Semantics, Science, and 10 Year Olds
Oren Etzioni, AI2

ALLEN INSTITUTE
for ARTIFICIAL INTELLIGENCE

Outline

I. Machine reading & Open IE
II. The Allen Institute for AI (AI2)

III. Arithmetic story problems
IV. Diagram understanding in plane geometry
V. Knowledge & inference in 4th grade science
VI. Semantic literature search

I. Machine Reading (Etzioni, AAAI 2006)

Auto-Text to Knowledge

Source: DARPA, Machine Reading initiative

Information Extraction

IE(sentence) = [tuple, confidence]

"Edison, by all accounts, was the inventor of the light bulb."
➔ **invented(Edison, light bulb), 0.98**

Typically, IE requires:

- Pre-specified relations
- Hand-labeled training examples
- Lexicalized features/patterns

NELL Knowledge Base

categories

- relatedto
  - visualartmoveme
  - personinacadem
  - generalizationof
    - athletessucha
    - academicfields
    - chemicaltypeh
    - astronautssuc
    - automobileeng
    - airportsuchasa
    - amphibiansuc
    - aquariumssuc
    - arteriessuchas
    - animaltypehas
      - animalsucha
        - animalsuc
        - animalsuc
      - animalsucha
      - inverseofarti
    - professiontype
    - productinstan
    - booksuchasbo
    - vehicletypehas
    - touristattractio
    - celebritysucha
    - archaeasucha

obama (politician)

literal strings: OBAMA, OBama, Obama, obama

Help NELL Learn!

NELL wants to know if these beliefs are correct.
If they are or ever were, click thumbs-up. Otherwise, click thumbs-down.

- obama is a politician
- obama is a male
- obama is a person who belongs to the organization house (governmentorganization)
- obama is a person born on the date n1961 (dateliteral)
- obama is a person who graduated from the university harvard (university)
- obama is a person who graduated from the university harvard_law_school (university)

categories

- **politician**(100.0%)
  - SEAL @9 (75.0%) on 13-jan-2010 [1 2] using obama
  - CPL @722 (88.0%) on 06-apr-2013 ["_ is a good president" "_ defends choice" "_ beats McCain" massive scope" "_ is the only President" "_ is the new Hitler" "_ is half-white" "_ is an empty suit" "_ President" "_ is leading McCain" "McCain trails _" "Kennedy endorsed _" "_ 's Inauguration Day" "_ campaign" "York Times endorses _" "_ 's Inaugural Address" "more popular votes than _" "_ 's his

NELL Lexical Patterns for “haswife”

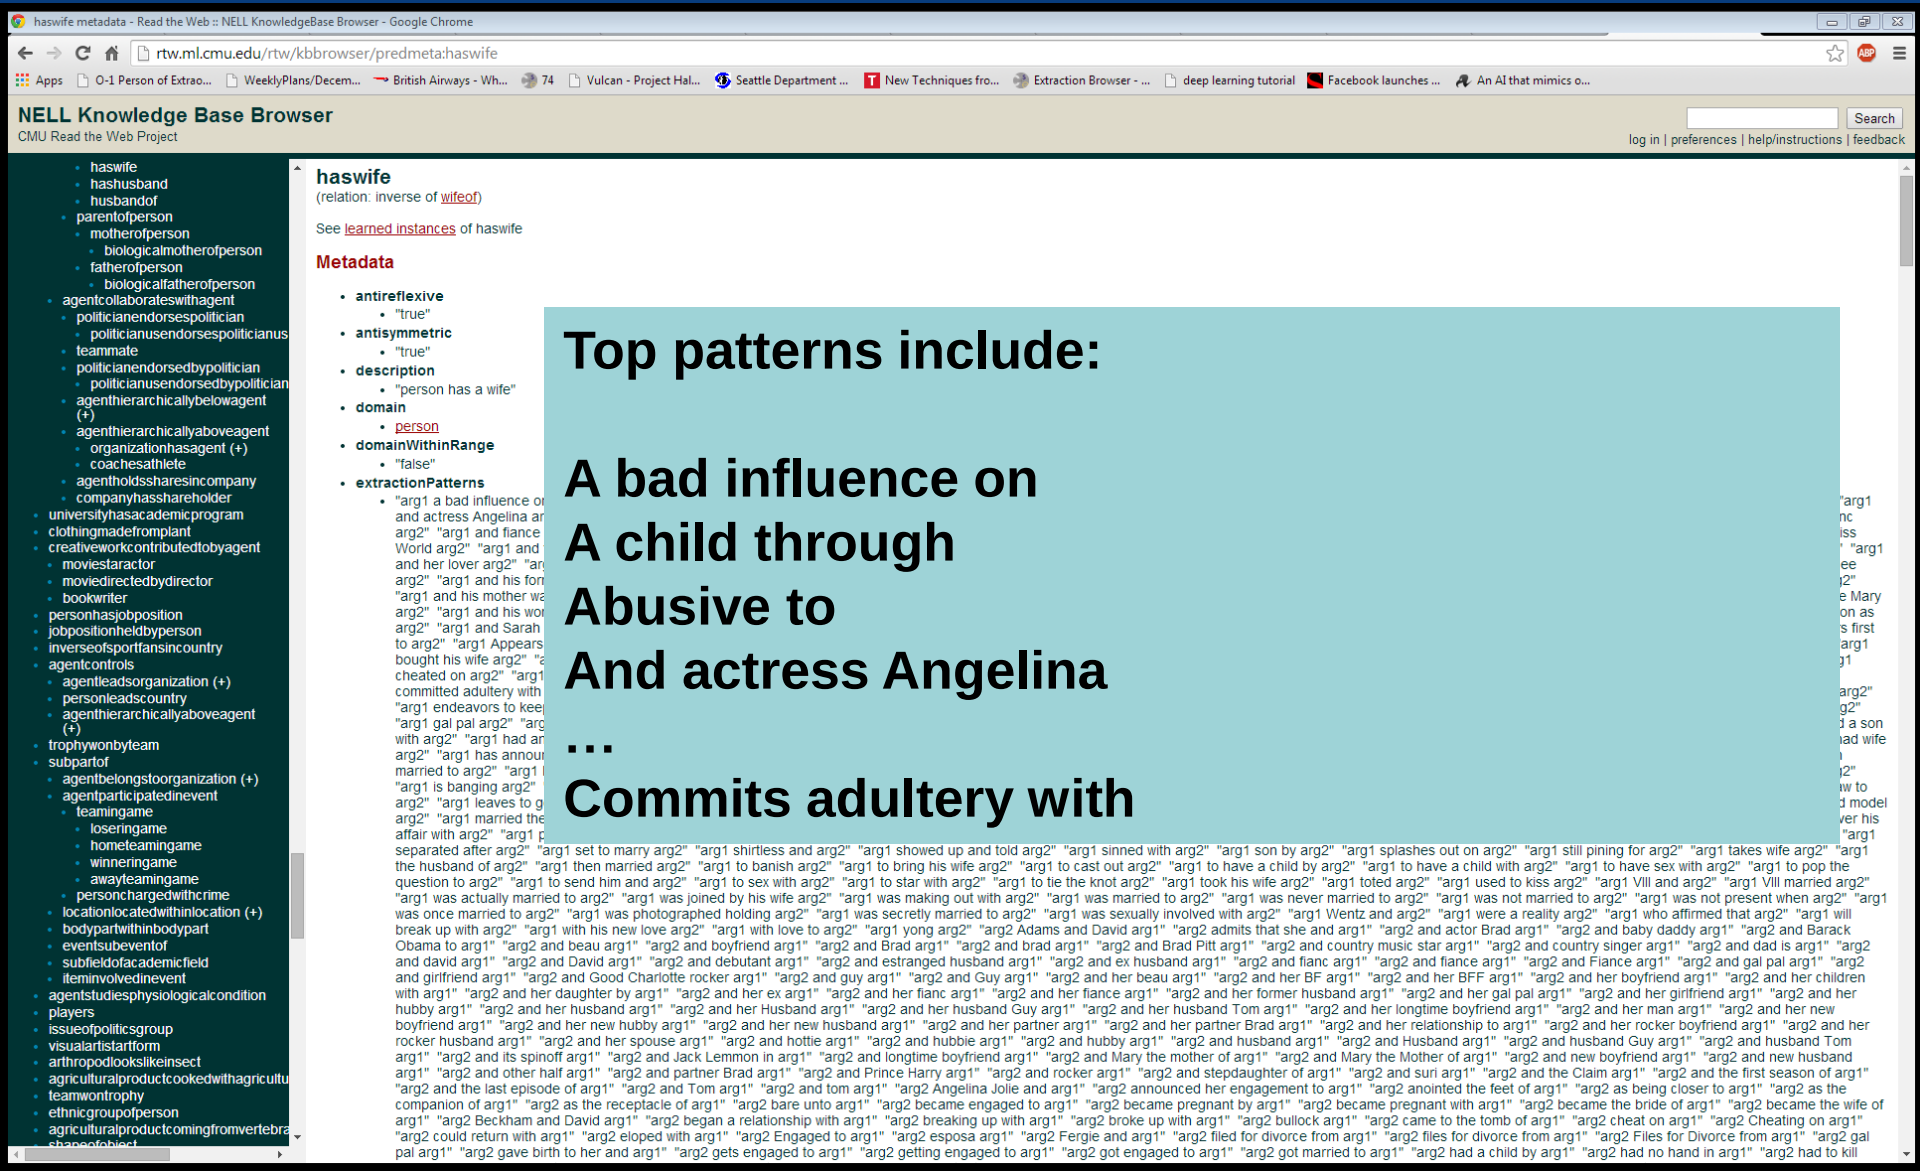

ALLEN INSTITUTE for ARTIFICIAL INTELLIGENCE

Open Information Extraction (2007)

Question: can we leverage regularities in language to extract information in a relation-independent way?

Relations often:

- anchored in verbs
- exhibit simple syntactic form

Virtues:

- No hand-labeled data
- "No sentence left behind"
- Exploit redundancy & serendipity of Web
- Robust to parser errors

History of Open IE

2007: 1.0 (TextRunner, auto-labeled examples, CRF)

2011: 2.0 (ReVerb, simple model of verb-based relations)

2012: 3.0 (Ollie, parser, verbs + nouns)

2013: 4.0 (semantic role labeling, n-ary relations)

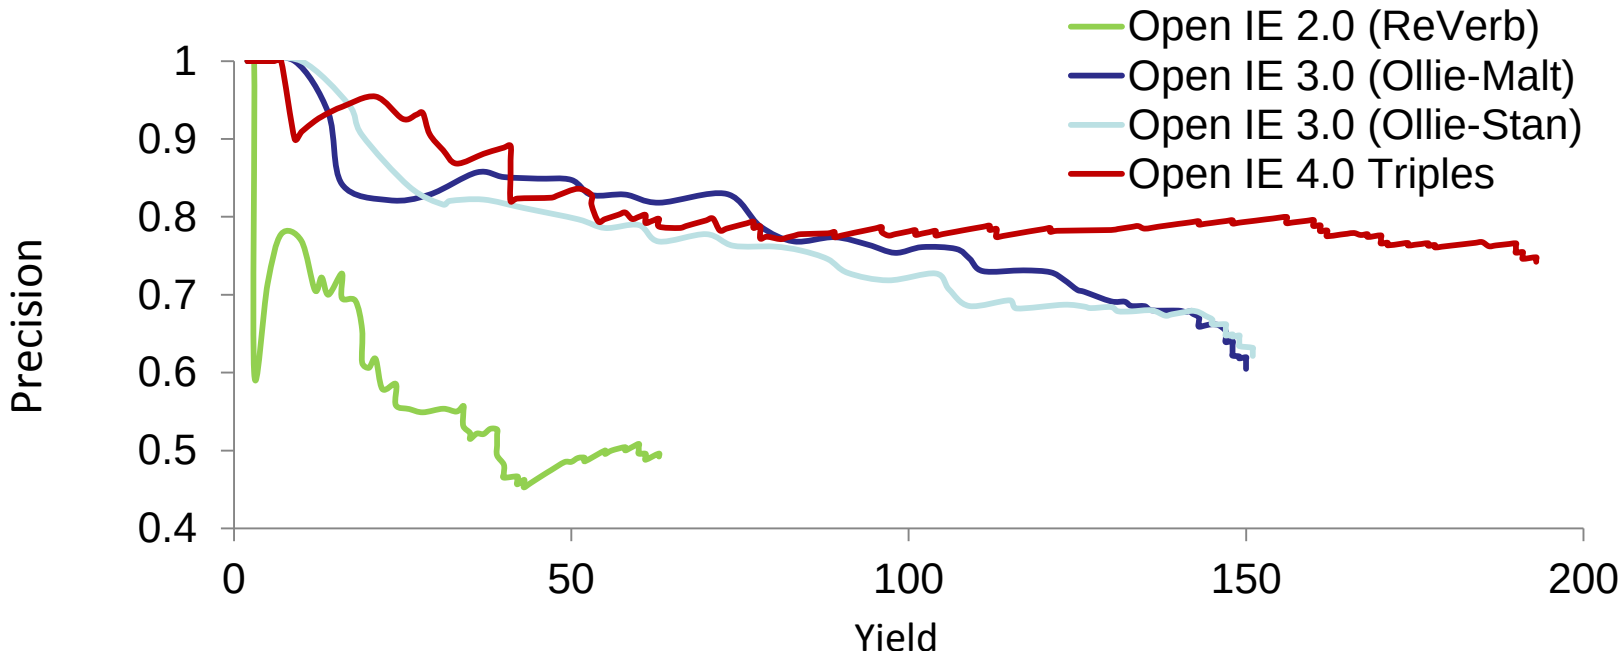

Open Information Extraction - Google Chrome

openie.cs.washington.edu/search?arg1=apple&rel=&arg2=&corpora=

Apps | O-1 Person of Extrao... | WeeklyPlans/Decem... | British Airways - Wh... | 74 | Vulcan - Project Hal... | Seattle Department ... | New Techniques fro... | Extraction Browser - ... | deep learning tutorial | Facebook launches ... | An AI that mimics o...

Open Information Extraction

Turing Center, University of Washington

Argument 1: apple Relation: Argument 2: All Search

408 answers from **11543 sentences** (cached)

Do you mean...

Apple Inc. (345)

Apple Inc. (NASDAQ: AAPL; formerly Apple Computer, Inc.) is an American multinational corporation that designs and sells consumer electronics, computer software, and personal computers. The company's best-known hardware products are the Macintosh line of computers, the iPod, the iPhone and the iPad. Its software includes the Mac OS X operating system; the iTunes media browser; the iLife suite of multimedia and creativity software; the iWork suite of productivity software; Aperture, a...read more

Apple (14)

The apple is the pomaceous fruit of the apple tree, species Malus domestica in the rose family (Rosaceae). It is one of the most widely cultivated tree fruits, and the most widely known of the many members of genus Malus that are used by humans. Apples grow on small, deciduous trees. The tree originated in Western Asia, where its wild ancestor, Malus sieversii, is still found today. Apples have been grown for thousands of years in Asia and Europe, and were brought to North America by...read more

IPad (4)

The iPad (/ˈaɪpæd/ EYE-pad) is a line of tablet computers designed and marketed by Apple Inc., primarily as a platform for audio-visual media including books, periodicals, movies, music, games, apps and web content. Its size and weight fall between those of contemporary smartphones and laptop computers. The iPad runs on iOS, the same operating system used on Apple's iPod Touch and iPhone, and can run its own applications as well as iPhone applications. Without modification or a developer...read more

Apple Store (1)

The Apple Retail Store is a chain of retail stores owned and operated by Apple Inc., dealing in computers and consumer electronics. The stores sell Macintosh personal computers, software, iPods, iPads, iPhones, third-party accessories, and other consumer electronics such as Apple TV. Many stores feature a Theatre for presentations and workshops and a Studio for training with Apple products; all stores offer a Genius Bar for technical support and repairs, as well as free workshops available...read more

Open Information Extraction - Google Chrome

openie.cs.washington.edu/results/?arg1=entity%3AApple&justResults=false

Apps | O-1 Person of Extrao... | WeeklyPlans/Decem... | British Airways - Wh... | 74 | Vulcan - Project Hal... | Seattle Department ... | New Techniques fro... | Extraction Browser - ... | deep learning tutorial | Facebook launches ... | An AI that mimics o...

Open Information Extraction

Turing Center, University of Washington

Argument 1: entity:Apple Relation: Argument 2: All Search

289 answers from **2008 sentences** (results truncated)

Apple

The apple is the pomaceous fruit of the apple tree, species Malus domestica in the rose family (Rosaceae). It is one of the most widely cultivated tree fruits, and the most widely known of the many members of genus Malus that are used by humans. Apples grow on small, deciduous trees. [illegible] Malus sieversii, is still found [illegible]

Open IE is scalable but shallow!

all | ingredient (37) | food (33) | [illegible]

- are tender (263)
- fall to the ground (68)
- are in bloom (68)
- grow on **Tree** (58)
- contain **Pectin** (53)
- are a good source of **Dietary fiber** (40)
- produces **Apple** (34)
- are rich in **Pectin** (33)
- are in season (32)
- are high in **Dietary fiber** (30)
- fall from **Tree** (24)
- is cut in half (22)

II. The Allen Institute for AI (AI2)

"In order to be truly intelligent, computers must understand– that is probably the critical word. It is one thing to feed *The Tale of Two Cities* into a computer. It's another to have the computer understand what's being said."

Source: (**Paul Allen**, Microcomputer Interface, **1977)**

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Mission: *contribute to the world through high-impact AI research and engineering, with emphasis on reasoning, learning, and reading capabilities.*

Scientific Advisory Board (SAB)

Adam Cheyer

Co-founder and VP Engineering at Siri, Inc.

Dan Roth

Professor at University of Illinois Urbana-Champaign, fellow of ACM, AAAI, and ACL, Associate Editor in Chief of JAIR

Eric Horvitz

Director of Microsoft Research (Redmond), fellow of AAAI and AAAS, AAAI President (2007-09)

Dan Weld

Professor at University Washington, fellow of ACM and AAAI

Tom Mitchell

Chair of Machine Learning Department, Carnegie-Mellon, fellow of AAAI and AAAS, AAAI Distinguished Service Award

AI2 Overview (Nonprofit Research Institute)

- Experimental science projects

- Sample of Research collaborations
  - Chris Manning, Stanford
  - Luke Zettlemoyer, UW
  - Andrew McCallum, Umass
  - Dan Roth, UIUC

- Open source initiatives
  - Open AI Resources

Discover & Discuss Open Source AI

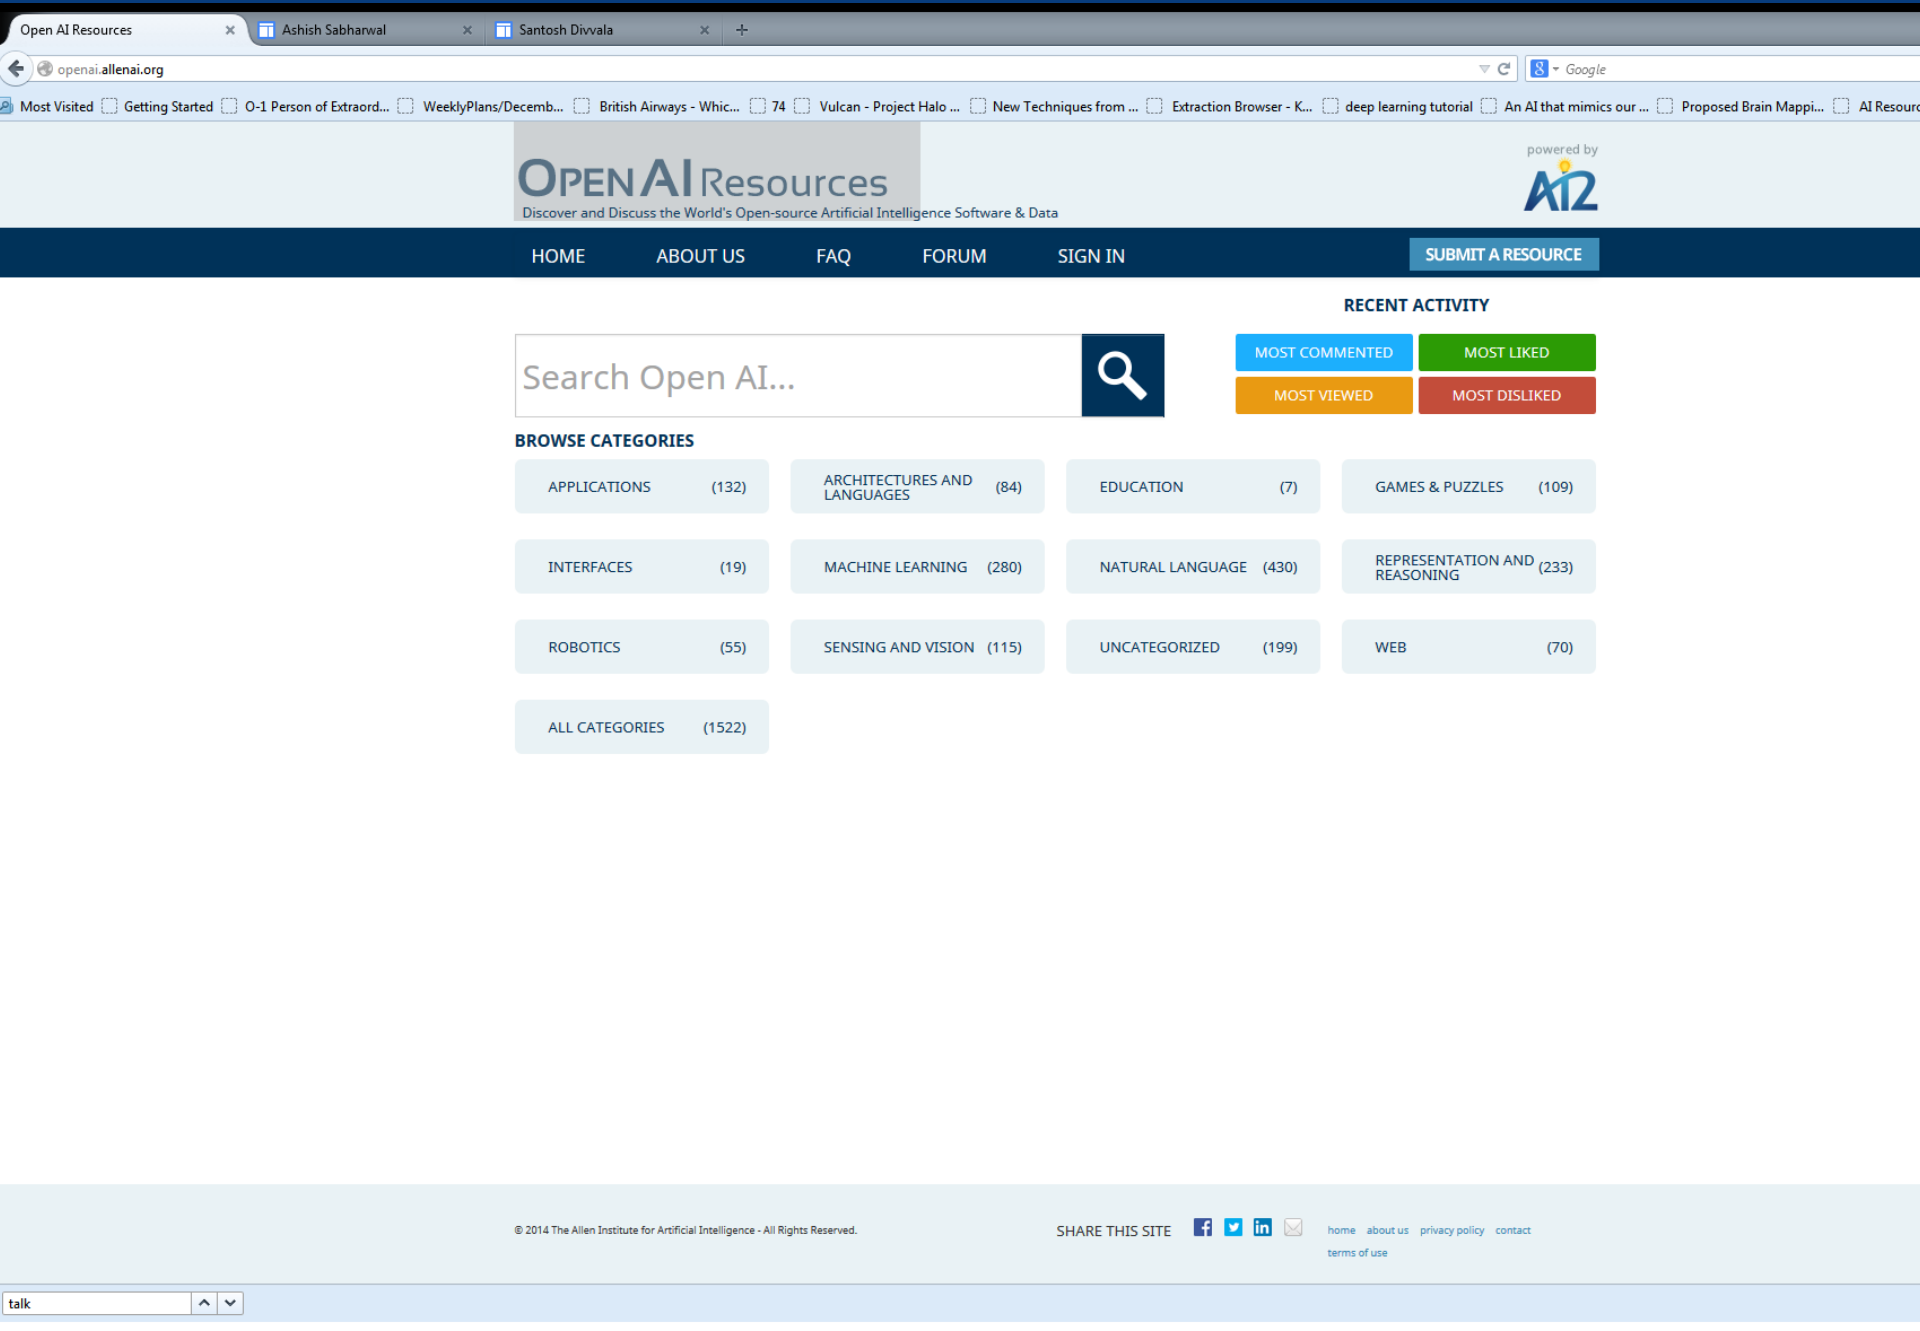

AI2 Chronology

Key Events	
AI2 launched	Jan. 2014
$8M AI ADI Program launched	March 2014
Team of 20	April 2014
Summer Intern Program	June 2014
Machine Vision Workshop	August 2014
Team of ~50	Dec. 2015

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Talent + Mission = Impact

ALLEN INSTITUTE
for ARTIFICIAL INTELLIGENCE

Flaghsip Project: Aristo

Project Methodology:

1. Externally-defined challenge tasks
2. Training data + **unseen** test data (“as is”)
3. Measurable progress, clear focus

Key differences with Watson:

1. Deeper semantics & inference
2. Open model: publish, collaborate, open source

Arithmetic word Problem

Liz had 9 black kittens. She gave some of her kittens to Joan. Joan now has 11 kittens. Liz has 5 kittens left and 3 have spots. How many kittens did Joan get?

State Transition

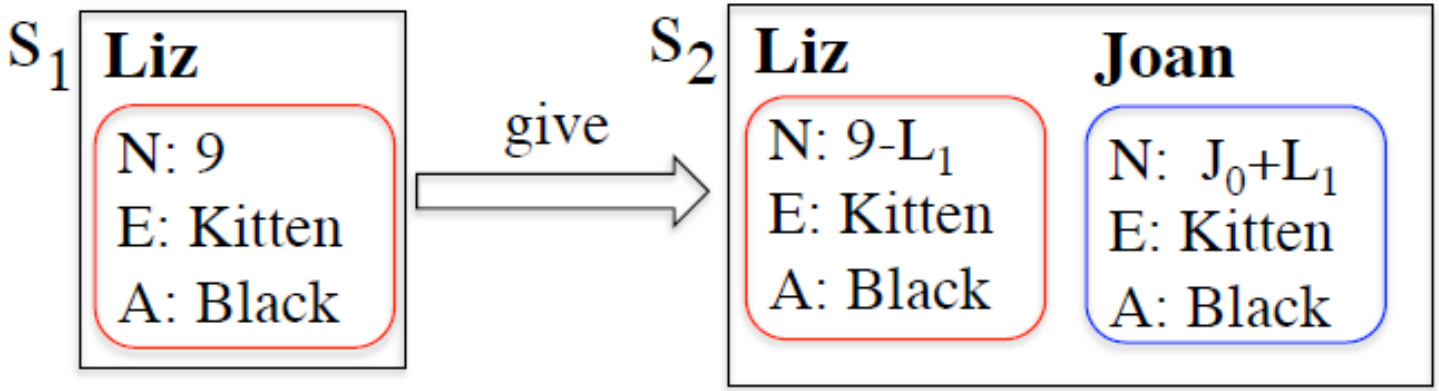

Liz gave some of her kittens to **Joan**.

Equation: $9 - x = 5$

Solution: $x = 4$ kittens

Verb Categorization Approach (submitted)

Category	Example
Observation	There **were** 28 bales of hay in the *barn*.
Positive	*Joan* **went** to 4 football games this year.
Negative	*John* **lost** 3 of the violet balloons.
Positive Transfer	*Mike's dad* **borrowed** 7 nickels from *Mike*.
Negative Transfer	*Jason* **placed** 131 erasers in the *drawer*.
Construct	*Karen* **added** 1/4 of a cup of walnuts to a *batch* of trail mix.
Destroy	The *rabbits* **ate** 4 of *Dan's* potatoes.

Table 1: Examples for different verb categories in sentences. Entities are underlined; *containers* are italic, and **verbs** are bolded.

81% categorization accuracy; 77% of test problems solved

IV. First Steps in the Geometry Domain

Combines vision, NLP, and simple semantics
Data set: 100 9th grade geometry problems
To Appear in AAAI '14 (*Seo, Hajishirzi, Farhadi*, Etzioni)

In the diagram, AB intersects circle O at D, AC intersects circle O at E, AE = 4, AC = 24, and AB = 16. Find AD.

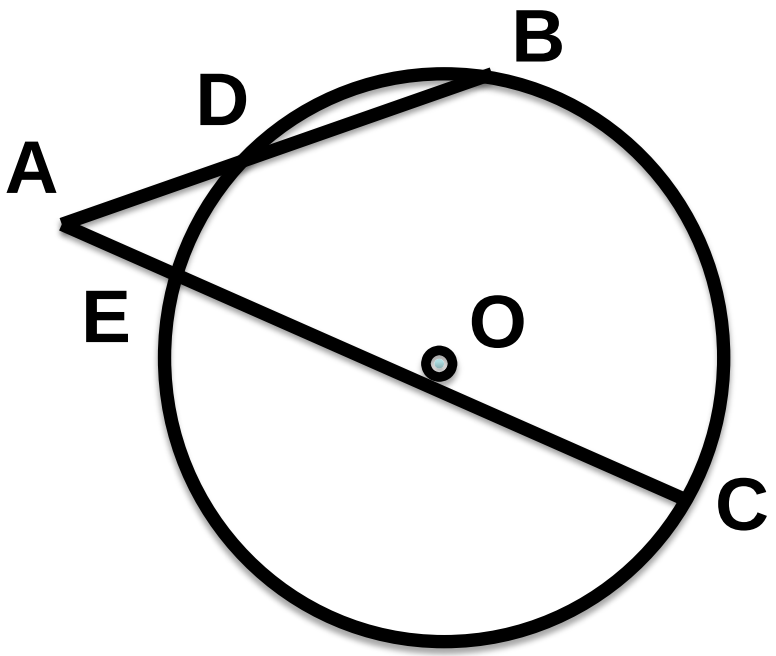

Diagram Understanding

- **Problem:** identify visual elements, their locations, their equations
- Couple with text

In the diagram, secant AB intersects circle O at D, secant AC intersects circle O at E, AE = 4, AC = 24, and AB = 16. Find AD.

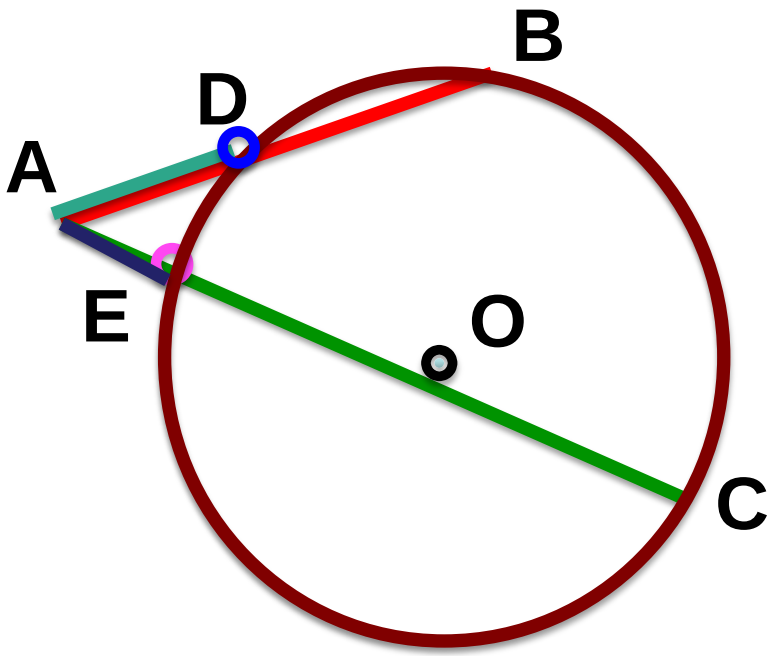

Diagram Understanding as Optimization

Objective function:

$$\hat{L} = \text{argmax}_L \, F(\hat{L}, D, T) =$$

$$P(D, \hat{L}) \quad \textbf{Coverage}$$

$$+ \, C(C, \hat{L}) \quad \textbf{Visual Coherence}$$

$$+ \, S(T, \hat{L}) \quad \textbf{Text Alignment}$$

- **Bad news: optimization requires 2^L operations**
- **Good news: function *F* is "submodular"**

We have an algorithm that returns a (1-1/e) approximation of the optimal solution!

Ablation Study

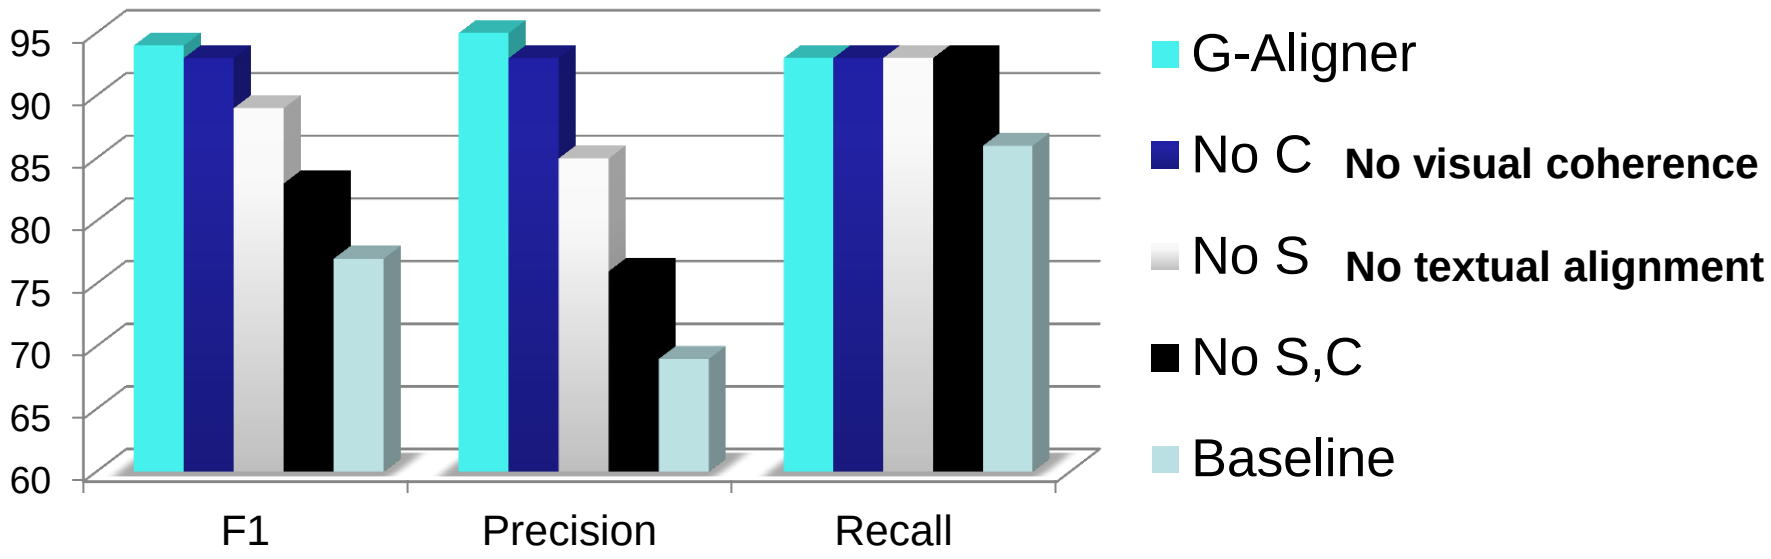

Diagram Demo (9th Grade geometry) G-aligner

G-Aligner Demonstration - Google Chrome

trusty.cs.washington.edu:9239/vpsel/#

Apps | O-1 Person of Extrao... | WeeklyPlans/Decem... | British Airways - Wh... | 74 | Vulcan - Project Hal... | New Techniques fro... | Extraction Browser - ... | deep learning tutorial | An AI that mimics o... | Proposed Brain Map... | AI Resources page - ...

Problem Definition

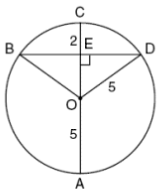

In the diagram at the left, circle O has a radius 5, and CE = 2. Diameter AC is perpendicular to chord BD at E. What is the length of BD?

Baseline Solution (Hough Transform)

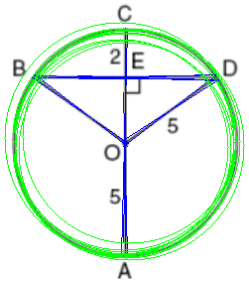

G-Aligner Solution

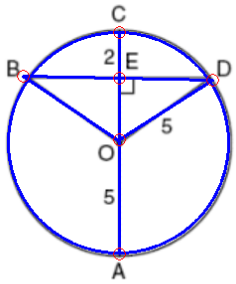

In the diagram at the left, circle O has a radius 5, and CE = 2. Diameter AC is perpendicular to chord BD at E. What is the length of BD?

ALLEN INSTITUTE for ARTIFICIAL INTELLIGENCE

V. 4th Grade "common-sense science"

THE UNIVERSITY OF THE STATE OF NEW YORK

GRADE 4

ELEMENTARY-LEVEL

SCIENCE TEST

WRITTEN TEST

JUNE 3, 2013

12 Which force causes a bicycle to slow down when the brakes are used?

A friction
B electricity
C gravity
D magnetism

More Example Questions

1 Which example describes an organism taking in nutrients?

A a dog burying a bone
B a girl eating an apple
C an insect crawling on a leaf
D a boy planting tomatoes in a garden

28 When a baby shakes a rattle, it makes a noise. Which form of energy was changed to sound energy?

A electrical
B light
C mechanical
D heat

Retrieval-based Approach (Baseline)

1 Which example describes an organism taking in nutrients?

A a dog burying a bone
B a girl eating an apple
C an insect crawling on a leaf
D a boy planting tomatoes in a garden

H_D A boy planting tomatoes in a garden is an example of an organism taking in nutrients?

Retrieval-based Approach Fails

1 Which example describes an organism taking in nutrients?

A a dog burying a bone
B a girl eating an apple
C an insect crawling on a leaf
D a boy planting tomatoes in a garden

H_D A boy **planting** tomatoes in a garden is an example **of** an organism **taking in nutrients**?

Best matching sentence

The roots **of** the **plant** **take in** water and **nutrients.**

Answer is D

Another Example

27 Fourth graders are planning a roller-skate race. Which surface would the best for this race?

A gravel
B sand
C blacktop
D grass

World knowledge necessary for NLP!

I. Winograd Schemas and Knowledge (1972)

The city councilmen refused the demonstrators a permit because **they** [feared/advocated] violence.

They = ?

Winograd Schema Challenge (Levesque, 2011)

The large ball crashed right through the table because **it** was made of ***styrofoam.***

It = table

The large ball crashed right through the table because **it** was made of ***steel.***

It = ball

Factual Knowledge for 4th Grade Science

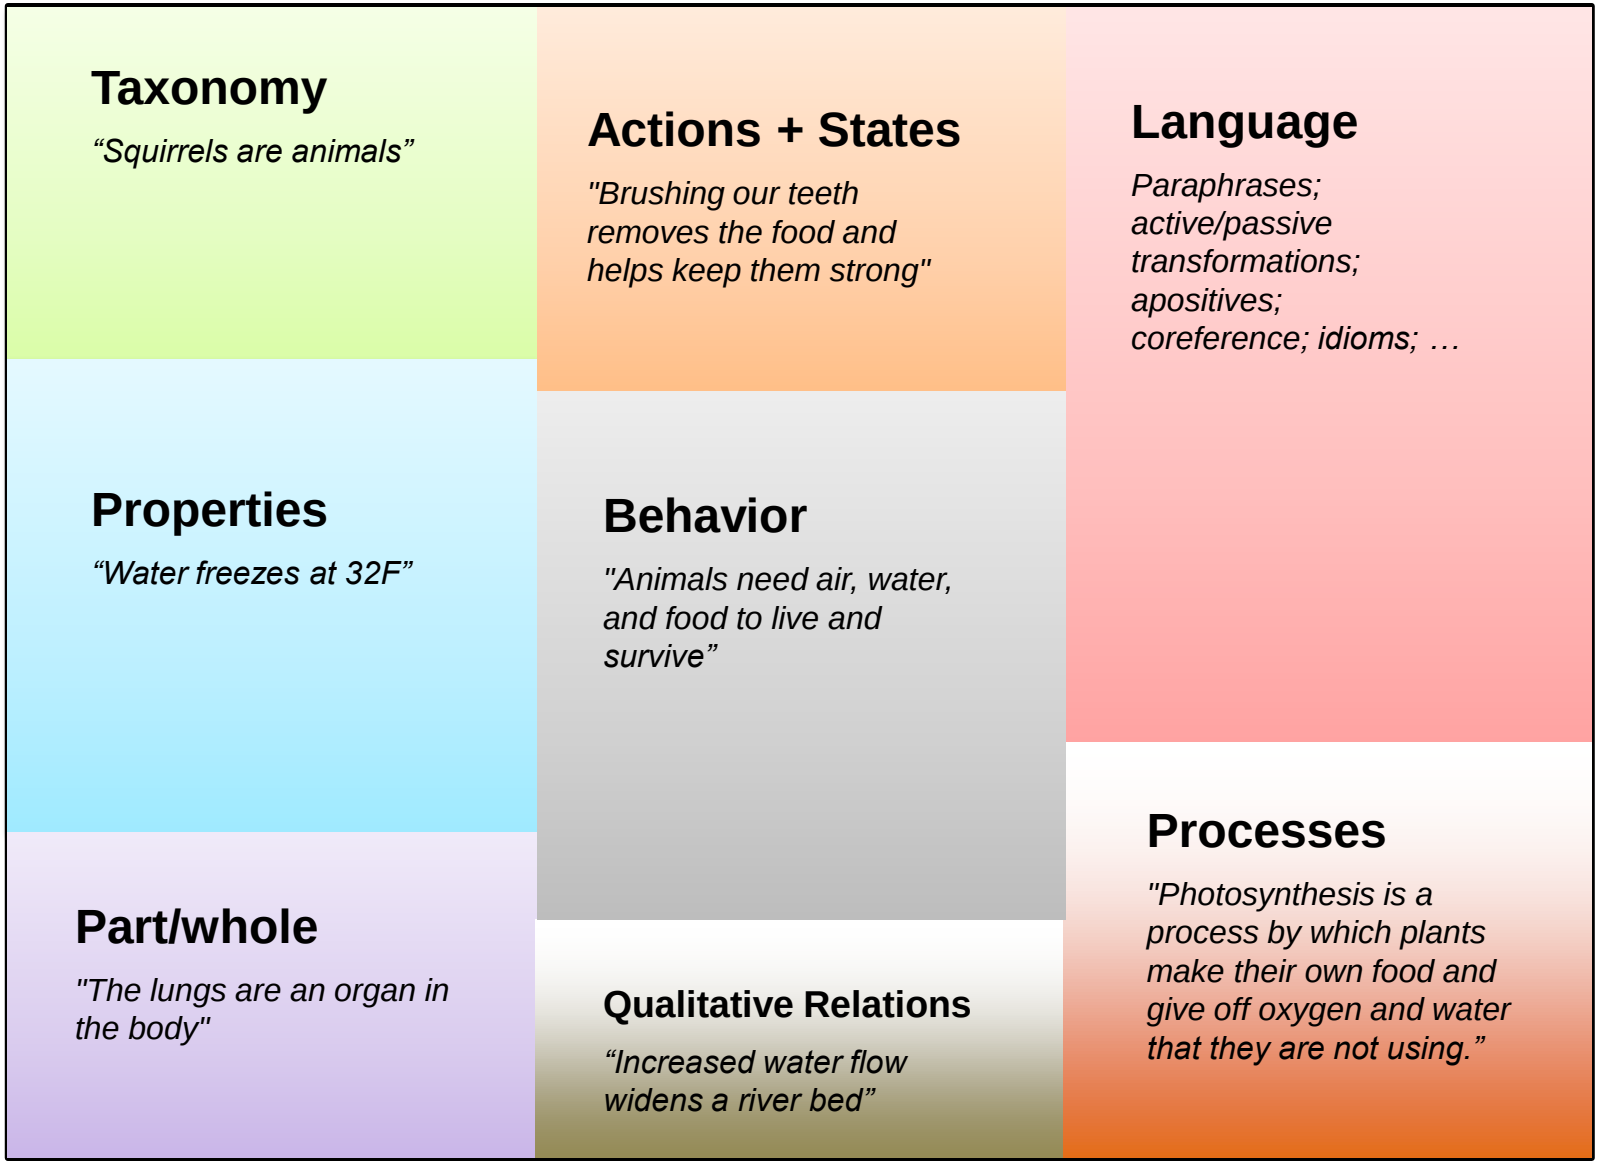

Sample Co-reference Challenges

Some birds fly south before winter each year. **This** is an example of (A) migration (B) hibernation (C) germination (D) evaporation

A student has a ball of clay that sinks when placed in a pan of water. **Which property** should he change to make the clay float? (A) color (B) texture (C) mass (D) shape

The diagram below shows an electrical circuit. The purpose of **the copper wire** is to (A) conduct electricity (B) produce electricity (C) store electricity (D) stop the flow of electricity

Entailment Challenges

- Absorb → take in
- 24 hours → day
- Lunch → Dinner
- Blue eyes → color of your eyes
- Disappear → evaporate
- Cold → decrease in temperature
- Gas → gaseous state
- Good health habit → be healthy

- **Verb synonym**
- **Phrase synonym**
- **“Siblings”**

Current Methodology

Give Aristo an exam

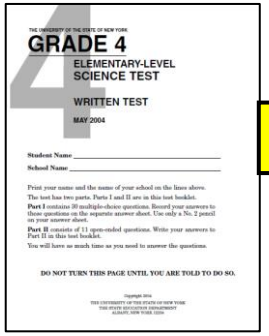

Mixture of:

Lexical Inference

{move,object,Earth}

{fall,ball,ground}

Logical Inference

move(X,to(Earth))

fall(X,to(ground))

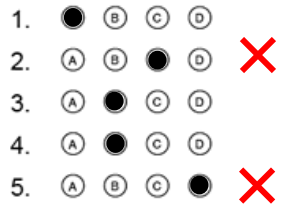

Identify where correct answer was not selected (inference path too weak)

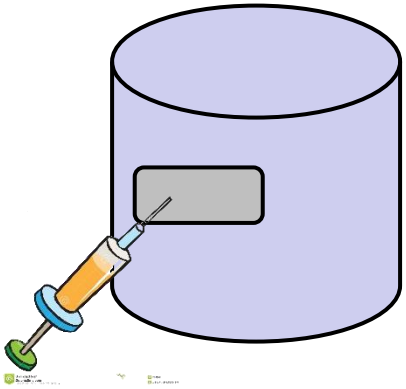

Inject new knowledge into the KB

Identify missing knowledge

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Is it all “about” Knowledge?

What about semantics?

Semantics: Treatment of Generics:

Fruits contain seeds **= ?**

- Two big issues: quantifier scope and strength
  - All (most? many?) fruits contain a seed?
  - All (most? many?) seeds are contained in a fruit?
  - All fruits contain all? some? seeds?
  - …
- Pragmatic approximation:
  - Take the most common (forall-exists) interpretation, and the reverse
    a. If it's a *fruit*, it contains a *seed* (with some confidence)
    b. If it's a *seed*, it's contained in a *fruit* (with some confidence)

Formal representation:

∀ Fruit isa(Fruit,"fruit") → ∃ Seed isa(Seed,"seed"), contain(Fruit,Seed) @ weight 1.0
∀ Seed isa(Seed,"seed") → ∃ Fruit isa(Fruit,"fruit"), contain(Fruit,Seed) @ weight 1.0

Semantics: (Word Senses)

What do words mean? “animal”, “cat”, “bank”, …

- **Traditional approach:**
  - map words/phrases to an ontology at reading time:
    - “animal” → **`Animal`**
    - “cat”, “kitty” → **`Cat`**
    - …
  - **but:** requires *anticipating all the needed distinctions* up front, before a particular task is encountered
- **Deferred decisions** (a lesson from Watson):
  - *preserve words/phrases* in the representation itself

∀ Fruit isa(Fruit,"fruit") → ∃ Seed isa(Seed,"seed"), contain(Fruit,Seed)

  - make ontology decisions as needed at runtime *in context*
    - “person” isa “animal”?
    - “water” isa “basic need”?
    - “shiny crimson” isa “bright color”?

Summary of our Approach

1. Knowledge + tractable reasoning is necessary for AI

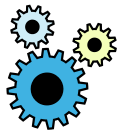

2. Knowledge acquisition has to be highly automated

3. Large bodies of high-quality knowledge can be acquired from text ("machine reading")

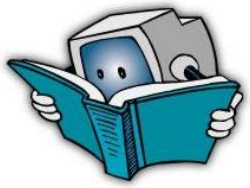

4. Aristo is tested on unseen, standardized, multiple choice tests in science, arithmetic, and geometry

Semantic challenges everywhere…

VI. Semantic Scholar Project

"Show all papers where x is implicated in y where the levels of z are elevated"…

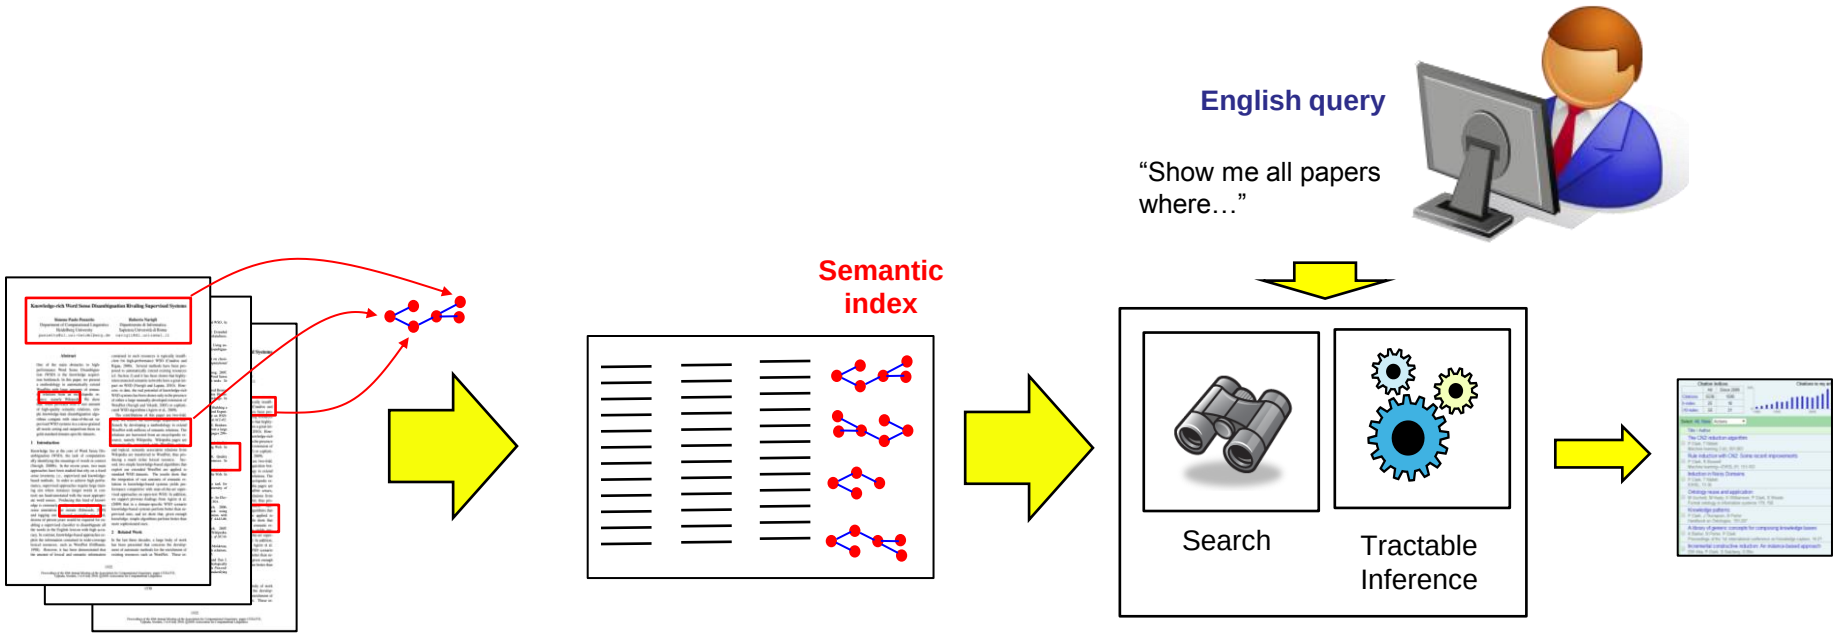

"lift all scientific boats"

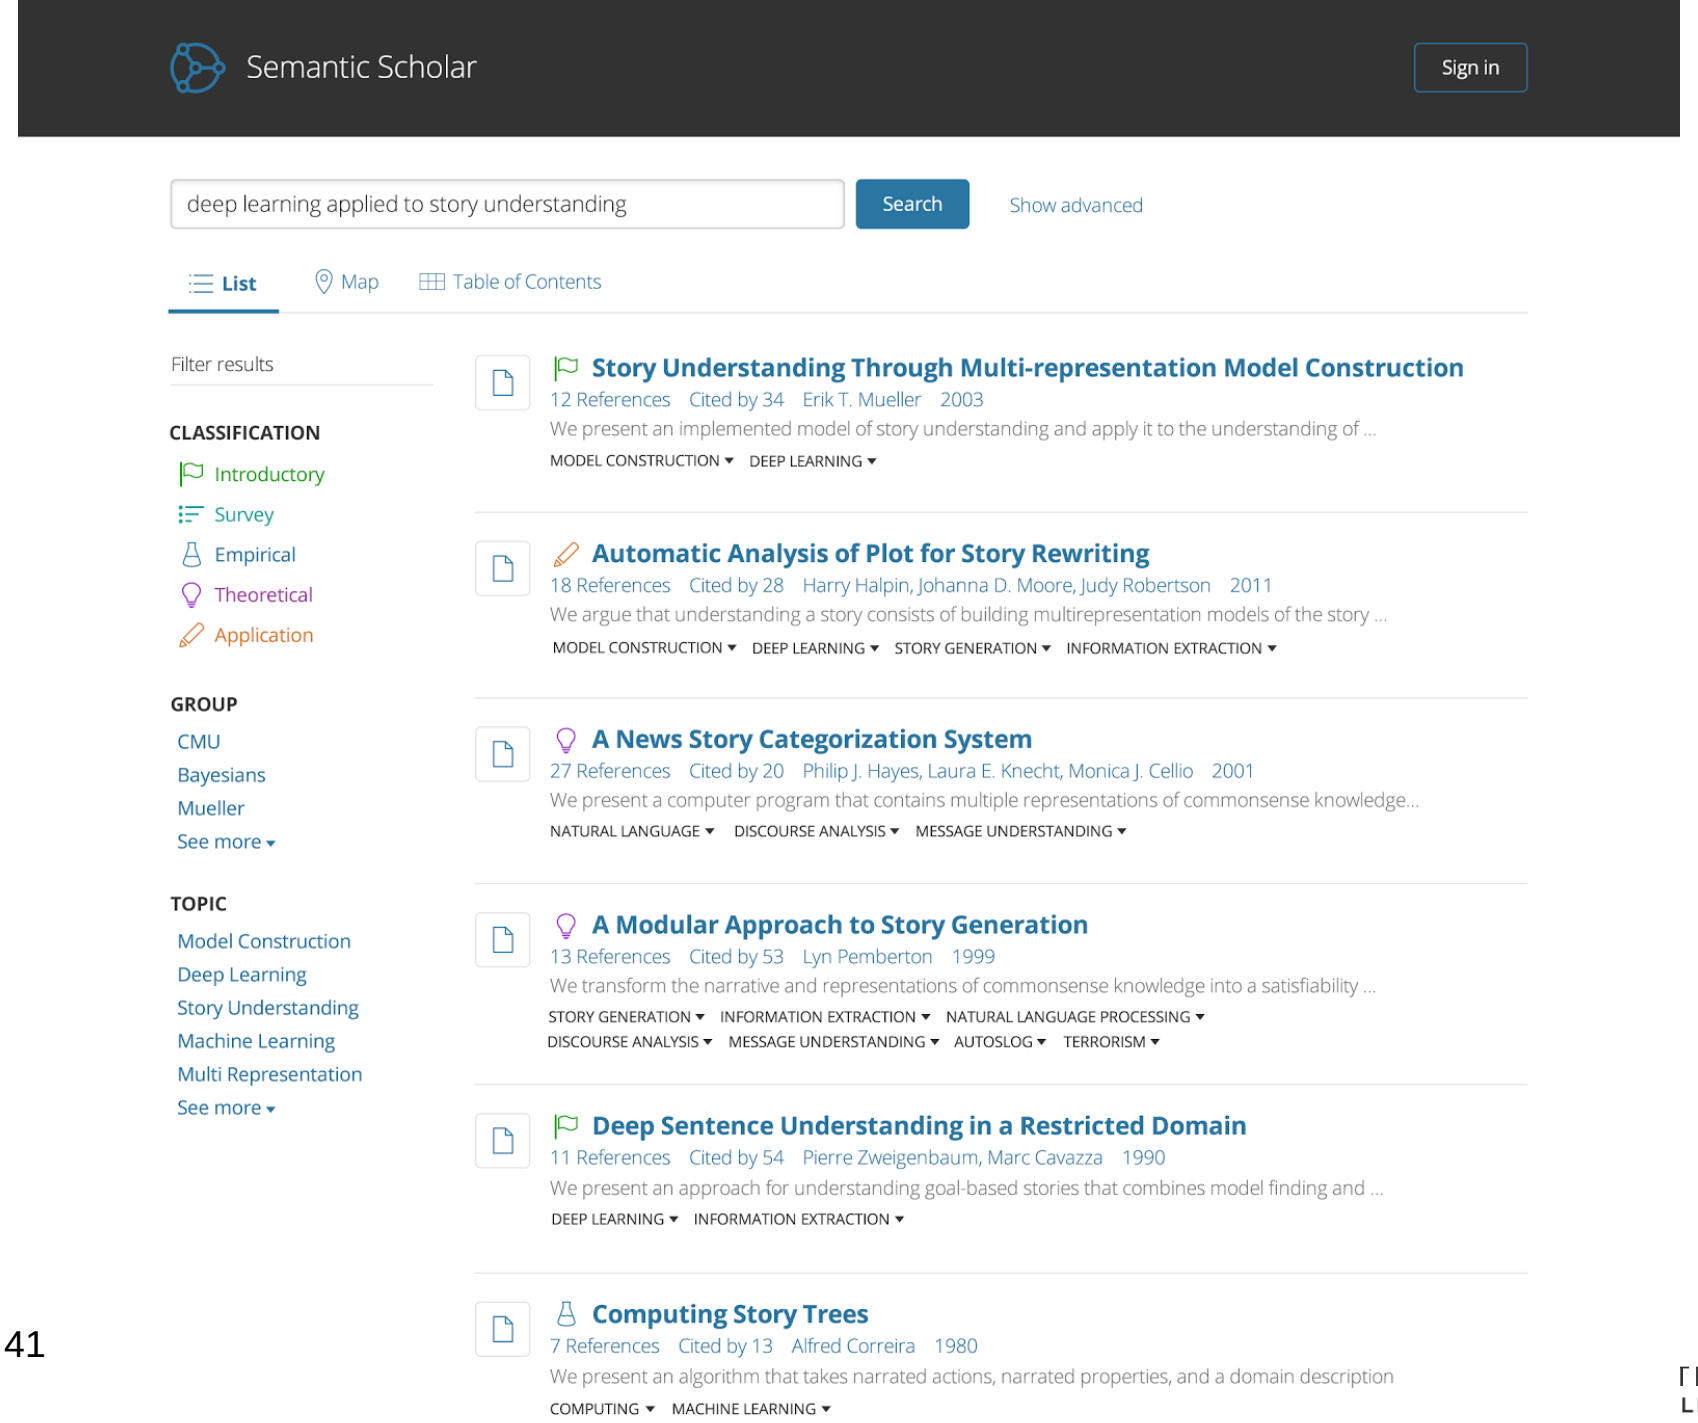
Semantic Scholar
Sign in
deep learning applied to story understanding
Search
Show advanced
List
Map
Table of Contents
Filter results
CLASSIFICATION
Introductory
Survey
Empirical
Theoretical
Application
GROUP
CMU
Bayesians
Mueller
See more
TOPIC
Model Construction
Deep Learning
Story Understanding
Machine Learning
Multi Representation
See more
Story Understanding Through Multi-representation Model Construction
12 References Cited by 34 Erik T. Mueller 2003
We present an implemented model of story understanding and apply it to the understanding of ...
MODEL CONSTRUCTION DEEP LEARNING
Automatic Analysis of Plot for Story Rewriting
18 References Cited by 28 Harry Halpin, Johanna D. Moore, Judy Robertson 2011
We argue that understanding a story consists of building multirepresentation models of the story ...
MODEL CONSTRUCTION DEEP LEARNING STORY GENERATION INFORMATION EXTRACTION
A News Story Categorization System
27 References Cited by 20 Philip J. Hayes, Laura E. Knecht, Monica J. Cellio 2001
We present a computer program that contains multiple representations of commonsense knowledge...
NATURAL LANGUAGE DISCOURSE ANALYSIS MESSAGE UNDERSTANDING
A Modular Approach to Story Generation
13 References Cited by 53 Lyn Pemberton 1999
We transform the narrative and representations of commonsense knowledge into a satisfiability ...
STORY GENERATION INFORMATION EXTRACTION NATURAL LANGUAGE PROCESSING
DISCOURSE ANALYSIS MESSAGE UNDERSTANDING AUTOSLOG TERRORISM
Deep Sentence Understanding in a Restricted Domain
11 References Cited by 54 Pierre Zweigenbaum, Marc Cavazza 1990
We present an approach for understanding goal-based stories that combines model finding and ...
DEEP LEARNING INFORMATION EXTRACTION
Computing Story Trees
7 References Cited by 13 Alfred Correira 1980
We present an algorithm that takes narrated actions, narrated properties, and a domain description
COMPUTING MACHINE LEARNING

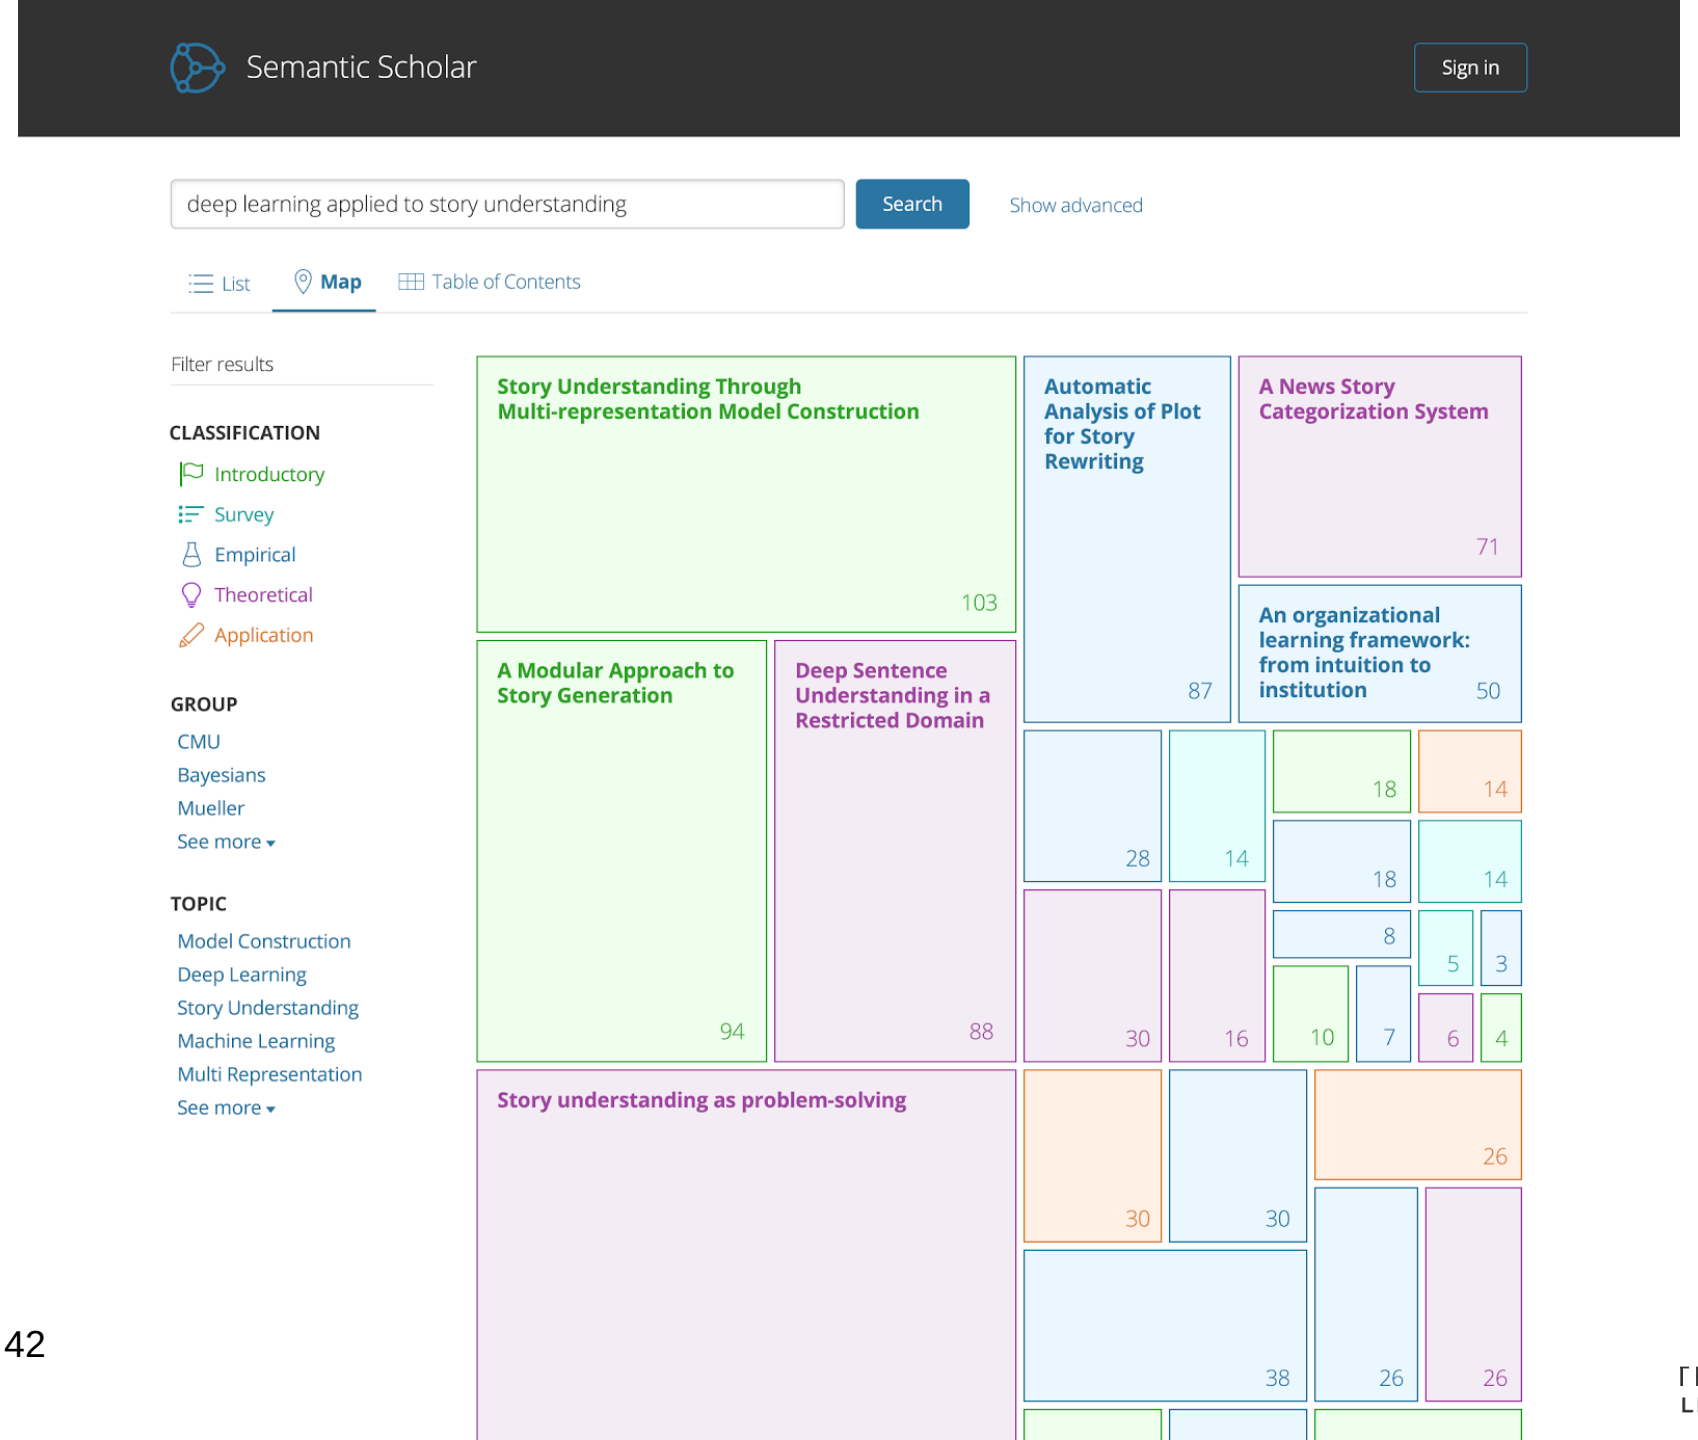
Semantic Scholar
Sign in
deep learning applied to story understanding
Search
Show advanced
List
Map
Table of Contents
Filter results
CLASSIFICATION
Introductory
Survey
Empirical
Theoretical
Application
GROUP
CMU
Bayesians
Mueller
See more
TOPIC
Model Construction
Deep Learning
Story Understanding
Machine Learning
Multi Representation
See more
Story Understanding Through Multi-representation Model Construction
103
Automatic Analysis of Plot for Story Rewriting
87
A News Story Categorization System
71
An organizational learning framework: from intuition to institution
50
A Modular Approach to Story Generation
94
Deep Sentence Understanding in a Restricted Domain
88
28
14
18
14
18
14
8
5
3
30
16
10
7
6
4
Story understanding as problem-solving
30
30
26
38
26
26

Technologies on our Roadmap

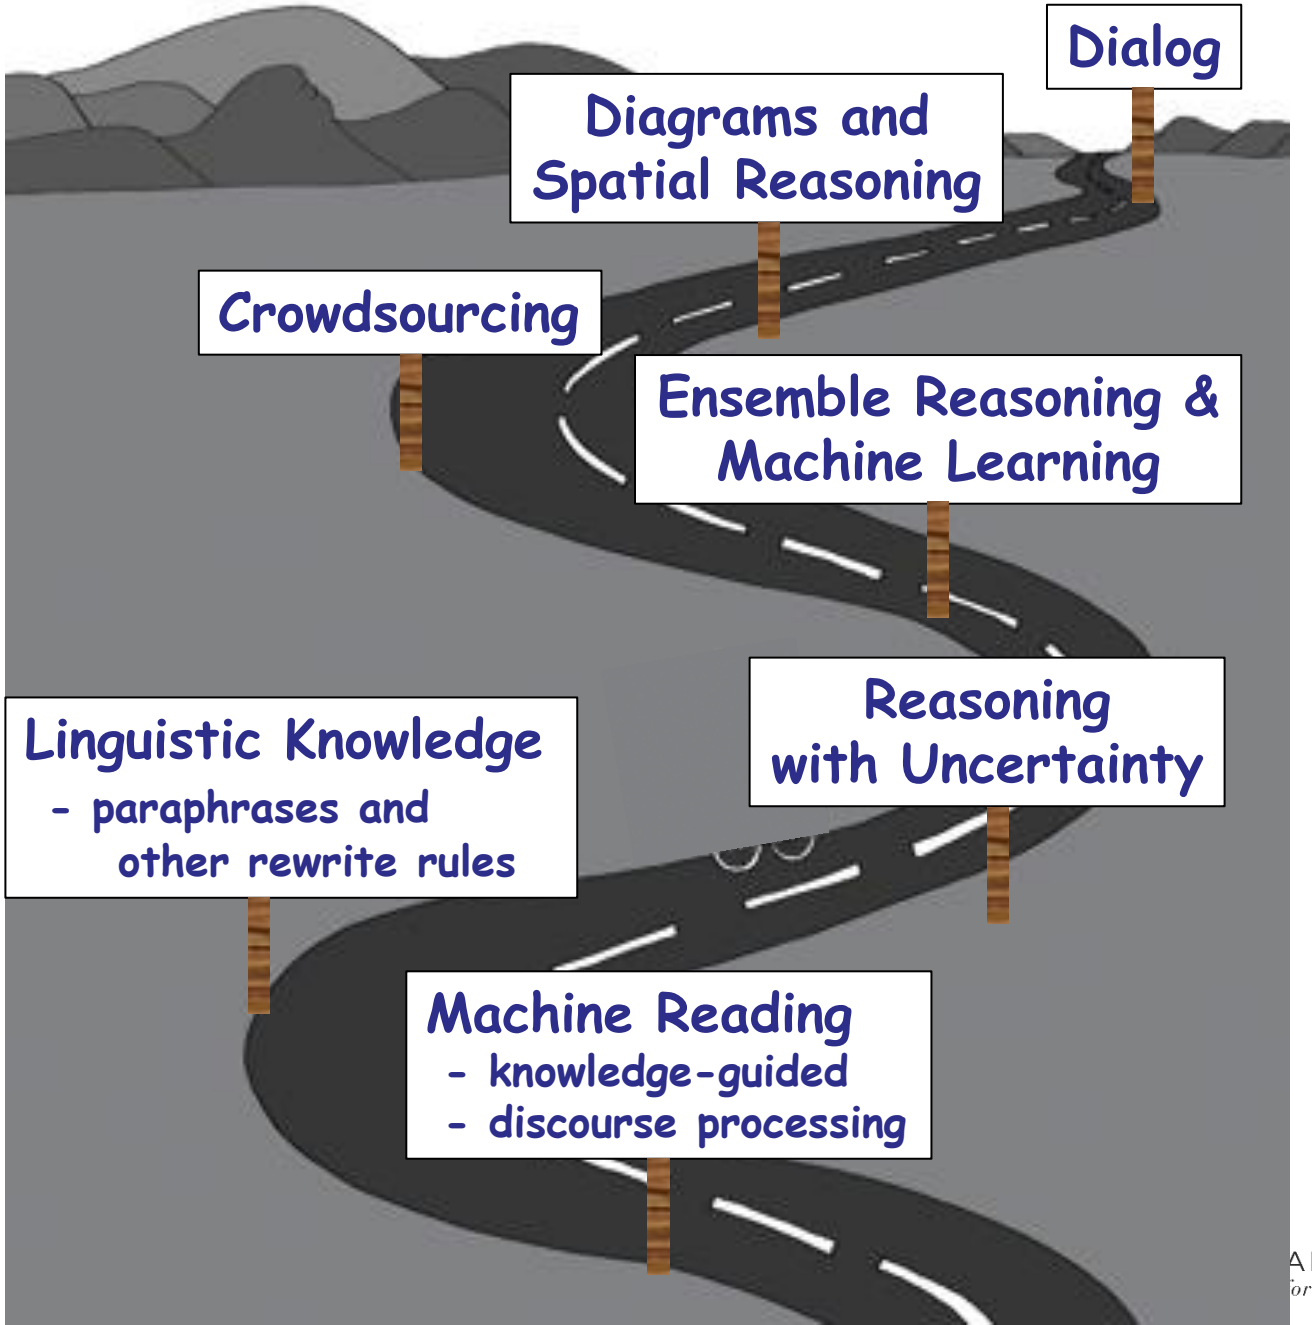

Conclusion: Join us!