


IBM Power Systems™

IBM Systems & Technology Group



Watson



Patrick O'Rourke
pmorour@us.ibm.com
Executive Briefing Center


Power is performance redefined

Deliver IT services faster,
with higher quality,
and with superior economics

www.ibm.com/power



From Science Fiction to Reality




Watson takes on Jeopardy!

Advanced computing system has potential to take business intelligence to a new level

- **Date: February 14 / 15 / 16 2011**
- **IBM Research project named “Watson”**
- **Competition with humans at the game of Jeopardy:**
 - **Human vs. Machine contest.**
- **Competition:**
 - **Ken Jennings & Brad Rutter**
 - **Two most successful Jeopardy contestants of all time**

Watson's Popularity...



2011 Webby Person of the Year!



EMC 2011 Data Hero Awards



EMC WORLD 2011—LAS VEGAS—May 9, 2011 — Technology Application award to IBM Watson

- Awards honor those delivering unique and innovative solutions and techniques for using Big Data to profoundly impact individuals, organizations, industries and the world.
- Awards are selected by a prestigious independent judge panel

The Jeopardy! Challenge: *A compelling and notable way to drive and measure the technology of automatic Question & Answering*

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

The first person mentioned by name in 'The Man in the Iron Mask' is this hero of a previous book by the same author.



Word Plays


Puns

Rhymes

Lyrics

Dual meanings

Watson Overview



Watson History.

- 3+ years development by IBM scientists
- Software: IBM Research Software Stack
- Hardware: Power Systems

Why Jeopardy?

- Grand challenge for a computing system
- Broad range of subject matter,
- Speed at which contestants must provide both accurate responses
- Determine a confidence they are correct

Tale of the Tape



Game	Chess
Architecture	Specialized Hardware and Power2 SC
# Cores	30 + 480 ASICs
Hardware	Highly Specialized
Processing	Mathematical / Sturctured
Data Analysis	<ul style="list-style-type: none"> ▪ Finite number of possible moves and countermoves ▪ Structured data ▪ Mathematical probability

Watson Info.....

Hardware:

- Cluster: 90 Power 750 (2880 Cores) @ 3.55 GHz
- 80 Teraflops
- 88 Compute nodes, 2 I/O nodes
 - ❖ 15TB of memory
- 4 SAS Storage drwr & 2 xCAT Servers

Software:

- SLES 11, JAVA, CNFS, GPFS, xCat, Apache Hadoop

Middleware:

- Apache UIMA (Open source)

Applications:

- DeepQA - Main analytical engine which ran on POWER 7
- Lenovo desktop : Voice synthesis, strategies for betting, buzzing in, clue selection & exchanging info with Jeopardy Computers
- Mac notebook: Avatar



Power 750 System



Watson Environment	
POWER7 Architecture	32 Cores @ 3.55 GHz 88 compute nodes 2 I/O nodes
DDR3 Memory	60 Nodes 128GB 30 Nodes 256GB
System Unit SFF Bays	6 HDD 146GB @ 15k
System Unit IO Expansion Slots	(1) FC 1983 Dual port 1Gb Ethernet (1) FC 5769 10Gb Fiber SR In the 2 I/O nodes: (1) FC 5903 SAS RAID Controller
Integrated Virtual Ethernet	Dual 10 Gb IVE / HEA
Storage Drawers	Four SAS enclosure FC 5888 <ul style="list-style-type: none"> ▪ Two per I/O node ▪ Twelve 300GB
IO Drawers	None

Watson's Sources of Information

Watson reads roughly 200 million pages of content (equivalent to one million books), written in natural human language ... in less than 3 seconds

Organized by topic, but additional information comes from actually reading the sentences


Unstructured Sources

Wikipedia (Full text)
IMDb
Encyclopedias
Dictionaries
Thesauri
Newswire Articles
Literary Works (including the Bible)

Structured Sources

DBpedia
Wordnet
YAGO

200 million pages structured and unstructured content
= 1 Million books
= 4 TB storage+16 TB memory



These structured data sources were typically used to obtain information about the English language, e.g. nouns, verbs, adjectives and adverbs grouped into sets of cognitive synonyms



Watson Odds and Ends...

22 Different Process Types

- Heavily Parallelized

389 Processes

- 199 C++
- 190 Java

Threads were heavily core multi-threaded

- Large memory: 256 GB System
- Smaller memory: 128 GB System


Watson Challenge.....




**How to Process / Analyze / Evaluate /
Prioritize all of this **Unstructured Data****

Data Problem.....

Structured Data



Unstructured Data



Easy Questions.....

$$\ln((12,546,798 * \pi)) ^ 2 / 305.992 = 1.0$$

Select *Payment* where *Owner*="David Jones" and *Type(Product)*="Laptop",

Owner	Serial Number
David Jones	45322190-AK

Serial Number	Type	Invoice #
45322190-AK	LapTop	INV10895

Invoice #	Vendor	Payment
INV10895	MyBuy	\$104.56

↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓
 David Jones
 David Jones =

↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓
 Dave Jones ≠
 David Jones

Hard Questions....

Computer programs are natively explicit, fast and exacting in their calculation..

Natural Language is implicit, highly contextual, ambiguous and often imprecise.

Structured

Person	Birth Place	Person	Organization
A. Einstein	ULM	J. Welch	GE

UnStructured

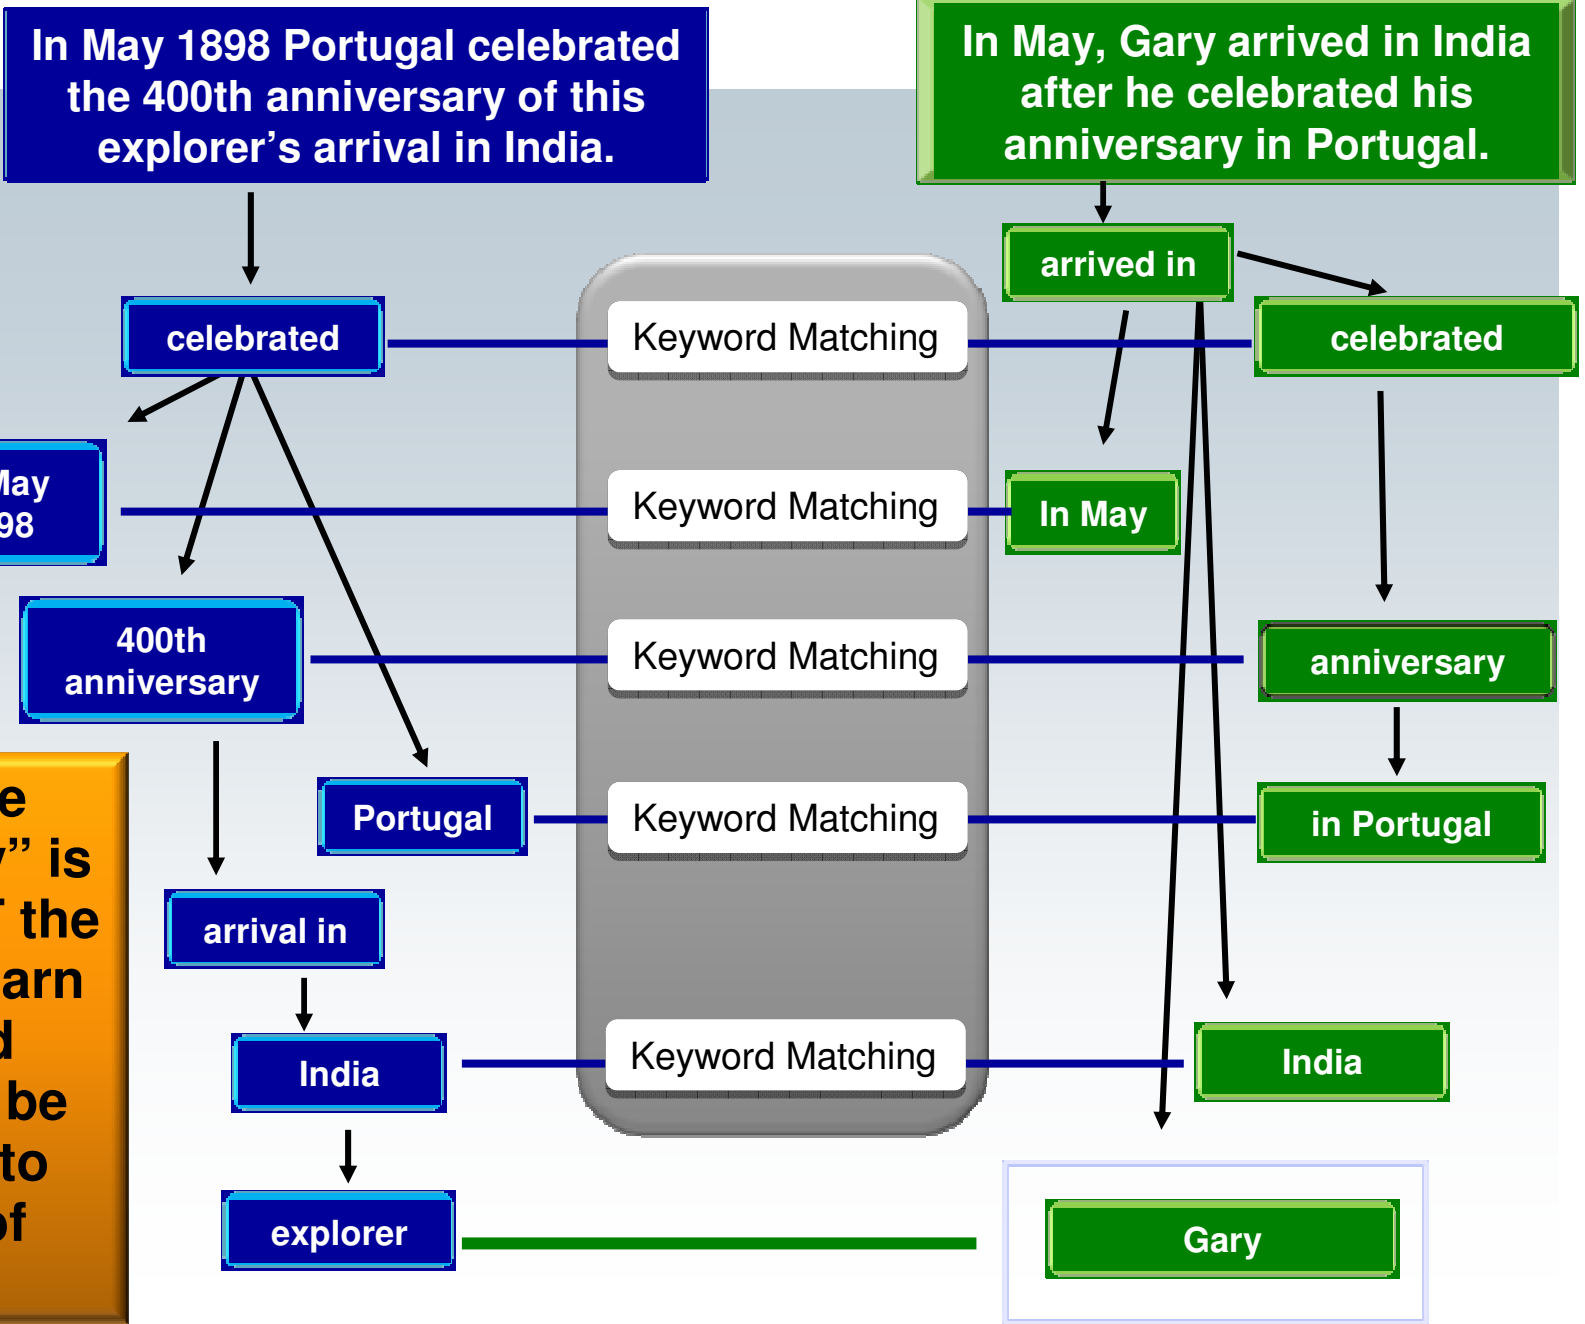
1. Where was X born?

Otto chose a water color of this city to send to Albert Einstein as a remembrance of Einstein's birthplace.

2. X ran this?

If leadership is an art then surely, Jack Welch proved himself a master painter during his tenure at this company..

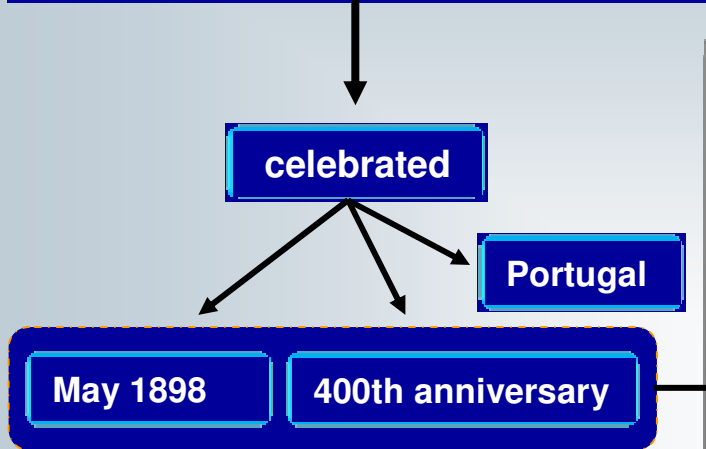
Keyword Evidence



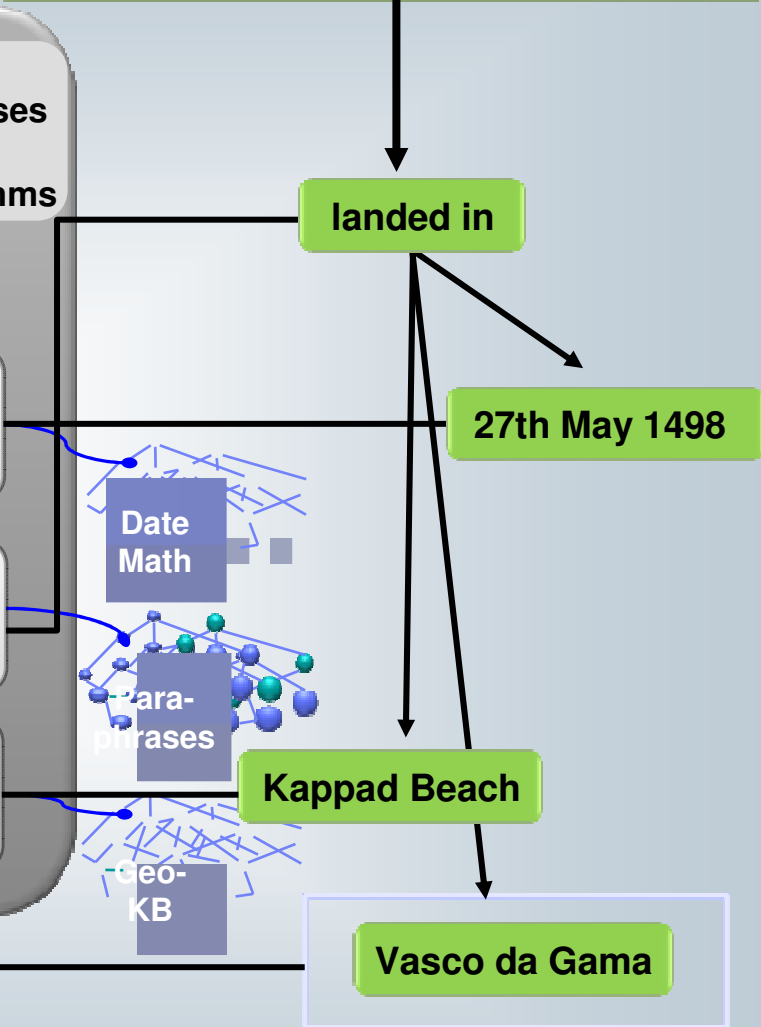
This evidence suggests "Gary" is the answer BUT the system must learn that keyword matching may be weak relative to other types of evidence

Deeper Evidence

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



On the 27th of May 1498, Vasco da Gama landed in Kappad Beach




- Search Far and Wide
- Explore many hypotheses
- Find Judge Evidence
- Many inference algorithms

Temporal Reasoning

Statistical Paraphrasing

GeoSpatial Reasoning



Stronger evidence can be much harder to find and score.




The evidence is still not 100% certain.

Best Human Performance.....

Each dot represents an actual historical human Jeopardy! game

Top human players are remarkably good.



Computers?

Winning Human Performance

Grand Champion Human Performance

2007 QA Computer System

More Confident

Less Confident

IBM Apache UIMA Hadoop, and DeepQA,

UIMA is the industry standard for Content Analytics

- Unstructured Information Management Architecture.

UIMA SDK was originally developed by IBM

- 2006: SDK available at alphaWorks®.
- 2008: IBM donated UIMA SDK to **Apache**
- Ongoing: Development by the Apache UIMA community

Apache UIMA add-on: **UIMA Asynchronous Scaleout (AS)**

- Provides the ability to scale out in a clustered environment

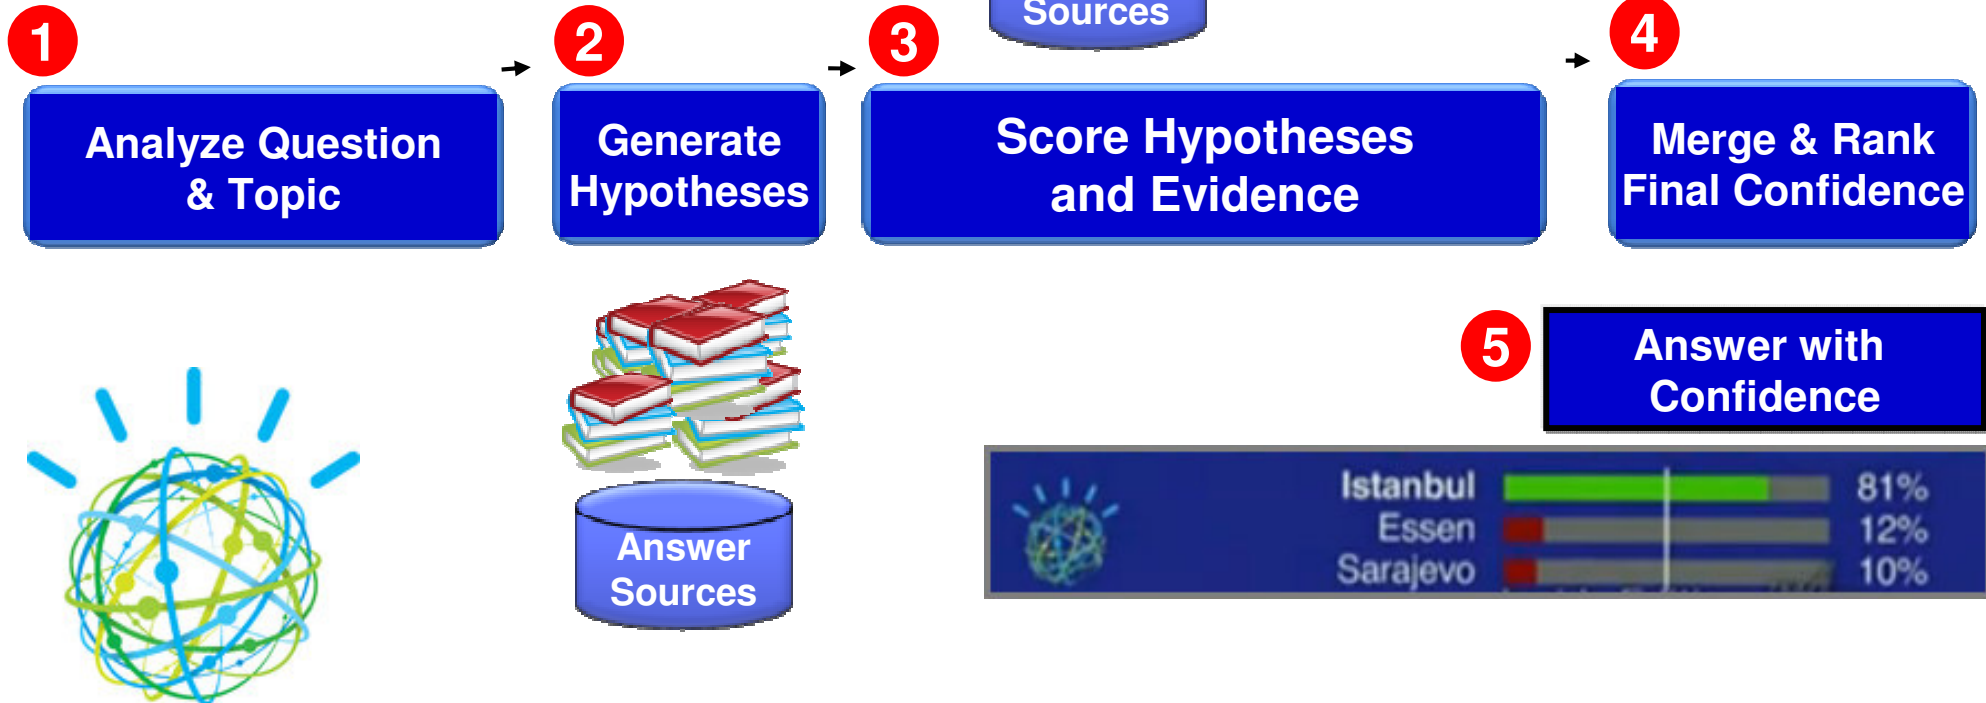
Apache Hadoop: Framework used by Watson to facilitate preprocessing the large volume of data, created in-memory datasets used at run-time

DeepQA: Collection of Algorithms

- Can be divided into independent parts, each executed by a separate processor / Computation is embarrassing parallel
- Gathers, evaluates, weighs and balances different types of evidence, delivering the answer with the best support it can find.

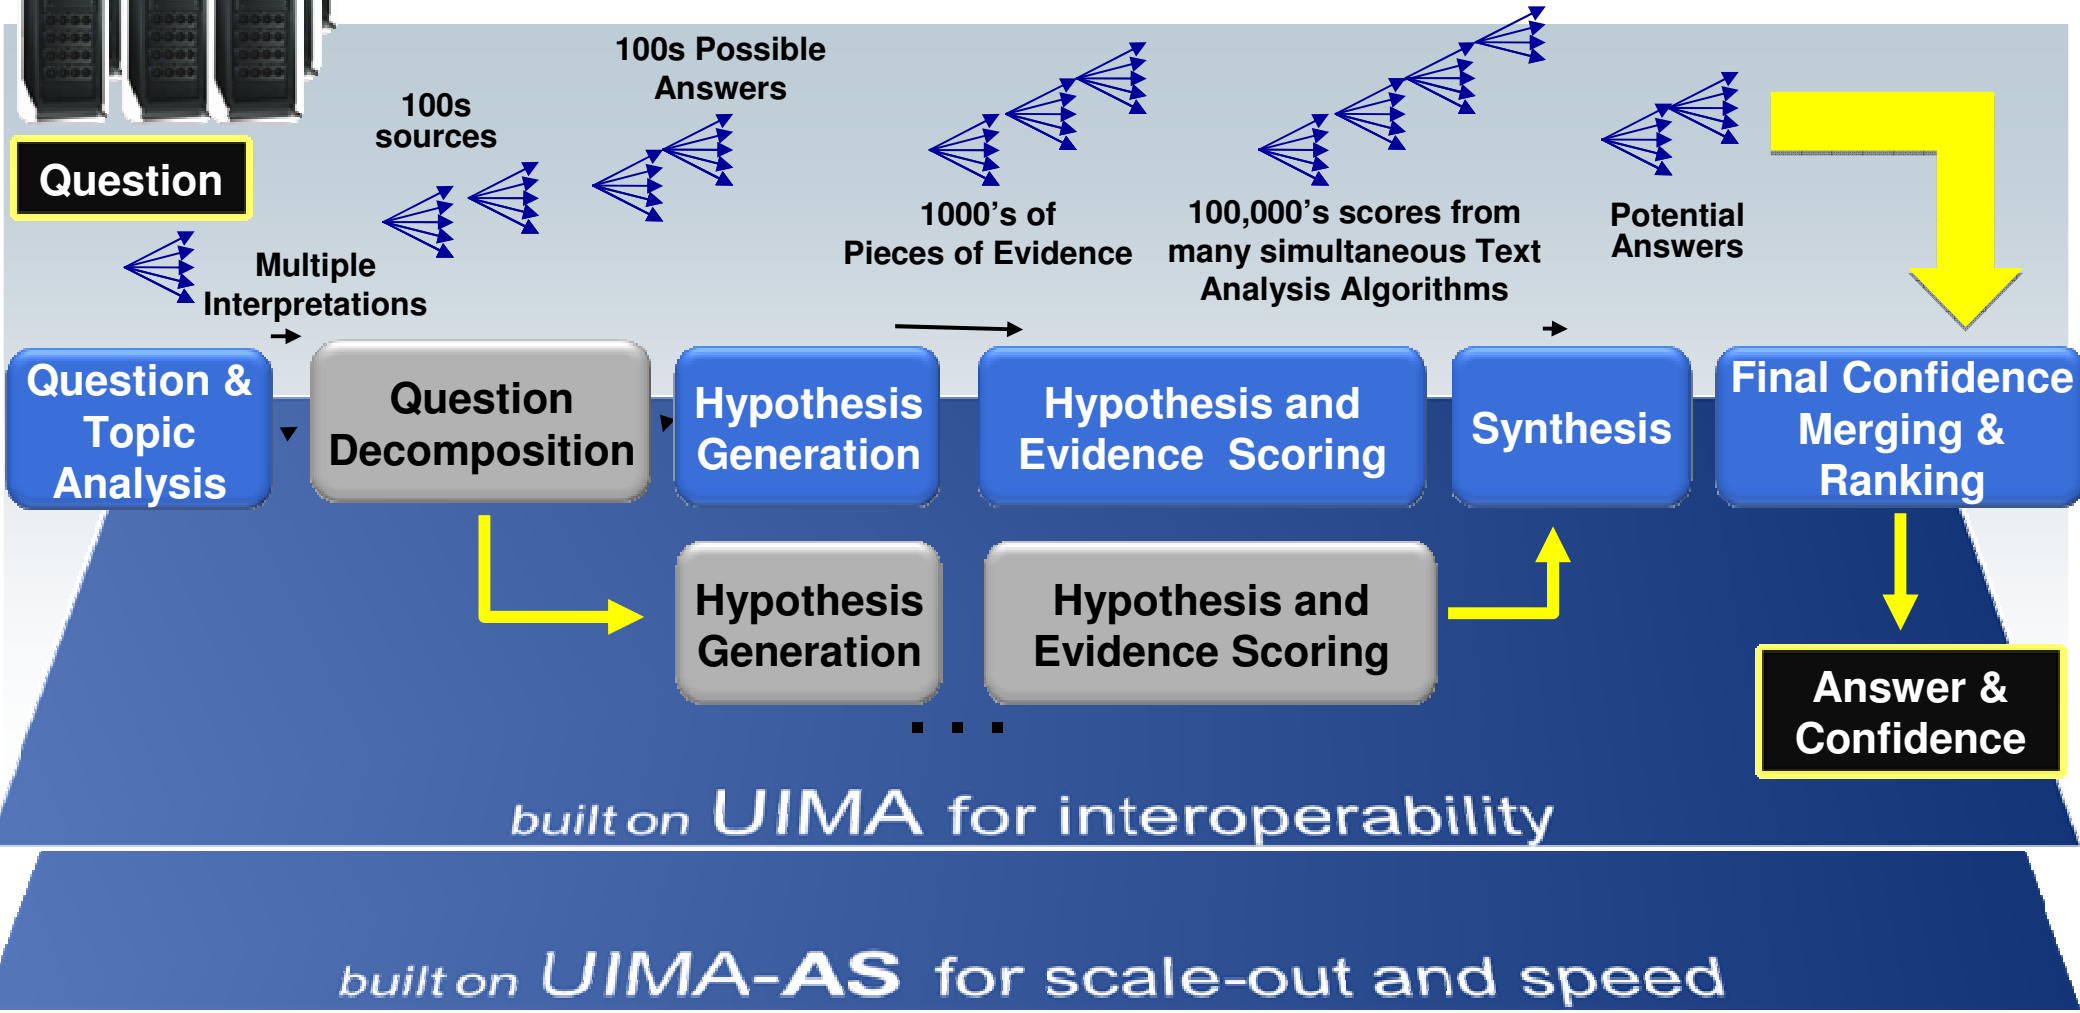
IBM DeepQA: How Watson answered questions

Each year the EU selects capitals of culture; one of the 2010 cities was this Turkish "meeting place of cultures"

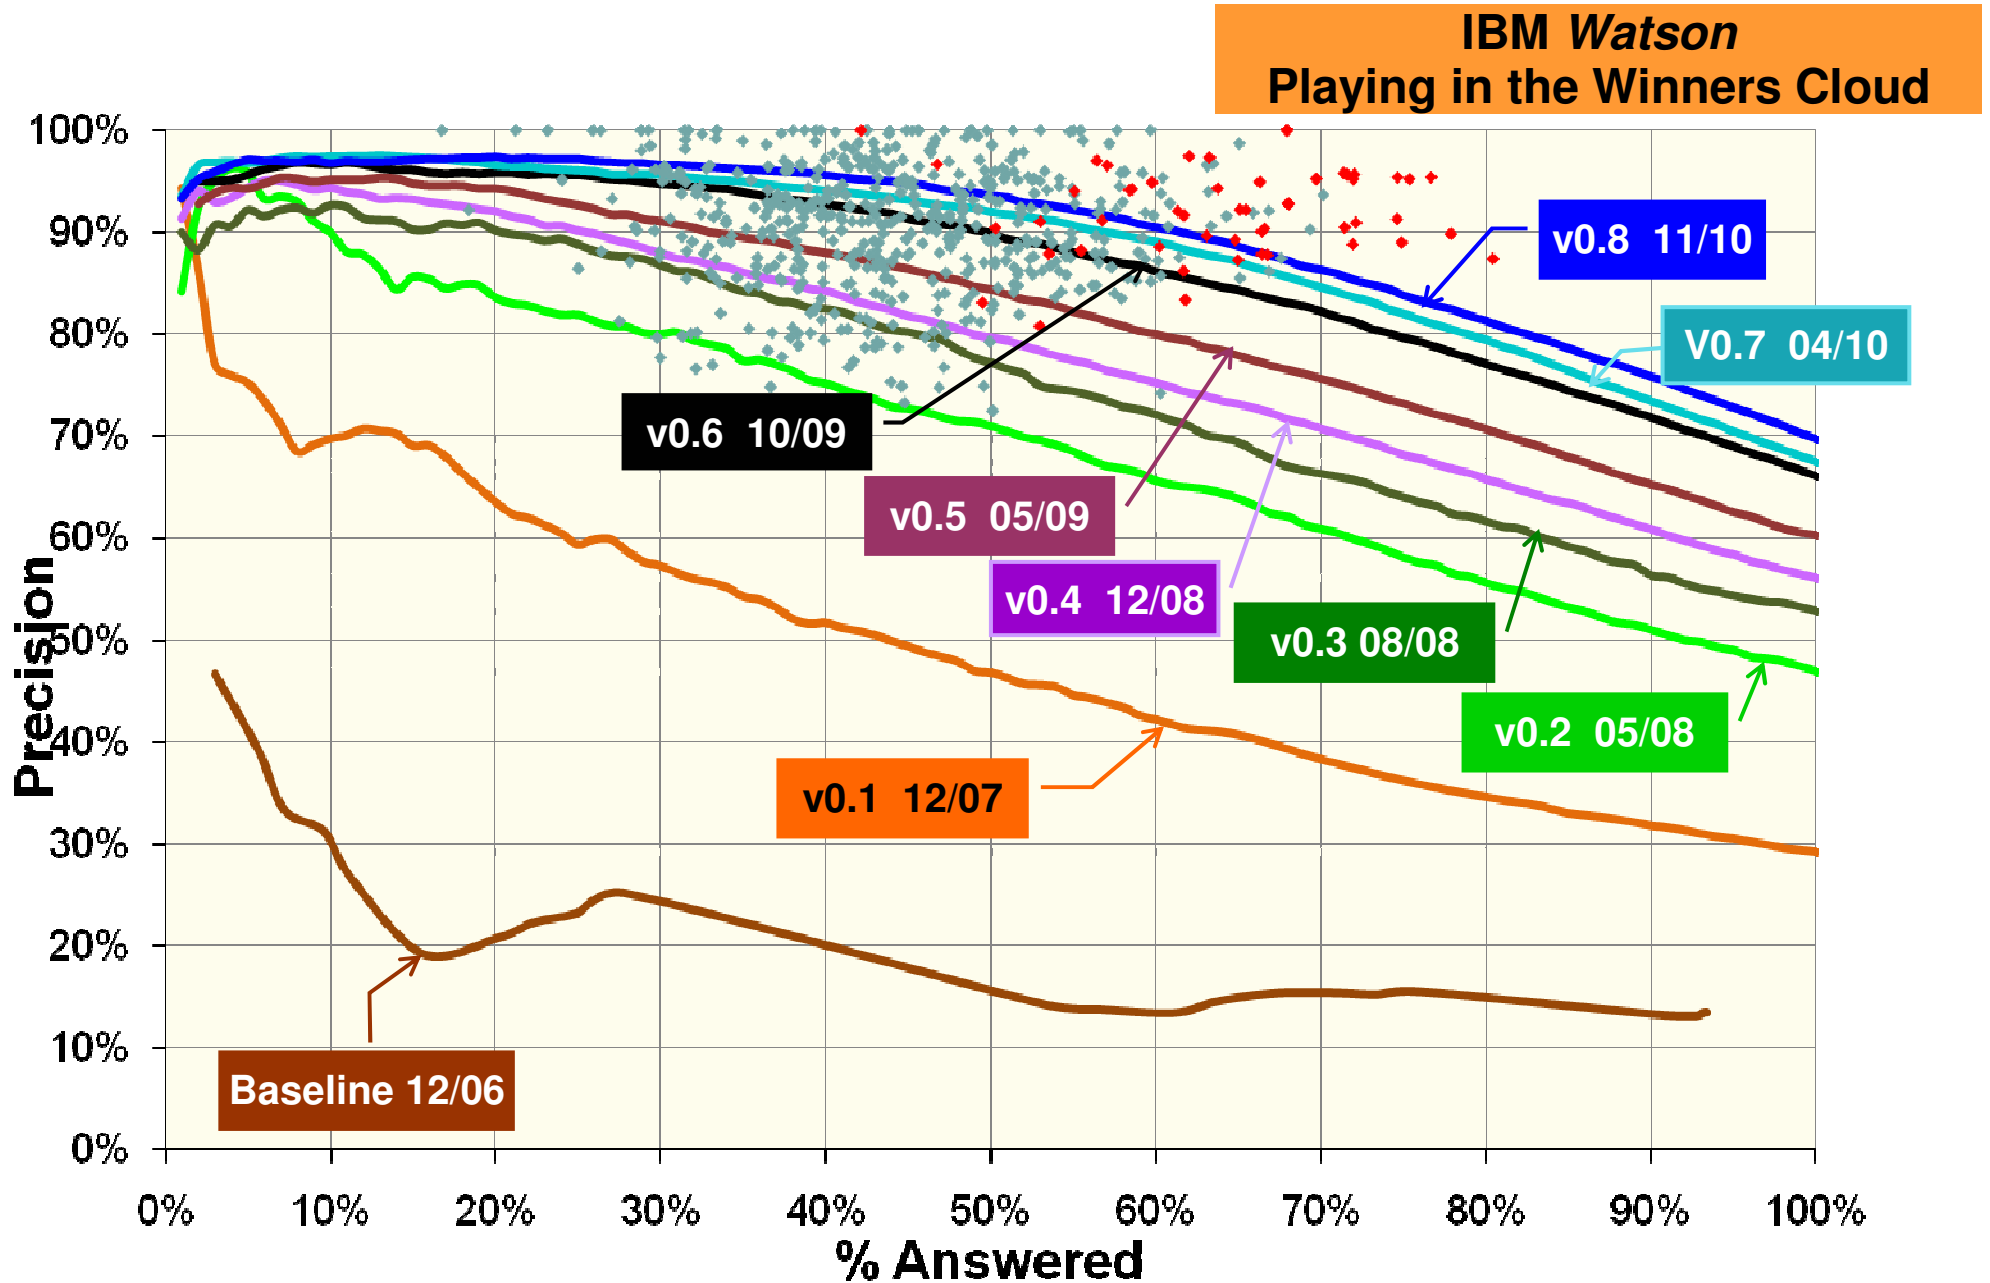


DeepQA: The Technology Behind Watson

*Generates and scores many hypotheses using a combination of 1000's **Natural Language Processing, Information Retrieval, Machine Learning and Reasoning Algorithms.***





DeepQA: Progress in Answering Precision 6/2007-11/2010



Precision / Confidence & Speed

Deep Analytics – Combining many analytics in a novel architecture, we achieved very high levels of *Precision* and *Confidence* over a huge variety of *as-is* content.



Speed – By optimizing Watson’s computation for Jeopardy! on over **2,800 POWER7** processing cores we went from **2 hours per question on a single CPU to an average of just 3 seconds.**



Results – in 55 real-time sparring games against former **Tournament of Champion Players last year**, Watson put on a very competitive performance in all games -- placing 1st in 71% of the them!

Watson has much to learn about Chicago!

Category names on Jeopardy! are tricky

“What US city” wasn’t in the question

Multiple cities named Toronto in the US

Toronto, Canada, has an American League baseball team

Watson found little evidence to connect either city’s airport to WWII

With Toronto at 14% confidence, would not have buzzed in

Chicago at 11% was a very close second on list of possible answers

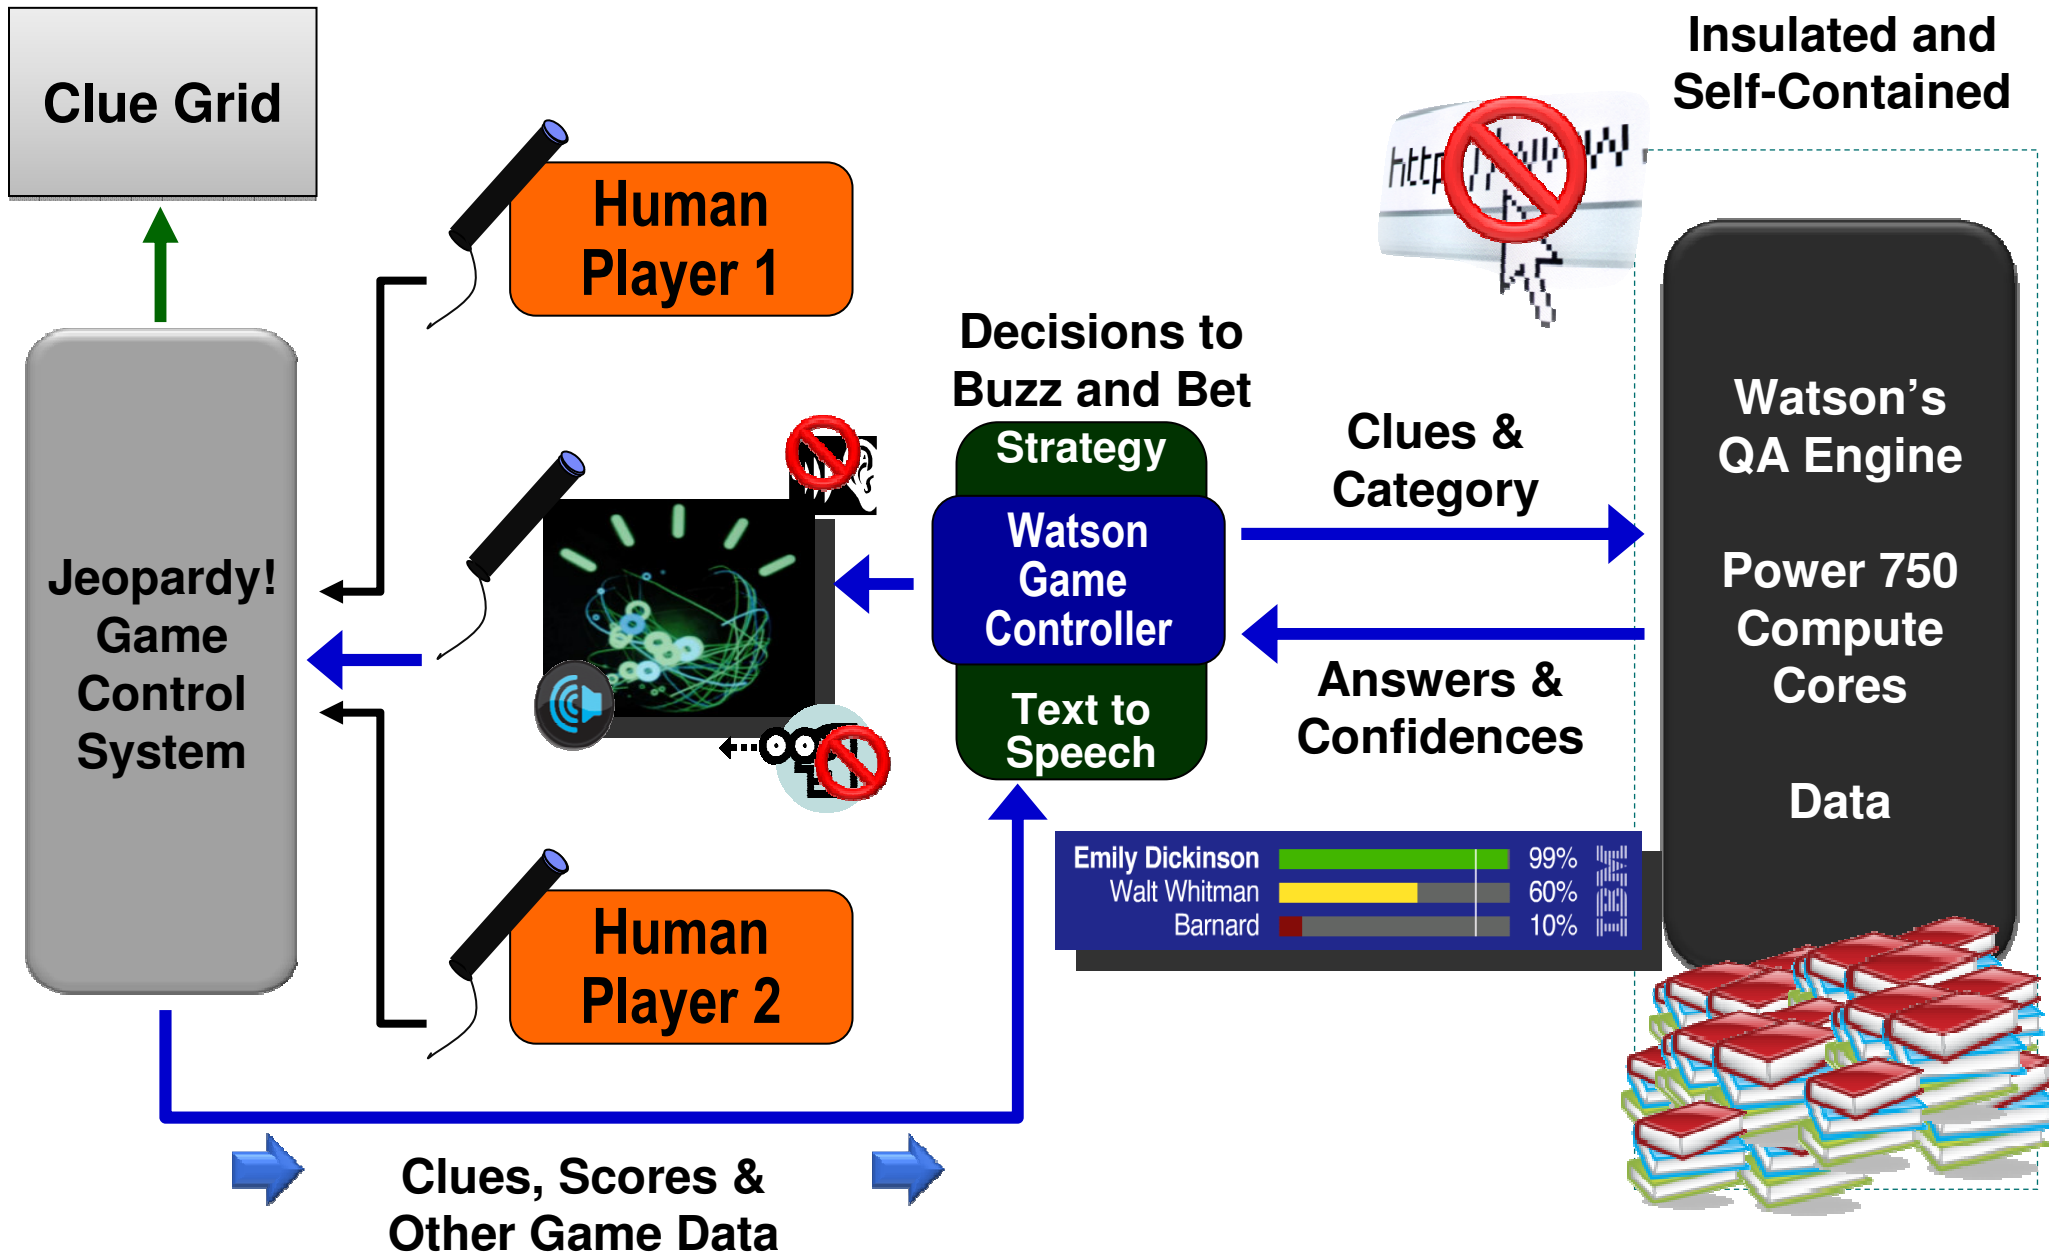
Easy for humans, difficult for Watson

US Cities

Its largest airport is named for a World War II hero, Its second largest for a World War II battle.

Watson: Toronto???
Brad: Chicago
Ken: Chicago

Real-Time Game Configuration



Trivia: Who is....

Who is Ed Toutant????

- Former Power Systems Engineer
- Former Game Show Contestant
- Watson Sparing Partner
- ?????



Watson Reflections



Next Steps for Watson / IBM?

Commercialization of Watson/Power

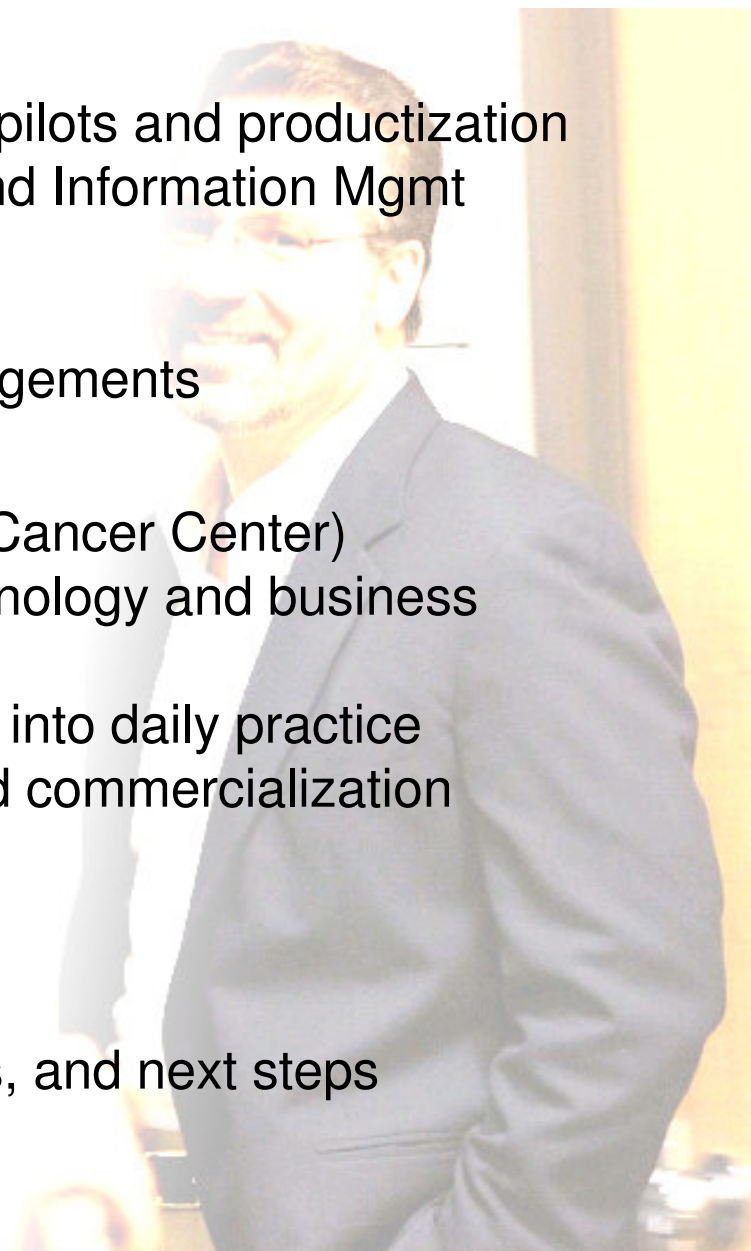
- Power as strategic development platform in early pilots and productization
- Vectoring clients to existing Business Analytics and Information Mgmt offerings
(e.g. pureScale, Cognos, Content Analytics ...)
- “Roadmap to Watson” for longer-term Client engagements

Deep Collaboration

- Healthcare Pilot Applications (e.g. MD Anderson Cancer Center)
- Target broad scale applicability, validation of technology and business model
- Consumability of Jeopardy! accelerating analytics into daily practice
- Technology transfer agents for other research and commercialization activity

Client Engagements

- Collaborative STG/SWG/GBS/Research focus
- Identify collaborative opportunities, client priorities, and next steps

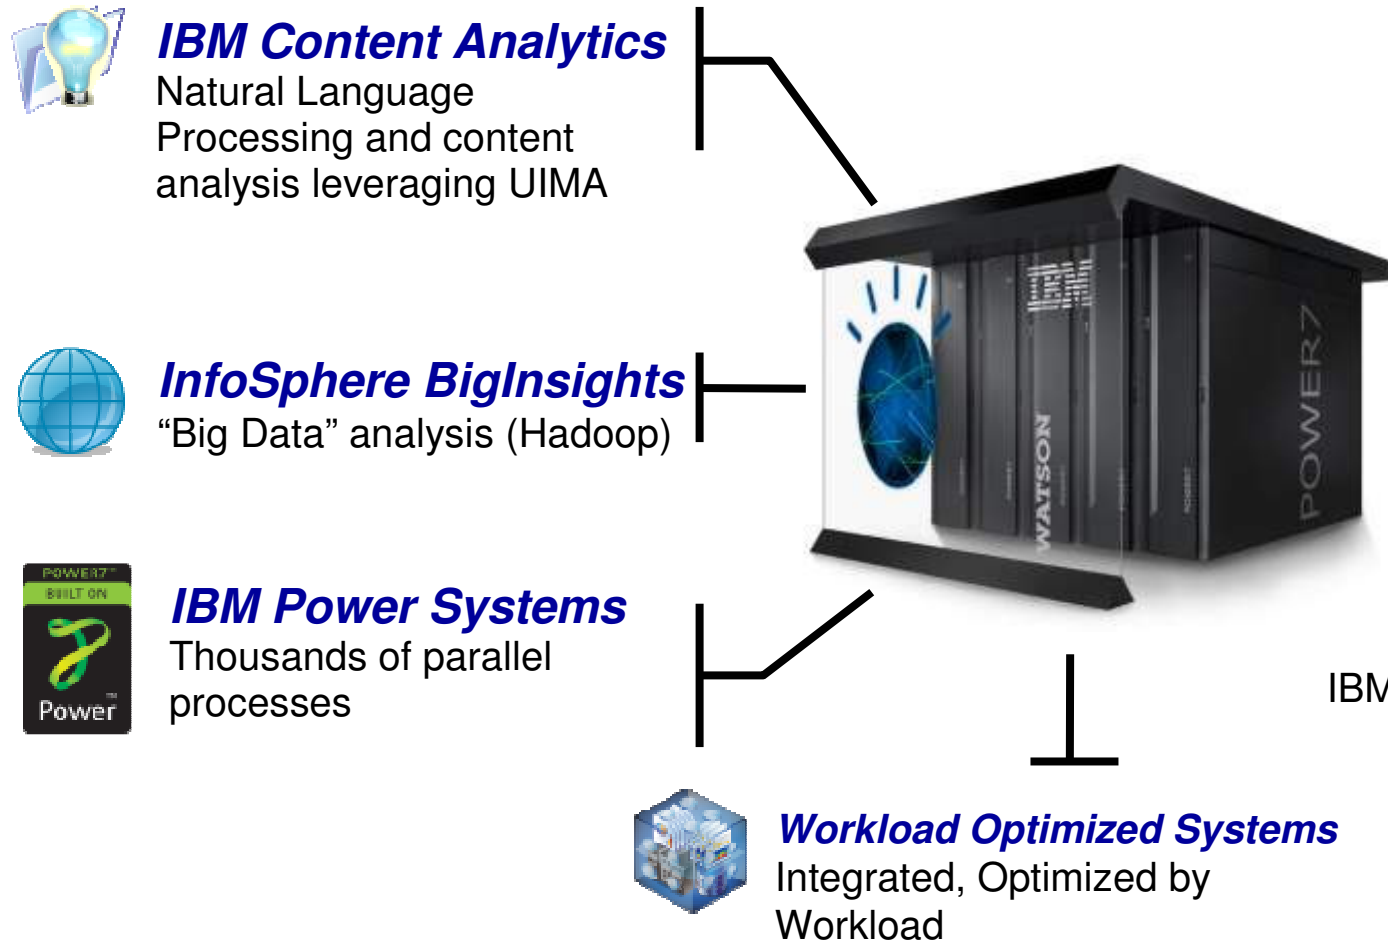


Bridging from Watson to the present

IBM Business Analytics and Optimization solutions

Used by Watson

Related Innovations



InfoSphere Warehouse
DB2, Informix, Netezza
Aggregating and storing data and content



InfoSphere Streams
Massively parallel analysis



Business Analytics
BI, Predictive Analytics and more



ECM Solutions
IBM eDiscovery Analyzer
IBM Classification Module
IBM OmniFind Enterprise Search

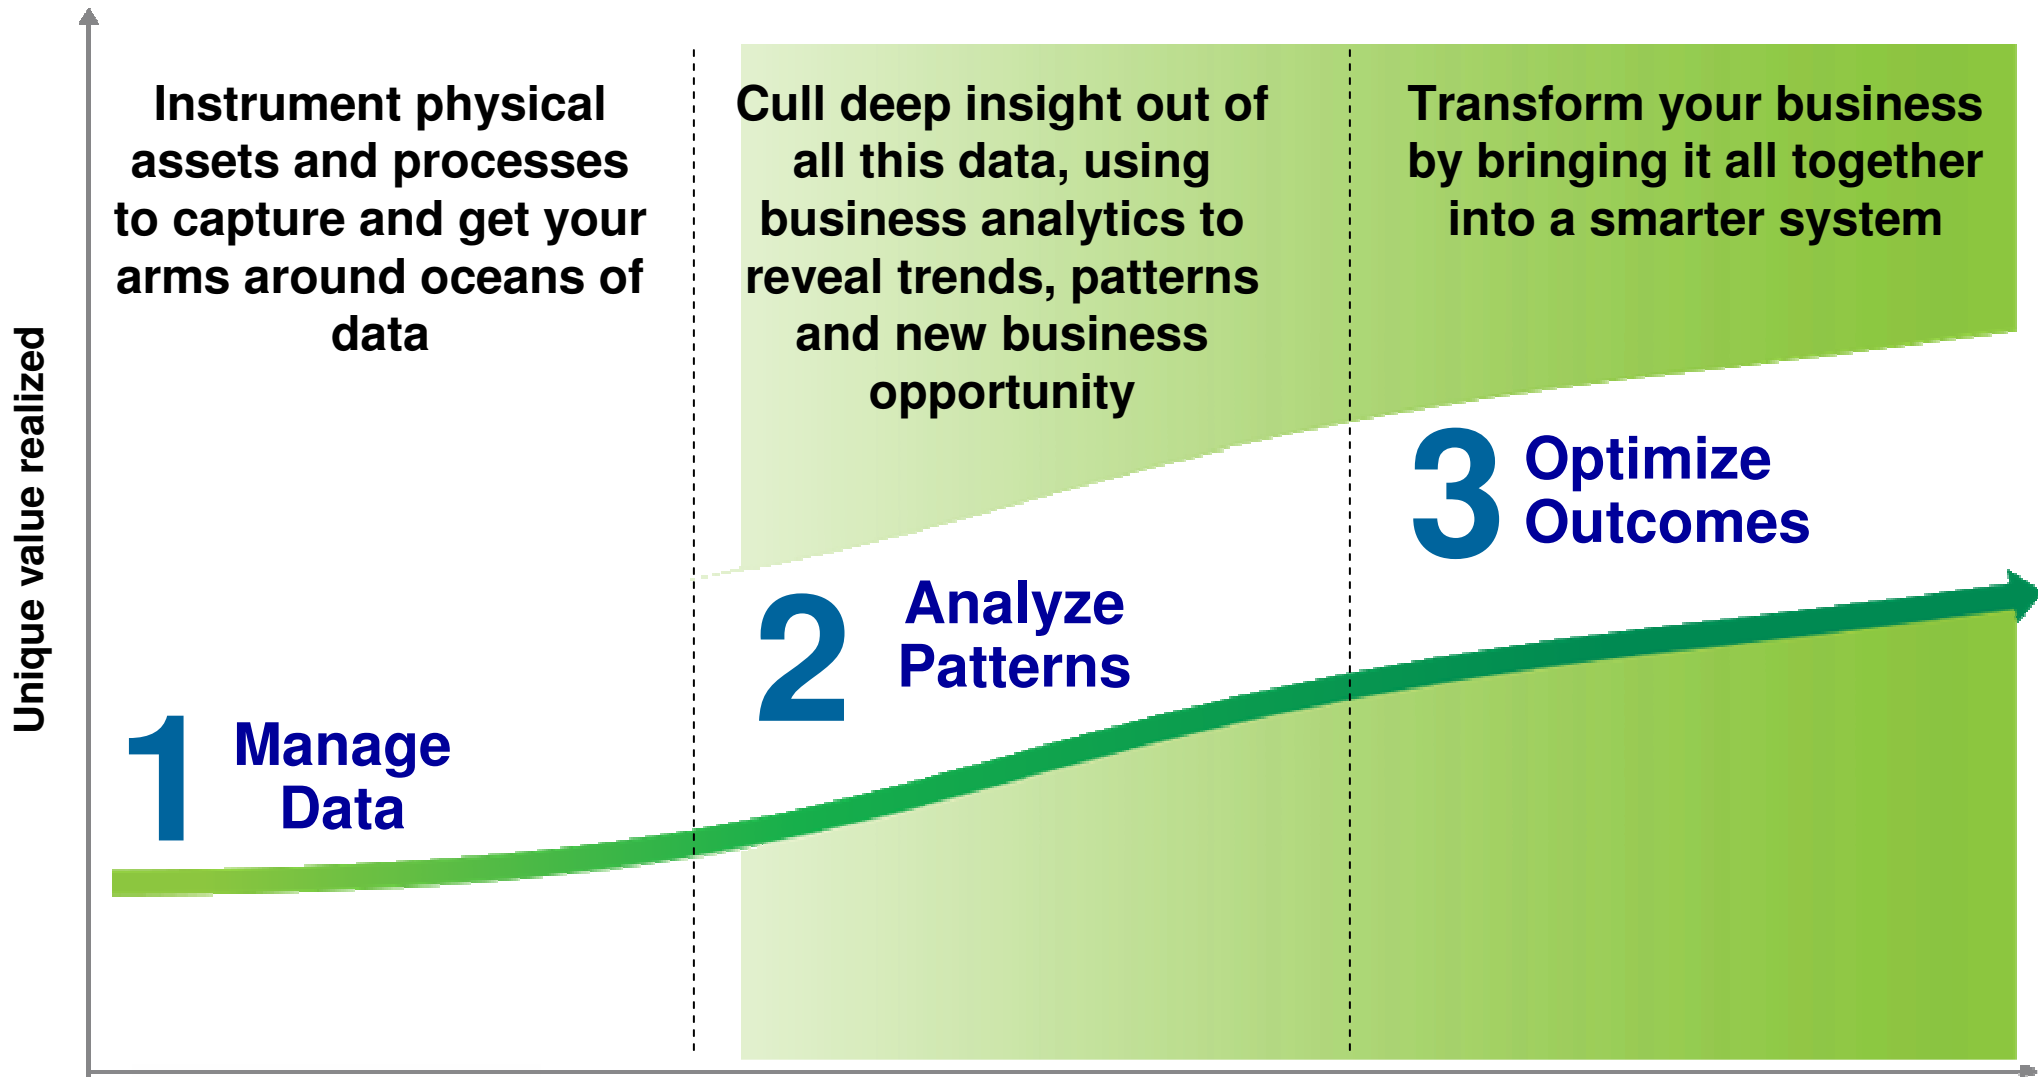


IBM Global Business Services
Research, expertise and analytical assets



Applying Watson's capabilities for business

Helping clients with deployments which tend to follow three phases:



Turning information into insights

Enterprise Content Management

- Advanced Case Management
- Content Analytics
- Document Imaging and Capture
- Information Lifecycle Governance
- Social Content Management



Data Management

- Database Software
- Database Management Tools

Business Analytics

- Analytic Applications
- Business Intelligence
- Predictive Analytics
- Financial Performance Management
- Governance, Risk & Compliance
- Web Analytics

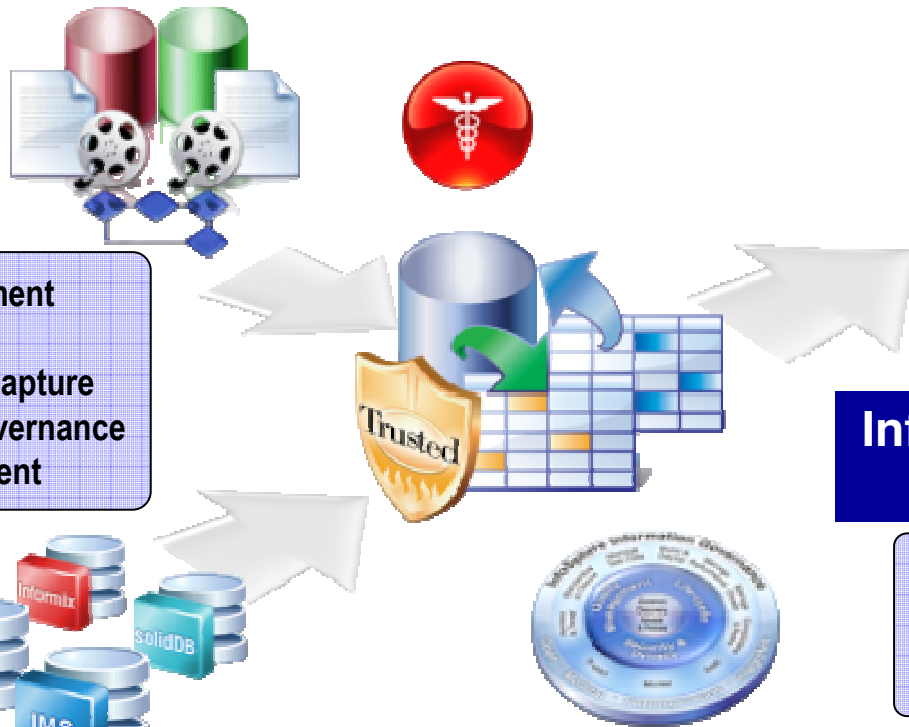


Information Integration & Federation

- Information Integration
- Master Data Management
- Data Warehousing
- Big Data and Streams

Information Governance

- Data Lifecycle Management
- Data Security and Privacy



Potential Business Applications



Healthcare / Life Sciences: Diagnostic Assistance, Evidenced-Based, Collaborative Medicine

Tech Support: Help-desk, Contact Centers



Enterprise Knowledge Management and Business Intelligence

Government: Improved Information Sharing and Security




Evidence Profiles from disparate data sources

Each dimension contributes to supporting or refuting hypotheses based on

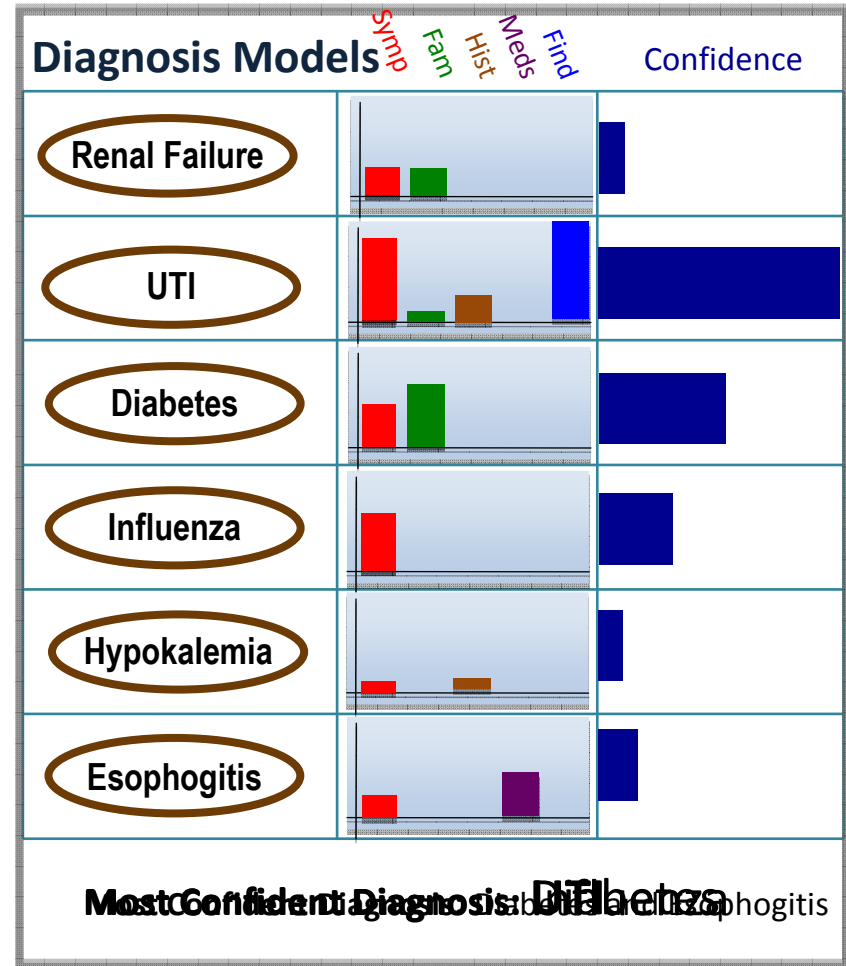
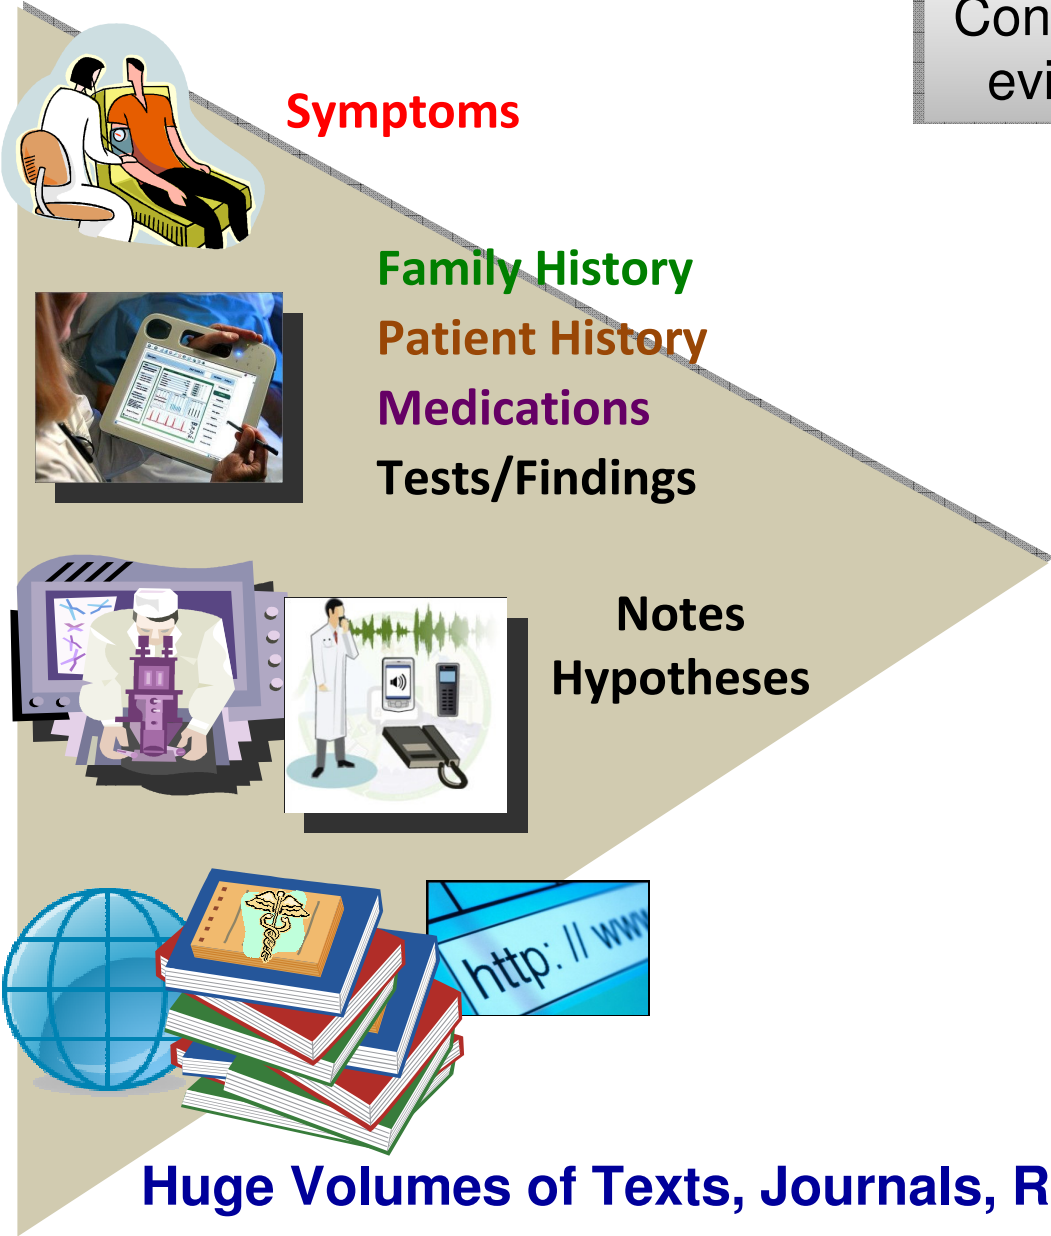
- **Strength of evidence**
- **Importance of dimension for diagnosis** (learned from training data)

Evidence dimensions are combined to produce an overall confidence



Continuous Evidence-Based Diagnostic Analysis



Considers and synthesizes a broad range of evidence improving quality, reducing cost



Huge Volumes of Texts, Journals, References, DBs etc.

Active Watson / DeepQA Engagements

Client inquiries arriving with a broad range of Watson use cases, from a broad range of industries

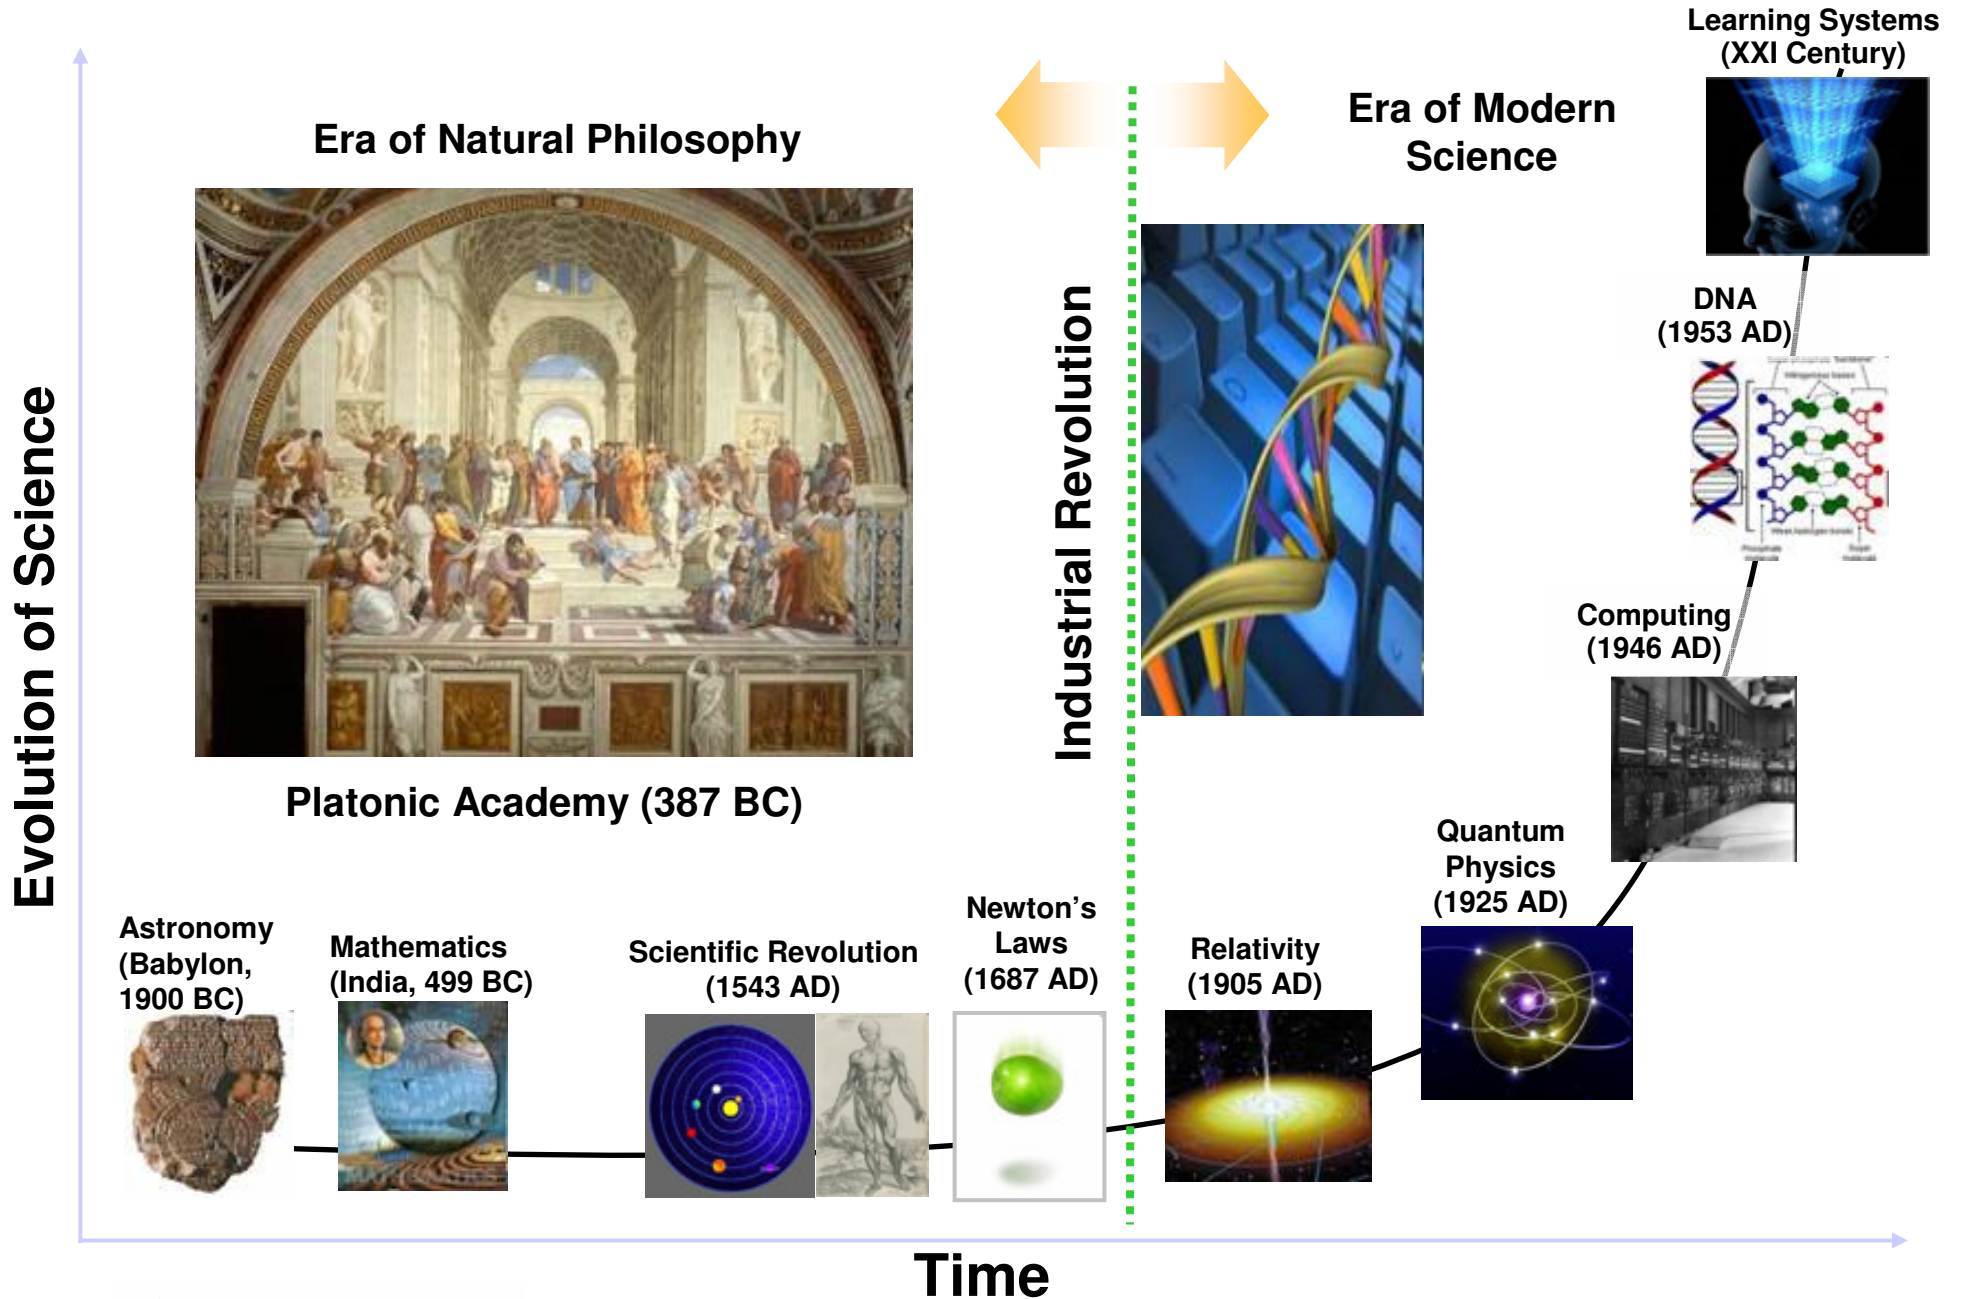


Past History....

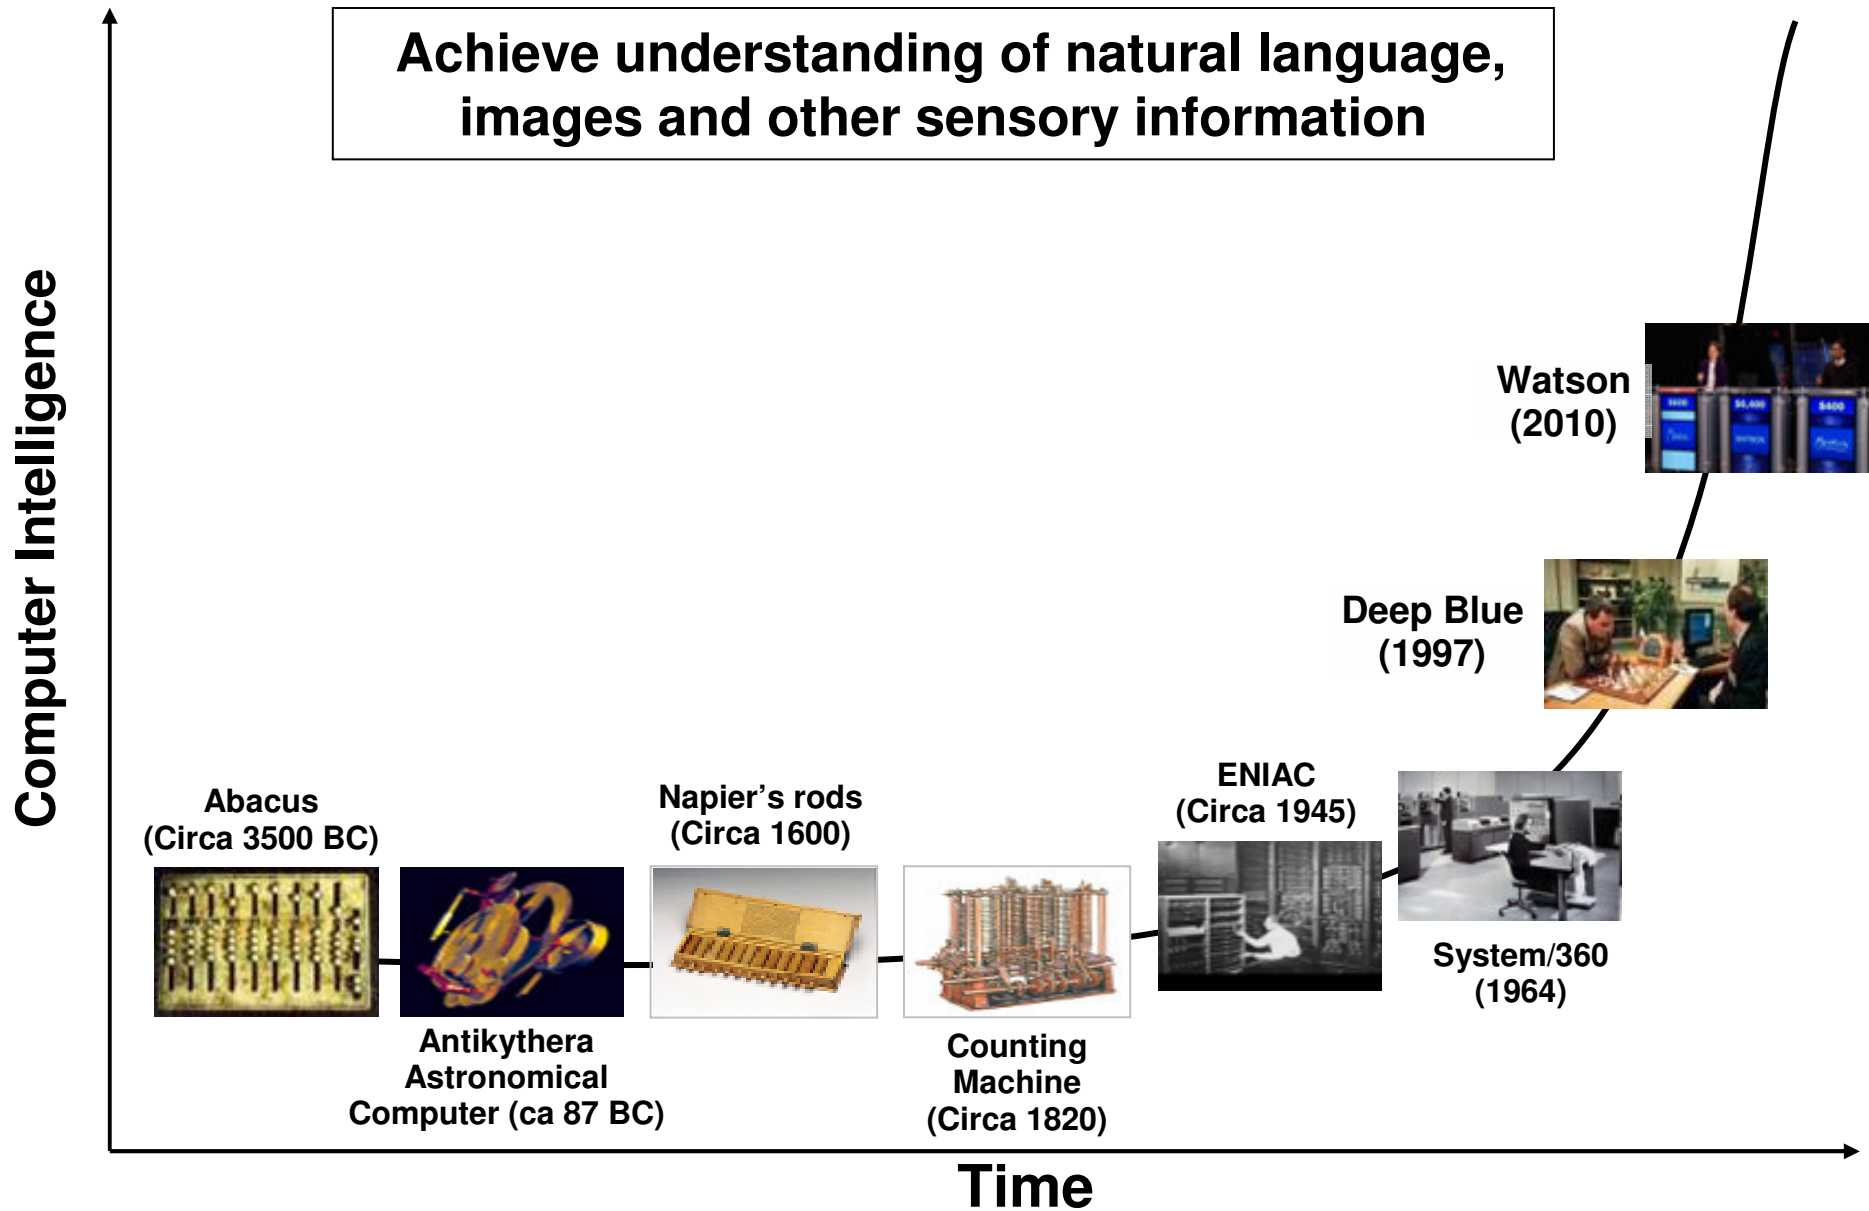


Technology that could compete against man

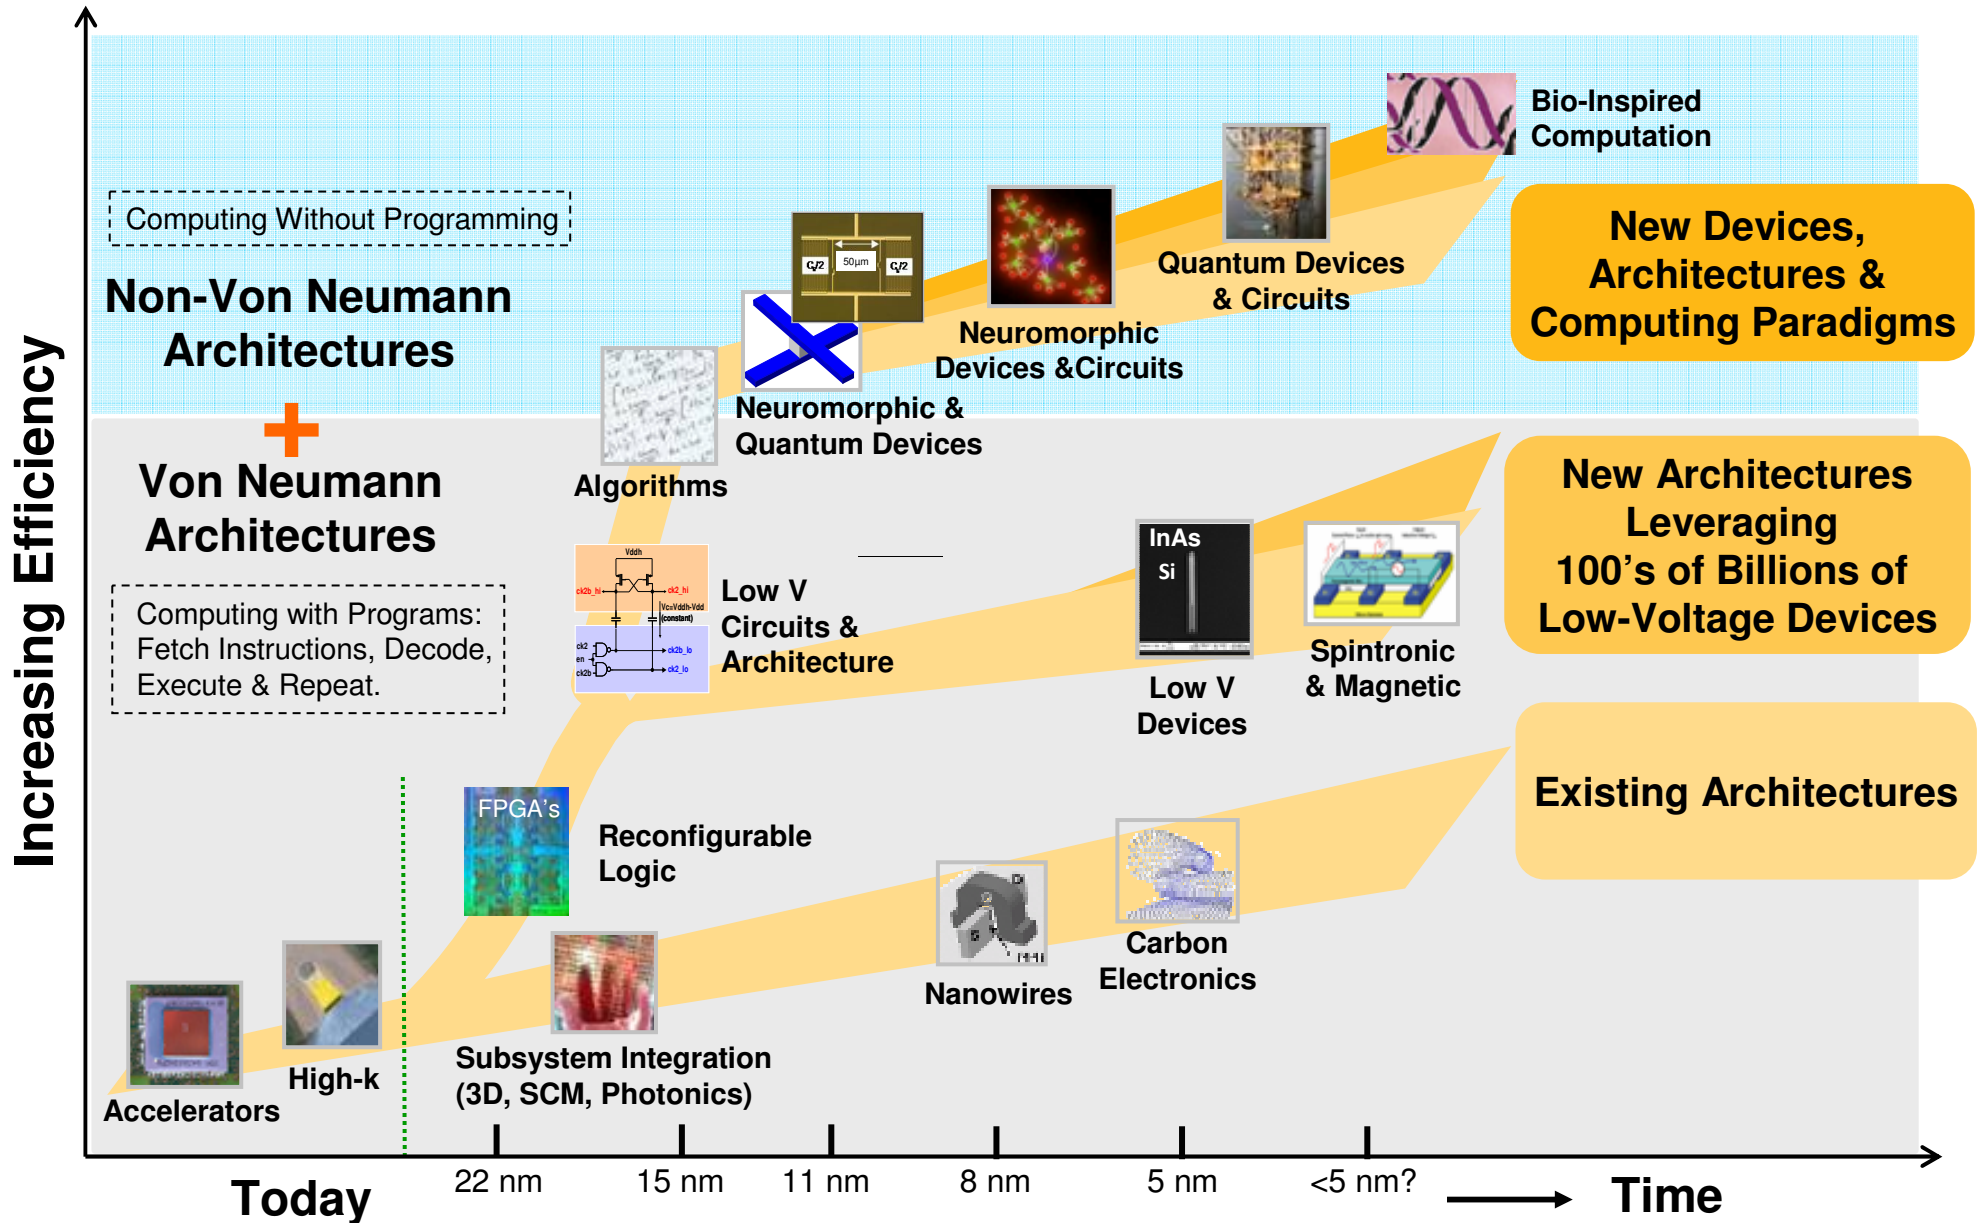
The Evolution of Science




The Evolution of Technology



Device and Technology Roadmap






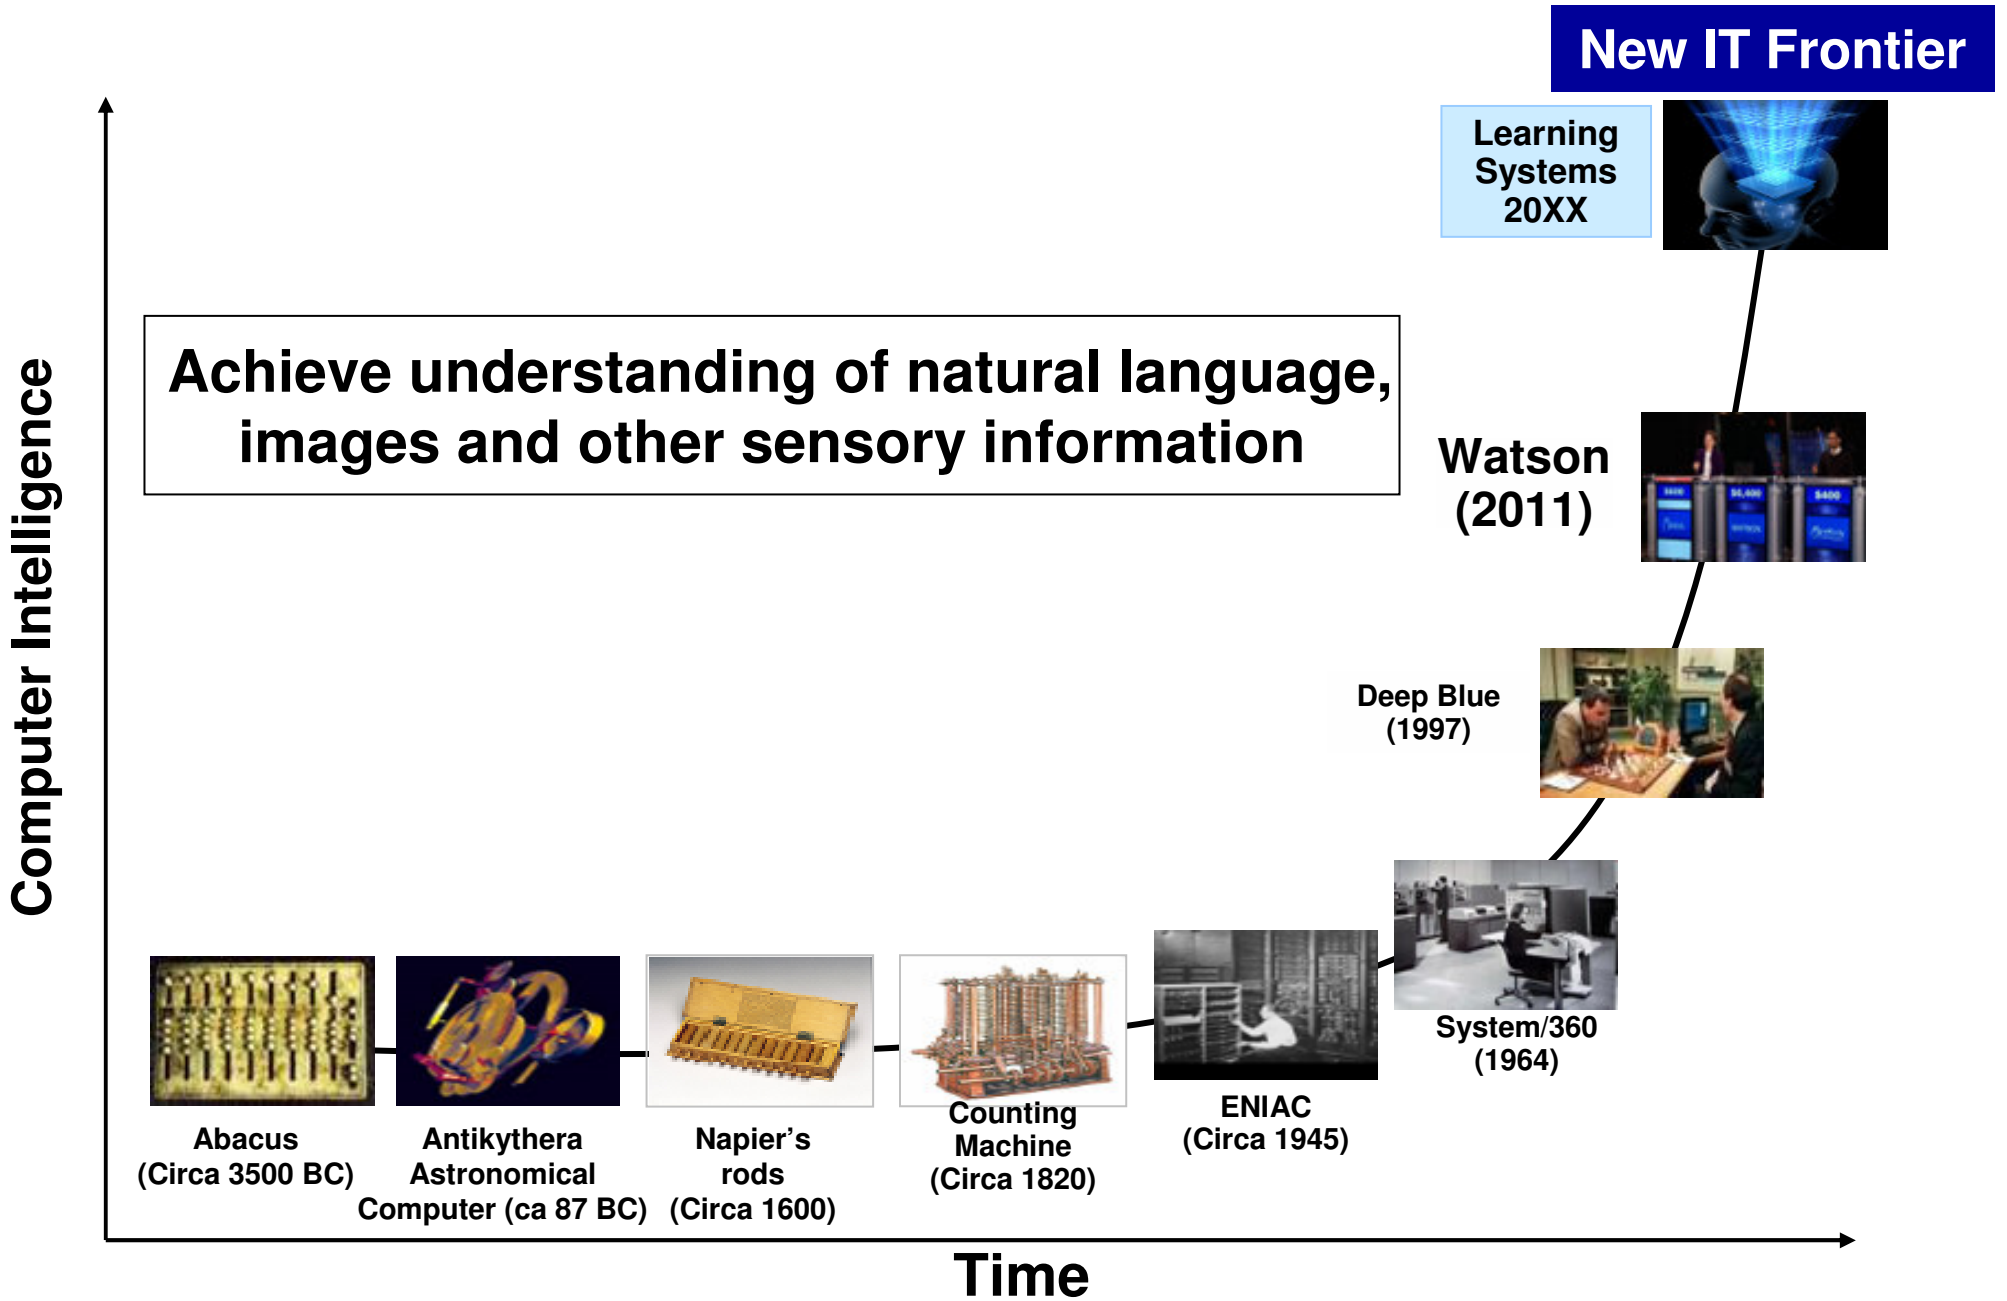
Technology that can think/reason like man

Data Problem of the future.....

Future Data Components

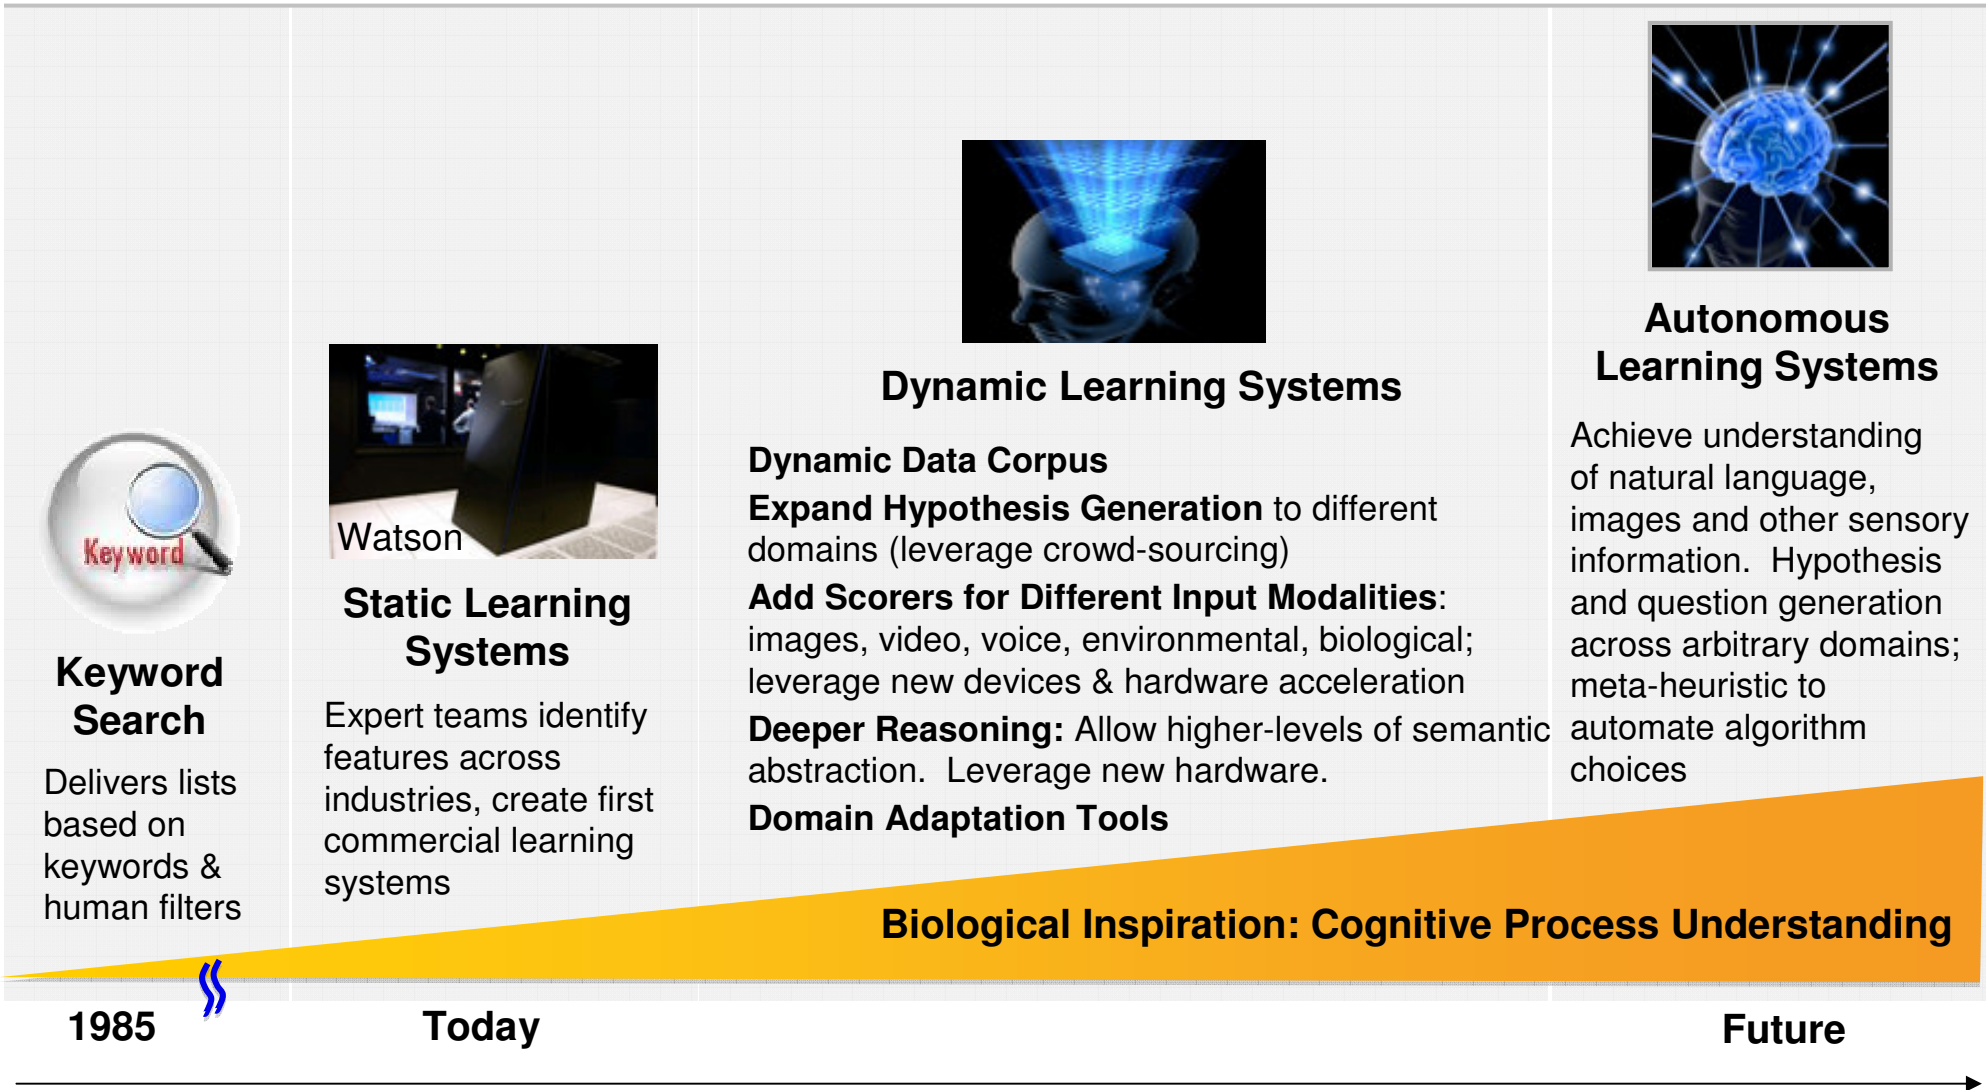


The Evolution of “Thinking” Machines



Learning Systems Roadmap to Meet the Challenge

The Research Division will ensure IBM is the world leader in Learning Systems. We will define & benchmark progress through a series of Grand Challenges.



The Future

Watson

It is just the beginning...