

SHALLOW QA

Joseph Casamento

<jcasamento2009@my.fit.edu>

Christopher Struttman

<cstruttman2010@my.fit.edu>

Faculty Advisor: Phillip Chan

<pkc@cs.fit.edu>

About Us

- Chris Struttmann
 - Florida Tech Senior (B.S. Software Engineering)
 - FIRST Robotics alumni (Melbourne Highschool, Team Voltage 386)
 - Systems Engineer (pending developer position) -- Tastemaker Labs, San Francisco, CA
- Joey Casamento
 - Florida Tech Senior (B.S. Computer Science)
 - Software Development Engineer at Microsoft (starting June 2013)

What is ShallowQA?

- Question/Answer System
 - Takes question posed in natural language as input.
 - Output is a hypothesis (potential answer)
- Inspired by IBM Watson/DeepQA
 - Big
 - Expensive
- ShallowQA doesn't rely on internal documents
 - Little storage needed

What's an algorithm?

Merriam-Webster says: a procedure for solving a mathematical problem (as of finding the greatest common divisor) in a finite number of steps that frequently involves repetition of an operation; broadly : a step-by-step procedure for solving a problem or accomplishing some end especially by a computer

Put Simply: a set of specific instructions to achieve a desired result

What is "Natural Language?"

- The language that we as humans use to communicate
- Unstructured
- Ambiguous
- Many interpretations for the same thing

Consider the following...

"I shot an elephant wearing my pajamas"

- NLP's favorite phrase
- Who was wearing the pajamas?
- What'd you shoot with? A camera?
- Can *YOU* write an algorithm to answer these questions given the above sentence?
 - You can come to F.I.T. and find out!

**Knowing all the challenges... how
does ShallowQA work?**

**Knowing all the challenges... how
does ShallowQA work?**

PROPRIETARY

**Knowing all the challenges... how
does ShallowQA work?**

PROPRIETARY

(KIDDING)

Knowing all the challenges... how does ShallowQA work?

- Three Basic Steps (Highest Level)
 - Question Analysis
 - What is the type of question?
(Who/What/Where/When)
 - Do we even know how to answer this kind of question? Should we even try.
 - What is the question asking?
 - Hypothesis Generation
 - Analyze Evidence
 - Generate hypotheses from evidence
 - "Hypothesis" ≈ Educated Guess
 - Ranking and Merging
 - Take all of the generated hypotheses, rank them.
 - The "most correct" answer should be the highest ranked.

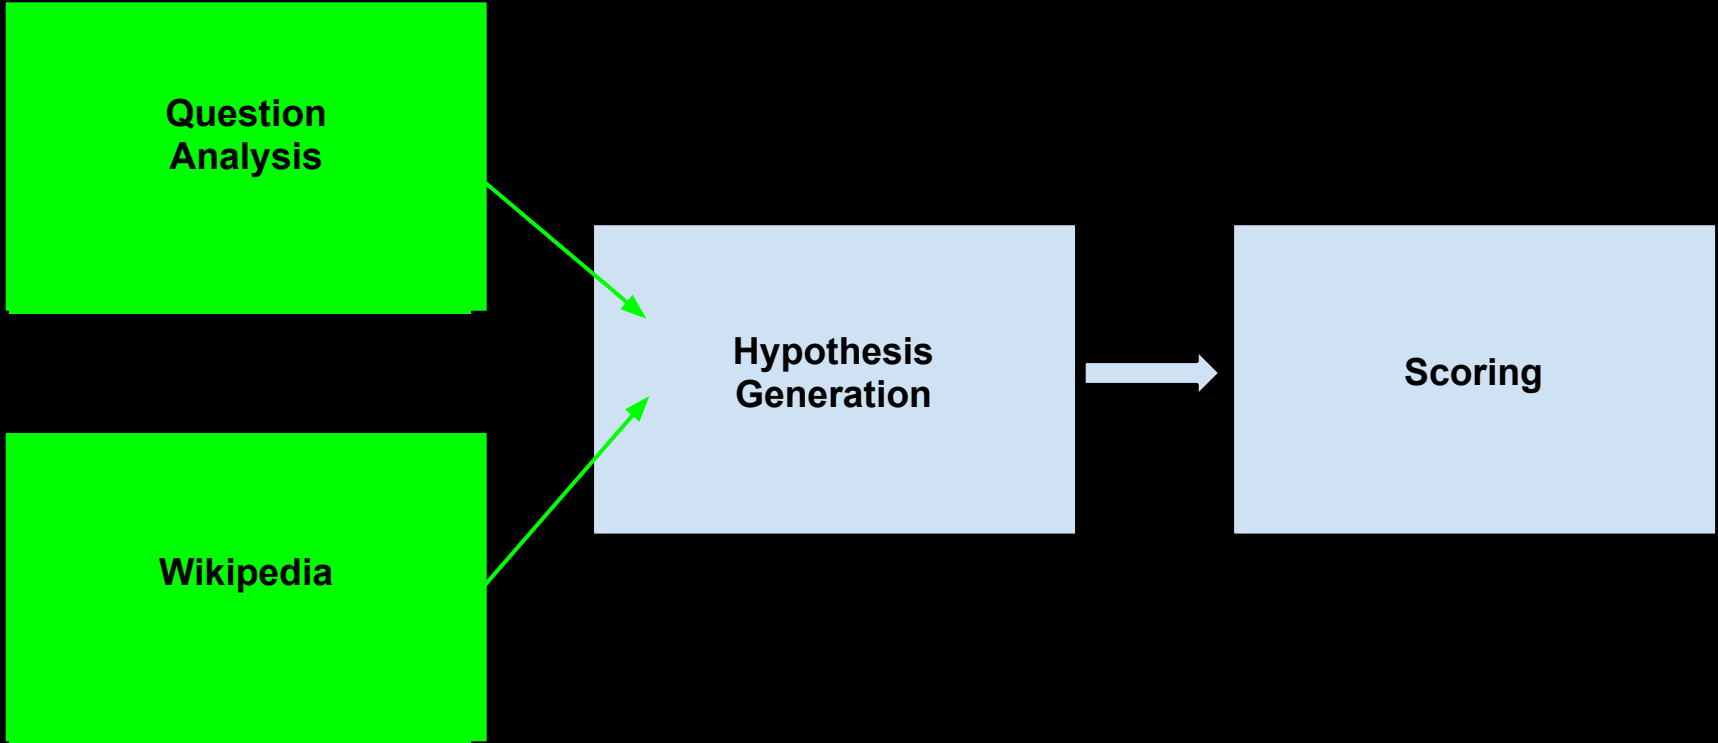
Toolbox

- Current

- Java7 JDK
- Stanford CoreNLP Library
 - Natural Language Parsing library
 - We abandoned the bigger project and raided smaller bits for performance reasons
- Princeton Wordnet
 - Lexicographical database

- Incoming

- Google Guava
 - Cool utility functions that are tedious to otherwise build and test
- Wordnet Blast
 - Ultra-fast access to the WordNet database



Step one: Question Analysis

Question:

"When was the original Macintosh Introduced?"

Question Analysis

Answer: Before I was born

Question Analysis

Also accepted: 1984.

Question Analysis: What do we need to know?

- What type of question is this?

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis: What do we need to know?

- What type of question is this?
 - "When"

Question Analysis: What do we need to know?

- What type of question is this?
 - "When"
- What's the question asking about

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis: What do we need to know.

- What type of question is this?
 - "When"
- What's the question asking about?
 - "Macintosh"

Question Analysis: What do we need to know.

- What type of question is this?
 - "When"
- What's the question asking about?
 - "Macintosh"
- What about the Macintosh?

Question Analysis

"When was the original **Macintosh** Introduced?"

Question Analysis: What do we need to know.

- What type of question is this?
 - "When"
- What's the question asking about?
 - "Macintosh"
- What about the Macintosh?
 - "Introduced"

Question Analysis: What do we need to know.

- What type of question is this?
 - "When"
- What's the question asking about?
 - "Macintosh"
- What about the Macintosh?
 - "Introduced"
- Anything else?

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis

"When was the original Macintosh Introduced?"

Question Analysis: What do we need to know.

- What type of question is this?
 - "When"
- What's the question asking about?
 - "Macintosh"
- What about the Macintosh?
 - "Introduced"
- Anything else?
 - The **original** macintosh

Question Analysis



"When was the original Macintosh Introduced?"

Question Analysis

"When was the original Macintosh Introduced?"



Question Analysis

"When was the original Macintosh Introduced?"



Question Analysis


"When was the original Macintosh Introduced?"



The diagram consists of a blue line above the sentence and a white line below it. The blue line starts at the beginning of the sentence, goes up at an angle, then horizontally across the top, and then down at an angle to the end of the sentence. The white line starts below the word 'Macintosh', goes down at an angle, then horizontally across the bottom, and then up at an angle to the end of the sentence. An orange arrow points downwards from the word 'original' to the white line.

(Remember sentence diagramming)

Question Analysis



Sentence Parsing

Interestingly Enough...

Evidence Analysis

Consider this hypothetical sentence from Wikipedia:

"The original Macintosh was introduced in 1984."

Evidence Analysis

"The original Macintosh was introduced in 1984."



Evidence Analysis

"The original Macintosh was introduced in 1984."



Question:



"When was the original Macintosh Introduced?"

Evidence:



"The original Macintosh was introduced in 1984."

But what if...

Instead of:

"The original Macintosh was introduced in 1984."

But what if...

The evidence says:

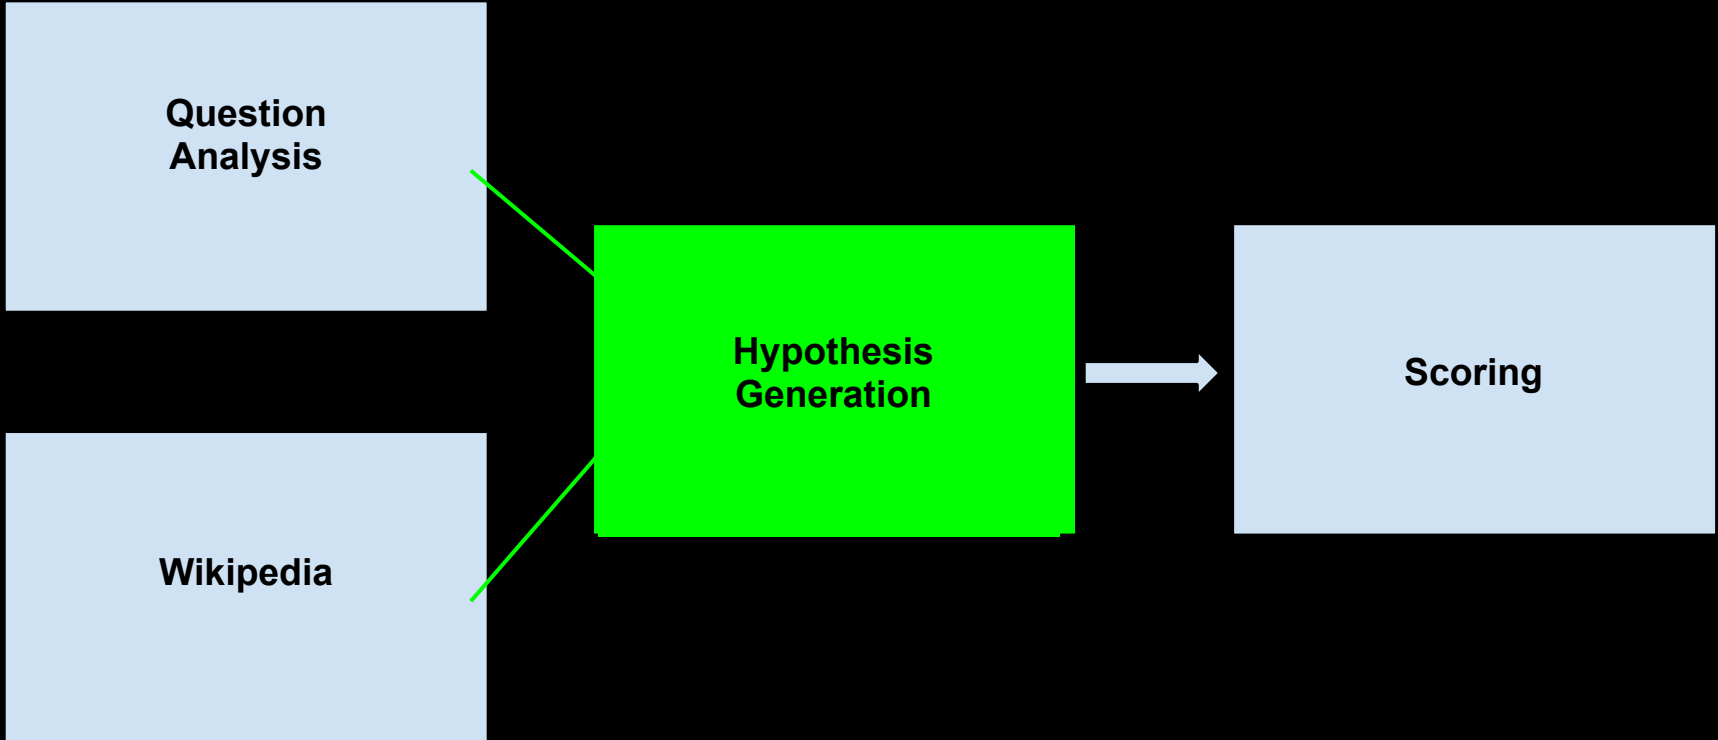
"The first Macintosh was unveiled in 1984."

But what if...

The evidence says:

"The first Macintosh was unveiled in 1984."

We have steps 2 & 3, and Joey to tell you about them



Step 2: Hypothesis Generation

- Identify potentially millions of possible answers
- Why is hypothesis generation important?
 - If the correct answer to the question isn't in the set of candidate hypotheses, then it will be impossible for ShallowQA to provide the correct answer.

Step 2: Hypothesis Generation

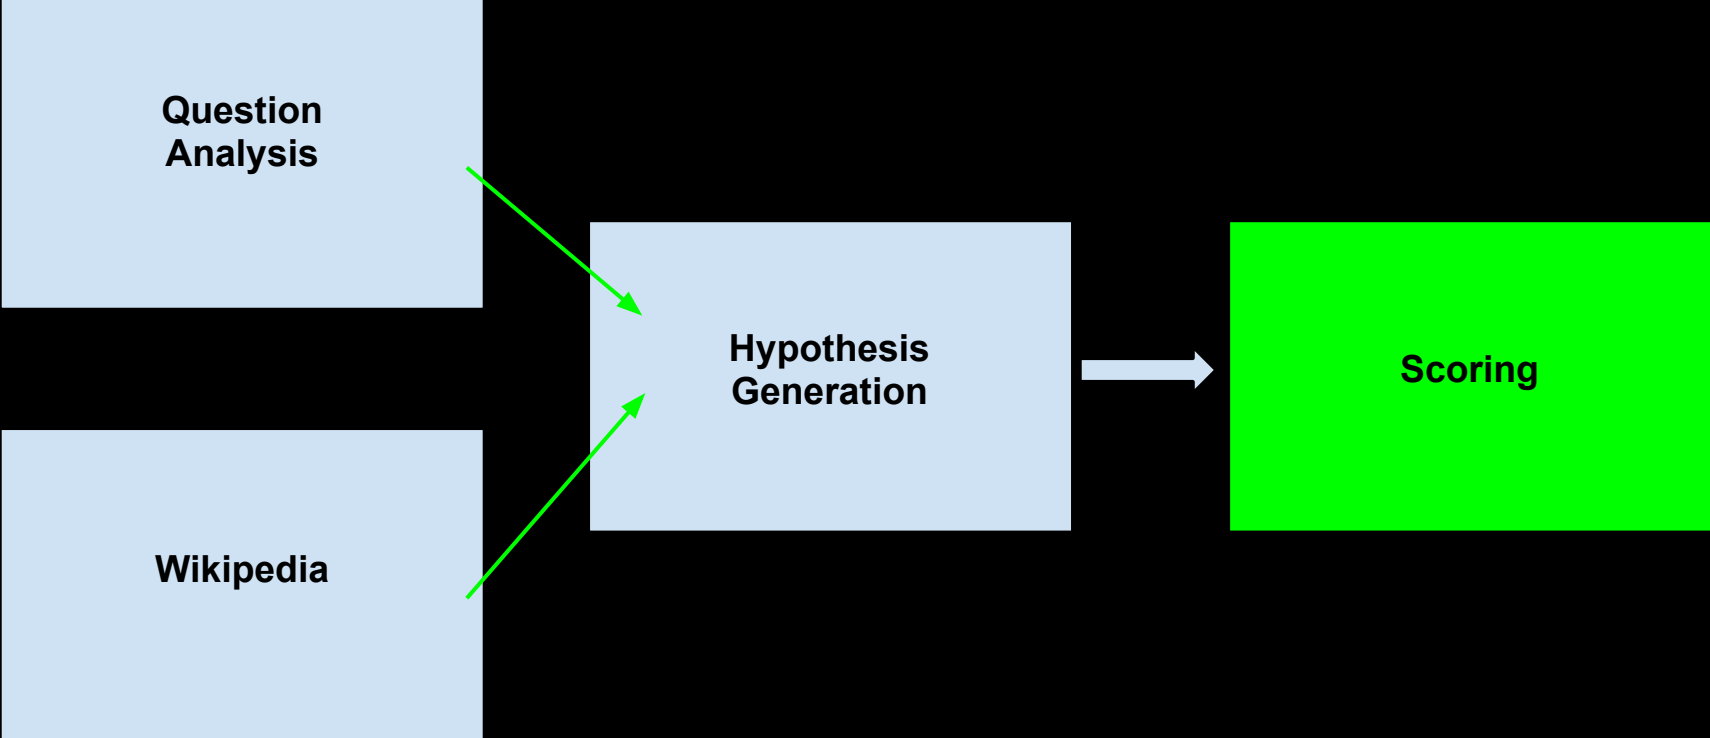
- So far we've
 - Parsed the question
 - Identified the question as a "when" question
 - Downloaded and parsed the corresponding Wikipedia page
- Because it's a "when" question, the answer will be a date
 - Identify evidence with supporting dates
 - Discard the remaining set

Step 2: Hypothesis Generation

- Example set of candidate hypotheses
 - Jobs resigned from Apple in 1985.
 - In 1991 Apple released the PowerBook line.
 - Apple Inc.'s then-chairman Steve Jobs introduced the first Macintosh on January 24, 1984.
 - Jobs resigned from Apple in 1985.

Step 2: Hypothesis Generation

- Example set of candidate hypotheses
 - Jobs resigned from Apple in 1985.
 - In 1991 Apple released the PowerBook line.
 - Apple Inc.'s then-chairman Steve Jobs introduced the first Macintosh on January 24, 1984.
 - Jobs resigned from Apple in 1985.



Step 3: Ranking

- The process of sorting through all possible answers and ordering them in order of likelihood of being the correct answer.
- If the correct answer isn't at the top of the list, ShallowQA provides the wrong answer
- Current development stage

Step 3: Ranking

- So far we've
 - Parsed the question
 - Identified the question as a "when" question
 - Downloaded and parsed the corresponding Wikipedia page
 - Generated a list of possible answers (candidate hypotheses)

Step 3: Ranking

- Consider our question
 - When was the original Macintosh Introduced?
- Consider this candidate hypothesis
 - In 1991 Apple released the PowerBook line.

Step 3: Ranking

When was the original Macintosh Introduced?

Very Close



In 1991 Apple released the PowerBook line.

Step 3: Ranking

When was the original Macintosh Introduced?

NOT the same



In 1991 Apple released the PowerBook line.

Step 3: Ranking

When was the original Macintosh Introduced?

Direct Match



Steve Jobs introduced the first Macintosh on January 24, 1984.

Step 3: Ranking

When was the original Macintosh Introduced?

Direct Match

Direct Match

Steve Jobs introduced the first Macintosh on January 24, 1984.

Step 3: Ranking

When was the original Macintosh Introduced?

Close synonym

Direct match

Direct match

Steve Jobs introduced the first Macintosh on January 24, 1984.

Step 3: Ranking

When was the original Macintosh Introduced?

Close synonym

Direct match

Direct match

Steve Jobs introduced the first Macintosh on January 24, 1984.

VERY Close Match

Final Remarks

- Unstructured data accounts for 70-80% of the world's information
- World's data doubles every ~2 years¹
- QA technology is part of an emerging technology area
 - Everyone's heard the "big data" buzzword.
- Uses beyond Siri
 - Medical (IBM)
 - Financial data
 - Research tools

¹ <http://www.dailytech.com/Worlds+Data+to+Reach+18+Zettabytes+by+2011/article11055.htm>

Thank You!

Live Demo

Questions?