


You Can't Play 20 Questions with Nature And Win

Oren Etzioni, CEO
Allen Institute for AI (AI2)




Big Data Tidal Wave

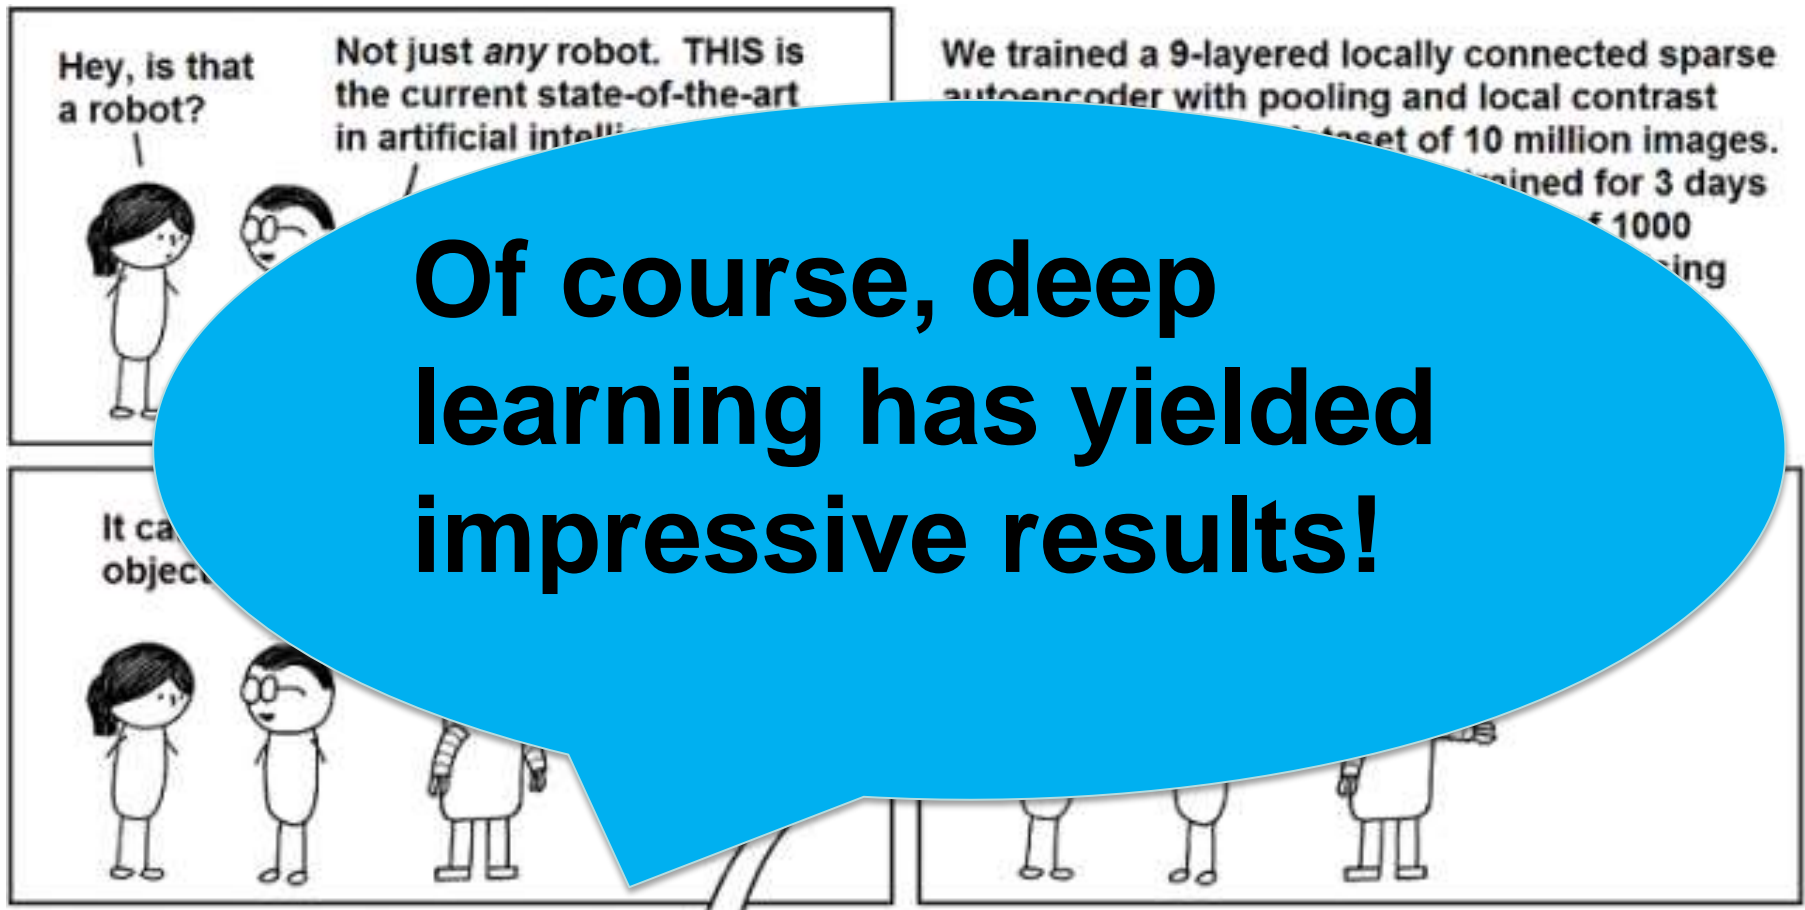




Deep Learning




Some Skepticism about Deep Learning



Outline

- I. Limits of current paradigm
- II. Today's KBs: fact rich but knowledge poor
- III. Machine Reading
- IV. Allen Institute for AI (AI2)



What are the limits of the classification paradigm?

Critique applies to both discriminative and generative 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?

Can you guess which tweet was retweeted more often?

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



wee

House of Cards
2013-2014

Bad, for a great
A ruthless politician
conquer Washington
Golden Globe winning p

Starring: Kevin Spacey, H

Based on your interest in: *Orange Is the New Black*

oren's rating
★★★★☆

Not Interested



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.




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 rules & procedures
 - 100% accuracy
 - Distribution insensitive learning
- **Playing Chess:** move choice requires search
Deep Blue
human grand masters


Reminiscent of...

- **System-one:** fast, associative, but biased, error prone.
- **System two:** deliberative, explicit, more reliable.



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 → reasoning



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

Allen Newell, 1973



**Key AI tasks cannot be reduced to
classification problems!
(using realistic feature sets)**

Natural Language Understanding

- Complex syntax, semantics, pragmatics, discourse, reference, context, etc.
- **Background knowledge**

The large table is on the table
because

**World knowledge is
necessary for elementary
understanding, so do we
have it?**

It

Simple Thought Experiment

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...

2014 Knowledge Tour



apple



+Oran

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

About 397,000,000 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, Secre...](#)
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](#)



"apple fruit"



+Oran


[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/1/apple - Encyclopaedia 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/h3177 - 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](#)

[More images](#)

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%	Vitamin C	14%
Calcium	1%	Iron	1%
Vitamin D	0%	Vitamin B-6	5%
Vitamin B-12	0%	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 [abryuch, 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/Mx8Ngh4rvVipdpwEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA]
OpenCyc (Versioned): [http://sw.opencyc.org/2012/05/10/concept/Mx8Ngh4rvVipdpwEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA]
OpenCyc (Readable): [http://sw.opencyc.org/2012/05/10/concept/en/Apple]

Search



OpenCyc Collection: apple

Unique ID: [Mx8Ngh4rvVipdpwEbGdrcN5Y29ycB4rvVjBnZwpEbGdrcN5Y29ycA]

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.



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

1. The truth is "out there"
2. Not specific to Apples
3. Famous KBs have gaping holes

What about obtaining knowledge from text?

III. Machine Reading



Auto-Text to Knowledge

Information Extraction (IE)

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

Consider Tom Mitchell's NELL System

- categories
- relations
- everypromotedthing
- abstracting
 - creativework
 - book
 - poem
 - lyrics
 - musicalbum
 - music song
 - 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
 - building
 - 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 Computer Inc.](#), [Apple](#), [MacApple](#), [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
- personinacademia
- generalizationof
 - athletessuchas
 - academicfields
 - chemicaltype
 - astronautsuccess
 - automobileeng
 - airportsuchasa
 - amphibianssuc
 - aquariumssuc
 - arteriessuchas
 - animaltypephas
 - animalsuchas
 - animalsuc
 - animalsuc
 - animalsuc
 - 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



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 (36)

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)

Knowing is not enough ,
We must APPLY.

Willing is not enough ,
We must DO.

- *B*ruce *L*ee



IV. The Allen Institute for AI (AI2)

Non-profit Research Institute:

- Basic & applied AI research
- Collaborations, internships, open source
- AI Allen Distinguished Investigator (ADI) program

**Small number of
High-impact projects**

Mission + Talent = Impact

Time Line

AI2 launched

Jan. 2014

Team of 30 + 15 interns

Oct. 2014


Team of 50

Dec. 2015



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

THE UNIVERSITY OF THE STATE OF NEW YORK

GRADE 4

ELEMENTARY-LEVEL SCIENCE TEST

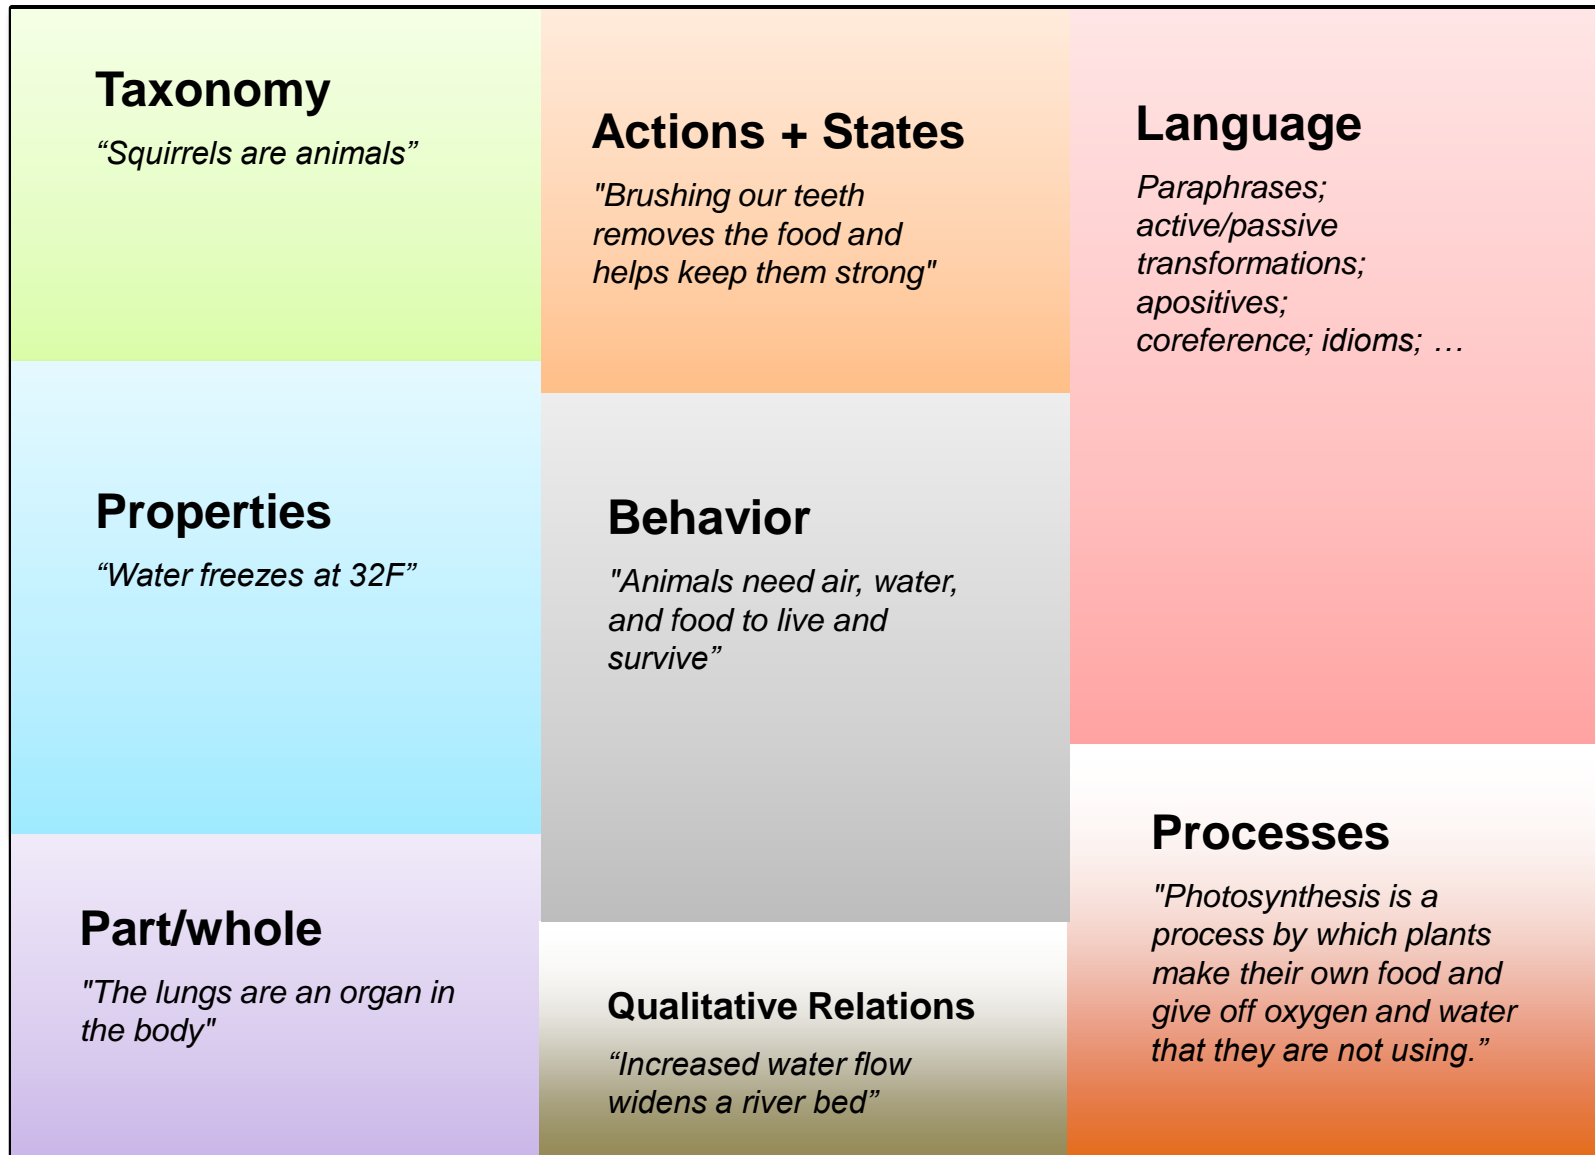
WRITTEN TEST

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

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



AI challenges everywhere...

THE WEB OF SCHOLARSHIP

Around 114 million English-language scholarly documents (including papers, books and technical reports) can be found on the web.





"information extraction"



Scholar

About 139,000 results (0.12 sec)

Articles

Case law

My library

[PDF] [Maximum Entropy Markov Models for **Information Extraction** and Segmentation.](#)

[A McCallum](#), [D Freitag](#), [FCN Pereira](#) - ICML, 2000 - [courses.ischool.berkeley.edu](#)

Page 1. 1 Maximum Entropy Markov Models for **Information Extraction** and Segmentation Andrew McCallum, Dayne Freitag, and Fernando Pereira ... Named entity recognition: <ORG>Mips</ORG> Vice President <PRS>John Hime</PRS> - **Information extraction**: ...

Cited by 1126 Related articles All 50 versions Cite Save More

Any time

Since 2014

Since 2013

Since 2010

Custom range...

[Incorporating non-local information into **information extraction** systems by gibbs sampling](#)

[JR Finkel](#), [T Grenager](#), [C Manning](#) - ... of the 43rd Annual Meeting on ..., 2005 - [dl.acm.org](#)

Abstract Most current statistical natural language processing models use only local features so as to permit dynamic programming in inference, but this makes them unable to fully account for the long distance structure that is prevalent in language use. We show how to ...

Cited by 1129 Related articles All 23 versions Cite Save

Sort by relevance

Sort by date

- include patents
- include citations

[PDF] [Learning dictionaries for **information extraction** by multi-level bootstrapping](#)

[E Riloff](#), [R Jones](#) - AAAI/IAAI, 1999 - [aaai.org](#)

Abstract **Information extraction** systems usually require two dictionaries: a semantic lexicon and a dictionary of extraction patterns for the domain. We present a multilevel bootstrapping algorithm that generates both the semantic lexicon and extraction patterns simultaneously. ...

Cited by 753 Related articles All 18 versions Cite Save More

Create alert

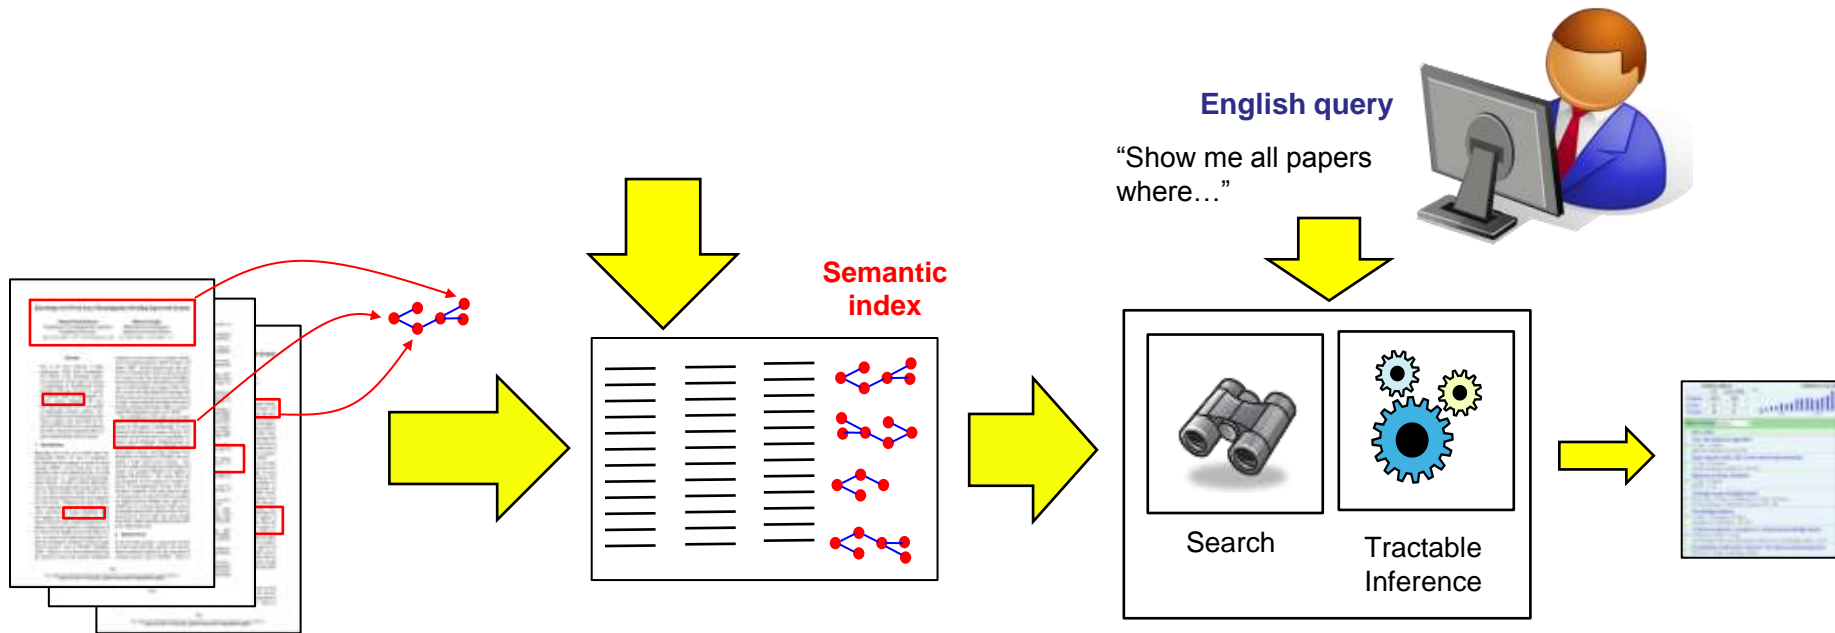
[PDF] [Open **information extraction** for the web](#)

[M Banko](#), [MJ Cafarella](#), [S Soderland](#), [M Broadhead](#)... - IJCAI, 2007 - [aaai.org](#)

Abstract Traditionally, **Information Extraction** (IE) has focused on satisfying precise, narrow, pre-specified requests from small homogeneous corpora (eg, extract the location and time of seminars from a set of announcements). Shifting to a new domain requires the user to ...

Cited by 806 Related articles All 31 versions Cite Save More

Leverage AI to combat information overload





Cut through the clutter.

Home in on key papers, citations, and results.

 Find it fast





Try: [Open information extraction](#) [POS tagging](#) [Dependency parsing](#)

Made By AI2 • powered by ACL Anthology Reference Corpus



FILTER RESULTS

CLASSIFICATION

-  Survey
-  Experimental
-  Theoretical
-  Software

YEAR

YYYY to YYYY

Update

VENUES (18)

TOPICS (92)

6,929 results

Showing 1-10

« ‹ 1 2 3 4 5 › »

Sort by: Relevance

On-Demand Information Extraction

Satoshi Sekine / ACL / 2006

Cited by 22 / Abstract / View PDF / Add to reading list

At present, adapting an **Information Extraction** system to new topics is an expensive and slow for each new topic. We propose a new paradigm of **Information Extraction** which operates 'on demand

Probabilistic Coreference In Information Extraction

Andrew Kehler / Conference On Empirical Methods In Natural Language Processing / 1997

Cited by 17 / Abstract / View PDF / Add to reading list

Certain applications require that the output of an **information extraction** system be probabilistic **Extraction** Abstract Certain applications require that the output of an **information...**

Multidocument Summarization Via Information Extraction

Michael White, Tanya Korelsky, Claire Cardie, Vincent Ng, David R. Pierce, Kiri Wagstaff / HLT / 2001

Cited by 11 / Abstract / View PDF / Add to reading list

We present and evaluate the initial version of RIPTIDES, a system that combines **information information extraction, extraction**-based summarization, and natural language generation to...

Confidence Estimation For Information Extraction

Aron Culotta, Andrew McCallum / NAACL / 2004

Cited by 19 / Abstract / View PDF / Add to reading list



216 results

Showing 1-10

« ‹ 1 2 3 4 5 › »

Sort by: Relevance

Clear all filters

FILTER RESULTS

CLASSIFICATION

- Survey
- Experimental
- Theoretical
- Software

YEAR

yyyy to yyyy

Update

VENUES (18)

TOPICS (92)

Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

Jenny Rose Finkel, Trond Grenager, Christopher D. Manning / ACL / 2005

Cited by 166 / Abstract / View PDF / In reading list / Software

use this technique to augment an existing CRF-based **information extraction** system with long-distance inference. We use this technique to augment an existing CRF-based **information extraction** system with...

Semantic Role Labeling for Open Information Extraction

Janara Christensen, undefined Mausam, Stephen Soderland, Oren Etzioni / Proceedings of the NAACL HLT 2010 First International Workshop on Formalisms and Methodology for Learning by Reading / 2010

Cited by 2 / Abstract / View PDF / Add to reading list / Software

Open **Information Extraction** is a recent paradigm for machine reading from arbitrary text. In Abstract Open **Information Extraction** is a recent paradigm for machine reading from arbitrary text

Multi-Field Information Extraction And Cross-Document Fusion

Gideon S. Mann, David Yarowsky / ACL / 2005

Cited by 9 / Abstract / View PDF / Add to reading list / Software

examples of fact extractions and train Rete, Naïve Bayes, and Conditional Random Field **extraction** Introduction Much recent statistical **information extraction** research has applied graphical models to...

An Entity-Level Approach to Information Extraction

Aria Haghighi, Dan Klein / ACL / 2010



216 results

Showing 1-10

Sort by: Relevance

Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

Scholar Rank:



Top 2%

Citation rate:



63.8

Cited by 166 papers:



7 are important

4.2%

TOPICS

CRF model Conditional Markov Models Monte Carlo algorithm Conditional Random Fields Viterbi algorithm

CITATION CONTEXTS

In this paper, Stanford Named Entity Recognizer (Finkel et al. 2005) is used to classify noun phrases into four semantic categories: PERSON, LOCATION, ORGANIZARION and MISC.

We utilized an offthe-shelf system, Stanford Named Entity Recognizer 4 (Finkel et al., 2005) for detecting entity mentions on the English sentences.

[View more citation contexts](#)

[Cited by 9](#) / [Abstract](#) / [View PDF](#) / [Add to reading list](#) / [Software](#)

examples of fact extractions and train Rote, Na" ive Bayes, and Conditional Random Field **extraction** Introduction Much recent statistical **information extraction** research has applied graphical models to...

An Entity-Level Approach to Information Extraction

Aria Haghighi, Dan Klein / ACL / 2010

Clear all filters

FILTER RESULTS

CLASSIFICATION

- Survey
- Experimental
- Theoretical
- Software ×

YEAR

yyyy to yyyy

Update

VENUES (18)

TOPICS (92)



Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

- Abstract
- View PDF
- In Reading List
- Software

Scholar Rank:



Citation rate:



Authors:

- Jenny Rose Finkel
- Trond Grenager
- Christopher D. Manning

Source:

ACL

Published:

2005

Topics:

- CRF Model
- Conditional Markov Models

Citing papers Cluster map

Cited by 166 papers 7 are important (4.2%)

Sort by: Importance

IMPORTANT

The Web is not a PERSON, Berners-Lee is not an ORGANIZATION, and African-Americans are not LOCATIONS: An Analysis of the Performance of Named-Entity Recognition

Robert Krovetz, Paul Deane, Nitin Madnani / Proceedings of the Workshop on Multiword Expressions: from Parsing and Generation to the Real World / 2011

Abstract / View PDF / Add to reading list

Citation contexts:

"The Stanford Tagger is based on Conditional Random Fields (Finkel et al., 2005)."

"Although Named Entity Recognition is reported in the literature to have an accuracy rate of 85-95% (Finkel et al., 2005; Ratinov and Roth, 2009), it was clear by inspection that both the Stanford and the LBJ tagger made a number of mistakes."

IMPORTANT

Guessing Parts-Of-Speech Of Unknown Words Using Global Information

Tetsuji Nakagawa, Yuji Matsumoto / ACL / 2006

Cited by 6 / Abstract / View PDF / Add to reading list

Citation contexts:

"Global information is known to be useful in other NLP tasks, especially in the named entity recognition task, and several studies successfully used global features (Chieu and Ng, 2002; Finkel et al., 2005)."

"Finkel et al. (2005) used simulated annealing with Gibbs sampling to find a solution in a similar situation."

TOPICS

- Relaxation
- Markov random fields
- Systems based on perceptron
- POS tagging
- Dual Decomposition
- Maximum Entropy Markov Models
- JNLPBA
- F-score
- L1 regularization
- CRF models
- Entity Recognition
- NER system
- Conditional Markov Models
- Gibbs sampling
- Information extraction
- POS tags of unknown words
- Global Information

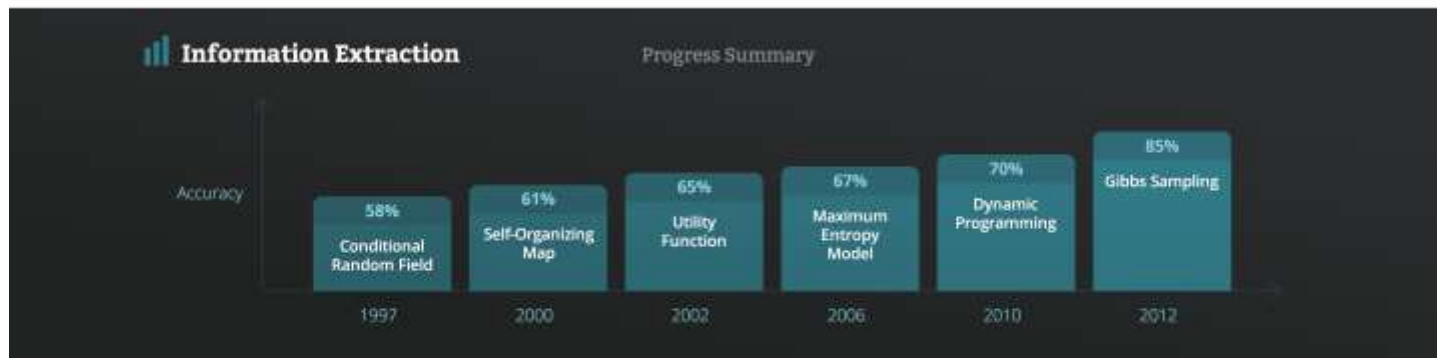
Cut through the clutter.

Home in on key papers, citations, and results.

 information extraction

information extraction progress summary

- **information extraction** using *gibbs sampling*
146 papers, 41.7 average rank
- **information extraction** using *dynamic programming*
598 papers, 41.7 average rank
- **information extraction** using *maximum entropy model*
1,086 papers, 32.4 average rank
- **information extraction** using *self-organizing map*
36 papers, 30.24 average rank
- **information extraction** using *conditional random field*
358 papers, 24.77 average rank
- **information extraction** using *preprocessing*
1,148 papers, 23.31 average rank



FILTER RESULTS

CLASSIFICATION

- Survey
- Experimental
- Theoretical
- Software

YEAR

YYYY to YYYY

VENUES (21)

- HLT
- NAACL
- Workshop On Automatic Summarization
- ACL
- Workshop On Operational Factors In Practical, Robust Anaphora Resolution For Unrestricted Tests

View all venues

6,962 results Sort by: Relevance

Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

Jenny Rose Finkel, Trond Grenager, Christopher D. Manning · ACL · 2005

Cited by 166 · Abstract · View PDF · Add to reading list

structure that is prevalent in language use. We show how to solve this dilemma with Gibbs sampling **information extraction** task. We show 10 runs of Gibbs sampling in the same CRF...

On-Demand Information Extraction

Satoshi Sekine · ACL · 2006

Cited by 22 · Abstract · View PDF · Add to reading list

At present, adapting an **Information Extraction** system to new topics is an expensive and slow for each new topic. We propose a new paradigm of **Information Extraction** which operates 'on demand

Multidocument Summarization Via Information Extraction

Michael White, Tanya Korelsky, and 4 others · HLT · 2006

Cited by 11 · Abstract · View PDF · Add to reading list

We present and evaluate the initial version of RIPTIDES, a system that combines **information extraction, extraction**-based summarization, and natural language generation to...

Confidence Estimation For Information Extraction

Aron Culotta, Andrew McCallum · NAACL · 2002

Cited by 19 · Abstract · View PDF · Add to reading list

Information extraction techniques automatically create structured databases from unstructured data Abstract **information extraction** techniques automatically create structured databases from

Using Predicate-Argument Structures For Information Extraction

Mihai Surdeanu, Sanda M. Harabagiu, John Williams, Paul Aarseth · ACL · 2000



Gibbs Sampling in Information Extraction

KEY PAPER

Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

Jenny Rose Finkel, Trond Grenager, Christopher D. Manning / ACL / 2012

An illustration of the effectiveness of **Gibbs sampling**, compared to Viterbi inference, for the two tasks addressed in this paper: the CoNLL named entity recognition task **which returned an accuracy rate of 85.54%**, and the CMU Seminar Announcements **information extraction** task. We show 10 runs of **Gibbs sampling** in the same CRF model that was used for Viterbi. For each run the sampler was initialized to a random sequence, and used a linear annealing schedule that sampled the complete sequence 1000 times. CoNLL performance is measured as per-entity, and CMU Seminar Announcements performance is measured as per-token.

FILTER RESULTS

CLASSIFICATION

- Survey
- Experimental
- Theoretical
- Software

YEAR

YYYY to YYYY

VENUES (21)

- HLT
- NAACL
- Workshop On Automatic Summarization
- ACL
- Workshop On Operational Factors In Practical, Robust Anaphora Resolution For Unrestricted Tests
- View all venues

TOPICS (215)

- Logistic Regression Model
- Log-Linear Models
- Expected F-Measure

On-Demand Information Extraction

Satoshi Sekine / ACL / 2006

Cited by 22 Abstract View PDF Add to reading list

At present, adapting an **Information Extraction** system to new topics is an expensive and slow for each new topic. We propose a new paradigm of **Information Extraction** which operates 'on demand

Multidocument Summarization Via Information Extraction

Michael White, Tanya Korelsky, and 4 others / HLT / 2006

Cited by 11 Abstract View PDF Add to reading list

We present and evaluate the initial version of RIPTIDES, a system that combines **information information extraction, extraction**-based summarization, and natural language generation to...

Confidence Estimation For Information Extraction

Aron Culotta, Andrew McCallum / NAACL / 2002

Cited by 19 Abstract View PDF Add to reading list

Information extraction techniques automatically create structured databases from unstructured data Abstract **information extraction** techniques automatically create structured databases from

Using Predicate-Argument Structures For Information Extraction

Mihai Surdeanu, Sanda M. Harabagiu, John Williams, Paul Aarseth / ACL / 2000

Cut through the clutter.

Home in on key papers, citations, and results.

● **gibbs sampling** overview of applications

- ↳ **gibbs sampling** in *information extraction*
 162 papers, 41.7 average rank
- ↳ **gibbs sampling** in *dependency parsing*
 105 papers, 18.71 average rank
- ↳ **gibbs sampling** in *parsing*
 159 papers, 32.4 average rank
- ↳ **gibbs sampling** in *machine translation*
 163 papers, 30.24 average rank
- ↳ **gibbs sampling** in *POS tagging*
 87 papers, 23.31 average rank



FILTER RESULTS

CLASSIFICATION

- Survey
- Experimental
- Theoretical
- Software

YEAR

YYY to YYY

VENUES (15)

- ACL
- Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data
- EMNLP
- CoNLL
- NAACL

View all venues

429 results Sort by: Relevance

Incorporating Non-Local Information Into Information Extraction Systems By Gibbs Sampling

Jenny Rose Finkel, Trond Grenager, Christopher D. Manning · ACL · 2012
 Cited by 166 · Abstract · View PDF · Add to reading list
 structure that is prevalent in language use. We show how to solve this dilemma with Gibbs sampling information extraction task. We show 10 runs of Gibbs sampling in the same CRF...

Not-So-Latent Dirichlet Allocation: Collapsed Gibbs Sampling Using Human Judgments

Jonathan Chang · Proceedings of the NAACL HLT 2010 Workshop on Creating Speech ... / 2010
 Cited by 1 · Abstract · View PDF · Add to reading list
 Probabilistic topic models are a popular tool for the unsupervised analysis of text, providing both ... and cluster that annotation. This task simulates the **sampling** step of the collapsed **Gibbs** sampler

Sampling Alignment Structure under a Bayesian Translation Model

John DeNero, Alexandre Bouchard-Côté, Dan Klein · EMNLP · 2008
 Cited by 31 · Abstract · View PDF · Add to reading list
 We describe the first tractable **Gibbs sampling** procedure for estimating phrase pair frequencies
 Abstract We describe the first tractable **Gibbs sampling** procedure for estimating phrase pair

A Gibbs Sampler for Phrasal Synchronous Grammar Induction

Philip Blunsom, Trevor Cohn, Chris Dyer, Miles Osborne · ACL · 2009
 Cited by 34 · Abstract · View PDF · Add to reading list
 a novel **Gibbs** sampler over synchronous derivations. This sampler side-steps the intractability -scale experiments and the short **sampling** runs. In this work we also propose a **Gibbs** sampler but...

Stochastic Attribute-Value Grammars

Mihal Sureshbabu, Sanda M. Harababu, John Williams, Paul Aarseth · ACL · 2003




Dependency Parsing using Gibbs Sampling

KEY PAPER

Unsupervised Dependency Parsing using Reducibility and Fertility features

David Mareček, Zdeněk Zabokrtský · NAACL · 2012

Inference	CoNLL	Seminar
Viterbi	85.51	91.85
Gibbs	85.54	91.85
Sampling	85.51	91.85
	85.49	91.85
	85.51	91.85
	85.51	91.85
	85.51	91.85
	85.51	91.85
	85.51	91.85
	85.51	91.85
Mean	85.51	91.85
Std. Dev.	0.01	0.004



Sort by: Relevance

- FILTER RE...
- CLASSIFIC...
- Survey
 - Exper...
 - Theor...
 - Softw...
- YEAR
- WWW

- VENUES (15)
- ACL
- Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data
- EMNLP
- CoNLL
- NAACL
- [View all venues](#)

- TOPICS (213)
- Stanford NER Tagger
- Stanford Named Entity Recognizer
- Krishnan And Manning
- Non-Local Features

Cited by 1 [Abstract](#) [View PDF](#) [Add to reading list](#)

Probabilistic topic models are a popular tool for the unsupervised analysis of text, providing both ... and cluster that annotation. This task simulates the **sampling** step of the collapsed **Gibbs** sampler

Sampling Alignment Structure under a Bayesian Translation Model

John DeNero, Alexandre Bouchard-Côté, Dan Klein · EMNLP · 2008

Cited by 31 [Abstract](#) [View PDF](#) [Add to reading list](#)

We describe the first tractable **Gibbs sampling** procedure for estimating phrase pair frequencies Abstract We describe the first tractable **Gibbs sampling** procedure for estimating phrase pair

A Gibbs Sampler for Phrasal Synchronous Grammar Induction

Philip Blunsom, Trevor Cohn, Chris Dyer, Miles Osborne · ACL · 2009

Cited by 34 [Abstract](#) [View PDF](#) [Add to reading list](#)

a novel **Gibbs** sampler over synchronous derivations. This sampler side-steps the intractability -scale experiments and the short **sampling** runs. In this work we also propose a **Gibbs** sampler but...


Stochastic Attribute-Value Grammars

Mihal Sierfanyi, Sanda M. Harabagiu, John Williams, Paul Aarseth · ACL · 2003

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



Parting Thought

"The best minds of my generation are thinking about how to make people click ads...

that sucks."



Jeff Hammerbacher

Our focus: AI for the Common Good

Thank you from AI2!




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 .

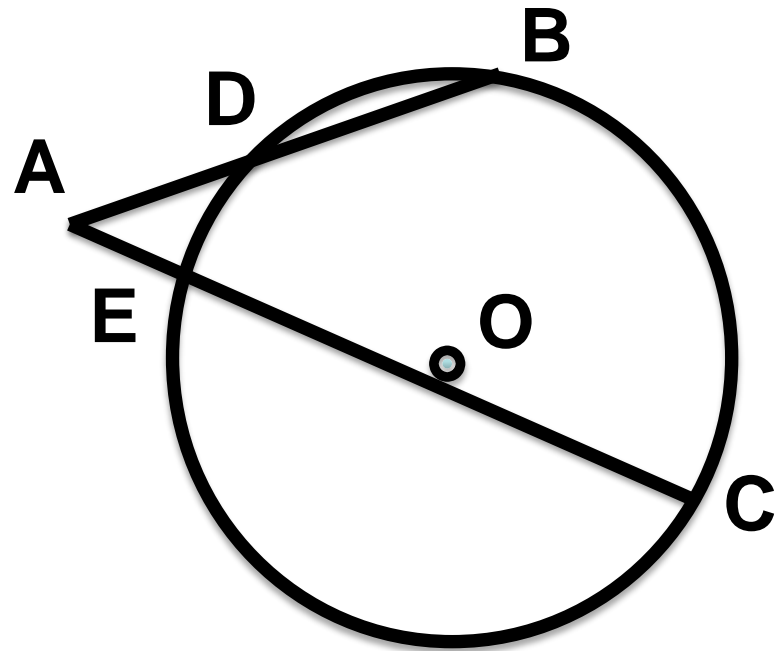


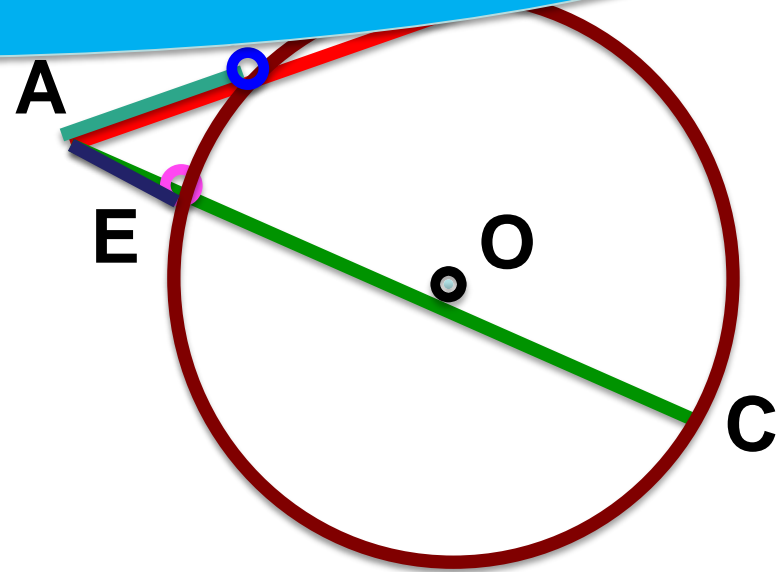
Diagram Understanding

Problem:

1. Identify visual elements
2. Ali

**Diagram understanding →
submodular optimization**

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



Experimental Study

