeddings, Demeester, Rocktaschel, Riedel, 2016

- Learning to represent and execute programs

  Neural Program Interpreters, Reed, et al, 2015

- Learning representations to induce logical rules and perform multi-hop reasoning

  End-to-end Differentiable Proving, Rocktaschel, Riedel, 2017

Want to do more?
Watson Developer Cloud

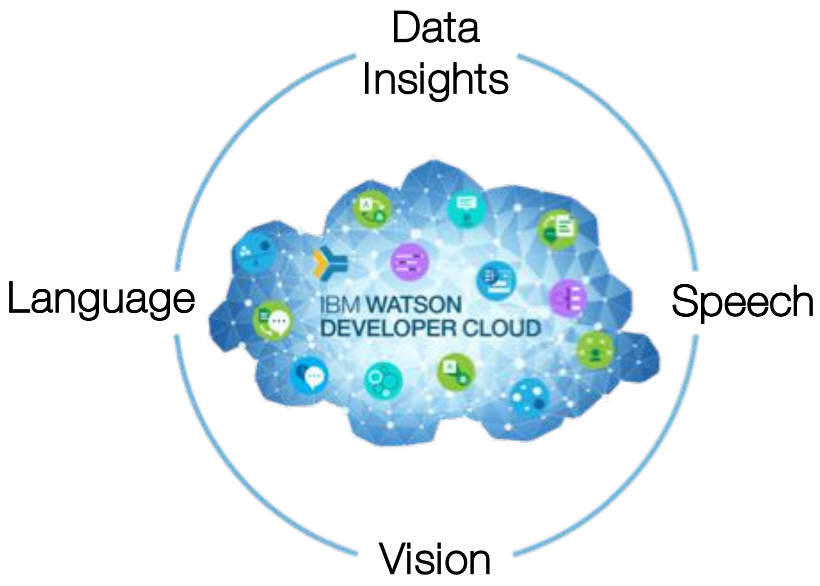

Language

– Natural Language Classifier
– Language Translator
– Personality Insights
– Tone Analyzer
– Natural Language Understanding

– Conversation – build chatbots and virtual agents across any channel and domain

Data Insights

– Discovery – aggregate and organize massive amounts of enterprise data, and answer questions in context

Speech

– Speech to Text – transcribe audio and take action
– Text to Speech – verbalize written text into understandable audio

Vision

– Visual Recognition – help people understand and take action from visual data

Thank you!