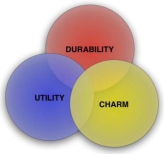

Digital Readiness: AI/ML
Finding a doing machine

Gregg Vesonder
Stevens Institute of Technology
http://vesonder.com

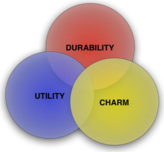

Three Talks

- ~~Digital Readiness: AI/ML, The thinking system quest.~~
  - ~~Artificial Intelligence and Machine Learning (AI/ML) have had a fascinating evolution from 1950 to the present. This talk sketches the main themes of AI and machine learning, tracing the evolution of the field since its beginning in the 1950s and explaining some of its main concepts. These eras are characterized as "from knowledge is power" to "data is king".~~
- Digital Readiness: AI/ML, Finding a doing machine.
  - In the last decade Machine Learning had a remarkable success record. We will review reasons for that success, review the technology, examine areas of need and explore what happened to the rest of AI, GOFAI (Good Old Fashion AI).
- Digital Readiness: AI/ML, Common Sense prevails?
  - Will there be another AI Winter? We will explore some clues to where the current AI/ML may reunite with GOFAI (Good Old Fashioned AI) and hopefully expand the utility of both. This will include extrapolating on the necessary melding of AI with engineering, particularly systems engineering.

Roadmap

- Systems
  - Watson
  - CYC
  - NELL
  - Alexa, Siri, Google Home
- Technologies
  - Semantic web
  - GPUs and CUDA
  - Back office (Hadoop)
  - ML Bias
- Last week’s questions

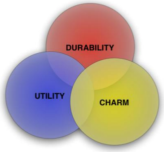

Winter is Coming?

- First Summer: Irrational Exuberance (1948 – 1966)
- First Winter (1967 – 1977)
- Second Summer: Knowledge is Power (1978 – 1987)
- Second Winter (1988 – 2011)
- Third Summer (2012 – ?)
- Why there might not be a third winter!

Henry Kautz – Engelmore Lecture

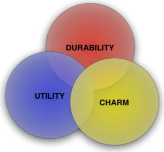

SYSTEMS

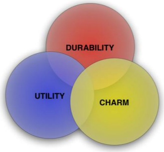

Winter 2 Systems

- Knowledge is power theme
- Influence of the web, try to represent all knowledge
  - Creating a general ontology organizing everything in the world into a hierarchy of categories
  - Successful deep ontologies: Gene Ontology and CML Chemical Markup Language
- Indeed extreme knowledge
  - CYC and Open CYC
  - IBM's Watson

Upper ontology of the world

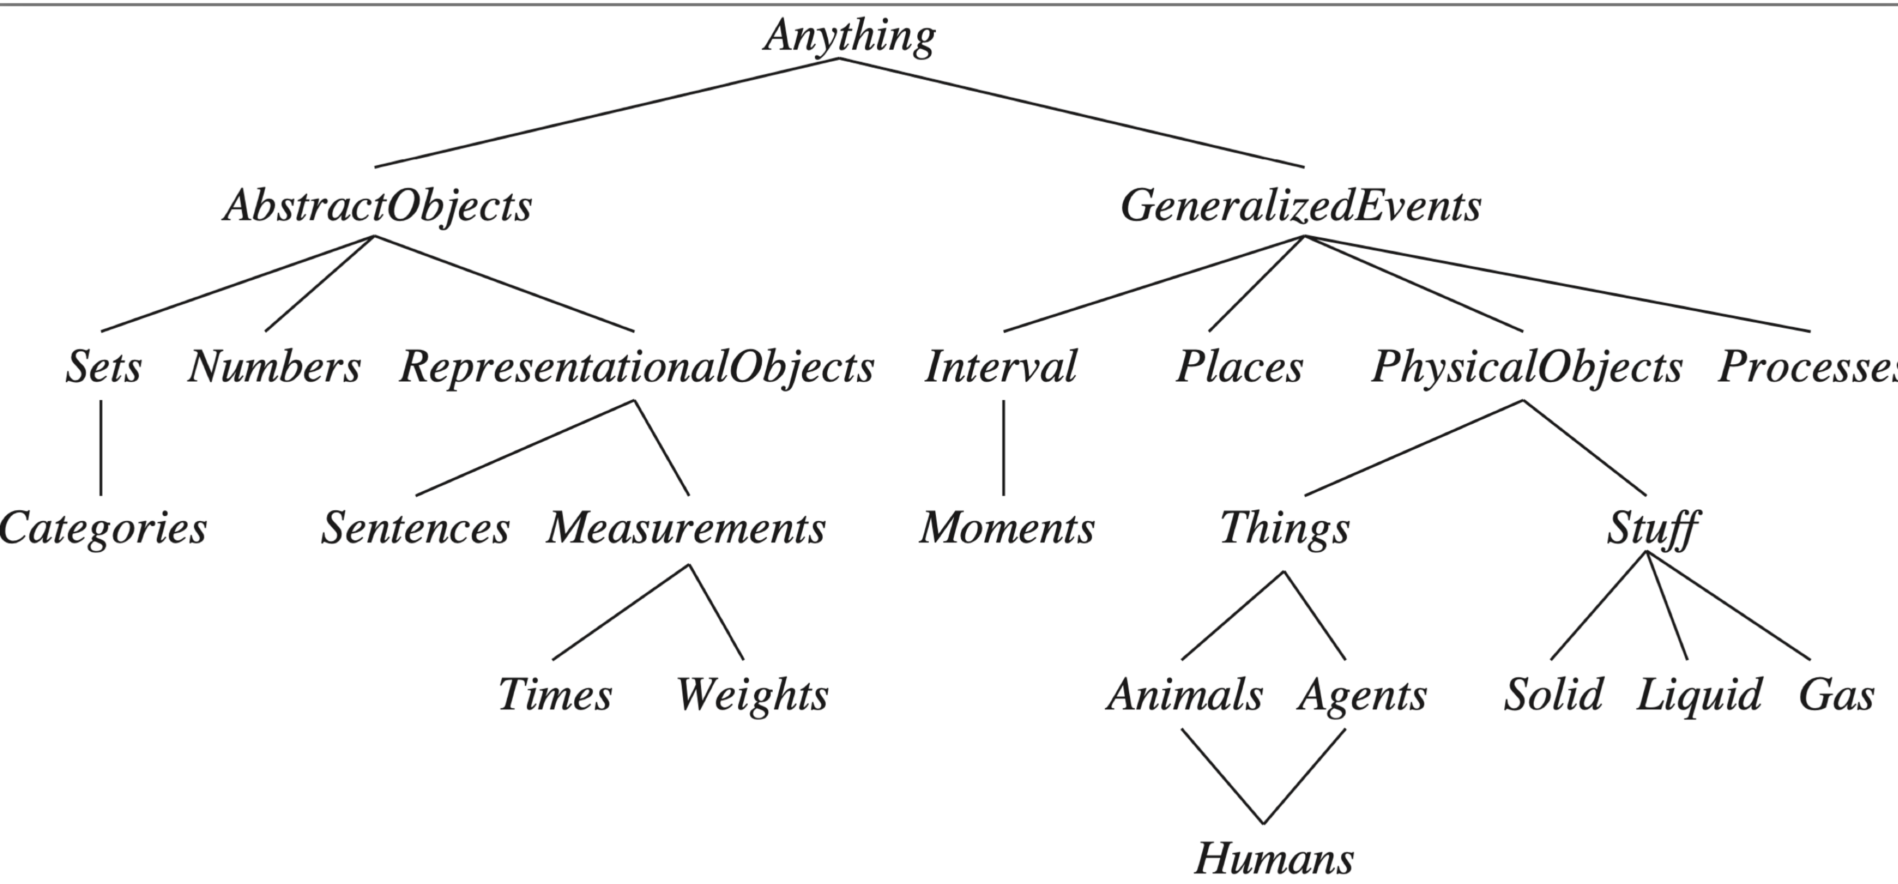

Russell and Norvig figure 12.1

Properties of a subject area and how they are related

Ferrucci, D., et.al. “Building Watson: An overview of the DeepQA Project”, AI Magazine, Fall 2010.

Watson

- IBM Watson question answering computer system
  - DeepQA project
- Famous for winning a Jeopardy contest against Ken Jennings (KenJen) and Brad Rudder, 2 champions
  - Jeopardy is difficult, ambiguous questions and average of 3 seconds to ring in
  - Excluded A/V questions and special instruction questions
- Applications: tax preparation, teaching assistant, medicine, …
- “ … Create general purpose, reusable natural language processing (NLP) and knowledge representation and reasoning (KRR) that can exploit as-is natural language resources and as-is structured knowledge rather than curate task-specific knowledge resources.”

Distribution of Jeopardy Questions

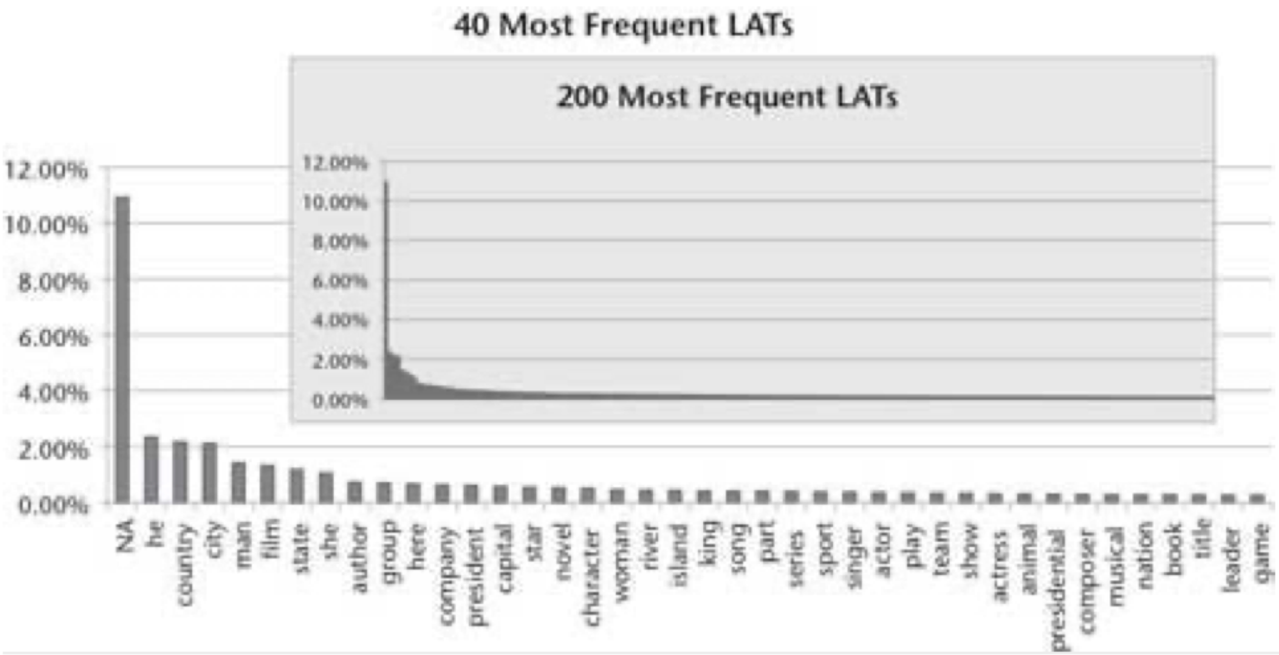

Figure 1. Lexical Answer Type Frequency.

KenJen

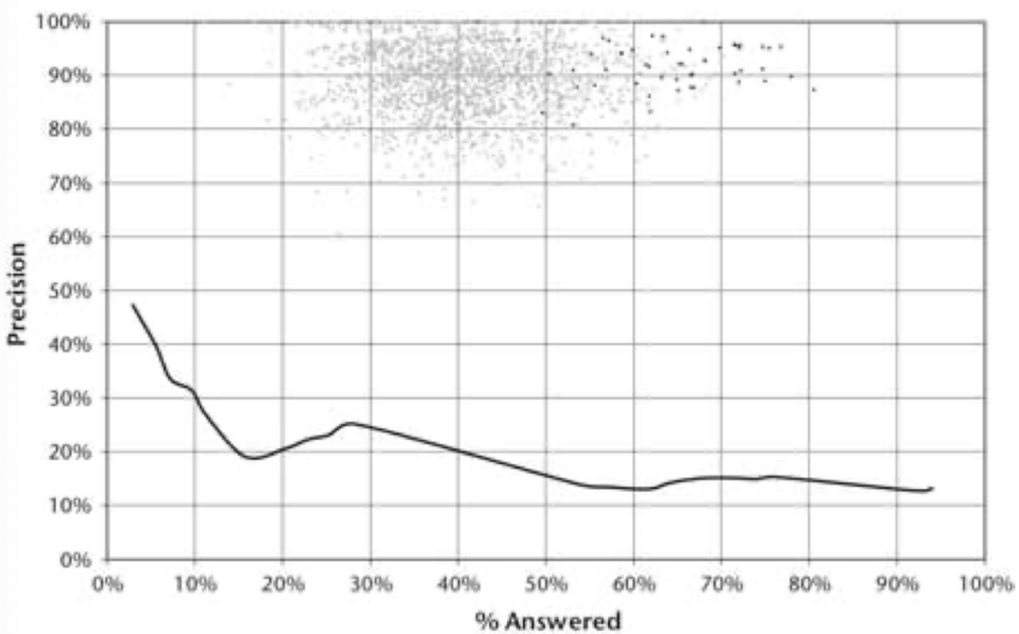

Figure 4. Baseline Performance.

The darker dots on the graph represent Ken Jennings's games. Ken Jennings had an unequaled winning streak in 2004, in which he won 74 games in a row. Based on our analysis of those games, he acquired on average 62 percent of the questions and answered with 92 percent precision. Human performance at this task sets a very high bar for precision, confidence, speed, and breadth.

TREC Text Retrieval Contest
https://trec.nist.gov/

- to encourage research in information retrieval based on large test collections;
- to increase communication among industry, academia, and government by creating an open forum for the exchange of research ideas;
- to speed the transfer of technology from research labs into commercial products by demonstrating substantial improvements in retrieval methodologies on real-world problems; and
- to increase the availability of appropriate evaluation techniques for use by industry and academia, including development of new evaluation techniques more applicable to current systems.

DeepQA Technology

- Used more than 100 different techniques
- Not specific to Jeopardy but specific to QA
- Technological Principles involved:
  - Massive parallelism
  - Many Experts
  - Pervasive confidence estimates
  - Integrate shallow and deep knowledge – leveraging many loosely formed ontologies
  - *No one algorithm dominates*

DeepQA High-Level Architecture

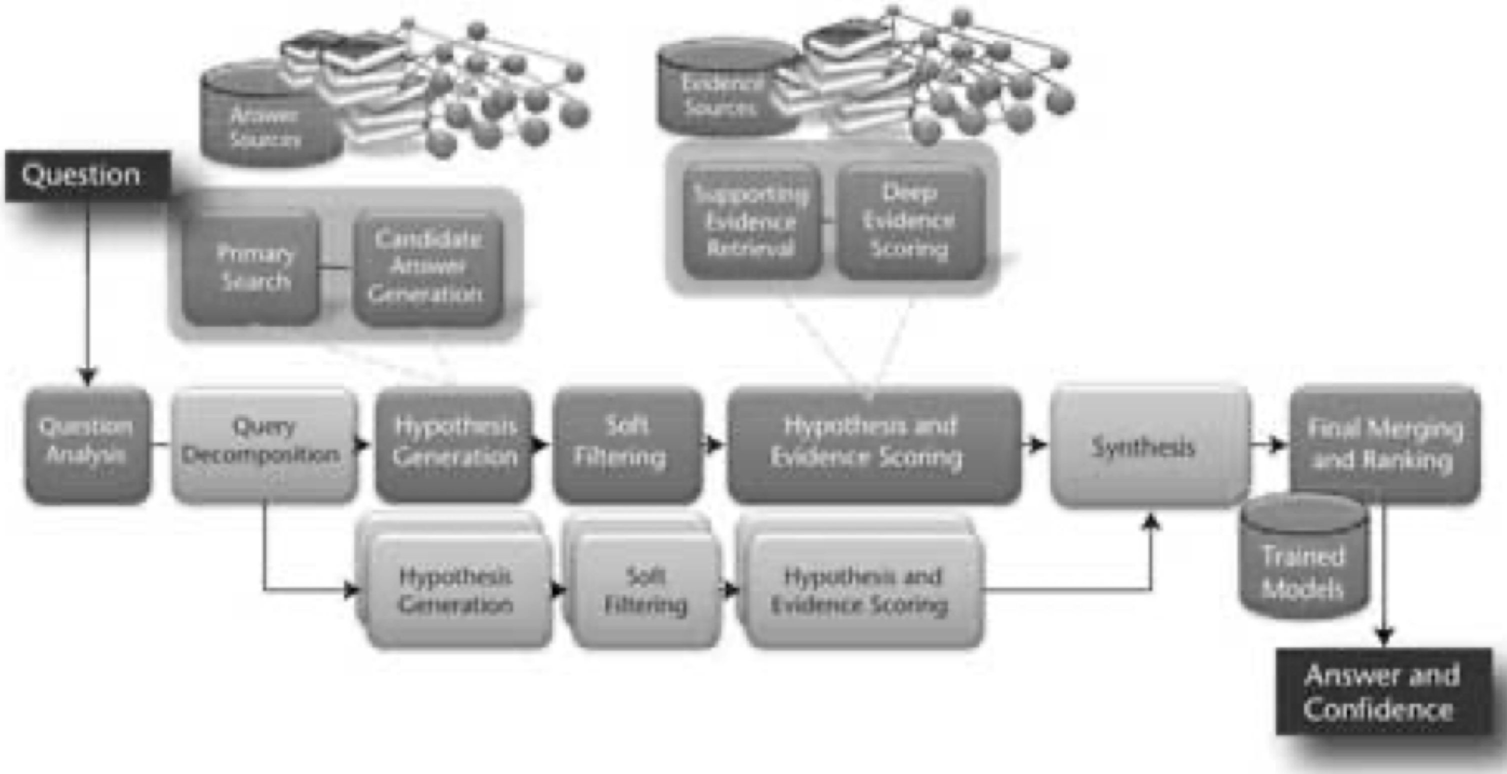

DeepQA - wikipedia

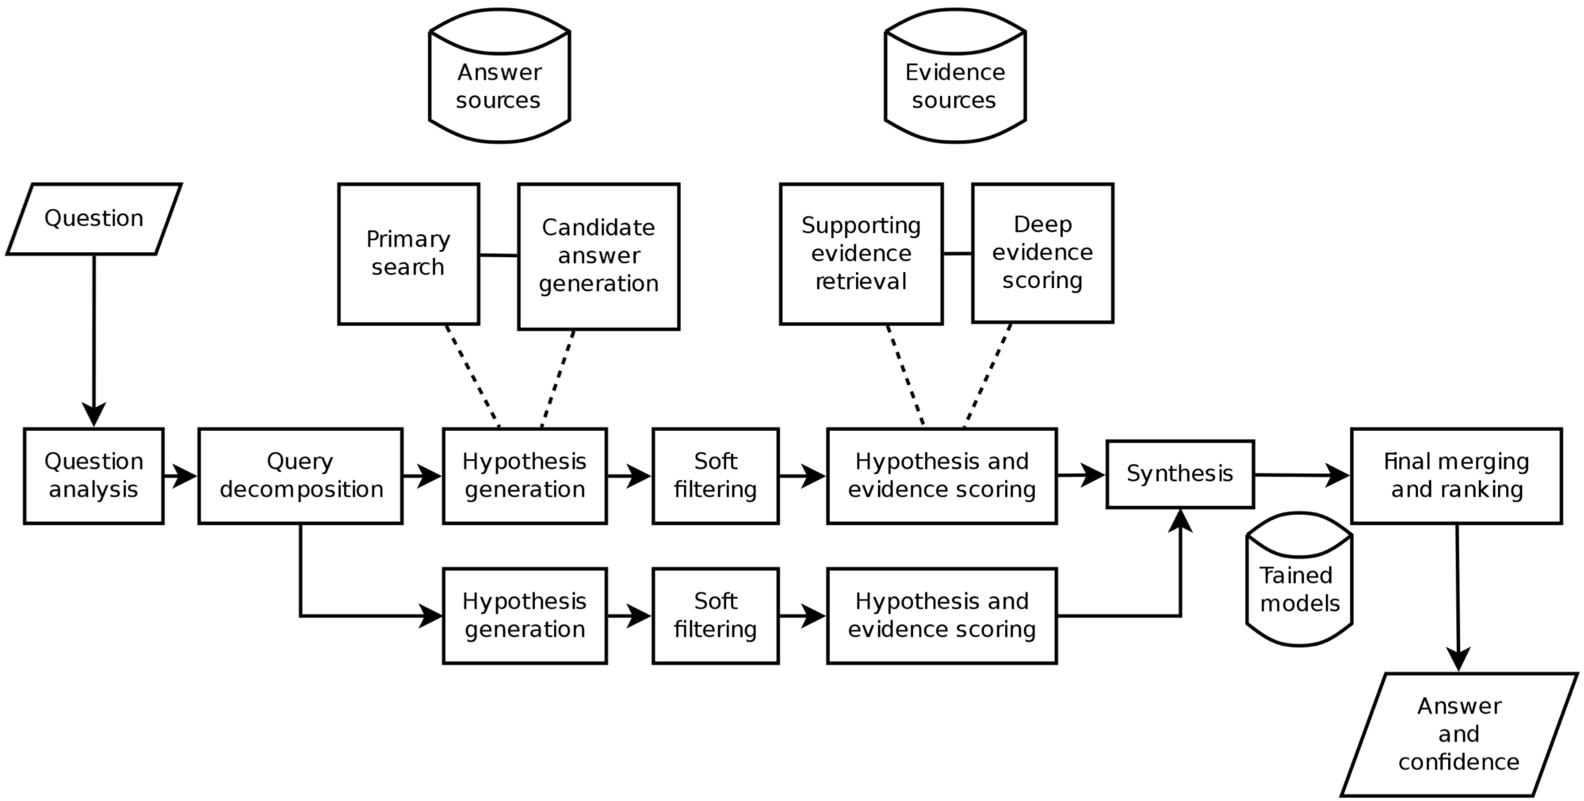

Deep QA/Watson

- Uses Hadoop for distributed computing
- Watson can process 500 gigabytes/sec, equivalent of 1 million books/sec
- Resources included taxonomies and ontologies including DBpedia, WordNet and YAGO

CYC

- Goal to assemble a comprehensive ontology and knowledge base that spans the basic concepts and rules about how the world works hoping to capture common sense and implicit/tacit knowledge.
- Combat brittle AI systems
- Began in 1984 by Doug Lenat
- Also OpenCyc (CYC) and ResearchCYC (Cyc)

Thanks Wikipedia

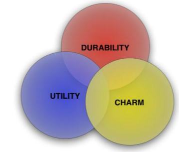

CYC factoids

- Ontology (2017) 1.5 Million terms:
  – 416K collections e.g., fish and types of fishing
- CYC KB of common-sense rules and assertions were largely hand-written 24.5 million taking >1K person years!
- Uses community of agents (blackboard) architecture
- ”every tree is a plant”, “plants die eventually”, asked whether trees die it will infer correctly

CYC Elements

- *Individual items* #$BillClinton or #$France (CycL)
- *Collections* #$Tree-ThePlant – contains all trees
- Functions produce new terms from given ones. #$FruitFn arg will return collection of fruits for a provided type or collection of plants

Applications of CYC

- Pharmaceutical Term Thesaurus Manager/Integrator (Glaxo)
- Terrorism Knowledge Base
- MathCraft – 6^{th} grade level help in understanding math

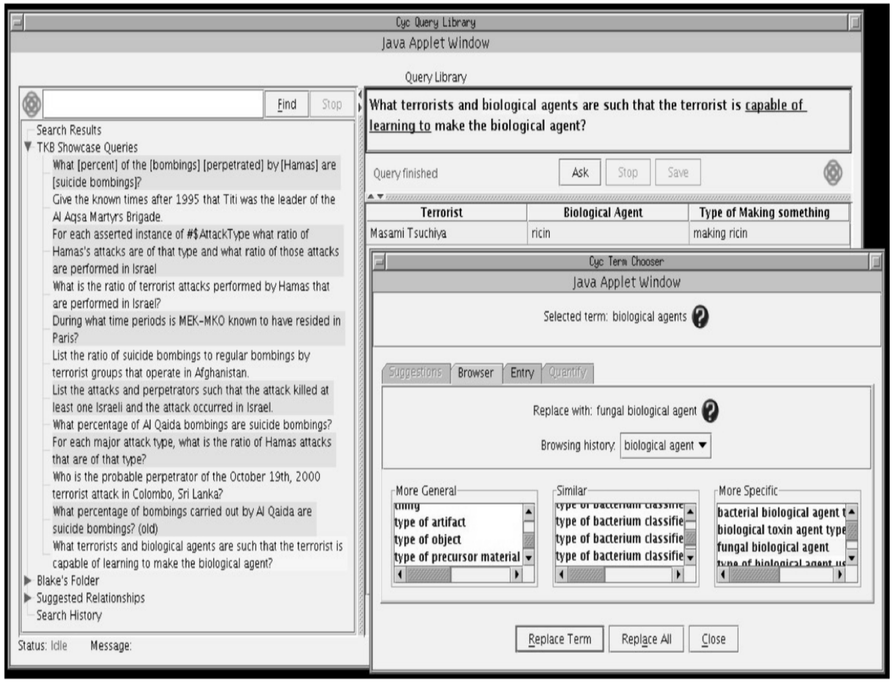

Figure 1. The Query Library

Lots of Opinions on CYC/cyc/Cyc
but true to GOFAI

NELL

- Never Ending Language Learner, Tom Mitchell, Carnegie Mellon University
- We will never truly understand machine or human learning until we can build computer programs that, like people:
  - learn many different types of knowledge or functions,
  - from years of diverse, mostly self-supervised experience,
  - in a staged curricular fashion, where previously learned knowledge enables learning further types of knowledge,
  - where self-reflection and the ability to formulate new representations and new learning tasks enable the learner to avoid stagnation and performance plateaus.
- T. Mitchell, et.al. AAAI 2015

NELL-2

- Given:
  - an initial ontology defining categories (e.g., Sport, Athlete) and binary relations (e.g., AthletePlaysSport(x,y)),
  - approximately a dozen labeled training examples for each category and relation (e.g., examples of Sport might include the noun phrases "baseball" and "soccer"),
  - the web (an initial 500 million web pages from the ClueWeb 2009 collection (Callan and Hoy 2009), and access to 100,000 Google API search queries each day),
  - occasional interaction with humans (e.g., through NELL's public website http://rtw.ml.cmu.edu);

NELL -3

- DO:
  - Run 24 hours/day, forever, and each day:
    - 1. read (extract) more beliefs from the web, and remove old incorrect beliefs, to populate a growing knowledge base containing a confidence and provenance for each belief,
    - 2. learn to read better than the previous day.

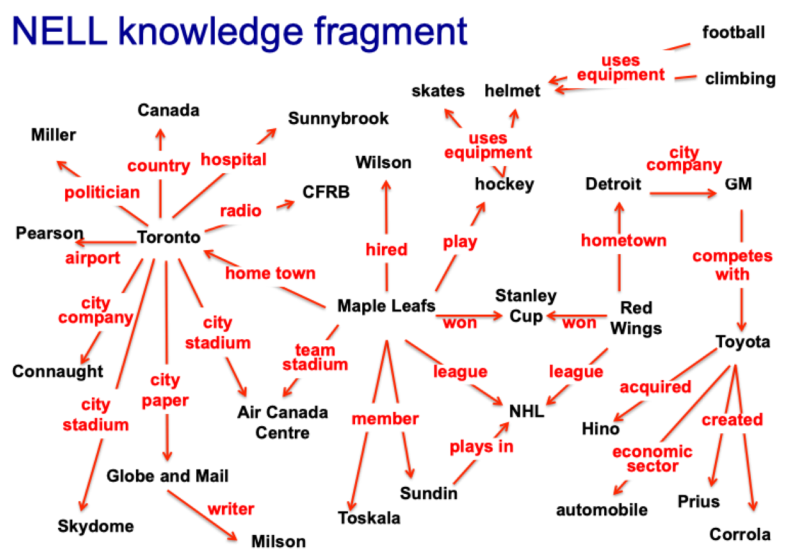

Figure 1: **Fragment of the 80 million beliefs NELL has read from the web.** Each edge represents a belief triple (e.g., play(MapleLeafs, hockey), with an associated confidence and provenance not shown here. This figure contains only correct beliefs from NELL's KB – it has many incorrect beliefs as well since NELL is still learning.

NELL-4

Knowledge Base uses
Blackboard Architecture

NELL Architecture

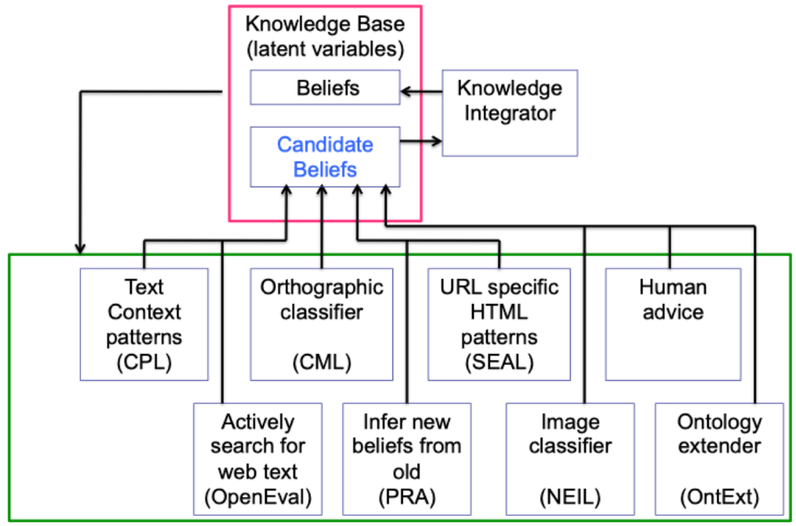

Figure 2: **NELL's software architecture.** NELL's growing knowledge base (KB) serves as a shared blackboard through which its various reading and inference modules interact. NELL's learning cycle iteratively retrains these software modules using the current KB, then updates the KB using these refined modules.

Mitchell, et.al. (2015)

NELL Results

High confidence indicates that
one of NELL's modules assigns
a 0.9 confidence to the belief
or multiple modules propose the belief

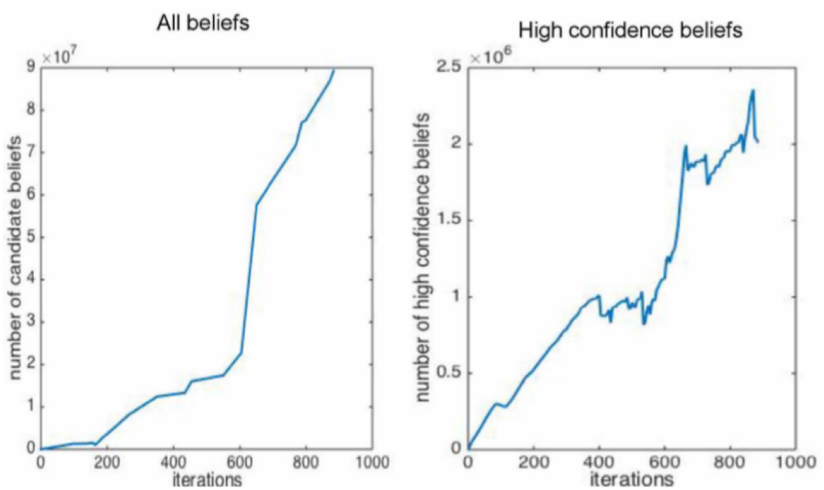

Figure 3: **NELL KB size over time.** Total number of beliefs (left) and number of high confidence beliefs (right) versus iterations. Left plot vertical axis is tens of millions, right plot vertical axis is in millions.

Mitchell, et.al. 2015

Alexa

- Siri, Google Assistant too

AWS

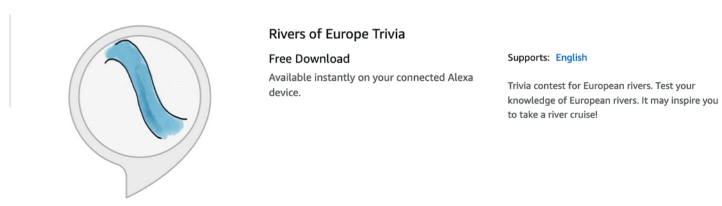

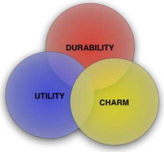

Technologies

Semantic Web

Semantic Web

Big Data, API/Platform, & M2M Intersection

•What is Semantic Web?

•A collaborate movement led by W3C to promote common data format aimed at converting the current web dominated by unstructured and semi-structured document into a "web of data"

•Components of Semantic Web

- Resource Description Framework (RDF)
- RDF Schema
- Web Ontology Language (OWL)

•Implications of Semantic Web

- **Link heterogeneous data** sources on web
- Drive **"meaning"** of information on the web
- Enable **machine-readable** web
- Enable **intelligent** services
- Improve quality of search, targeted ads, etc.

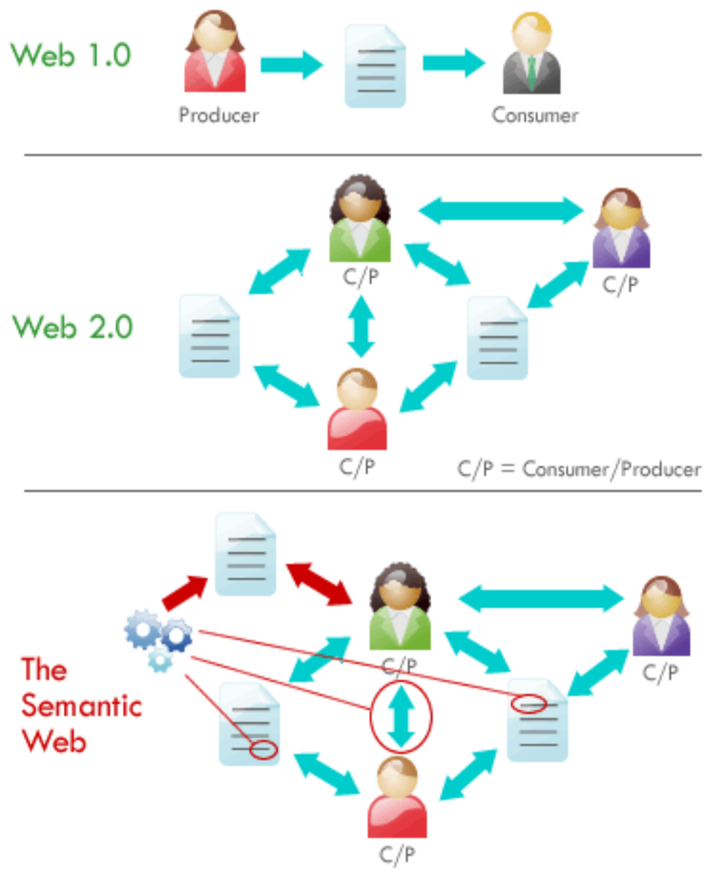

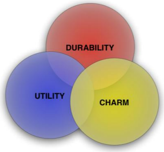

Early Adopters of Semantic Web

Synergy with Effortless Customer Experince

schema.org project: a joint effort of Google, Yahoo!, and Microsoft documents the concepts and attributes that these three search engines will look for when Web publishers include Semantic Web data on their sites. That is RDF like data.

Typical Applications

- Optimized search engine
- Knowledge graph
- Knowledge management
- Speech recognition
- Content management
- Customer/product driven analytics
- Customer acquisition & retention
- Predictive analytics for security
- Visual discovery

Early Adopters

BBC WORLD NEWS, YAHOO!, BEST BUY, bing, Facebook, IBM Watson, Microsoft, Siri, Google

- 30% increase in traffic (Best Buy). 15% increase in CTR (Yahoo!).

External Adopters - Semantic Technology In The Wild

- Department of Defense
- US Intelligence Agencies
- Library of Congress
- Wikipedia - DBpedia
- NASA
- Amdocs
- BestBuy
- BBC
- MITRE

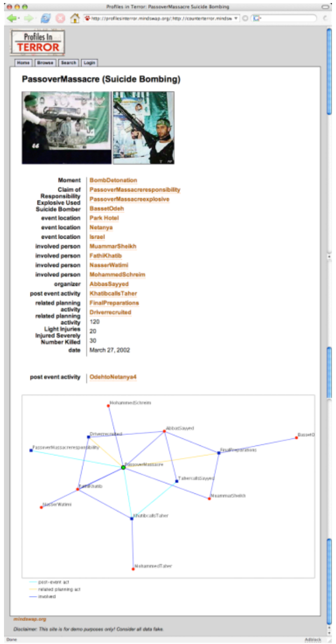

Internet Semantic Web Linked Data

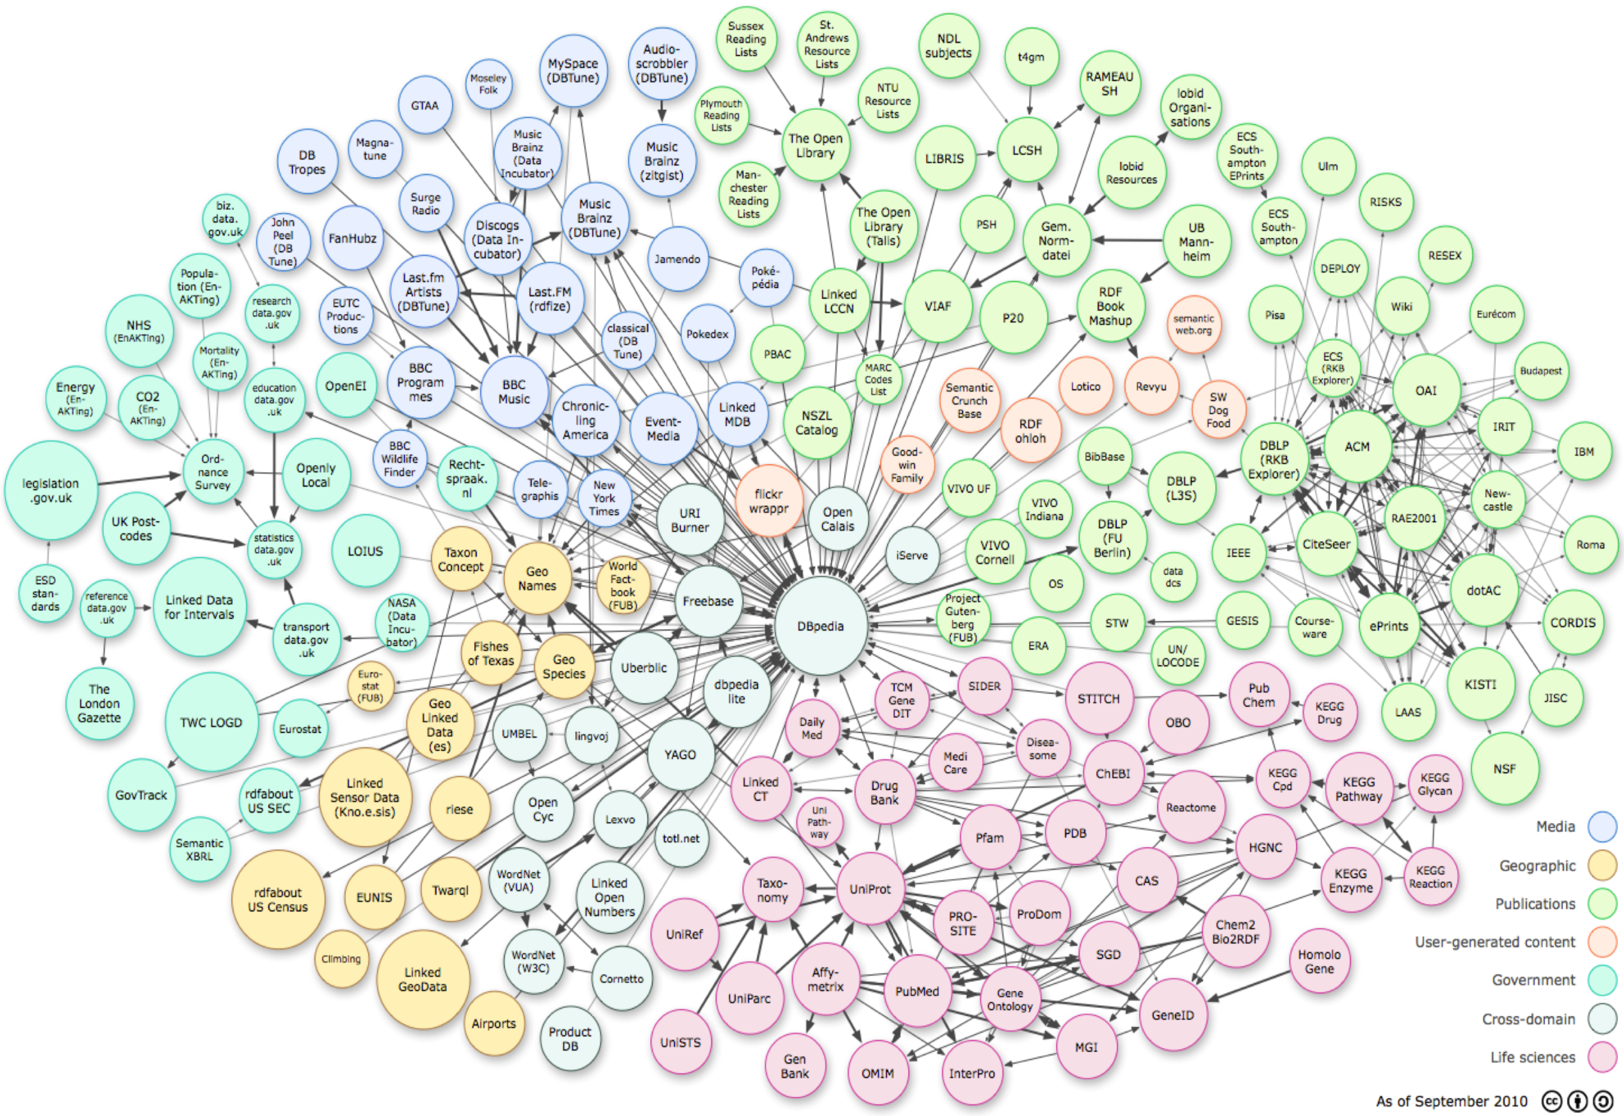

Semantic Web – Background

- Tim Berners-Lee 1999
- ***"I have a dream for the Web [in which computers] become capable of analyzing all the data on the Web – the content, links, and transactions between people and computers. A 'Semantic Web', which should make this possible, has yet to emerge, but when it does, the day-to-day mechanisms of trade, bureaucracy and our daily lives will be handled by machines talking to machines. The 'intelligent agents' people have touted for ages will finally materialize"***

- The dream is beginning to emerge and it's all based on W3C Standards

- Alias: Web of Data Web 3.0
- Linked Data Machine Readable Web

- Our goal is to understand this concept, it's formats and technologies and begin to explore some of the uses and power of this technology.

SemTech Terms

- Resource Description Framework (RDF)
- RDF Schema (RDFS)
- OWL (Web Ontology Language)
- SPARQL
- TripleStore - GraphDatabase
- Inferring
- Metadata
- RDFa
- Microformats
- Ontology Modeling

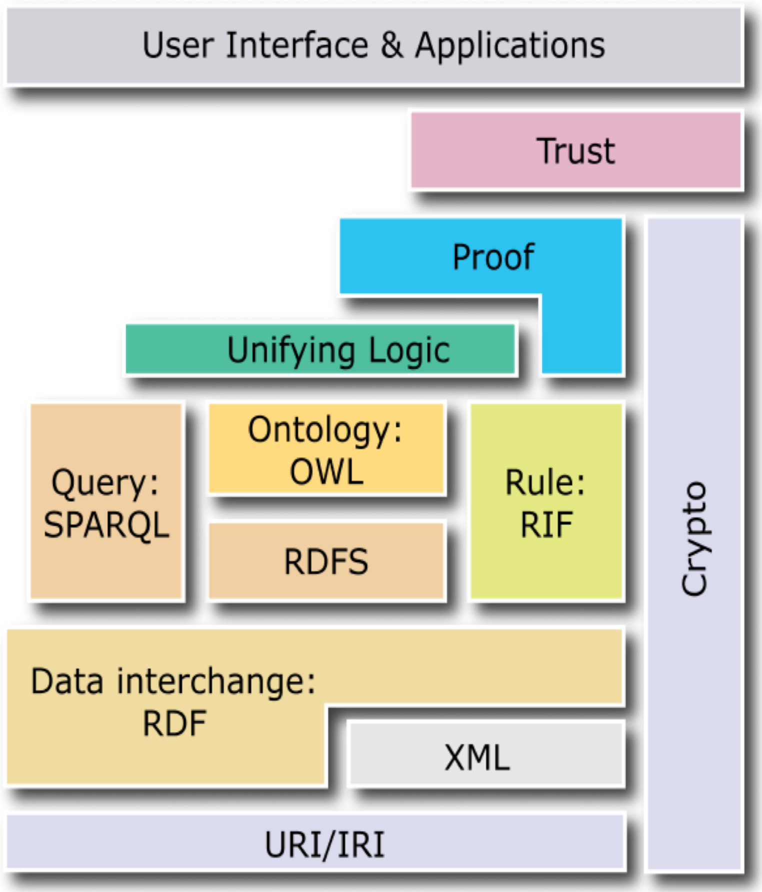

Today – Data and Information

- **Web Search:**
- ❑ Keyword Based "Results"
- ❑ Provides a List of Links
- ❑ Answer to Question?

- **Business Applications:**
- ❑ Query based results
- ❑ Range of information available limited due to "silo" data
- ❑ Good applications setup join keys and cross applications but this is limited and difficult

- **New Goals:**
- ❑ Queries provided answers – not links
- ❑ Introduction of Intelligent Applications
- ❑ Allow for discovery and exploration of information
- ❑ Breakdown data silo's barriers

Promise of the Semantic Technology

- Data becomes self describing
  - ❑ Metadata extends “knowledge” beyond rows and columns
  - ❑ Intranet and Internet web of data

- Linked data extends data accessibility
  - ❑ Data is shared and related across silo’s
  - ❑ Linked Open Data from the Internet
  - ❑ Inclusion of non-structured data

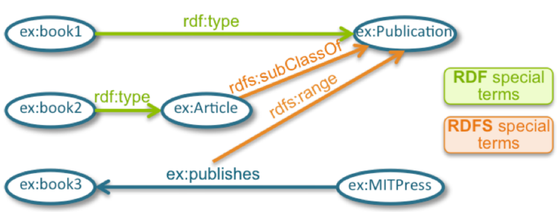

- Emergence of intelligent applications
  - ❑ Computer based infering
  - ❑ Graph based pattern analysis

- Intelligent agents – 21^{st} Century Business
  - ❑ SemTech and Intelligent Agents may enable a new online services and commerce
  - ❑ Data Pull (not push) will be a disruptive force

Example of Linked Open Data (Dbpedia for Acxiom)

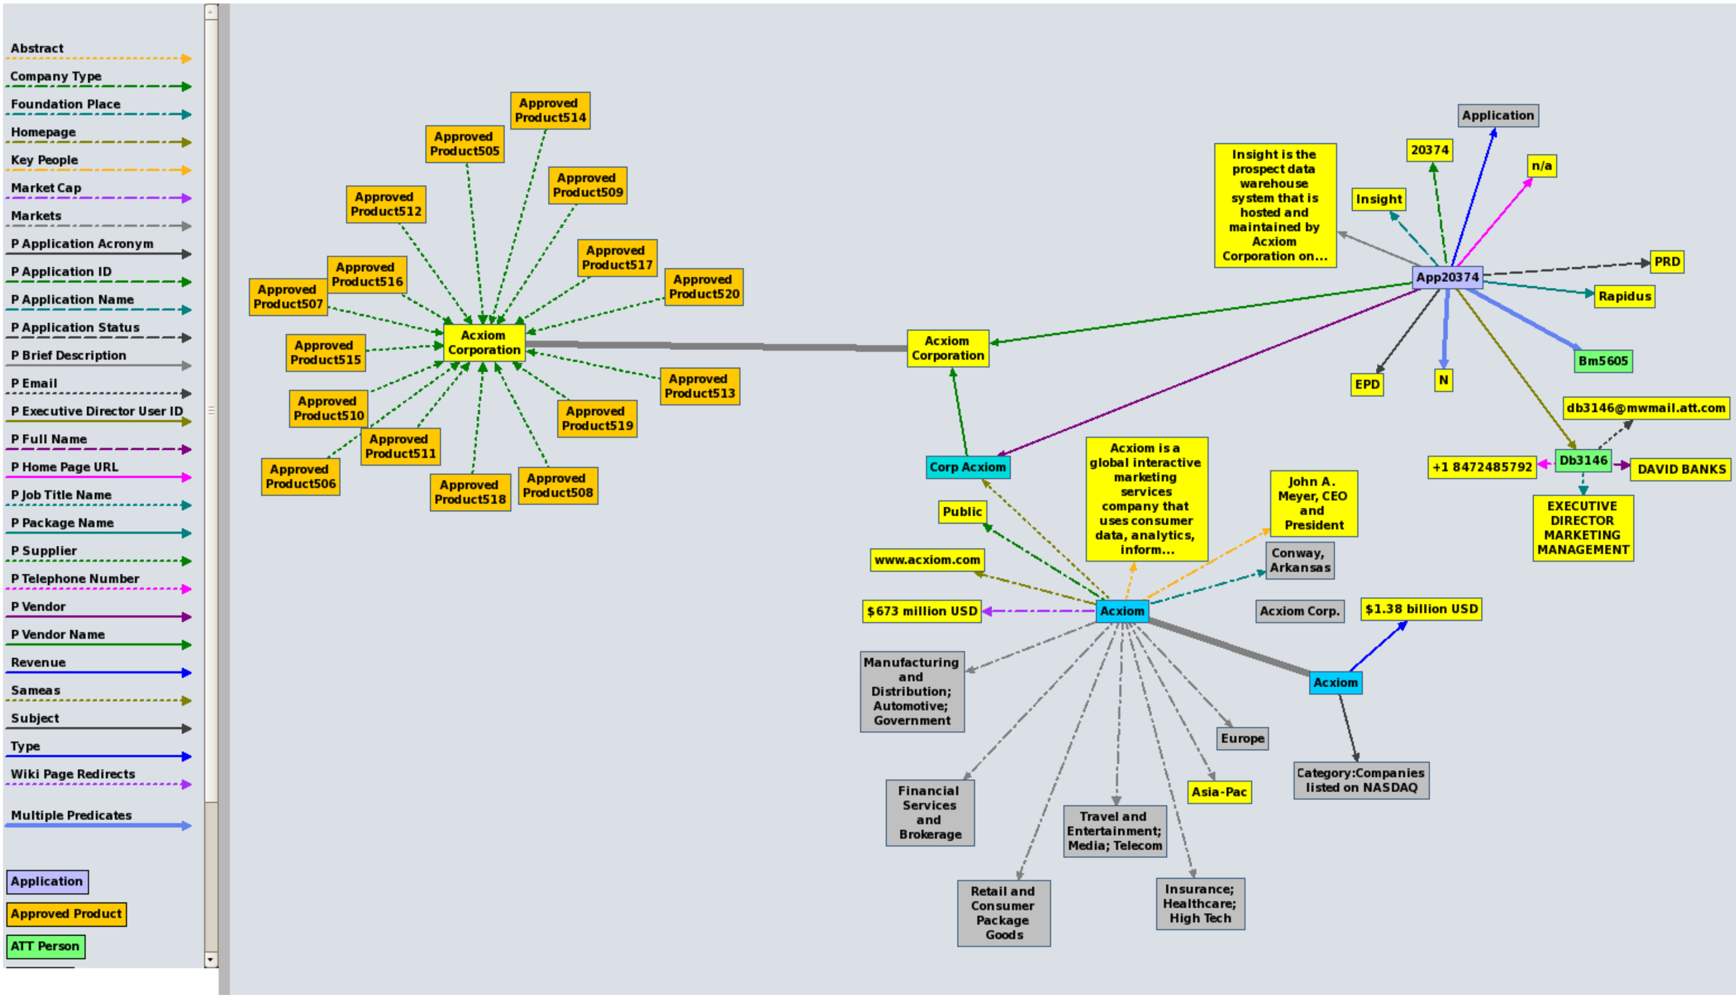

schema.org

- Mission: "create, maintain, and promote schemas for structured data on the Internet, on web pages, in email messages, and beyond.”
- Founded by Google, Microsoft, Yahoo and Yandex
- Search Engine Optimization (SEO)

AI/ML Infrastructure

GPUs and CUDA

- GPU = Graphics Processing Unit, CUDA = Compute Unified Device Architecture
- CUDA, Ian Buck Stanford grad student now with INVIDIA
- Began with video games and graphic techniques
- Threads execute a kernel program – the identified paralizable parts in your program
- C, C++, python plus API

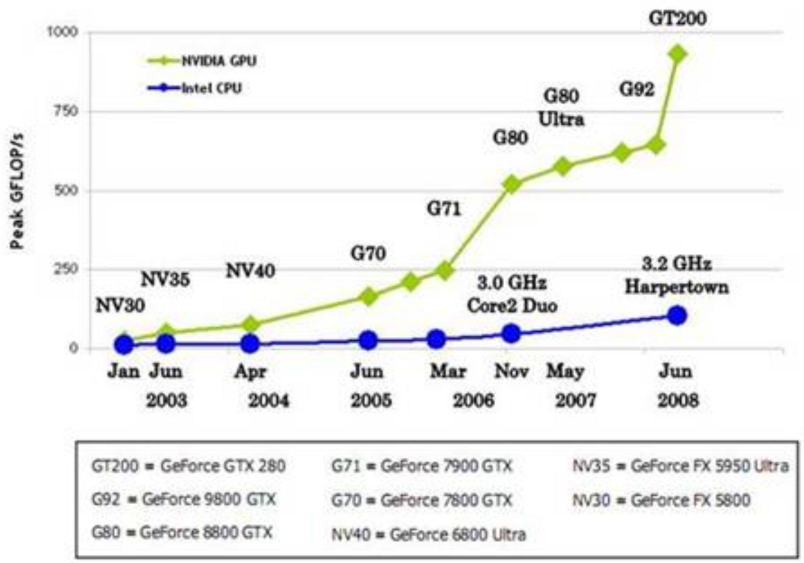

http://cuda.ce.rit.edu/cuda_overview/cuda_overview.htm

Hardware

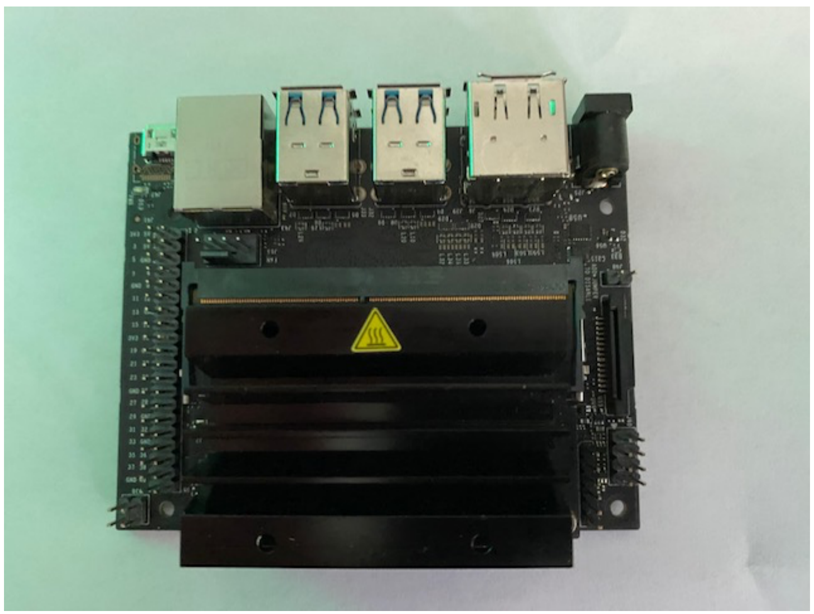

NVIDIA Jetson

NVIDIA Volta™ architecture
with 384 NVIDIA® CUDA® cores and 48 Tensor cores

Hadoop

- Hadoop
  - Distributed file system
  - Map reduce
  - Hadoop common – tools to read data stored under Hadoop file system
  - YARN – resource manager
  - Offered by a variety of vendors: Microsoft, Hortonworks

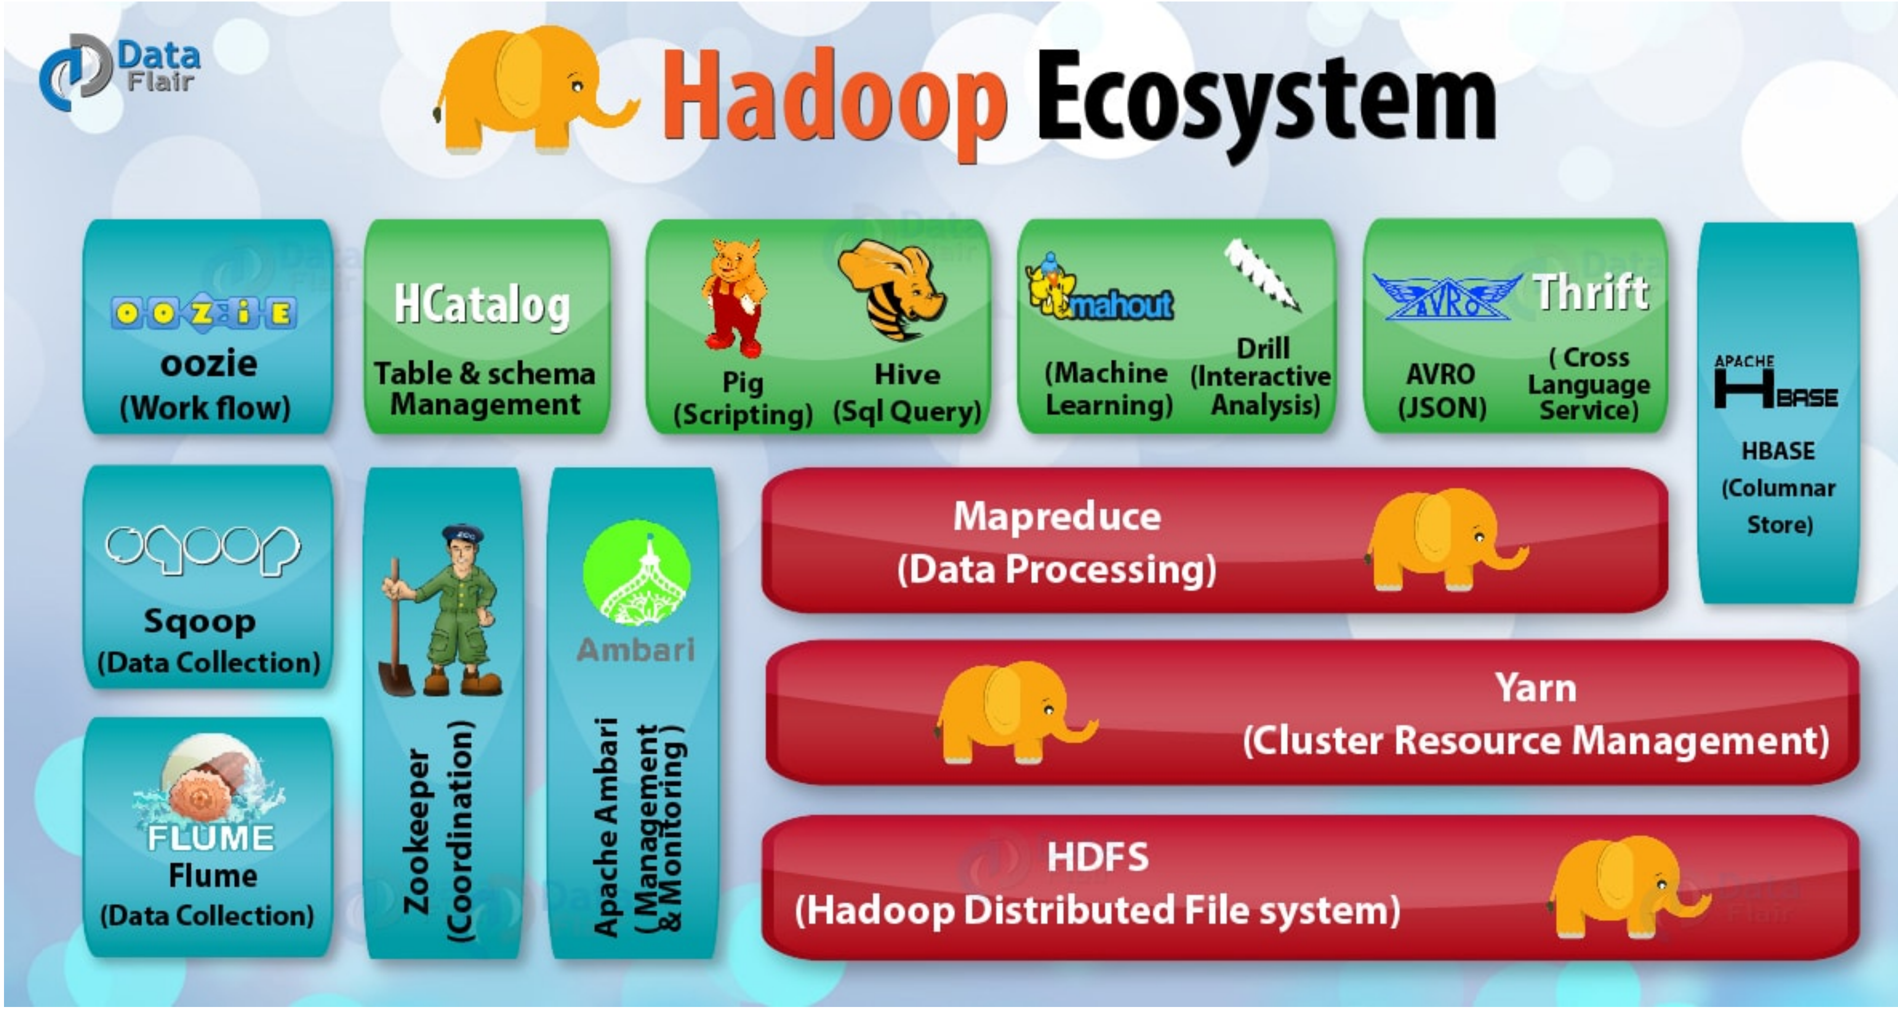

https://data-flair.training/blogs/hadoop-ecosystem-components/

Mapreduce

- Takes a set of input key-value pairs and produces a set of output key-value pairs
  - map does this in parallel
  - reduce consolidates the results

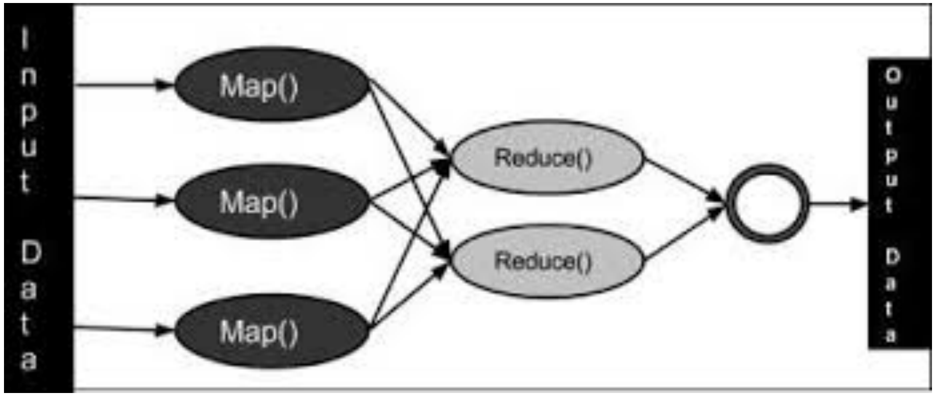

(write (mapcar '1+ '(23 34 45 56 67 78 89)))
When you execute the code, it returns the following result –
(24 35 46 57 68 79 90)

Tutorialspoint.com

Storage

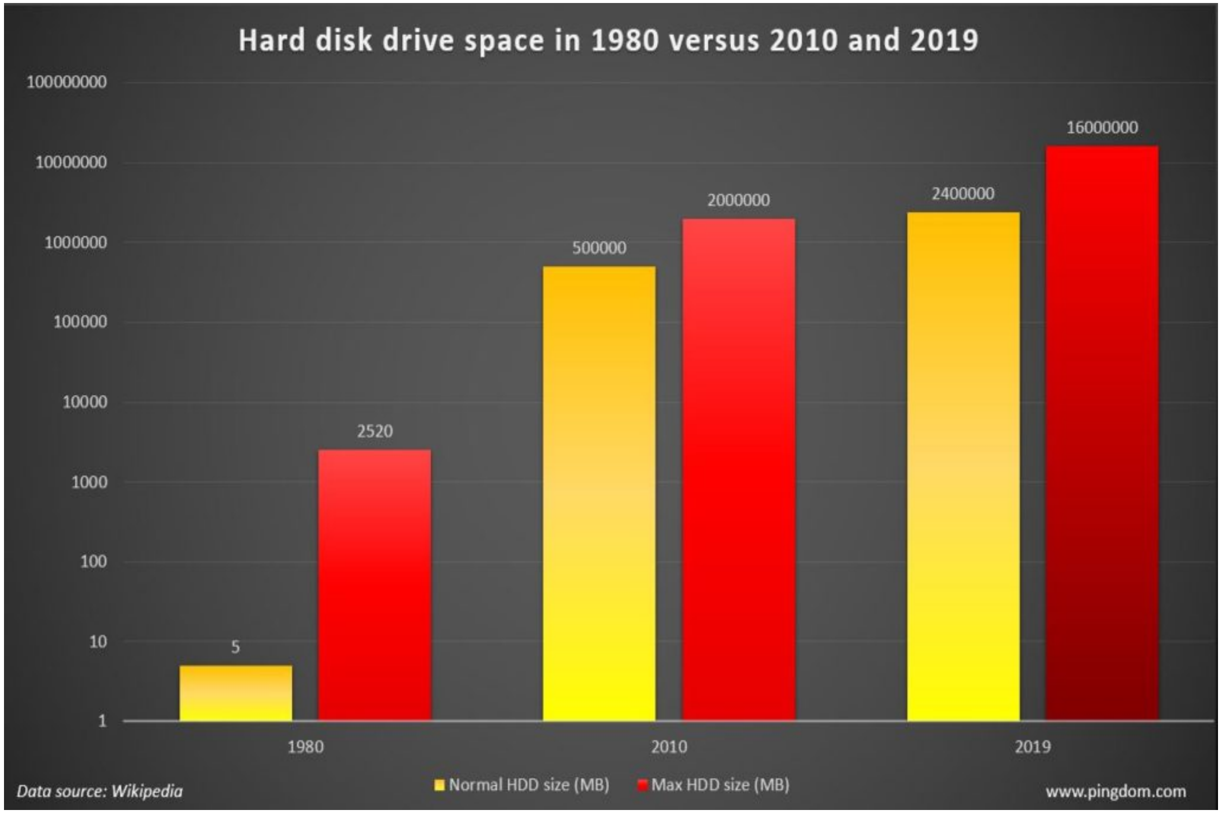
Hard disk drive space in 1980 versus 2010 and 2019
100000000
10000000
1000000
100000
10000
1000
100
10
1
5
2520
500000
2000000
2400000
16000000
1980
2010
2019
Normal HDD size (MB)
Max HDD size (MB)
Data source: Wikipedia
www.pingdom.com

Compression

- Still data is so large – difficult to transport it over long distances
- Lossy and lossless compression
- Some can create a factor of 200 or more compression
- Depends on structure of the data
- Data lives on the disk compressed only uncompressed when processed
- Vo and Korn RFC 3284

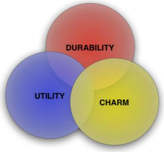

Bias

Fairness

- A Survey on Bias and Fairness in Machine Learning by A Survey on Bias and Fairness in Machine Learning N. Mehrabi, F. Morstatter, N. Saxena,K. Lerman, and A. Galstyan, - USC
- Growing awareness – also part of all research and tech: Rosenthal and Hawthorne effects

Types of Bias -1

- Historical Bias – image search for Fortune 500 CEOs (male)
- Representation Bias - how we define sample (no geo diversity)
- Measurement Bias - prior arrests not convictions was basis
- Evaluation Bias – skewed samples
- Aggregation Bias – covid vaccine tested only on males, 20-30 age group, …
- Population Bias – when statistics, demographics, representatives and user characteristics are different in the user population represented in the dataset or platform from the original target population
- Simpson’s Paradox subgroups may have different underlying trends
- Longitudinal data fallacy treat cross-sectional data as if it is longitudinal, rather than a sampling

Types of Bias - 2

- Sampling bias – non random sampling of subgroups
- Behavioral bias arises from different user behavior across platforms, contexts or different data sets
- Content production bias arises from structural, lexical, semantic and syntactic differences in the content generated by users.
- Linking bias arises when network attributes obtained from user connections, activities or interactions differ and misrepresent the true behavior of the users
- Temporal bias arises in difference from populations or behaviors over time
- Popularity bias – items that are more popular tend to be exposed more or manipulated by social bots or fake reviews

Types of Bias - 3

- Algorithmic bias is added by the algorithm
- User interaction bias user interface can influence behavior includes presentation bias how items are presented and ranking bias (primacy and recency effects)
- Social bias happens when other people's actions or content coming from them can affect our judgment
- Emergent bias users differ and population and cultures change over time
- Self selection bias subjects of the research select themselves, not random
- Omitted variable bias occurs when one or more important variables are left out of the model

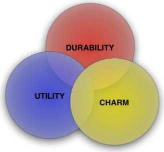

Types of Bias - 4

- Cause – effect bias as a fallacy that correlation implies causality
- Observer bias researchers subconsciously project their bias
- Funding bias results reported (*or nuanced*) to satisfy the funding agency

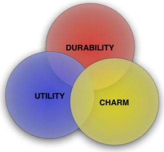

Next Time

- Safety in AI
  - Robustness in AI
  - Software Engineering
  - Systems Engineering
- Transparency
- Augmentation
- True AI

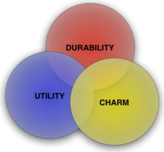

Three Talks

- ~~Digital Readiness: AI/ML, The thinking system quest.~~
  - ~~Artificial Intelligence and Machine Learning (AI/ML) have had a fascinating evolution from 1950 to the present. This talk sketches the main themes of AI and machine learning, tracing the evolution of the field since its beginning in the 1950s and explaining some of its main concepts. These eras are characterized as "from knowledge is power" to "data is king".~~
- ~~Digital Readiness: AI/ML, Finding a doing machine.~~ –
  - ~~In the last decade Machine Learning had a remarkable success record. We will review reasons for that success, review the technology, examine areas of need and explore what happened to the rest of AI, GOFAI (Good Old Fashion AI).~~
- Digital Readiness: AI/ML, Common Sense prevails?
  - Will there be another AI Winter? We will explore some clues to where the current AI/ML may reunite with GOFAI (Good Old Fashioned AI) and hopefully expand the utility of both. This will include extrapolating on the necessary melding of AI with engineering, particularly systems engineering.

Last Week's Questions