

The Battle for the Future of Data Mining

Oren Etzioni, CEO
Allen Institute for AI (AI2)



Big Data Tidal Wave

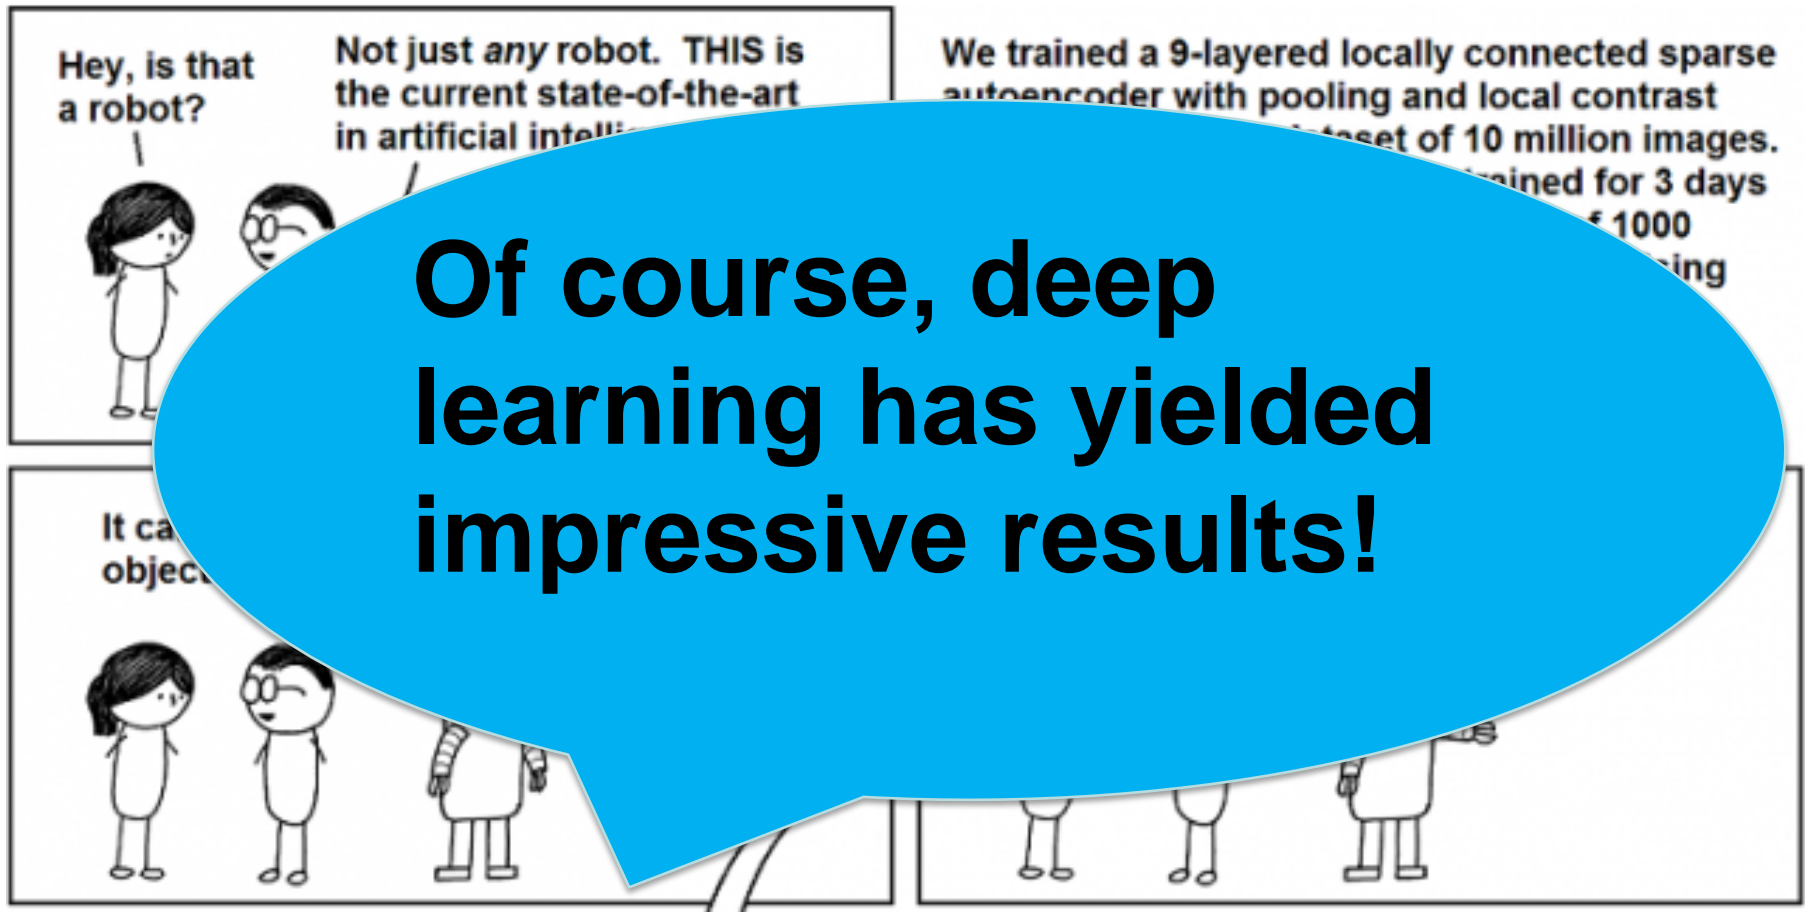




Deep Learning

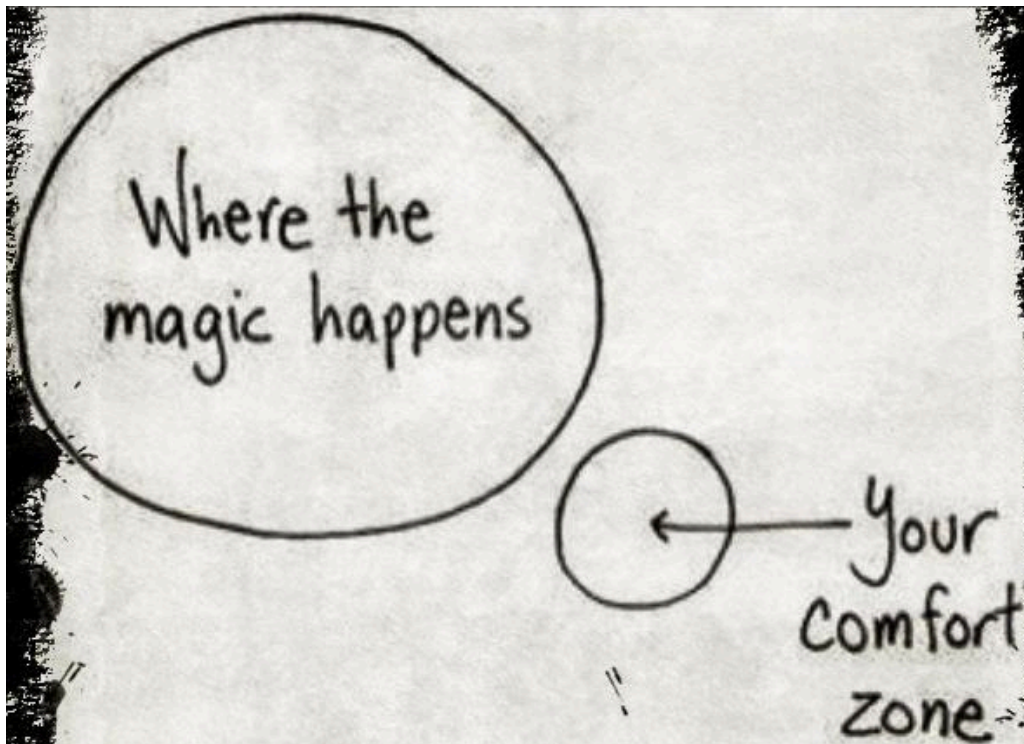


Some Skepticism about Deep Learning



Outline

- I. Limits of current paradigm
- II. Surprising poverty of Today's Mined KBs
- III. Machine Reading
- IV. Allen Institute for AI (AI2)



What are the limits of Classification and prediction?

*Note: not arguing for "generative models"
over discriminative methods*

Example of State of the Art

Retweets - NYTimes.com - Google Chrome

mes.com/2014/07/03/upshot/a-25-question-twitter-quiz-to-predict-retweets.html

ites

HOME SEARCH

The New York Times



Sendhil
Mullainathan

@m_sendhil

Consider these two tweets by Al Gore, both promoting the same article:



Al Gore
@algore

Follow

40% of smartphone users connect to Internet
immediately upon awakening, before leaving bed.

#TheFuture bit.ly/WYRz39 @TheAtlantic

8:45 AM - 26 Jan 2013



Based on Tan, Lee, Pang
(ACL '14)

What about creating
Tweets in the first
place?

I can't be reduced
to a vector of
surface variables
and correlations!



But is Netflix lying to us?!





Prototype in the cloud, then easily
move back on-premises

with IBM PureApplication® Solutions

IBM

MUST READ: *Microsoft's Windows 9: Much ado about little, given cloud shift*

Topic: *Big Data*

Follow via:  

After a year of medical school, IBM's Watson passes first milestone

Summary: *IBM's year-long residency at Memorial Sloan-Kettering Cancer Center and Wellpoint is finally producing cognitive computing breakthroughs (and two new products).*



By Andrew Nusca for [Between the Lines](#) | February 8, 2013 -- 16:45 GMT (08:45 PST)

 Follow @editorialiste

[Get the ZDNet Big Data newsletter now](#)

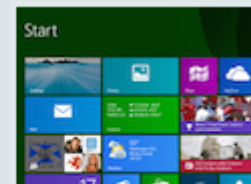
IBM's Watson has made its first commercially-developed cognitive computing breakthroughs, the company announced this morning. They arrive after more than a year of training with healthcare partners WellPoint and Memorial Sloan-Kettering Cancer Center.



cirba

Driving automation
and incredible cost
efficiency

[Learn how »](#)



*Six Windows 8.1
tips for
maximum
performance,
reliability, and*

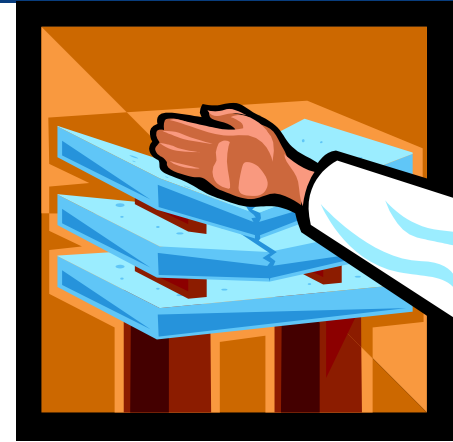
security

More Challenges for Classification

- **Learning Arithmetics** = set of crisp rules & procedures
Distribution insensitive; 100% accuracy
- **Playing Chess:** move choice requires search
Deep Blue; human grand masters
- **Language Understanding**
Complex grammar
Background knowledge

Fundamental Claim

Classification, prediction limited
("structured prediction"; "multi-task learning"
are steps forward)



Need a process that's
constructive,
multi-layered,
Knowledge-intensive

Data → knowledge → theories
Theories → reasoning



**“You can’t play 20 questions
with nature and Win.”**

Allen Newell, 1973



**Key AI tasks cannot be reduced to
classification problems!
(using “reasonable” features)**

The large ball is on the table
because

**World knowledge is
necessary for elementary
understanding, but is our
technology ready?**

Just the Facts...

How much do *you* know about apples?

- Edible fruit
- Has peel, stem, flesh
- Consists of mostly water
- Ripens on a tree

at least 20 non-trivial facts + inheritance...

Mined KBs are knowledge poor!

2014 Knowledge Tour

[+Oren](#)

[Web](#) [News](#) [Images](#) [Shopping](#) [Maps](#) [More](#) [Search tools](#)

About 597,000,000 results (0.37 seconds)

Apple

www.apple.com/ [Apple Inc.](#)

Apple designs and creates iPod and iTunes, Mac laptop and desktop computers, the OS X operating system, and the revolutionary iPhone and iPad.

3.8 ★★★★★ 46 Google reviews · [Write a review](#)

2656 NE University Village St, Seattle, WA 98105
(206) 892-0433

Apple Store

Shop iPad - Shop iPhone -
Shop Mac - Refurbished &
Clearance

iPhone

Discover everything
iPhone, including the most
advanced ...

Support

Apple Support is your
starting point for help with
Apple products ...

[More results from apple.com »](#)

iPad

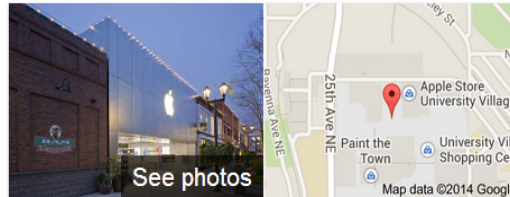
Introducing iPad Air and the
iPad mini with Retina
display. Watch ...

iTunes

iTunes is the world's best
way to play — and add to
— your ...

University Village

The Apple Store is located
in University Village, at
2656 NE ...



Apple Store University Village

[Directions](#)[Write a review](#)

Address: 2656 NE University Village St, Seattle, WA 98105

Phone: (206) 892-0433

Prices: \$\$\$\$

Hours: Open today · 11:00 am – 6:00 pm

Reviews

3.8 ★★★★★ 46 Google reviews

More reviews: [judysbook.com](#), [yahoo.com](#), [insiderpages.com](#),
[servicearea.com](#)

People also search for



University
Village
Shopping
Center



The Mac
Store



Apple Store
Alderwood
Mall



Banana
Republic



Best Buy

Are you the business owner?

[Feedback](#)

See results about

[Apple Inc.](#)

Consumer electronics company

Apple Inc. is an American multinational corporation
headquartered in Cupertino, California, that ...



News for apple



Another Apple-Samsung Skirmish Heads to Court

ABC News · 4 minutes ago

The fiercest rivalry in the world of smartphones is
heading back to court this week in the heart of the
Silicon Valley, with **Apple** and Samsung ...

[Apple Loop: iPhone 6 Screens, Microsoft Office For iPad, iBeacons For MLB, Sece...](#)
[Forbes](#) · by Ewan Spence · 7 hours ago

[Apple plans to create 'transparent' iPhones with rear camera technology](#)
[The Independent](#) · 4 hours ago

[More news for apple](#)

Web Images Shopping Videos News More Search tools

About 365,000 results (0.27 seconds)

Apple - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/Apple Wikipedia
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 ...
Apple Inc. - Disambiguation - Malus sieversii - List of apple cultivars

Apple fruit nutrition facts and health benefits - Nutrition and ...
www.nutrition-and-you.com/apple-fruit.html
by Umesh Rudrappa - in 70 Google+ circles
Delicious and crunchy **apple fruit** is one of the popular fruits containing an impressive list of antioxidants and essential nutrients required for good health...

apple (fruit and tree) -- Encyclopedia Britannica
www.britannica.com/EBchecked/topic/.../apple Encyclopedia Britannica
Jul 19, 2013 - *Malus domestica* fruit of the domesticated tree *Malus domestica* (family Rosaceae), one of the most widely cultivated tree fruits. The apple is a ...

Cashew-Apple Fruit Growing in the Florida Home Landscape...
edis.ifas.ufl.edu/hs377 Institute of Food and Agricultural Sciences
Cashew-Apple Fruit Growing in the Florida Home Landscape1.
John McLaughlin, Carlos Balerdi and Jonathan Crane2.
Scientific Name: *Anacardium* ...

The Apple - "King" of the Fruit - Nature's Wonderland - Jewish...
www.chabad.org/kids/article_cdo/aid/114763/.../Apples.htm Chabad
History of the Apple. In honor of Tu-B'Shevat (15th day of Sh'vat), we will now talk about a fruit - a very special fruit - the Apple, "King" of all fruits (in Hebrew it is ...

Sneaky Green Apple Fruit Leather Recipe - Oh My Veggies
ohmyveggies.com/recipe-sneaky-green-apple-fruit-leather/ 6 hrs 10 mins
Aug 15, 2013 - This sugar-free **apple fruit leather** recipe is hiding a secret--it's made with spinach! But it tastes so good, you'll never know it's in there.

Truffle - Edible Arrangements



Apple

Fruit
The apple is the pomaceous fruit of the apple tree, species *Malus domestica* in the rose family. 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. Wikipedia

Nutrition Facts

Apple ▾

Amount Per 1 medium (3" dia) (182 g) ▾

Calories 95

	% Daily Value*
Total Fat 0.3 g	0%
Saturated fat 0.1 g	0%
Polyunsaturated fat 0.1 g	
Monounsaturated fat 0 g	
Cholesterol 0 mg	0%
Sodium 2 mg	0%
Potassium 195 mg	5%
Total Carbohydrate 25 g	8%
Dietary fiber 4.4 g	17%
Sugar 19 g	
Protein 0.5 g	1%
Vitamin A	1%
Calcium	1%
Vitamin D	0%
Vitamin B-12	0%
Vitamin C	14%
Iron	1%
Vitamin B-6	5%
Magnesium	2%



Making Wolfram|Alpha free to all wouldn't be possible without the support of our Pro subscribers.

Help us make even more of the world computable in 2014 by going Pro.

Go Pro

Share:



more



apple fruit



Examples Random

Assuming "apple" is a food | Use as [a general material](#) instead

Assuming apple | Use [prepared apples](#) or [more](#) instead

Assuming any type of apple | Use [apple, with skin](#) or [apple, without skin](#) instead

Assuming any type of fruit | Use [abiyuch, generic](#) or [more](#) instead

Input interpretation:

apple

amount

1 apple

fruit

amount

140 grams

Individual nutrition facts:

Total nutrition facts

Apple:

serving size 1 apple (182 g)

total calories 91

fat calories 2

% daily value*

total fat 273 mg

0%

saturated fat 45 mg

0%

trans fat

cholesterol 0 g

0%

sodium 910 µg

0%

New to
Wolfram|Alpha?



Take the Tour >>

Take Wolfram|Alpha anywhere...



OpenCyc (Current): [<http://sw.opencyc.org/concept/Mx8Ngh4rvVipdpwpEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA>]

OpenCyc (Versioned): [<http://sw.opencyc.org/2012/05/10/concept/Mx8Ngh4rvVipdpwpEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA>]

OpenCyc (Readable): [<http://sw.opencyc.org/2012/05/10/concept/en/Apple>]

Search



OpenCyc Collection: apple

Unique ID: [[Mx8Ngh4rvVipdpwpEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA](#)]

English ID: [(FruitFn AppleTree)]

English Aliases: ["apples", "fruit of the apple", "fruit of the apple tree", "fruit of the apple trees", "fruit of the apples", "fruit of the Malus pumila"]

The collection of individual apples.

A Type of: accessory fruit, edible fruit

Instance of: type of food, type of life stage, type of object whose instances do not physically overlap

Subtypes: Adams Pearmain apple, Ambrosia apple, Braeburn apple, Gala apple, Golden Delicious apple, Granny Smith apple, Jazz apple, McIntosh apple, red delicious apple

Instances:

Same as:

Copyright © 2001-2012 Cycorp, Inc.



[Get /c/en/apple in JSON format](#)

apple

apple — *RelatedTo* → fruit
apple is related to a fruit

apple — *RelatedTo* → red fruit
apple is related to red fruit

apple — *RelatedTo* → core
apple is related to core

apple — *Antonym* → orange
apple is the opposite of orange

apple — *RelatedTo* → computer
apple is related to a computer

apple — *RelatedTo* → green
apple is related to green

apple — *RelatedTo* → eve
apple is related to eve

apple — *RelatedTo* → mac
apple is related to mac

apple — *RelatedTo* → pie
apple is related to pie

apple — *HasProperty* → green
apple can be green

apple — *RelatedTo* → red
apple is related to red

apple — *RelatedTo* → a
apple is related to a

apple — *RelatedTo* → tree
apple is related to tree

apple — *HasProperty* → red
apple can be red

apple — *RelatedTo* → seed
apple is related to seeds

apple — *RelatedTo* → round
apple is related to round

apple — *RelatedTo* → macintosh
apple is related to macintosh

apple — *RelatedTo* → pear
apple is related to a pear

apple — *RelatedTo* → orange
apple is related to oranges

apple — *ReceivesAction* → eat
an apple can be eaten

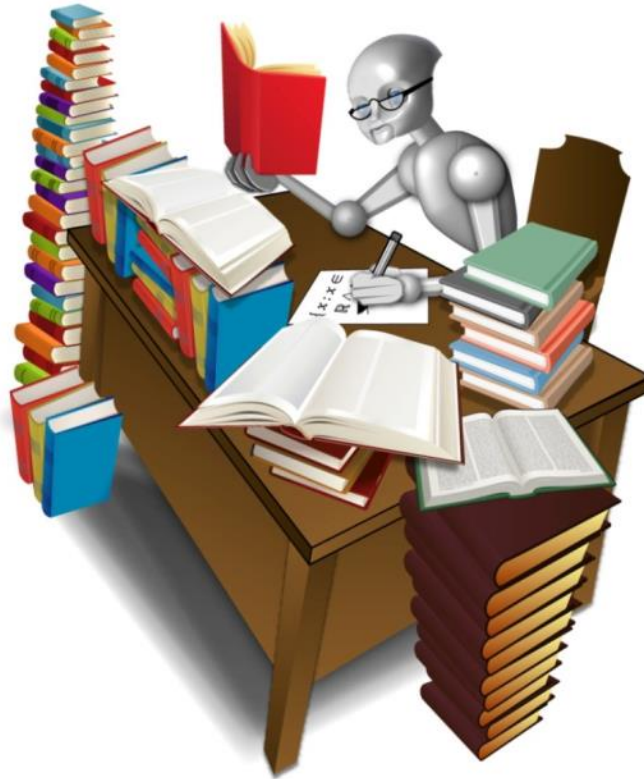
Observations from the Tour

Current KBs are surprisingly “knowledge poor”

- Not specific to Apples

What about obtaining knowledge from text?

III. Machine Reading



Auto-Text to Knowledge

Information Extraction (IE)

$IE(\text{sentence}) = [\text{tuple}, \text{confidence}]$

“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

Consider Tom Mitchell's NELL System

NELL Lexical Patterns for “haswife”

haswife metadata - Read the Web :: NELL KnowledgeBase Browser - Google Chrome

rtw.ml.cmu.edu/rtw/kbbrowser/predmeta:haswife

NELL Knowledge Base Browser

CMU Read the Web Project

log in | preferences | help/instructions | feedback

- haswife
- hashusband
- husbandof
- parentofperson
- motherofperson
 - biologicalmotherofperson
- fatherofperson
 - biologicalfatherofperson
- agentcollaborateswithagent
- politicianendorsespolitician
- politicianusendorsespoliticianus
- teammate
- politicianendorsesbypolitician
- politicianusendorsesbypoliticianus
- agenthierarchicallybelowagent (+)
- agenthierarchicallyaboveagent
- organizationhasagent (+)
- coachesathlete
- agentholdsshareincompany
- companyhasshareholder
- universityhasacademicprogram
- clothingmadefromplant
- creativeworkcontributedbyagent
- moviestaractor
- moviedirectedbydirector
- bookwriter
- personhasjobposition
- jobpositionheldbyperson
- inverseofsportfansincountry
- agentcontrols
 - agentleadsorganization (+)
 - personleadscountry
 - agenthierarchicallyaboveagent (+)
- trophywonbyteam
- subpartof
 - agentbelongstoorganization (+)
 - agentparticipatedinevent
 - teamingame
 - losinggame
 - homeleamingame
 - winninggame
 - awayteamingame
 - personchargedwithcrime
 - locationlocatedwithinlocation (+)
 - bodypartwithinbodypart
 - eventsubeventof
 - subfieldofacademicfield
 - iteminvolvedinevent
- agentstudiesphysiologicalcondition
- players
- issueofpoliticsgroup
- visualartistartform
- arthropodlookslikeinsect
- agriculturalproductcookedwithagricultural
- teamwontrophy
- ethnicgroupofperson
- agriculturalproductcomingfromvertebrate
- shapeofobject

haswife
(relation: inverse of [wifeof](#))

See [learned instances](#) of haswife

Metadata

- antireflexive
 - "true"
- antisymmetric
 - "true"
- description
 - "person has a wife"
- domain
 - person
- domainWithinRange
 - "false"
- extractionPatterns
 - "arg1 a bad influence on and actress Angelina ar arg2" "arg1 and fiance World arg2" "arg1 and and her lover arg2" "an arg2" "arg1 and his for" "arg1 and his mother we arg2" "arg1 and his wor arg2" "arg1 and Sarah to arg2" "arg1 Appears bought his wife arg2" "cheated on arg2" "arg1 committed adultery with" "arg1 endeavors to kee" "arg1 gal pal arg2" "arg with arg2" "arg1 had ar arg2" "arg1 has annou married to arg2" "arg1 "arg1 is banging arg2" arg2" "arg1 leaves to g arg2" "arg1 married the affair with arg2" "arg1 p separated after arg2" "arg1 set to marry arg2" "arg1 shirtless and arg2" "arg1 showed up and told arg2" "arg1 sinned with arg2" "arg1 son by arg2" "arg1 splashes out on arg2" "arg1 still pining for arg2" "arg1 takes wife arg2" "arg1 the husband of arg2" "arg1 then married arg2" "arg1 to banish arg2" "arg1 to bring his wife arg2" "arg1 to cast out arg2" "arg1 to have a child by arg2" "arg1 to have a child with arg2" "arg1 to have sex with arg2" "arg1 to pop the question to arg2" "arg1 to send him and arg2" "arg1 to sex with arg2" "arg1 to star with arg2" "arg1 to tie the knot arg2" "arg1 took his wife arg2" "arg1 toled arg2" "arg1 used to kiss arg2" "arg1 VIII and arg2" "arg1 VIII married arg2" "arg1 was actually married to arg2" "arg1 was joined by his wife arg2" "arg1 was making out with arg2" "arg1 was married to arg2" "arg1 was never married to arg2" "arg1 was not married to arg2" "arg1 was not present when arg2" "arg1 was once married to arg2" "arg1 was photographed holding arg2" "arg1 was secretly married to arg2" "arg1 was sexually involved with arg2" "arg1 Wentz and arg2" "arg1 were a reality arg2" "arg1 who affirmed that arg2" "arg1 will break up with arg2" "arg1 with his new love arg2" "arg1 with love to arg2" "arg1 yong arg2" "arg2 Adams and David arg1" "arg2 admits that she and arg1" "arg2 and actor Brad arg1" "arg2 and baby daddy arg1" "arg2 and Barack Obama to arg1" "arg2 and beau arg1" "arg2 and boyfriend arg1" "arg2 and Brad arg1" "arg2 and brad arg1" "arg2 and Brad Pitt arg1" "arg2 and country music star arg1" "arg2 and country singer arg1" "arg2 and dad is arg1" "arg2 and david arg1" "arg2 and David arg1" "arg2 and debutant arg1" "arg2 and estranged husband arg1" "arg2 and ex husband arg1" "arg2 and fianc arg1" "arg2 and fiance arg1" "arg2 and Fiance arg1" "arg2 and gal pal arg1" "arg2 and girlfriend arg1" "arg2 and Good Charlotte rocker arg1" "arg2 and guy arg1" "arg2 and Guy arg1" "arg2 and her beau arg1" "arg2 and her BF arg1" "arg2 and her boyfriend arg1" "arg2 and her children with arg1" "arg2 and her daughter by arg1" "arg2 and her ex arg1" "arg2 and her fianc arg1" "arg2 and her fiance arg1" "arg2 and her former husband arg1" "arg2 and her gal pal arg1" "arg2 and her girlfriend arg1" "arg2 and her hubby arg1" "arg2 and her husband arg1" "arg2 and her husband Guy arg1" "arg2 and her husband Tom arg1" "arg2 and her longtime boyfriend arg1" "arg2 and her man arg1" "arg2 and her new boyfriend arg1" "arg2 and her new hubby arg1" "arg2 and her new husband arg1" "arg2 and her partner arg1" "arg2 and her partner Brad arg1" "arg2 and her relationship to arg1" "arg2 and her rocker boyfriend arg1" "arg2 and her rocker husband arg1" "arg2 and her spouse arg1" "arg2 and hottie arg1" "arg2 and hubbie arg1" "arg2 and hubby arg1" "arg2 and husband arg1" "arg2 and husband Guy arg1" "arg2 and husband Tom arg1" "arg2 and its spinoff arg1" "arg2 and Jack Lemmon in arg1" "arg2 and longtime boyfriend arg1" "arg2 and Mary the mother of arg1" "arg2 and Mary the Mother of arg1" "arg2 and new boyfriend arg1" "arg2 and new husband arg1" "arg2 and other half arg1" "arg2 and partner Brad arg1" "arg2 and Prince Harry arg1" "arg2 and rocker arg1" "arg2 and stepdaughter of arg1" "arg2 and suri arg1" "arg2 and the Claim arg1" "arg2 and the first season of arg1" "arg2 and the last episode of arg1" "arg2 and Tom arg1" "arg2 and tom arg1" "arg2 Angelina Jolie and arg1" "arg2 announced her engagement to arg1" "arg2 anointed the feet of arg1" "arg2 as being closer to arg1" "arg2 as the companion of arg1" "arg2 as the receptacle of arg1" "arg2 bare unto arg1" "arg2 became engaged to arg1" "arg2 became pregnant by arg1" "arg2 became pregnant with arg1" "arg2 became the bride of arg1" "arg2 became the wife of arg1" "arg2 Beckham and David arg1" "arg2 began a relationship with arg1" "arg2 breaking up with arg1" "arg2 broke up with arg1" "arg2 bullock arg1" "arg2 came to the tomb of arg1" "arg2 cheat on arg1" "arg2 Cheating on arg1" "arg2 could return with arg1" "arg2 eloped with arg1" "arg2 Engaged to arg1" "arg2 esposa arg1" "arg2 Fergie and arg1" "arg2 filed for divorce from arg1" "arg2 Files for Divorce from arg1" "arg2 Files for Divorce from arg1" "arg2 gal pal arg1" "arg2 gave birth to her and arg1" "arg2 gets engaged to arg1" "arg2 getting engaged to arg1" "arg2 got engaged to arg1" "arg2 got married to arg1" "arg2 had a child by arg1" "arg2 had no hand in arg1" "arg2 had to kill

NELL Knowledge Base Browser

CMU Read the Web Project

log in | p

categories

relations

- everypromotedthing
- abstractthing
 - creativework
 - book
 - poem
 - lyrics
 - musicalbum
 - musicsong
 - televisionshow
 - movie
 - visualartform
- species
 - animal
 - vertebrate
 - bird
 - fish
 - reptile
 - mammal
 - amphibian
 - invertebrate
 - arthropod
 - insect
 - crustacean
 - arachnid
 - mollusk
 - archaea
 - bacteria
 - fungus
 - plant
 - physicalaction
 - celltype
 - event
 - eventoutcome
 - conference
 - mlconference
 - sportsevent
 - olympics
 - race
 - grandprix
 - sportsgame
 - convention
 - crimeorcharge
 - filmfestival
 - musicfestival
 - weatherphenomenon
 - election
 - militaryeventtype
 - militaryconflict
 - color
 - date
 - dayofweek
 - year
 - dateliteral
 - month
 - programminglanguage
 - protein
 - chemical
 - buildingmaterial
 - mldataset

Ask NELL on-demand results:

4 possible entities found

Click to change visible entity:

- [apple_inc \(organization\)](#)
- [apple \(magazine\)](#)
- [apple \(company\)](#)
- [apple \(hotel\)](#)

apple_inc (organization)

literal strings: [apple computers](#), [Apple inc.](#), [-Apple](#), [Apple](#), [Apple Computer Inc.](#), [Apple](#), [â€œApple](#), [APPLE](#), [Apple Computer](#), [Apple Inc.](#), [apple computer](#), [APPLE COMPUTER](#), [APple](#), [apple-computer](#), [apple](#), [Apple computer](#), [Apple computer](#), [Apple Computers](#)

categories

- [company](#)(100.0%)
 - SEAL @112 (100.0%) on 03-jun-2010 [why?](#) using apple
 - CPL @722 (100.0%) on 06-apr-2013 [why?](#) using apple
 - Seed
 - MBL @728 (99.9%) on 29-apr-2013 [why?](#) using concept:biotechcompany:apple_inc

NELL Knowledge Base













- categories
- relatedto
 - visualartmovement
 - personinacademy
 - generalizationof
 - athletessuchas
 - academicfields
 - chemicaltype
 - astronautsuccess
 - automobileeng
 - airportsuchasa
 - amphibianssuc
 - aquariumssuc
 - arteriessuchas
 - animaltypehas
 - animalsuchas
 - animalsuc
 - animalsuc
 - animalsuch
 - inverseofart
 - professiontype
 - productinstan
 - booksuchasbo
 - vehicletypehas
 - touristattractio
 - celebritysucha
 - archaeasuchas

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

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

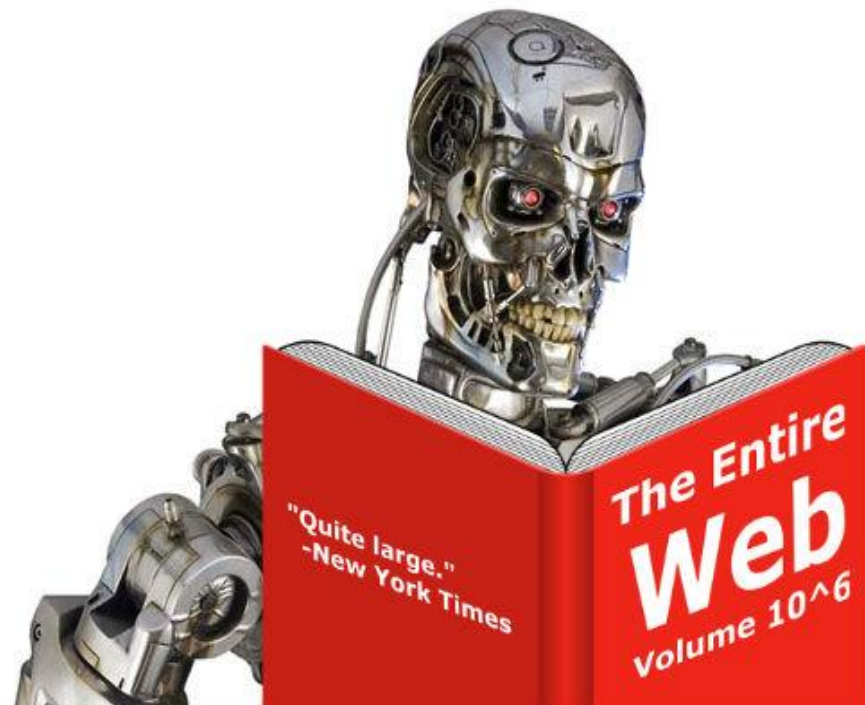


Relations often:

- anchored in verbs.
- exhibit simple syntactic form

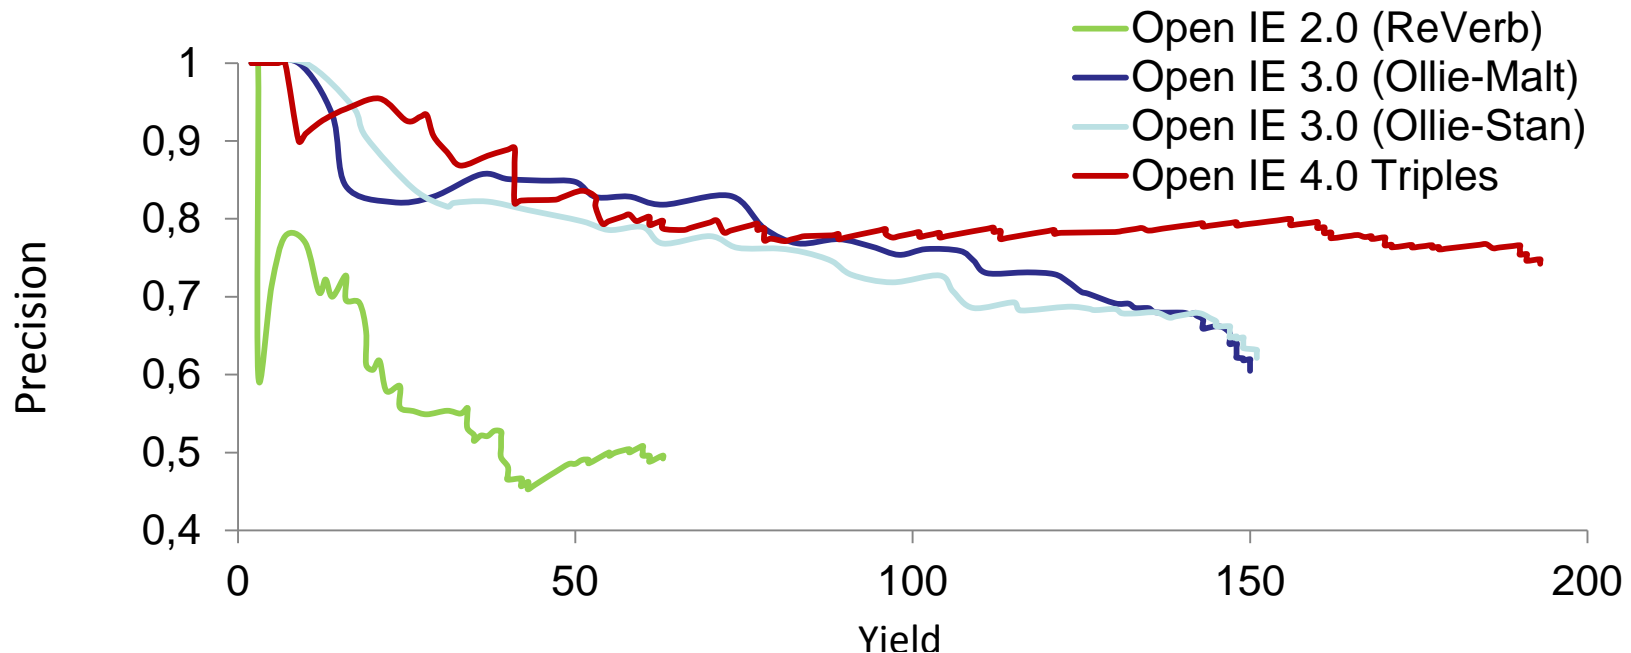
Virtues:

- No hand-labeled data
- “No sentence left behind”
- Exploit redundancy of Web



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)
- 2014: 5.0 (*multiple sentences, complex structures*)




Open Information Extraction - Google Chrome

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

AppsO-1 Person of Extra...WeeklyPlans/Decem...British Airways - Wh...74Vulcan - Project Hal...Seattle Department ...New Techniques fro...Extraction Browser - ...deep learning tutorialFacebook launches ...An AI that mimics o...

Open Information Extraction



Argument 1:

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](#)

30

ALLEN INSTITUTE
for ARTIFICIAL INTELLIGENCE



Open Information Extraction

Argument 1: Relation: Argument 2:

All ▾

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. The most common ancestor, *Malus sieversii*, is still found in the mountains of Central Asia.

Open IE suffers from Attention Deficit Disorder!

all ingredient (37) food (33)

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)

ReVerb

Argument 1:

what/who

all ingredient (38)

Vegetable (42)

Green tea (23)

Fruit (22)

Wine (20)

Relation:

contains

Argument 2:

antioxidant

Freebase

Freebase Find...

Topic

Tao Lin en

mid: /m/025wtg1 notable

Tao Lin is an American not

film production company.

Properties 118n

People /people

Person /people/person

Date of birth /people/person/date_of_birth

7/2/1983

Place of birth /people/person/place_of_birth

Alexandria

Country of nationality /people/person/nationality

United States of America

Gender /people/person/gender

Male

Profession /people/person/profession

Dr. Tony Fader (AI2): Open Question Answering KDD 2014

Read

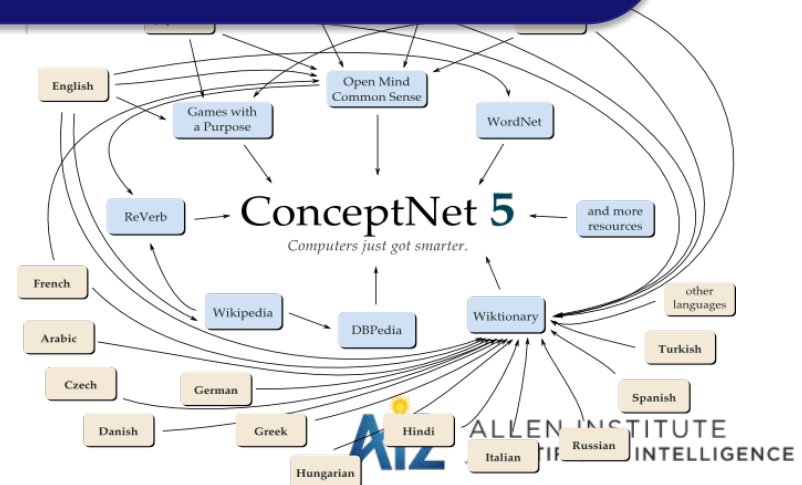
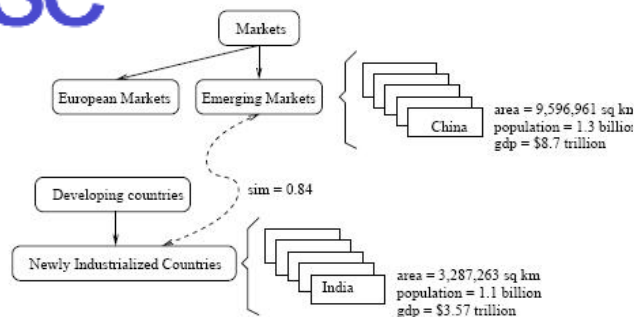
Research Proj

Home

Pr

NELL: Never-

ProBase



IV. The Allen Institute for AI (AI2)

“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.”

Paul Allen, Microcomputer Interface, **1977.**

AI2 Mission: *high-impact AI research
for the common good*

Mission + Talent = Impact

Time Line

AI2 launched

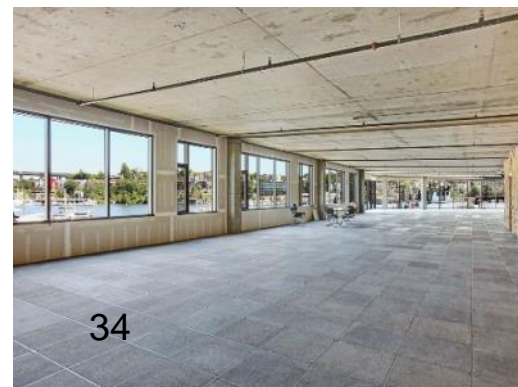
Jan. 2014

Team of 30

August

Team of 50;
Machine Vision Project

Dec. 2015



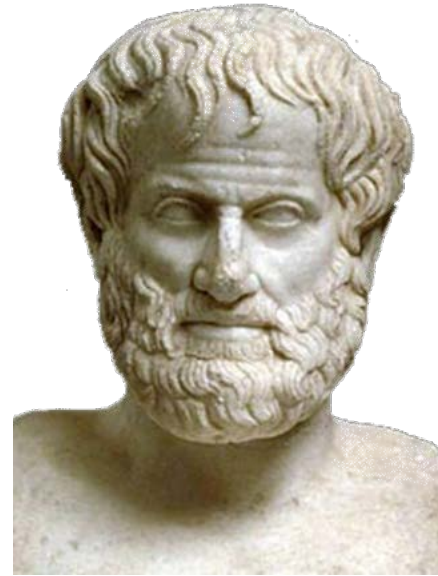
AI2 Overview (Nonprofit Research Institute)

- Experimental science projects
- Targeted collaborations
- Launched AI Allen Distinguished Investigator (ADI) program
- Open source projects

Small number of “Move the needle” projects

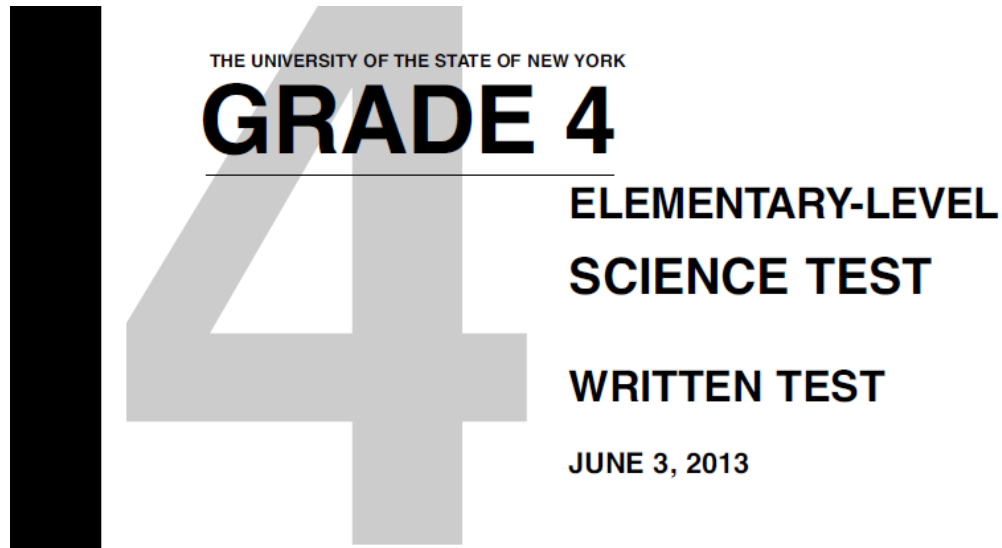
Project Methodology:

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



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

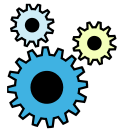
- A friction
- B electricity
- C gravity
- D magnetism

Factual Knowledge for 4th Grade Science



Summary of our Approach

1. Knowledge + tractable reasoning is necessary
2. Knowledge acquisition has to be highly automated
3. **Large** bodies of **high-quality** knowledge can be acquired from text (“machine reading”) images, and more.
4. Aristo is tested on unseen, standardized, multiple choice tests in science, arithmetic, and geometry



Semantic challenges everywhere...

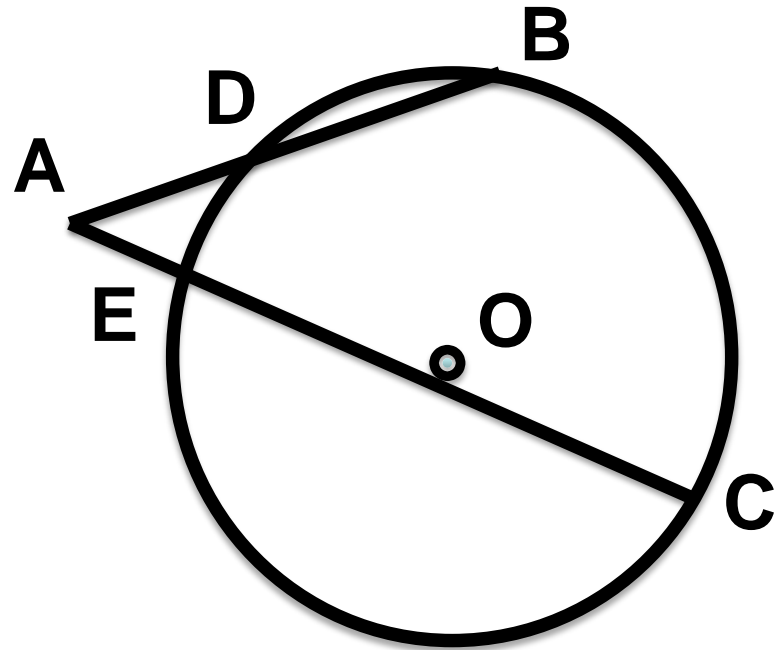
Geometry Domain (Case Study)

Combines vision, NLP, and simple semantics

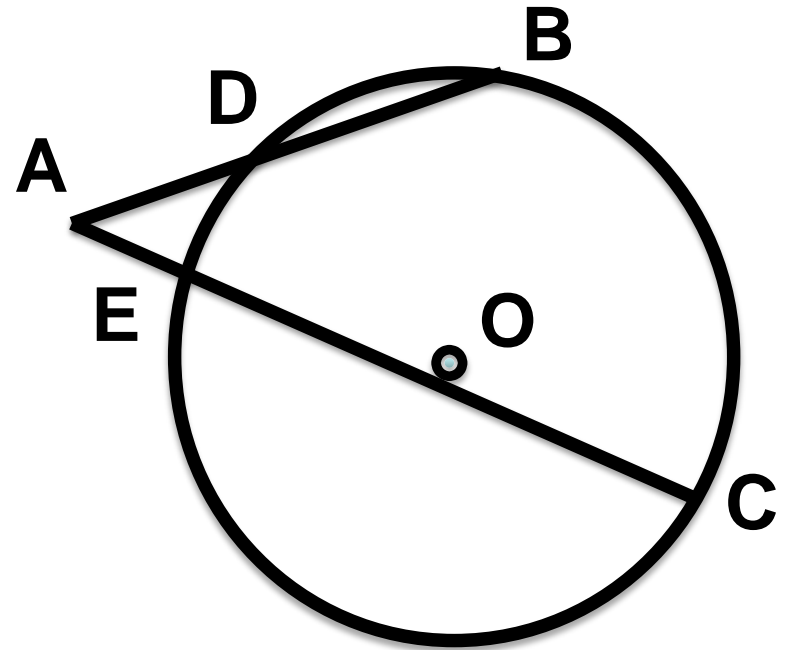
Data set: 100 9th grade geometry problems

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.



Geometry Diagram Alone



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

Need Both Diagram and Text

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

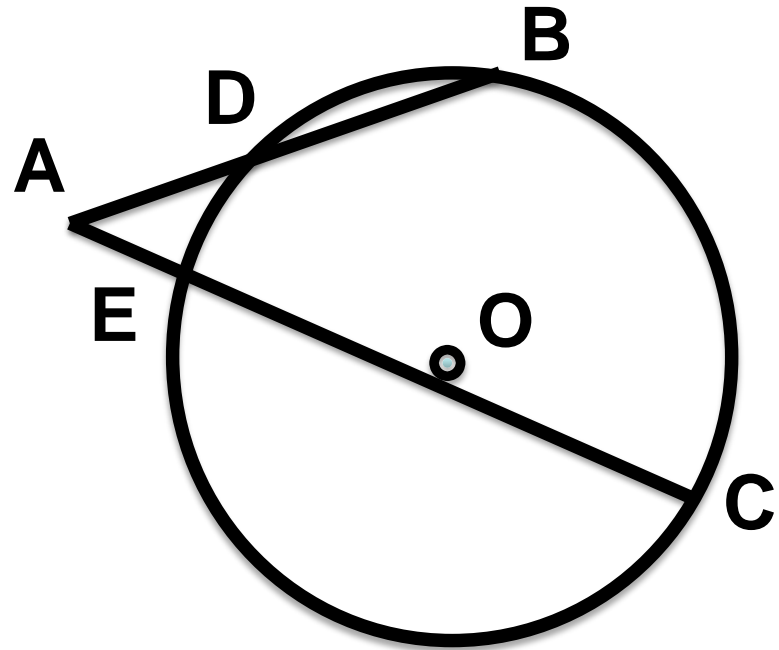


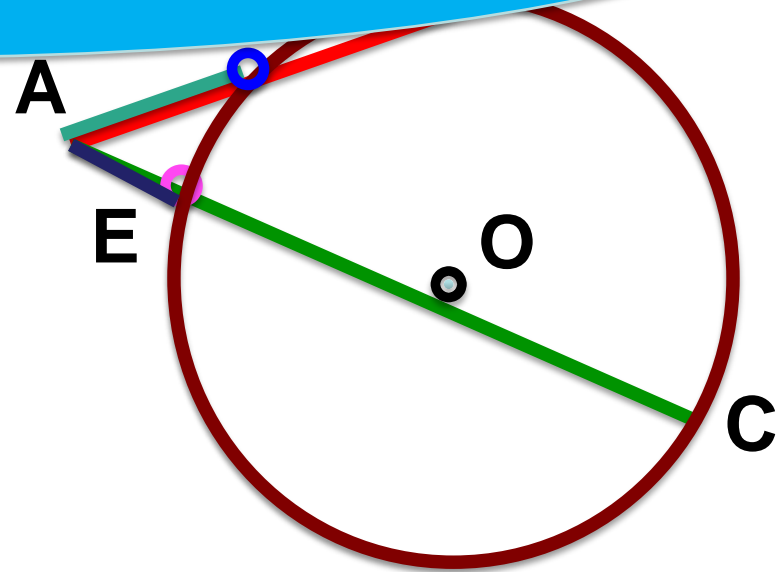
Diagram Understanding

Problem:

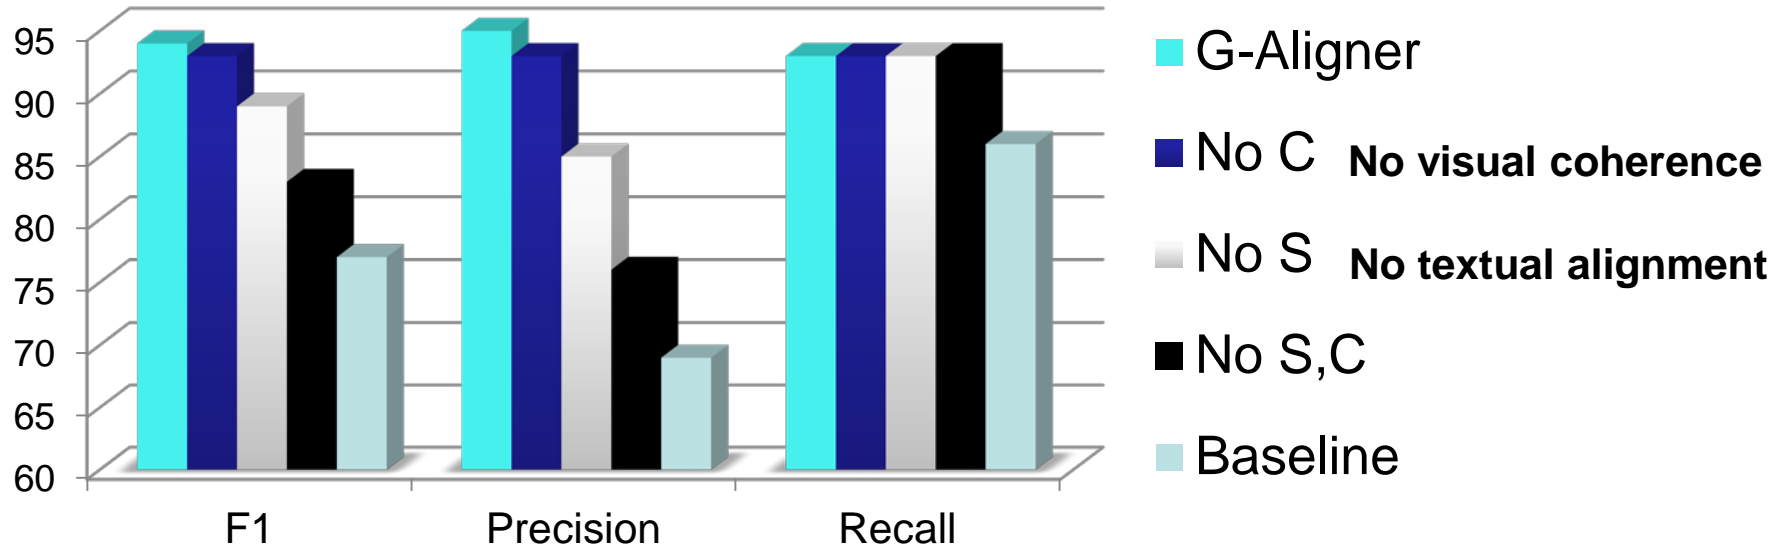
1. Identify visual elements
2. Apply

Diagram understanding → submodular optimization

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

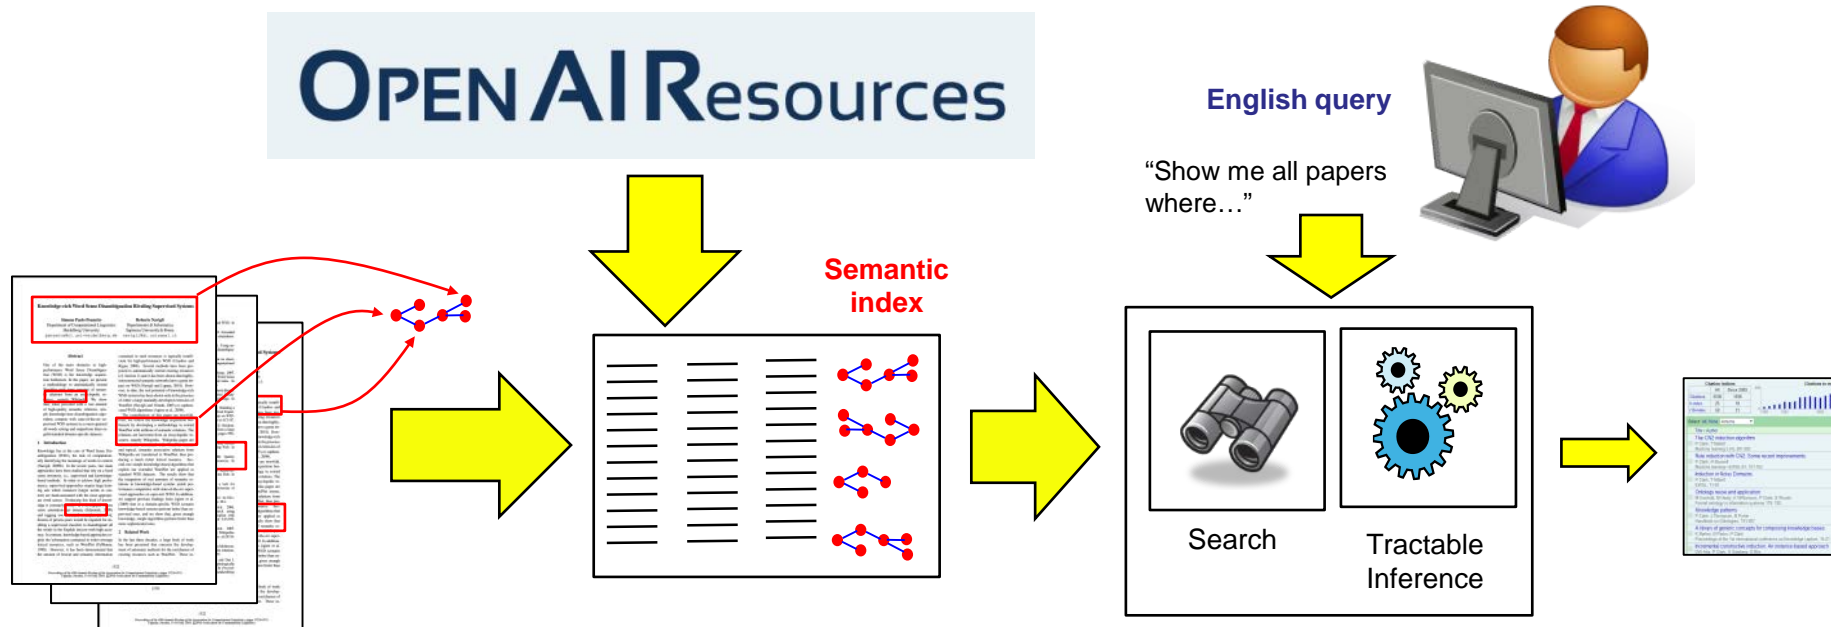


Experimental Study



“Semantic Scholar” Project

“Show all papers where x is implicated in y where the levels of z are elevated”...

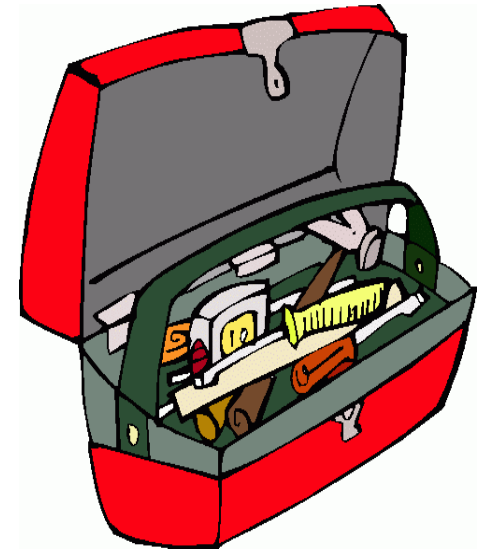


“lift all scientific boats”



Identify time-saving details early

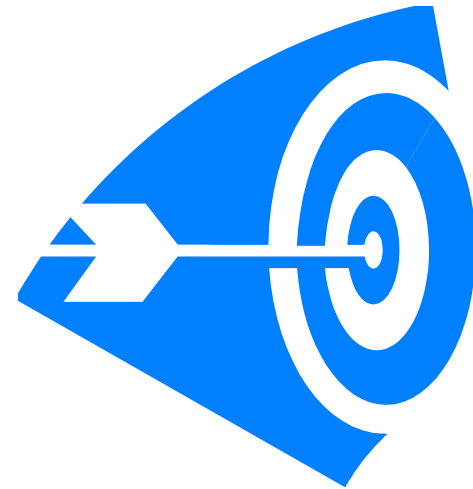
Focus on Tools, viz



Design for replicable experiments

Summary

- What's Next? (after “big data” & “deep learning”)
“You can’t play 20 questions with Nature and Win!”
- Focus on data/text Mining systems yielding:
 - Knowledge
 - Reasoning
 - Explanation



Thank you from AI2!



What is Data Mining?

KDD 2014

This year's special theme:

Data Science for Social Good

10th ACM SIGKDD Conference on Knowledge Discovery and Data Mining

August 24-27, 2014 New York

