
Learning from the World



Mihai Surdeanu

msurdeanu@email.arizona.edu

Natural Language Processing (NLP):
computer process that extracts
meaningful information from natural
language input.

NLP ≠ Neuro-linguistic Programming





NLP with machine learning =
computers learn a task based
on empirical data

Supervised Learning



- **Systems duplicate correct analyses from training data**
- **Why?**
 - Predictive accuracy by exploiting human task supervision
- **Why not?**
 - Time-consuming
 - Experts are expensive
 - Training data needs updating

Unsupervised Learning



- **Systems take raw data and automatically detect patterns**
- Why?
 - Zero cost for data annotation
 - Ability to train on big data
- Why not?
 - It's really hard to get something for nothing
 - It's an interesting research track, but the results are not great

Information from the World



The very fact that problems are of interest means that people have seen a need to produce structured data

Information from the World



Barack Hussein Obama II (ⓘ[ⓘ]/bəˈrɑːk huːˈseɪn ʊsˈbɑːmə/; born August 4, 1961) is the 44th and current President of the United States. He is the first African American to hold the office. Obama previously served as a United States Senator from Illinois from January 2005 until he resigned following his victory in the 2008 presidential election.

Born in Honolulu, Hawaii, Obama is a graduate of Columbia University Law School, where he was the president of the *Harvard Law Review* and a community organizer in Chicago before earning his law degree. He worked as a rights attorney in Chicago and taught constitutional law at the University of Chicago Law School from 1992 to 2004. He served three terms representing the state in the Illinois Senate from 1997 to 2004.

Following an unsuccessful bid against the Democratic incumbent for a seat in the United States House of Representatives in 2000, Obama ran for United States Senate in 2004. Several events brought him to national attention during the campaign, including his victory in the March 2004 Illinois Democratic primary for Senate election and his keynote address at the Democratic National Convention in July 2004. He won election to the U.S. Senate in Illinois in November 2004. His presidential campaign began in February 2007, and after a close campaign in the 2008 Democratic Party presidential primaries against Hillary Rodham Clinton, he won his party's nomination. In the 2008 presidential election, he defeated Republican nominee John McCain, and was inaugurated as president on January 20, 2009. In October 2009, Obama was named the 2009 Nobel Peace Prize laureate.

As president, Obama signed economic stimulus legislation in the form of the American Recovery and Reinvestment Act in February 2009 and the Tax Relief, Unemployment Insurance Reauthorization, and Job Creation Act in December 2010. Other domestic policy initiatives include the Patient Protection and Affordable Care Act, the Dodd–Frank Wall Street Reform and Consumer Protection Act, the Don't Ask, Don't Tell Repeal Act and the Budget Control Act of 2011. In foreign policy, he has increased troop levels in Afghanistan and withdrawn troops from Iraq.

Team Stat Leaders

Category	Player	Sea	SF	Yards
Passing Yards	Tarvaris Jackson	Sea		197 yds
	Alex Smith	SF		124 yds
Rushing Yards	Marshawn Lynch	Sea		33 yds
	Frank Gore	SF		59 yds
Receiving Yards	Doug Baldwin	Sea		83 yds
	Vernon Davis	SF		47 yds

Scoreboard

Week 1	Final	Box Score
Thursday, Sep 8		
New Orleans	34	
Green Bay	42	
Sunday, Sep 11		
Philadelphia	31	
St. Louis	13	
Buffalo	41	
Kansas City	7	

Information from the World

Barack Hussein Obama II (b) /bəˈrɑːk huːˈseɪn ʊsˈbɑːmə/; born August 4, 1961; the 44th and current President of the United States; to hold the office. Obama

Show me another »

Resolved Question

Who do you think John Galt really is?

I read the book, and now I wonder what the interweb truly thinks of Jonh Galt. I would love to have a description of who this man "Idea" truly means to you.

6 months ago

Report Abuse

Best Answer - Chosen by Voters

John Galt is the imaginary dream of social Darwinists who really do wish that they could just "shrug" off the rest of society.

Unfortunately, those ideals would never work in real life, mostly because the people who would be able to "shrug" would be too busy infighting and squabbling for power to actually get anything accomplished.

6 months ago

Report Abuse

100% 1 Vote

1 thumbs up 0 thumbs down

Sea	197 yds	
SF	124 yds	
Sea	33 yds	
SF	59 yds	
Sea	83 yds	
SF	47 yds	
sunday, Sep 8		
reans	34	Final Box Score
i Bay	42	
Sunday, Sep 11		
Philadelphia	31	Final Box Score
St. Louis	13	
Buffalo	41	Final Box Score
Kansas City	7	

Information from the World

Barack Hussein Obama II (a) /bəˈrɑːk huːˈseɪn ʊsˈbɑːmə/; born August 4, 1961) is the 44th and current President of the United States. He was the first African American to hold the office. Obama

Resolved Question

Show me another »

Who do you think John Galt really is?

- Why?

- Big data + low cost + human supervision

- Why not?

- Messy, low quality data
- Data may not model exactly what you want

100% 1 Vote

1  0 

City	Final	Box Score
Philadelphia	31	
St. Louis	13	
Buffalo	41	
Kansas City	7	

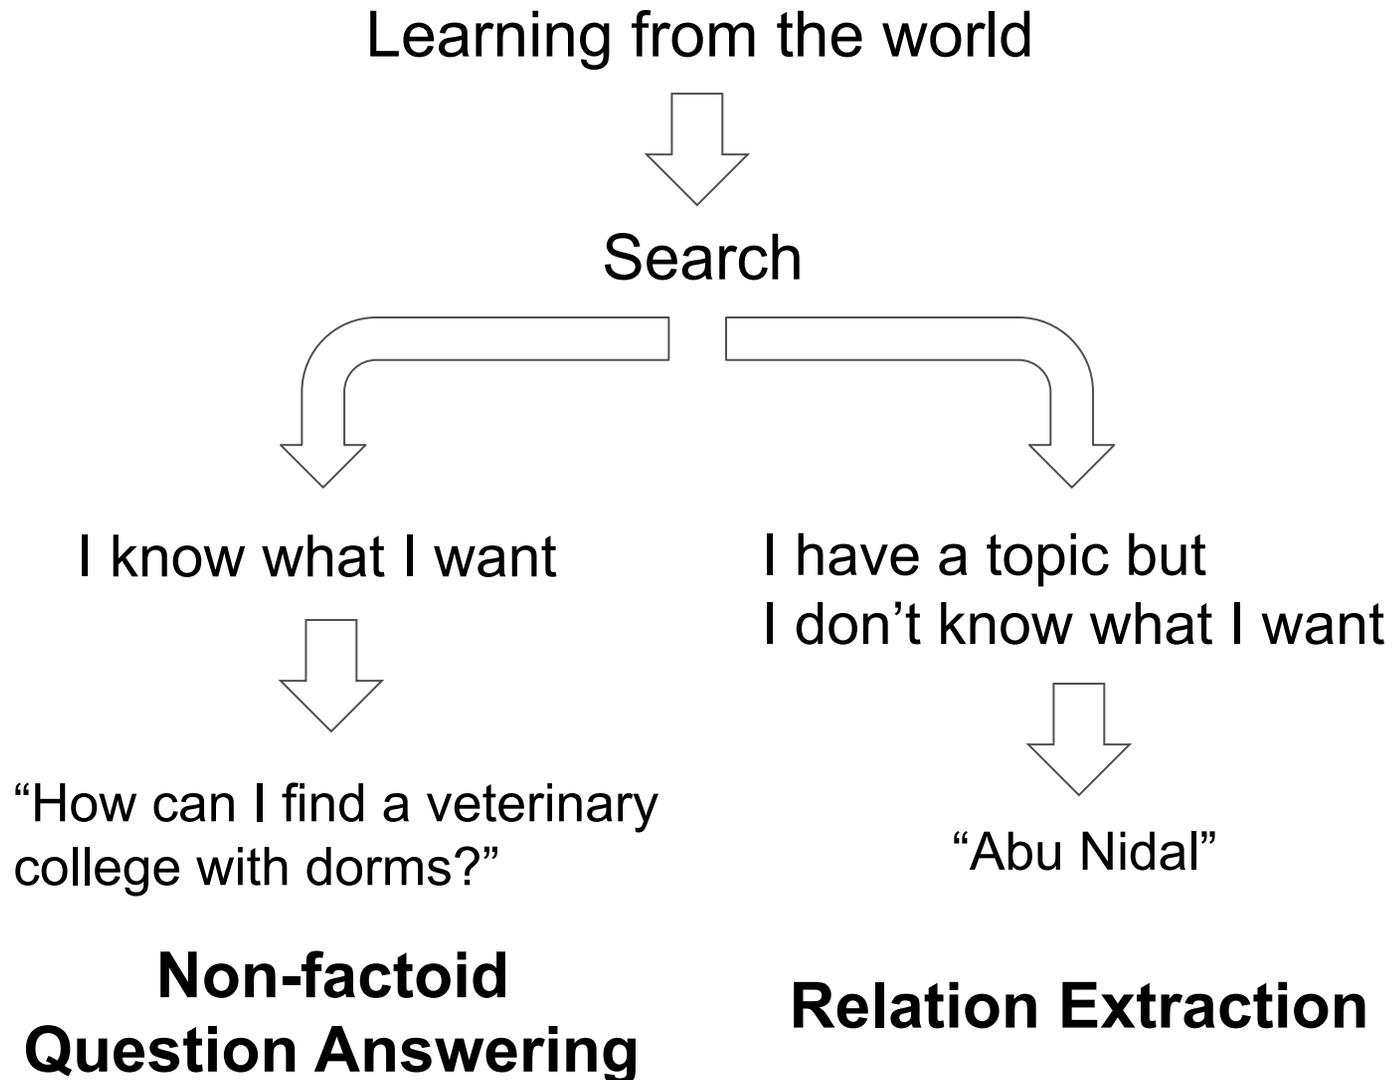
City	Final	Box Score
Sea SF	197 yds	124 yds
Sea SF	33 yds	59 yds
Sea SF	83 yds	47 yds

Information from the World



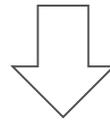
- Learning from the world outside of NLP = data science
 - Political analysis
 - Business intelligence
 - Computational advertising
 - ...

Roadmap



Roadmap

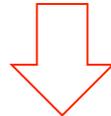
Learning from the world



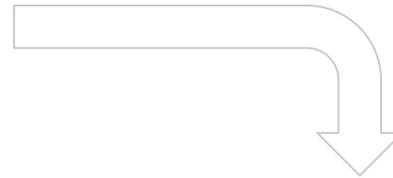
Search



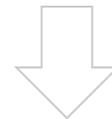
I know what I want



**Non-factoid
Question Answering**



I have a topic but
I don't know what I want



Relation Extraction

What is QA?

QA: Task of answering *natural language questions* with small fragments of *text*.

- Factoid QA: answer is a noun phrase
 - Lot of research and progress, e.g., Watson, Google

what is barack obama's date of birth

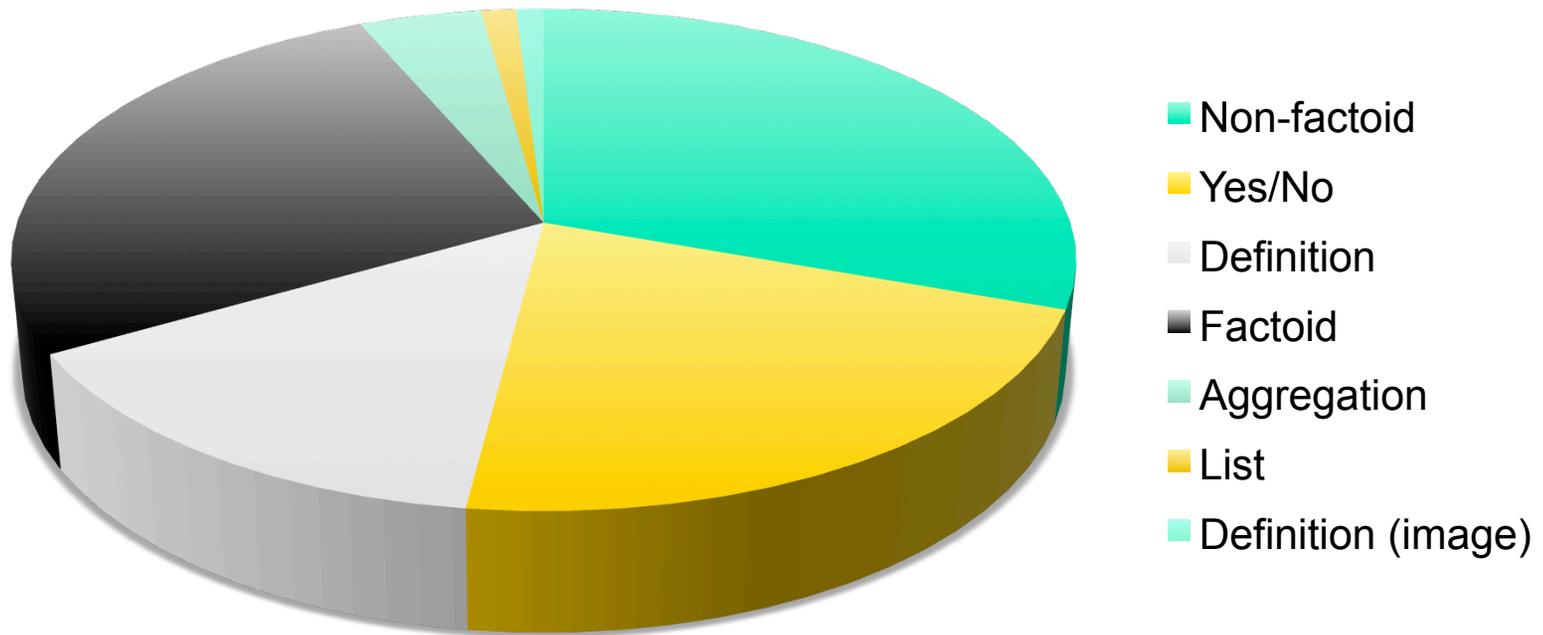


▶ Best guess for Barack Obama Date of birth is **August 4, 1961**

Mentioned on at least 6 websites including wikipedia.org, imdb.com and biography.com

- Non-factoid QA: answer is short text, possibly aggregated from multiple sources
 - Very little research to date
 - More common than factoid QA
-

Distribution of Questions for Textbook Search



Application of Non-Factoid QA: Snippet Extraction



how can i find a veterinary college with dorms? 

About 6,060,000 results (0.40 seconds)

► [Housing : Cummings School of Veterinary Medicine at Tufts University](#)

www.tufts.edu/vet/studentlife/housing.html 

Information on student **housing** near the Grafton Campus.

[Rules & Regulations](#)

www.sgu.edu/accepted-students/housing-rules-regulations.html 

St. George's University is a Caribbean Medical & **Veterinary School** with extensive ...
Students are eligible to seek off-campus **housing** if they are married with ...

[Visit a Few Vet Tech Schools | Veterinary Technician](#)

www.veterinarytechnician.com/visit-a-few-vet-tech-schools/ 

Anyone considering pursuing a career as a **veterinary** technician should consider ... Are
their **school**-sponsored/approved off campus **housing** opportunities? ...

[Are there any vet assistant colleges in pa that have dorms ...](#)

[answers.yahoo.com > ... > Higher Education \(University +\)](#) 

1 answer - Oct 5, 2006

Top answer: You do NOT need to go to **college** to be a **vet** assistant. Any **schools**
offering it are offering a useless degree - **vet** assistants get on the job training, and start
out ...

None of these
are good answers.
But they should be.

User-Generated Data Is Available



Yahoo! Answers has more than 360K answered “how to” questions

Q How do you quiet a squeaky door?

A Spray WD-40 directly onto the hinges of the door. Open and close the door several times. Remove hinges if the door still squeaks. Remove any rust, dirt or loose paint. Apply WD-40 to removed hinges. Put the hinges back, open and close door several times again.

Q How does a helicopter fly?

A A helicopter gets its power from rotors or blades. So as the rotors turn, air flows more quickly over the tops of the blades than it does below. This creates enough lift for flight.

But It Is Noisy

Simple ML

- Correctness is subjective
 - Q: How to extract html tags from an html documents with c++?
 - A: very carefully
- Grammar and spelling are optional
 - Q: Am I really that cute? girl said that I am 10/10 on heere? :/?
 - A: boy u is so hot lol

Robust features

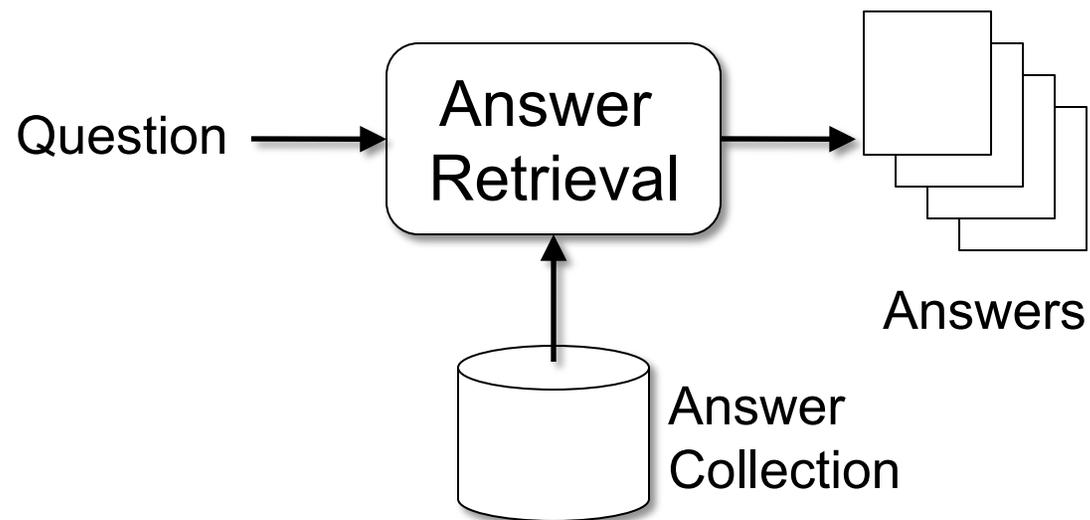
Objective



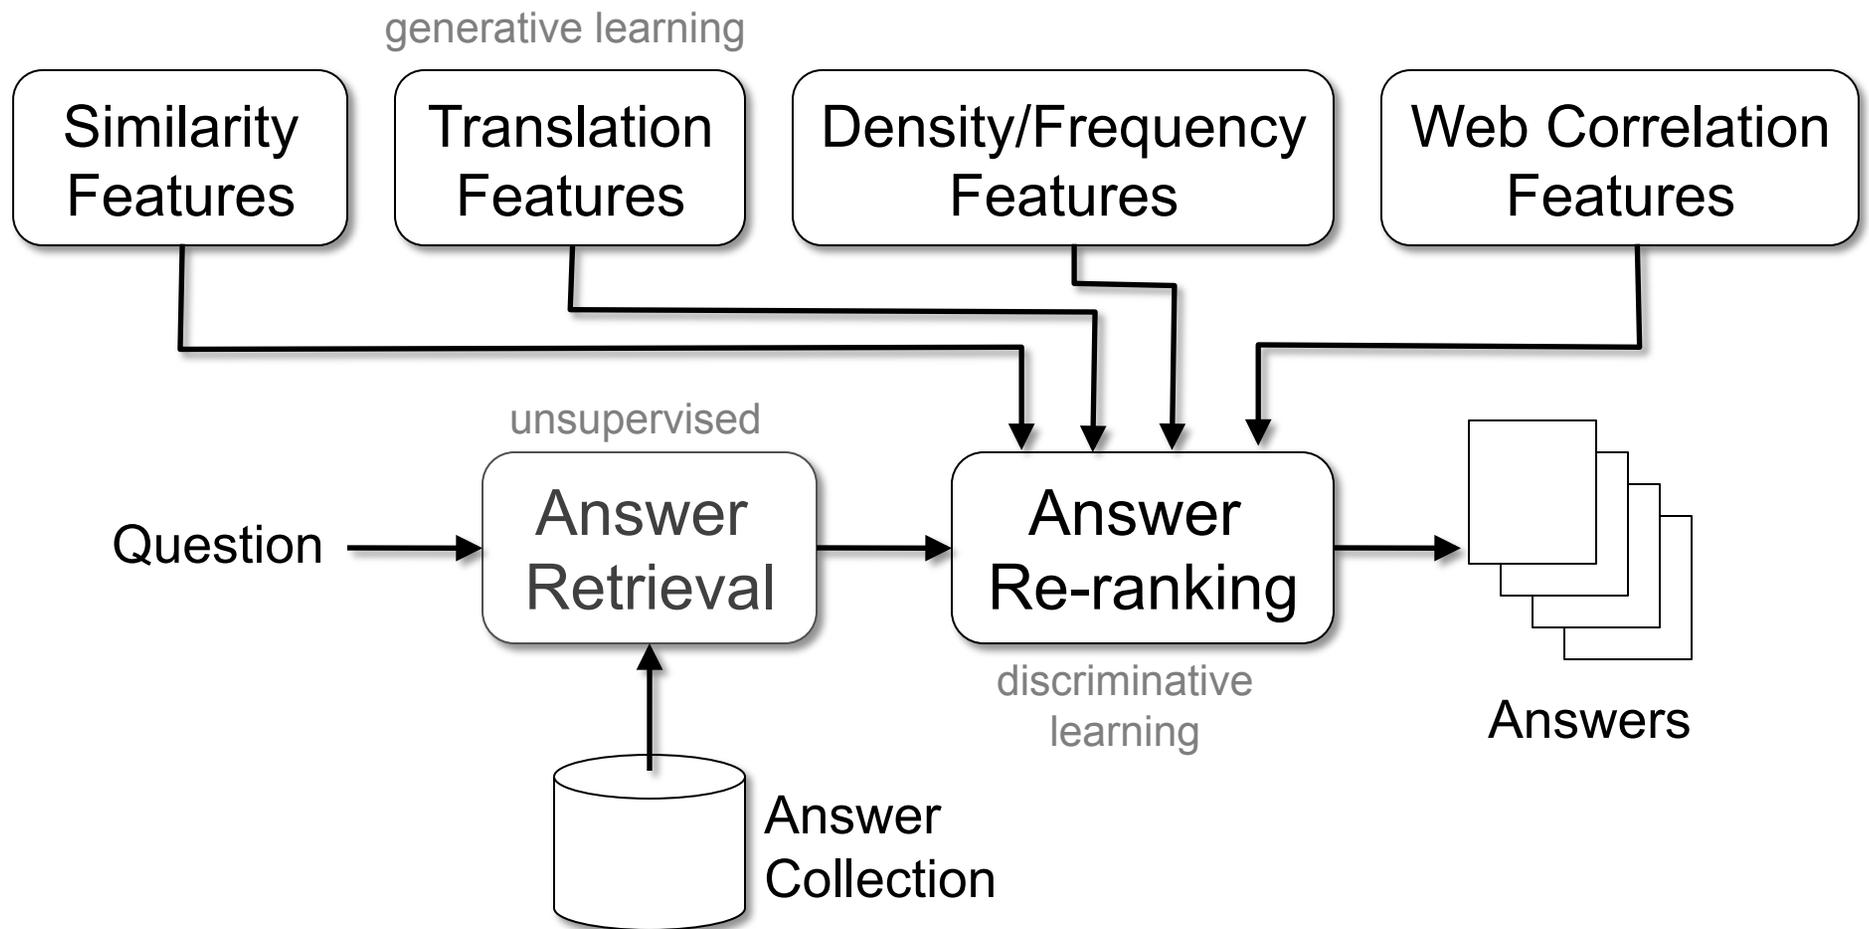
1. Is it possible to learn an answer ranking model for non-factoid questions, in a completely automated manner, using data available in online social QA sites?
 2. Which features and models are more useful in this context, i.e., ample but noisy data?
-

APPROACH

Architecture



Architecture



FEATURES AND MODEL

Features (1/4)

- FG1: similarity features
 - *Intuition: a better answer will reuse words from the question*

Question:

How to
make a
bonsai
plant?

Answer:

to **make** a **bonsai** takes many many years of clipping, wiring, bending and shaping. many online places sell **bonsai** kits but all you will get is dirt, a pot and a couple of seeds.

Features (2/4)

- FG2: translation features
 - *Intuition: A is a new language translated from Q. A better answer will be a more likely translation of the question.*

Question:

How can
I cook
grouse?

Answer:

I've never cooked a grouse, but poultry is poultry... You could **salt** it and **pepper** it, put some **flour** on a **plate** and roll it in the **flour** to coat it lightly, then heat a few **tablespoons** of **olive oil** in a **skillet** and **pan-fry** it.

Features (3/4)

- FG3: density/frequency features
 - *Intuition: a better answer will reuse patterns from the question*

Question:

How **to**
make a
bonsai
plant?

Answer:

to make a bonsai takes many
many years of clipping, wiring,
bending and shaping...

Features (4/4)

- FG4: web correlation features
 - *Intuition: words from a good answer will appear jointly with question words in other web pages or web queries*

Question:

how did I do
for make a
call and that
the other dont
see my
number?

Answer:

to make a call so that the other
person cant see the number...
dial *67 and wait for the three
beeps.. then **dial** the number

Google hits for “number dial” = 156M

Google hits for “call dial” = 137M

All these features can be computed for
different representations of the content

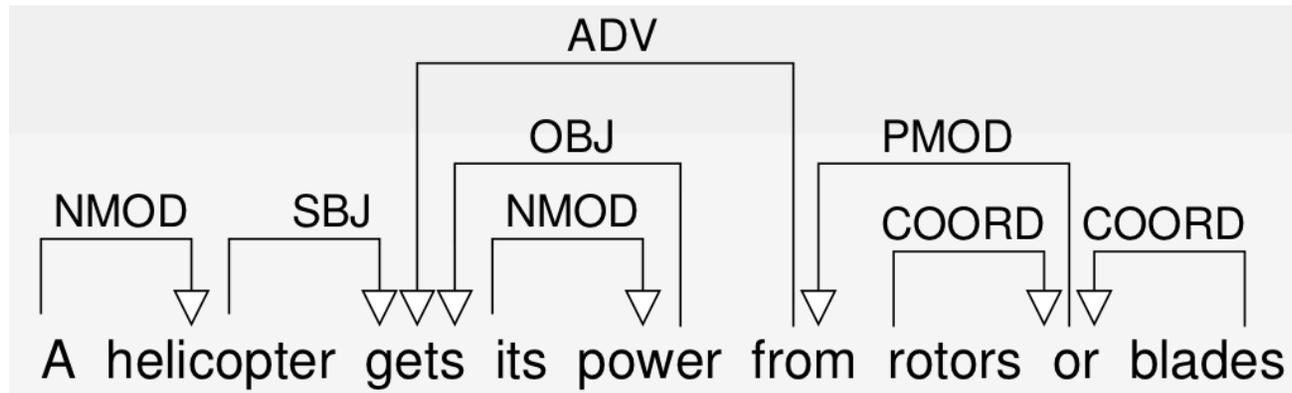
Surface Lexical Representations



“A helicopter gets its power from its rotor or blades.”

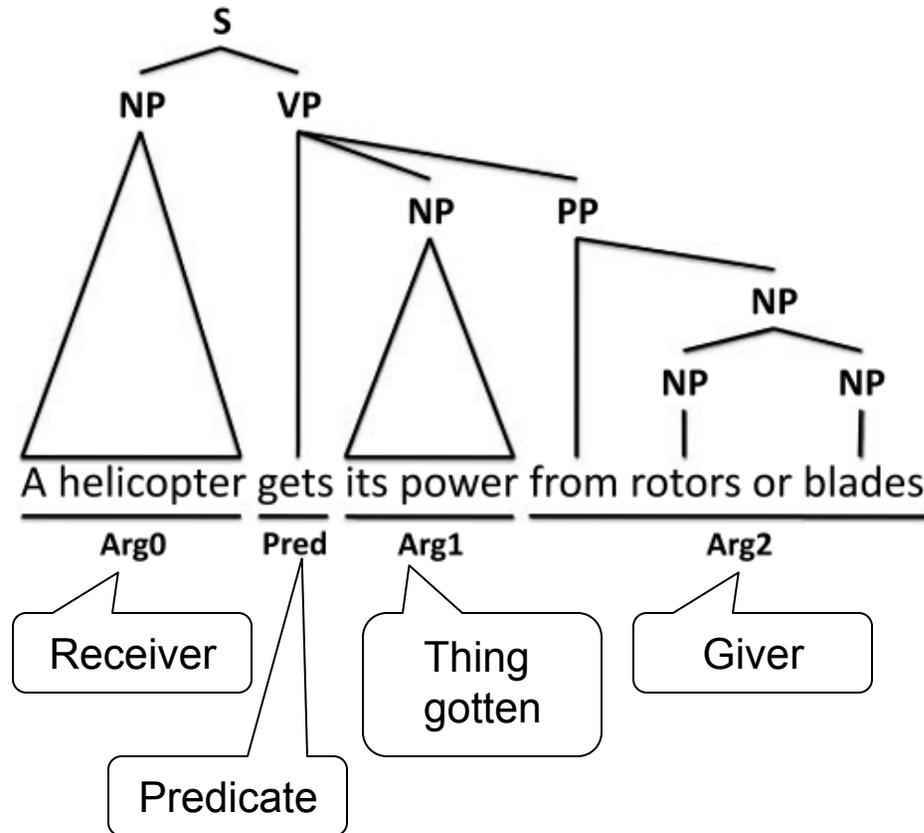
- Bag of words (W)
 - *A, helicopter, gets, its, power, from, its, rotor, or, blades*
- Bag of n -grams (N)
 - *A-helicopter, helicopter-gets, gets-its, its-power, power-from, from-its, its-rotor, rotor-or, or-blades*

Syntactic Representation



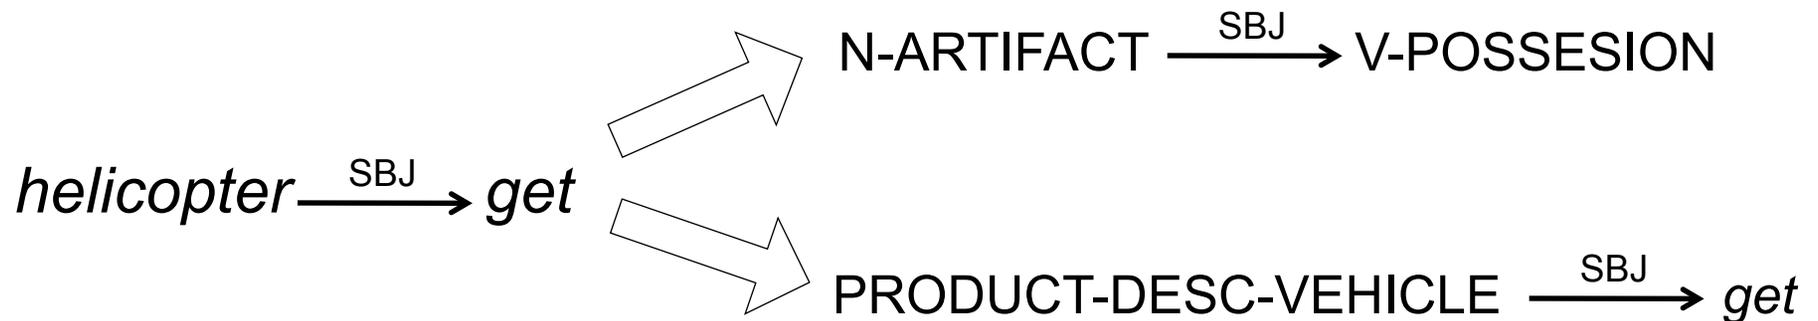
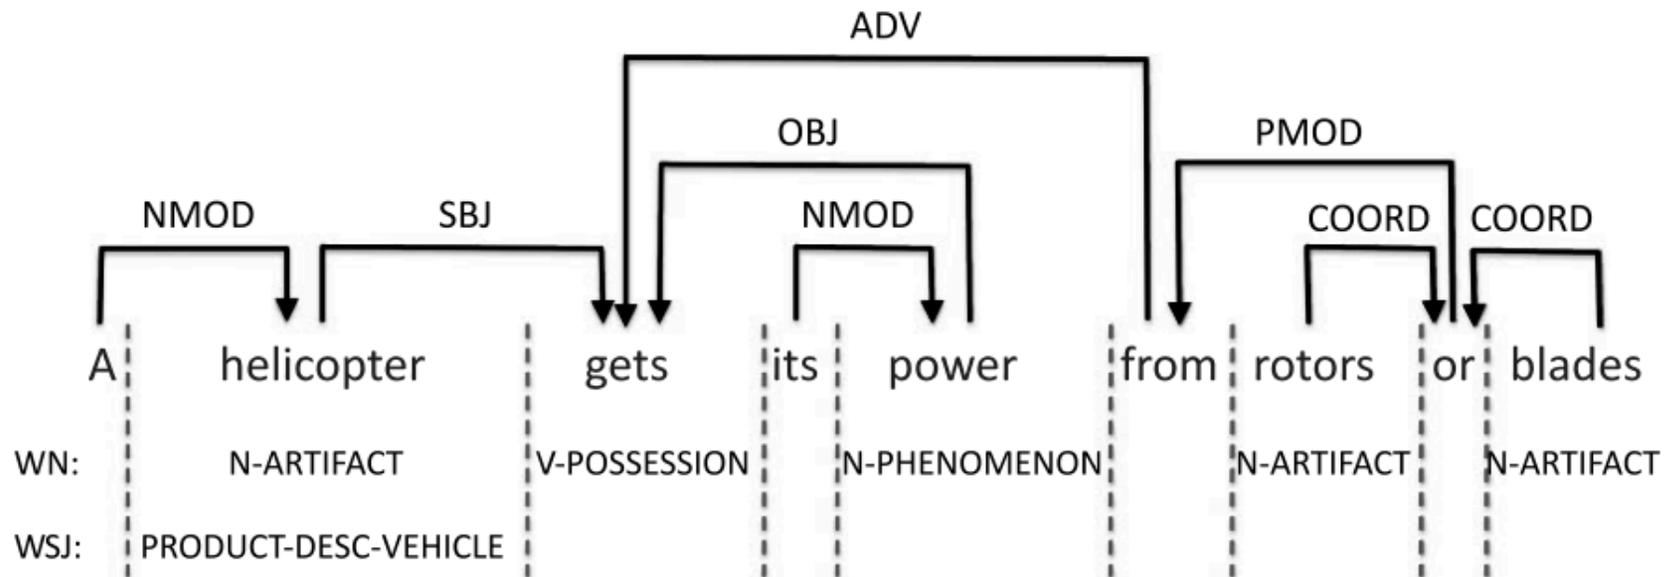
- Bag of syntactic dependencies (D)
 - $A \xrightarrow{\text{NMOD}} \textit{helicopter}$
 - $\textit{helicopter} \xrightarrow{\text{SBJ}} \textit{get}$
 - $\textit{power} \xrightarrow{\text{OBJ}} \textit{get}$
 - $\textit{from} \xrightarrow{\text{ADV}} \textit{get}$
 - $\textit{its} \xrightarrow{\text{NMOD}} \textit{power}$
 - $\textit{rotor} \xrightarrow{\text{COORD}} \textit{or}$
 - $\textit{blade} \xrightarrow{\text{COORD}} \textit{or}$
 - $\textit{or} \xrightarrow{\text{PMOD}} \textit{from}$

Semantic Representation



- Bag of semantic dependencies (R)
 - $get \xrightarrow{\text{Arg0}} \textit{helicopter}$
 - $get \xrightarrow{\text{Arg1}} \textit{power}$
 - $get \xrightarrow{\text{Arg2}} \textit{from-rotors}$

Degree of Lexicalization



16 different representations of content

All generated using open-source tools,
e.g., *<http://www.surdeanu.name/mihai/swirl>* for SRL

Model

- Linear models
 - score of an answer = $w_1 \times f_1 + \dots + w_n \times f_n$
- Re-ranking large-margin online learner (in house)
 - Updates model when an incorrect answer has a higher score than the correct one
 - For every question, it inspects all pairs of correct answer and one incorrect answer
 - Large margin: static margin parameter (tuned in development)
- SVM-rank – structural SVM tailored for ranking problems

EXPERIMENTS

Corpus

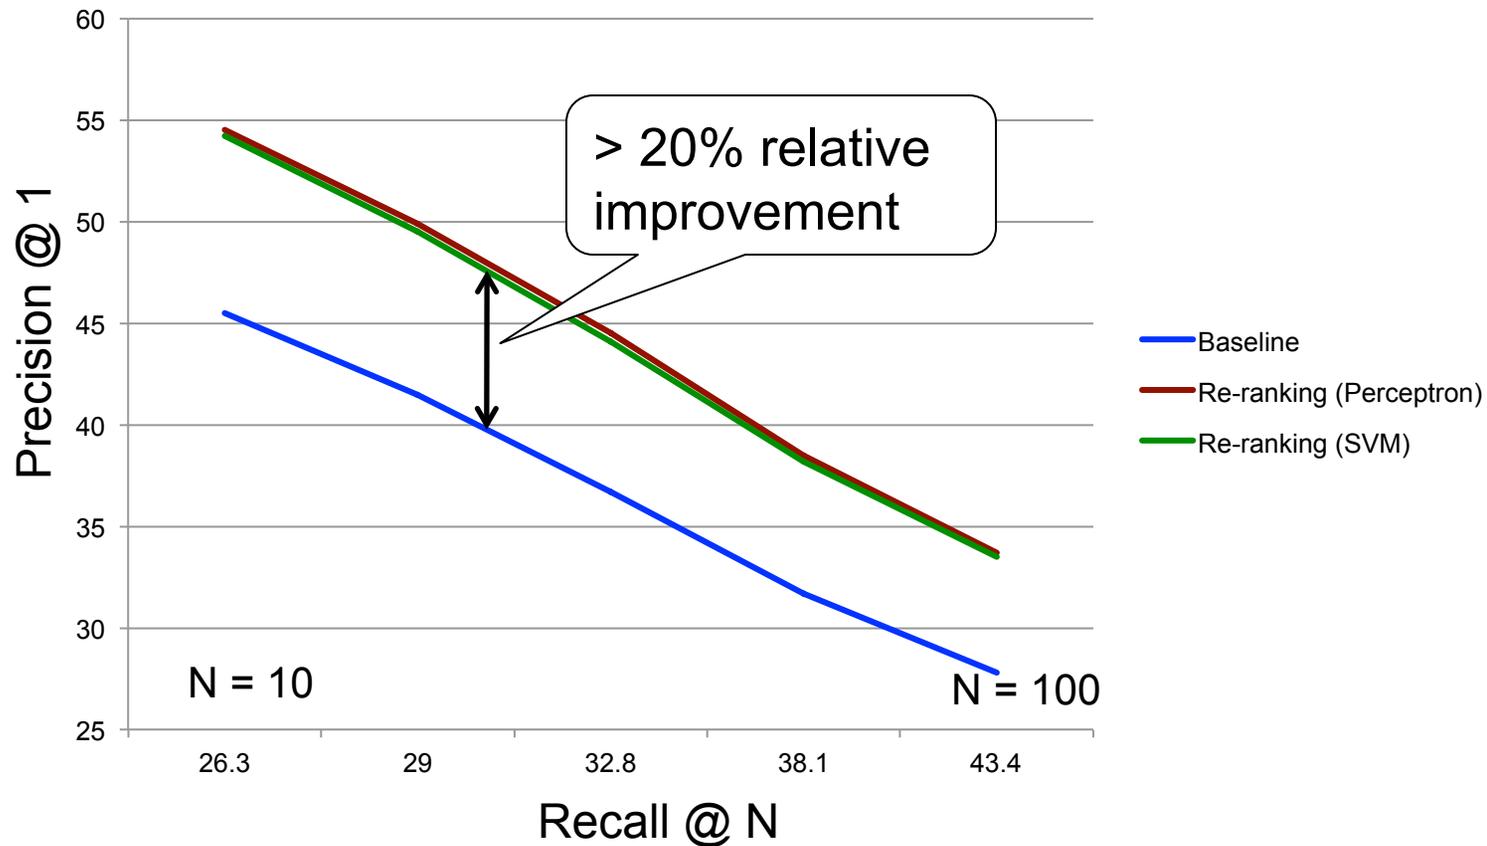


- Subset of “how to” questions from Yahoo! Answers
 - Clean questions, with a chosen/voted best answer
 - “Yahoo! Answers Manner Questions, version 1.0”. Freely available through Yahoo! Webscope.
 - 142,627 (Q, best A) pairs
 - We index all answers in this set as the answer collection
 - Partitioning of questions: 60% training, 20% development, 20% testing
-

Evaluation Measures

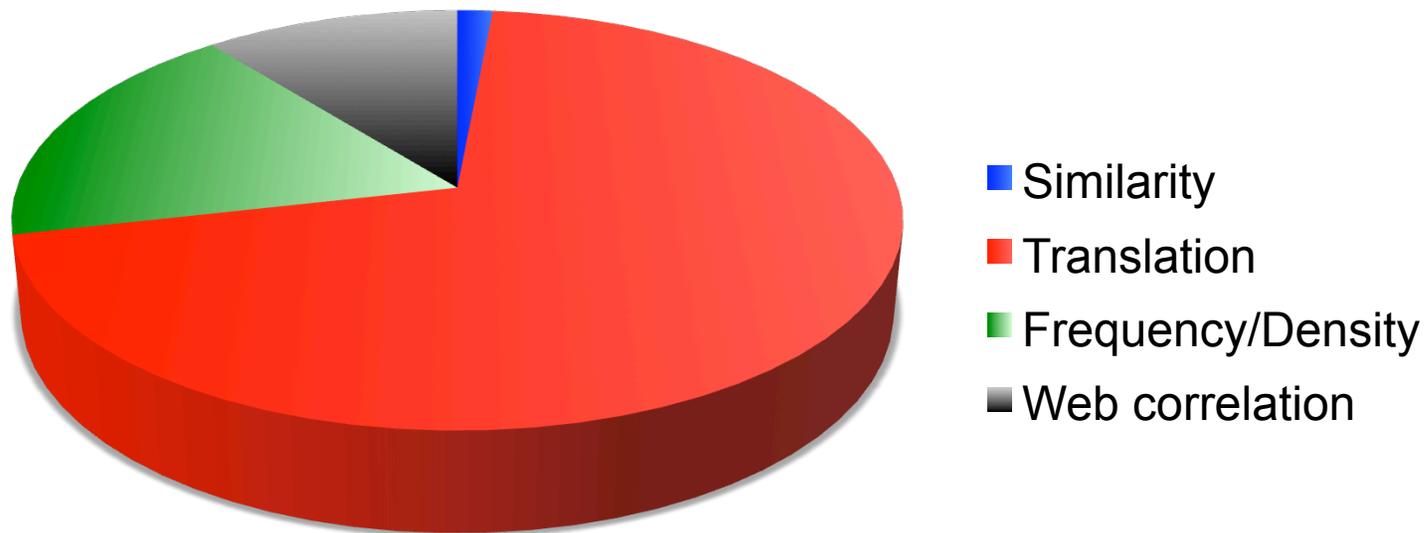
- Notations:
 - N – how many answer candidates we retrieve for each question
 - Q – all queries in the collection
 - Q^N – subset of queries for which the candidate answer pool contains the correct answer
- Recall @ $N = Q^N / Q$
- Precision at rank 1 ($P@1$) – percentage of questions in Q^N with the correct answer on the first position

Overall Results (P@I)



Baseline – No re-ranking, IR using BM25

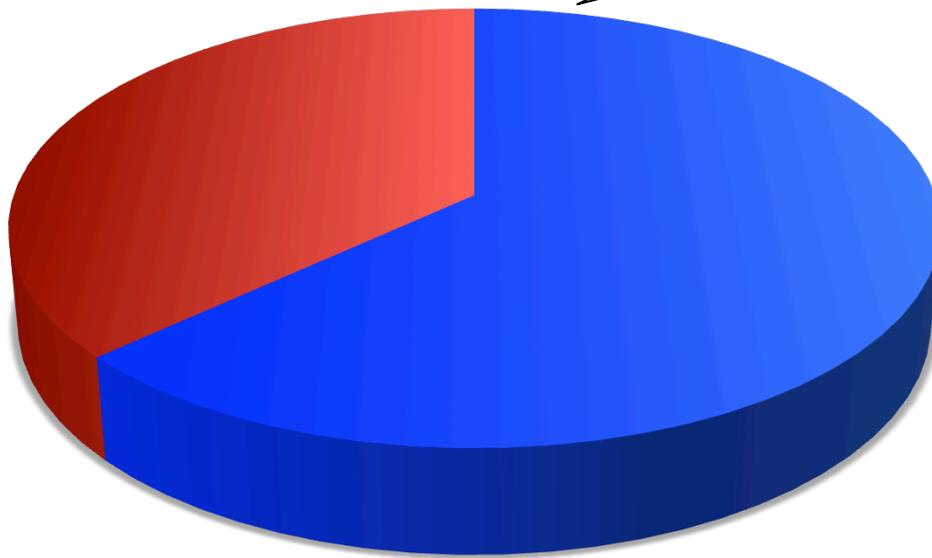
Contribution of Each Feature Group



Translation models complement well the IR baseline

Contribution of NLP

50% of improvement comes from SRL!



■ NLP
■ No NLP

NLP = syntax +
SRL +
NER +
WSD

NLP features complement well the bag-of-words baseline

ERROR ANALYSIS

Improved Questions

how can i cook grouse quick with normal household spices w/o going out to buy stuff?

Baseline	Re-rank	Correct?	Answer Text
----------	---------	----------	-------------

10

1

yes

I've never *cooked a grouse*, but poultry is poultry... You could ~~salt it~~ and pepper it, put some flour on a plate and roll it in the flour to coat it lightly, then heat a few tablespoons of olive oil in a skillet and pan-fry it. (If you have no olive oil, use a little vegetable oil plus a pat of butter – the oil ~~is to~~ keep the butter from burning.)
Some ~~add~~ a few drops of lemon juice over it if you want.
You can skip the flour. Salt and pepper the *grouse*. Pan-fry...

Shared structures:
Arg1(*cook, grouse*)

Many words that translate to *cook*:
salt, pepper, flour, tablespoons, oil, skillet, etc.

Well, a *grouse* is a prey animal. If there was a decline in the population of *grouse*, then the animals that usually prey on the *grouse* - coyotes, owls, etc - would probably start eating other prey animals, like the pheasants and squirrels.

Improved Questions

how did I do for make a call and that the other dont see my number?

Baseline	Re-rank	Correct?	Answer Text
----------	---------	----------	-------------

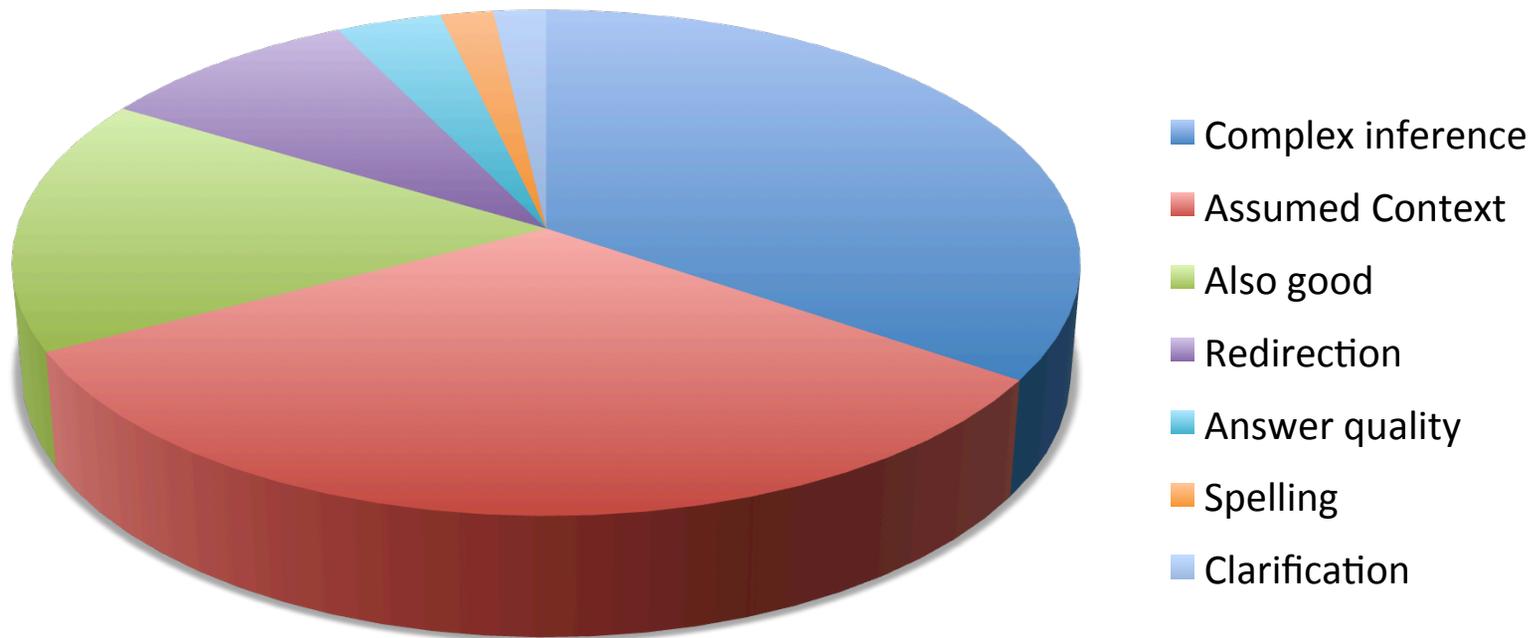
2	1	yes	to <i>make a call</i> so that the other person cant see the <i>number</i> ... dial *67 and wait for the three beeps.. then <i>dial the number</i>
---	---	-----	---

Shared structures:
Arg1(*make, call*)
Arg1(*see, number*)

Translated structures:
Arg1(*make, call*) to
Arg1(*dial, number*)

Oneday out of the blue *call* her. If u dont have her *number*, when u see her ask her if she wanted to go out oneday then get her *number*. When u talk on the phone get to know her. But dont ask her out too soon because she may not feel the same way. After a couple of days or weeks taking to her let her know how u felt about her since the first time u met her.

Questions Worsened by Re-ranking



Questions Worsened by Re-ranking

Questions Requiring Complex Inference

how to deal with a person in denial with M.P.D.?

Baseline	Re-rank	Correct?	Answer Text
1	6	yes	First, i would find out if <i>MPD</i> has been diagnosed by a professional. In current terminology, <i>MPD</i> is considered a part of Dissociative Personality Disorder. In any case, it would be up to the professionals to help this person because you could cause further problems by confronting this person with what you think the problem is. If this person is a family member, you could ask for a consultation with the psychiatric professional who is treating him/her. Please, please, just do you best to be supportive without being confrontational since that might make things even worse for that person.

being supportive and being confrontational are a form of dealing with somebody

Questions Worsened by Re-ranking Answers With Assumed Context

How do mineral ions affect the biology and survival of a pond organism?

Baseline	Re-rank	Correct?	Answer Text
1	3	yes	Some <i>mineral ions</i> are fertilizer and will increase vegetative growth while others are poisons.

Questions Worsened by Re-ranking

Model Selects Answers that Are Also Good

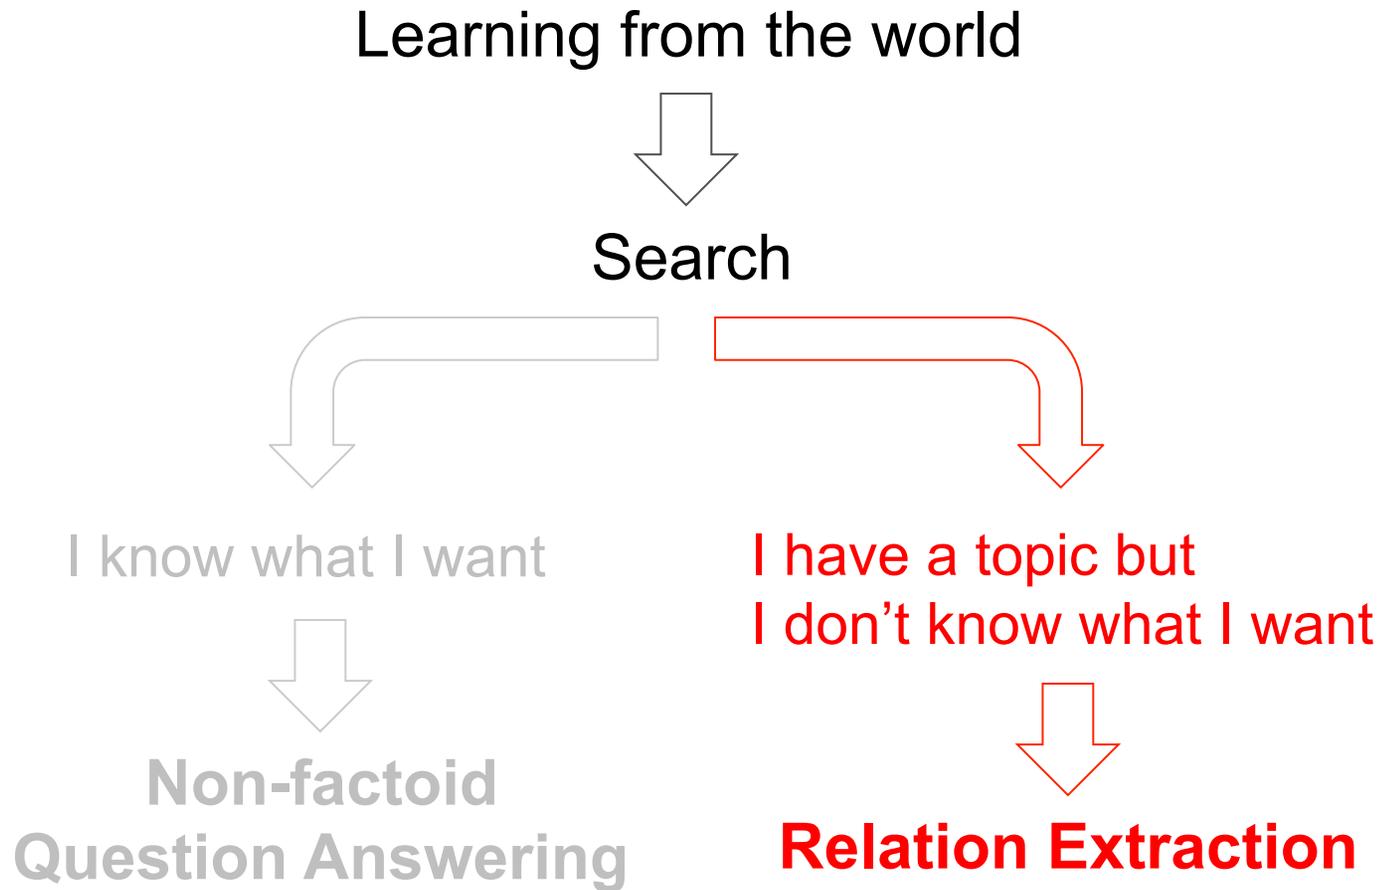
How to learn the British accent?

Baseline	Re-rank	Correct?	Answer Text
4	5	yes	Get a dictionary where there is a pronunciation guide which gives the pronunciation in <i>British</i> English. Watch <i>british</i> movies and imitate what you can. Then just practice, practice practice. But before you go about <i>learning</i> accents, slangs or dialects, make sure you brush up on your basic grammar.
3	1	no	You can do one of two things: first, go to a local bookstore, like Barnes and Noble. They sell cd's with different accents from around the world, accompanied by a book that phonetically spells the words. This is designed for actors/actresses who need to learn different accents. Also, go rent a bunch of <i>british</i> movies, or watch <i>british</i> television. Continually pause and repeat common phrases and words.

Conclusion - QA

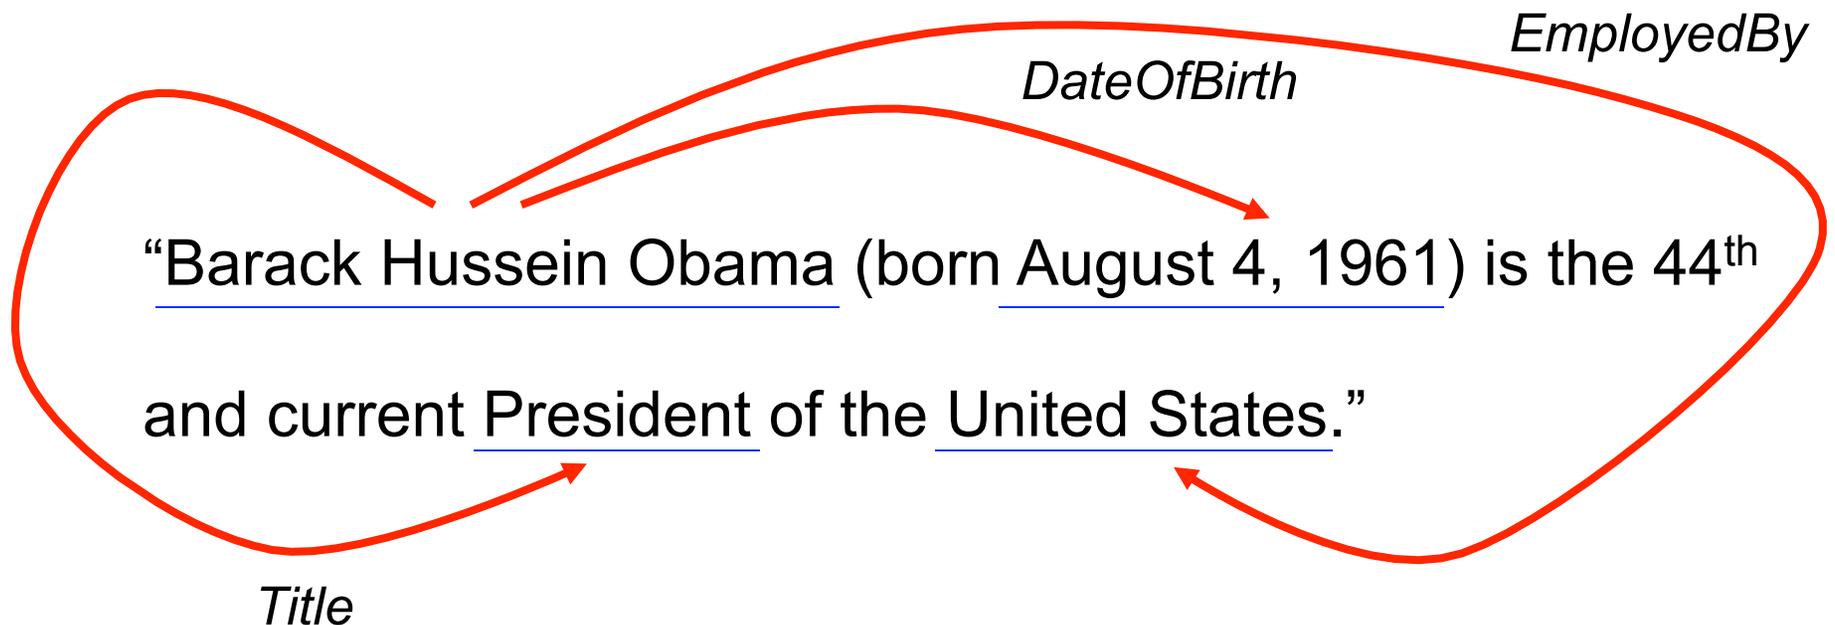
- Non-factoid QA model learned from *low quality* data
 - Key elements:
 - Powerful multi-module system combination (similarity, translation, patterns, web correlation)
 - Data representation is key to address noise. Semantic representations built with NLP (SRL, NER, WSD)
-

Roadmap

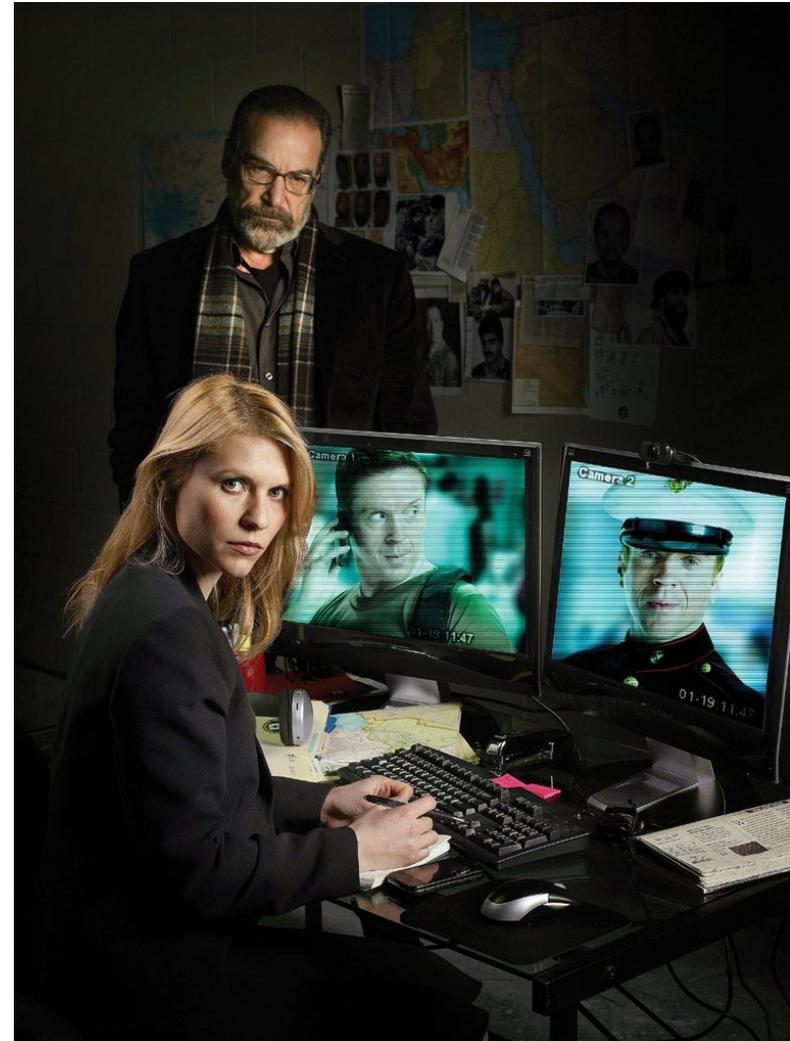
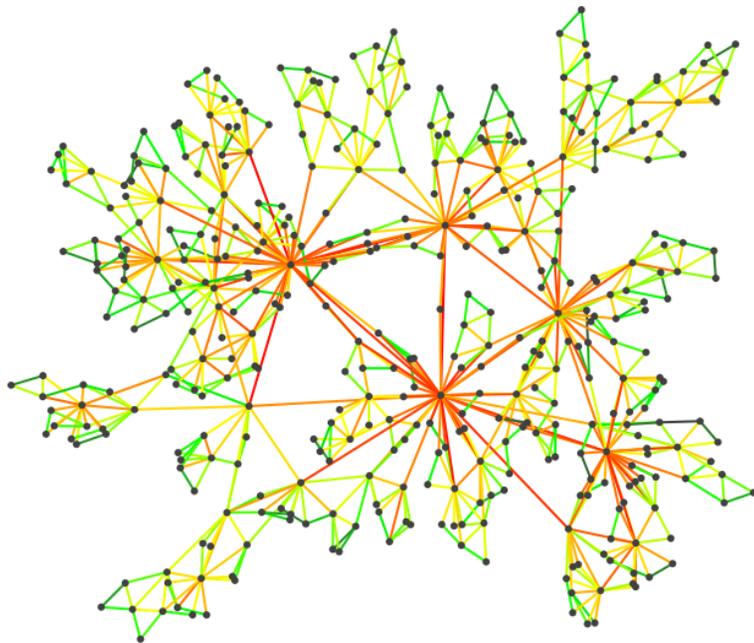


What is Relation Extraction?

RE: Task of answering extracting labeled relations between pairs of named/numeric entities.



Application of Relation Extraction



User-Generated Data Is Available



44th President of the United States

Incumbent

Assumed office
January 20, 2009

Vice President [Joe Biden](#)

Preceded by [George W. Bush](#)

United States Senator
from Illinois

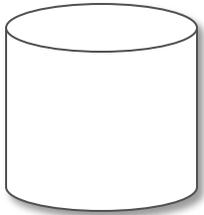
In office
January 3, 2005 – November 16, 2008

Preceded by [Peter Fitzgerald](#)

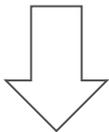
Succeeded by [Roland Burris](#)

- Millions of infoboxes in Wikipedia
- Over 22 million entities in Freebase. Each one has multiple relations

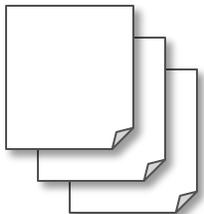
Can be aligned to text (distant supervision)



EmployedBy(Barack Obama, United States)

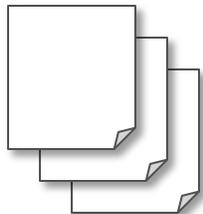
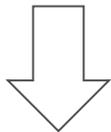
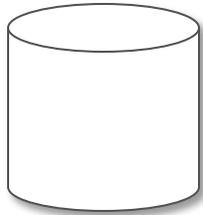


Barack Obama is the 44th and current President of the United States.



United States President Barack Obama meets with Chinese Vice President Xi Jinping today.

Not so fast!



EmployedBy(Barack Obama, United States)

BornIn(Barack Obama, United States)

Barack Obama is the 44th and current President of the United States.

United States President Barack Obama meets with Chinese Vice President Xi Jinping today.

Obama was born in the United States just as he has always said.

Obama ran for the United States Senate in 2004.

EmployedBy

EmployedBy

BornIn

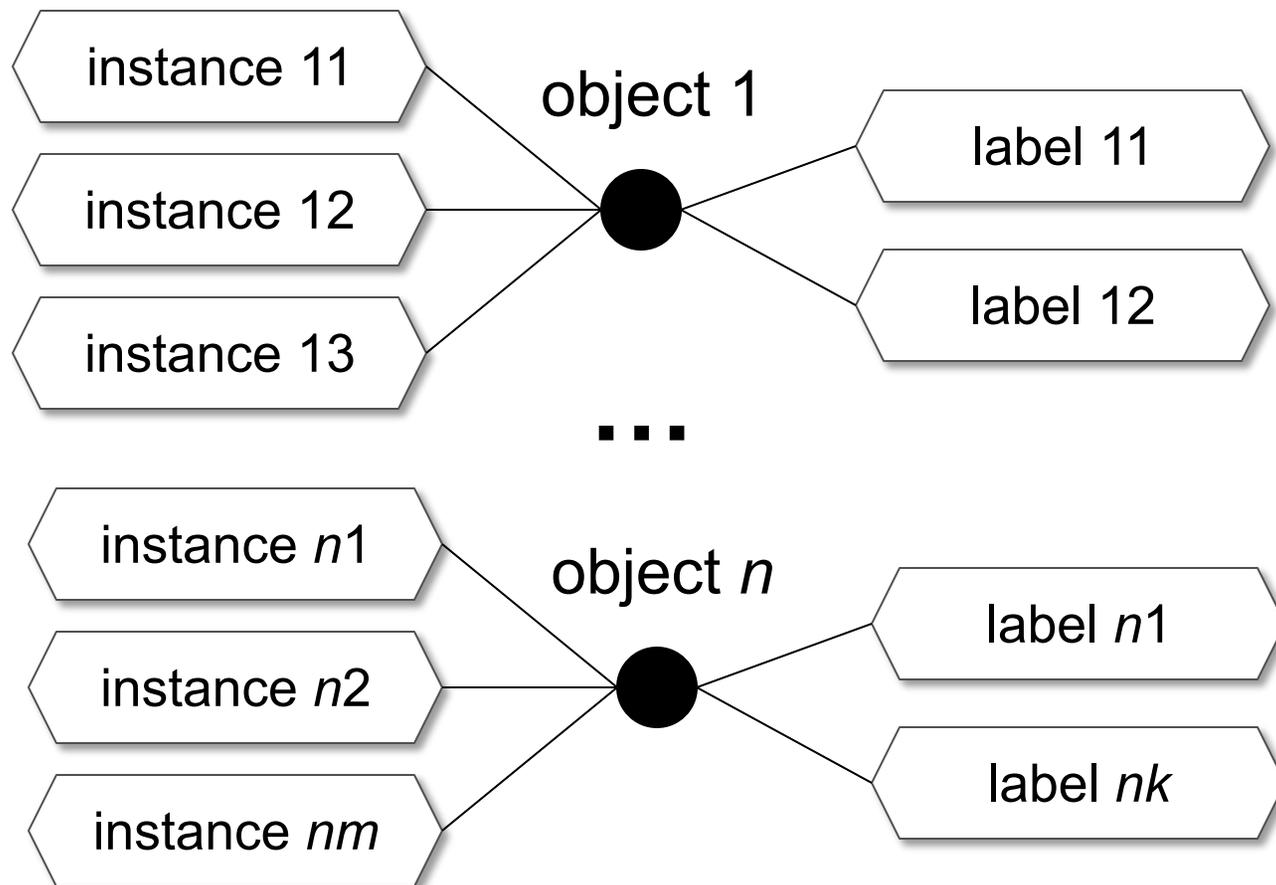
--



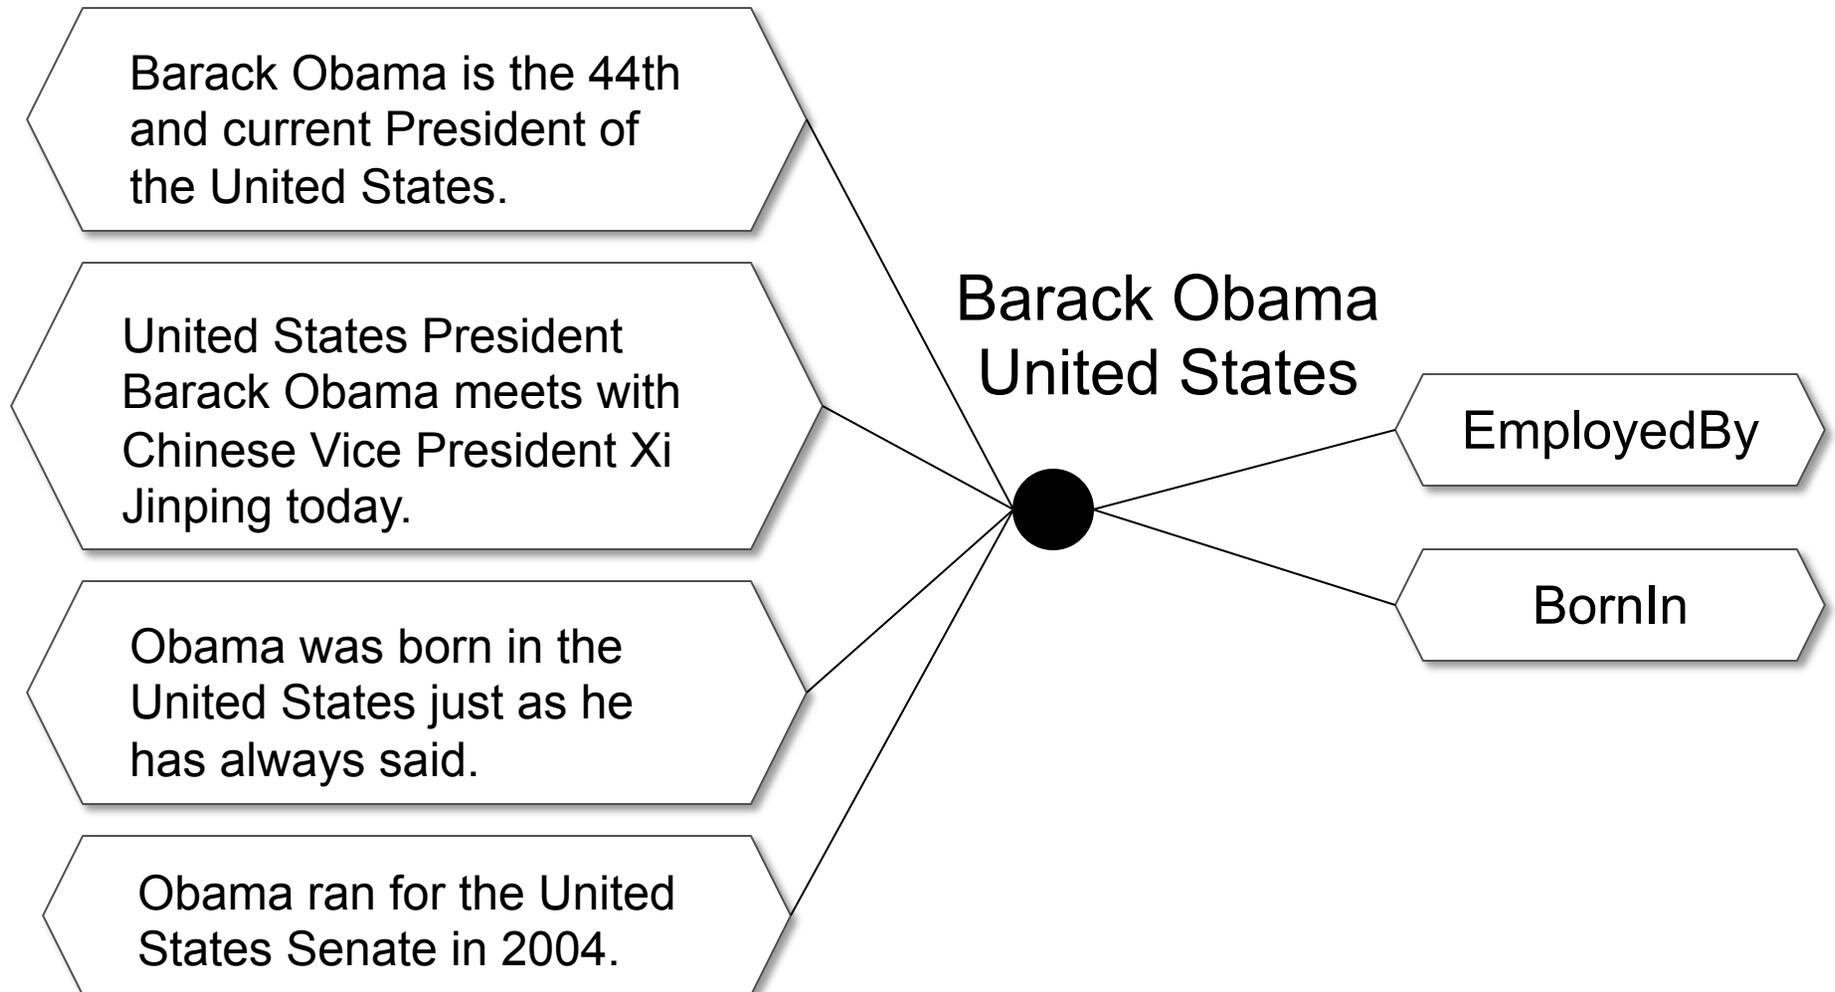
Traditional Supervised Learning



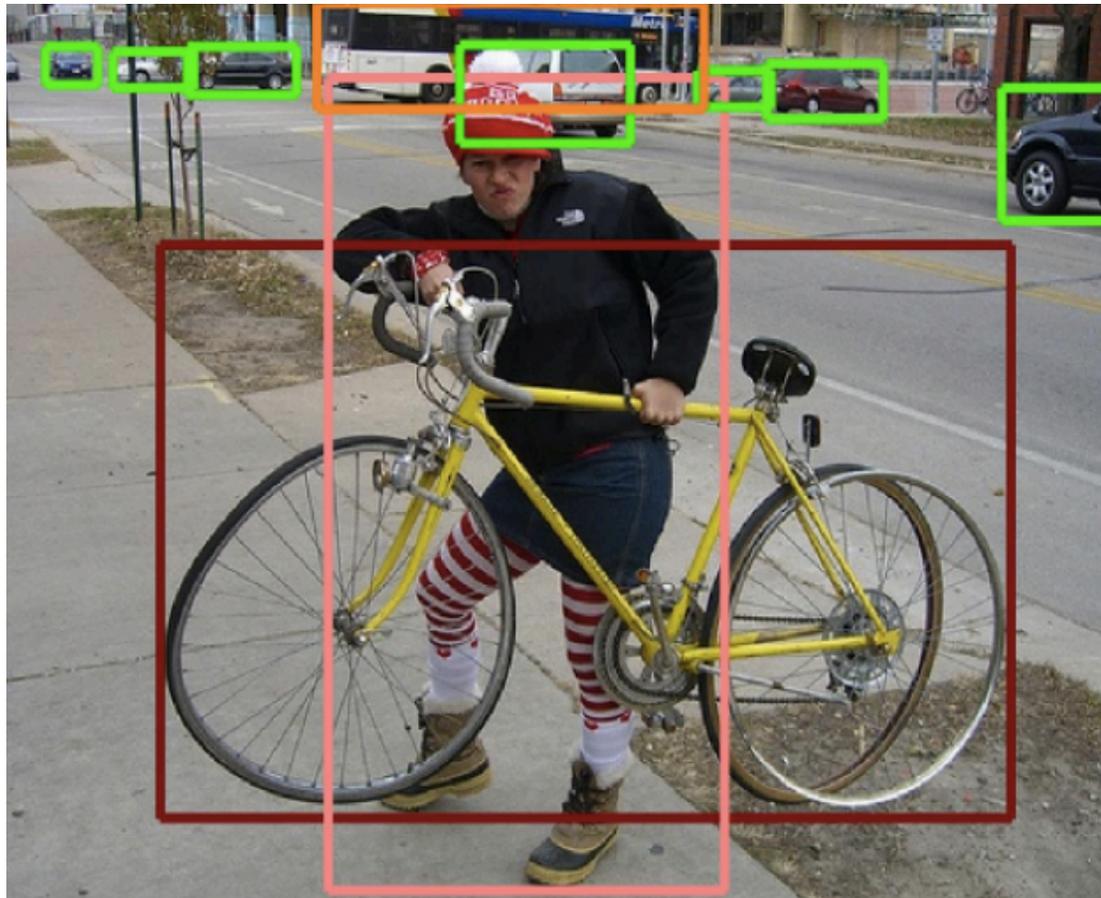
Multi-instance Multi-label Learning



MIML for Relation Extraction



MIML in Other Fields: Object Classification



“This picture contains *cars*, one *person* and a *bicycle*.”

Objective

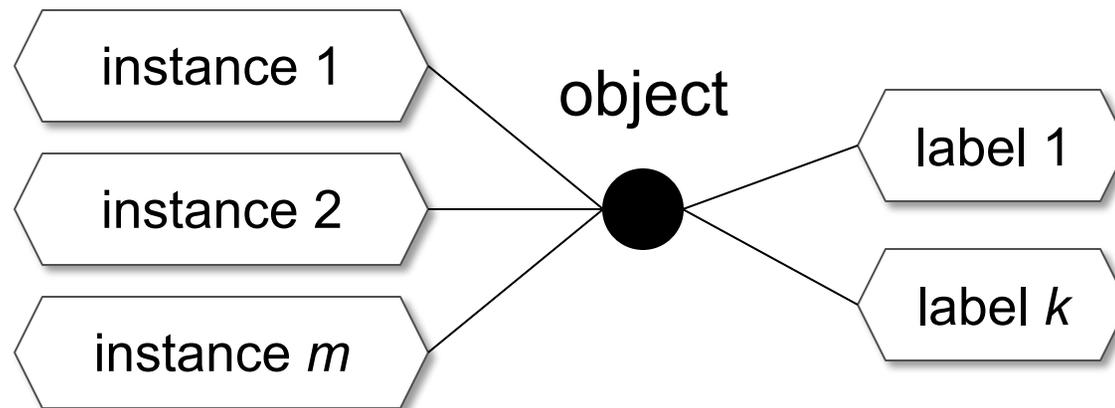


- Learn in the MIML scenario
 - Model multiple instances
 - Some instances are false positives
 - Some instances are false negatives
 - Model multiple labels
 - There are dependencies between labels

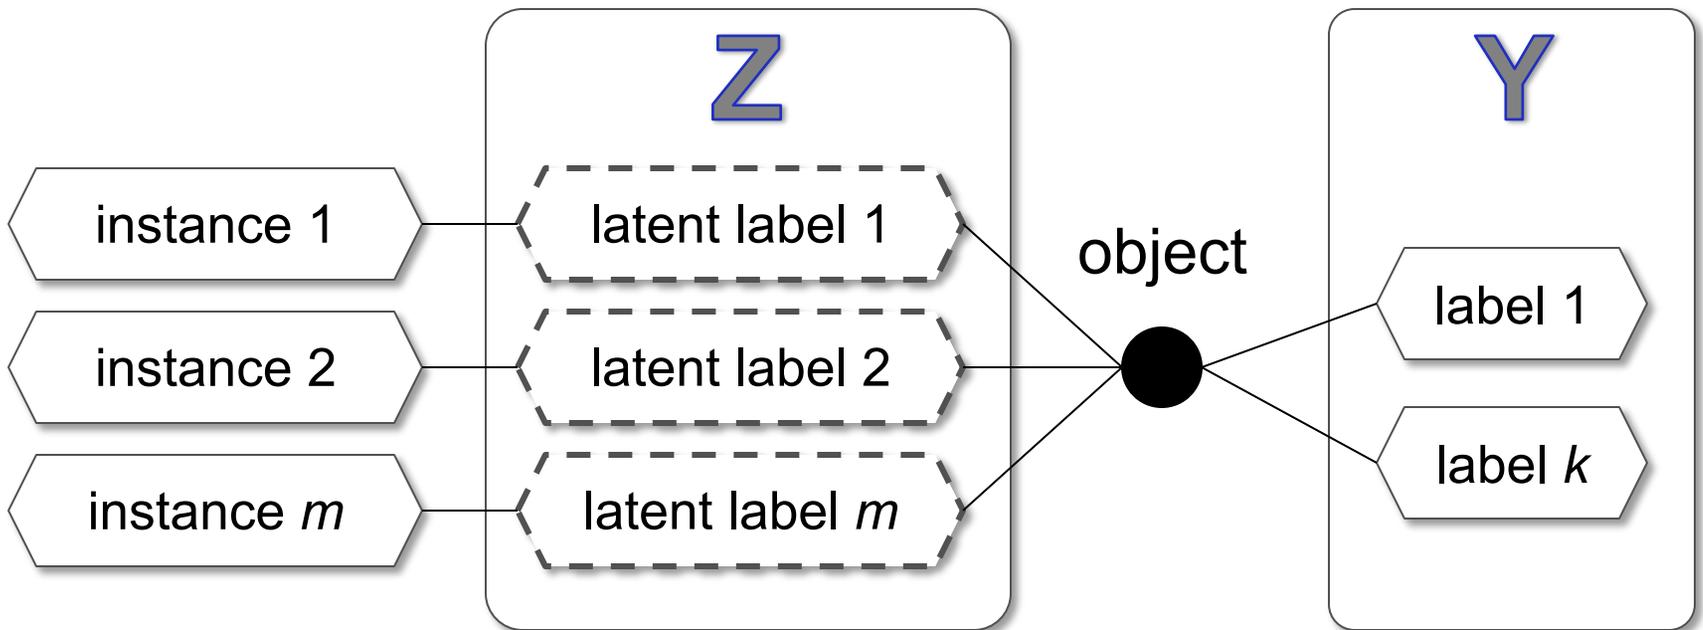


APPROACH

Model Intuition

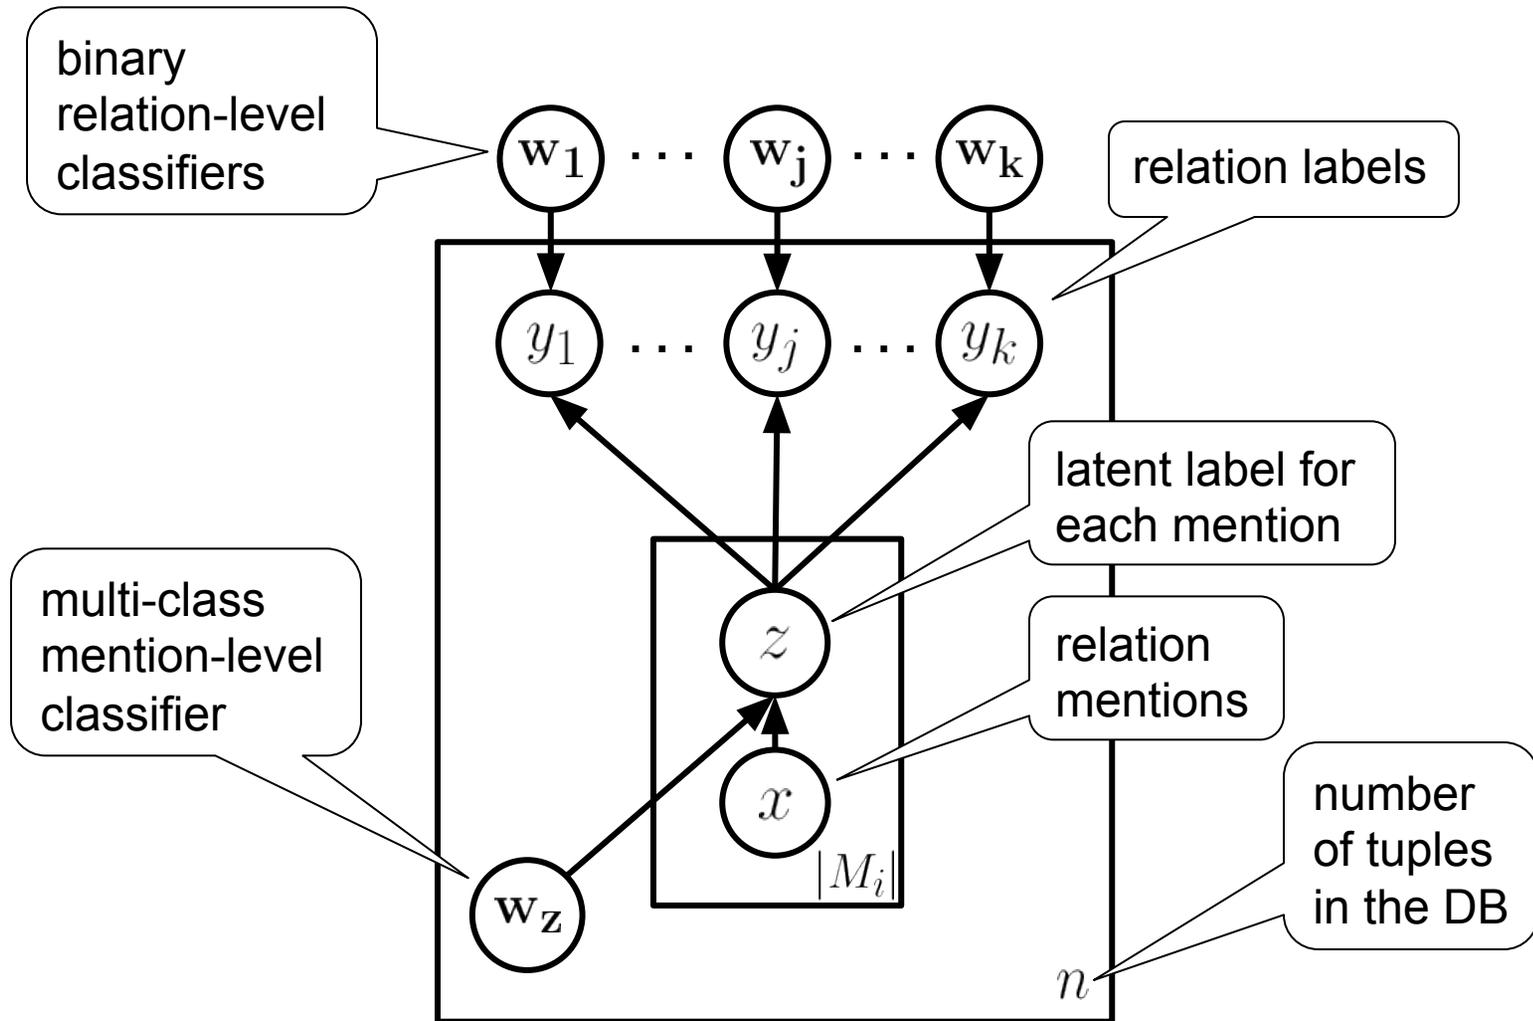


Model Intuition



- Jointly trained using discriminative EM:
 - E-step: assign latent labels using current θ_z and θ_y .
 - M-step: estimate θ_z and θ_y using the current latent labels.

Plate Diagram



E-step: Assign Latent Labels

- For the entire entity tuple i :

$$\mathbf{z}_i^* = \arg \max_{\mathbf{z}} p(\mathbf{z} | \mathbf{y}_i, \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z)$$

- For each individual mention m in tuple i :

$$p(z_i^{(m)} | \mathbf{y}_i, \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z)$$

probability given by the mention-level model

$$p(z_i^{(m)} | \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z)$$

probabilities that the correct relation labels are assigned and the incorrect ones are not

$$\approx p(z_i^{(m)} | x_i^{(m)}, \mathbf{w}_z) p(\mathbf{y}_i | \mathbf{z}_i, \mathbf{w}_y)$$

$$= p(z_i^{(m)} | x_i^{(m)}, \mathbf{w}_z) \prod_{r \in P_i \cup N_i} p(y_i^{(r)} | \mathbf{z}_i', \mathbf{w}_y^{(r)})$$

replaced mention (m) with $z_i^{(m)}$

M-step: Maximize Log Likelihood

$$\begin{aligned} LL(\mathbf{w}_y, \mathbf{w}_z) &= \sum_{i=1}^n \log p(\mathbf{y}_i | \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z) \\ &= \sum_{i=1}^n \log \sum_{\mathbf{z}_i} p(\mathbf{y}_i, \mathbf{z}_i | \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z) \end{aligned}$$

maximize this probability

M-step: Maximize Log Likelihood

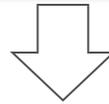
one multi-class logistic regression

$$\begin{aligned} \log p(\mathbf{y}_i, \mathbf{z}_i | \mathbf{x}_i, \mathbf{w}_y, \mathbf{w}_z) \\ = \sum_{m \in M_i} \log p(z_i^{(m)} | x_i^{(m)}, \mathbf{w}_z) + \\ \sum_{r \in P_i \cup N_i} \log p(y_i^{(r)} | \mathbf{z}_i, \mathbf{w}_y^{(r)}) \end{aligned}$$

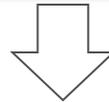
set of binary logistic regressions

Inference

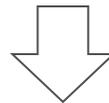
entity tuple



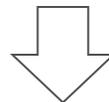
relation mentions



$$z_i^{(m)*} = \arg \max_z p(z | x_i^{(m)}, \mathbf{w}_z)$$



$$y_i^{(r)*} = \arg \max_{y \in \{0,1\}} p(y | \mathbf{z}_i^*, \mathbf{w}_y^{(r)})$$



relation labels

FEATURES

Features

- Z layer:
 - Features that model the two entities
 - Features that model syntactic context
 - Features that model surface context
- Y layer:
 - At least one mention with my label?
 - Co-occurrence of my label with the other $K - 1$ labels

EXPERIMENTS

Corpora



- Riedel:
 - DB: Freebase
 - Text: NY Times
 - Evaluation: a fragment of the DB
 - But the DB is incomplete!
 - Knowledge Base Population (KBP):
 - DB: Wikipedia infoboxes
 - Text: newswire, blogs, telephone conversations, Wikipedia
 - Development: 40 queries from the 2010 and 2011 eval
 - Evaluation: 160 queries from the 2010 and 2011 eval
 - This is complete!
-

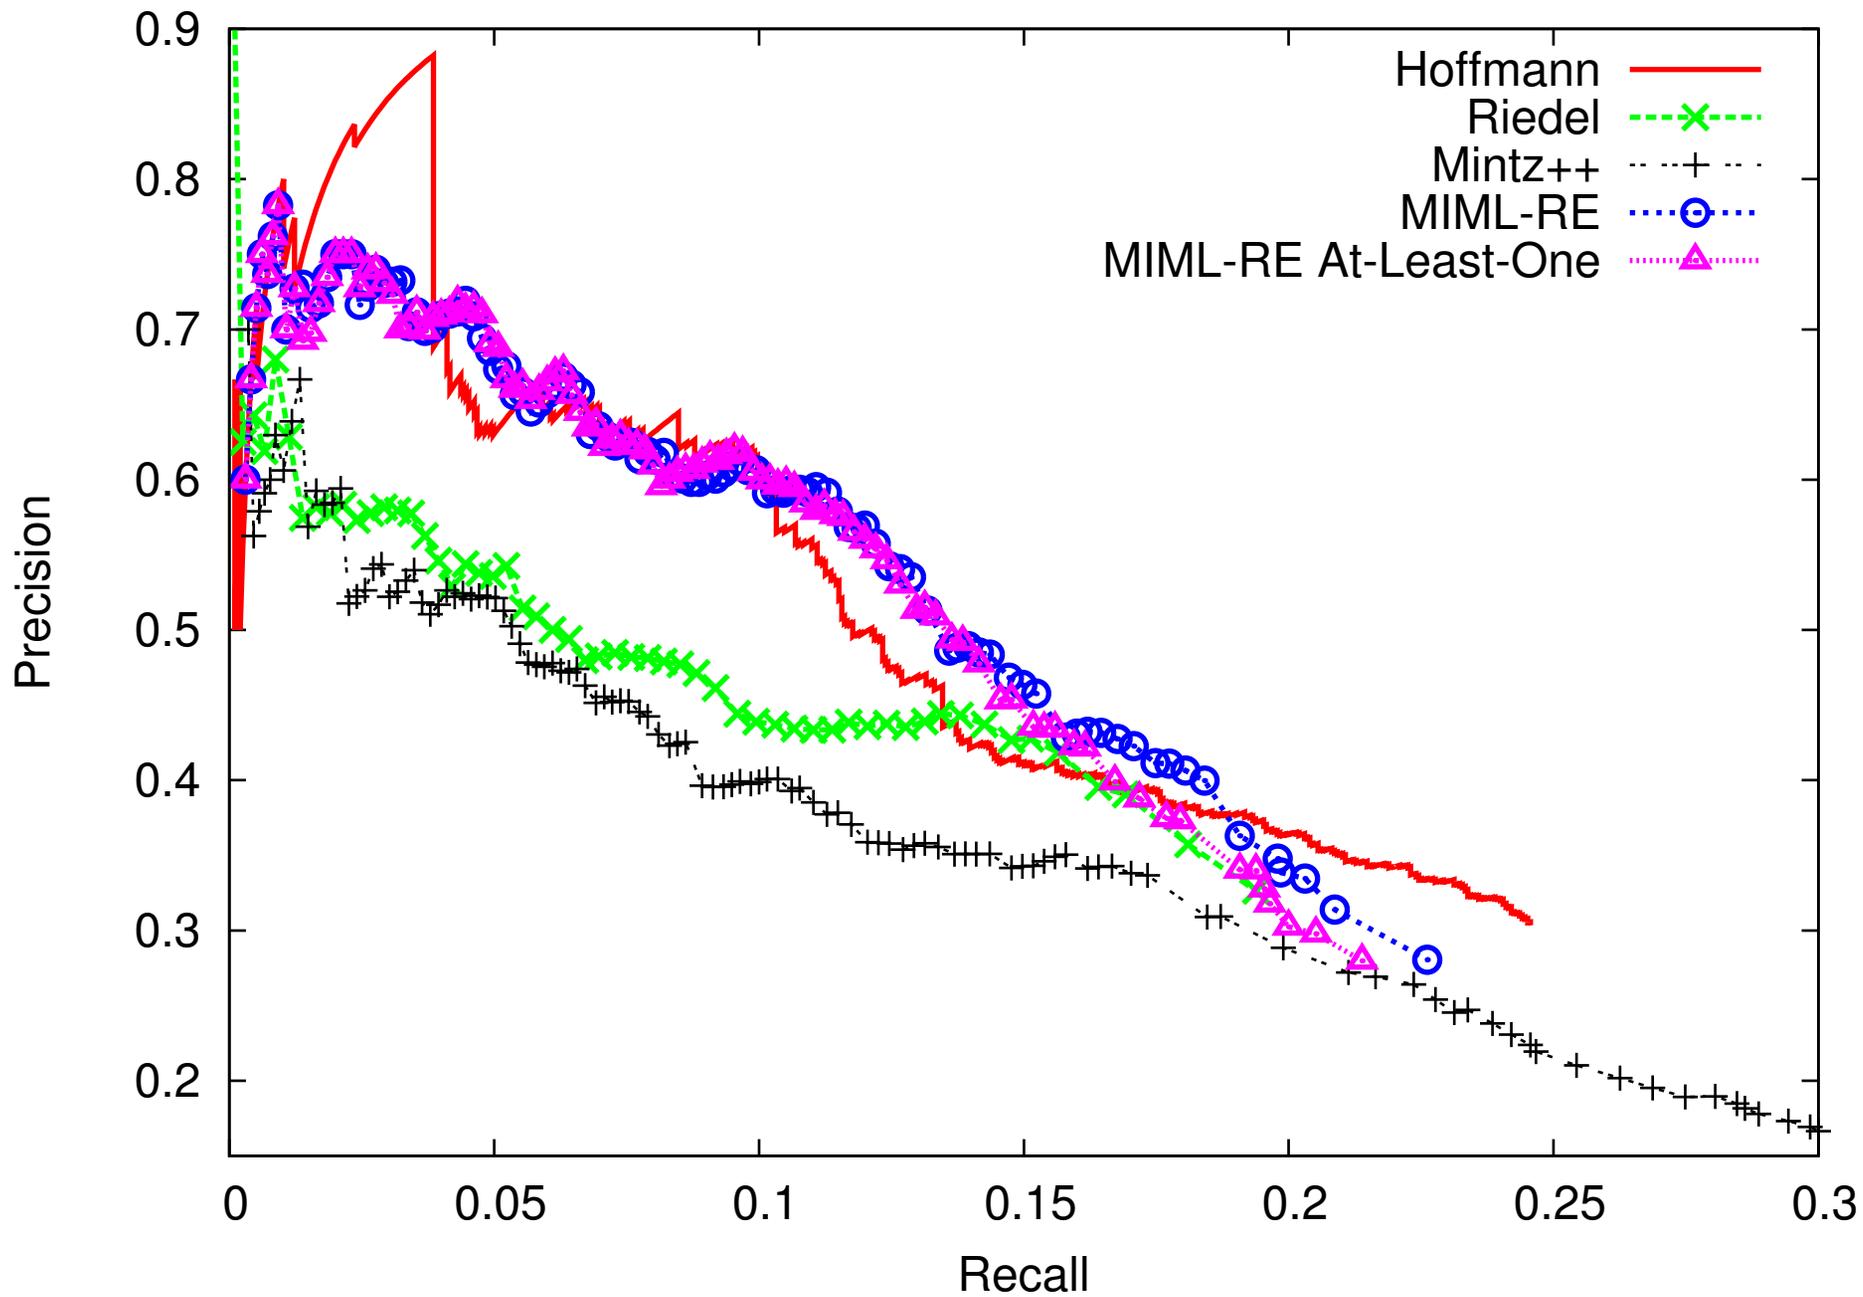
Corpora

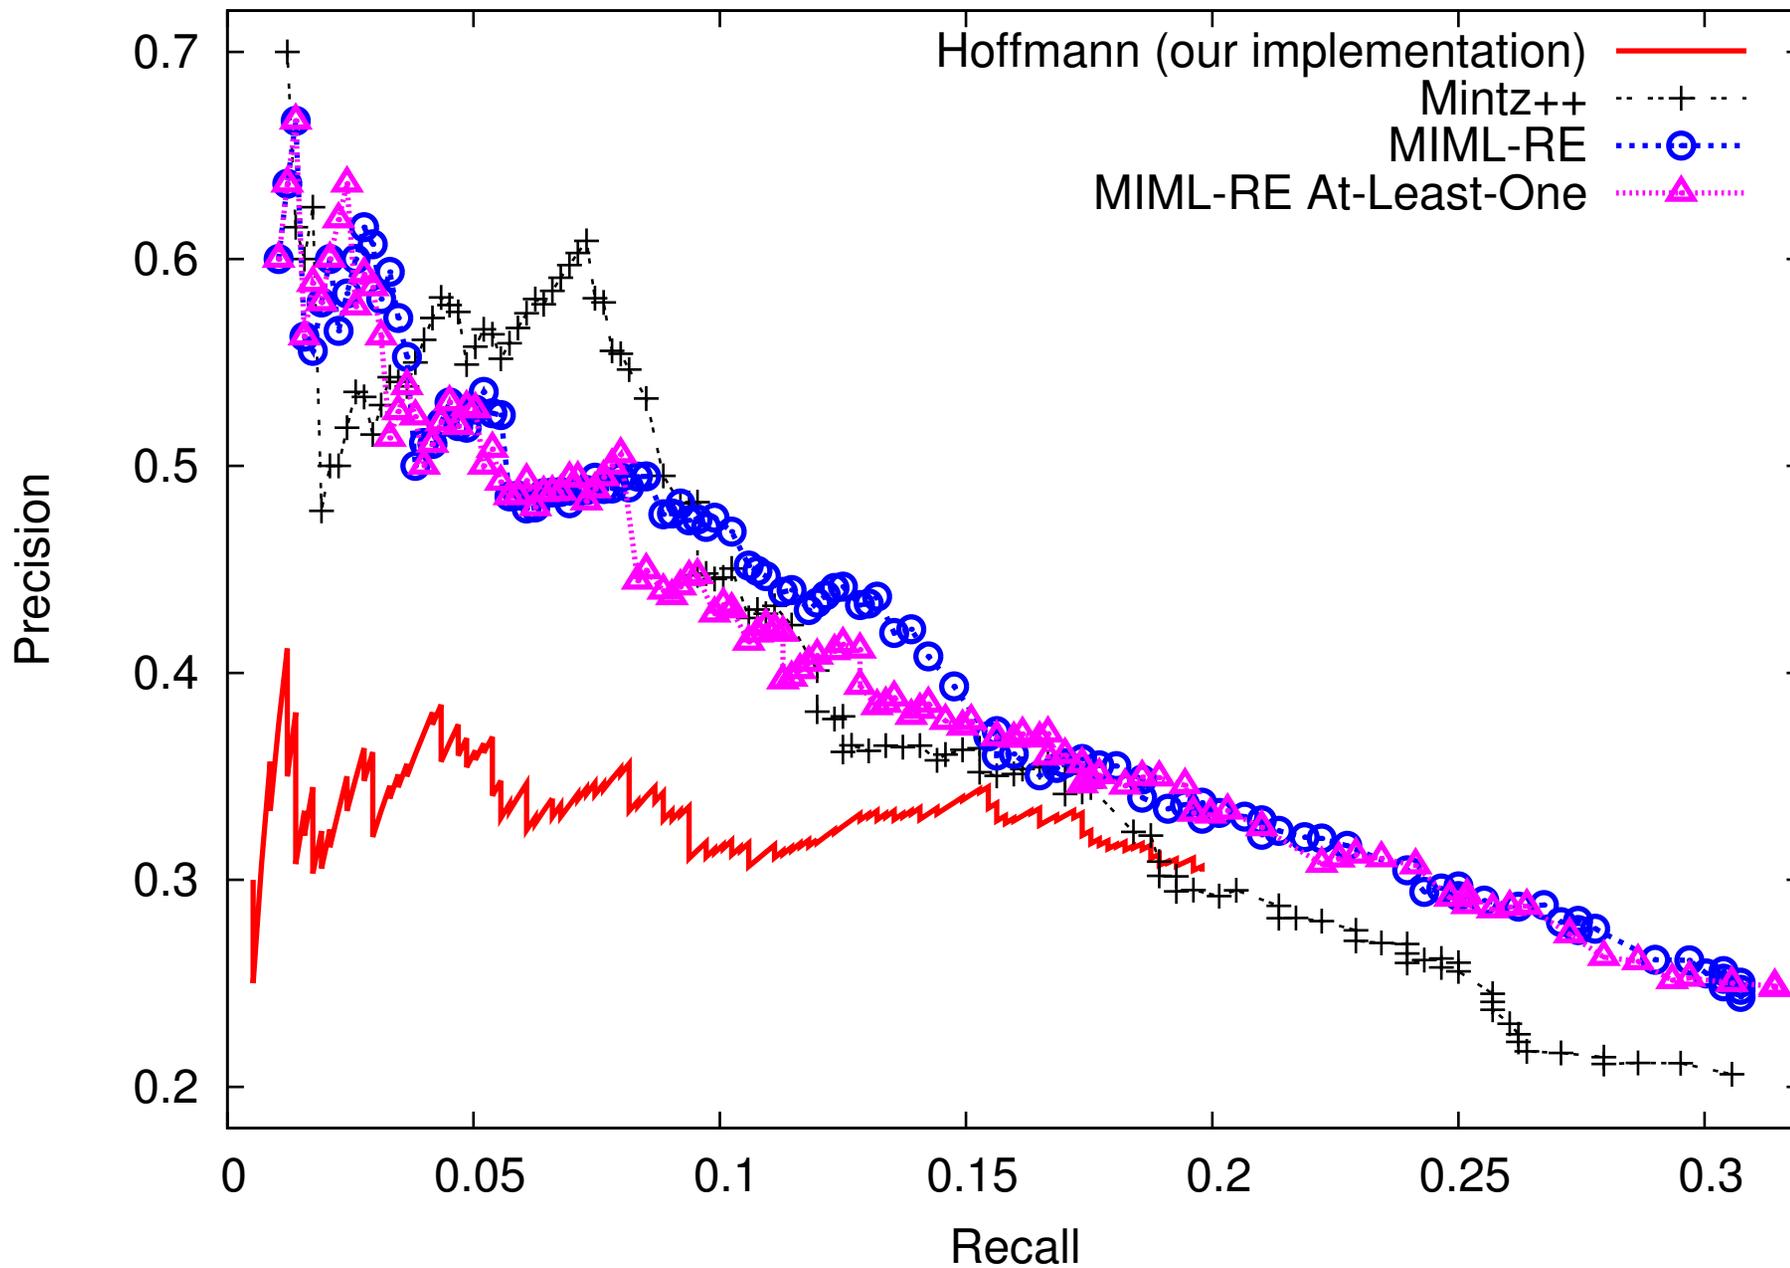


	# of gold relations in training	# of gold relations in testing	# of relation labels	% of gold entity tuples with multiple mentions in text (training)	% of mentions that do not express their relation	% of gold entity tuples with more than one label (training)
Riedel	4,700	1,950	51	46.4%	up to 31%	7.5%
KBP	183,062	3,334	41	65.1%	up to 39%	2.8%

Baseline/State of the Art

- Mintz++ (ACL, 2009)
 - “Traditional” distant supervision
 - No inter-label dependencies, no modeling of incorrect labels
 - Our extensions: multi-label predictions, bagging
 - Riedel (ECML, 2010)
 - Multi-instance single-label
 - Hoffmann (ACL, 2011)
 - Multi-instance multi-label
 - Does not model inter-label dependencies
-

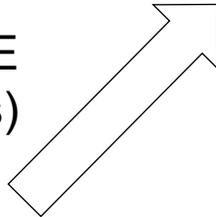




Example Predictions

“Mexico” is_country_of “Mexico City”

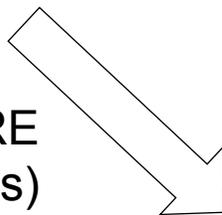
MIML-RE
(wo/ dependencies)



~~*“Mexico City” contains “Mexico”*~~

(Mexico City, Mexico)

MIML-RE
(with dependencies)



“Mexico” is_country_of “Mexico City”

Example: Dependencies Learned

Loc2 is_country_of Loc1



Loc1 contains Loc2

Person1 lived_in Loc2



Person1 was_born_in Loc2

Conclusion - RE

- First true multi-instance multi-label approach that models distant supervision for RE
 - State of the art results on two difficult domains
 - Data that does not express exactly the problem of interest can be modeled using anomaly-aware ML
-

Summary

- Exploiting non-expert annotations in the “big data” era
 - Two important issues
 - Quality of the data
 - Annotations may not model exactly the same problem
-

Learning from the World



Pick any three

Acknowledgements



- QA
 - Joint work with Massimiliano Ciaramita and Hugo Zaragoza (Computational Linguistics, 2011)
 - RE
 - Joint work with Ramesh Nallapati, Julie Tibshirani, and Chris Manning (EMNLP, 2012)
-

Thank You! Questions?



RunnerJo...

Your Resolved Question

[Show me another »](#)

How do you respond to a question when you don't know the answer?



Annmaree

Just say...I dont know..follow it up with a blank stare.
I would give a stupid , yet hum...

how the smart community does it.



mukwatha...

Best Answer - Chosen by Asker

It is simple ,Ido not know the answer.Doing otherwise will leave you appearing awkward.It is good to be genuine in whatever you do.

kuy4p1n0...

APPENDIX

Features

- FG1: similarity features
 - *Intuition: a better answer will reuse words from the question*
 - FG2: translation features
 - *Intuition: Q is a new language translated from A. A better answer will generate a more likely translation for the question.*
 - FG3: density/frequency features
 - *Intuition: a better answer will reuse patterns from the question*
 - FG4: web correlation features
 - *Intuition: words from a good answer will appear jointly with question words in other web pages or web queries*
-

FG I: Similarity Features



Intuition: a better answer will reuse words from the question

- Feature value: the similarity between Q and A using IR metrics
- We used length-normalized BM25 and cosine similarity using *tf-idf* weights

FG2: Translation Features

Intuition: similarity models perform poorly because they fail to “bridge the lexical chasm” between Q and A.

Intuition: Q is a new language translated from A. A better answer will generate a more likely translation.

- Feature value: $P(Q|A)$

$$P(Q|A) = \prod_{q \in Q} P(q|A)$$

Tuned to optimize retrieval

$$P(q|A) = (1 - \lambda)P_{ml}(q|A) + \lambda P_{ml}(q|C)$$

$$P_{ml}(q|A) = \sum_{a \in A} (T(q|a)P_{ml}(a|A))$$

Adjusted to make sure $T(w|w) > T(w'|w)$

FG3: Density/Frequency Features (1/2)



Intuition: a better answer will reuse patterns from the question

- **Same word sequence**: number of question words in the same order in the answer
- **Answer span**: largest distance between two question words appearing in the answer
- **Informativeness**: number of new NN/VB/JJ in the answer

FG3: Density/Frequency Features (2/2)



Intuition: a better answer will reuse patterns from the question

- **Same sentence match:** number of question words matched in the same sentence in the answer
- **Overall match:** number of question words matched in the complete answer
- **Tree kernels:** how many dependency trees are shared between question and answer?
 - Largest value between any two sentences
 - Average of all computed kernel values

FG4: Web Correlation Features

Intuition: words from a good answer will appear jointly with question words in other web pages or web queries

- **Web correlation:**
 - Feature value = Corrected Conditional Probability (CCP)
 - $CCP(Q, A) = \text{hits}(Q + A) / (\text{hits}(Q) \text{hits}(A)^{2/3})$
 - Needs query relaxation to work for non-factoid QA
- **Query-log correlation:**
 - Largest/average PMI and X^2 values between any two words in the question and answer
 - Number of (q, a) word pairs that appear in the top 10, 5, and 1 percentile of PMI and X^2 values

Parameters

Labels of Relations

helicopter $\xrightarrow{\text{SBJ}}$ *get* vs. *helicopter* \longrightarrow *get*



Feature Selection

Iter.	Feature Set	P@1
0	BM25(W)	41.1
1	+ translation (R)	46.3
2	+ translation (N)	48.0
3	+ overall match (D_{WNSS})	48.9
4	+ translation (W)	49.1
5	+ query-log avg (PMI)	49.6
6	+ overall match (W)	49.7
7	+ overall match, normalized by Q size (W)	49.9
8	+ same word sequence, normalized by Q size (W)	49.9
9	+ BM25 (N)	50.0
10	+ informativeness: verb count	50.0
11	+ query log max (PMI)	50.3
...		

Contribution of NL Structures

	Similarity	Translation
W	0	+4.2
N	-14.0	+3.0
N _g	-18.7	+3.0
D	-15.2	+1.5
D _g	-19.3	+3.4
R	-27.6	+0.3
R _g	-28.3	+3.5

Contribution of NL Structures

	Similarity	Translation
W	0	+4.2
N	-14.0	+2.5
N _g	-18.7	+3.6
D	-15.2	+1.5
D _g	-19.3	+3.4
R	-27.6	+0.3
R _g	-28.3	+3.5
W + N + N _g + D + D _g + R + R _g	+1.65	+6.3

M-step: Maximize Log Likelihood

one multi-class logistic regression

$$\mathbf{w}_{\mathbf{z}}^* = \arg \max_{\mathbf{w}} \sum_{i=1}^n \sum_{m \in M_i} \log p(z_i^{(m)*} | x_i^{(m)}, \mathbf{w})$$

$$\mathbf{w}_{\mathbf{y}}^{(\mathbf{r})*} = \arg \max_{\mathbf{w}} \sum_{1 \leq i \leq n \text{ s.t. } r \in P_i \cup N_i} \log p(y_i^{(r)} | \mathbf{z}_i^*, \mathbf{w})$$

set of binary logistic regressions

Riedel Dataset with Groups with 10+ Mentions

	P	R	F1
Hoffmann (our implementation)	48.6	29.8	37.0
Mintz++	43.8	36.8	40.0
MIML-RE	56.1	32.5	41.1
MIML-RE (w/ dependencies)	64.8	31.6	42.6

Future Work

- How robust is “learning from the world” to domain transfer?
 - Learn from community QA data, apply to textbooks? Mobile?
 - Can we extract complex structures from this type of data?
 - Event extraction
-