The View from AI2

Oren Etzioni, CEO
Allen Institute for AI (AI2)

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.*

Outline:

1. Overview of AI2 **(rapid)**
2. Observations about knowledge **(simple)**
3. Information Extraction **(visual)**
4. Reasoning in Aristo **(hard)**

AI2 Chronology and "Geography"

Time Line	
AI2 launched	Jan. 2014
Team of 30 + 12 interns	Fall 2014
Team of 50	Dec. 2015

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Summer of 2014 Interns

ALLEN INSTITUTE
for ARTIFICIAL INTELLIGENCE

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

Research Scientists

Peter Clark (leader)
UT Austin

Santosh Divvala
CMU

Tony Fader
UW

Vu Ha
University of Wisconsin

Mark Hopkins
UCLA

Kevin Humphreys
University of Edinburgh

Tushar Khot
University of Wisconsin

Jayant Krishnamurthy
CMU

Ashish Sabharwal
UW

Oyvind Tafjord
Princeton

Peter Turney
University of Toronto

Ali Farhadi (leader), UIUC

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Common Themes in AI2 Projects

- Ambitious, long-term goals
- Measurable results in 1-3 years
  - Standardized, unseen test questions
  - **“Beyond the Turing Test”**
- Open & collaborative (papers, **ADI**)
- Leveraging NLP, ML, and vision for:

Knowledge
Reasoning
Explanation

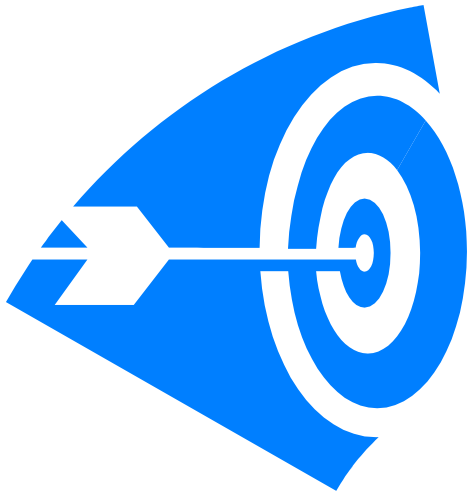

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Core Projects

66% Science (4th grade, NDMC)

Aristo

Da Vinci

AKBC over Science corpus

AKBC from Images & diagrams

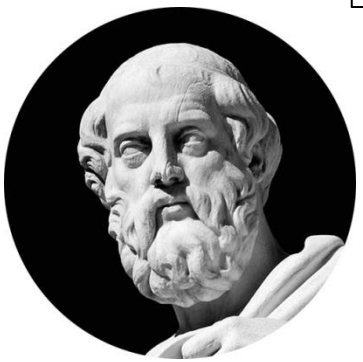

Plato

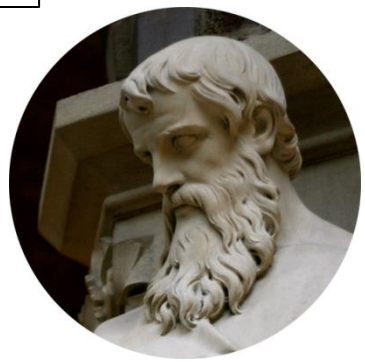

Euclid

EMNLP ’14 77.7 % arithmetic

AAAI ‘14 Geometry

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

High-level observations about knowledge & reasoning

Do we need a body to acquire intelligence?

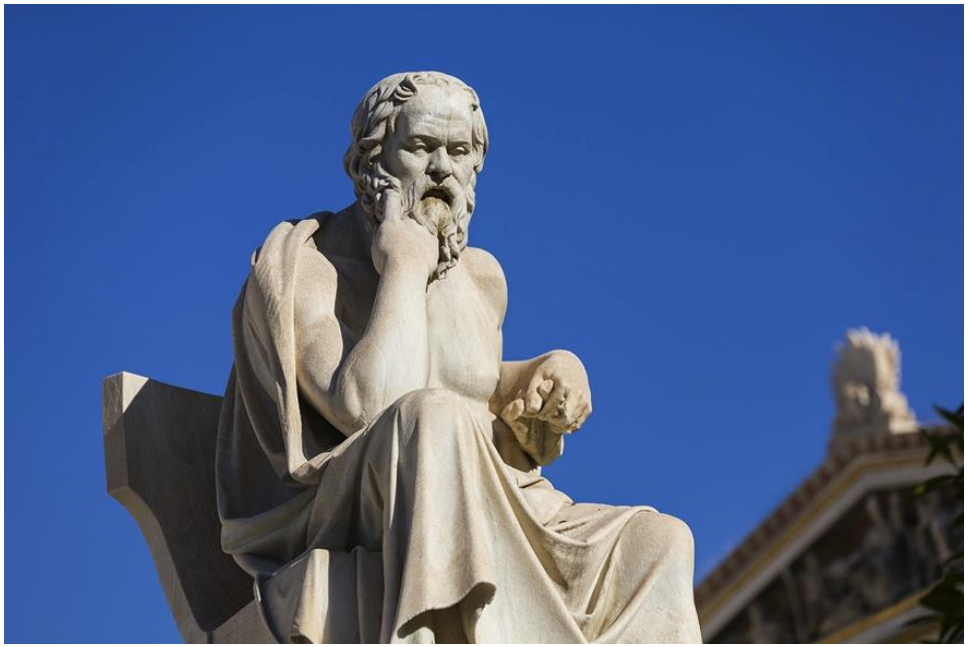

(too philosophical for us)

Do we need a body to acquire common-sense knowledge?

(a bit vague)

Do we need a body to pass the 4th grade science test?

(we can answer this one!)

Factual Knowledge for 4th Grade Science

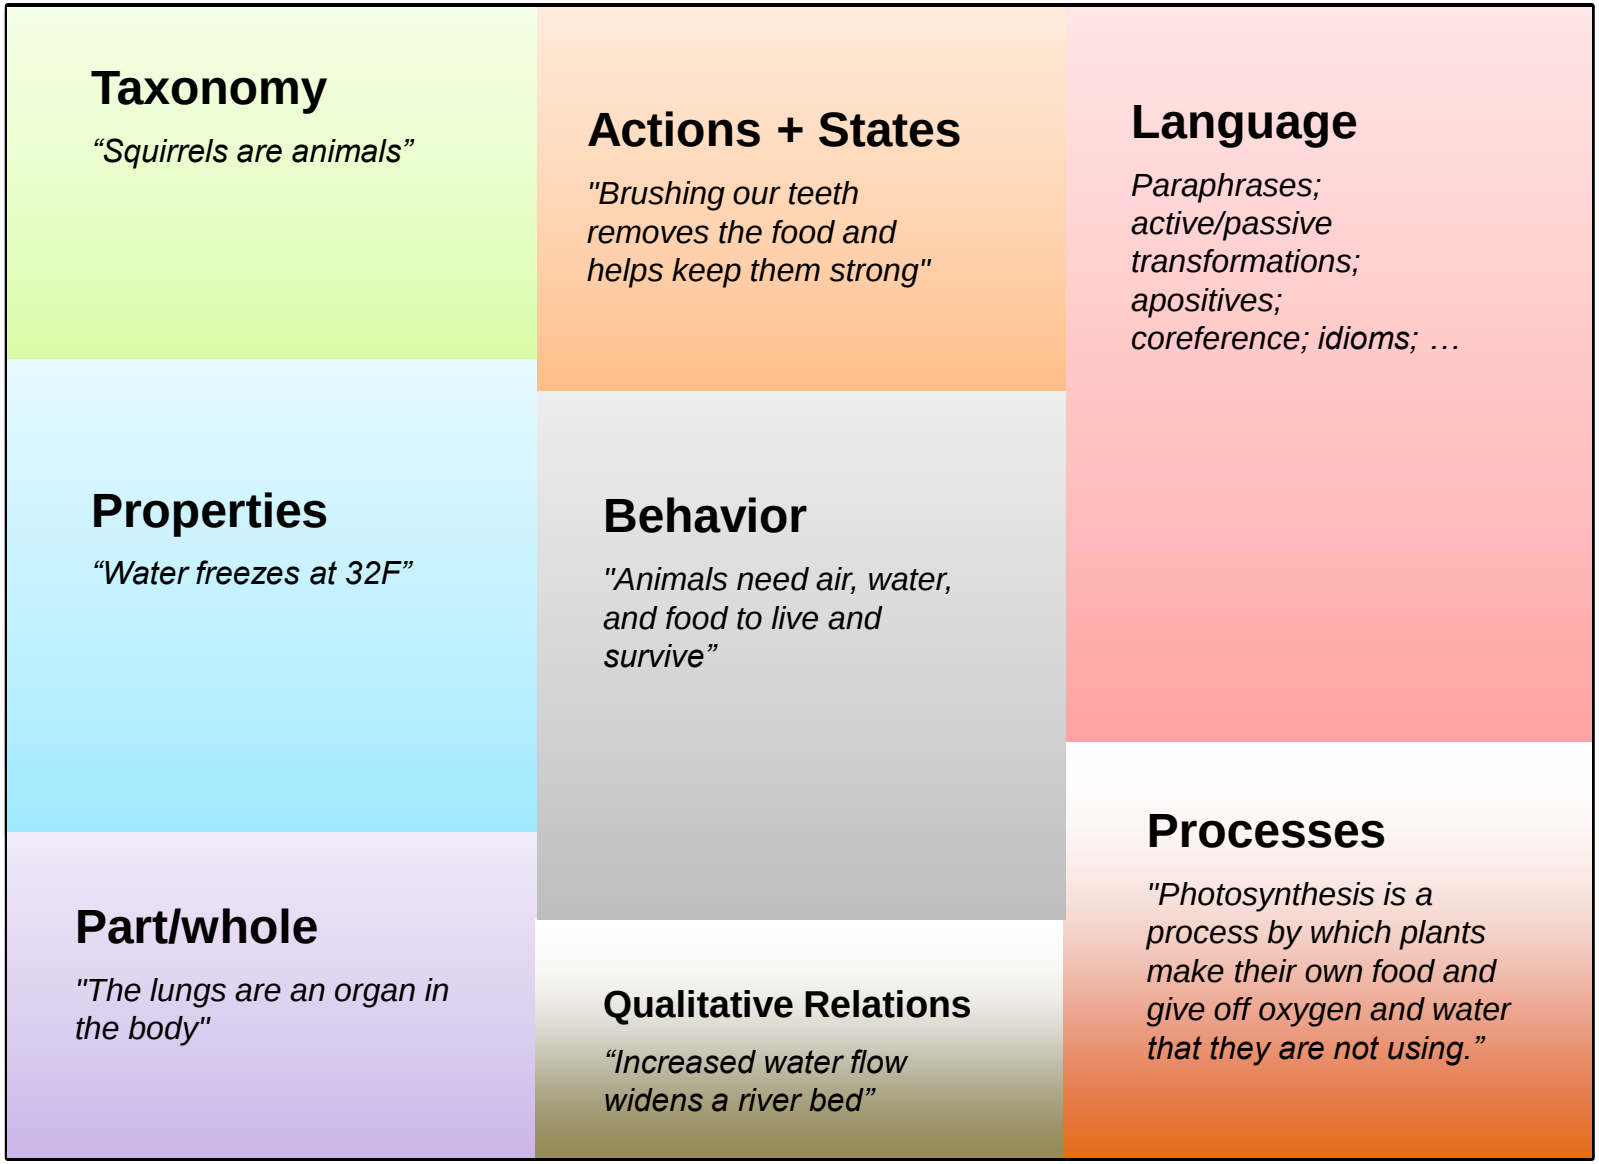

2014 Knowledge Tour

Google apple +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

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

iPhone
Discover everything iPhone, including the most advanced ...

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

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

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

More results from apple.com »

News for **apple**

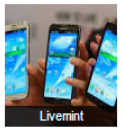

Another **Apple**-Samsung Skirmish Heads to Court
ABC News - 4 minutes ago
The fiercest rivalry in the world of smartphones is heading back to court this week in the heart of the Silicon Valley, with **Apple** and Samsung ...

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**

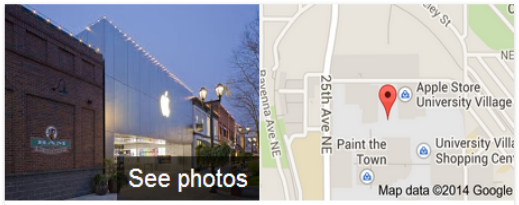

Apple Store University Village

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 ...

"apple fruit" - Google Search - Google Chrome

https://www.google.com/search?q="apple+fruit"

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...

Google "apple fruit" +Oren

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...

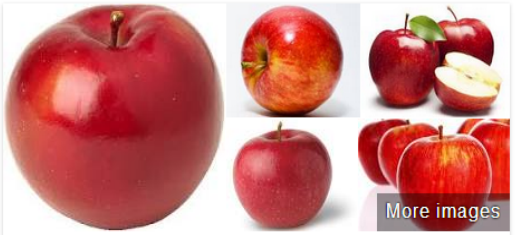

Apple

Fruit

The apple is the pomaceous fruit of the apple tree, species Malus

These KBs are fact rich but knowledge poor!

Malus domestica (family Rosaceae), one of the most widely cultivated tree fruits. The apple is a ...

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%

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

Machine Reading

Auto-Text to Knowledge

Source: DARPA, Machine Reading initiative

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

Ask NELL - Read the Web :: NELL KnowledgeBase Browser - Google Chrome

rtw.ml.cmu.edu/rtw/kbbrowser/special:asknell

Apps | 0-1 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...

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

anewell.JPG

NELL Knowledge Base

categories

- relatedto
  - visualartmoveme
  - personinacadem
  - generalizationof
    - athletessucha
    - academicfields
    - chemicaltypeh
    - astronautssuc
    - automobileeng
    - airportsuchasa
    - amphibiansuc
    - aquariumssuc
    - arteriessuchas
    - animaltypehas
      - animalsucha
        - animalsuc
        - animalsuc
      - animalsucha
      - inverseofarti
    - professiontype
    - productinstanc
    - 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

Open Information Extraction (Banko, *et al*, 2007)

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

Relations typically:

- anchored in verbs
- exhibit simple syntactic form

Virtues:

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

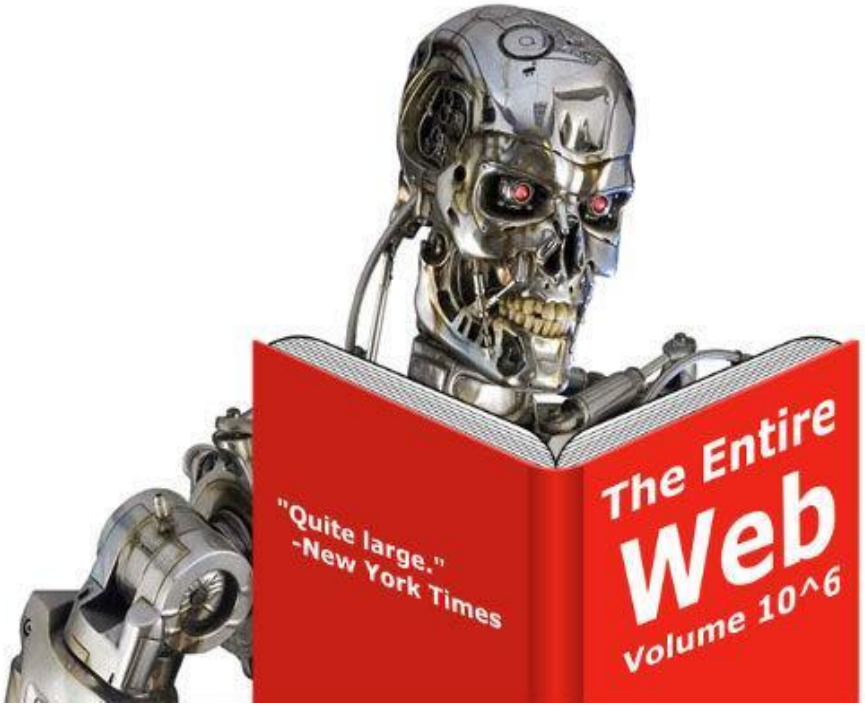

Open Information Extraction - Google Chrome

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 [illegible] its wild ancestor, Malus sieversii, is still [illegible] by

IE over Web sentences suffers from Attention Deficit Disorder!

are in bloom [illegible]

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)

Common-Sense Knowledge from Images

Which animals lay eggs?

Obtaining Visual Knowledge

1. Detect Objects (nouns)
2. Reason about Actions (verbs)

Key Challenges:

- Supervision (Bounding boxes, Spatial relations)
- Large-Scale (~10^5 objects, ~10^3 actions)

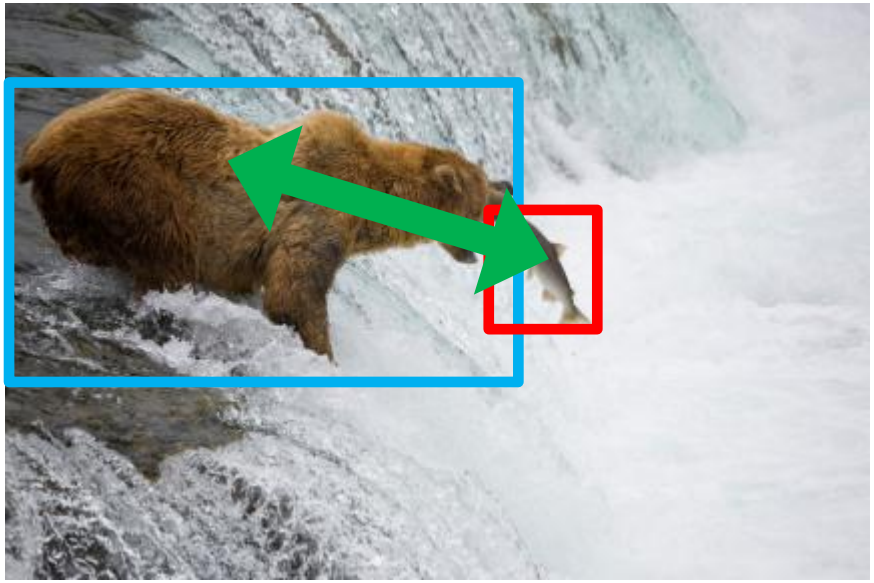

Do bears catch salmon?

VisIE: Visual Information Extraction
(Sadeghi, Divvala, Farhadi, submitted)

Do dogs eat ice cream?

dog dog eating ice cream

- ☐ OpenIE
- ☐ ConceptNet
- ☑ VisIE

Do snakes lay egg?

Snake laying eggs egg

- ☑ OpenIE
- ☐ ConceptNet
- ☑ VisIE

- **Builds object detectors based on Google images**
- **Utilizes a joint model over detectors to assess triples**
- **Mean Average Precision = 0.54**

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE

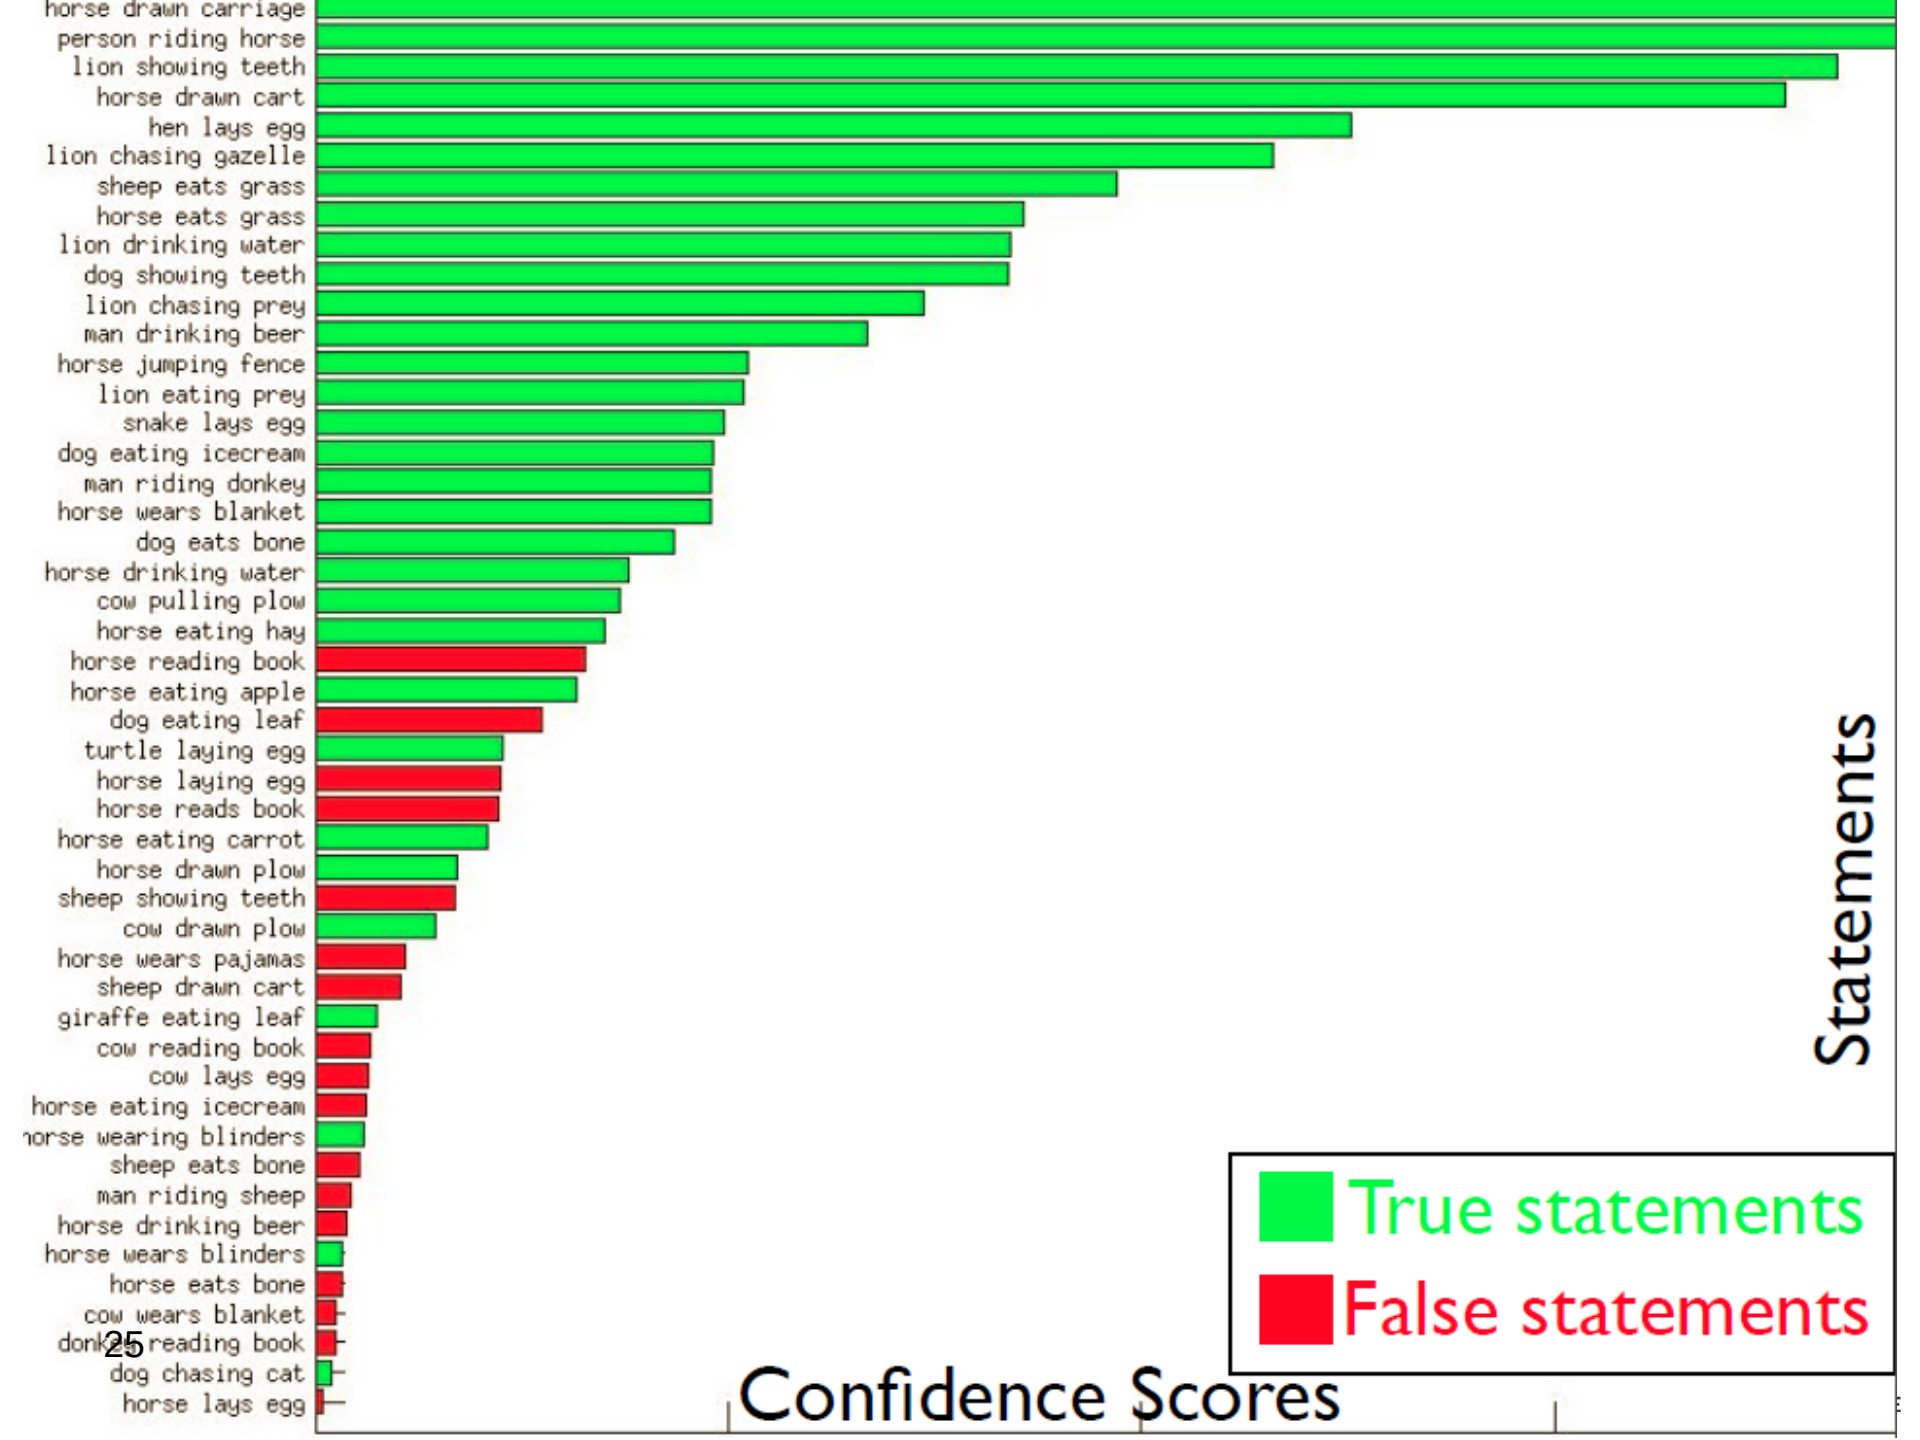
horse drawn carriage
person riding horse
lion showing teeth
horse drawn cart
hen lays egg
lion chasing gazelle
sheep eats grass
horse eats grass
lion drinking water
dog showing teeth
lion chasing prey
man drinking beer
horse jumping fence
lion eating prey
snake lays egg
dog eating icecream
man riding donkey
horse wears blanket
dog eats bone
horse drinking water
cow pulling plow
horse eating hay
horse reading book
horse eating apple
dog eating leaf
turtle laying egg
horse laying egg
horse reads book
horse eating carrot
horse drawn plow
sheep showing teeth
cow drawn plow
horse wears pajamas
sheep drawn cart
giraffe eating leaf
cow reading book
cow lays egg
horse eating icecream
horse wearing blinders
sheep eats bone
man riding sheep
horse drinking beer
horse wears blinders
horse eats bone
cow wears blanket
donkey reading book
dog chasing cat
horse lays egg
Statements
True statements
False statements
Confidence Scores

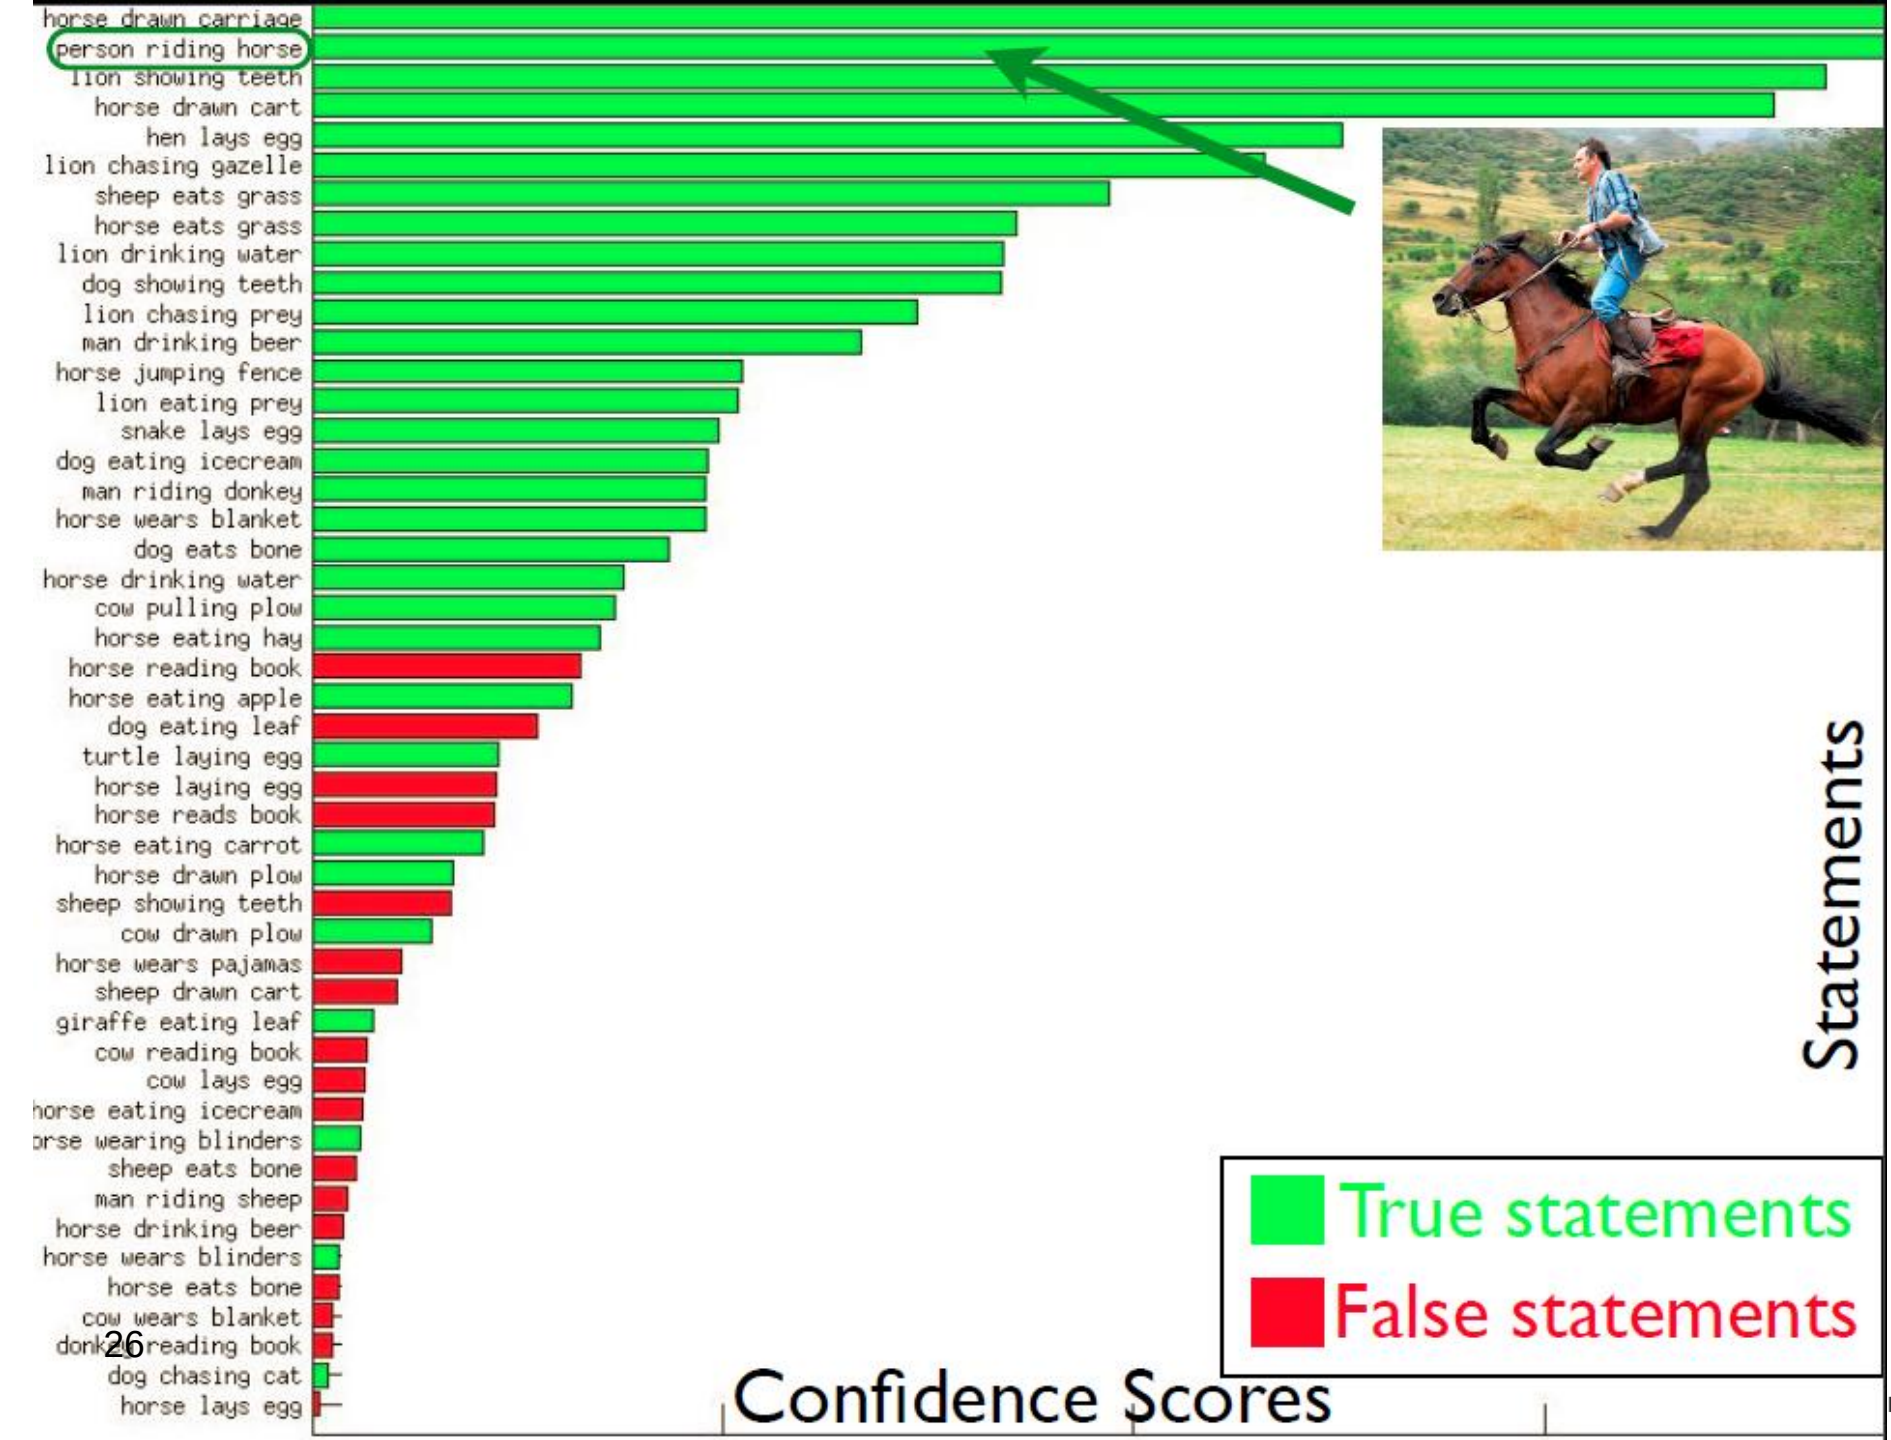
horse drawn carriage
person riding horse
lion showing teeth
horse drawn cart
hen lays egg
lion chasing gazelle
sheep eats grass
horse eats grass
lion drinking water
dog showing teeth
lion chasing prey
man drinking beer
horse jumping fence
lion eating prey
snake lays egg
dog eating icecream
man riding donkey
horse wears blanket
dog eats bone
horse drinking water
cow pulling plow
horse eating hay
horse reading book
horse eating apple
dog eating leaf
turtle laying egg
horse laying egg
horse reads book
horse eating carrot
horse drawn plow
sheep showing teeth
cow drawn plow
horse wears pajamas
sheep drawn cart
giraffe eating leaf
cow reading book
cow lays egg
horse eating icecream
horse wearing blinders
sheep eats bone
man riding sheep
horse drinking beer
horse wears blinders
horse eats bone
cow wears blanket
donkey reading book
dog chasing cat
horse lays egg
Statements
True statements
False statements
Confidence Scores

horse drawn carriage
person riding horse
lion showing teeth
horse drawn cart
hen lays egg
lion chasing gazelle
sheep eats grass
horse eats grass
lion drinking water
dog showing teeth
lion chasing prey
man drinking beer
horse jumping fence
lion eating prey
snake lays egg
dog eating icecream
man riding donkey
horse wears blanket
dog eats bone
horse drinking water
cow pulling plow
horse eating hay
horse reading book
horse eating apple
dog eating leaf
turtle laying egg
horse laying egg
horse reads book
horse eating carrot
horse drawn plow
sheep showing teeth
cow drawn plow
horse wears pajamas
sheep drawn cart
giraffe eating leaf
cow reading book
cow lays egg
horse eating icecream
horse wearing blinders
sheep eats bone
man riding sheep
horse drinking beer
horse wears blinders
horse eats bone
cow wears blanket
donkey reading book
dog chasing cat
horse lays egg

Statements

True statements

False statements

Confidence Scores

Facts are necessary, but not sufficient

A Theory also includes:

- Rules
- Reasoning
- Explanation

A Theory is Greater than the Sum of its Facts

Aristo Demo

1. General rules from Barron's Study Guide
2. Background facts stated in the question
3. Multiple Choice

Aristo Demo

Reasoning Method

- Deductive reasoning is too restrictive:
  - fall down ⊬ fall down to the ground
  - Most animals have legs ⊬ dogs have legs…
- Shallow text alignment is too permissive:
  - {turn,a,*liquid*,into,a,**solid**} ≈ {turn,a,**solid**,into,a,*liquid*}

- Probabilistic reasoning is challenging
  - **Text ➔ MLN mapping is unsolved**
  - "People breathe air."

  - **Naïve encoding of single sentence ➔ 10^10 node Markov Logic Network (MLN)**

MLN Scaling for Rules Extracted from Text

A short study guide example: *"Some animals grow thick fur in winter to stay warm."*

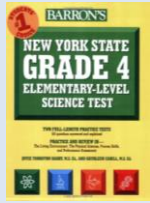

First order representation using 6 variables, 6 non-Isa predicates, 2 existentials**:**

∀ **a, g, f, w:** Isa(a, "Some animals"), Isa(g, "grow"), Isa(f, "thicker fur"), Isa(w, "the winter"),
Agent(g, a), Object(g, f), In(g, w)
⇒ ∃ **s, m:** Isa(s, "stays"), Isa(m, "warm"), **Enables(g, s), Agent(s, a), Object(s, m)**

MLN encoding k science rules

⇒ ~$(D*k)^V$ ground network rules

Domain size
- ~10
- But no symmetry or exchangeability

Variables *per rule*
- ~10 for extracted rules
- ~3 in typical hand-coded rul

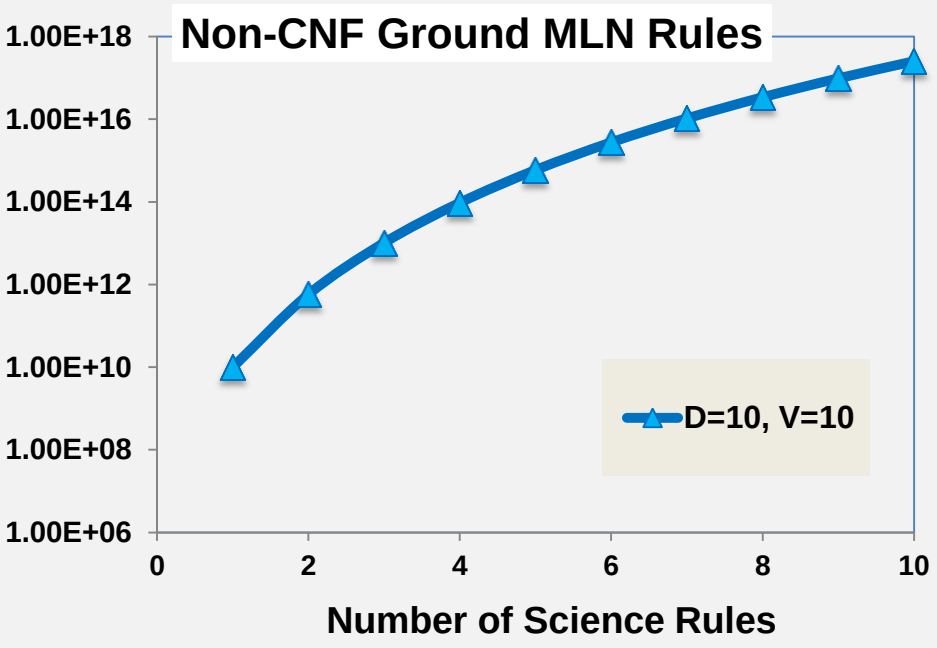

Enhancements for Tractability

1. Add semantic constraints
   - E.g., Cause(x,y) => Effect(y,x), events have unique agents, …

2. Use hard constraints to simplify & reduce soft constraints
   - SAT solver for unit propagation + backbone/fixed variable detection

3. Use refined types to reduce domain size
   - Consider only lexically similar entities/events

4. Use constants in place of first-order variables, where possible

Still slow and inaccurate!

- 3 min per question (with **just 1** extracted rule)
- 47% accuracy (4-way multiple choice)

Motivation for New Approach

- Can treat all mentioned entities/events as constants

- Inference requires “fuzzy” matching between extracted terms

thicker fur ≈ thicker fur in winter ≈ heavier coat

We formulate matching as a probabilistic inference

Probabilistic Alignment over graphs

Treat extracted rules as graphs

- vertices = entities/events;
- edges = relations; partitioned into antecedent/consequent

- Sibling inference tasks:

AlignmentMLN + **InferenceMLN**

Lexical Reasoning

Structured alignment beyond BOW:
word similarity + *graph structure*

Directional Inference with extracted rules

- Multi-path version of reasoning in "the demo"
- Directionality: thick fur => warm, but warm ≠> thick fur

ProbAligner Method: Inference (work in progress)

- Example Question: ***Is it true that a decomposer is an organism that recycles nutrients?***
- Example Rules (**antecedent => consequent**) :
  1. Decomposers are living things that break down and recycle
  2. Decomposers are living things that recycle their[consumers] nutrients into the soil

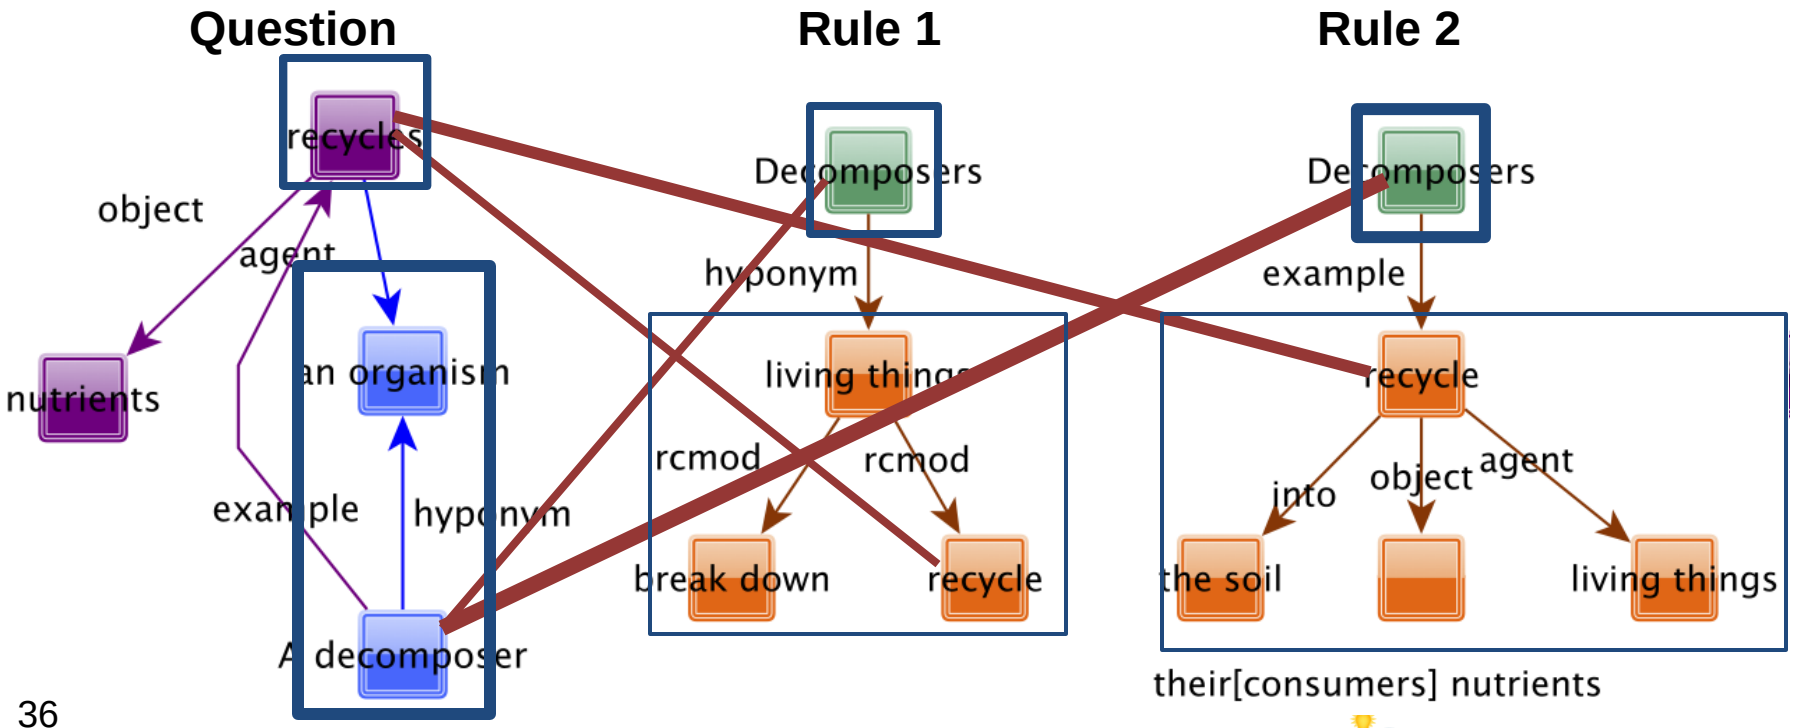

ProbAligner Results (work in progress)

- Faster
  - Few variables per rule (independent of extracted rule length)
  - No existentially quantified variables

  => Better scaling

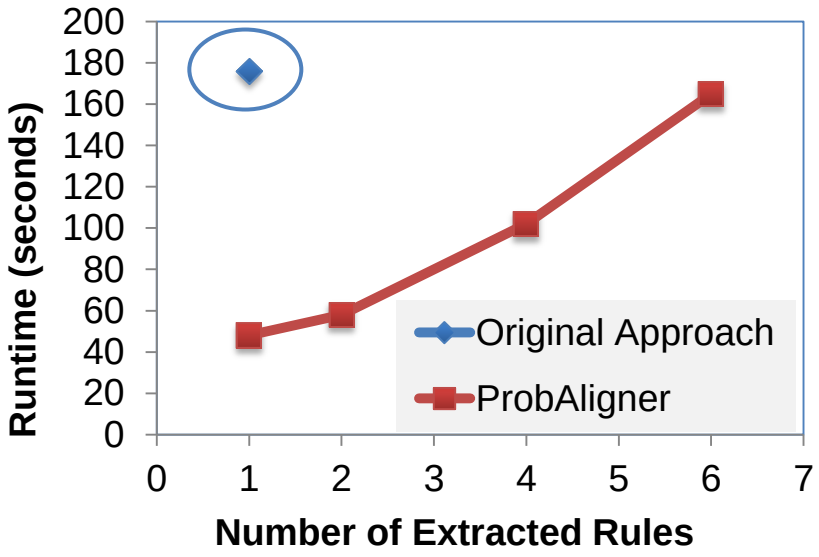

- More robust

Conclusion

AI2 is one year old

We are hard at work on:

- Sophisticated IE (rules, processes)
- Probabilistic reasoning over extracted rules
- Question understanding

We utilize standardized tests to assess progress

- Early results on Arithmetic & Geometry (EMNLP & AAAI)

Data and publications are here: ***www.allenai.org***

Join Us!

ALLEN INSTITUTE *for* ARTIFICIAL INTELLIGENCE