Feature Driven Question Answering

Benjamin Van Durme

Question Answering (QA)

Natural Language Interface to Databases

Information Retrieval (+ Information Extraction)

Question Answering (QA)

Natural Language Interface to Databases

Information Retrieval (+ Information Extraction)

1960’s, 1970’s, …

build semantic parsers by hand

Baseball (Green et al., 1961)

"Where did the Red Sox play on July 7?"

[Where] did [the Red Sox] play (on [July 7])?

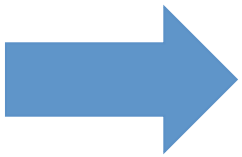

[ANSWER]

JOHNS HOPKINS UNIVERSITY

Lunar (Woods, 1977)

"What is the concentration of silicon in S10046?"

```
(For The X / (Dataline S10046 Overall SIO2) : T ;
    (For Every Y / Thing : (Equal X Y) ;
        (Printout Y) ))
```

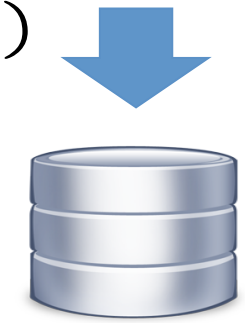

[ANSWER]

JOHNS HOPKINS UNIVERSITY

[fast forward]

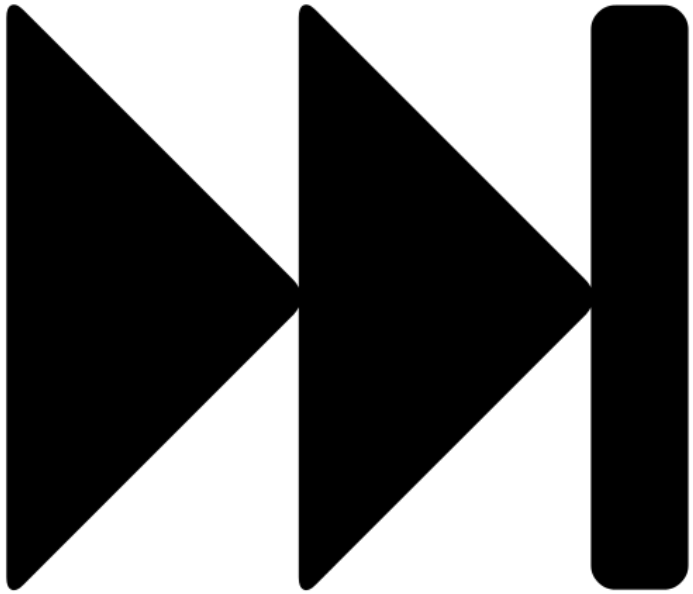

JOHNS HOPKINS
UNIVERSITY

Difference?

1960's, 1970's, …

build grammars by hand

1990's, 2000's

induce grammars from annotations

Zelle and Mooney (1996)

"What is the capital of the state with the largest population?"

```
answer(C,(capital(S,C),
              largest(P, (state(S), population(S,P)))))
```

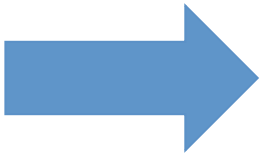

[ANSWER]

JOHNS HOPKINS UNIVERSITY

Zettlemoyer and Collins (2005)

"What states border texas?"

$\lambda x.\text{state}(x) \wedge \text{borders}(x, \text{texas})$

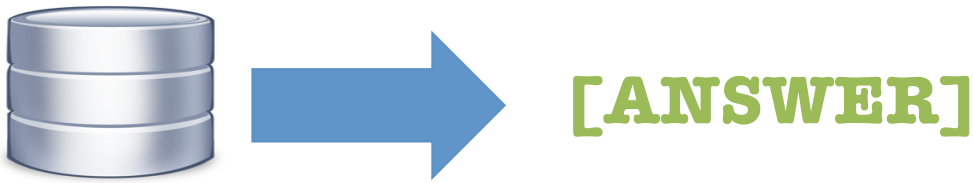

1960's, 1970's, …

build grammars by hand

1990's, 2000's

induce grammars from annotated questions

2010's

induce grammars based on questions + answers

Clarke *et al.* (2010)

"What is the largest state that borders Texas?"

JOHNS HOPKINS UNIVERSITY

Clarke *et al.* (2010)

"What is the largest state that borders Texas?"

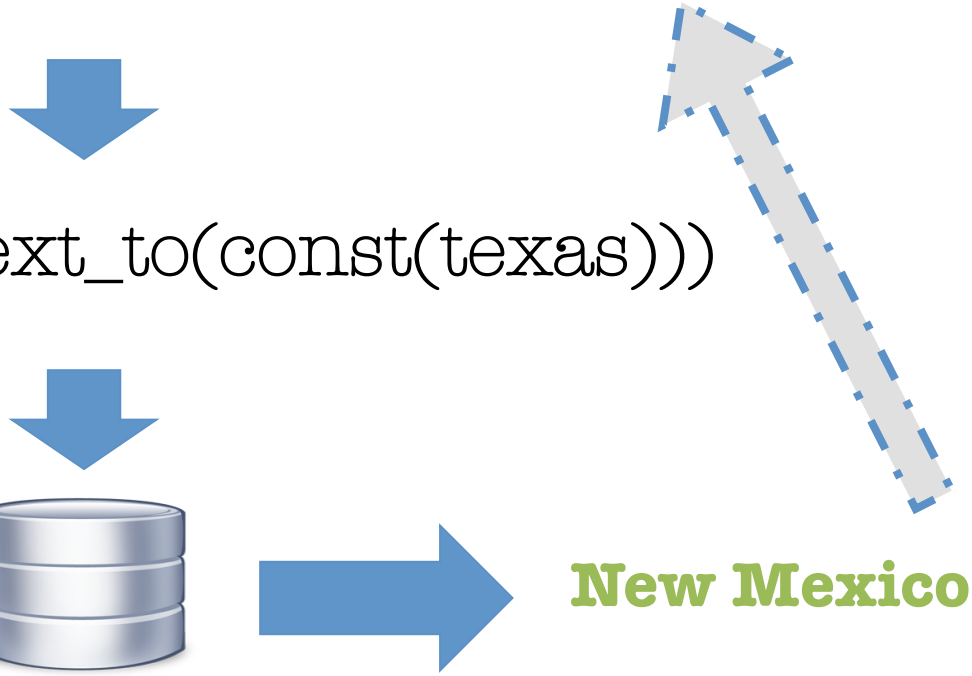

JOHNS HOPKINS UNIVERSITY

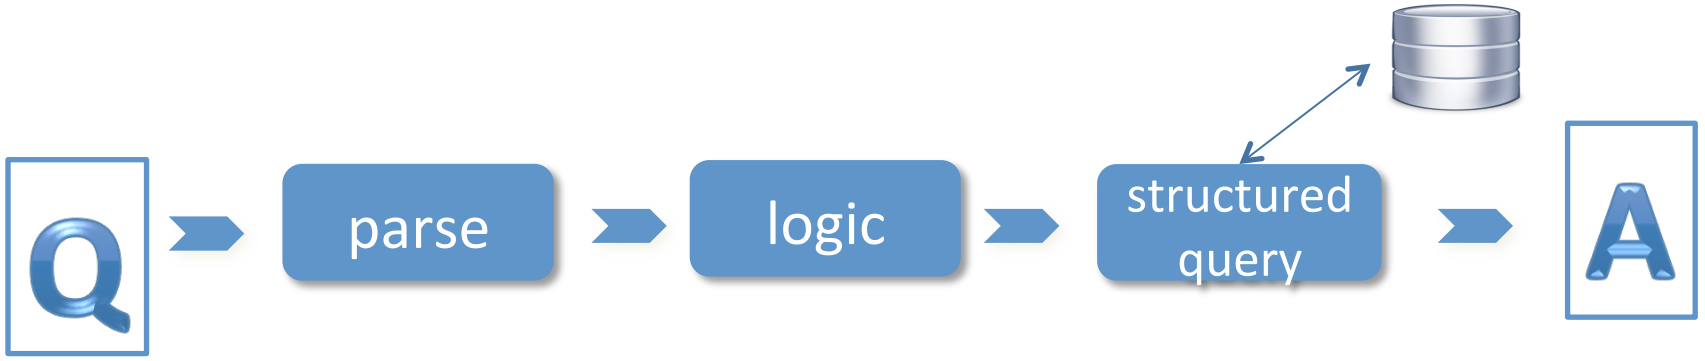
Q
parse
logic
structured query
A

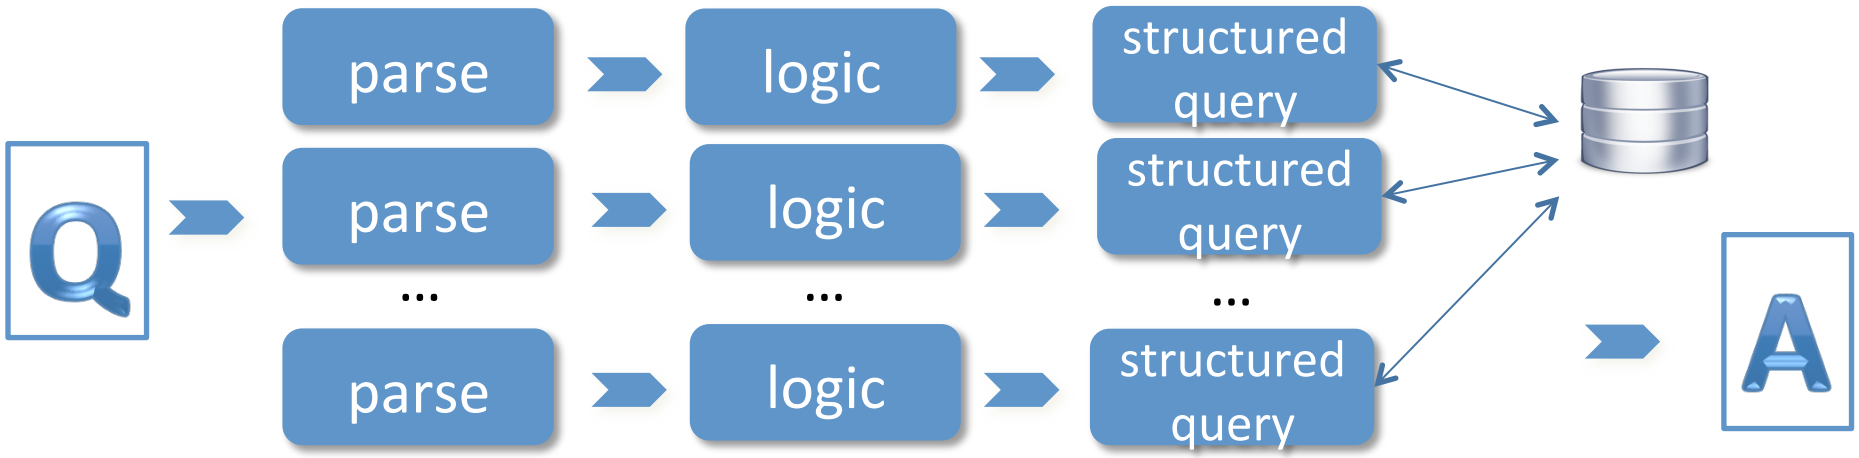
Q
parse
parse
...
parse
logic
logic
...
logic
structured query
structured query
...
structured query
A

Question Answering (QA)

Natural Language Interface to Databases

Information Retrieval (+ Information Extraction)

Question Answering (QA)

Natural Language Interface to Databases

Information Retrieval (+ Information Extraction)

"What is the something something something?"

`Lambda x [constraint1(x), constraint2(x), ...]`

"What is the something something something?"

`Lambda x [constraint1(x), constraint2(x), ...]`

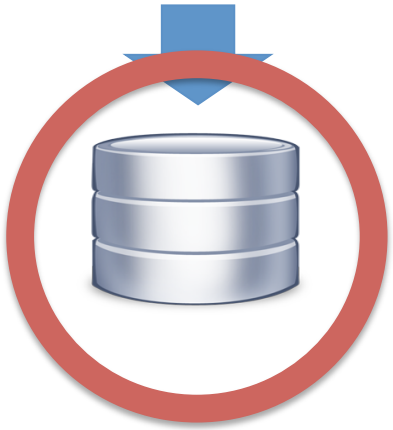

JOHNS HOPKINS
UNIVERSITY

"What is the something something something?"

`Lambda x [constraint1(x), constraint2(x), ...]`

"What is the something something something?"

`Lambda x [constraint1(x), constraint2(x), ...]`

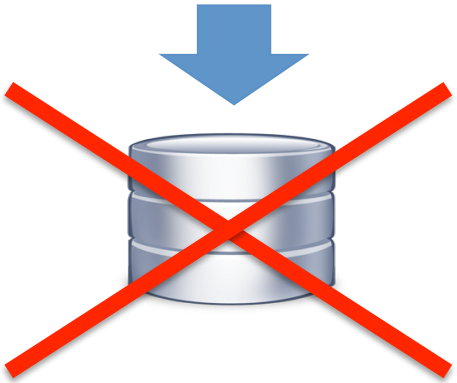

JOHNS HOPKINS
UNIVERSITY

“What is the something something something?”

[ANSWER]

JOHNS HOPKINS
UNIVERSITY

NIST TREC-8 QA

(Voorhees 1999)

200 questions, each with “snippet” answers

“How many calories in a Big Mac?”

There are 549 calories in 1 burger (7.6 oz) of McDonald's Big Mac Burger. You'd need to walk …

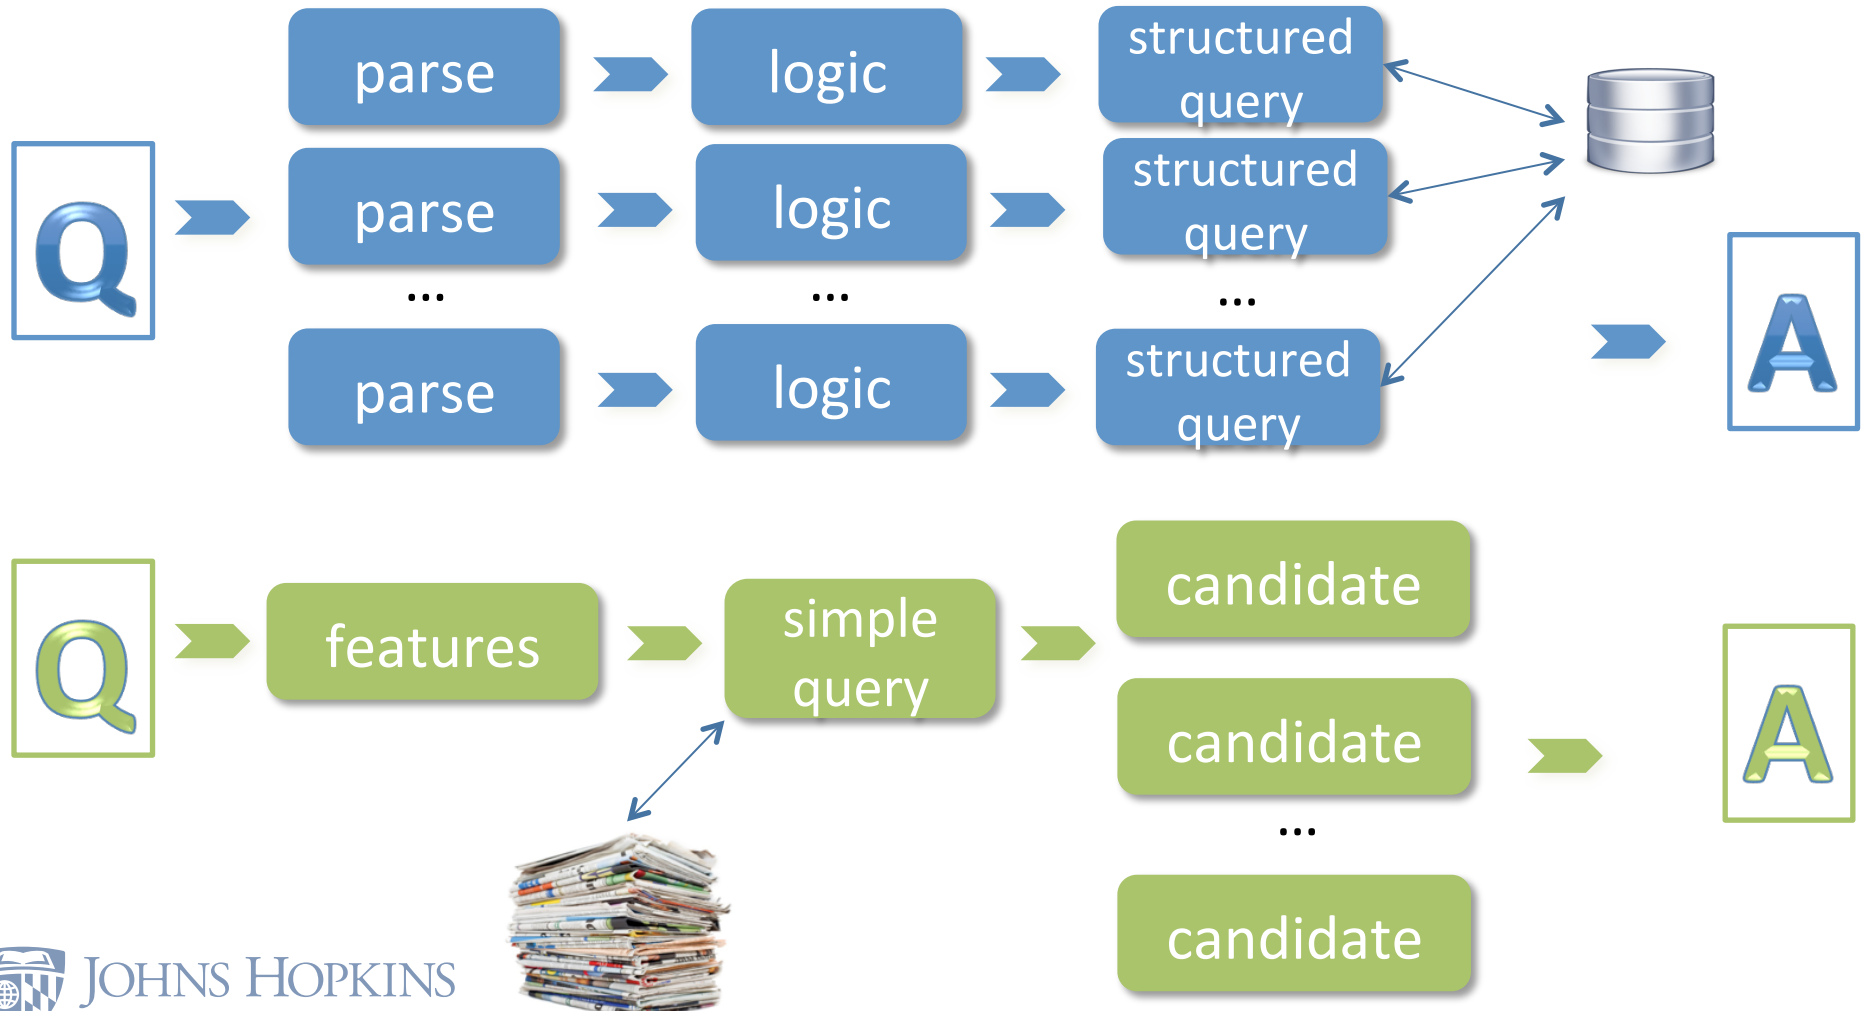

JOHNS HOPKINS UNIVERSITY

Chen and Van Durme (in-progress)

Yao and Van Durme (2014)

Yao *et al.* (2014)

Yao *et al.* (2013a,b,c)

Van Durme *et al.* (2003)

Nyberg *et al.* (2003)

Xuchen Yao

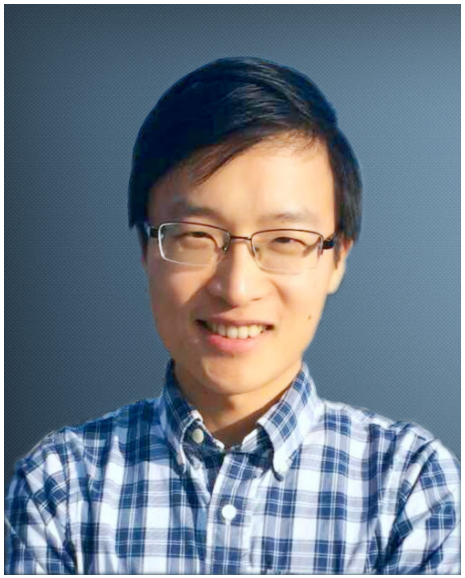

Chen and Van Durme (in-progress)

Yao and Van Durme (2014)

Yao *et al.* (2014)

Yao *et al.* (2013a,b,c)

Van Durme *et al.* (2003)

Nyberg *et al.* (2003)

JOHNS HOPKINS UNIVERSITY

KITT·AI

AI2

Question Answering (QA)

Information Retrieval (+ Information Extraction)

- Feature driven answer extraction
- Feature driven retrieval
- Feature driven QA on a KB

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

“What sport does Jennifer Capriati play?”

[ANSWER]

JOHNS HOPKINS
UNIVERSITY

"What sport does Jennifer Capriati play?"

"Capriati hasn't played on the tour since …"

"Tennis player Jennifer Capriati is 23"

"Jennifer Capriati enjoys sports"

…

JOHNS HOPKINS UNIVERSITY

"What sport does Jennifer Capriati play?"

"Capriati hasn't played on the tour since …"

"Tennis player Jennifer Capriati is 23"

"Jennifer Capriati enjoys sports"

…

JOHNS HOPKINS
UNIVERSITY

"What sport does Jennifer Capriati play?"

"Capriati hasn't played on the tour since …"

"Tennis player Jennifer Capriati is 23"

"Jennifer Capriati enjoys sports"

…

JOHNS HOPKINS
UNIVERSITY

"What sport does Jennifer Capriati play?"

Tennis player Jennifer Capriati is 23

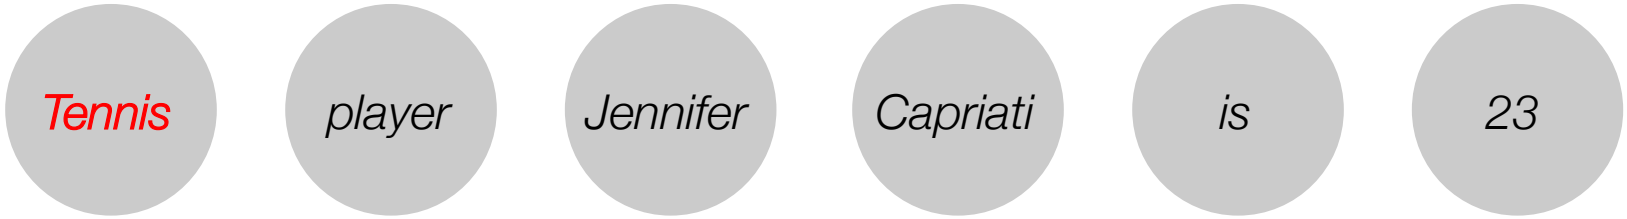

JOHNS HOPKINS
UNIVERSITY

"What sport does Jennifer Capriati play?"

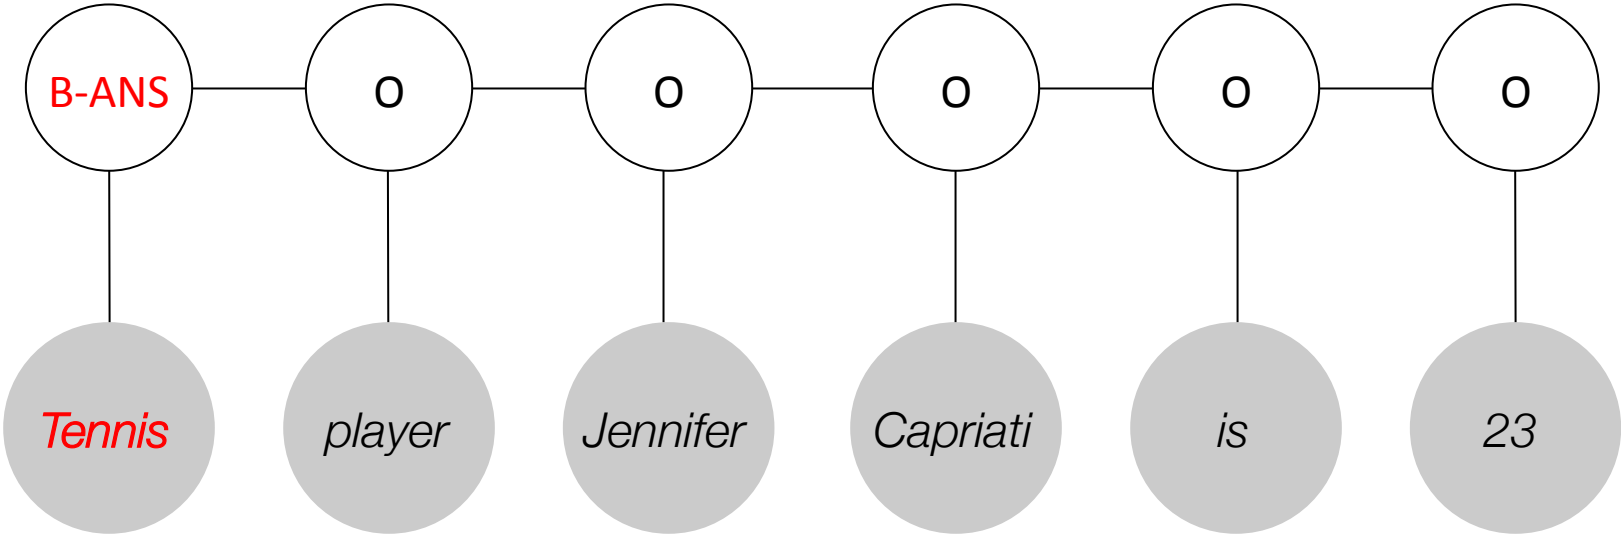

"BIO" tagging, 3 types of hidden states:

- B-ANS (beginning of answer)
- I-ANS (inside of answer)
- O (outside of answer, i.e., not an answer)

JOHNS HOPKINS UNIVERSITY

Answer Extraction as Sequence Tagging

using a linear-chain Conditional Random Field (CRF)

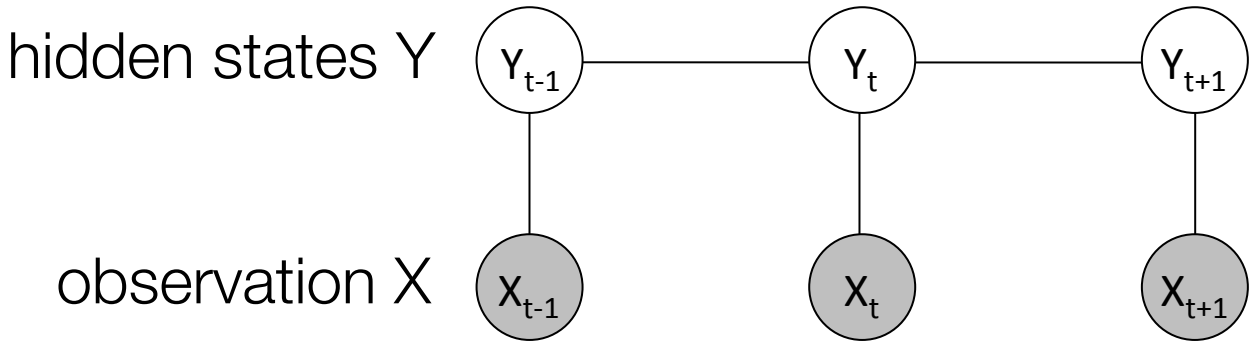

$$p(\mathbf{y} \mid \mathbf{x}) = \frac{1}{Z(\mathbf{x})} \prod_{t=1}^{T} \exp\{\sum_{k=1}^{K} \theta_k f_k(y_t, y_{t-1}, \mathbf{x}_t)\}$$

conditional model p(**y**|**x**) :

$f(y_t, y_{t-1}, \mathbf{x}_t)$: **feature functions**

θ : feature weights (to learn)

JOHNS HOPKINS UNIVERSITY

Answer Extraction as Sequence Tagging

using a linear-chain Conditional Random Field (CRF)

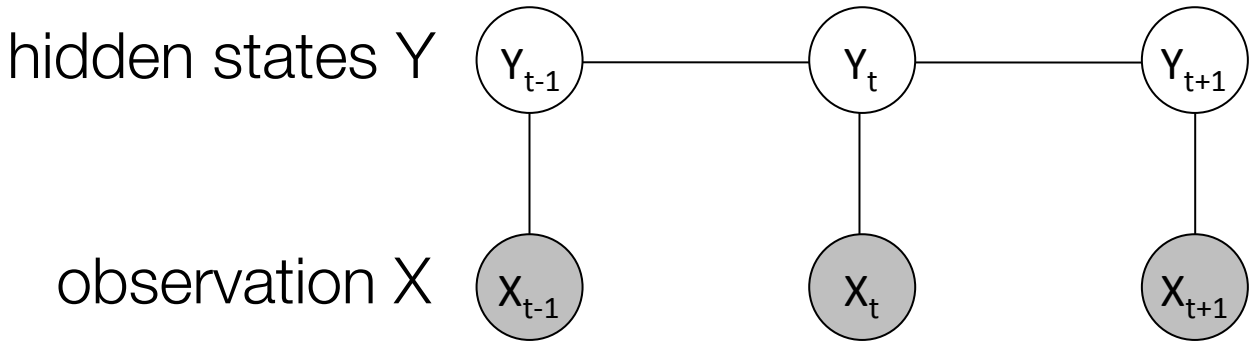

conditional model $p(\mathbf{y}|\mathbf{x})$:

$f(y_t, y_{t-1}, \mathbf{x}_t)$: feature functions

θ : feature weights (to learn)

JOHNS HOPKINS UNIVERSITY

NLP is full of features

does question starts with “what”, “where”, “who”, … ?

what parts of speech?

any named entities, like a PERSON or a LOCATION ?

automatic alignments between question and passage

(tree) edit distance between question and passage

…

Joint Features

Question starts with (“*who*”)

AND Passage contains entity of type (PERSON)

Question contains word (“*born*”)

AND Question starts with (“*where*”)

AND Passage contains word (“*birthplace*”)

Ranking Sentence Candidates

System	MAP	MRR
Wang et al. (2007)	0.6029	0.6852
Heilman and Smith (2010)	0.6091	0.6917
Wang and Manning (2010)	0.5951	0.6951
this paper (48 features)	0.6319	0.7270
+WNsearch	**0.6371**	0.7301
+WNfeature (11 more feat.)	0.6307	**0.7477**

JOHNS HOPKINS UNIVERSITY

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

“When was Alaska purchased?”

"When was Alaska purchased?"

JOHNS HOPKINS
UNIVERSITY

"When was Alaska purchased?"

JOHNS HOPKINS
UNIVERSITY

"When was Alaska purchased?"

Standard search query: (*alaska*, *purchased*)

JOHNS HOPKINS
UNIVERSITY

"When was Alaska purchased?"

"Eventually Alaska Airlines will allow all travelers who have purchased electronic tickets through any means …"

JOHNS HOPKINS UNIVERSITY

Expected Answer Type

work	question/answer types	domain
Harabagiu et al. (2000)	27 named entity labels with many-to-many mapping to answer types of 15 top nodes	TREC
Hovy et al. (2000, 2002)	topology of 94 nodes (47 leaf nodes) analyzing 17,384 questions	answers.com
Hermjakob (2001)	122 question targets from above questions	answers.com
Pasca and Harabagiu (2001)	taxonomy connecting 153 WordNet sub-hierarchies	TREC
Li and Roth (2002)	6 coarse classes and 50 fine classes	TREC
Prager et al. (2006)	flat hierarchy with 80 semantic classes	TREC

Expected Answer Type

work	question/answer types	domain
Harabagiu et al. (2000)	27 named entity labels with many-to-many mapping to answer types of 15 top nodes	TREC
Hovy et al. (2000, 2002)	topology of 94 nodes (47 leaf nodes) analyzing 17,384 questions	answers.com
Hermjakob (2001)	122 question targets from above questions	answers.com
Pasca and Harabagiu (2001)	taxonomy connecting 153 WordNet sub-hierarchies	TREC
Li and Roth (2002)	6 coarse classes and 50 fine classes	TREC
Prager et al. (2006)	flat hierarchy with 80 semantic classes	TREC

"When was Alaska purchased?"

Important learned QA feature

question starts with ("when")
AND
passage contains entity of type(DATE)

"When was Alaska purchased?"

Standard search query: (*alaska*, *purchased*)

Improved query: (*alaska*, *purchased*) & contains(DATE)

"When was Alaska purchased?"

Improved query: (*alaska*, *purchased*) & contains(DATE)

Results

improved search (17% relative MRR)

improved extraction (20% relative F1 for top candidate)

JOHNS HOPKINS UNIVERSITY

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

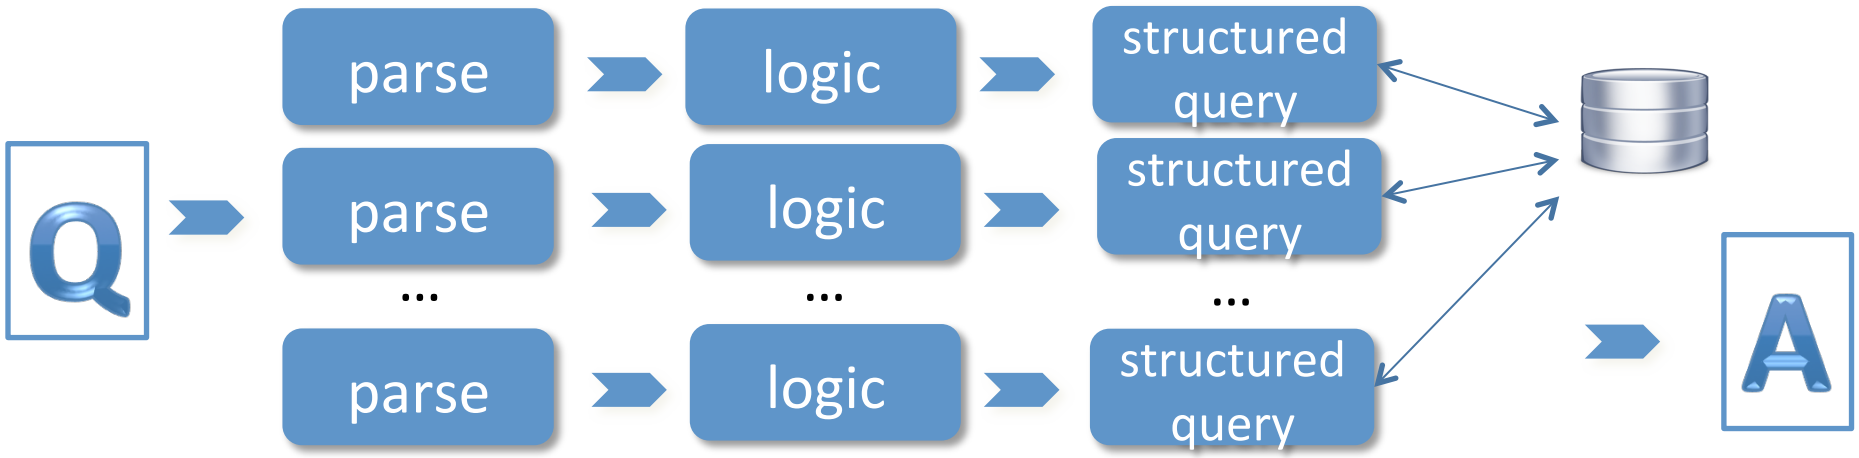
Q
parse
parse
...
parse
logic
logic
...
logic
structured query
structured query
...
structured query
A

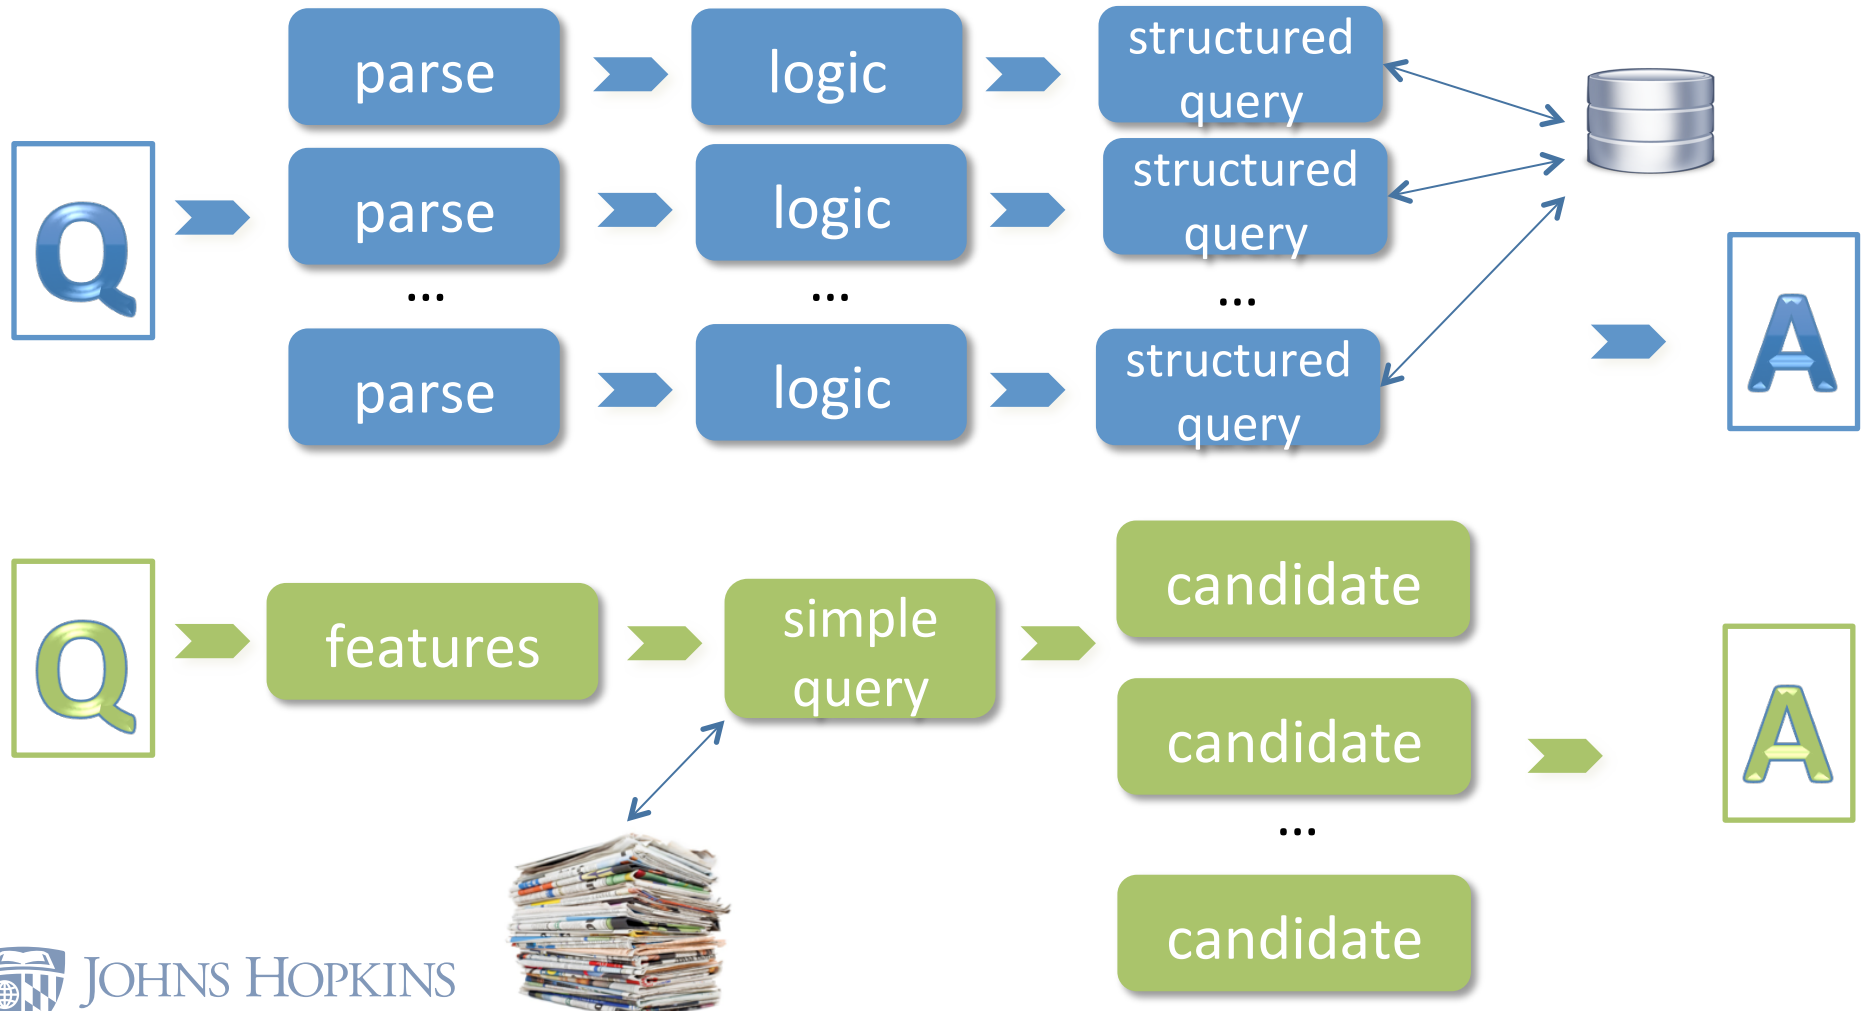

JOHNS HOPKINS UNIVERSITY

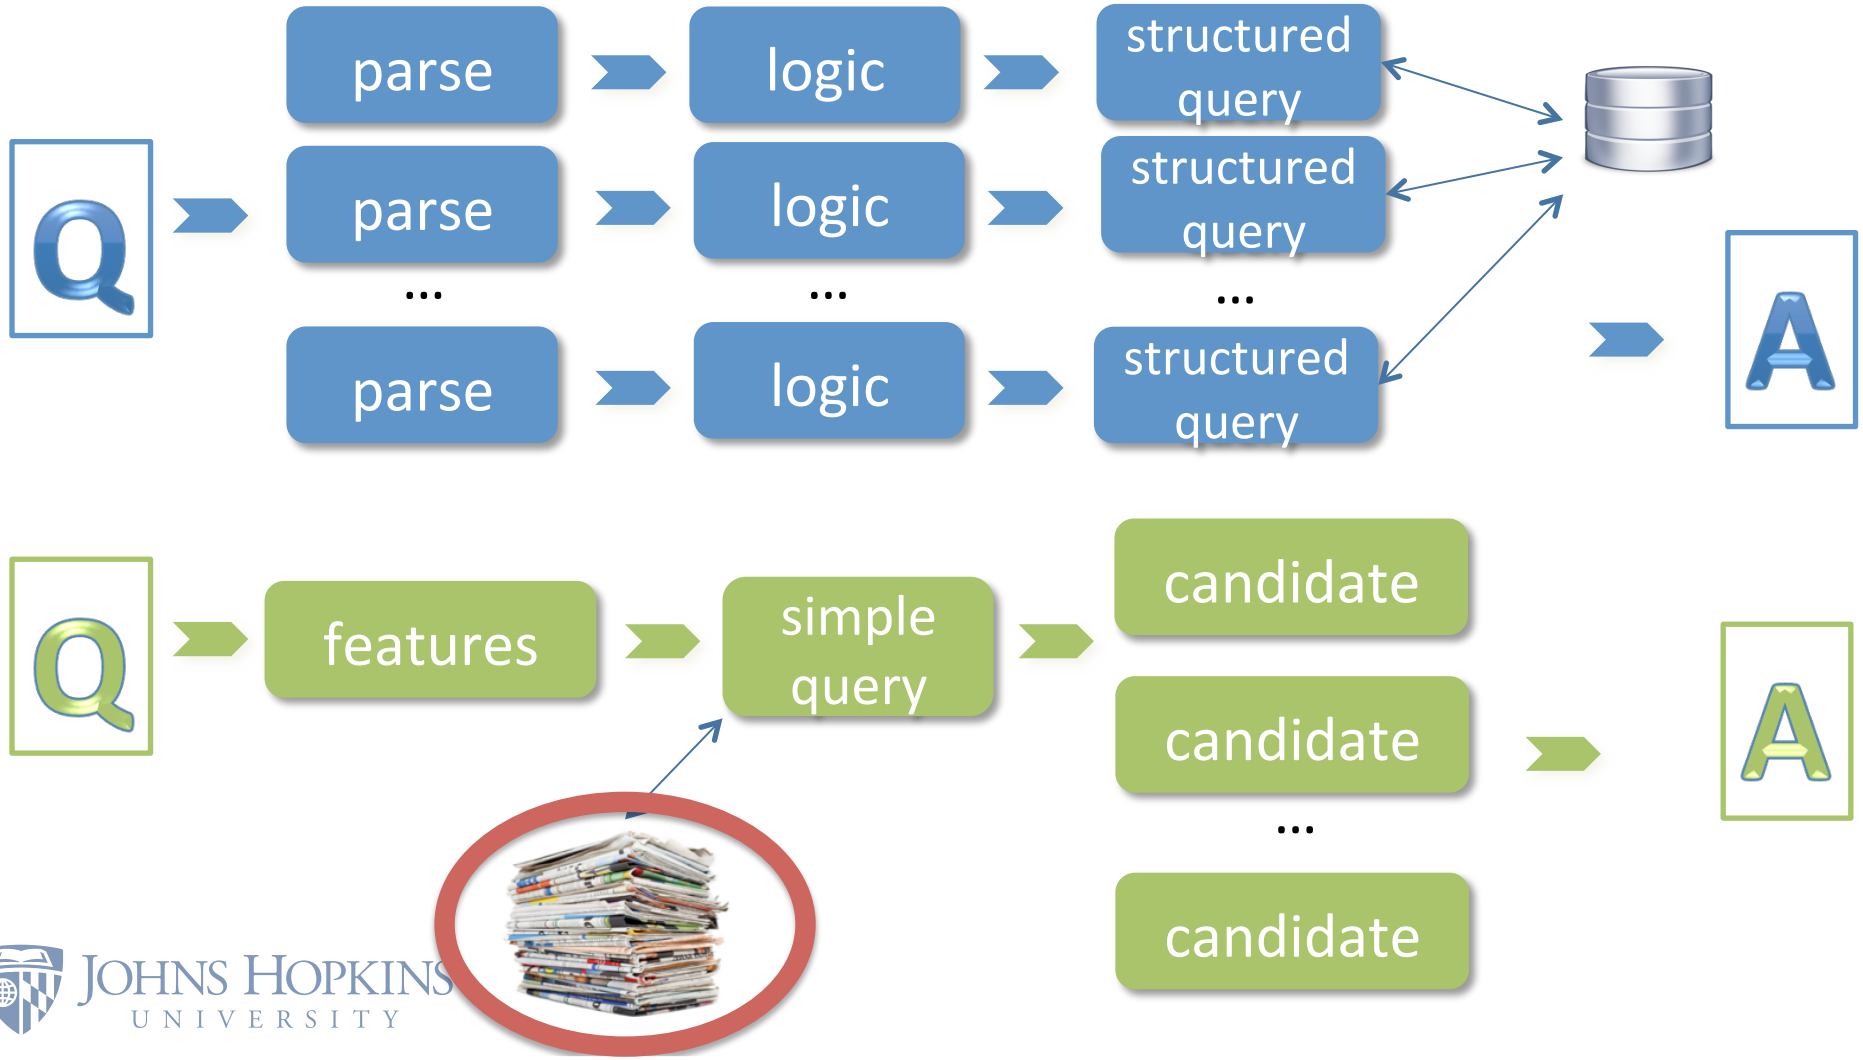
Q
parse
parse
...
parse
logic
logic
...
logic
structured query
structured query
...
structured query
A
Q
features
simple query
candidate
candidate
...
candidate
A
JOHNS HOPKINS
UNIVERSITY

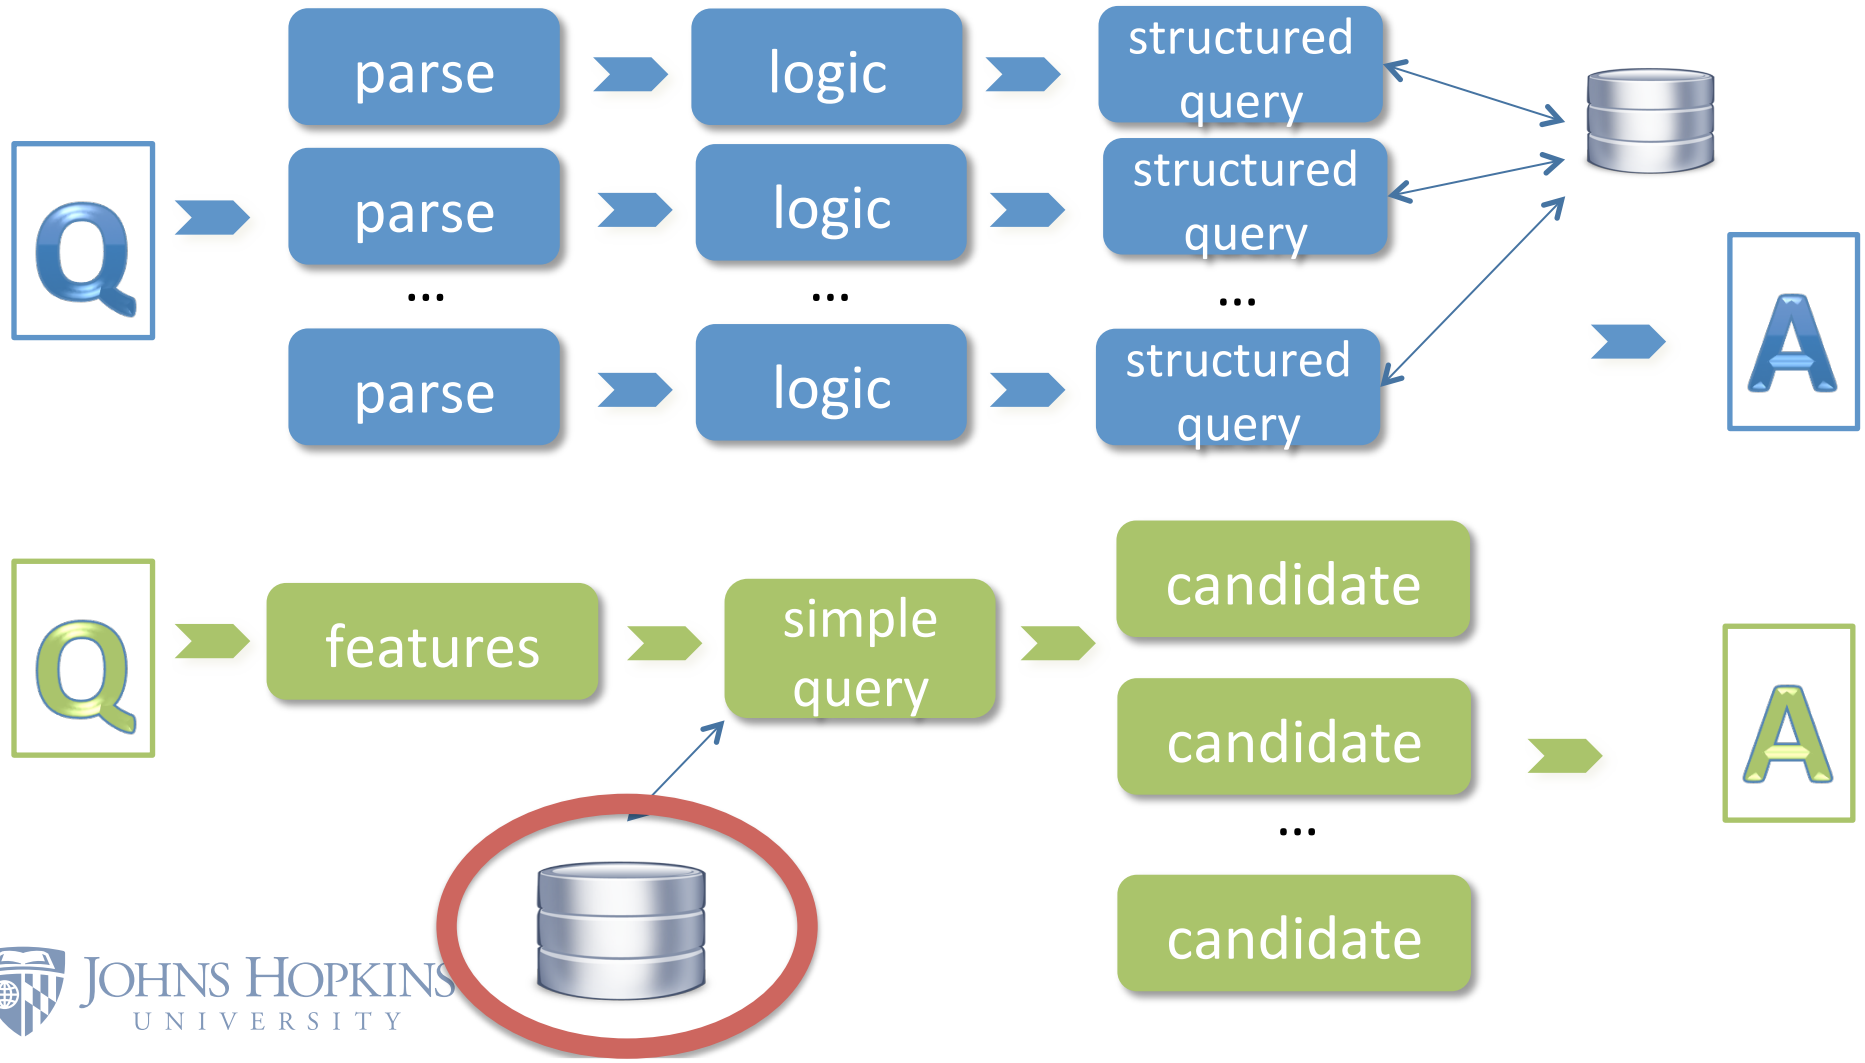

JOHNS HOPKINS UNIVERSITY

Question Answering on FreeBase

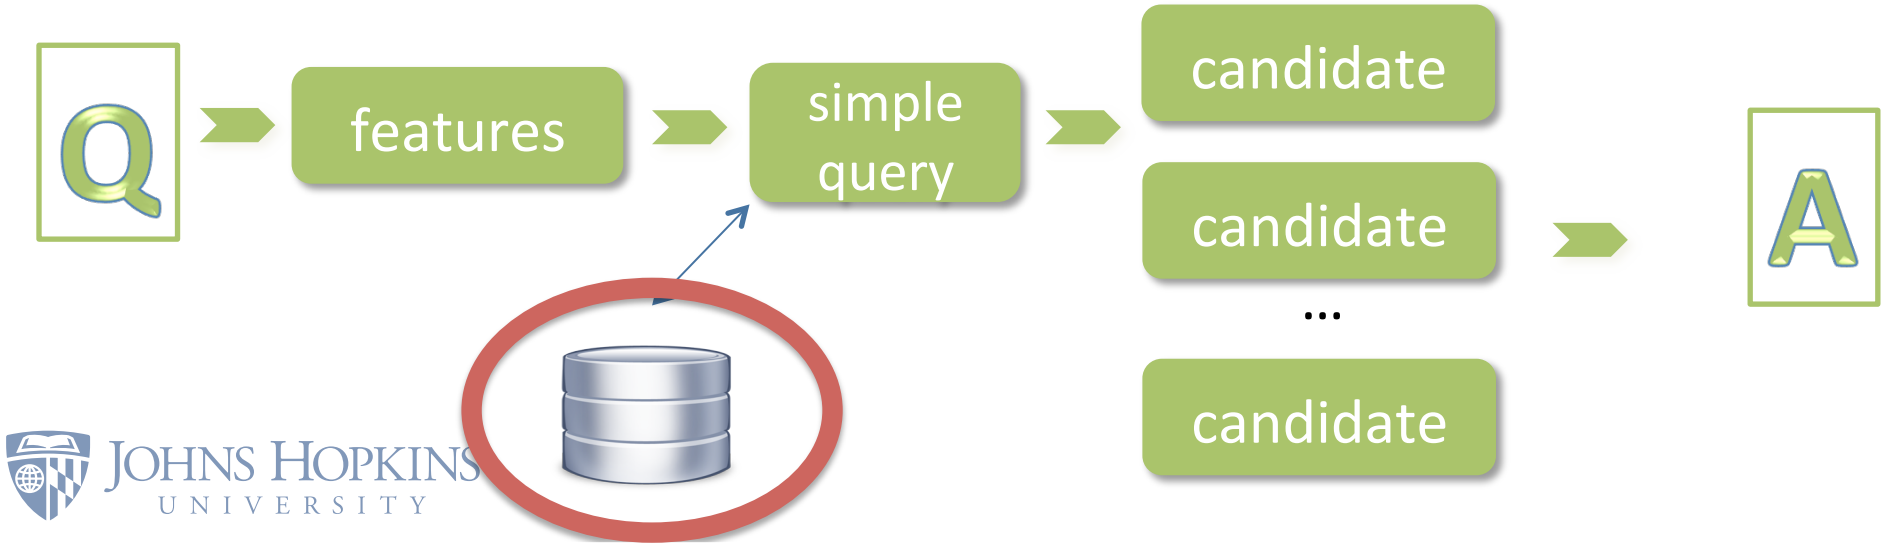

JOHNS HOPKINS UNIVERSITY

who is the brother of Justin Bieber?

who is the brother of Justin Bieber?

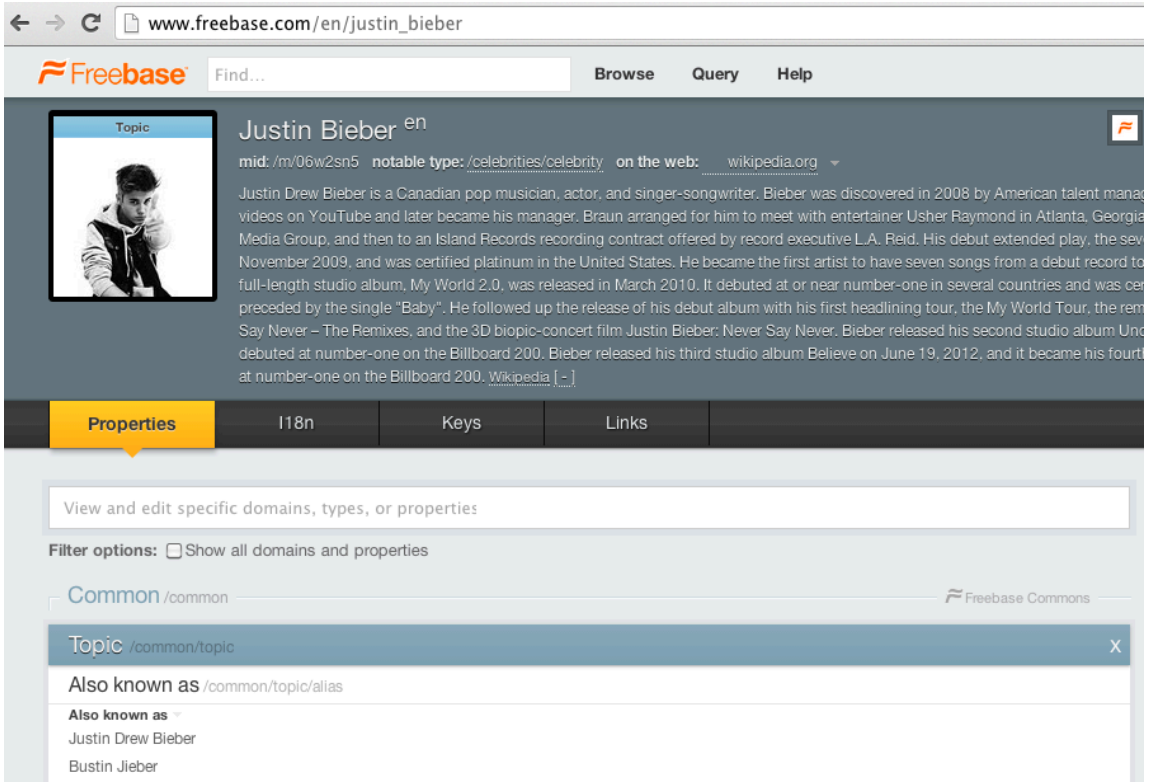

who is the brother of Justin Bieber?

JOHNS HOPKINS UNIVERSITY

who is the brother of Justin Bieber?

JOHNS HOPKINS
UNIVERSITY

features

who is the brother of Justin Bieber?

JOHNS HOPKINS
UNIVERSITY

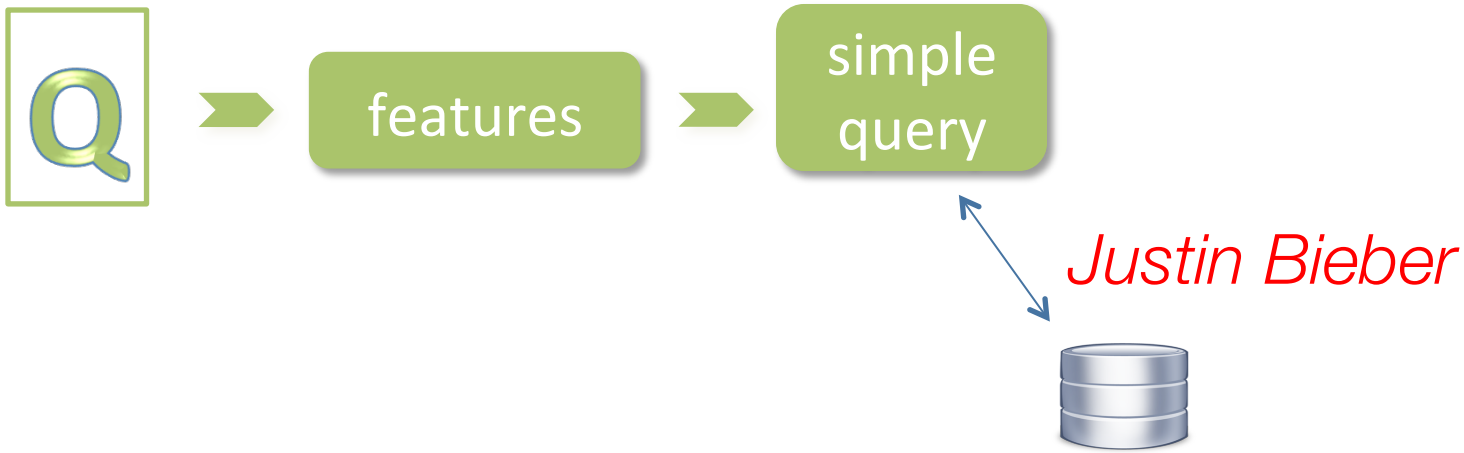
Q
features
simple query
Justin Bieber

Freebase Topic Graph

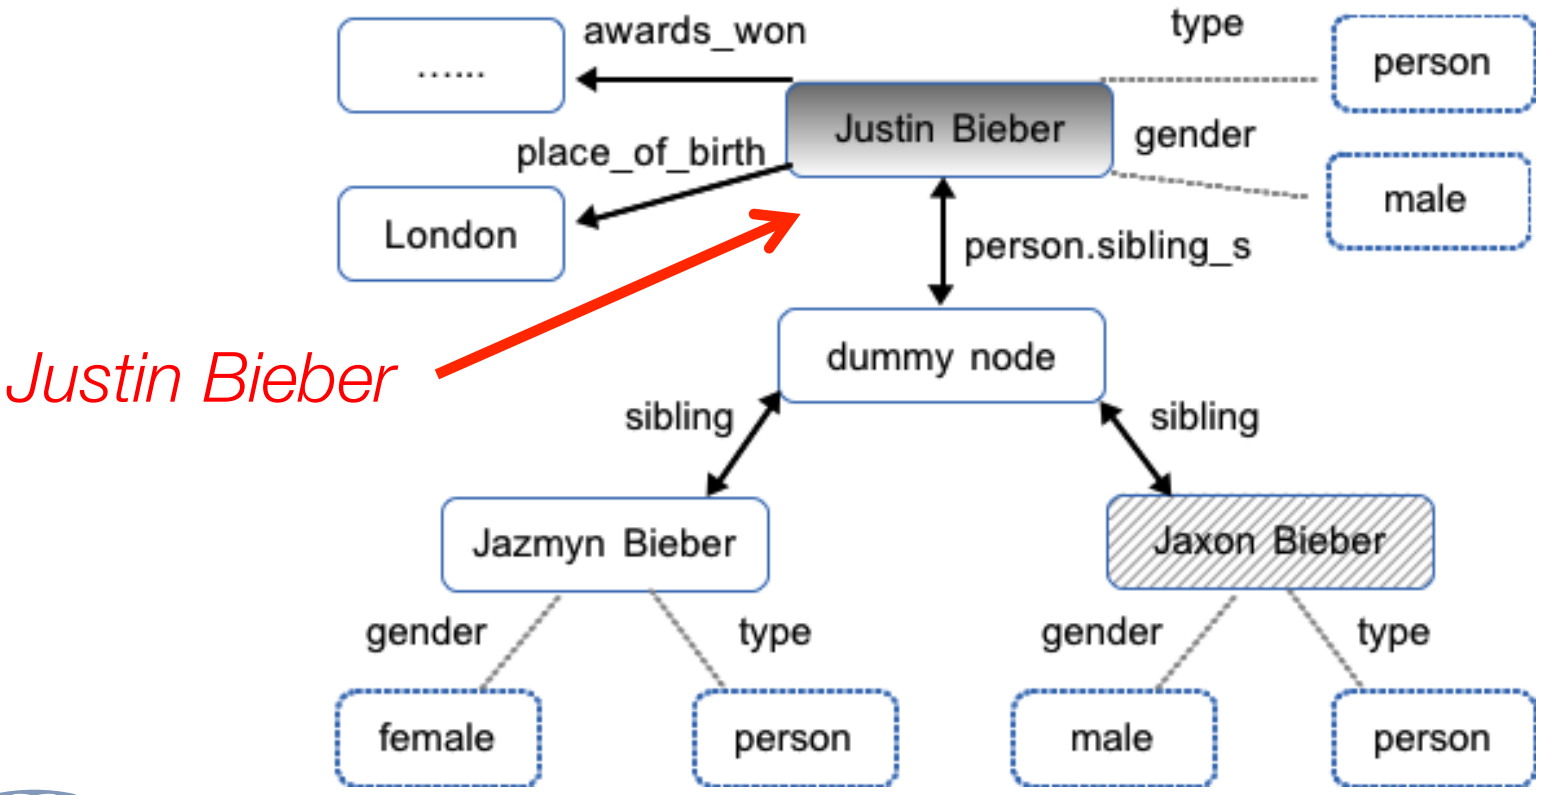

Freebase Topic Graph

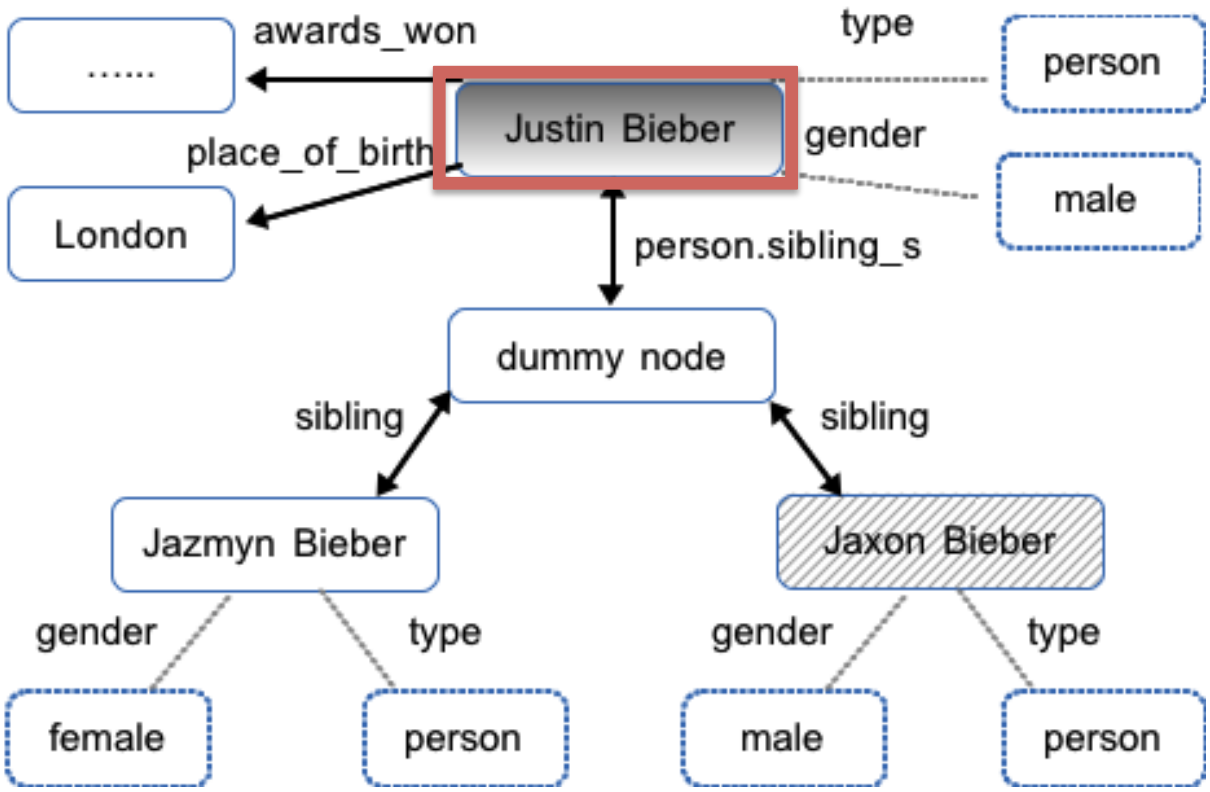

Freebase Topic Graph

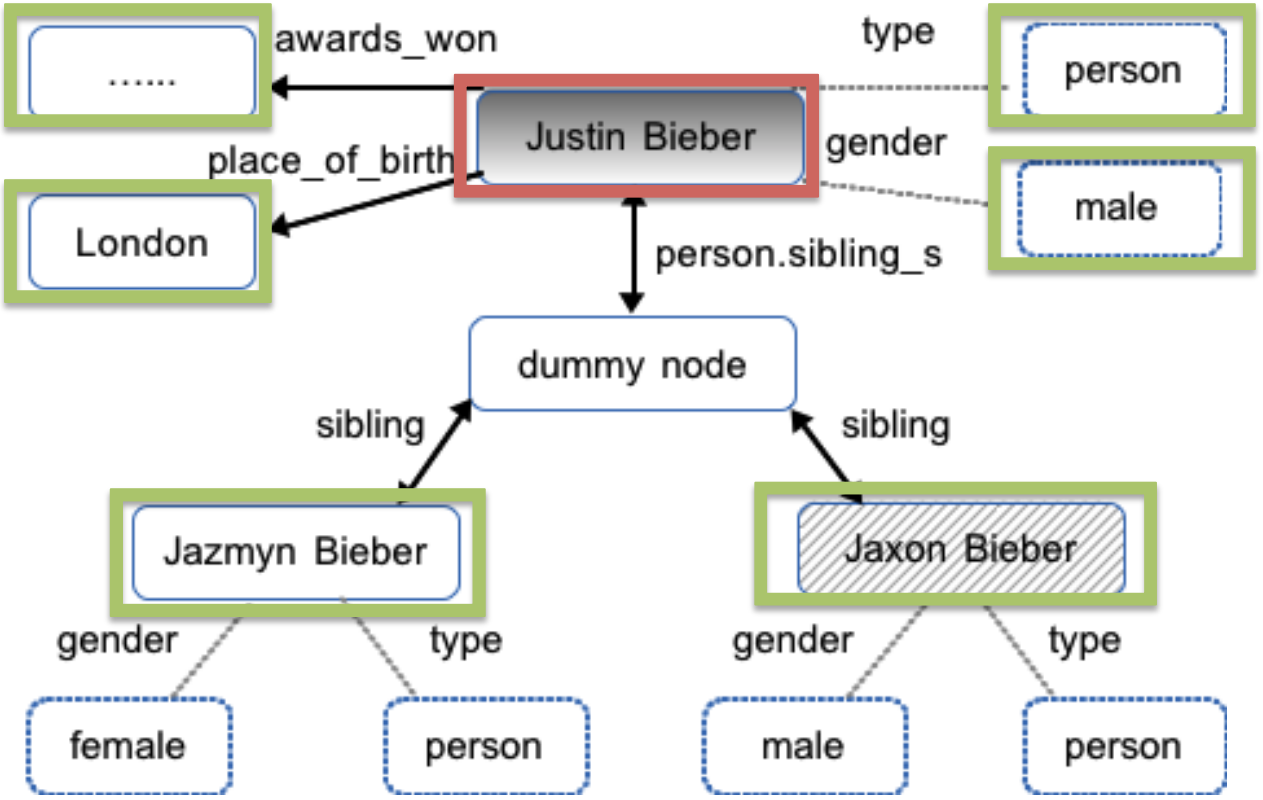

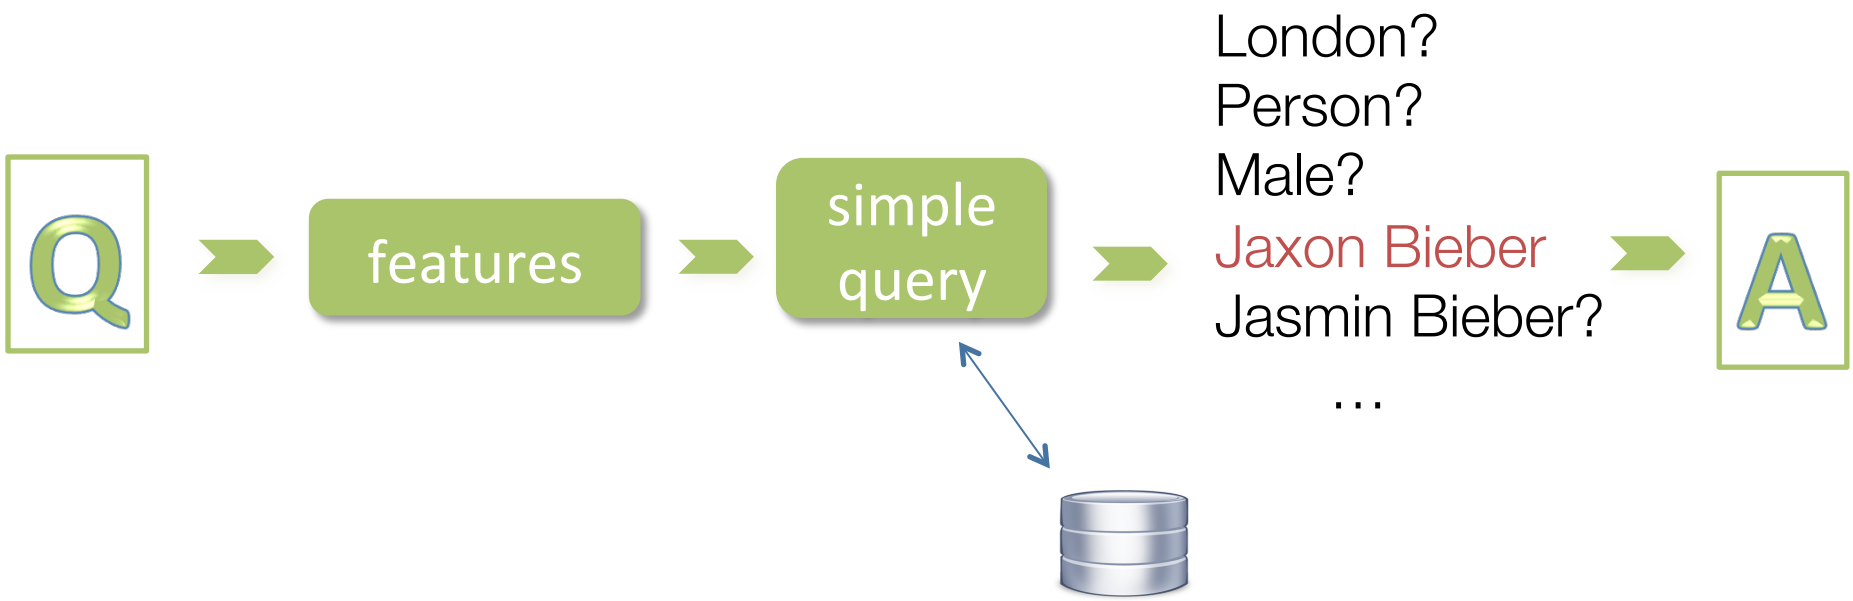

JOHNS HOPKINS UNIVERSITY

Features on Graph

extract features for *each* node

Justin Bieber	**Jazmyn Bieber**	**Jaxon Bieber**
has:awards_won	has:sibling	has:sibling
has:place_of_birth	gender:female	gender:male
type:person	type:person	type:person
...	...	...

Features on Question

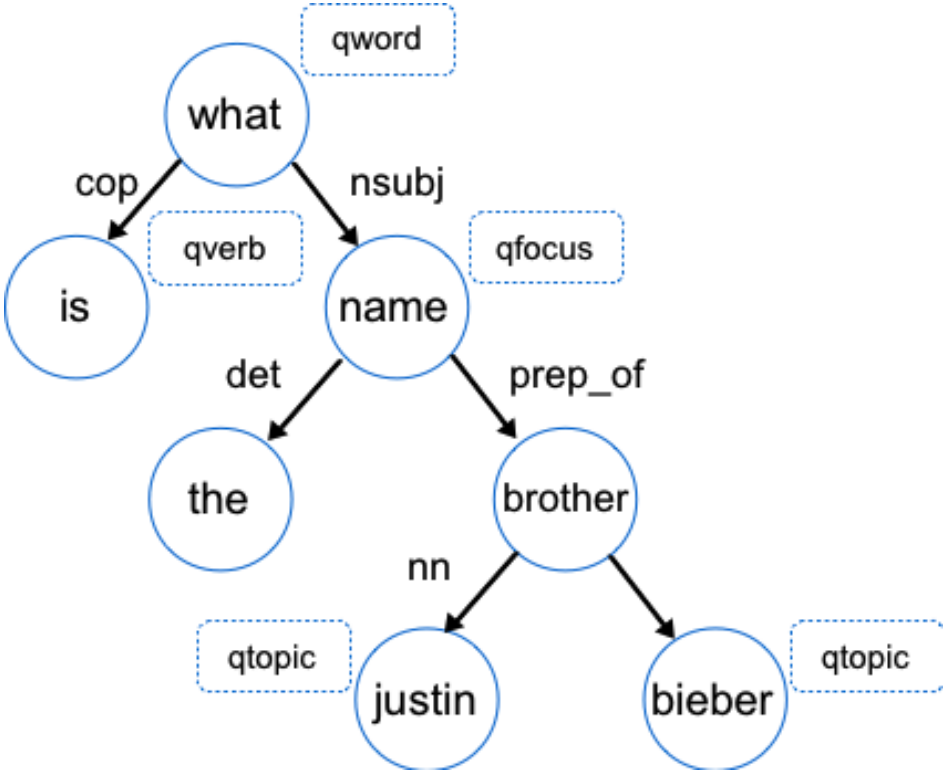

JOHNS HOPKINS UNIVERSITY

Joint Features

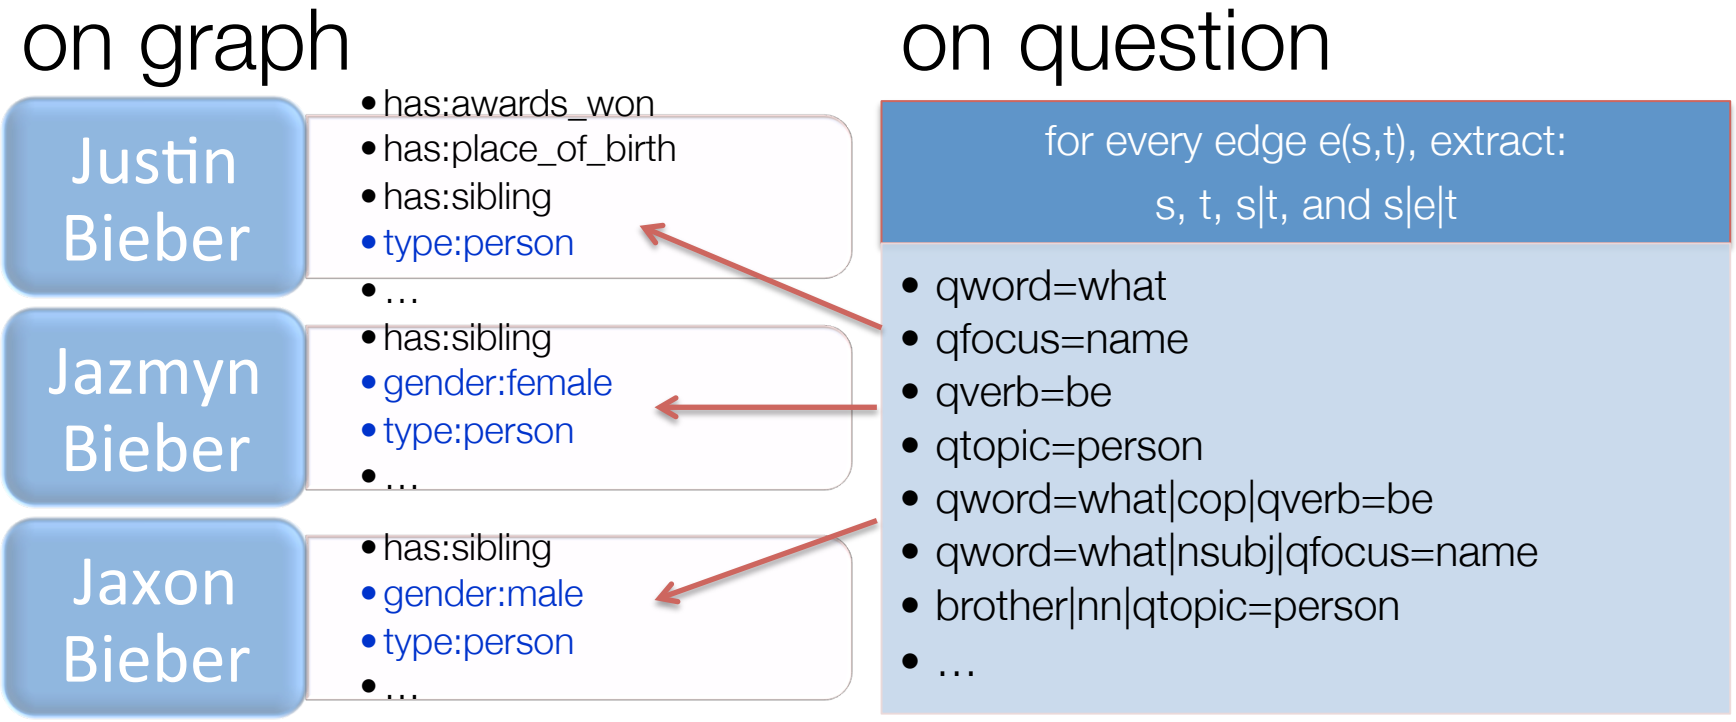

Evaluation

Data: WebQuestions

Berant, et. al. (2013)

5810 questions crawled from Google Suggest

Evaluation

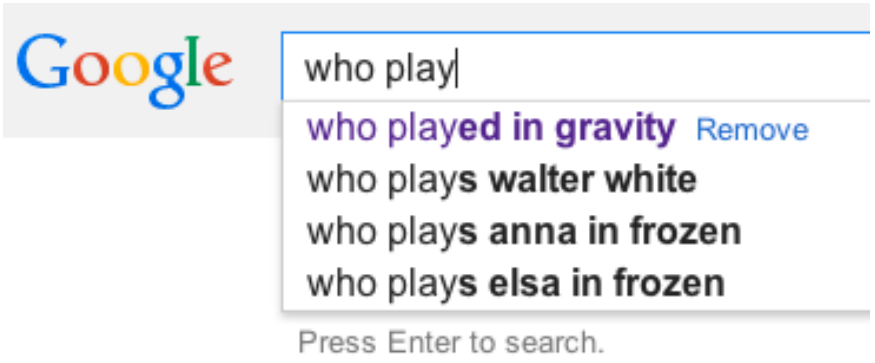

Data: WebQuestions

5810 questions crawled from Google Suggest

Berant *et al.* (2013)

Evaluation

Data: WebQuestions

5810 questions crawled from Google Suggest

Berant *et al.* (2013)

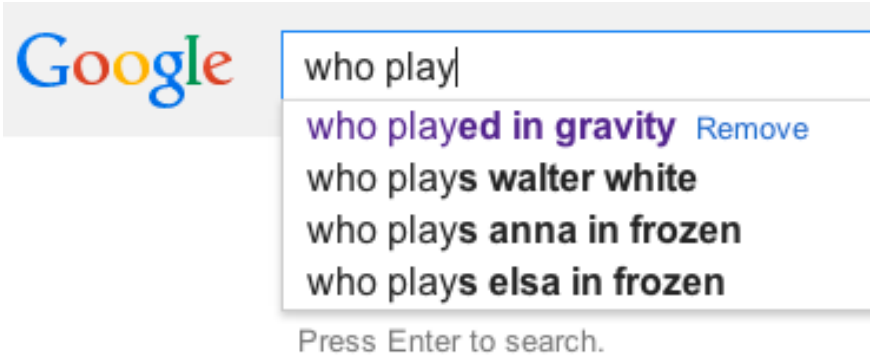

which states does the connecticut river flow through?

who does david james play for 2011?

what date was john adams elected president?

what kind of currency does cuba use?

who owns the cleveland browns?

…

Model

L1 regularized Logistic Regression

high-performance toolkit: Classias (Okazaki, 2009)

Start with: 7 million feature types extracted

Train for: 4 hours

Result: 30 thousand features with non-zero weight

Our Results

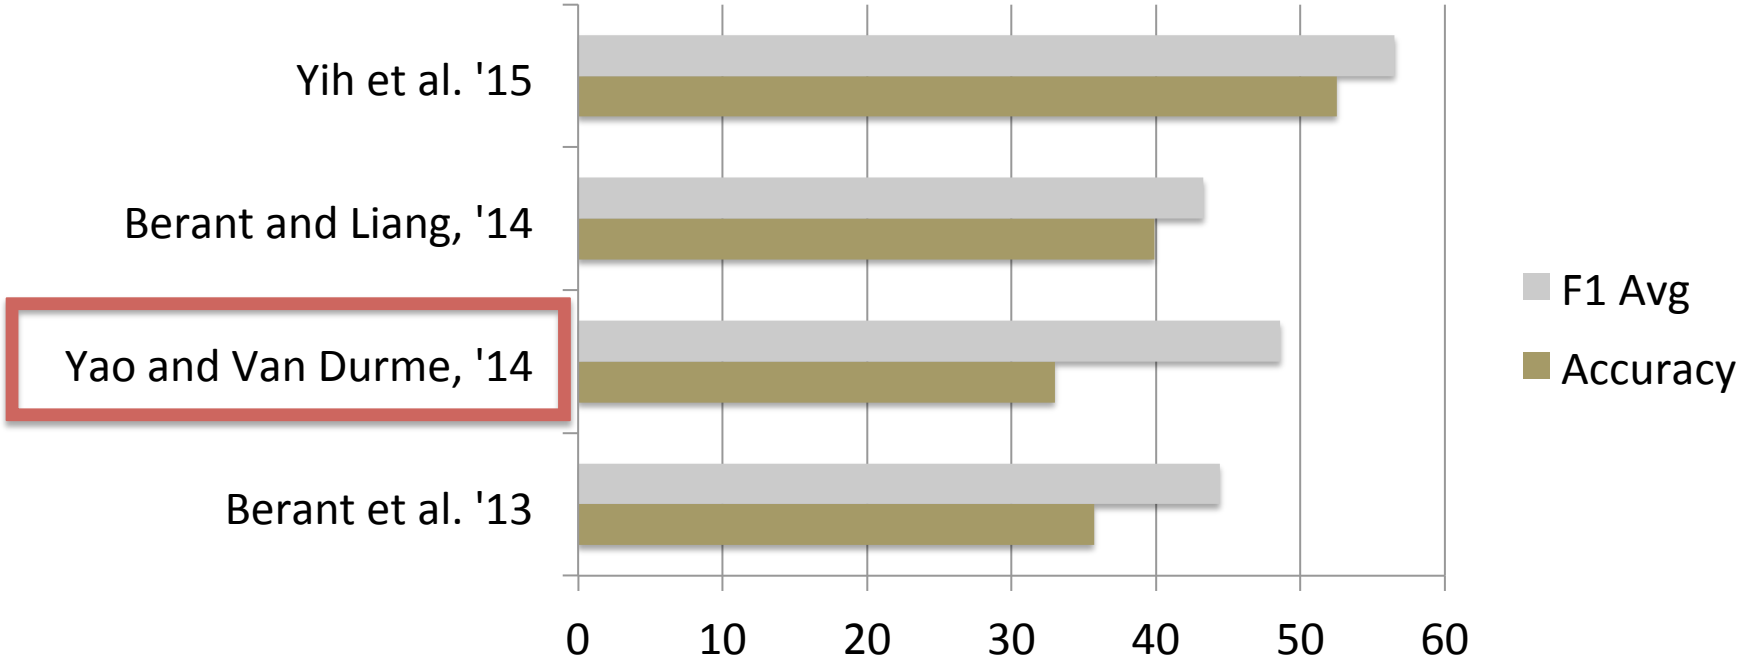

JOHNS HOPKINS UNIVERSITY

Stanford

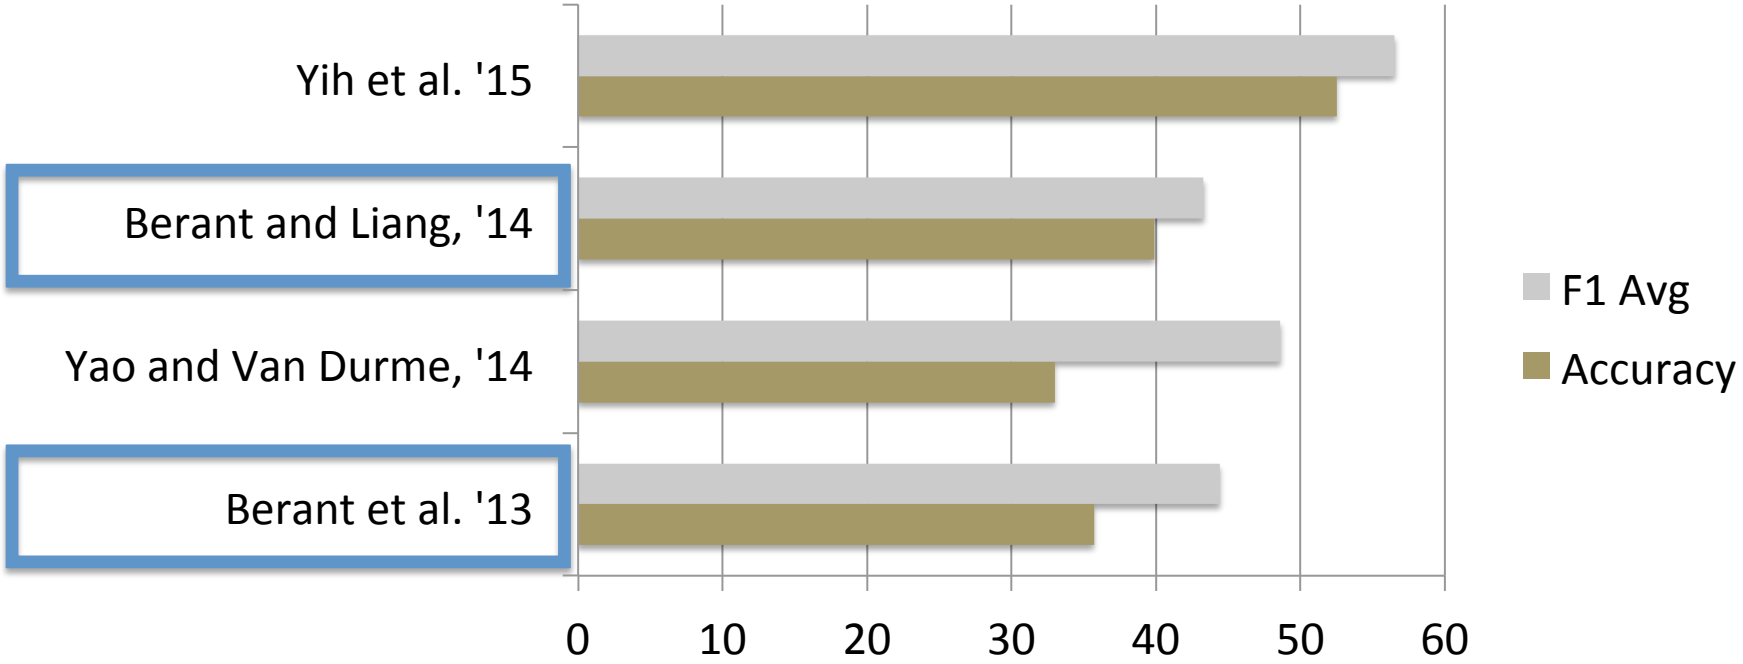

JOHNS HOPKINS UNIVERSITY

Microsoft Research

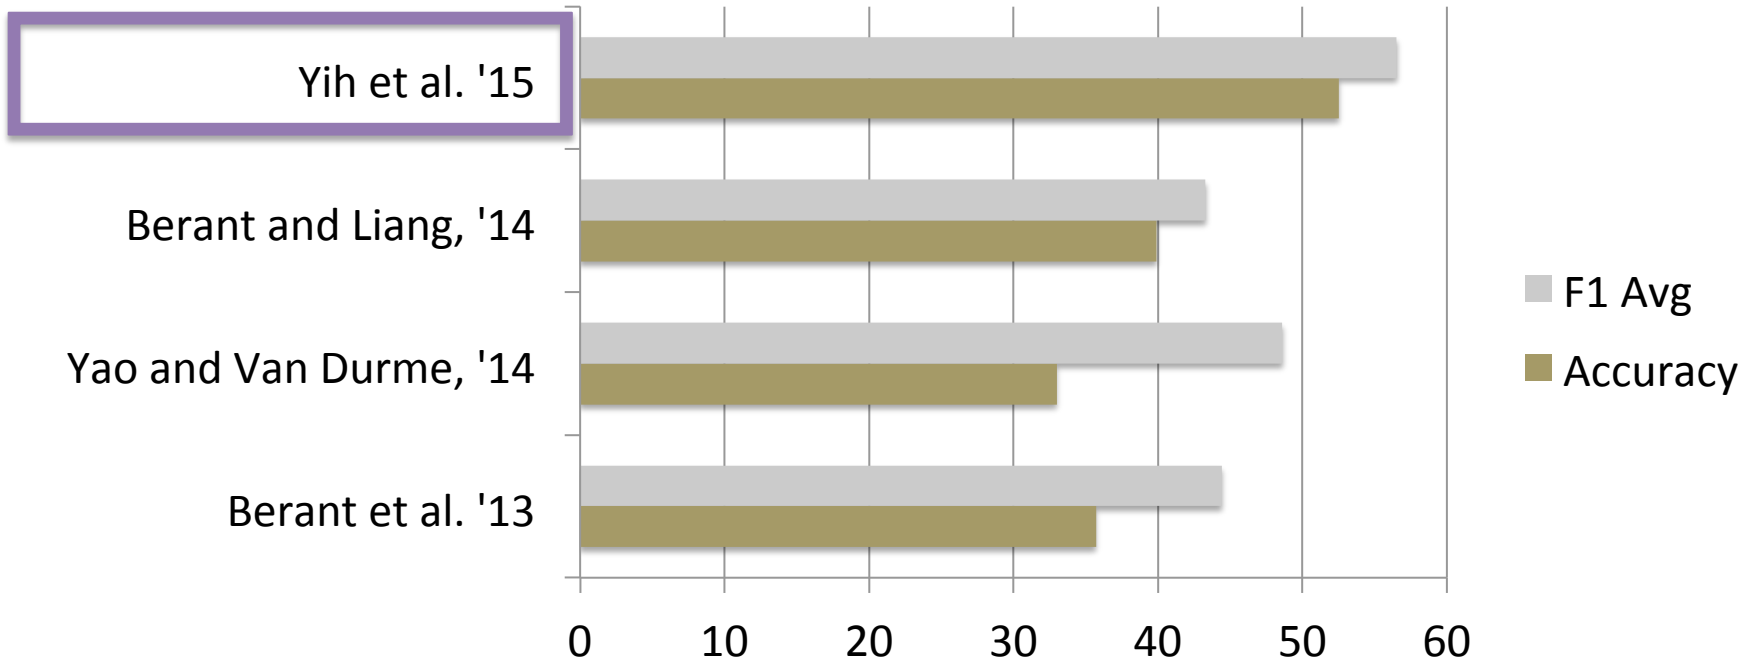

JOHNS HOPKINS UNIVERSITY

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

Feature driven answer extraction

Feature driven retrieval

Feature driven QA on a KB

This work supported by …

Google

AI2

DARPA

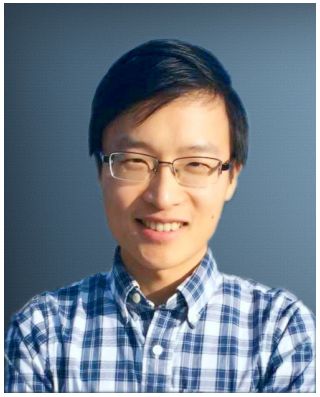

JOHNS HOPKINS UNIVERSITY