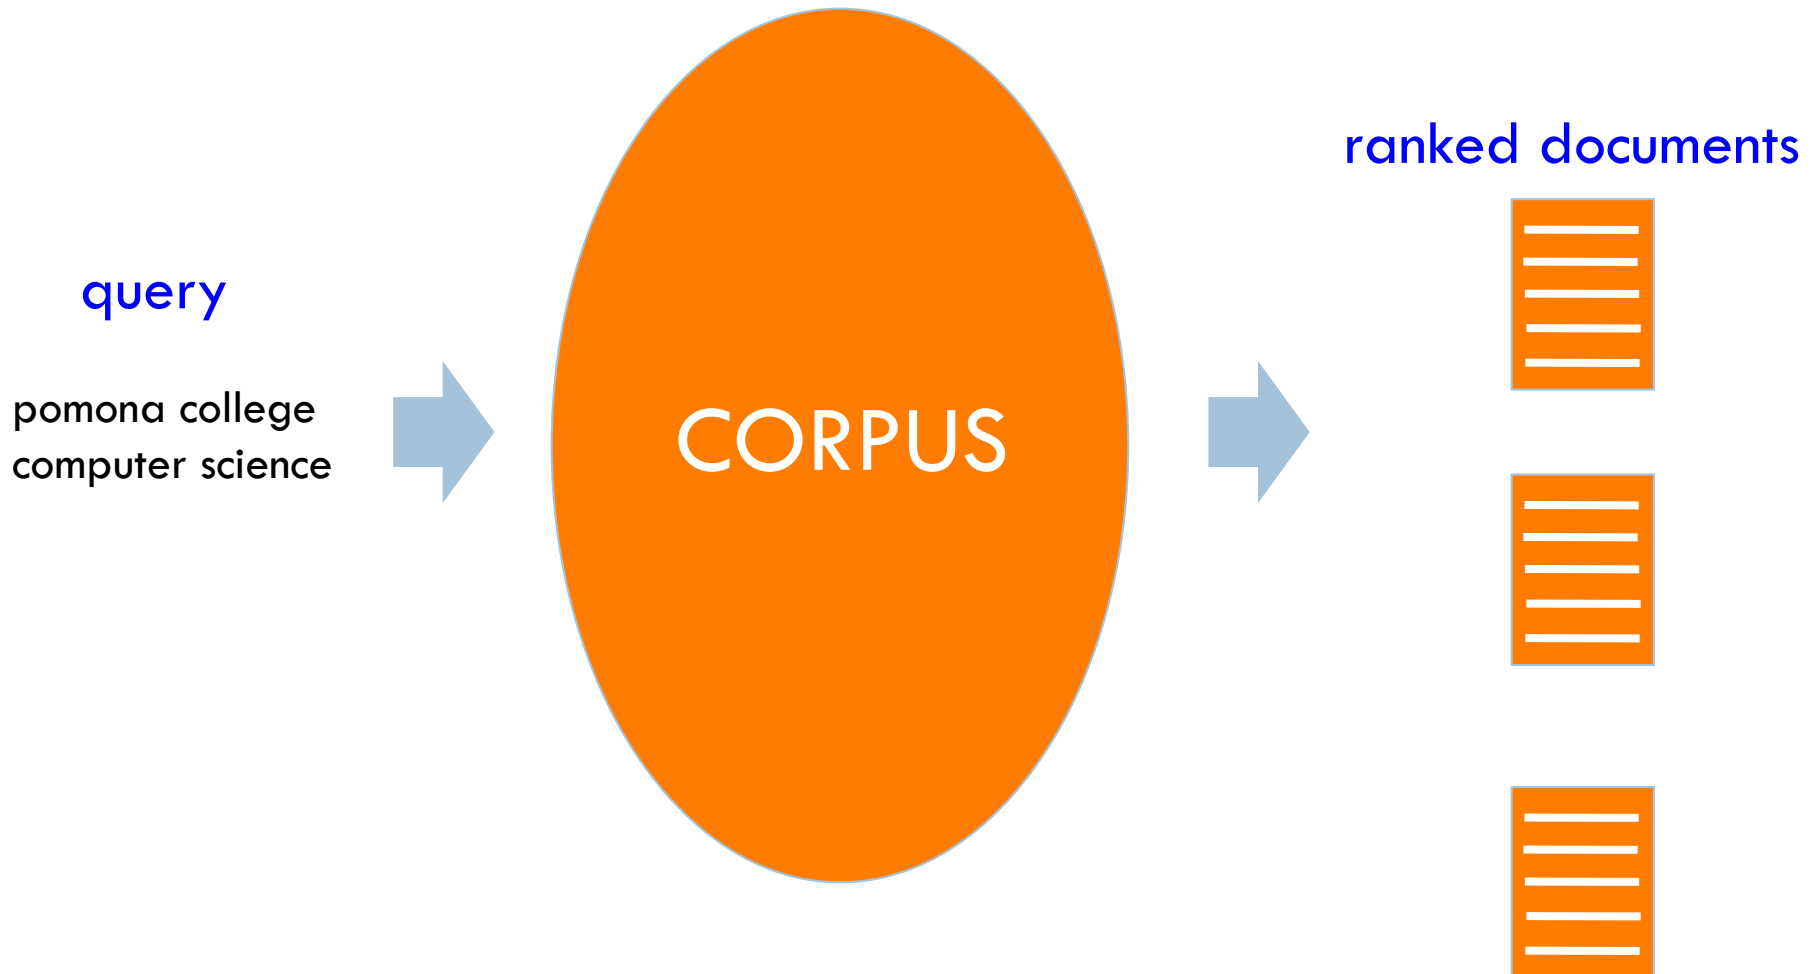


QUESTION ANSWERING

Günter Neumann, LT-Lab, DFKI


Information retrieval





Question answering

question

where is pomona college located?



answer

 **Pomona College** 
[Place page](#)

550 North College Avenue
Claremont, CA 91711-6301
(909) 621-8000

Train: [Claremont Metrolink Station](#)
[Get directions](#) - [Is this accurate?](#)

★★★★☆ 184 reviews - [Write a review](#)

"Claremont is absolutely beautiful. Come mid-January, you'll be back to the ..."

- [collegeprowler.com](#)



Question answering

Number of Documents:	979,000
Megabytes of Text:	3033
Document Sources:	AP, WSJ, Financial Times, San Jose Mercury News, LA Times, RBIS
Number of Questions:	682
Question Sources:	Encarta log, Excite log

Sample questions:

- How much folic acid should an expectant mother get daily?
- Who invented the paper clip?
- What university was Woodrow Wilson president of?
- Where is Rider College located?
- Name a film in which Jude law acted.
- Where do lobsters like to live?

Other variants...

Mars Polar Lander- Where Are You?

(January 18, 2000) After more than a month of searching for a single sign from NASA's Mars Polar Lander, mission controllers have lost hope of finding it. The Mars Polar Lander was on a mission to Mars to study its atmosphere and search for water, something that could help scientists determine whether life even existed on Mars. Polar Lander was to have touched down on December 3 for a 90-day mission. It was to land near Mars' south pole. The lander was last heard from minutes before beginning its descent. The last effort to communicate with the three-legged lander ended with frustration at 8 a.m. Monday. "We didn't see anything," said Richard Cook, the spacecraft's project manager at NASA's Jet Propulsion laboratory. The failed mission to the Red Planet cost the American government more the \$200 million dollars. Now, space agency scientists and engineers will try to find out what could have gone wrong. They do not want to make the same mistakes in the next mission.

- ❑ When did the mission controllers lost Hope of communication with the Lander?
- ❑ Who is the Polar Lander's project manager?
- ❑ Where on Mars was the spacecraft supposed to touch down?
- ❑ What was the mission of the Mars Polar Lander?

Other variants...

Mars Polar Lander- Where Are You?

(January 18, 2000) After more than a month of searching for a single sign from NASA's Mars Polar Lander, mission controllers have lost hope of finding it. The Mars Polar Lander was on a mission to Mars **to study its atmosphere and search for water**, something that could help scientists determine whether life even existed on Mars. Polar Lander was to have touched down on December 3 for a 90-day mission. It was to land **near Mars' south pole**. The lander was last heard from minutes before beginning its descent. The last effort to communicate with the three-legged lander ended with frustration at **8 a.m. Monday**. "We didn't see anything," said **Richard Cook, the spacecraft's project manager** at NASA's Jet Propulsion laboratory. The failed mission to the Red Planet cost the American government more the \$200 million dollars. Now, space agency scientists and engineers will try to find out what could have gone wrong. They do not want to make the same mistakes in the next mission.


- ❑ When did the mission controllers lost Hope of communication with the Lander?
(Answer: **8AM, Monday Jan. 17**)
- ❑ Who is the Polar Lander's project manager?
(Answer: **Richard Cook**)
- ❑ Where on Mars was the spacecraft supposed to touch down?
(Answer: **near Mars' south pole**)
- ❑ What was the mission of the Mars Polar Lander?
(Answer: **to study Mars' atmosphere and search for water**)

Question answering



challenges? how is this different than IR?

question

where is pomona college located?



answer

 **Pomona College** 
[Place page](#)

550 North College Avenue
Claremont, CA 91711-6301
(909) 621-8000
Train: [Claremont Metrolink Station](#)
[Get directions](#) - [Is this accurate?](#)

★★★★☆ 184 reviews - [Write a review](#)
"Claremont is absolutely beautiful. Come mid-January, you'll be back to the ..."
- collegeprowler.com



Challenges



Question: “When was Wendy’s founded?”

Challenges

Question: “When was Wendy’s founded?”

Passage candidate:

“The renowned Murano glassmaking industry, on an island in the Venetian lagoon, has gone through several reincarnations since it was founded in 1291. Three exhibitions of 20th-century Murano glass are coming up in New York. By Wendy Moonan.”

Answer: 20th Century or 1291?

Wendy's

Fast food company

Wendy's is an international fast food chain restaurant founded by Dave Thomas on November 15, 1969, in Columbus, Ohio, United States. The company moved its headquarters to Dublin, Ohio, on January 29, 2006. [Wikipedia](#)

Founded: 1969
Founders: [John T. Schuessler](#), [Dave Thomas](#)

Recent posts on Google+

 **Wendy's**
538,557 followers • Shared publicly Follow

 Most buns can only dream of holding a sandwich like this. 12 hours ago

Challenges

Question: *When was Microsoft established?*

Challenges

Question: *When was Microsoft established?*

Microsoft tends to establish lots of things...

Microsoft plans to establish manufacturing partnerships in Brazil and Mexico in May.

Solution?

Challenges

Question: *When was Microsoft established?*

Microsoft tends to establish lots of things...

Microsoft plans to establish manufacturing partnerships in Brazil and Mexico in May.

Need to be able to detect sentences in which `Microsoft' is **object** of `establish' or close synonym.

Matching sentence:

Microsoft Corp was founded in the US in 1975, incorporated in 1981, and established in the UK in 1982.

Challenges

Question: *What is the occupation of Bill Clinton's wife?*

No documents contain these keywords plus the answer

Ideas?

Challenges

Question: *What is the occupation of Bill Clinton's wife?*

No documents contain these keywords plus the answer

One solution: decompose into two questions:

Who is Bill Clinton's wife? = X

What is the occupation of X?



AI/planning/logic!

Relevance of Answer Source

- The structure of the answer source might effect the complexity of the whole QA process.

"Where did Bill Gates go to college?"

"What is the rainiest place on Earth?"



"Harvard"

"Mount Waialeale"

"...Bill Gates, Harvard dropout and founder of Microsoft..." (Trec-Data)

"... In misty Seattle, Wash., last year, 32 inches of rain fell. Hong Kong gets about 80 inches a year, and even Pago Pago, noted for its prodigious showers, gets only about 196 inches annually. (The titleholder, according to the National Geographic Society, is Mount Waialeale in Hawaii, where about 460 inches of rain falls each year.) ..." (Trec-Data; but see [Google-retrieved Web page.](#))

Question answering

How are we going to do it?

Information
retrieval

parsing

text
similarity

language
modeling

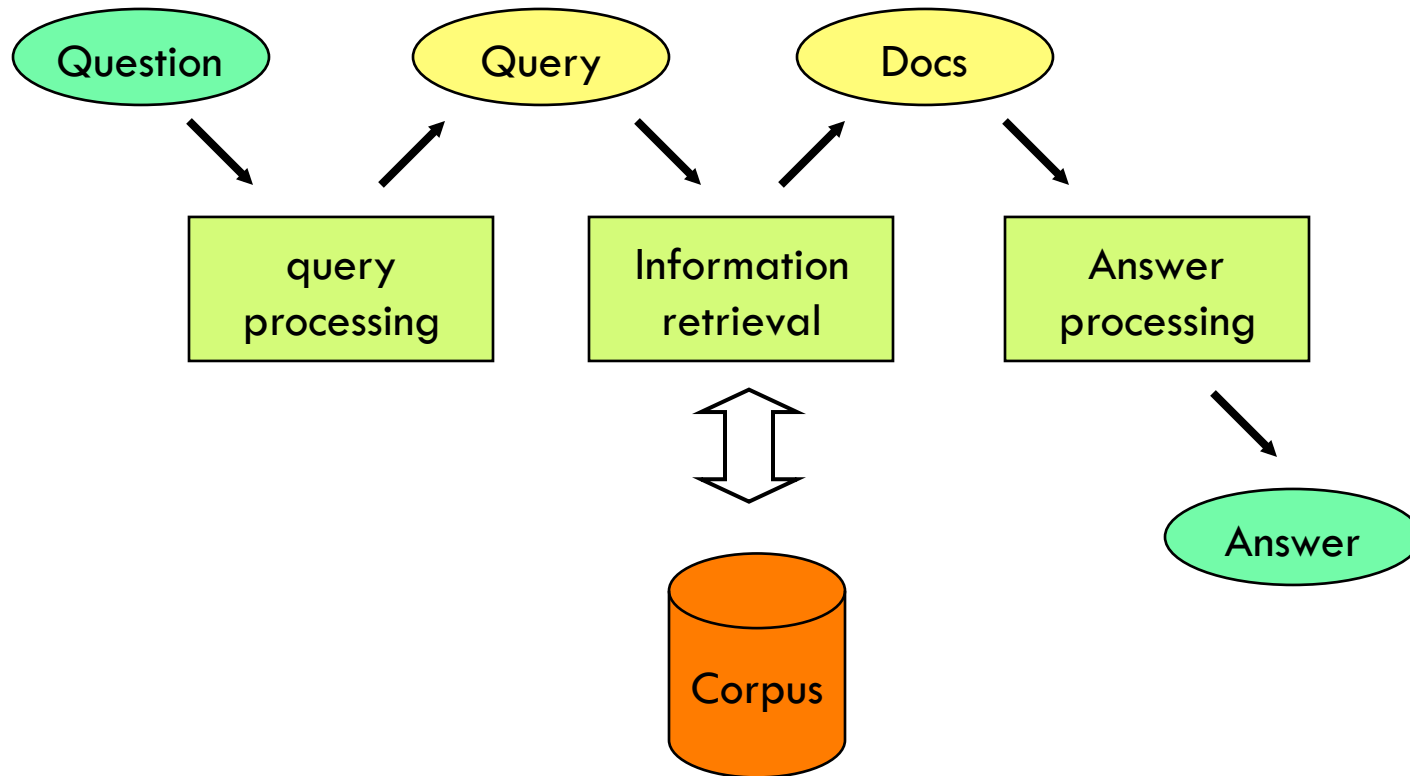
machine
learning

information
extraction

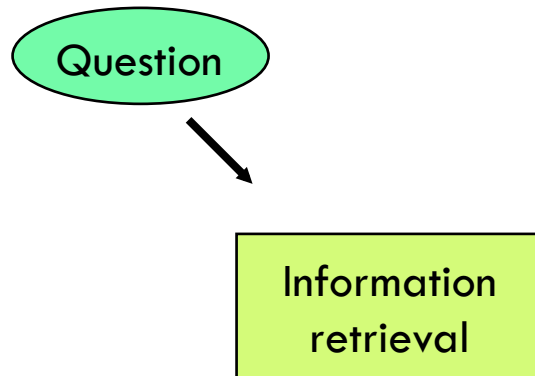
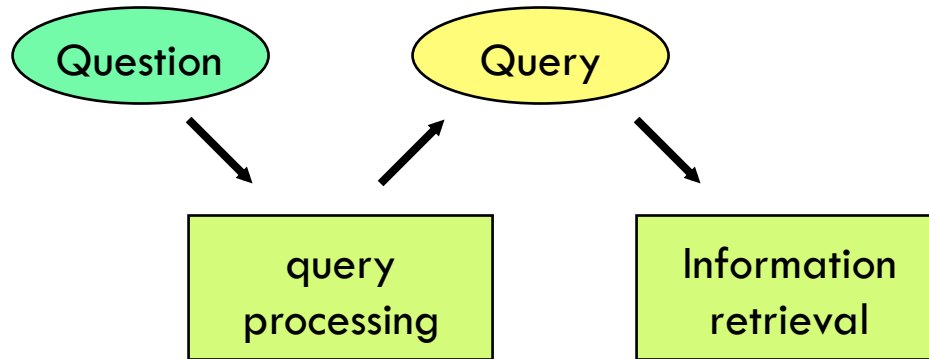
POS
tagging



QA basic framework



Query processing



Why do we need query processing?

What types of things might we need to do?

Query processing

- If we're going to feed the question into an IR system, often need to reformulate the query
 - ▣ Some things may already be done by the IR system
 - remove stopwords
 - remove question words
 - query expansion:
 - who is the president of the US?: US -> United States
 - ▣ Formulate as a statement rather than a question:

When was the paper clip invented?

the paper clip was invented

what is the tallest building?

the tallest building is

Query processing

- Analyze the query: what questions might we ask about the query that might help limit our set of answers?

Query processing

- Analyze the query: **what questions might we ask about the query that might help limit our set of answers?**
- What is the expected answer type?

	Q-Type	A-Type
When did the Titanic sink ?	event-completion	time
Who was Darth Vader's son?	concept-completion	person-name
What is thalassemia ?	definition	definition

Answer type hierarchies

Class	#	Class	#
ABBREVIATION	18	term	19
abbreviation	2	vehicle	7
expression	16	word	0
DESCRIPTION	153	HUMAN	171
definition	126	group	24
description	13	individual	140
manner	7	title	4
reason	7	description	3
ENTITY	174	LOCATION	195
animal	27	city	44
body	5	country	21
color	12	mountain	5
creative	14	other	114
currency	8	state	11
disease/medicine	3	NUMERIC	289
event	6	code	1
food	7	count	22
instrument	1	date	146
language	3	distance	38
letter	0	money	9
other	19	order	0
plant	7	other	24
product	9	period	18
religion	1	percent	7
sport	3	speed	9
substance	20	temperature	7
symbol	2	volume/size	4
technique	1	weight	4

Question classification

	Q-Type	A-Type
When did the Titanic sink ?	event-completion	time
Who was Darth Vader's son?	concept-completion	person-name
What is thalassemia ?	definition	definition

How can we identify this?

Question classification

	Q-Type	A-Type
When did the Titanic sink ?	event-completion	time
Who was Darth Vader's son?	concept-completion	person-name
What is thalassemia ?	definition	definition

□ Hand-written rules

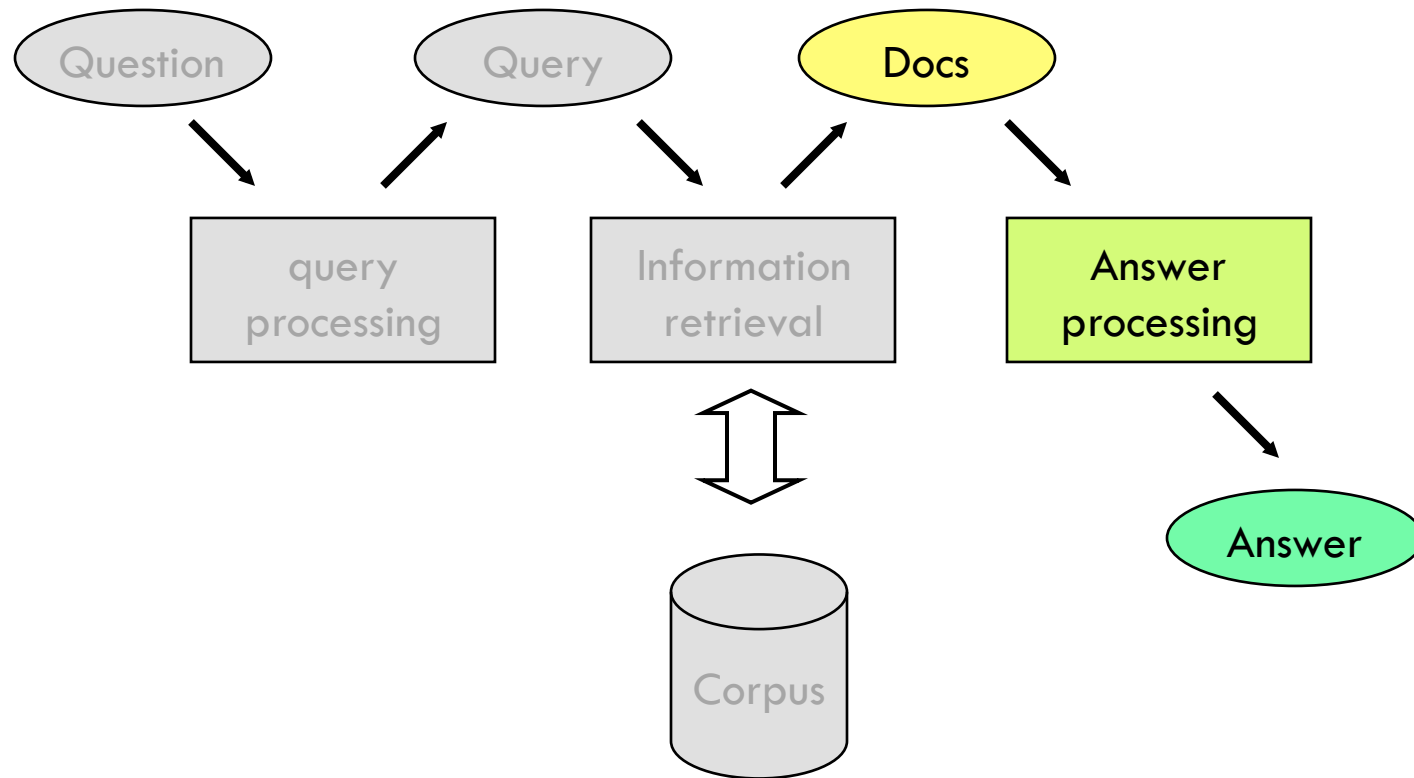
- **WHO:** person, organization, or country.
- **WHERE:** location (specific or general)
- **WHEN:** date or time period
- **HOW MUCH:** an amount
- **HOW MANY:** a number
- **WHICH CITY:** a city

Question classification

	Q-Type	A-Type
When did the Titanic sink ?	event-completion	time
Who was Darth Vader's son?	concept-completion	person-name
What is thalassemia ?	definition	definition

- Better: machine learning!
- lots of *clues*
 - What *city* is pomona college in?
 - Where is pomona college *located*?
 - What is the tallest *building* in the world?

QA basic framework



Answer processing

What is the tallest building in the world?


problem?



The *world's tallest* structure is the 828 m (2,717 ft) tall Burj Khalifa in Dubai, United Arab Emirates. The *building* gained the official title of "*Tallest Building in the World*" at its opening on 4 January 2010. It is taller than any other man-made structure ever built.



The Council on Tall *Buildings* and Urban Habitat, an organization that certifies buildings as the "*World's Tallest*", recognizes a *building* only if at least fifty percent of its height is made up of floor plates containing habitable floor area.[1] Structures that do not meet this criterion, such as the CN Tower, are defined as "towers".




There are dozens of radio and television broadcasting towers which measure over 600 meters (about 2,000 ft) in height, and only the *tallest* are recorded in publicly available information sources.

Answer processing

The *world's tallest* structure is the 828 m (2,717 ft) tall Burj Khalifa in Dubai, United Arab Emirates. The *building* gained the official title of "*Tallest Building in the World*" at its opening on 4 January 2010. It is taller than any other man-made structure ever built.

The Council on Tall *Buildings* and Urban Habitat, an organization that certifies buildings as the "*World's Tallest*", recognizes a *building* only if at least fifty percent of its height is made up of floor plates containing habitable floor area.[1] Structures that do not meet this criterion, such as the CN Tower, are defined as "towers".

There are dozens of radio and television broadcasting towers which measure over 600 meters (about 2,000 ft) in height, and only the *tallest* are recorded in publicly available information sources.



Burj Khalifa

ideas?

NLP to the rescue!

- Need more information than just the text
 - ▣ named entity recognition/information extraction
 - ▣ parsing
 - ▣ relationships between entities
 - Microsoft established vs. Microsoft was established
- Pair this information with information we extracted from the query
 - ▣ What are we looking for? A person? A date? ...?

Extracting entities

President George Bush announced a new bill that would send \$1.2 million dollars to Miami Florida for a new hurricane tracking system.

After applying a Named Entity Tagger, the text might look like this:

<Person="President George Bush"> announced a new bill that would send <MONEY="\$1.2 million dollars"> to <LOCATION="Miami Florida:"> for a new hurricane tracking system.

Who announced the new bill?

Where will the money from the new bill go?

How much will be provided for the new bill?

Combining and normalizing evidence

Q: *What year did the Titanic sink?*

A: 1912

Supporting evidence:

It was the worst peacetime disaster involving a British ship since the Titanic sank on the 14th of April, 1912.

The Titanic sank after striking an iceberg in the North Atlantic on April 14th, 1912.


The Herald of Free Enterprise capsized off the Belgian port of Zeebrugge on March 6, 1987, in the worst peacetime disaster involving a British ship since the Titanic sank in 1912.

What happened?

- **Different formats** for answer candidates detected, normalized and combined:
 - `April 14th, 1912`
 - `14th of April, 1912`
- **Supporting evidence** detected and combined:
 - `1912` supports `April 14th, 1912`
- **Structure** of date expressions understood and correct piece output:
 - `1912` rather than `April 14th, 1912`
- **Most frequent answer** candidate found and output:
 - `April 14th, 1912` rather than something else.

NLP is expensive!


- We need to run:



- sentence identifier

- tokenizer


- IE



- parsers

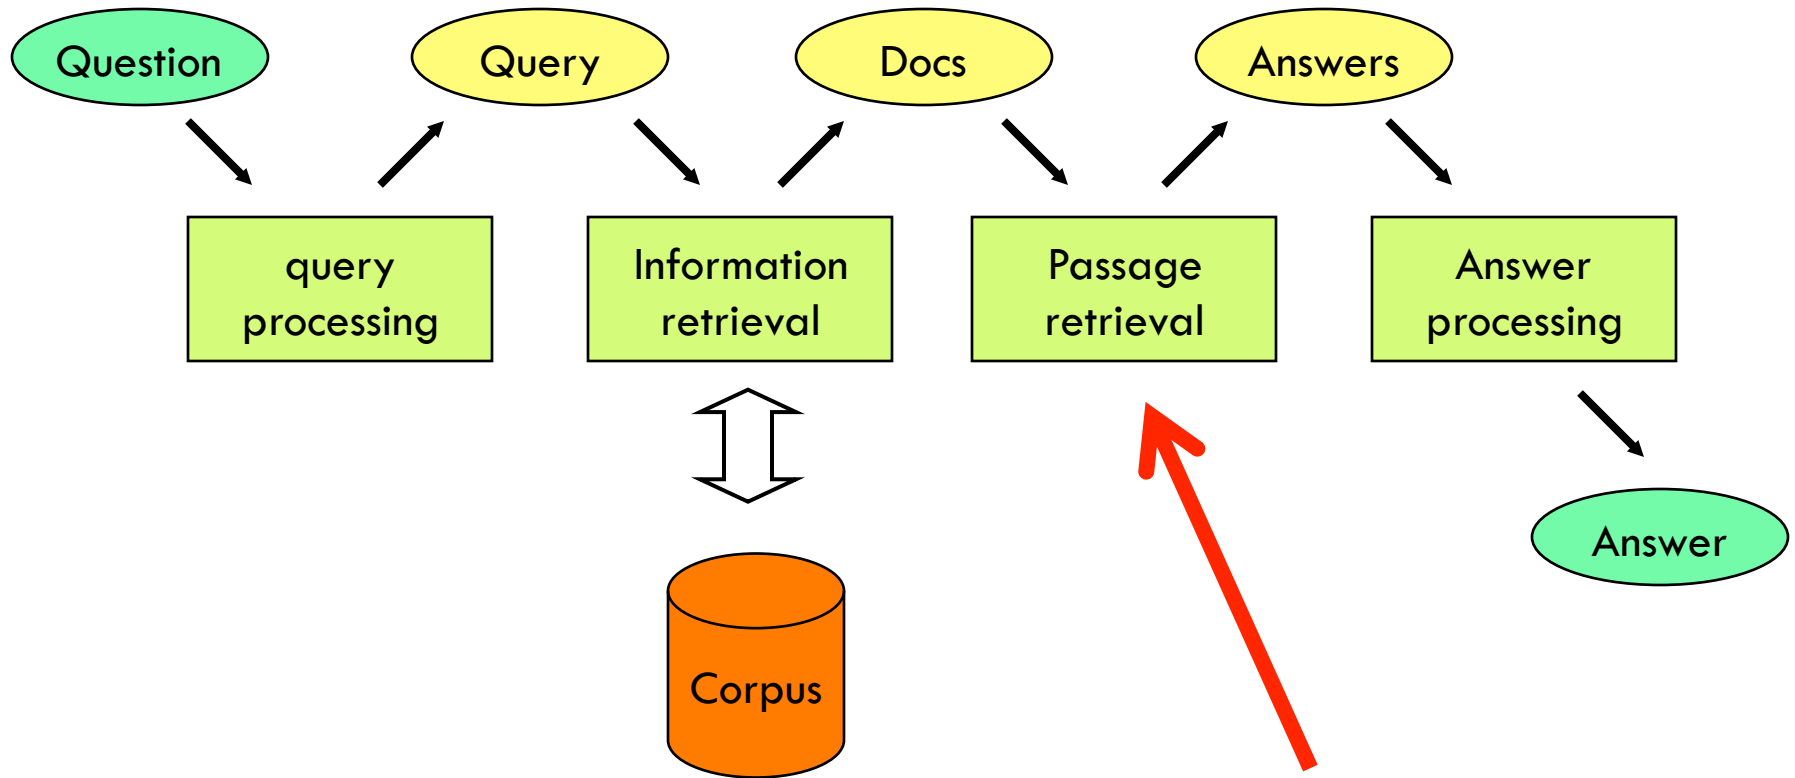
- ...

- These are expensive! (and there may be many documents)



- Can we reduce this time?

QA actual framework



find relevant passages
before doing NLP

Passage retrieval

- *IR Systems* return the N most relevant documents
- *Passage retrieval systems* then return the N most relevant *passages*
- The answer processing is then only run on these passages



Evaluation

- How can/should we evaluate the system?
- Ideas:
 - ▣ Accuracy
 - ▣ Precision/Recall (often we allow it to say “No Answer”)
 - ▣ Produce multiple answers and use IR like measures

Reciprocal ranking scheme



For the purposes of evaluating the systems, several answer candidates are often ranked by confidence.

Reciprocal ranking scheme

For the purposes of evaluating the systems, several answer candidates are often ranked by confidence.

Reciprocal Ranking Scheme: the score for a question is $1/R$, where R is a rank of the first correct answer in the list.

Q: What is the capital of Utah?

A1: Ogden

A2: Salt Lake City

A3: Provo

A4: St. George

A5: Salt Lake

RRS?

Reciprocal ranking scheme

For the purposes of evaluating the systems, several answer candidates are often ranked by confidence.

Reciprocal Ranking Scheme: the score for a question is $1/R$, where R is a rank of the first correct answer in the list.

Q: What is the capital of Utah?

A1: Ogden

A2: Salt Lake City

A3: Provo

A4: St. George

A5: Salt Lake

The score for the question Q would be $1/2$.

Sample Results

Javelin Batch Test results (04/16)

Question	Score	Answers	Answer Pattern
290: When was Hurricane Hugo?	0.0	Nov. 30 June 1	1989
327: When was the San Francisco fire?	0.0	# Error message: RF server EXCEPTION: No answer found!	82\s+years\s+ago 1906 1849 april\s*18 '06
372: When was the Triangle Shirtwaist fire?	1.0	1911 (Correct) March 25, 1911 (Correct) 1988 Dec. 31, 1986 May 28, 1977 Oct. 24, 1976 1985	(March\s*25\s*,?\s*)?1911
416: When was CNN's first broadcast?	0.5	1988 June, 1980 (Correct)	June\s*,?\s+1980
		May (Correct)	

Does it work?



how tall is the tallest building



Input interpretation:

tallest building

by height

Result by height:

[More](#)

1	Burj Khalifa	2717 ft	<input type="text"/>
2	CN Tower	1814 ft	<input type="text"/>
3	Ostankino Tower	1772 ft	<input type="text"/>
4	Taipei 101	1670 ft	<input type="text"/>
5	Shanghai World Financial Center	1614 ft	<input type="text"/>

[Units »](#)

Computed by [Wolfram Mathematica](#)

[Source information »](#)

Download as: [PDF](#) | [Live Mathematica](#)

start.csail.mit.edu/answer.php?query=How+tall+is+the+tallest+building+%3F

S7 Nachrichten aus P... SPIEGEL ONLINE ... H2020 MDFramework JTIG DFKI Mathe4ML

START

Natural Language Question Answering System

How tall is the tallest building ? [Ask Question >](#)

====> How tall is the tallest building ?

I am sorry to say I don't know how tall the tallest building is.

[Back to home page](#)

Does it work?



The screenshot shows the WolframAlpha logo at the top, with the text "computational... knowledge engine" to its right. Below the logo is a search bar containing the text "what is the capitol of Utah?". Underneath the search bar is a section labeled "Input Interpretation:" which shows the words "Utah" and "capital" in separate boxes. Below that is a section labeled "Result:" which displays the text "Salt Lake City, Utah, United States".

START's reply

==> what is the capitol of utah?

[Salt Lake City](#) is the capital of [Utah](#).

Source: START KB

Does it work?



who invented the light bulb?



↳ Using closest Wolfram|Alpha interpretation: **light bulb**



More interpretations: **the light**

Definition:

noun electric lamp consisting of a transparent or translucent glass housing containing a wire filament (usually tungsten) that emits light when heated by electricity

Does it work?

START's reply

====> who invented the lightbulb?

[Electric lamp](#)

[Electric lamp:](#)

(arc lamp) Sir Humphrey Davy, England, 1801; (fluorescent lamp) A.E. Becquerel, France, 1867; (incandescent lamp) Sir Joseph Swann, England, [Thomas A. Edison](#), U.S., contemporaneously, 1870s; (carbon arc street lamp) [Charles F. Brush](#), U.S., 1879; (first widely marketed incandescent lamp) Thomas A. Edison, U.S., 1879; (mercury vapor lamp) Peter Cooper Hewitt, U.S., 1903; (neon lamp) Georges Claude, France, 1911; (tungsten filament) [Irving Langmuir](#), U.S., 1915.



[Thomas Alva Edison](#)
(1847–1931)
Library of Congress

Source: [Infoplease.com](#)

Does it work?



who is the best natural language processing professor?



↳ Using closest Wolfram|Alpha interpretation: **best natural language processing**



More interpretations: [the best professor](#) | [professor](#) | [the best](#) | [natural language processing](#)

Natural Language Processing



Development of this topic is under investigation...

Leave your email address to
show your interest.

send

Does it work?

- Where do lobsters like to live?
 - ▣ on a Canadian airline
- Where do hyenas live?
 - ▣ in Saudi Arabia
 - ▣ in the back of pick-up trucks
- Where are zebras most likely found?
 - ▣ near dumps
 - ▣ in the dictionary
- Why can't ostriches fly?
 - ▣ Because of American economic sanctions
- What's the population of Maryland?
 - ▣ three

Text REtrieval Conference (TREC)

*...to encourage research in information retrieval
from large text collections.*

<http://trec.nist.gov>

Overview

**Other
Evaluations**

Publications

**Information
for Active
Participants**



**Frequently
Asked
Questions**

Tracks

Data

**Past TREC
Results**

**Contact
Information**

NIST

National Institute of Standards and Technology

Question Answering at TREC

- Question answering competition at TREC consists of answering a set of 500 fact-based questions, e.g., “When was Mozart born?”.
- For the first three years systems were allowed to return 5 ranked answer snippets (50/250 bytes) to each question.
 - IR think
 - Mean Reciprocal Rank (MRR) scoring:
 - 1, 0.5, 0.33, 0.25, 0.2, 0 for 1, 2, 3, 4, 5, 6+ doc
- Mainly Named Entity answers (person, place, date, ...)
- From 2002 the systems were only allowed to return a single exact answer and the notion of confidence has been introduced.

The TREC Document Collection

- The current collection uses news articles from the following sources:
 - ▣ AP newswire,
 - ▣ New York Times newswire,
 - ▣ Xinhua News Agency newswire,
- In total there are 1,033,461 documents in the collection.
3GB of text
- Clearly this is too much text to process entirely using advanced NLP techniques so the systems usually consist of an initial information retrieval phase followed by more advanced processing.
- Many supplement this text with use of the web, and other knowledge bases

Sample TREC questions

- Who is the author of the book, "The Iron Lady: A Biography of Margaret Thatcher"?
- What was the monetary value of the Nobel Peace Prize in 1989?
- What does the Peugeot company manufacture?
- How much did Mercury spend on advertising in 1993?
- What is the name of the managing director of Apricot Computer
- Why did David Koresh ask the FBI for a word processor?
- What debts did Qintex group leave?
- What is the name of the rare neurological disease with symptoms such as: involuntary movements (tics), swearing, and incoherent vocalizations (grunts, shouts, etc.)?

Top Performing Systems

- Currently the best performing systems at TREC can answer approximately 70% of the questions
- Approaches and successes have varied a fair deal
 - ▣ Knowledge-rich approaches, using a vast array of NLP techniques got the best results in 2000, 2001
 - ▣ AskMSR system stressed how much could be achieved by very simple methods with enough text (and now various copycats)
 - ▣ Middle ground is to use large collection of surface matching patterns (ISI)

ISI: Surface patterns approach

- ISI's approach
- Use of Characteristic Phrases
- "When was <person> born"
 - ▣ Typical answers
 - "Mozart was born in 1756."
 - "Gandhi (1869-1948)..."
 - ▣ Suggests phrases like
 - "<NAME> was born in <BIRTHDATE>"
 - "<NAME> (<BIRTHDATE>-"
 - ▣ as Regular Expressions can help locate correct answer

Use Pattern Learning

- Example:
 - “The great composer Mozart (1756-1791) achieved fame at a young age”
 - “Mozart (1756-1791) was a genius”
 - “The whole world would always be indebted to the great music of Mozart (1756-1791)”
- ▣ Longest matching substring for all 3 sentences is “Mozart (1756-1791)”
- ▣ Suffix tree would extract “Mozart (1756-1791)” as an output, with score of 3
- Reminiscent of IE pattern learning

Algorithm 1 for Pattern Learning

- Select an example for a given question
 - Ex. for BIRTHYEAR questions we select “Mozart 1756” (“Mozart” as the question term and “1756” as the answer term).
- Submit the question and the answer term as queries to a search engine. Download the top 1000 web documents
- Apply a sentence breaker to the documents. Retain only those sentences that contain both the question and the answer term.
- Tokenize and pass each retained sentence through a suffix tree constructor. This finds all substrings, of all lengths, along with their counts. For example consider the

Example

- Given 3 sentences:
 - “The great composer Mozart (1756–1791) achieved fame at a young age”
 - “Mozart (1756–1791) was a genius”.
 - “The whole world would always be indebted to the great music of Mozart (1756–1791)”.
- The longest matching substring for all 3 sentences is “Mozart (1756–1791)”, which the suffix tree would extract as one of the outputs along with the score of 3.
- Pass each phrase in the suffix tree through a filter to retain only those phrases that contain both the question and the answer term. For the example, we extract only those phrases from the suffix tree that contain the words “Mozart” and “1756”.
- Replace the word for the question term by the tag “<NAME>” and the word for the answer term by the term “<ANSWER>”.

Pattern Learning (cont.)

- Repeat with different examples of same question type
 - “Gandhi 1869”, “Newton 1642”, etc.
- Some patterns learned for BIRTHDATE
 - a. born in <ANSWER>, <NAME>
 - b. <NAME> was born on <ANSWER> ,
 - c. <NAME> (<ANSWER> -
 - d. <NAME> (<ANSWER> -)

Algorithm 2: Calculating Precision of Derived Patterns

- Query the search engine by using only the question term (in the example, only “Mozart”). Download the top 1000 web documents
- Segment these documents into sentences. Retain only those sentences that contain the question term.
- For each pattern obtained from Algorithm, check the presence of each pattern in the sentence obtained from above for two instances:
 - Presence of the pattern with <ANSWER> tag matched by any word.
 - Presence of the pattern in the sentence with <ANSWER> tag matched by the correct answer term.

Algorithm 2: Example

- For the pattern “<NAME> was born in <ANSWER>” check the presence of the following strings in the answer sentence
 - Mozart was born in <ANY_WORD>
 - Mozart was born in 1756
- Calculate the precision of each pattern by the formula $P = C_a / C_o$ where
 - C_a = total number of patterns with the answer term present
 - C_o = total number of patterns present with answer term replaced by any word
- Retain only the patterns matching a sufficient number of examples (> 5).

Experiments

- 6 different Q types
 - ▣ from Webclopedia QA Typology (Hovy et al., 2002a)
 - BIRTHDATE
 - LOCATION
 - INVENTOR
 - DISCOVERER
 - DEFINITION
 - WHY-FAMOUS

Experiments: pattern precision

- BIRTHDATE table:
 - 1.0 <NAME> (<ANSWER> -)
 - 0.85 <NAME> was born on <ANSWER> ,
 - 0.6 <NAME> was born in <ANSWER>
 - 0.59 <NAME> was born <ANSWER>
 - 0.53 <ANSWER> <NAME> was born
 - 0.50 - <NAME> (<ANSWER>
 - 0.36 <NAME> (<ANSWER> -
- INVENTOR
 - 1.0 <ANSWER> invents <NAME>
 - 1.0 the <NAME> was invented by <ANSWER>
 - 1.0 <ANSWER> invented the <NAME> in

Experiments (cont.)

□ DISCOVERER

- 1.0 when <ANSWER> discovered <NAME>
- 1.0 <ANSWER>'s discovery of <NAME>
- 0.9 <NAME> was discovered by <ANSWER> in

□ DEFINITION

- 1.0 <NAME> and related <ANSWER>
- 1.0 form of <ANSWER>, <NAME>
- 0.94 as <NAME>, <ANSWER> and

Experiments (cont.)

□ WHY-FAMOUS

- 1.0 <ANSWER> <NAME> called
- 1.0 laureate <ANSWER> <NAME>
- 0.71 <NAME> is the <ANSWER> of

□ LOCATION

- 1.0 <ANSWER>'s <NAME>
- 1.0 regional : <ANSWER> : <NAME>
- 0.92 near <NAME> in <ANSWER>

- Depending on question type, get high MRR (0.6–0.9), with higher results from use of Web than TREC QA collection

Shortcomings & Extensions

- Need for POS &/or semantic types
 - "Where are the Rocky Mountains?"
 - "Denver's new airport, topped with white fiberglass cones in imitation of the Rocky Mountains in the background , continues to lie empty"
 - <NAME> in <ANSWER>
- NE tagger &/or ontology could enable system to determine "background" is not a location

Shortcomings... (cont.)

- Long distance dependencies
 - "Where is London?"
 - "London, which has one of the most busiest airports in the world, lies on the banks of the river Thames"
 - would require pattern like:
<QUESTION>, (<any_word>)*, lies on <ANSWER>
- Abundance & variety of Web data helps system to find an instance of patterns w/o losing answers to long distance dependencies

Shortcomings... (cont.)

- System currently has only one anchor word
 - ▣ Doesn't work for Q types requiring multiple words from question to be in answer
 - "In which county does the city of Long Beach lie?"
 - "Long Beach is situated in Los Angeles County"
 - required pattern:
<Q_TERM_1> is situated in <ANSWER>
<Q_TERM_2>



Say hello to Watson


Meet IBM Watson, a cognitive system that enables a new partnership between people and computers that enhances and scales human expertise.



How Watson Works





Developer Cloud, your gateway
to cognitive computing



Read more about how Watson is
transforming travel

IBM Watson's project started 2007

- Project started in 2007, lead David Ferrucci
- Initial goal: create a system able to process natural language & extract knowledge faster than any other computer or human
- Jeopardy! was chosen because it's a huge challenge for a computer to find the questions to such "human" answers under time pressure
- Watson was NOT online!
- Watson weighs the probability of his answer being right – doesn't ring the buzzer if he's not confident enough
- Which questions Watson got wrong almost as interesting as which he got right!



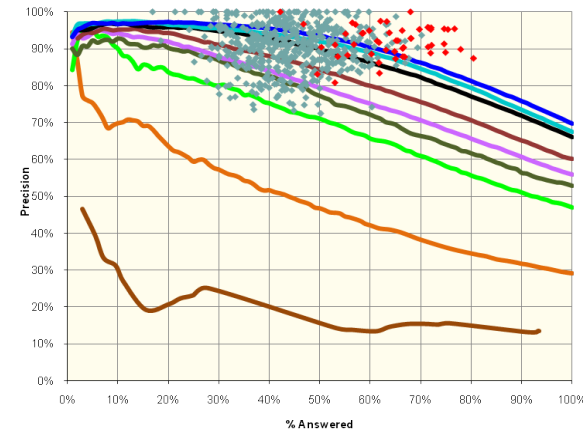
Watson – a Workload Optimized System

- 90 x IBM Power 750 servers
- 2880 POWER7 cores
- POWER7 3.55 GHz chip
- 500 GB per sec on-chip bandwidth
- 10 Gb Ethernet network
- 15 Terabytes of memory
- 20 Terabytes of disk, clustered
- Can operate at 80 Teraflops
- Runs IBM DeepQA software
- Scales out with and searches vast amounts of unstructured information with UIMA & Hadoop open source components
- Linux provides a scalable, open platform, optimized to exploit POWER7 performance
- 10 racks include servers, networking, shared disk system, cluster controllers



Watson: Performance

- **Deep Analytics** – IBM Watson achieved champion-levels of *Precision* and *Confidence* over a huge variety of expression
- **Speed** – By optimizing Watson’s computation for Jeopardy! on **2,880 POWER7** processing cores we went from **2 hours per question on a single CPU** to **an average of just 3 seconds** – fast enough to compete with the best.
- **Results** – in 55 real-time sparring against former **Tournament of Champion Players last year**, Watson put on a very competitive performance, winning 71%. In the final Exhibition Match against Ken Jennings and Brad Rutter, Watson won!





Siri.

Your wish is
its command.

Siri lets you use your voice to send messages, schedule meetings, place phone calls, and more. Ask Siri to do things just by talking the way you talk. Siri understands what you say, knows what you mean, and even talks back. Siri is so easy to use and does so much, you'll keep finding more and more ways to use it.



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

About 14,700,000 results (0.33 seconds)

The Burj Khalifa in Dubai, United Arab Emirates. The **world's tallest** man-made structure is the 829.8 m (2,722 ft) tall Burj Khalifa in Dubai, United Arab Emirates. The building gained the official title of "**Tallest Building in the World**" at its opening on January 4, 2010.



[List of tallest buildings and structures in the world - Wikipedia](https://en.wikipedia.org/wiki/List_of_tallest_buildings_and_structures_in_the_world)
[en.wikipedia.org/.../List_of_tallest_buildings_and_structures_in_the_world](https://en.wikipedia.org/wiki/List_of_tallest_buildings_and_structures_in_the_world)

Feedback