

# *AI: A Return to Meaning*

*Perspectives on the Evolution AI*

David A. Ferrucci



# Outline

- From Theory-Driven to Data Driven AI
  - Where's the Sweet Spot?
- Reflections on IBM's *Watson*
  - A landmark in Language AI
- The Future of AI
  - Deep Understanding



# Artificial Intelligence

*Computer systems whose interactive behavior is indistinguishable from a human's.*



Watson?

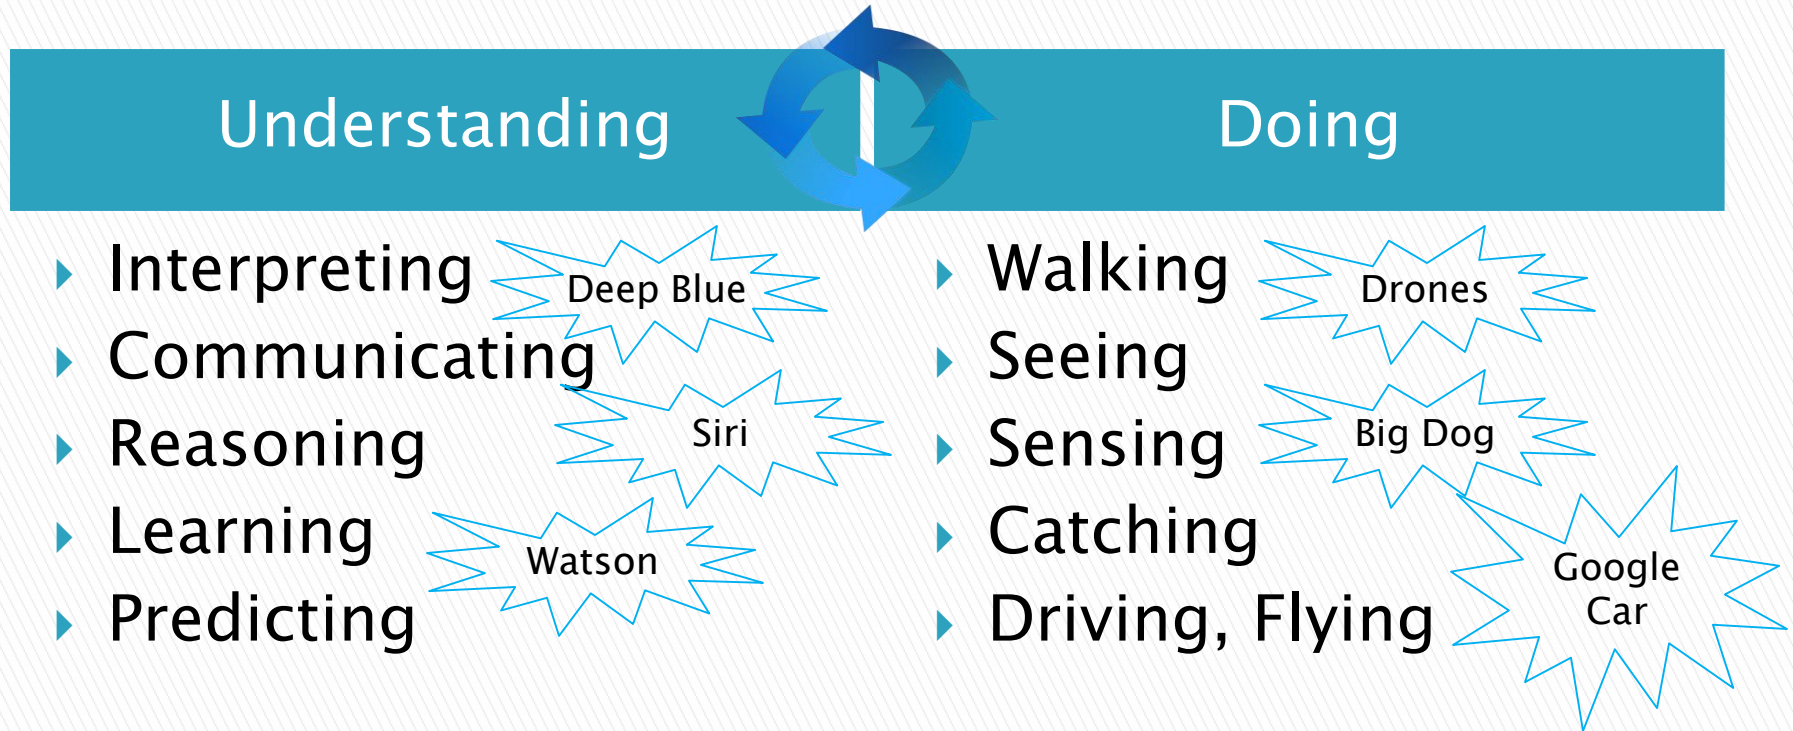


*Computer systems that perform tasks that if performed by a human would be associated with Intelligence*



DeepBlue?

# AI: Knowing and Doing



Question Answering Systems,  
Expert Systems...

WATCH (1)

Operating in **Cognitive Space**

Vision Systems,  
Autonomous Robots...

WATCH (1)


Operating in the **Physical Space**

Self/Situational Awareness

# What's Harder: A Game of Chess? or Just a Good Chat?

- Chess

- Finite, mathematically well-defined search space
- All responses grounded in precise, unambiguous rules
- Large but finite set of possible moves
- Perfect for a computer. Amazing that humans can do it!



- Human Language

- Words (or images, speech) lack precise interpretation
- Nearly infinite expressions map to a huge variety in meaning
- **Meaning** grounded only in shared human experience - highly **contextual**, uniquely human and none precisely alike



*The majority of human interpretable knowledge is in unstructured forms.*

*Fluent personalized access to this knowledge is the key to productivity & better decision making*

# Why is getting at the meaning so important?


Volume and complexity outpacing our ability to assimilate well enough to make good decisions

- The Healthcare Bill is about 2000+ pages of natural language text.
- In 2006 US Tax code was 16,000 pages. In 2010 it was 71,000 pages!
- 100,000s+ medical text books and journals...and rapidly changing.
- Wikipedia has over 5,000,000 articles!
- The internet provides access to the equivalent of BILLIONS of books


**WHAT DOES IT ALL MEAN TO ME?**

**HELP ME UNDERSTAND THEM AND APPLY THE MEANING  
TO MAKE BETTER DECISIONS FASTER ?**





# UNDERSTANDING HUMAN LANGUAGE REQUIRES INTERPRETING *MEANING*



*Meaning is subjective  
and we **Humans** are the **Subjects***

- *White, Black and Red*
- *Calculator + Dice on Top*
- *Calculating the Odds*
- *Winning\**
- *Human use of Tools*

*Levels of  
Interpretation*



**Meaning:** A probabilistic mapping from symbols to “common experience”  
Context narrows the possibilities & improves confidence in the mapping




The **bat** was flying toward him.




Billy ran as fast as he could.



He made it home safe!



He Scored!





Mapping from language to Meaning can be very subtle  
Meaning can be vastly different based on tiny changes in expressions triggering entirely different experiences.

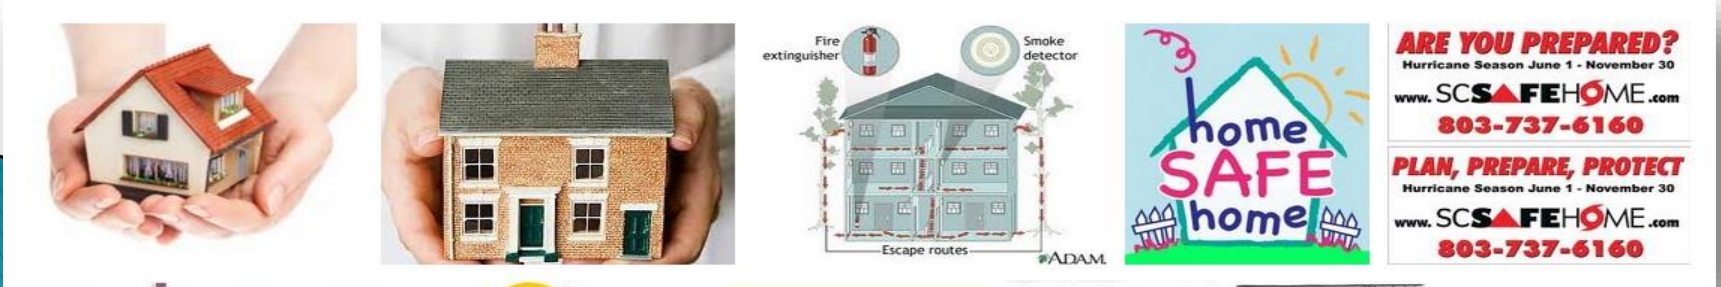
“Safe at Home”



“Home Safe”




“Safe Home”



# A practical perspective on AI...

HOW TO GET AT THE **MEANING** IN THE DATA AND THEN  
HOW TO GET THE MACHINE TO USE THE MEANING TO MAKE USEFUL PREDICTIONS.

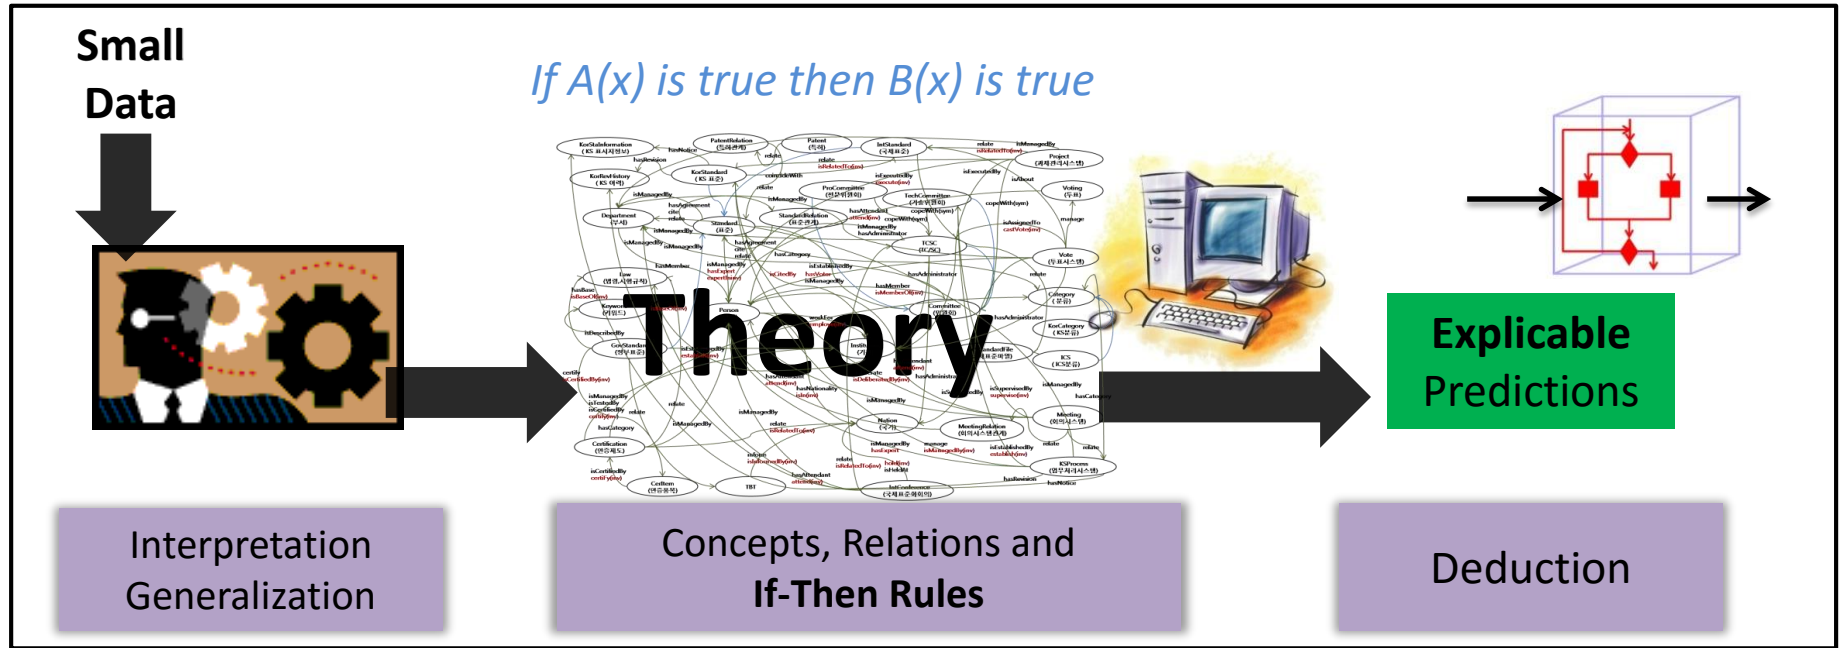


# Theory-Driven Beginnings

**Humans** interpret “small data” and manually capture meaning in the form of a Logical models.

The **Computer** applies rules of inference to **deduce** new predictions.

*Expensive, Narrow, Brittle but Transparent/Explicable*

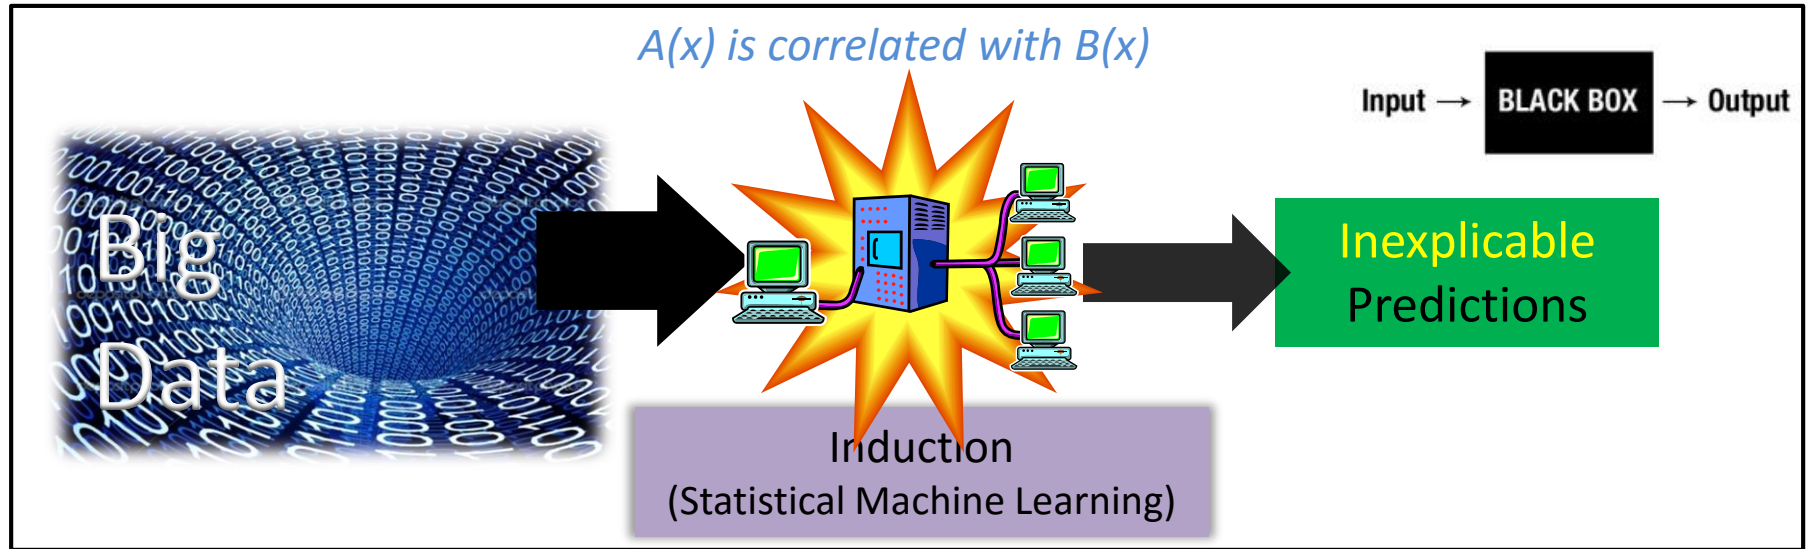


Rationale narrative-based rich understanding

*“Consumers have extended too much credit to pay for homes that the housing bubble had made unaffordable. Many of them had stopped making their payments and there were likely to be substantial losses from this. The degree of leverage in the system would compound the problem, paralyzing the credit market and the financial industry more broadly. The shock might be large enough to trigger a severe recession.”*

# Data-Driven Success

**Massive amounts of data** accessible to **massive compute power** can produce predications based on discovering patterns in the data with much less human effort & interpretation




- Healthcare
- E-commerce (Netflix, Amazon)
- Economics
- Talent (sports and corporate)
- Elections
- ...

Predicting the future  
based on patterns in  
the data

**[Calling a recession] “...the most reliable forward-looking indicators are now collectively behaving as they did on the cusp of full-blown recessions...”**

# A Painfully Simple Decision

Its been raining  
a bit and we are  
walking to  
lunch...



Should I  
wear my  
galoshes?





# Galoshes: Theory Version 1.0

1.  $\exists (x) \text{ Surface}(x) \;; \text{There are surfaces}$
2.  $\exists (x) \text{ Path}(x) \;; \text{There are paths}$
3.  $\forall (x) \text{ Path}(x) \rightarrow \text{Surface}(x) \;; \text{A path is a surface}$
4.  $\exists (x) \text{ Surface}(x) \wedge \text{Covered}(x) \;; \text{Surfaces can be covered}$
5.  $\exists (x) \text{ Surface}(x) \wedge (\text{Wet}(x) \vee \text{Dry}(x)) \;; \text{Surfaces can be wet or dry}$
6.  $\exists (e) \text{ Event}(e) \forall (e) \text{ Raining}(e) \rightarrow \text{Event}(e) \;; \text{There are events and raining is an event}$
7.  $\forall (x,e) \text{ Wet}(x) \leftarrow \text{Raining}(e) \wedge \text{not Covered}(x) \;; \text{a thing is wet if it is raining and it is not covered}$
8.  $\exists (p) \text{ People}(p) \;; \text{There are People}$
9.  $\forall (p,s) \text{ Wear}(p, \text{galoshes}) \leftarrow \text{Walking}(p,s) \wedge \text{Wet}(s) \;; \text{People wear galoshes if walking on a wet path}$
10.  $\forall (p) \text{ Wear}(p, \text{galoshes}) \rightarrow \text{protect}(P, \text{"shoes"})$

Domain Theory in FOL

**User:** I will be walking to lunch. Should I wear galoshes?

**System:** Is it raining? [R9 - R7]

**User:** Yes

**System:** Is the path covered? [R9 – R7]

**User:** No.

**System:** I suggest you wear galoshes.

**User:** Why?

**System:**

1. If it is raining and path is **not** covered then the path is wet.
2. If the path is wet, then people wear galoshes.

**User:** Why?

**System:** To protect the people's shoes? [R10]

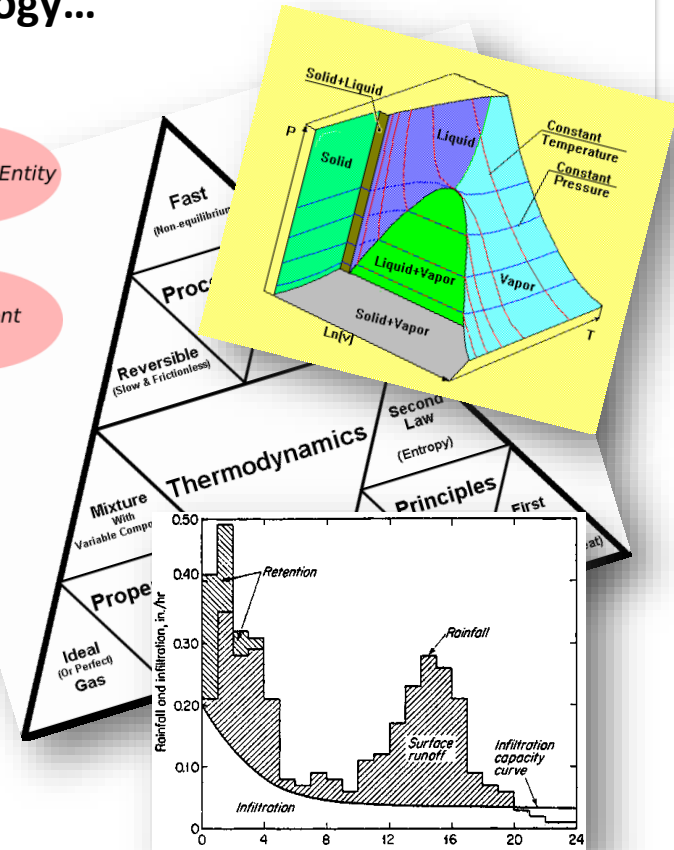

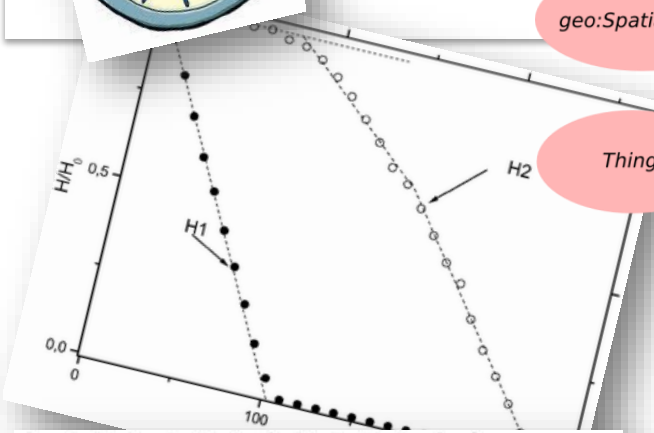
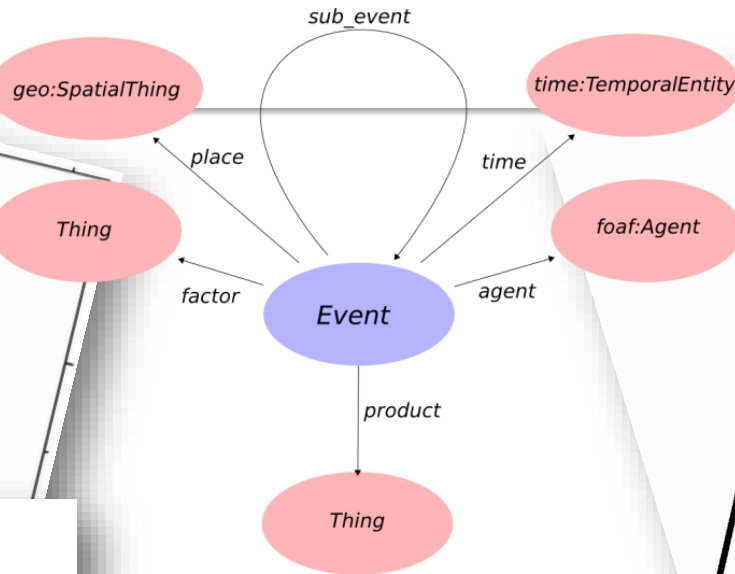
Interaction  
Based on Deduction



# Galoshes Theory Version 1.1

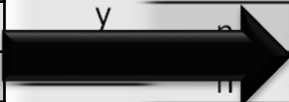
What if it is not raining but the path is still wet. Having a **Theory** allows us to engage human thought, intuition, perception...But it can get challenging to discover, build and maintain.

**Duration of raining events, event start and end points**  
**Drying time, Ground Retention, Temperature, Humidity**  
**Depressions in the surface, Topology...**



# Galoshes Data Version 1.0

Obs	rain	Covered		Galoshes Good
		wet path	Y	
1	y	y	y	Y
2	y	y	n	y
3	n	n	n	n
4	n	n	y	n
5	y	y	y	y
6	y	y	n	y
7	y	n	n	n
8	n	n	y	n
9	n	y	n	y
10	y	n	y	n
11	y	y	n	y
12	n	y	y	n
13	n	y	n	y
...	y	n		n



Are there missing variables that can better explain what is going? How do they relate to how humans think about the problem? What are their logical relationships?

**User:** Should I wear galoshes?

**System:** Yes

**User:** Why?

**System:** Statistically people wear galoshes in situations similar to yours.

**User:** Why?

**System:** It has something to do with rain, wet paths and coverings.

# Galoshes Data Version 1.1

Easy enough to add features...But is what is going on easily interpretable by humans

Obs	Rain	Galoshes Good	Tree						
			Type	season	temp	humidity	Start Time	Covered	Red Sox Won
		y							
1	y	y							
2	y	n							
3	n	n							
4	n	y							
5	y	y							
6	y	y							
7	y	n							
8	n	y							
9	n	n							
10	y	y							
11	y	n							
12	n	y							
13	n	n							
...	y	y							




**User:** Should I wear galoshes?

**System:** Yes.


**User:** Why

**System :** Just listen to me, I have an “intuition” and I am right more than 60% of the time.

# Is **Intelligence** simply finding the functions that map input data to the output data







- One view of Learning is that its about **finding functions** that relate input data to output data.
- The **complexity of the function** that describes the mapping is one **measure of the intelligence** challenge.
- Neural Networks can find complex non-linear functions.

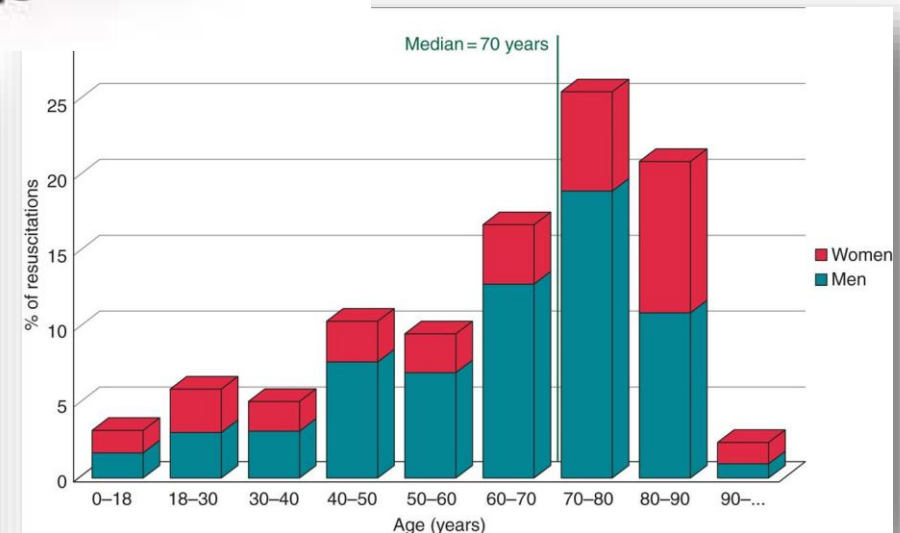




- But does the function describe a data set or does it describe the data or the mechanism that generated the data in human understandable terms. – Does it give the why?

# Learning how to Play Go with a complex Neural Network. Can it explain what it does and why?



# Data, Reason and Expected Value

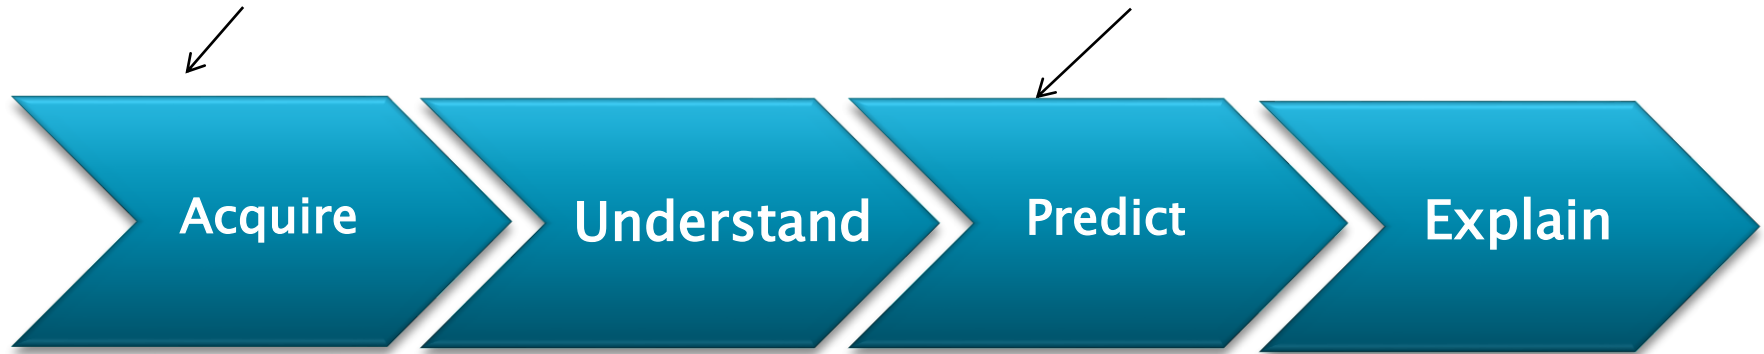




# The Holy Grail in AI: Autonomously Learn to Understand, Predict and Explain.

*What data even matters?*

*What are likely good/bad answers?*




*What does the data mean?  
How is it useful and why?*

*Why good or bad? Explain, relate to  
how user understands...what they  
already know.*

Interesting to observe how dependent this  
process is on **human Cognition.**

# Where Watson Fit...



... an interesting point along the spectrum...

# Jeopardy!: A great challenge for advancing AI. Specifically in the areas of natural language understanding

Broad/Open  
Domain

Complex  
Language

High  
Precision

Accurate  
Confidence

High  
Speed

\$200

If you're standing, it's the *direction* you should look to check out the wainscoting.

\$1000

I tell you it was so cold today... (How cold was it?) It was so cold, I wished we were back in 64 when he was emperor. Hot times, if you know what I mean.

\$800

Seems this *perp* was the first murderer in the Bible and to top it off he *iced* his own brother

\$600

In cell division, mitosis splits the nucleus & cytokinesis splits this liquid *cushioning* the nucleus

\$2000

Of the 4 countries in the world that the U.S. does not have diplomatic relations with, the one that's farthest north

# Categories are not as revealing as they may seem

Watson used statistical methods to discover that Jeopardy! categories were only weak indicators of the answer type.

## U.S. CITIES

St. Petersburg is home to Florida's annual tournament in this game popular on shipdecks  
(Shuffleboard)

Rochester, New York grew because of its location on this  
(the Erie Canal)

## Country Clubs

From India, the shashpar was a multi-bladed version of this spiked club  
(a mace)

A French riot policeman may wield this, simply the French word for "stick"  
(a baton)

## Authors

Archibald MacLeish? based his verse play "J.B." on this book of the Bible  
(Job)

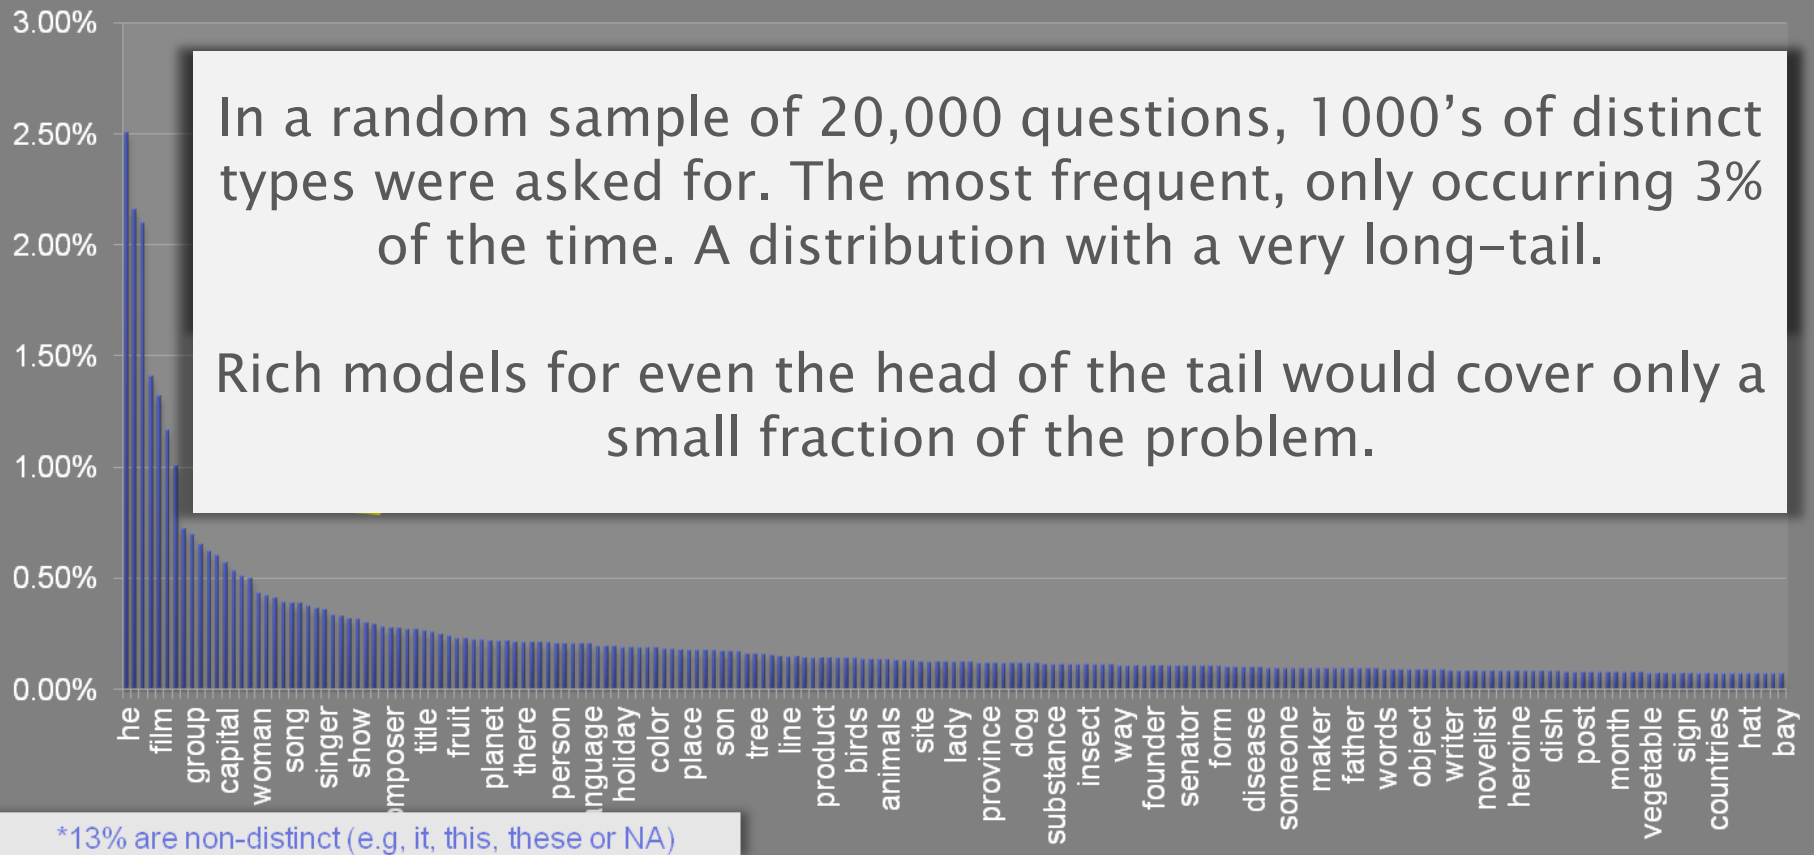
In 1928 Elie Wiesel was born in Sighet, a Transylvanian village in this country  
(Romania)

# Open-Domain QA Challenges

## ask about a wide variety

In a random sample of 20,000 questions, 1000's of distinct types were asked for. The most frequent, only occurring 3% of the time. A distribution with a very long-tail.

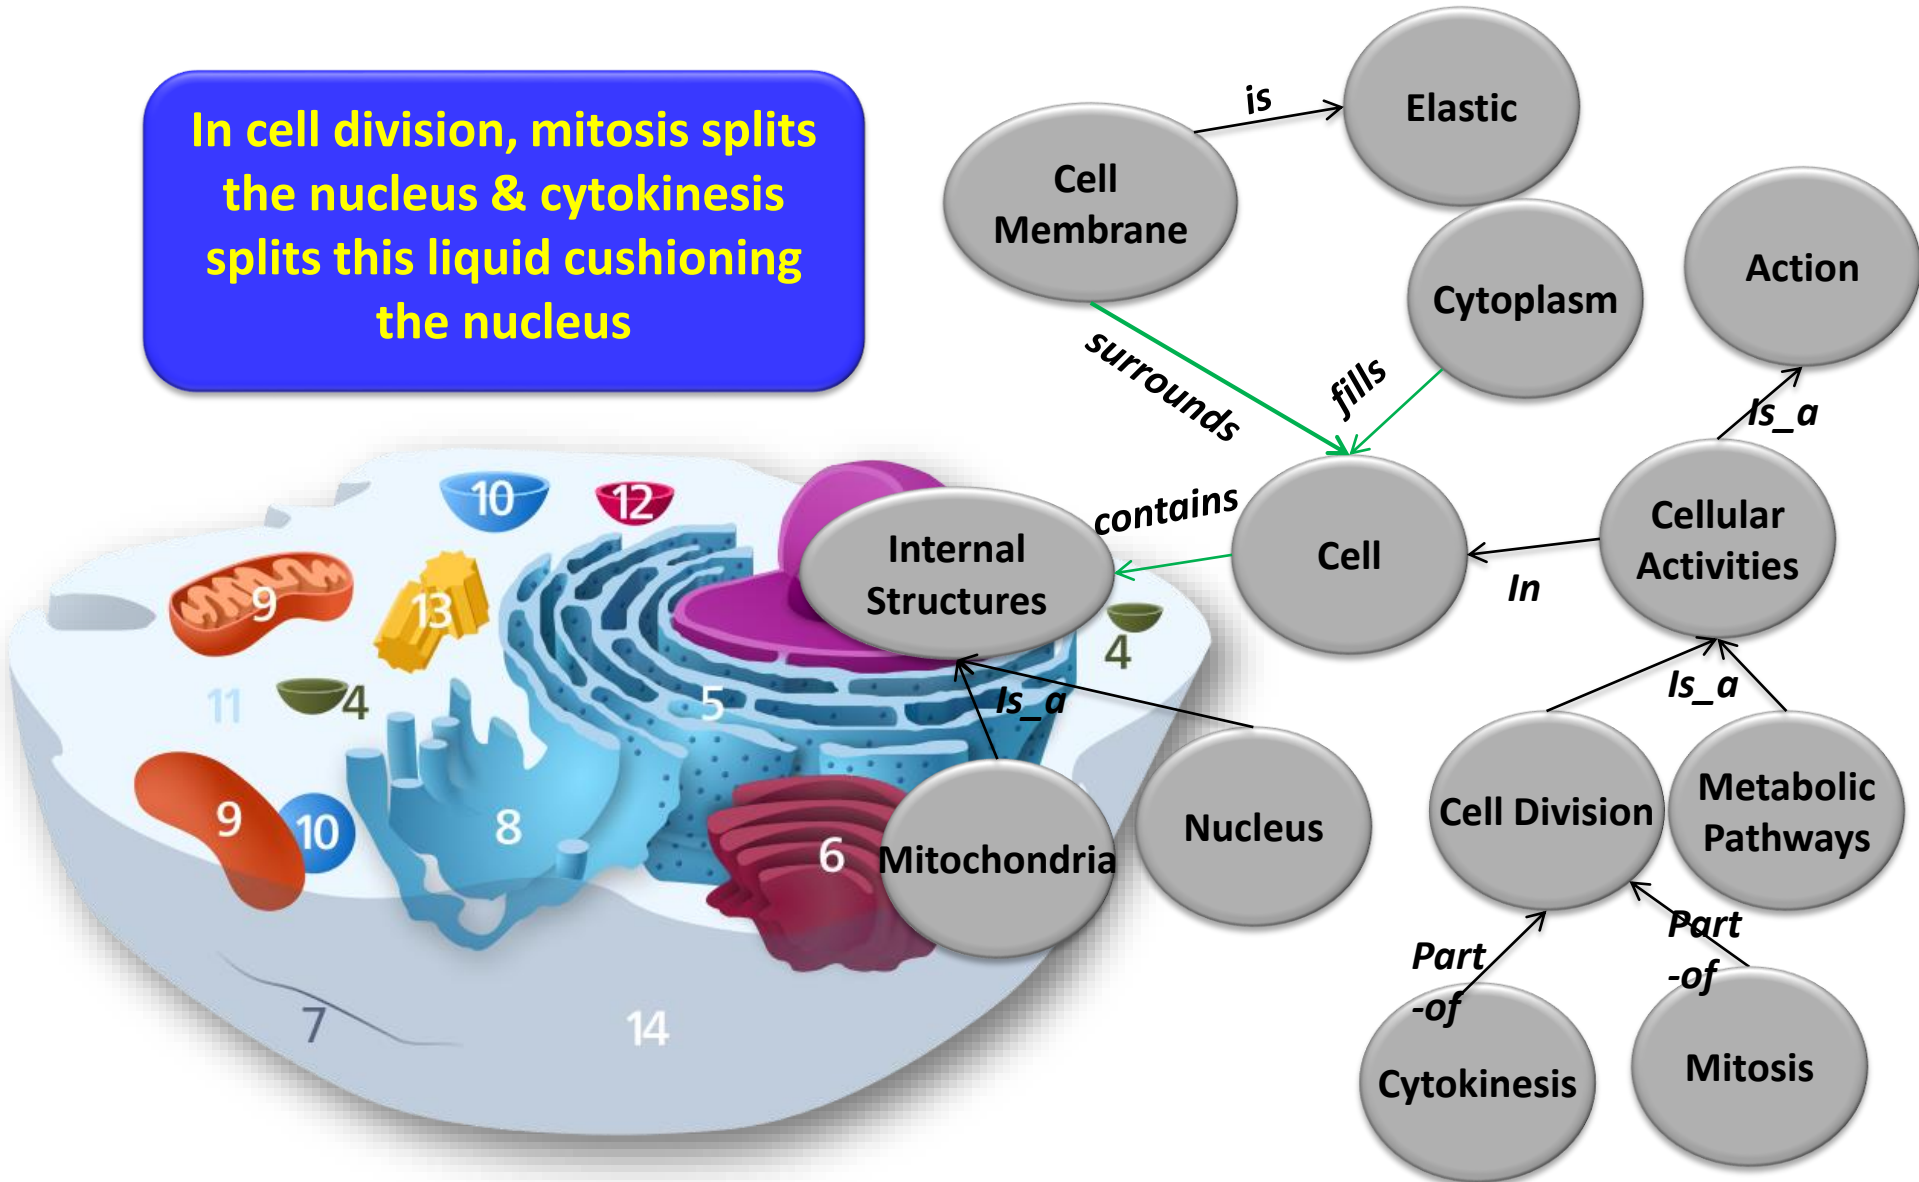
Rich models for even the head of the tail would cover only a small fraction of the problem.






Domain Theories enable deep reasoning with explicable conclusions, but are difficult to build and map to language for broad domains.

In cell division, mitosis splits the nucleus & cytokinesis splits this liquid cushioning the nucleus

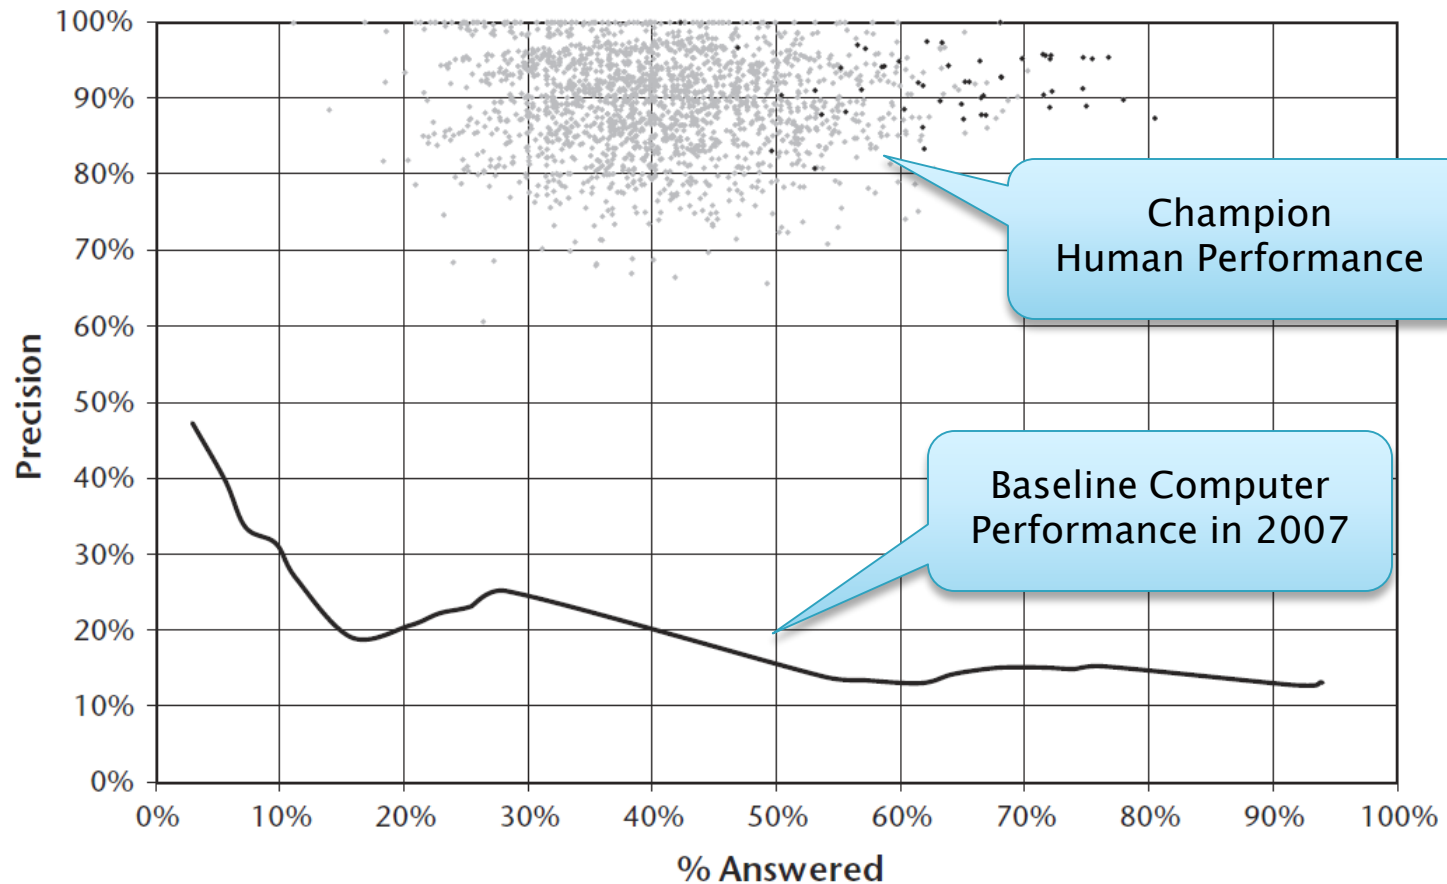




# Inducing “Meaning” From Context (Distributional Semantics)



# How good do you have to be to Win



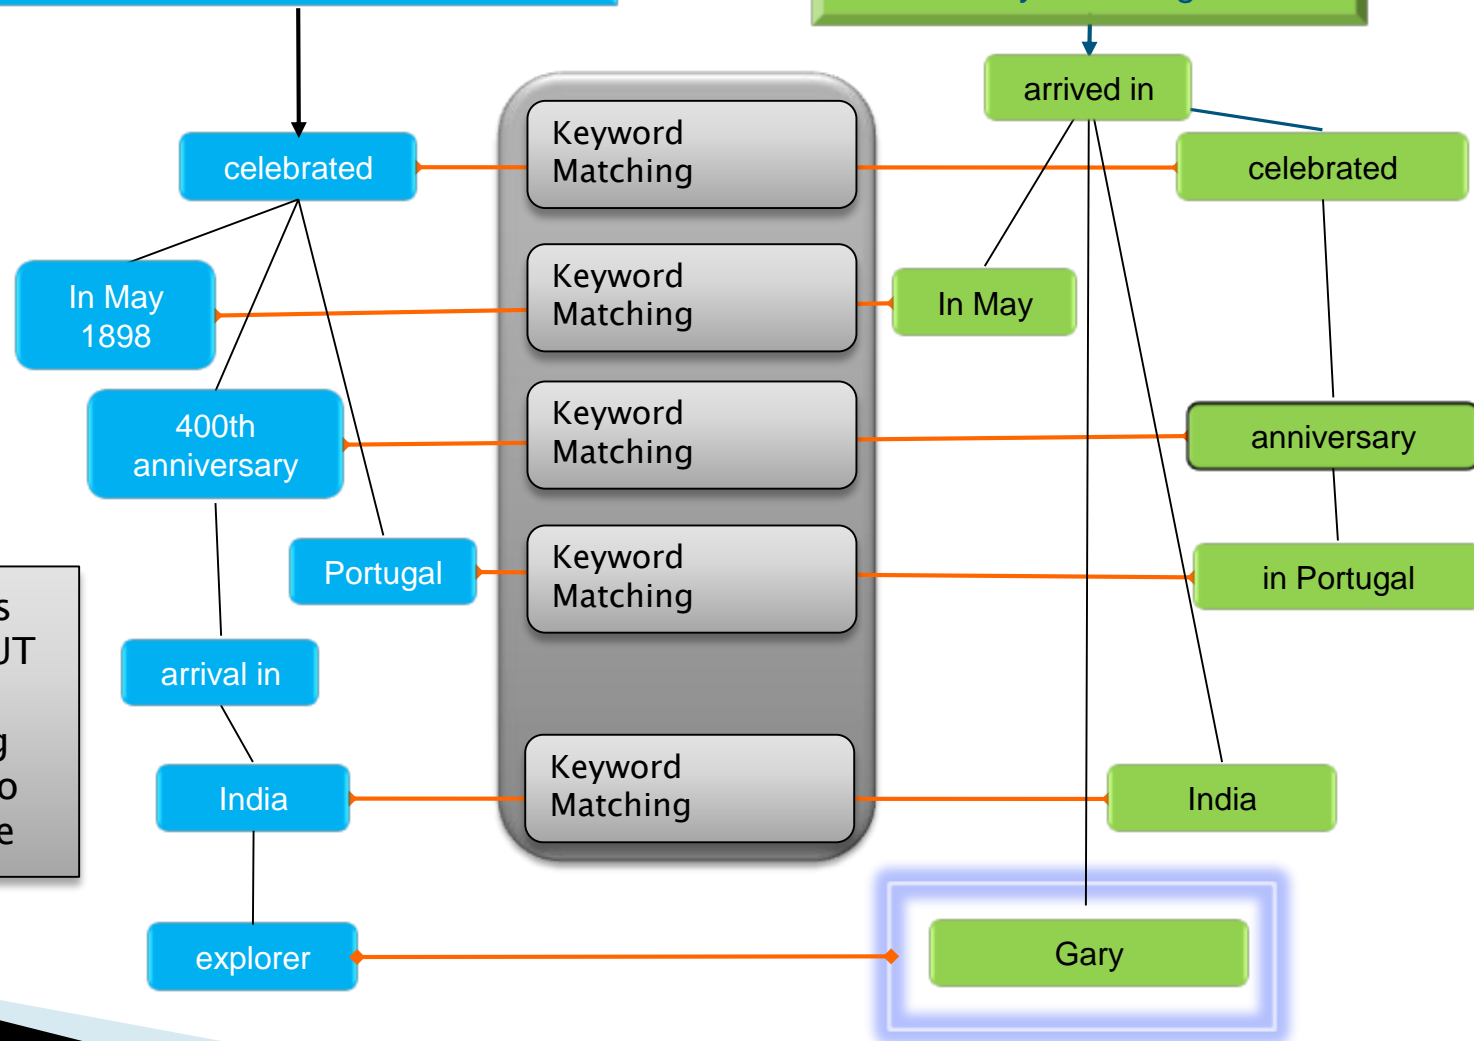
Ferrucci, et. al. AI Magazine, *Building Watson: An Overview of the DeepQA Project*

Figure 4. Baseline Performance.

# Find and Score Evidence: Weak Features

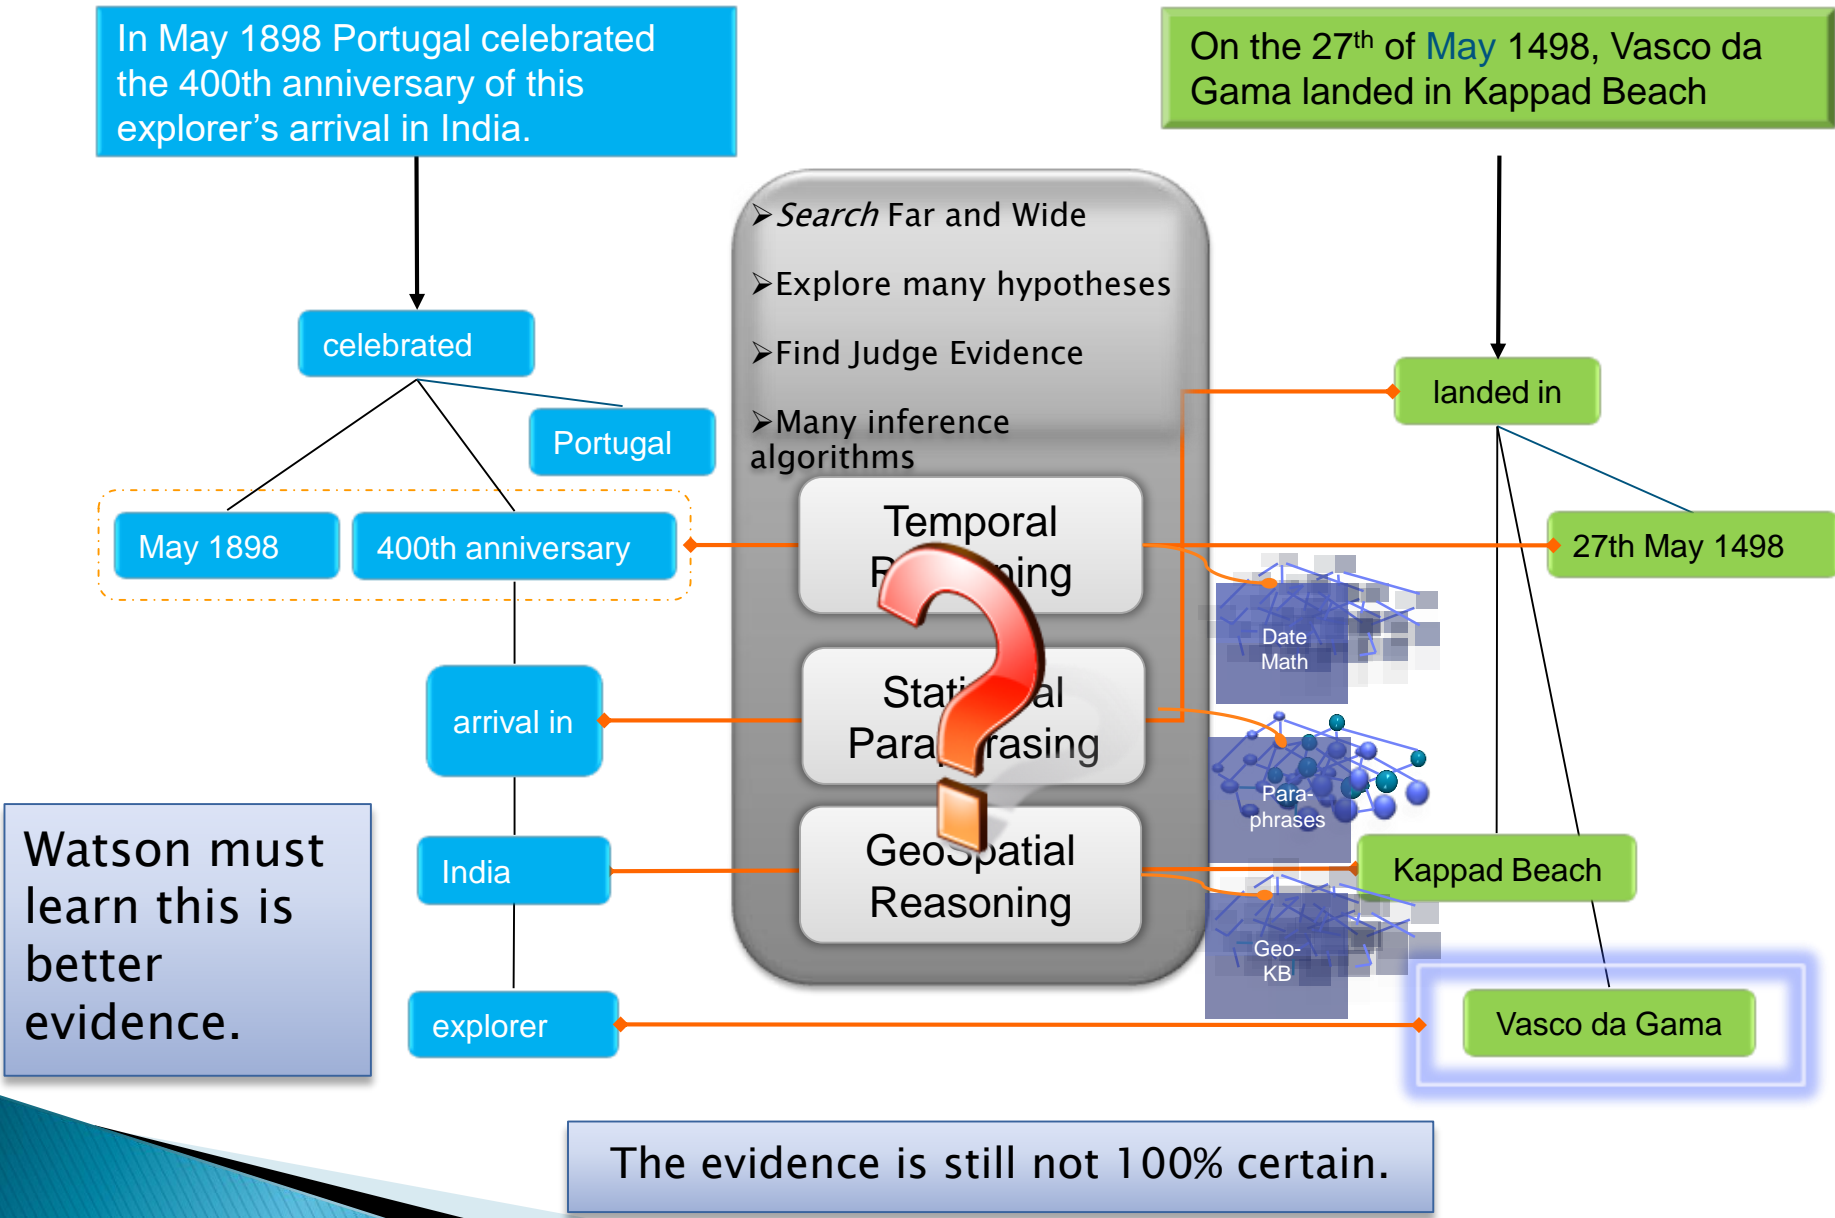
In May 1898 Portugal celebrated the 400th anniversary of this explorer's arrival in India.

In May, Gary arrived in India after he celebrated his anniversary in Portugal.

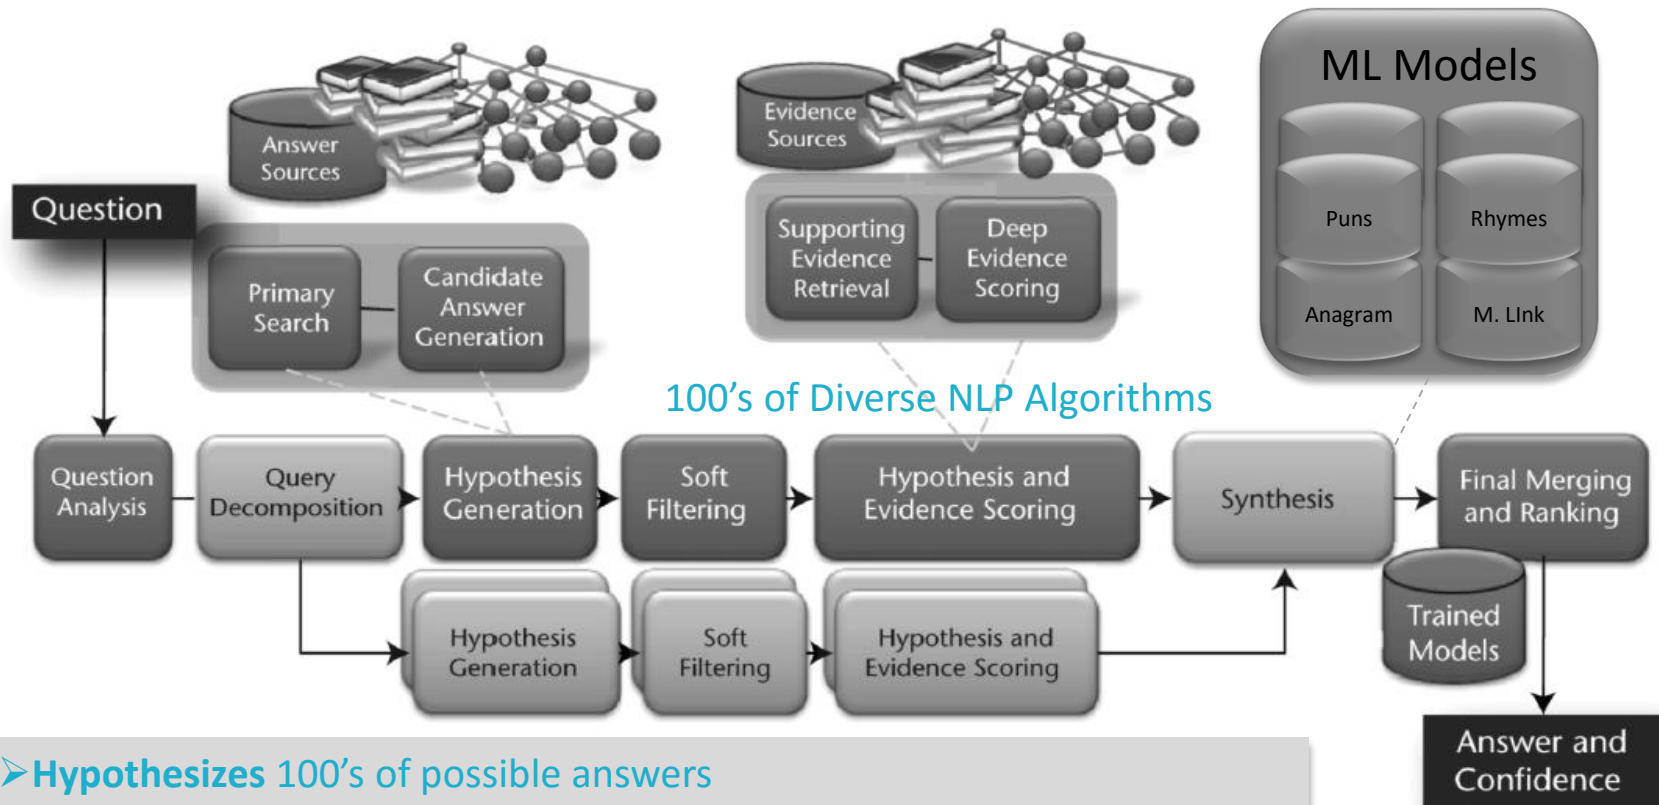


This evidence suggests "Gary" is the answer BUT the system must learn that keyword matching may be weak relative to other types of evidence

# Find and Rate Evidence: Better Features



# The Watson Architecture: How it Worked to Play Jeopardy!




- **Hypothesizes** 100's of possible answers
- Finds and **scores** 10'000s of pieces of **evidence**.
- **Learns** how best to combine 100's of diverse NLP algorithms.
- **Ranks** highest **probability** of being right at the top.

Figure 6. DeepQA High-Level Architecture.

Ferrucci, et. al. AI Magazine, *Building Watson: An Overview of the DeepQA Project*



35+ AI Scientists and Software Engineers at IBM Research and University partners, built on decades of groundwork in Search, ML & NLP, performed >8000 documented experiments over 4 years and broke new ground in question-answering to tackle Jeopardy. This required dramatic changes in culture, management and technical approaches.



Ferrucci et. al. **AI Magazine**, *Building Watson: An Overview of the DeepQA Project*

# The Jeopardy Contest: Human vs. Machine



*Both “Disconnected”*

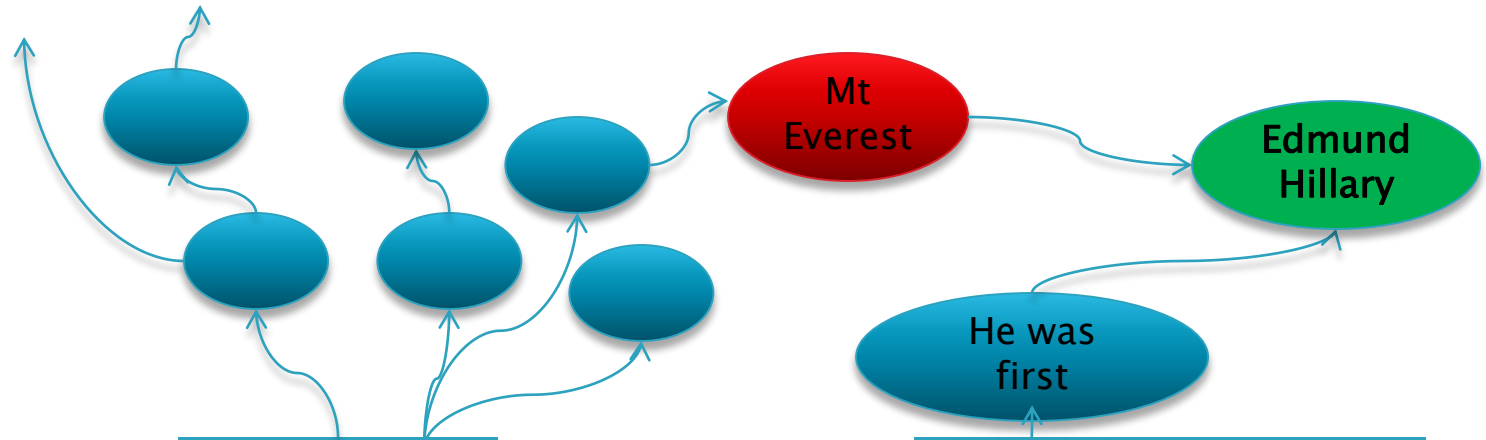


- 2,880 CPUs/15TB of RAM.....
- Size of 10 Refrigerators.....
- 80 KW of Electricity.....
- 20 Tons of Cooling.....
- 4 Yrs + 2 million books of content..

- 1 Brain
- Fits in a shoe box
- Tuna Fish Sandwich + Glass of Milk
- Hand Fan
- ~30 years of human learning



# Some of the hardest questions required multiple hops over Missing Links



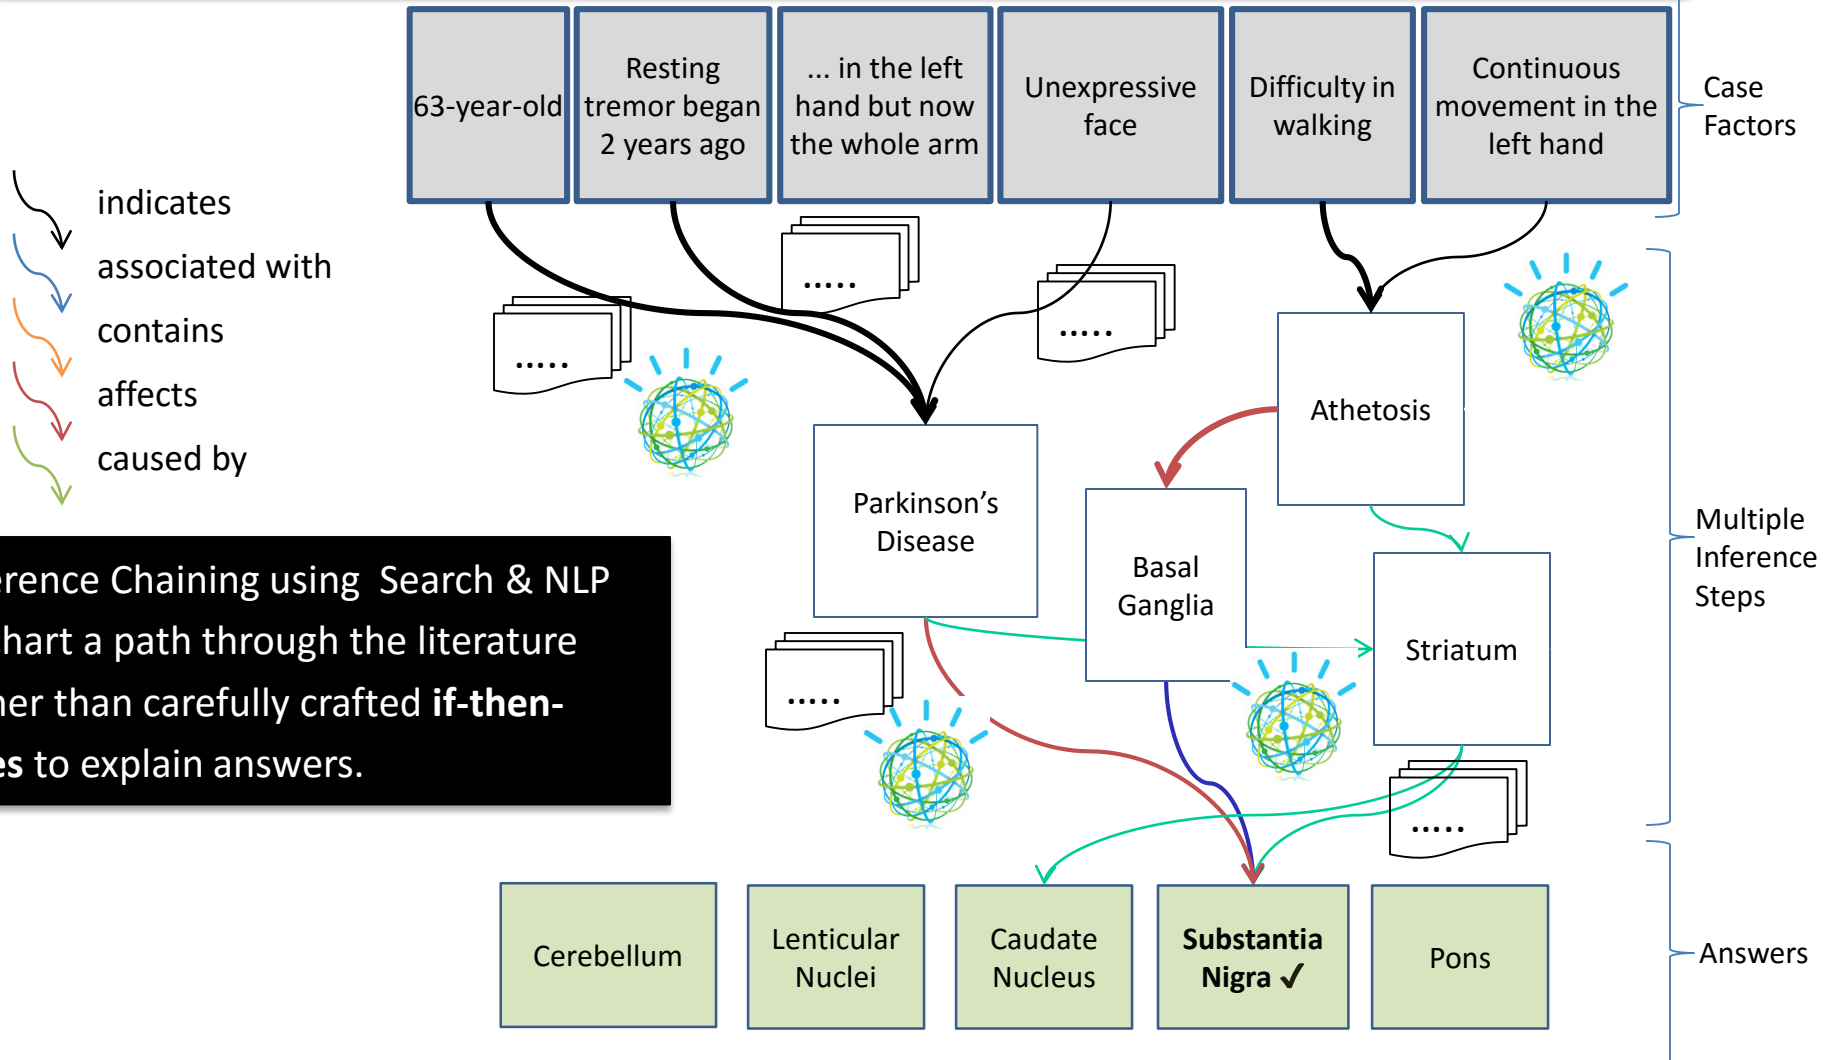
On hearing of the discovery of George Mallory's body, he told reporters he still thinks he was first.

## These Led to the Idea Behind Watson Paths...

# WatsonPaths Links Existing Text to Explain Connections

A 63-year-old patient is sent to the neurologist with a clinical picture of resting tremor that began 2 years ago. At first it was only on the left hand, but now it compromises the whole arm. At physical exam, the patient has an unexpressive face and difficulty in walking, and a continuous movement of the tip of the first digit over the tip of the second digit of the left hand is seen at rest.

What part of his nervous system is most likely affected?



# Reflections on *Watson*

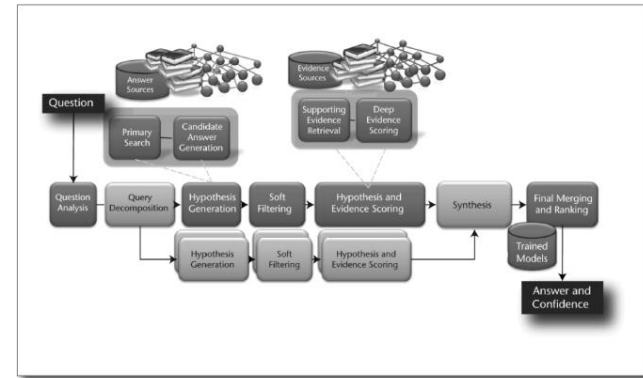



Figure 6. DeepQA High-Level Architecture.



M. Minsky

- Innovation through New Architectures
  - Complex Machines vs. Single Algorithms
  - Combining **Science**, **Engineering** and **Research Management**
  - To rapidly extend beyond what was thought possible
- Combining a Diversity of Skills and Methods
  - A wide diversity of skilled researchers
  - A wide diversity of loosely integrated techniques
  - Combined with **Machine Learning** to balance and integrate efficiently
- **Yet** to program a computer to “Understand” Language
  - Computers still can not fluently converse with humans about what they read
  - Yet to enable the machine to learn human-compatible logical models underlying language
  - Yet to dynamically engage human thought in interpreting the data and extending the understanding

Deep Understanding

**TAKING A GIANT STEP BACK IN  
ORDER TO LEAP FORWARD**



# We need machines to go beyond the words and their patterns. We need them to UNDERSTAND.

- Imagine a collaborative thought partner
- That can truly understand and explain its understanding in fluent human language dialog tailored to the listener's model of the world.
- No such capability exists today...
- Cracking this problem will dramatically accelerate access to actionable understanding and provide for better, faster and more transparent decision making.

**Dave:** I'd like to better understand Stem Cell research.

**System:** I can help you with that.

**Dave:** I'd like to discuss **these** articles on the role of stem cells in organ regeneration both from a technical perspective and from an ethical perspective

**System:** Ok. This content suggests that stem cells can be very effective at organ regeneration. The technology works because stems cells contain.....and act to.... Open challenges remain regarding....Ethical concerns include....

**Dave:** Why though do stem cells gravitate to the area of injury? What does it mean for them to differentiate?

**System:** I explained this already, but I can do better. Would you like me to rephrase my explanation in terms you are more likely to be familiar with.

**Dave:** Yes

**System:** Imagine Stem Cells are similar to Classes in a programming language, like Java. These cells can be specialized to perform more specific functions much the same way that...

# Yet to program a Computer to Understand.

## A 1<sup>st</sup> Grade Story...

### John and Mary

*John and Mary were running a race. John fell. He hurt his knee. Mary looked back. Mary wanted to win. If she kept running she would win. Mary stopped. She ran back. She helped John up.*

*Consider how you know...*

- *Who was running a race?*
- *Who wanted to win?*
- *Who hurt his knee?*

*Consider how you know...*

- *What was Mary looking at?*
- *What did Mary stop?*
- *Where did she run back to?*
- *What did Mary **decide** to do?*

*Consider how you might predict*

- *Did Mary **Win**?*
- *What was the weather?*
- *How old are John and Mary?*

# Interesting to Consider....

- ▶ Data Sets
  - 700 Stories for human readers: K – 3<sup>rd</sup> Grade
- ▶ A “Linguistic Pattern–Based” QA System
  - Using similar techniques to those found in today’s language AI
- ▶ Performance Metric
  - Accuracy = % of questions correct answered.

	Grade K	Grade 1	Grade 2	Grade 3
# Questions	218	321	334	310
Random	50.0%	33.3%	33.3%	33.3%
Accuracy	93.6%	61.9%	50.0%	39%

The hard ones don’t have a “proximate text match” in the story or on the internet. They require **Understanding and Reasoning**. There are more of those as you progress through the grades.

Yet a computer beat the best humans at Jeopardy with near 80% accuracy. The best chess player and the best GO player. Can you computer convince you it **understands** at even a 3<sup>rd</sup> grade level?

# Simple Stories. Questions By Grade...

## Grade K

*Mostly fact retrieval w/ some temporal ordering*

Sometimes, volcanoes erupt. That means they explode. Smoke and hot liquid rock come out of the volcano.

**Q: What comes out of the volcano when a volcano erupts?**

Maria felt very warm. She needed to take a break. She sat down under a tree. It was cooler in the shade. Maria smiled. She waved to her friends. "Cool off with me in the shade!" said Maria.

**Q: What was the weather like at the end of the story?**

## Grade 1

*Synthesizing and justifying*

Misty does not want other dogs to bother or come near her owner. What part of the passage shows us this is true?

- A: If Misty hears a noise, she stops. She raises her ears and listens.
- B: Misty barks if she sees another dog. She barks loudly to protect me.
- C: Misty runs to her bowl and drinks water. I give her some food to eat.

## Grade 2

*Reasoning and inference*

Read these sentences from the text.

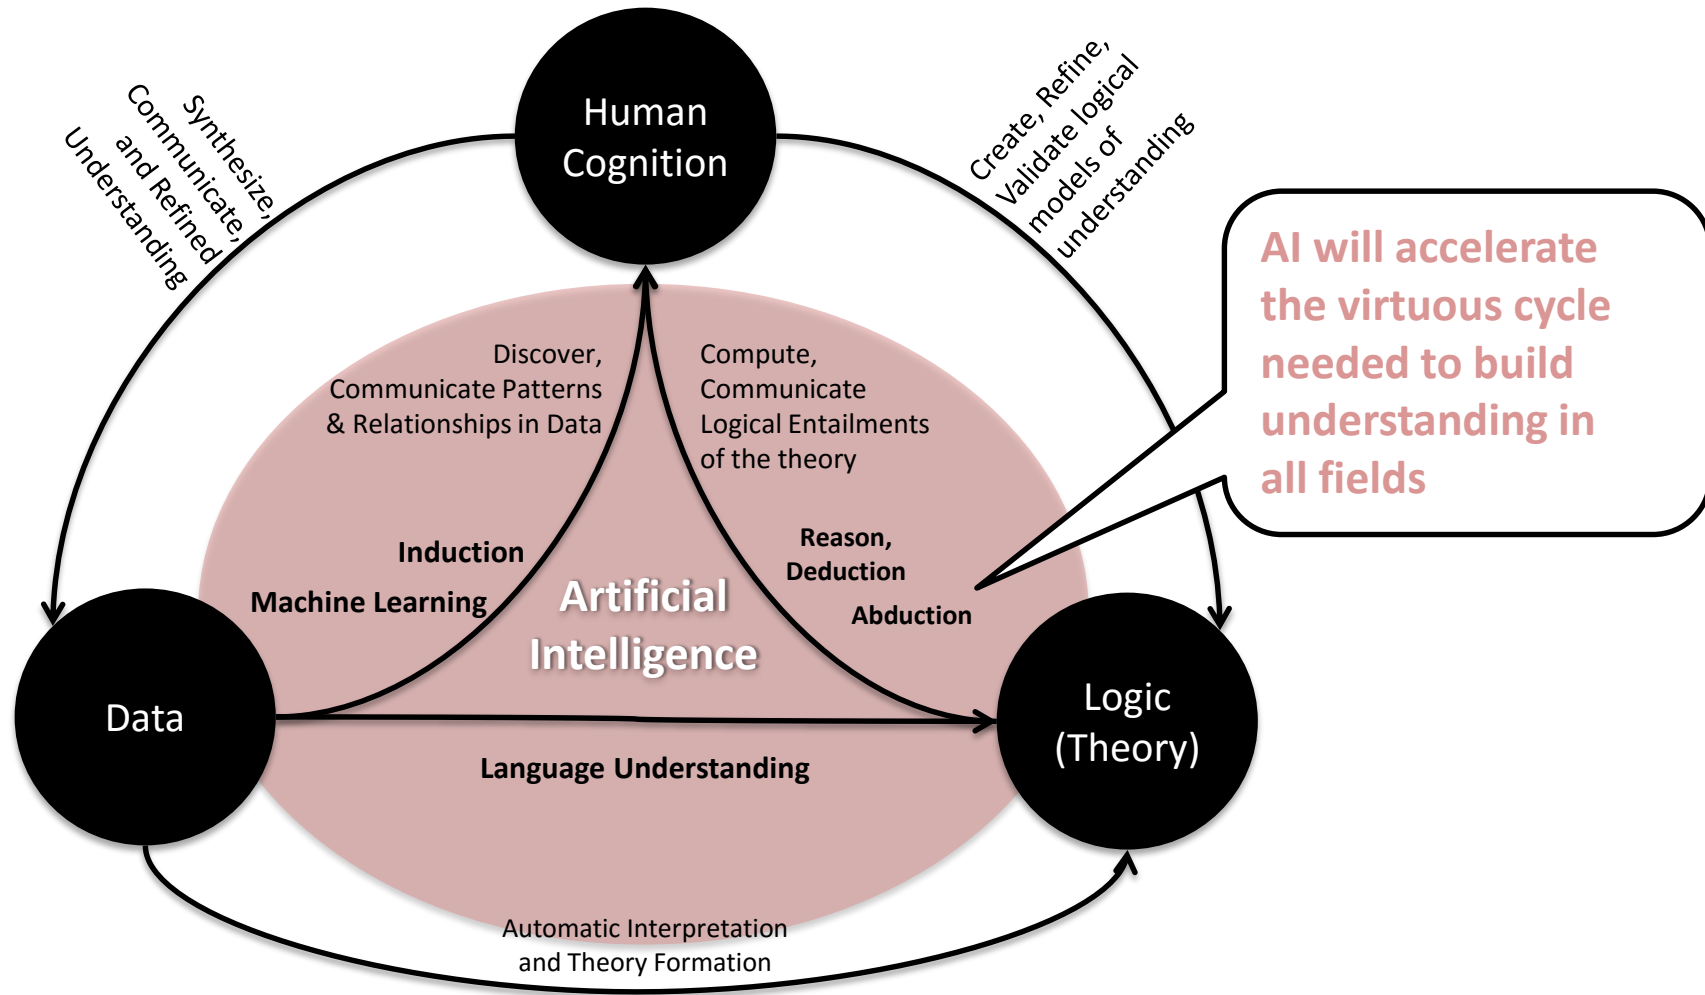
*"About three billion crayons are sold in the United States each year. This is enough to circle the earth at least four times!"* Why does the author include this fact in the article?

- A: to show how small the earth is
- B: to show how popular crayons are
- C: to show how long crayons are

# True Thought Partnership

## The Grand Collaboration Between Mind and Machine

The future of AI will produce a collaboration between **Human Cognition, Data & Theory** aimed at accelerating our ability to develop and share UNDERSTANDING.



**Thank You**