


IBM Watson

and Cognitive Computing



Christopher Codella, Ph.D.
IBM Distinguished Engineer, Public Sector CTO, Watson Group



Cognitive Systems...

... learn and interact naturally with people to extend what humans can do on their own. They help us solve problems by penetrating the complexity of Big Data.

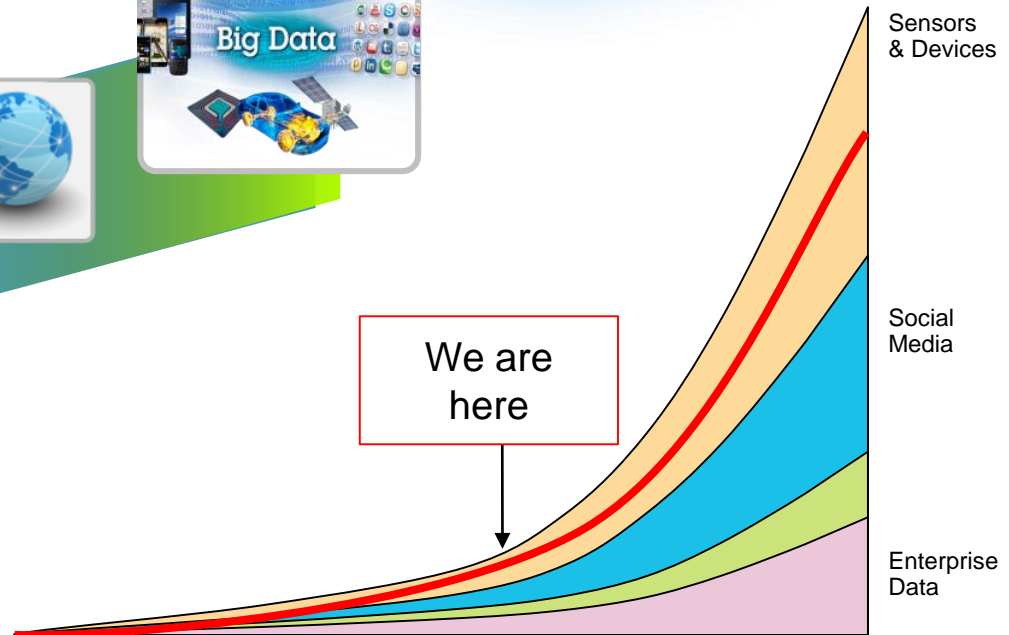
Tabulating Systems Era



Programmable Systems Era




Cognitive Systems Era



A Cognitive System is Not...

A black box



Rather, it possesses a model of its domain, a model of its users, and a degree of self-knowledge that contributes to conversational discovery and decision-making

Statically programmed



Instead, it learns (although domain adaptation is a very hard problem)

Conscious or sentient



However, the more interesting cognitive systems grow via unassisted learning, can possess a degree of self-directed action, and may also engage with the world with an element of autonomy

Automatic Open-Domain Question Answering

A Long-Standing Challenge in Artificial Intelligence to emulate human expertise

■ Given

- Rich **Natural Language Questions**
- Over a **Broad Domain of Knowledge**

■ Deliver

- **Precise Answers:** Determine what is being asked & give precise response
- **Accurate Confidences:** Determine likelihood answer is correct
- **Consumable Justifications:** Explain why the answer is right
- **Fast Response Time:** Precision & Confidence in <3 seconds

A Grand Challenge Opportunity

- Capture the imagination
 - The Next *Deep Blue*
- Engage the scientific community
 - Envision new ways for computers to impact society & science
 - Drive important and measurable scientific advances
- Be Relevant to IBM Customers
 - Enable better, faster decision making over unstructured and structured content
 - Business Intelligence, Knowledge Discovery and Management, Government, Compliance, Publishing, Legal, Healthcare, Business Integrity, Customer Relationship Management, Web Self-Service, Product Support, etc.



Real Language is Real Hard

■ Chess

- A finite, mathematically well-defined search space
- Limited number of moves and states
- Grounded in **explicit, unambiguous** mathematical rules



■ Human Language

- Ambiguous, contextual and implicit
- Grounded only in **human cognition**
- Seemingly infinite number of ways to express the same meaning

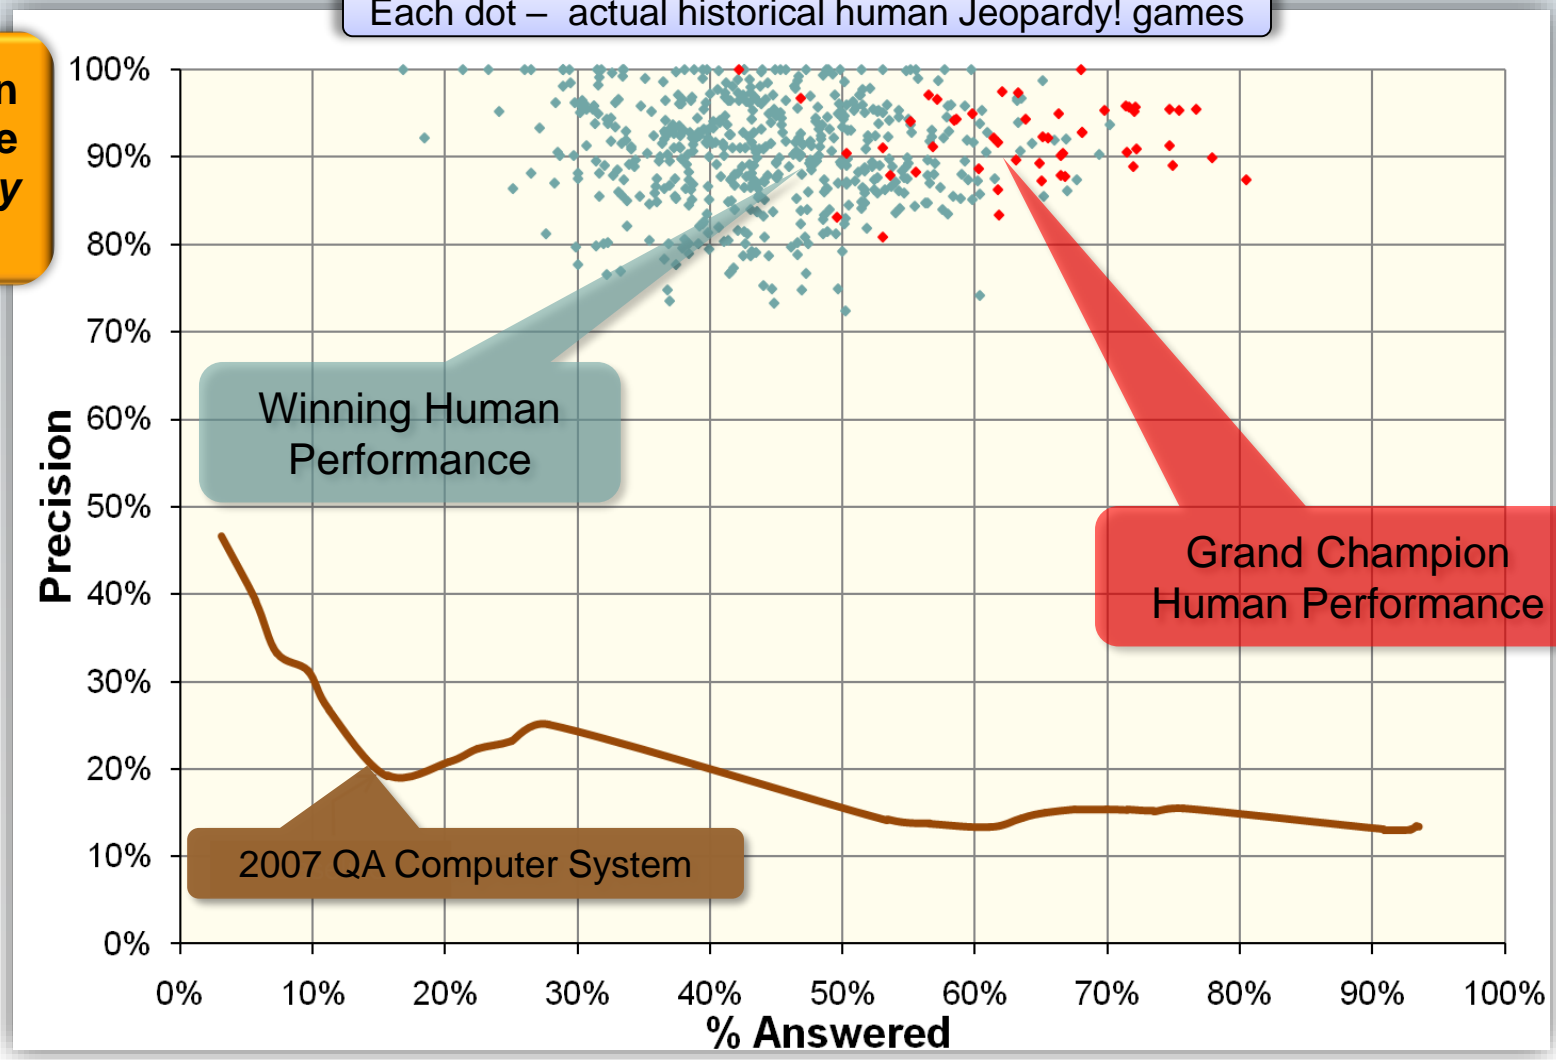


What It Takes to compete against Top Human Jeopardy! Players

Our Analysis Reveals the Winner's Cloud

Top human players are remarkably good.

Each dot – actual historical human Jeopardy! games



What Computers Find Easy (and Hard)

$$(\ln(12,546,798 * \pi)) ^ 2 / 34,567.46 = \mathbf{0.00885}$$

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

Owner	Serial Number
David Jones	45322190-AK

Invoice #	Vendor	Payment
INV10895	MyBuy	\$104.56

Serial Number	Type	Invoice #
45322190-AK	LapTop	INV10895

David Jones
 ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓ ↓
 David Jones

=

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

≠

Traditional computing methods make it hard for computers to understand us

Structured

?
J. Welch
ran this

Person	Organization
L. Gerstner	IBM
J. Welch	GE
W. Gates	Microsoft

Unstructured

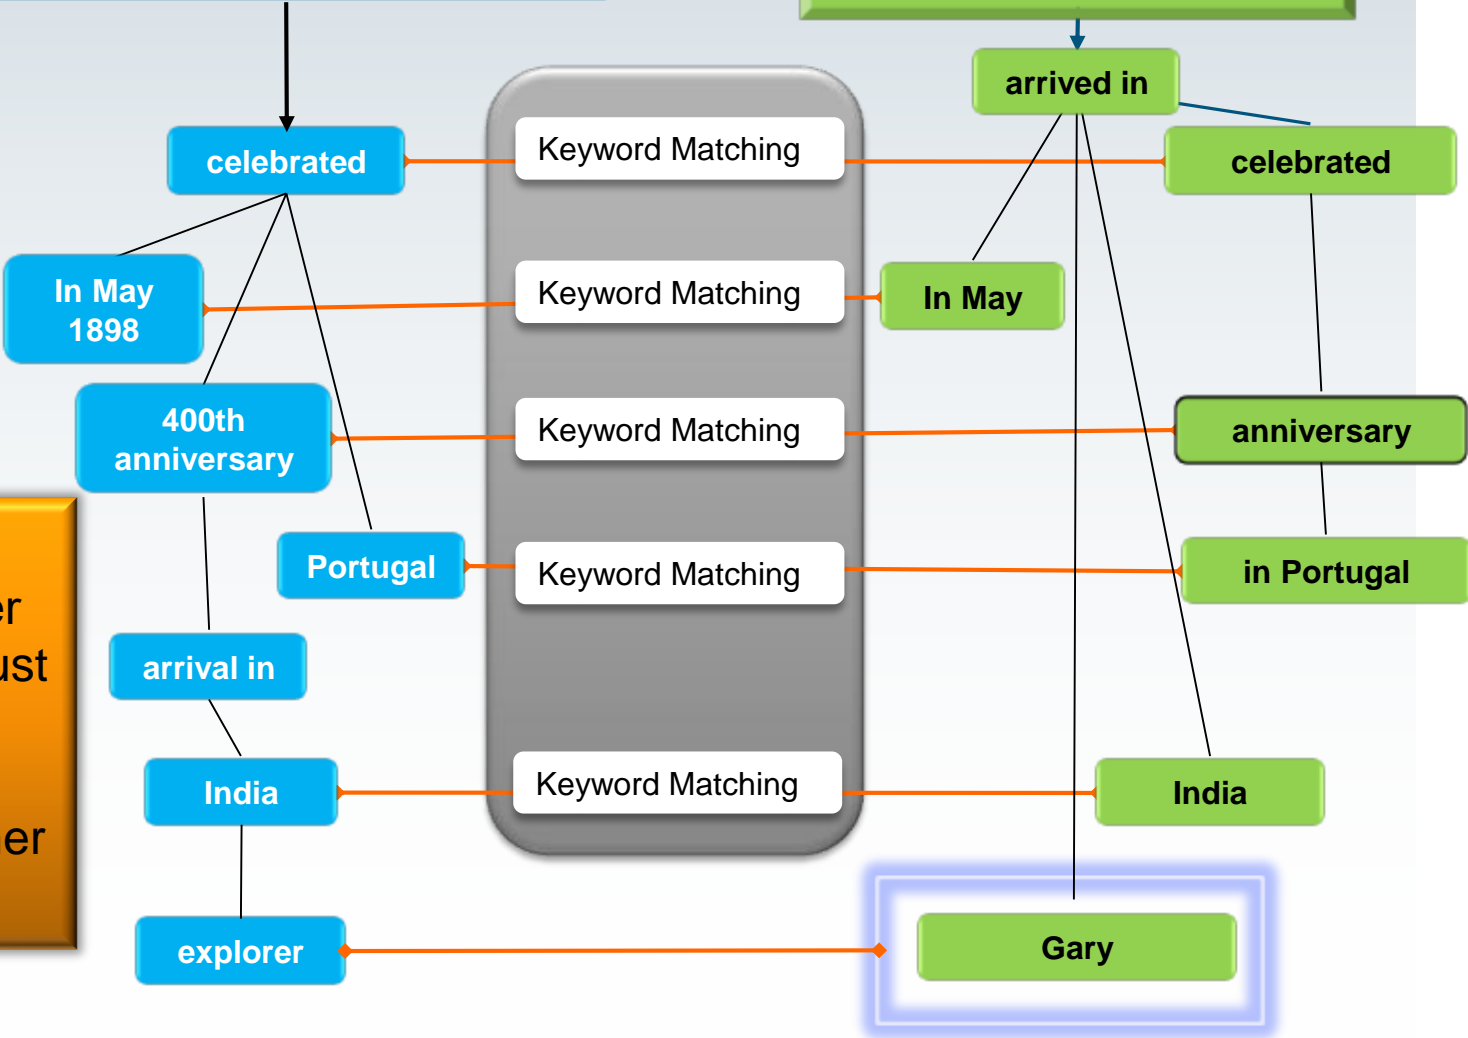
“If leadership is an art then surely Jack Welch has proved himself a master painter during his tenure at GE.”

- Noses that run and feet that smell?
- How can a house burn up as it burns down?
- Does CPD represent a complex comorbidity of lung cancer?
- What mix of zero-coupon, non-callable, A+ munis fit my risk tolerance?
- I will have coke and ice.

Different Types Of Evidence: Keyword Evidence

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.



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

Different 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

celebrated

Portugal

May 1898

400th anniversary

arrival in

India

explorer

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

Temporal Reasoning

Statistical Paraphrasing

Geospatial Reasoning

landed in

27th May 1498

Kappad Beach

Vasco da Gama

Date Math

Para-phrases

Geo-KB

Stronger evidence can be much harder to find and score.

The evidence is still not 100% certain.

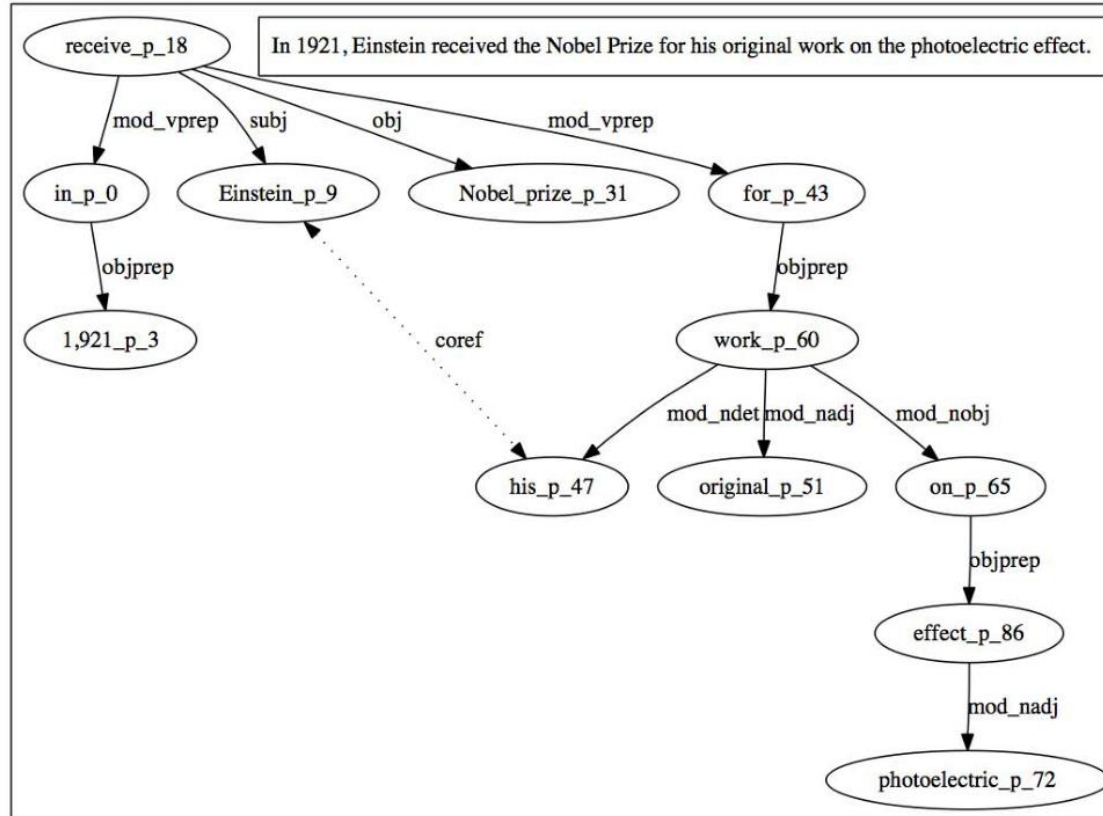
Watson starts by breaking a sentence into pieces

▪ Corpus Processing:

- Syntactic and Semantic relations drive everything else.

▪ Example:

- In 1921, Einstein received the Nobel Prize for his original work on the photoelectric effect



Automatic Learning For “Reading”


Sentence
Parsing

Generalization &
Statistical
Aggregation



Volumes of Text

Syntactic Frames

Semantic Frames



subject verb object

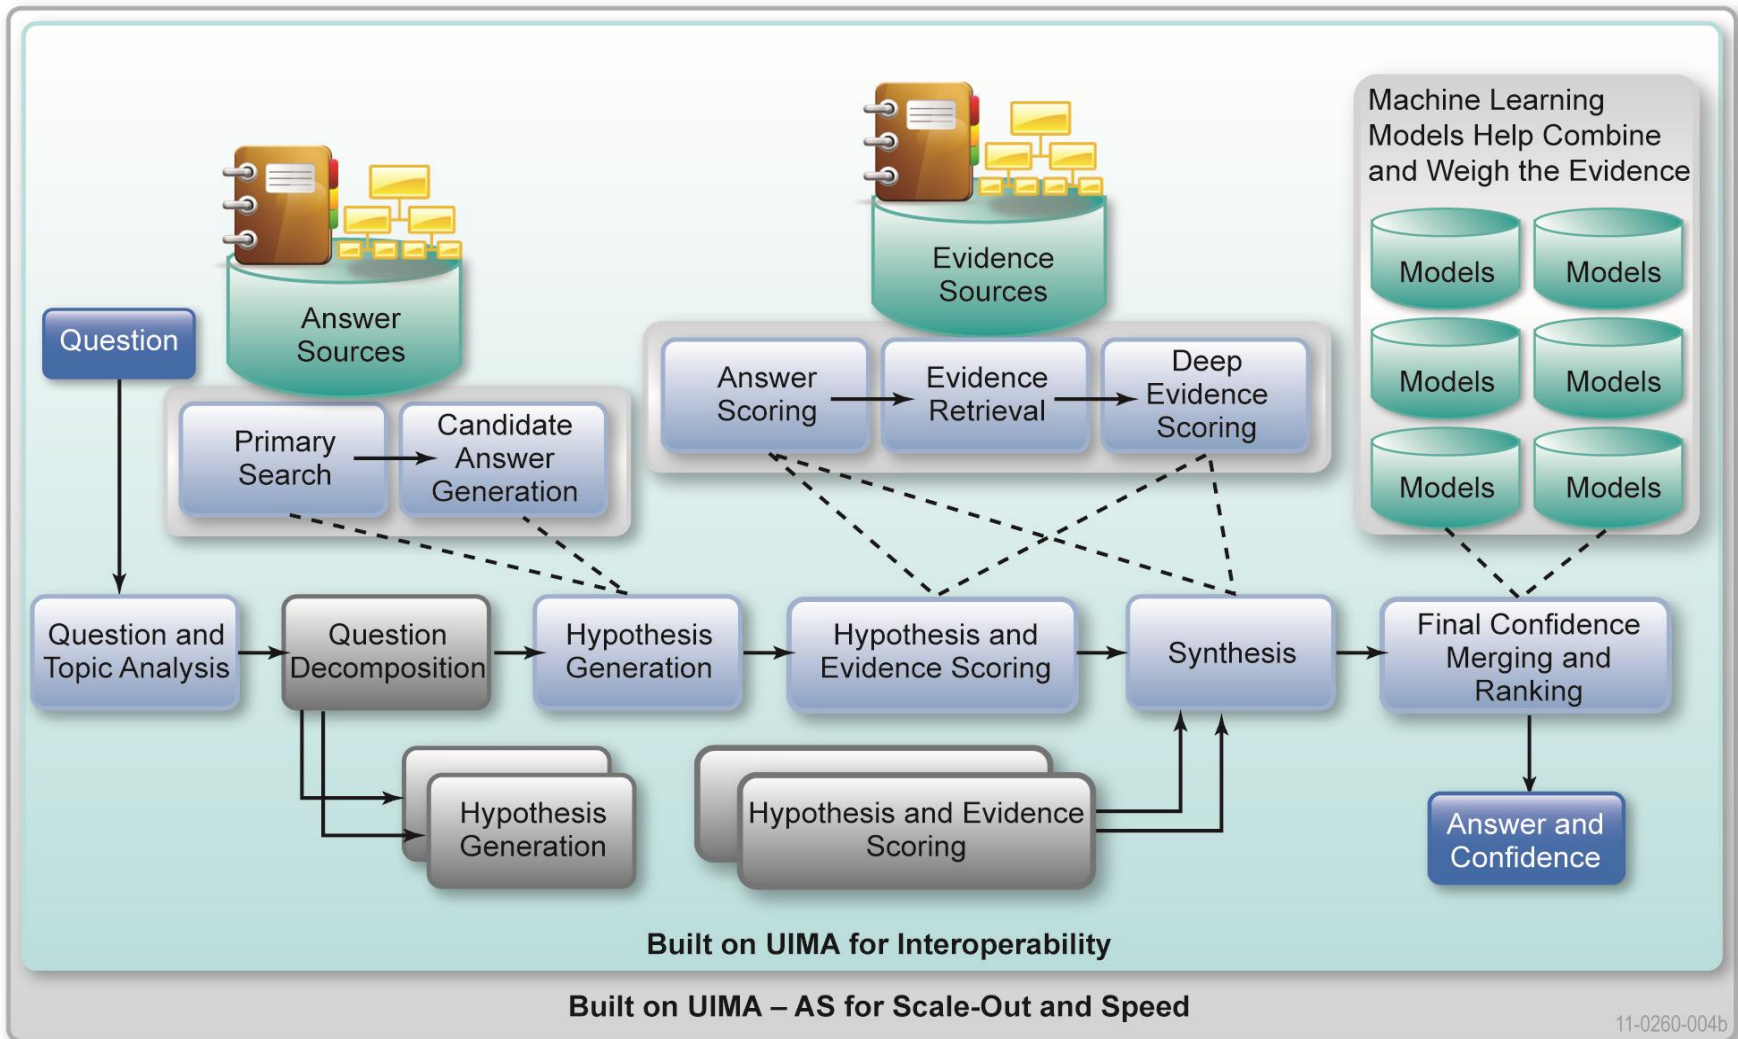


Inventors patent inventions (.8)
 Officials Submit Resignations (.7)
 People earn degrees at schools (0.9)
 Fluid is a liquid (.6)
 Liquid is a fluid (.5)
 Vessels Sink (0.7)
 People sink 8-balls (0.5) (in pool/0.8)




Inside Watson

Massively Parallel Probabilistic Evidence-Based Architecture



Jeopardy! - Incremental Progress in Precision and Confidence



Watson

Final Score: \$ 24,000 \$ 77,147 \$ 21,600

Contestant	Prize Amount	Final Score
Ken	\$300,000	\$ 24,000
Watson	\$1,000,000	\$ 77,147
Brad	\$200,000	\$ 21,600

Informed Decision Making: Search vs. Expert Q&A

Decision Maker

Has Question

Distills to 2-3 Keywords

Reads Documents, Finds Answers

Finds & Analyzes Evidence

Search Engine

Finds Documents containing Keywords

Delivers Documents based on Popularity

Expert

Understands Question

Produces Possible Answers & Evidence

Analyzes Evidence, Computes Confidence

Delivers Response, Evidence & Confidence

Decision Maker

Asks NL Question

Considers Answer & Evidence

Use Case Categories

Exploration

Collect the information that you need to explore your problem area better

Engagement

Dialog with end users to answer the questions needed around products and services

Discovery

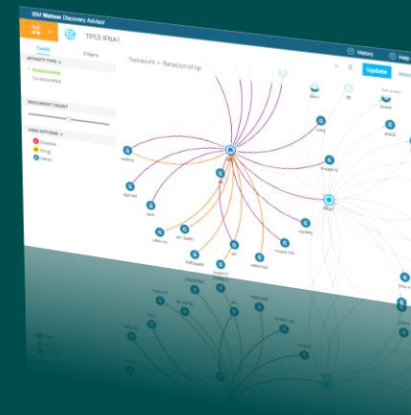
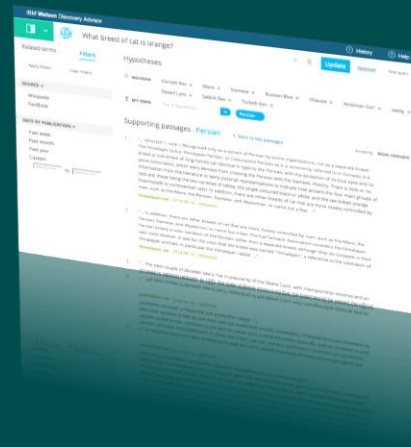
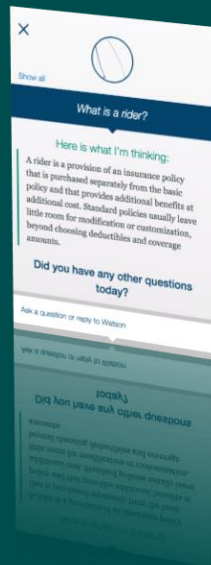
Help find the questions you're not thinking to ask and connect the dots that you're missing that will lead to new inspiration

Decision

Assess the choices that enable you to make better decisions

Policy


Test conformance to a set of written policy conditions




Dimensions for classifying use-cases

- Corpus size and complexity
- Special user interface
- Integration with external analytics
 - Use results from Watson (e.g. visualization)
 - Provide knowledge to Watson (e.g. from streams, RDBs)
 - Trigger a question to be asked
- Time sensitivity, volatility, dependence
- Question type
 - Simple, free-form question or assertion
 - Question with context (e.g. a case file)
 - Standing (persistent or watched) question
 - Template question
- Number of users (e.g. 10s, 100s, thousands, millions)
- Dialog capability necessary

Watson R&D Advisor – The Discovery Pattern



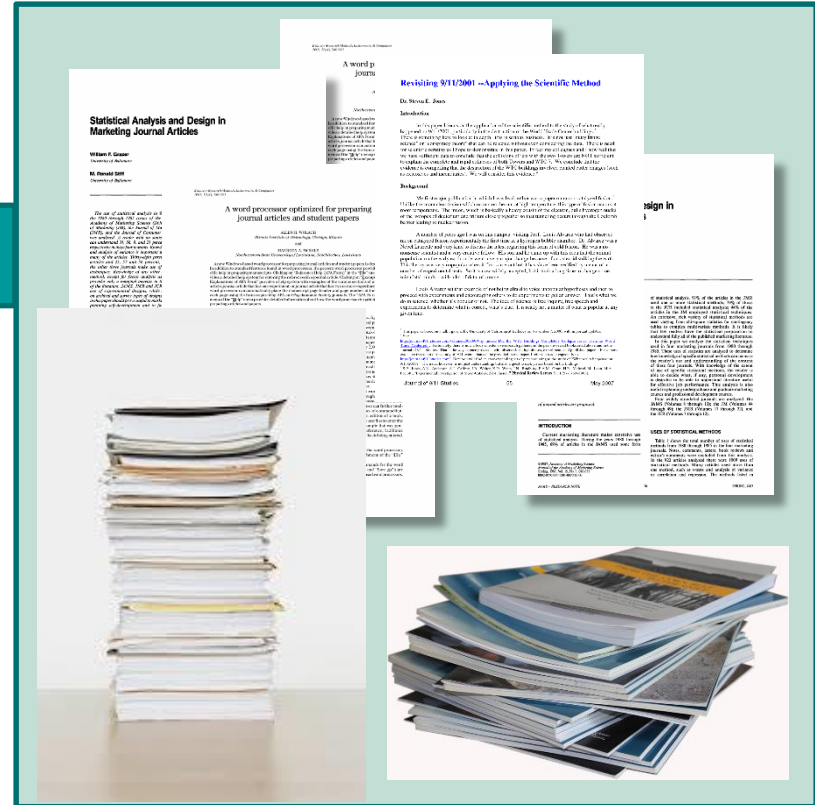
← Ingest



↕ Consult



Scientists, Engineers, Planners,
Project Management, Attorneys,
Scholars, Economists,
Legislators, Analysts, ...



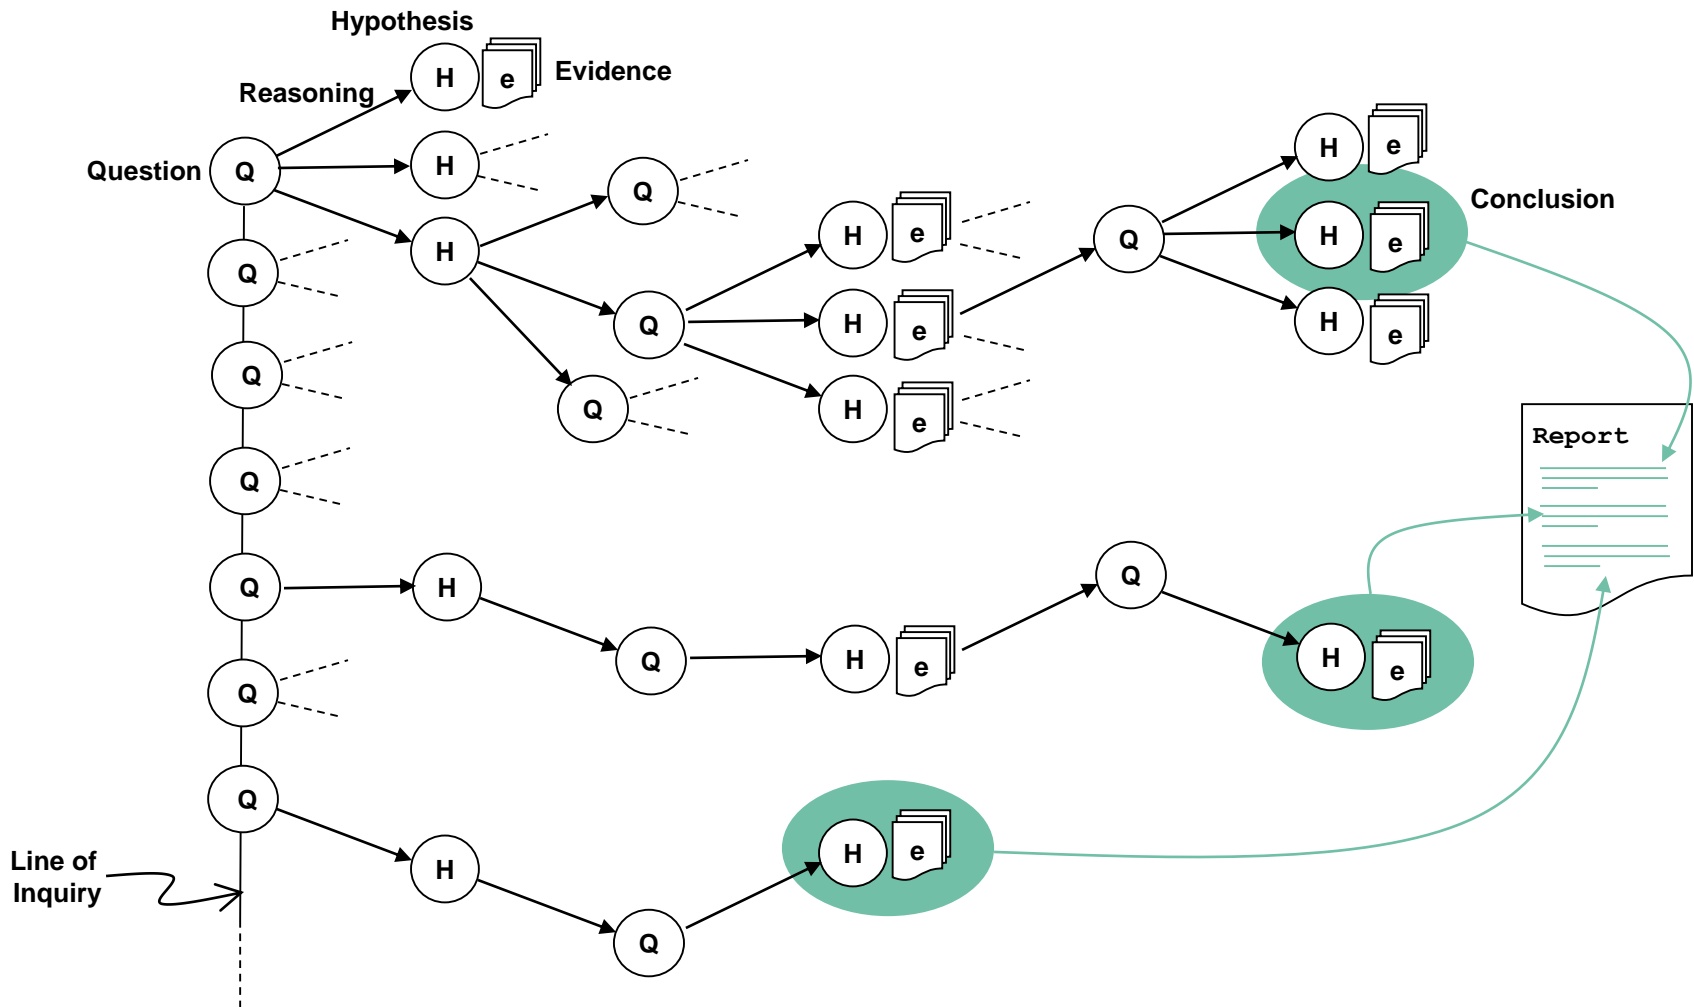
Watson Discovery

A trusted assistant

- Problem characteristics:
 - 10s of millions of natural language documents must be considered
 - Unbounded variety of questions and domains
 - No single answer is sought – competing hypotheses must be considered
 - Exploration of evidence takes precedence over answers

- Watson:
 - Analyzes and extracts knowledge from an unlimited number of documents
 - Understands natural language questions and assertions
 - Responds with competing hypotheses and evidence behind each
 - Provides relevant passages of evidence
 - Provides links to original documents
 - Enables analysts to collaborate on a line of inquiry or investigation
 - **Assists analysts to navigate from what is known to what is knowable**
 - **Enables analysts to more rapidly develop new insights**
 - **Amplifies and augments the analyst's inherently human cognitive ability, intuition, experience, and judgment**

Discovery: Navigating and Exploring the Information Space with Watson



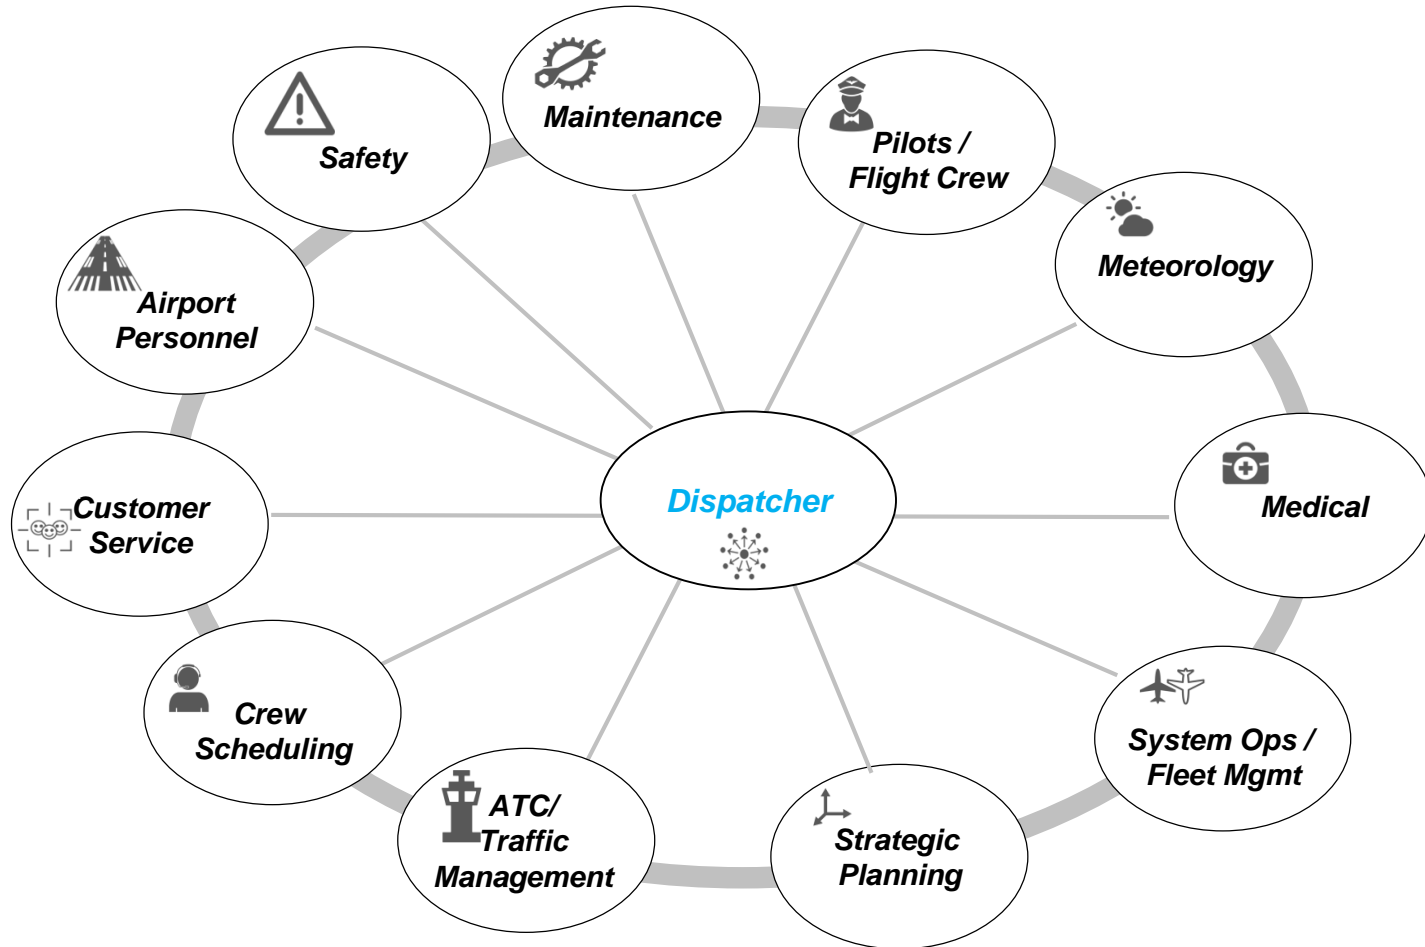
How Can Watson Help Decision Making in Operations Centers?





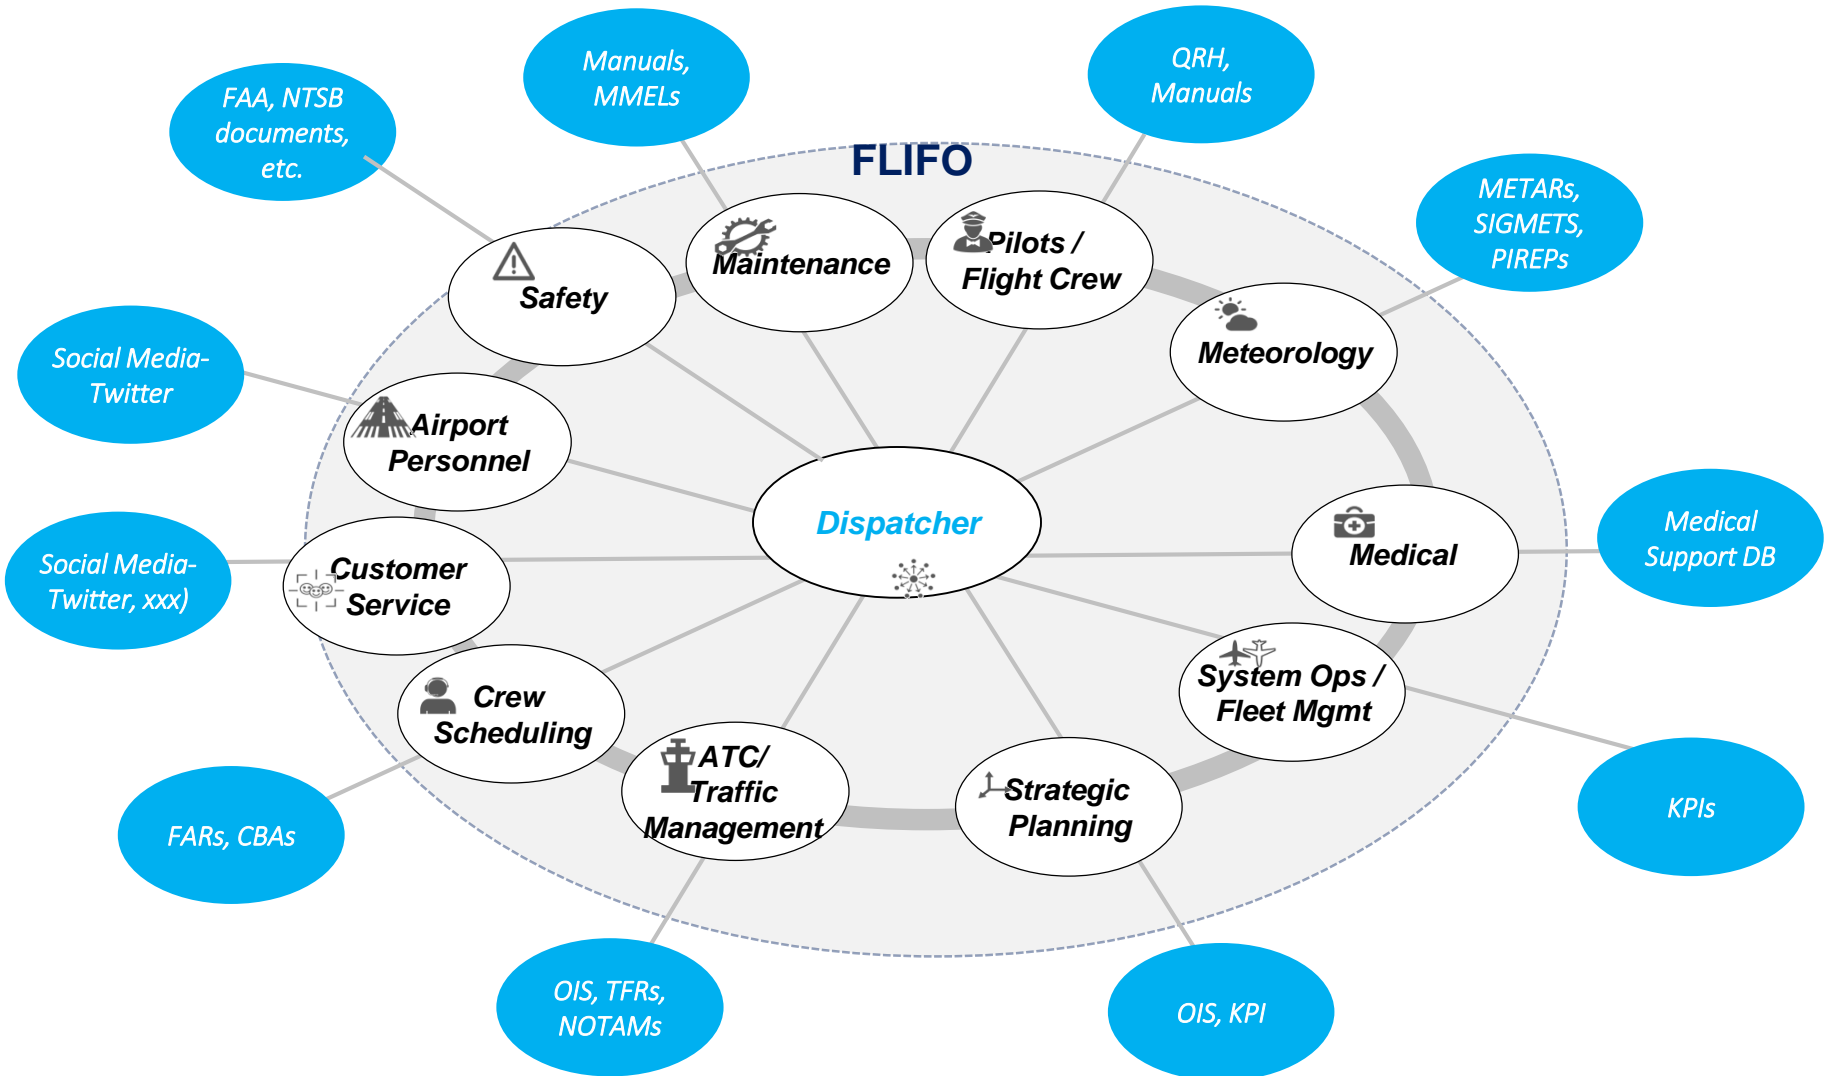


There are many different actors within the AOC*

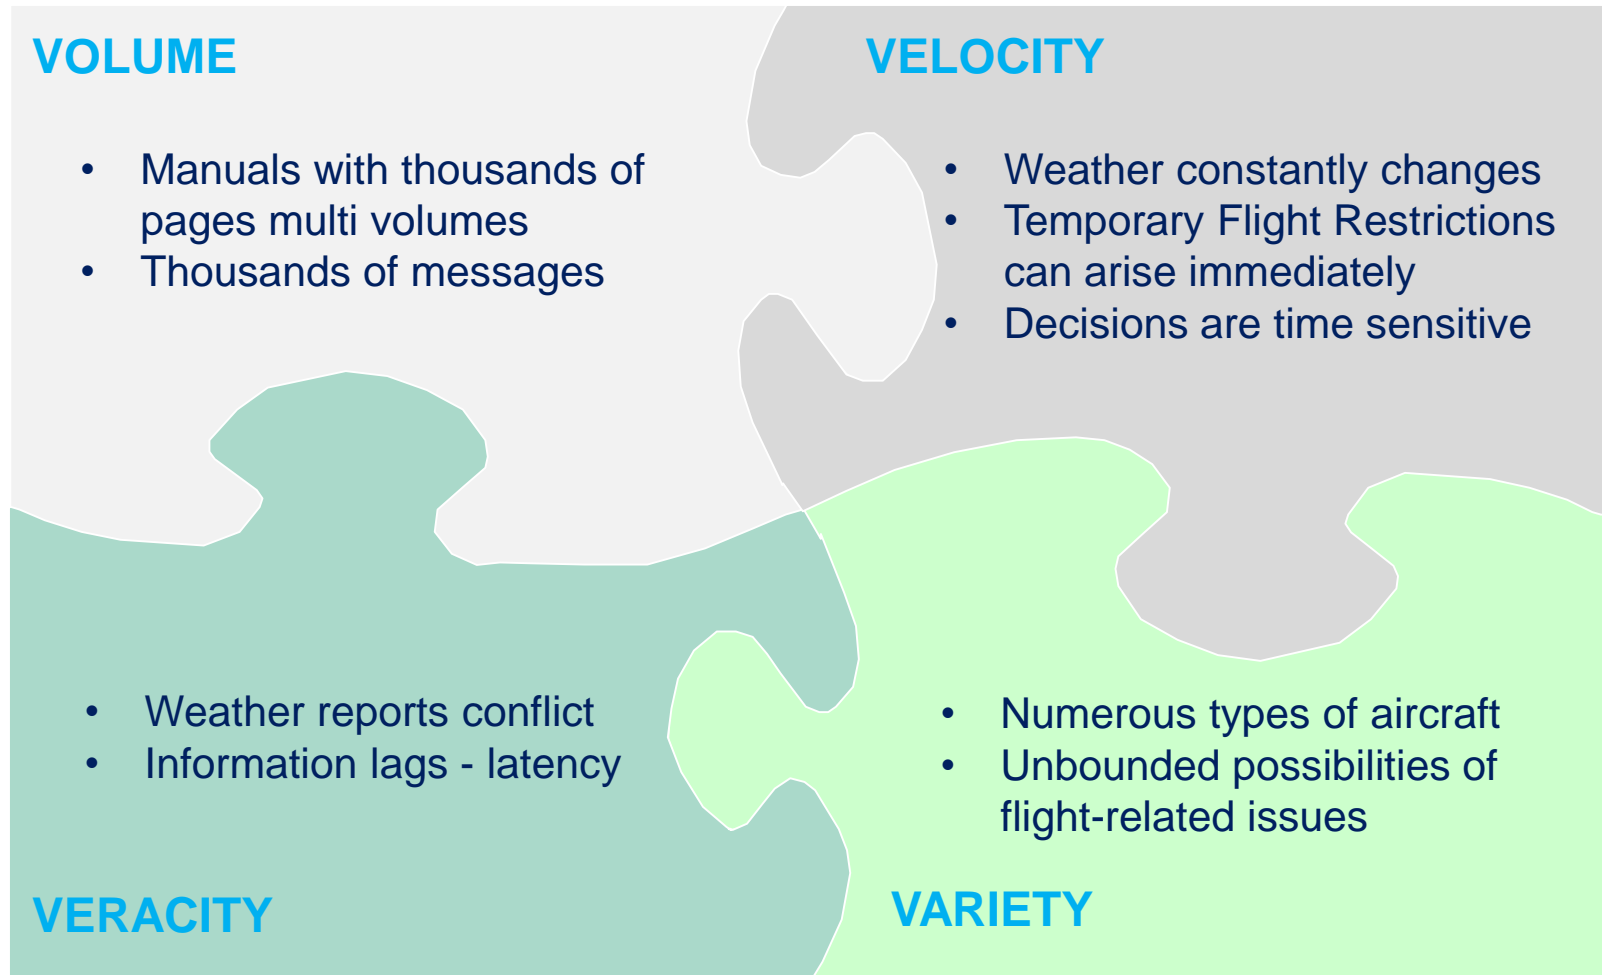


* Note- this is a representative sample, many other Actors (Security, Payroll, etc.) interact with Dispatchers.

Actors consult numerous different inputs for decision making




Big data is overwhelming the AOC






Process for Building a Watson System

- Selection and ingestion of documents to build a corpus
 - Pre-processing conditioning (if necessary)
 - Watson ingests documents to create its knowledge base
- Domain Training
 - Training sets (Q/A pairs) are created by subject matter experts
 - Training sets are used to train the machine learning models
 - Special lexicons are established (if any)
- User pilots
 - System is tested in mission domain by a small set of end users
 - Modifications are made to enhance effectiveness
- Staged deployment to production level
 - Expanded user base
 - Expand infrastructure to meet scaling objectives
 - Periodically expanded corpus
 - Integrate with existing analytics and other tools



Three deployment models



1. Managed Service

IBM owned and managed service hosted in US



2. "In Geography" Managed Service

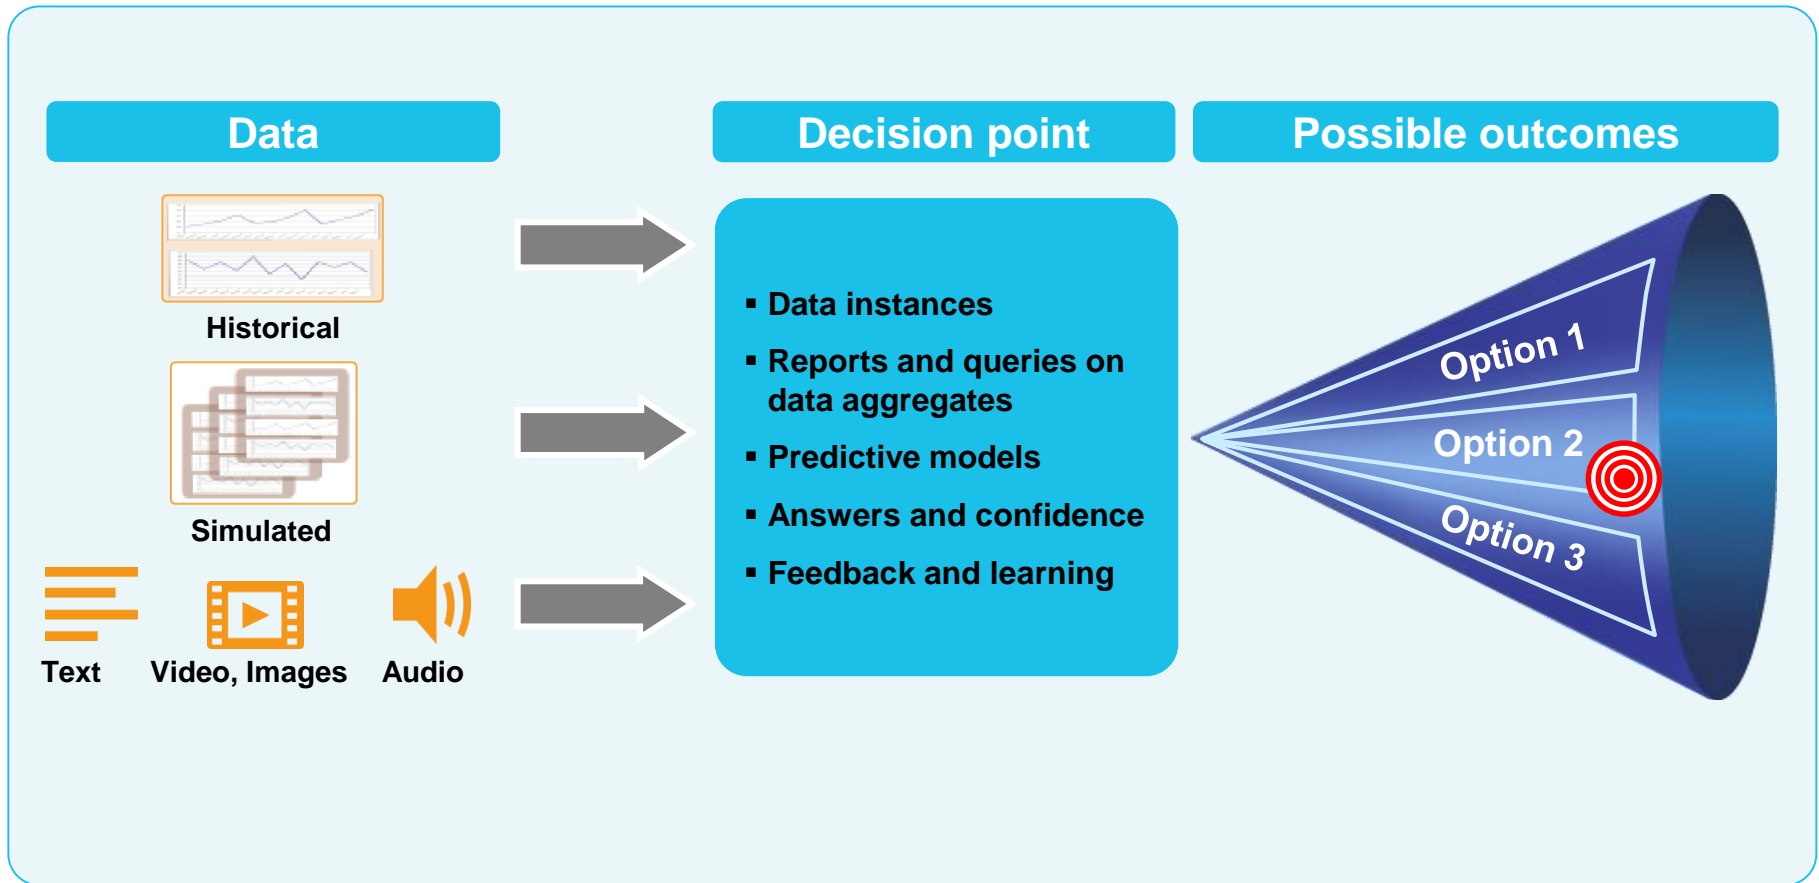
IBM owned and managed service hosted in client geography




3. On-premise Managed Service

IBM owned and managed service hosted in client data center

Broadening Capability in Decision Support



Watson is Cognitive Computing



Watson Understands Language

- Reads news, policies, information
- Interacts with language

Watson Learns with Experience

- Trains with experts and practice
- Improves with experience & feedback

Watson Describes Evidence

- Provides reasons behind thinking
- Increases trust and confidence

