Artificial Intelligence
a Rupture Technology for Innovation

James L. Crowley
Professor Grenoble INP
Grenoble Informatics Laboratory (LIG)
INRIA Grenoble Rhone-Alpes
Univ. Grenoble Alpes

Artificial Intelligence
a Rupture Technology for Innovation

Outline:

- Turing's Definition for Intelligence
- History of Paradigms for AI
- Barriers and Enabling Technologies
- Potential for Innovations
- Current open problems

Artificial Intelligence (AI)

The Turing Test: An imitation game

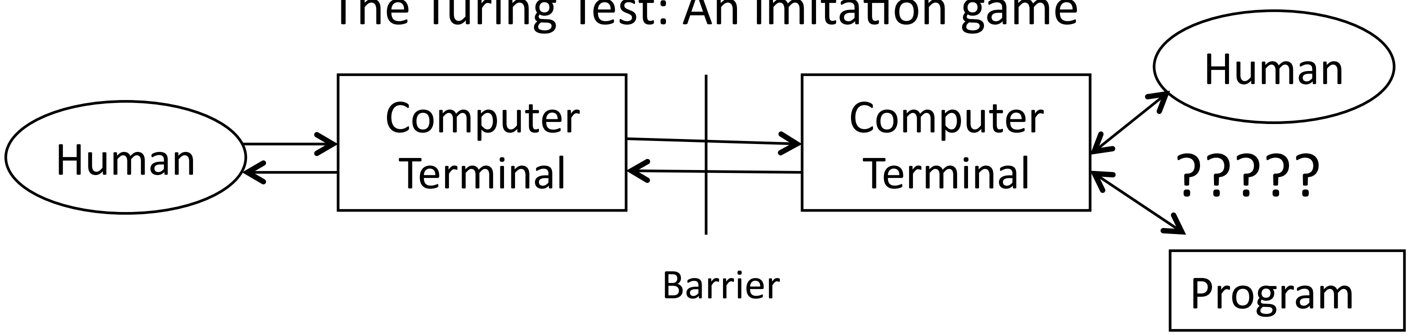

Artificial Intelligence according to Turing:

Human-level performance at social interaction.

Artificial Intelligence (AI)

The Turing Test: An imitation game

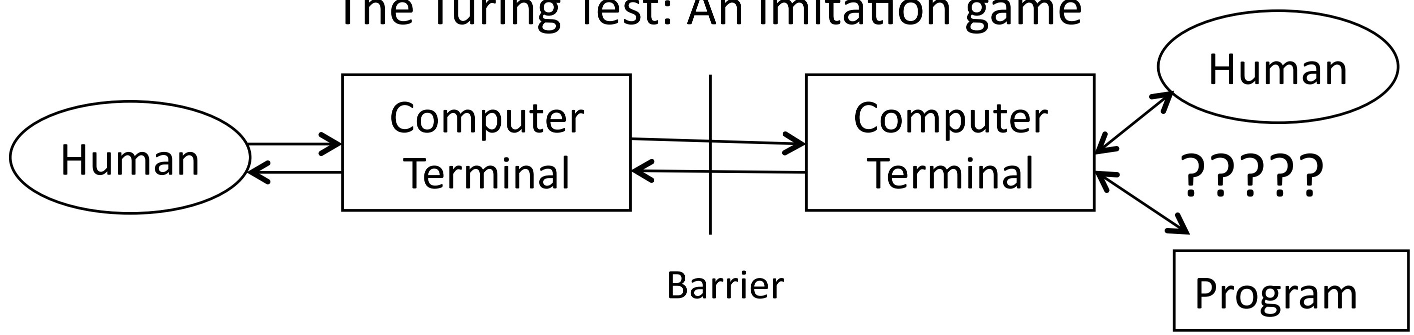

Modern technologies allow us to extend Turing's definition to interaction with people, systems, or the physical world.

Intelligence: Human-level performance at <u>interaction</u>.
(using perception, action, communication or cognition)

AI as a Modern Scientific Discipline

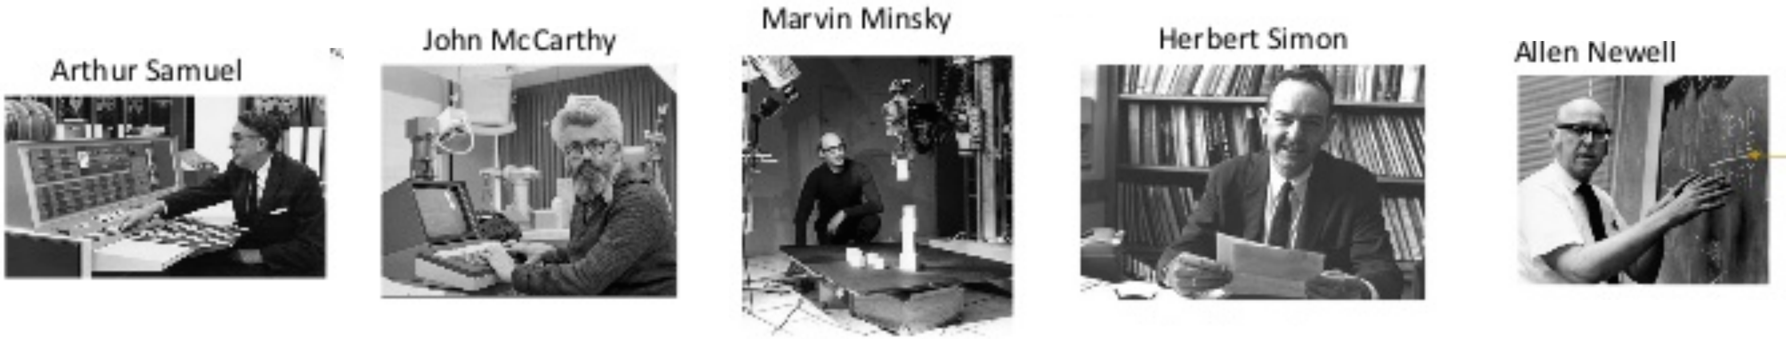

AI Pioneers at the Dartmouth Symposium (1956)

The modern scientific domain emerged in the 1960s as a convergence of Cognitive Science, Logic, Planning, Pattern Recognition, Image Processing and other fields, driven by the emergence of Computer Science.

Evolution of Artificial Intelligence
From Pattern Recognition to Deep Learning

Dominant Paradigms for Artificial Intelligence:

- Pre-1960: Automata and Pattern Recognition
- 1960-1985: Planning, Problem Solving
- 1975-1990: Expert Systems
- 1985-2000: Logic Programming
- 1995-2010: Bayesian Methods, Semantic Web

Three Fundamental Barriers to AI:

(1) Insufficient Labeled Data for Learning.

(2) Insufficient Computing Power.

(3) Prohibitive Cost of Encoding Domain Knowledge.

Expert System Design Process (1980)

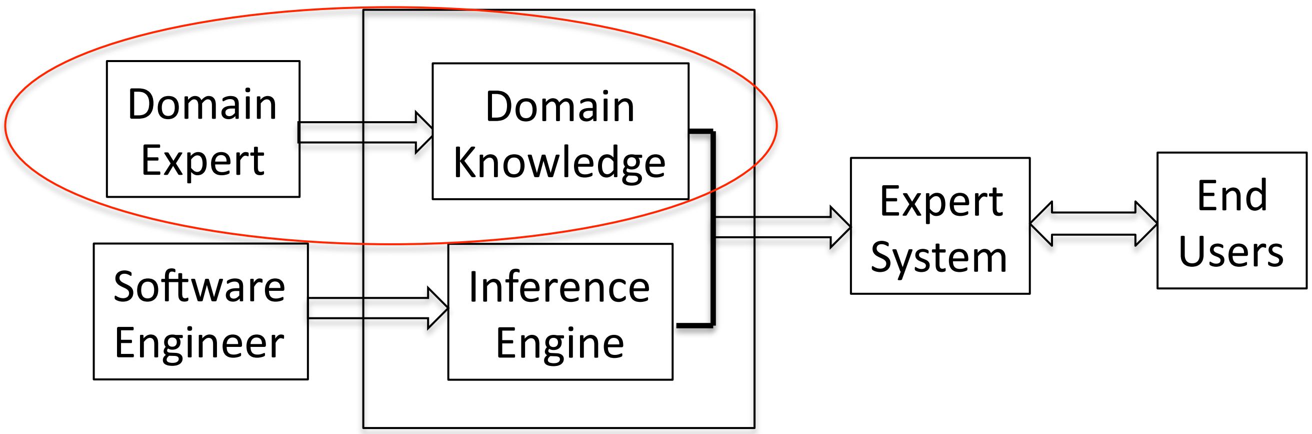

Example: MYCIN – Antibiotic Therapy Advisor (Feigenbaum et al 1980).
Domain expert worked with Software Engineer to build system.

Fundamental Problem:
Prohibitive cost of generating Domain Knowledge.

Three fundamental Barriers to AI

(1) Insufficient training data
(2) Insufficient computing power
(3) Prohibitive cost of encoding domain knowledge

AI Enabling Technologies

Overcoming the three fundamental Barriers:

(1) Insufficient training data

$\Rightarrow$ Planetary scale data from the internet and the WWW

(2) Insufficient computing power

$\Rightarrow$ Moore's Law, GPUs, Cloud Computing

(3) Prohibitive cost of encoding knowledge

$\Rightarrow$ Deep Learning

Deep Learning:
A major scientific breakthrough

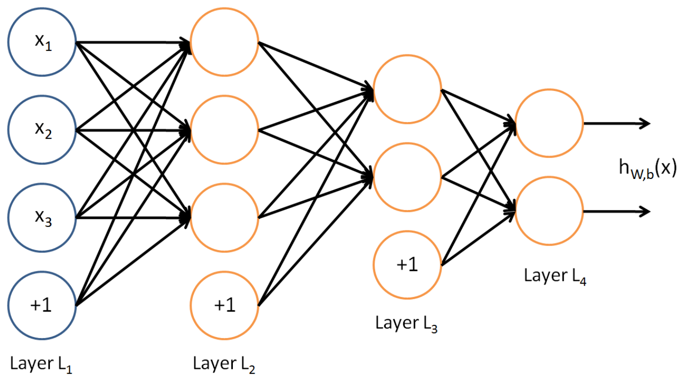

Since 2012, Deep Learning has been found to provide reliable solutions to longstanding problems in perception and problem solving.

AI is now a rupture technology driven by the convergence of Deep Learning, Super Computing, and Planetary-Scale Data.

The impact on human society is expected to be on a scale comparable to electricity or the printing press.

Rosenblatt's Perceptron (1958)

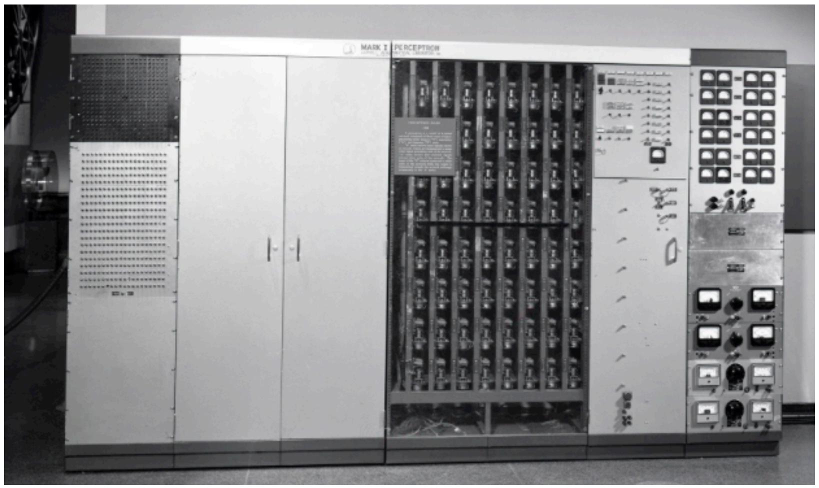

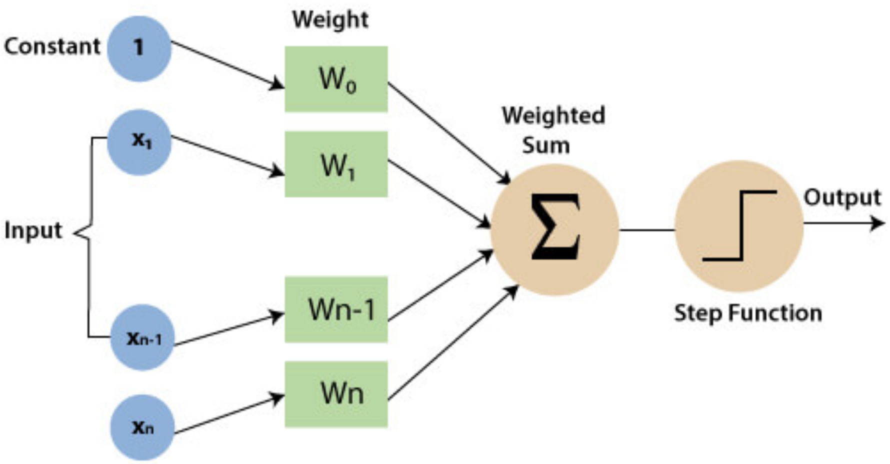

Perceptron: Learning algorithm for a linear decision surface.

Problems: (1) Could only classify patterns
(2) Required labeled training data
(3) Required linearly separable properties for classes.

If the training data was not linearly separable, the algorithm would not terminate

Artificial Neural Networks (1975-1990)
Multi-layer Perceptrons with Back Propagation Learning

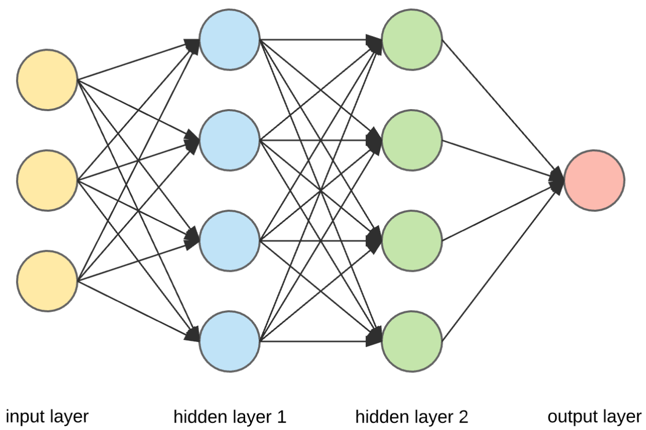

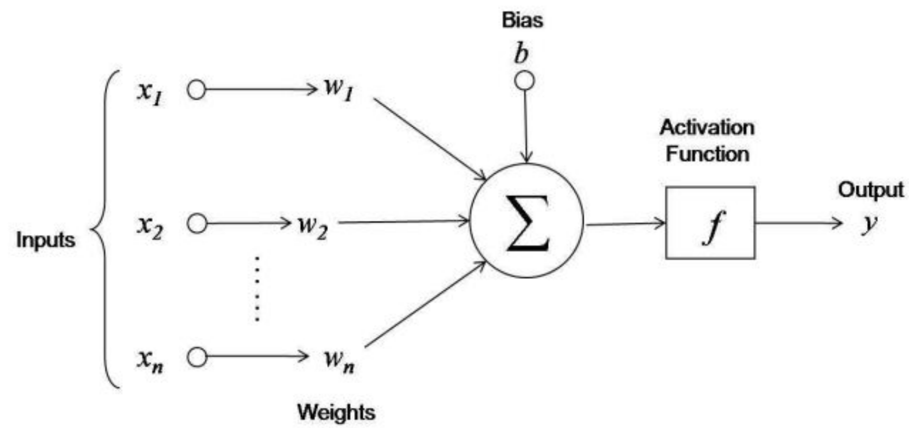

Artificial Neural Networks (1975-1990) – Two innovations

1) Multi-layer perceptrons with soft decision surface
2) Learning with Back-Propagation (Gradient Descent).

<u>Advantages</u>: Simple solution to hard AI recognition problems

<u>Problems</u>:

1) Black Box (unexplainable, unpredictable behavior)
2) Difficult to reproduce
3) Cost of Learning (data and computation) grew with number of Layers

Le Net5 - 1994

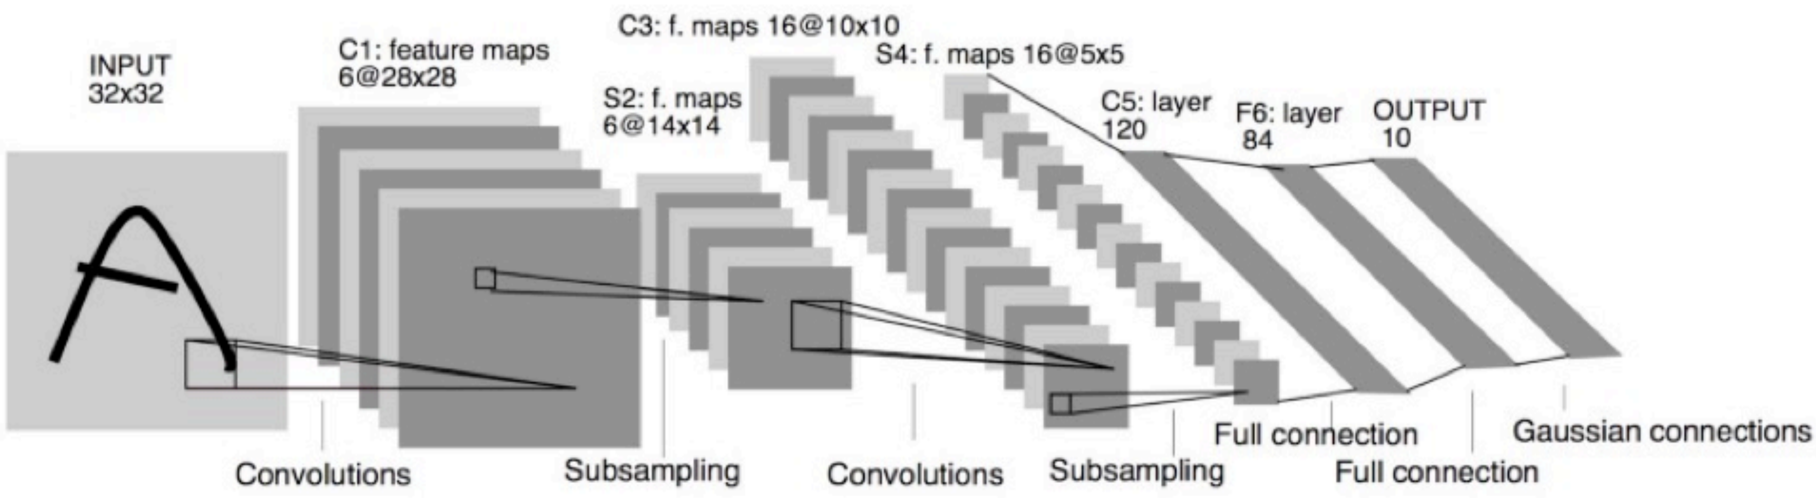

7-level convolutional network by Yann LeCun in 1998.
State of the art for recognizing hand-written numbers on checks.

Ignored by the Machine Learning and Computer Vision communities until around 2010.

AlexNet 2012

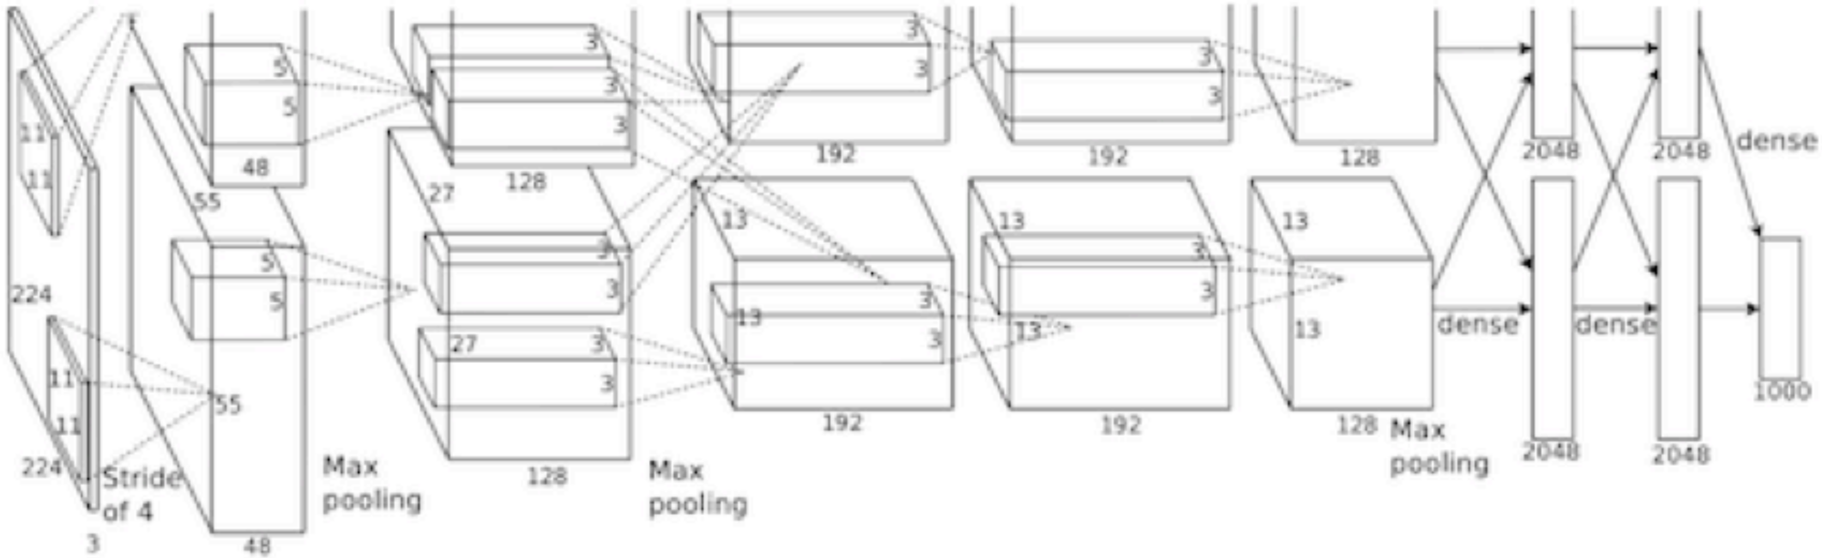

Created by Alex Krizhevsky and Geoff Hinton
Won the ImageNet Large Scale Visual Recognition Challenge in 2012 with an error of 15%
Triggered a Paradigm shift for Computer Vision, Speech Recognition and Machine Learning.

VGG 2015

Karen Simonyan and Andrew Zisserman, Oxford **Visual Geometry Group**
Published at ICLR 2015, Available in Github, Tensorflow, Keras
Simple and effective workhorse for Transfer Learning

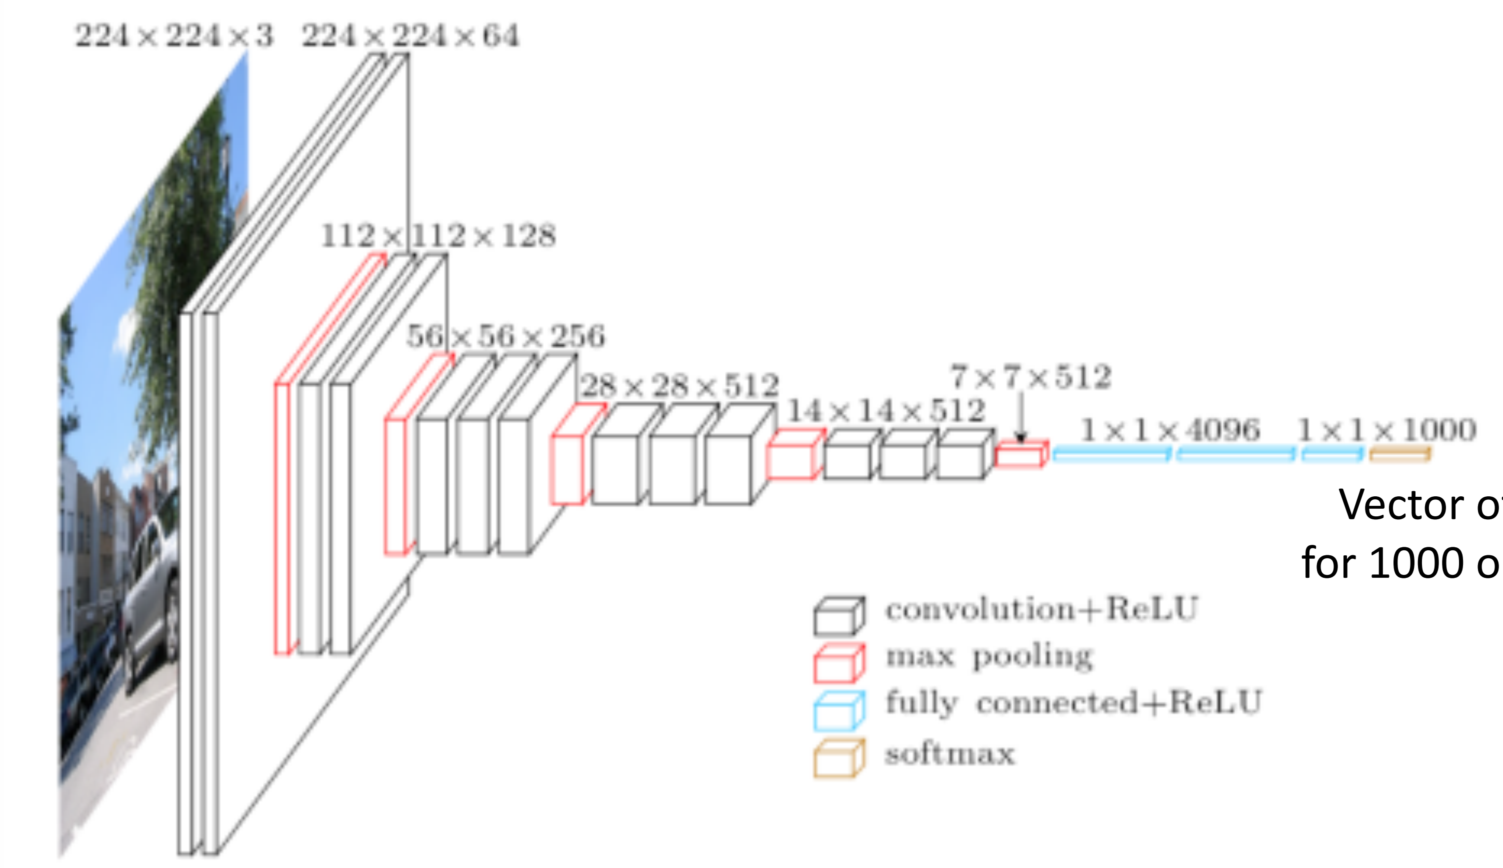

Output:
Vector of Likelihoods
for 1000 object classes

Artificial Intelligence

Enabling Technological Rupture throughout Society

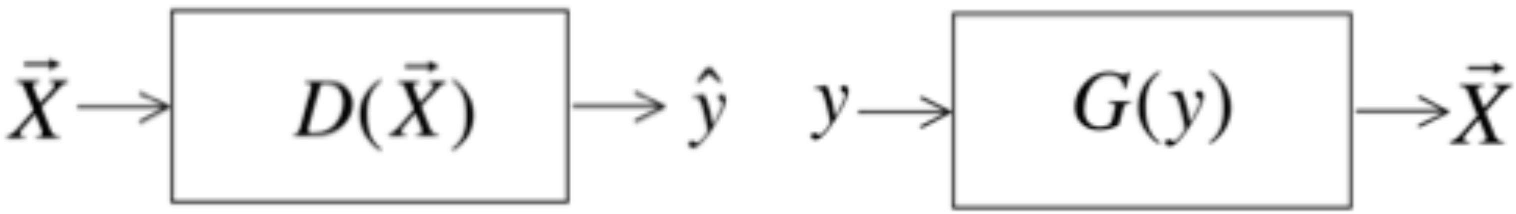

Discriminative Networks:
Does data X contain class y?

Generative Networks:
Generate pattern X for class y

Deep learning was originally invented for Recognition.

The same technology can be used for generation.

Examples:
- Natural sounding speech
- Natural Language
- Synthetic images
- Robot animation
- Realistic talking heads (Deep Fake!)

Hourglass model

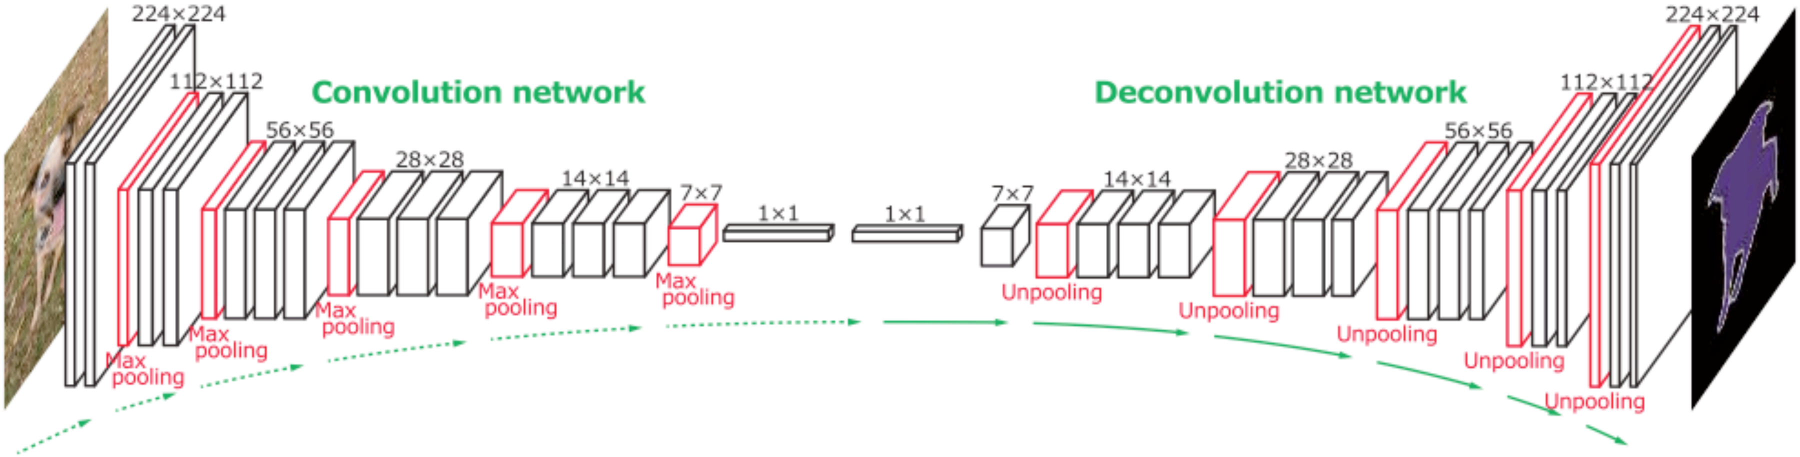

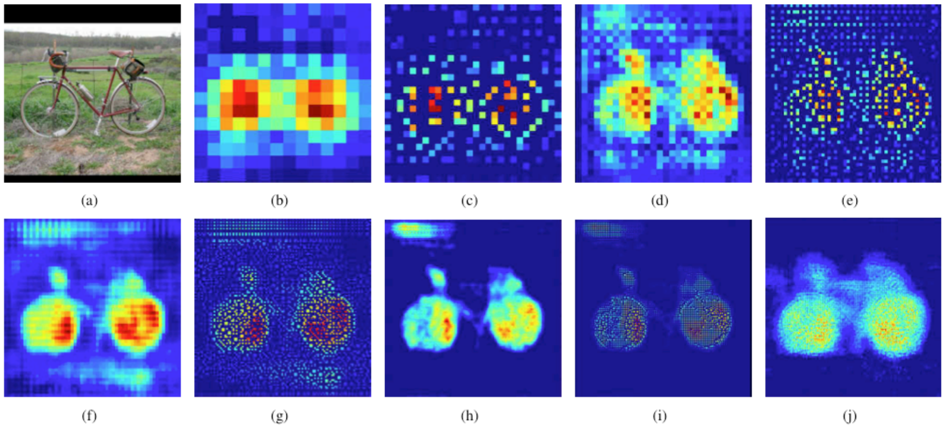

Current Research Trends in AI

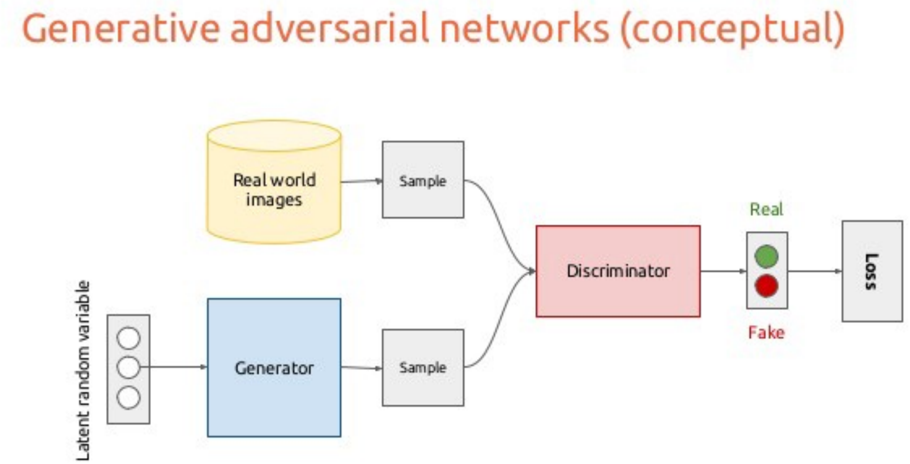

Hot Topics at the International Joint Conference on AI (IJCAI) in July 2018.

1. Generative Adversarial Networks: Unsupervised learning by a competition by Generative and a Discriminative network (ex DeepFake, Realistic Speech synthesis, photo Realistic images)
2. Deep Reinforcement Learning: Unsupervised learning using interaction with a domain model to train a generative network. (ex Games, Robotics)
3. Convergence of Deep Learning and Cognitive Systems: Combination of Deep Networks for Perception and Action with Cognitive Computing for conceptual reasoning and natural language expression. (ex: IBM Watson)

Word2Vec – Linguistic Encoding

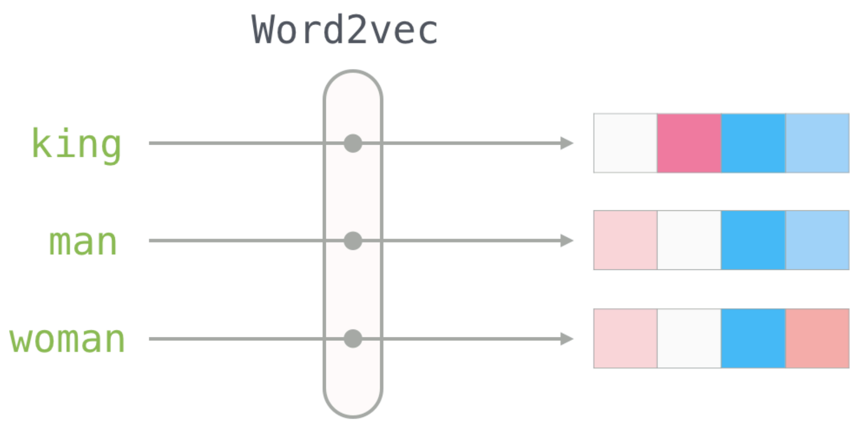

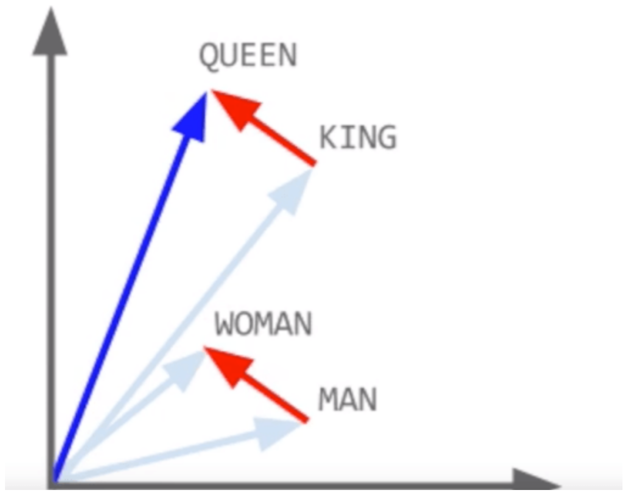

- Any document is a context, represented by word frequencies.
- Word2Vec projects each word onto a Context vector.
- Context vectors can be compared with cosine angles
- Context vectors can be used with Deep Learning!

Applications: Chat-bots, Data mining on the internet

Chatbot Example: Weabo

Example: Cleverbot – online chatbot (http://www.cleverbot.com)

User: Good afternoon.

Cleverbot: Good afternoon.

User: Tell me about yourself.

Cleverbot: I'm planning to do something big.

User: Good for you.
What kind of big thing do you plan?

Cleverbot: I'm taking that as a compliment.

User: Fine. But what are you up to?

Cleverbot: Nothing today.

User: Tell me how you work.

Cleverbot: Huh?

Chatbots

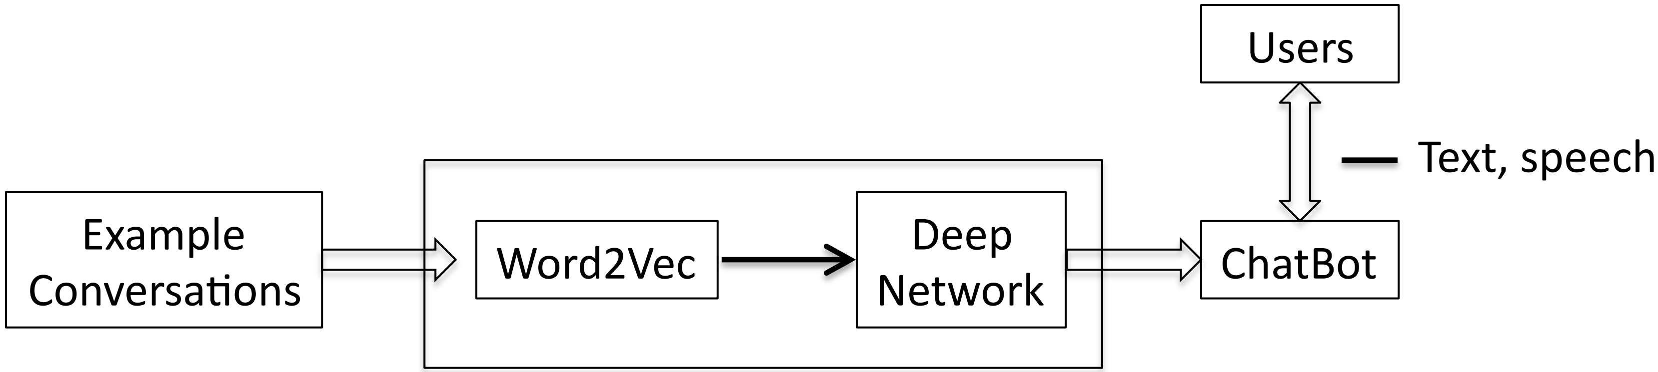

Problem: Chatbots do not "understand" the conversation
Chatbots operate by Stimulus - Response (like a parrot).
(Not really a solution to the Turing Test – but close).

Current research: Use word2Vec to generate logic forms, use logic programming to reason, convert results with Generative net.

Cognitive Computing

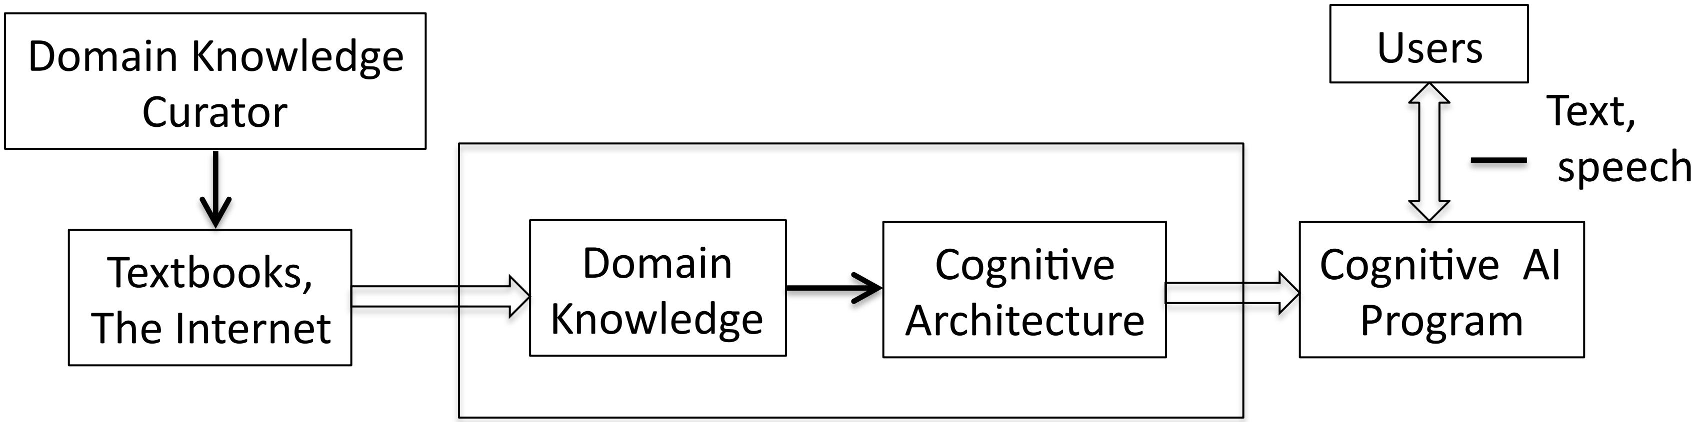

Cognitive Computing encodes knowledge from any written source (textbooks, literature, the internet) to generate a domain expert advisor program (an expert system!)

Example domains: Medical, Legal, Financial…

Cognitive Computing Example:
IBM Watson

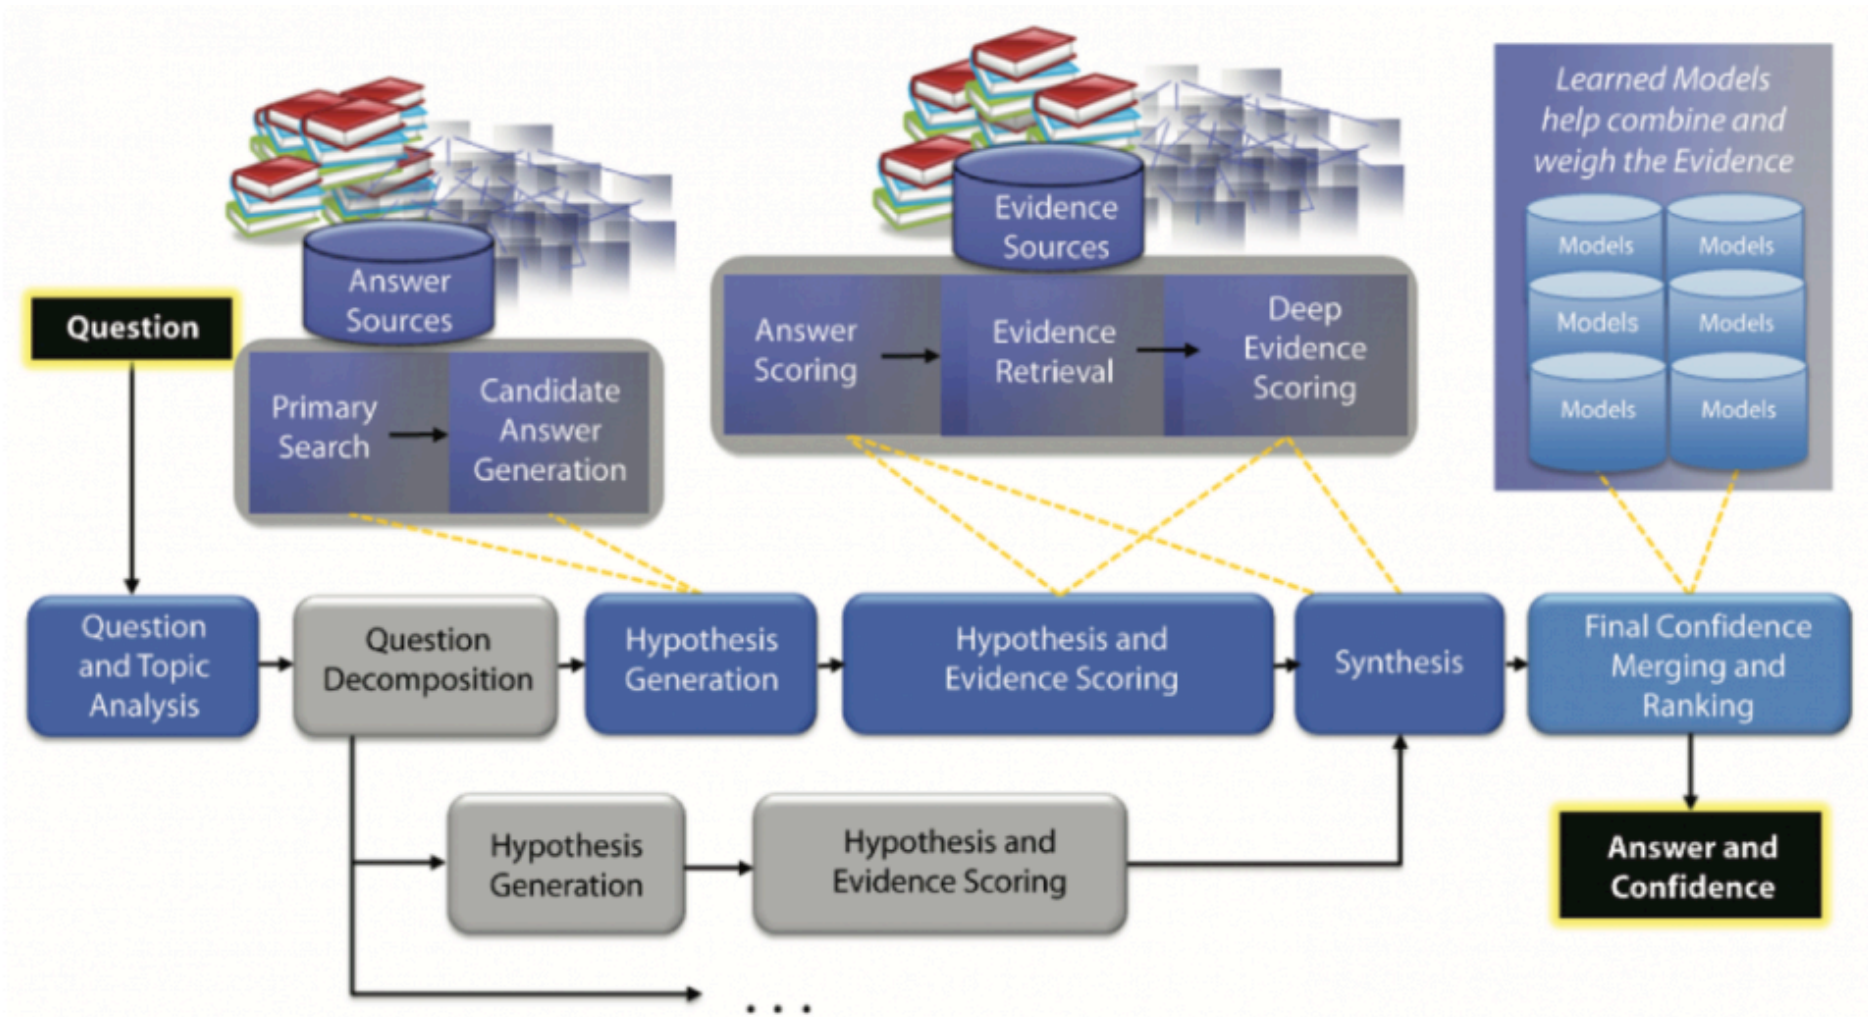

Source: IBM

Pepper – Animated by Watson

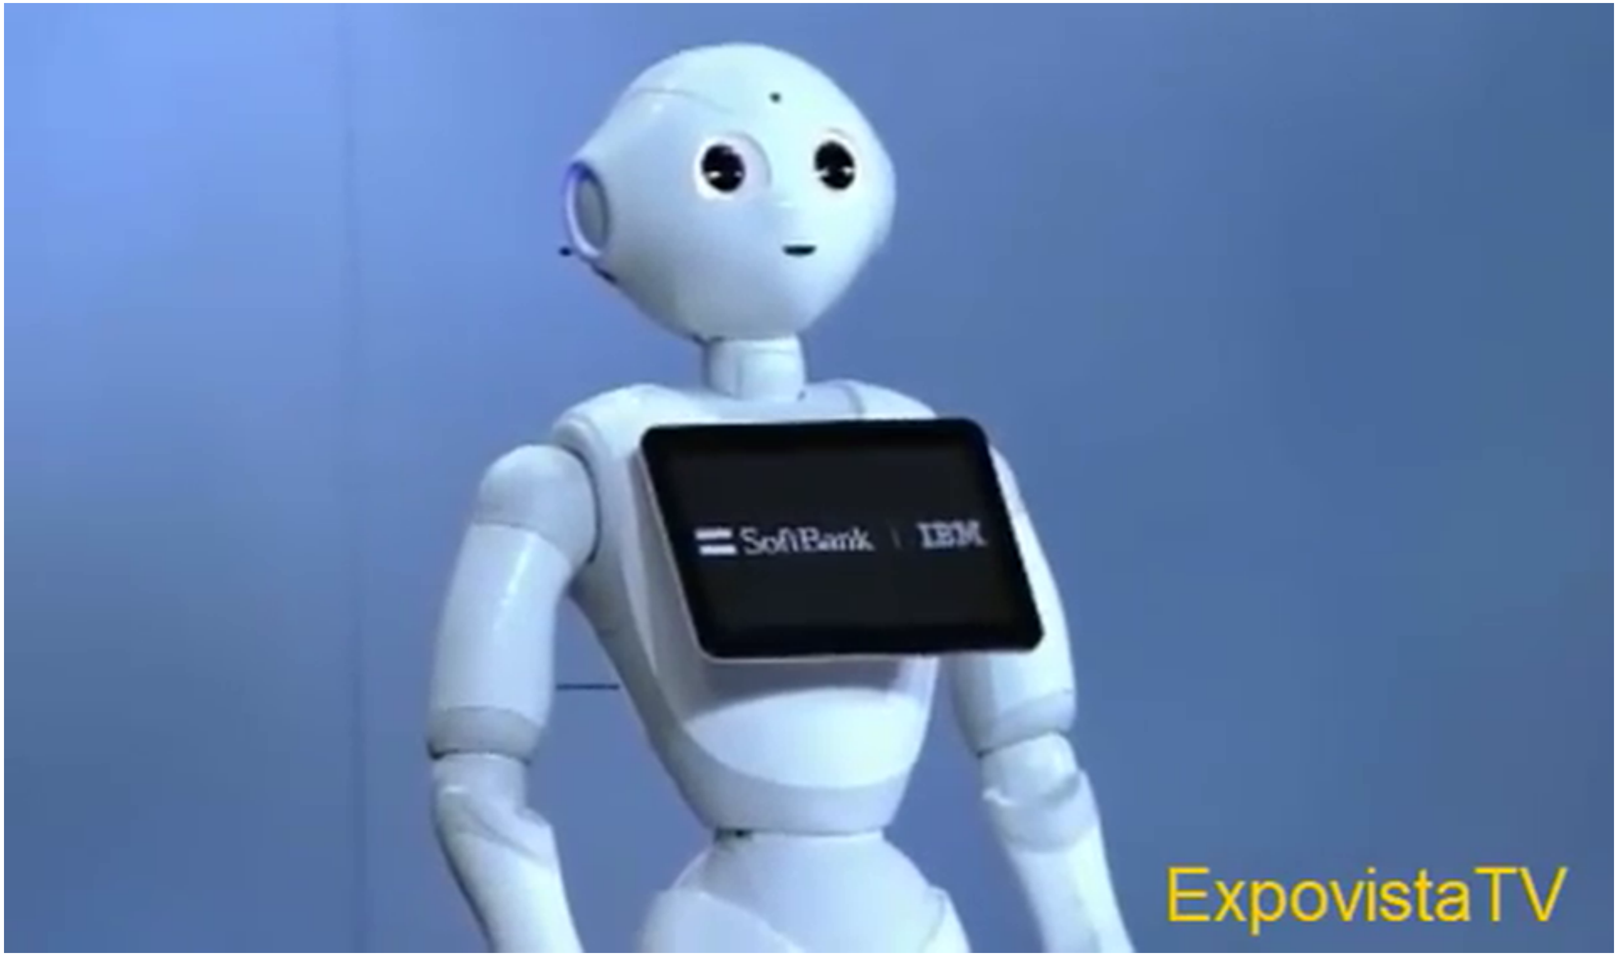

Pepper (Softbank Robotics) and Watson at CES 2016

AI as a Service on the Cloud

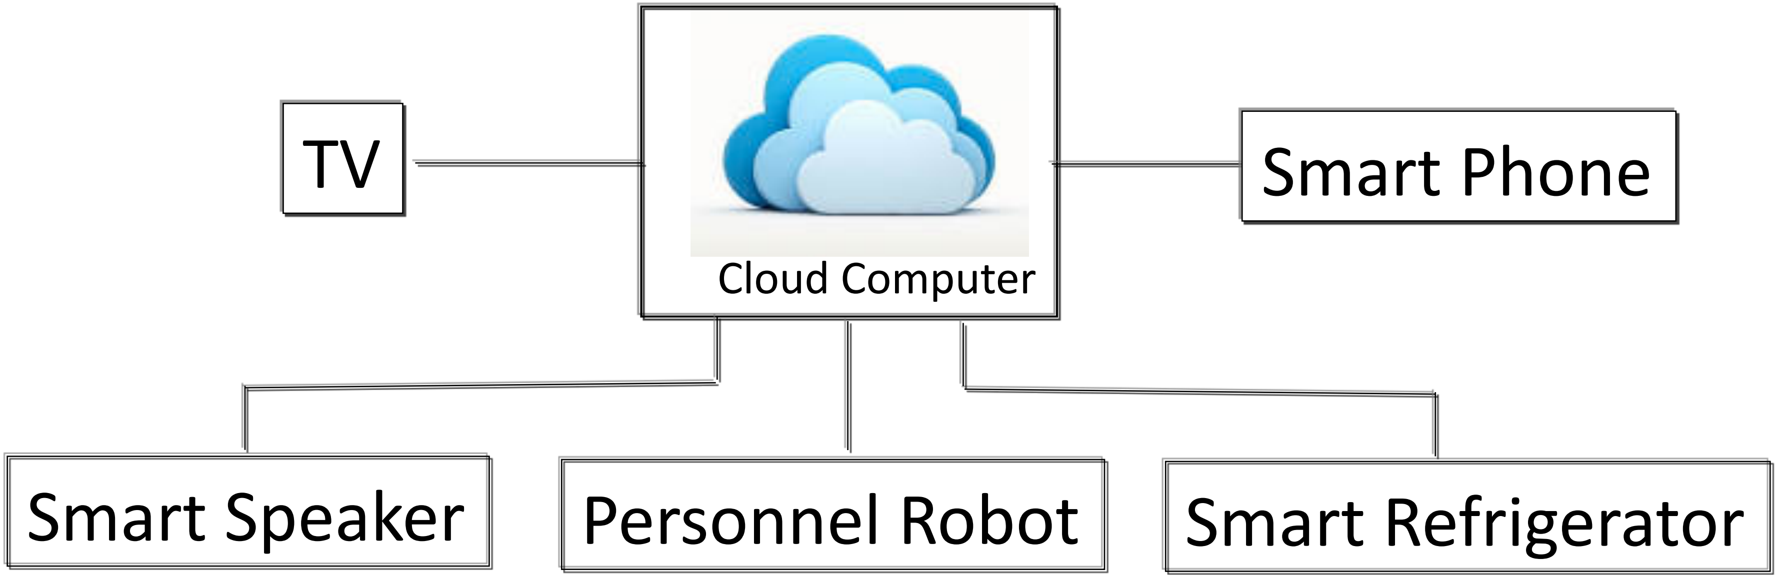

Software Engineering Process

Classic Software Engineering:

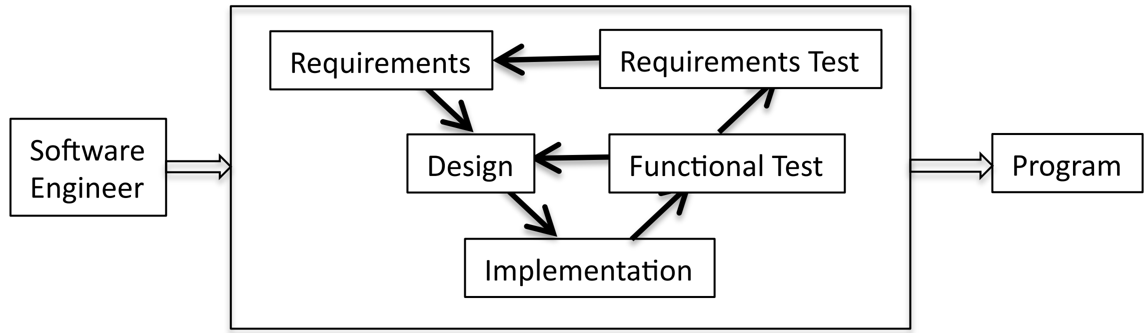

Machine Learning:

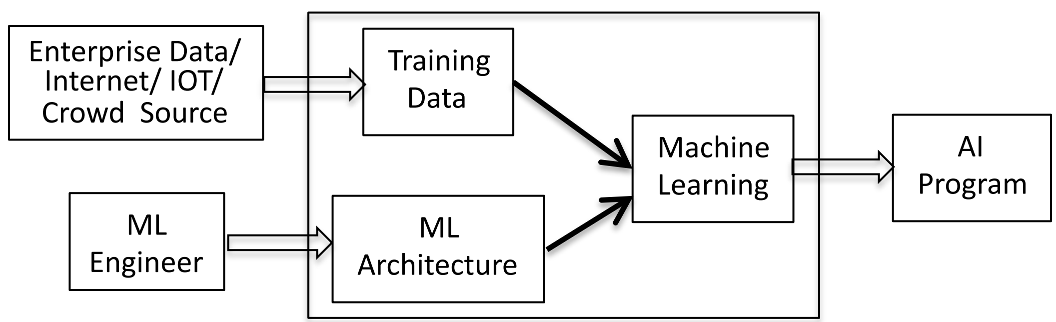

With Machine Learning, the data becomes the code.

<u>Problems</u>: Verification: How can you Certify correct behavior?
Explanation: How can you explain output from network?

Potential Innovations from AI

How can we use Machine Learning and AI for technological rupture in various domains?

What are the potential domains?

AI is the fire. Data is the fuel.

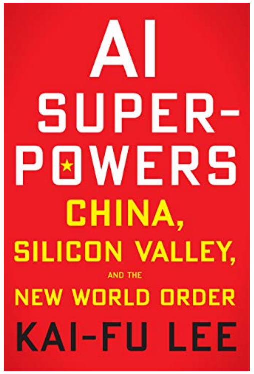

To predict AI innovation, look for the data (Kai Fu Lee).

Five Waves of rupture from innovation through AI

1. Internet AI and "AI as a Service" (2015 – 2025) (US and China)
2. Enterprise AI (2015 – 2025) (US leads)
3. Mobile AI using Smart Phones (2015 – 2025) (China leads)
4. Ubiquitous Perception and Interaction (2020 – 2030)
5. Autonomous AI Systems. (2025 – 2035)

USA, China, and Europe are unevenly positioned to profit or suffer from each wave.

AI4EU Pilot Projects

AI4Citizens	Personal Assistant for Public Services
AAI4Robotics	Intelligent Performance Analytics for Industrial Robots
AI4Industry	AI-Driven Digital Companion for Production Facility
AI4Healthcare	Improve quality of medical images reports
AI4Media	AI-Based 3D-Generation of Animated Video
AI4Agriculture	Crop quality assessment through computer vision
AI4IOT	Air Quality monitoring
AI4Cyber-Scecurity	AI-driven attack learning

Potential Innovations from AI

AI: Human level ability at interaction

Interaction with the Physical World:
=> Robotics, Transportation, Manufacturing, ...
Interaction with People:
=> Commerce, Education, Entertainment, Well Being
Interaction with Internet and other Systems:
=> Virtual Personal Assistant, Smart Buildings...
Interaction with Technology:
=> AI for design, AI for programming...

Interaction with the Physical World

Adult Supervision Required!

For the next decade, AI will enhance human productivity

Transportation

Manufacturing

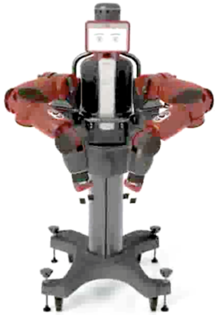

Human Environments

Nest

Roomba

Interaction with the People

Advisors

Propose courses of actions.
Completely obedient. Do not act.
Avoid unwanted distractions.
Example: GPS Navigation system

Affectors

Inspire affection.
Compensate for a loss of social contact.
Examples: Aibo, Nao, Paro, ...

Media

Extend human perception and experience.
Can be interactive or peripheral
Provide a sense of immersion.
Examples: Ambient Orb (Rose 14)

Collaborative Intelligent Systems

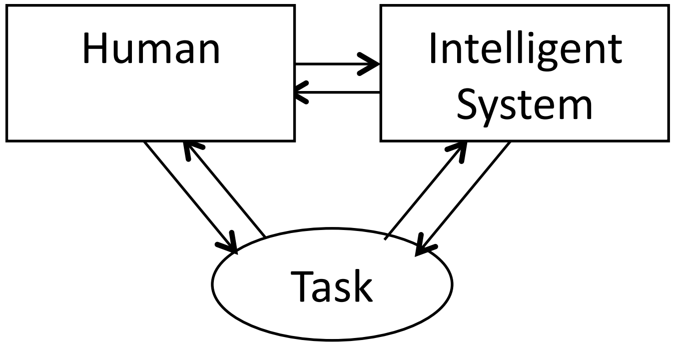

Collaboration is a process where two or more actors (agents) work together in order to achieve some shared goals.

Collaborative Intelligent Systems are intelligent systems that work with humans as partners to achieve a common goal, sharing a mutual understanding of the abilities and respective roles of each other.

Collaborative Intelligent Assistant

An intelligent system that

- Assists an expert at complex, critical tasks.
- Provides assistance and advice on demand
- Monitors actions and attention
- Anticipate problems, Warns of errors
- Can perform delegated tasks

Examples: Virtual Copilot, Radiologist Assistant, Intelligent Tutors, Legal advisor, financial advisor

Principle: The human retains control and responsibility
The system provide assistance and empowerment.

Example: Intelligent Math Tutor

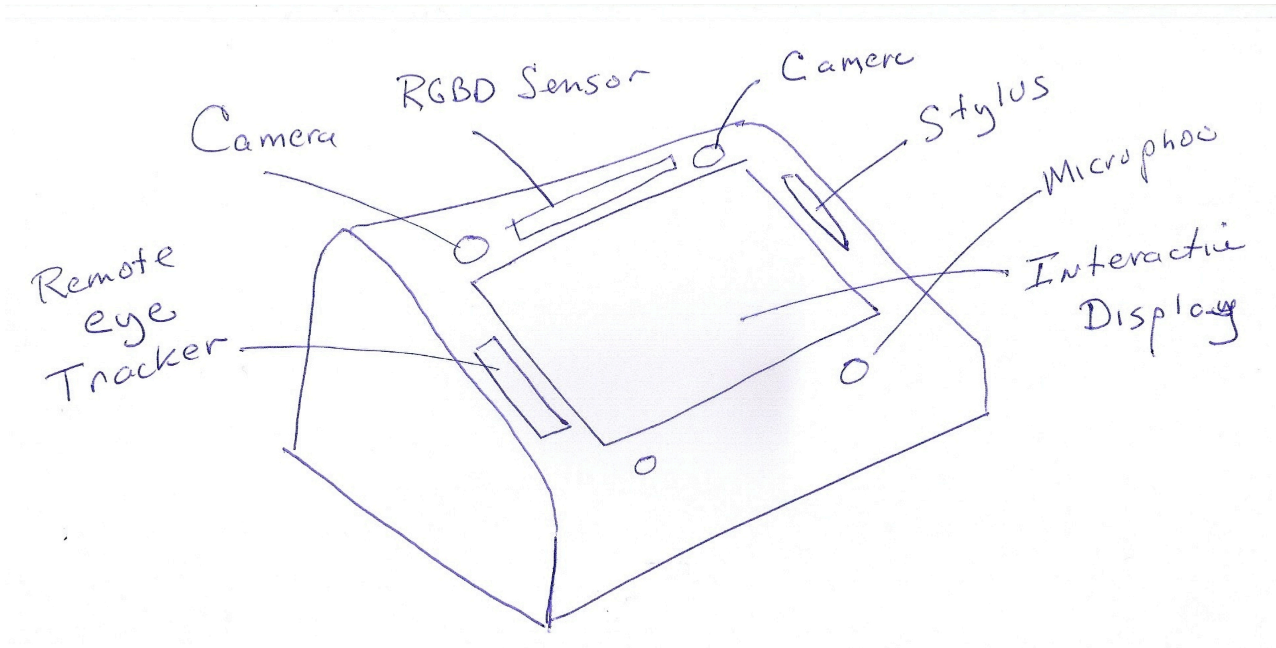

Monitors awareness and emotion from eye-gaze, speech, posture, face expression, cardiac rate, etc.   Guides student through concept acquisition and helps develop problems solving skills.

Collaborative Intelligent Systems

AI technologies
that enhance human abilities and extend human intelligence.

Research Challenge: build systems that:

1. Understand and interact humans as individuals, adapting to individual human abilities and limits, and complying with human social, ethical, and cultural norms
2. Anticipate human needs to provide information and services synergistically work with humans
3. Provide explainability, accountability, and compliance with ethical, legal, social, and cultural values

Explainable AI

Definition: (wikipedia): methods and techniques in … AI such that the results of the solution can be understood by human experts.

There are several different AI explainability problems.

1) Explain the output from a deep network (games, robots, chatbots, etc) to a user.
2) Explain the advice given by an AI system.
3) Explain why one network architecture is better than another.

Explainable AI – Personal View

Explanation: a reason for an action or a deviation from expectation.

Humans provide “narratives” as explanations.

A narrative interprets a sequence of events as a story, placing the events in a context.

Human narratives are often generated after the fact, often simplistic or just plain wrong, but they are credible and thus believed.

Explainable AI will require the ability to generate credible narratives to explain actions, decisions and consequences of systems.

Certifiable AI for Critical Systems

Classic Software Process

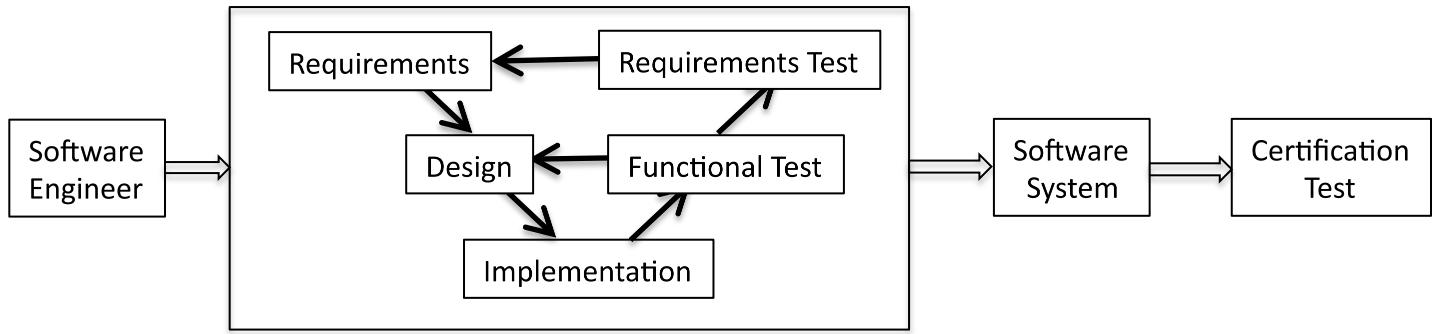

Classic Process: Certification guarantees compliance with specifications.
In case of Failure, the engineer is responsible.

Machine Learning:

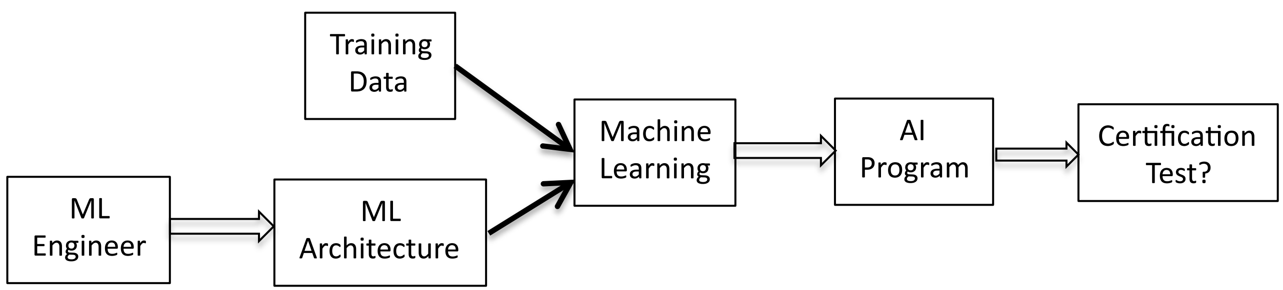

The AI Program Is a black box. What are the specifications? What do we test?
Who is responsible if the system fails?
Can we "qualify" AI systems as we do with programmers drivers, Airline pilots?

Trustworthy AI

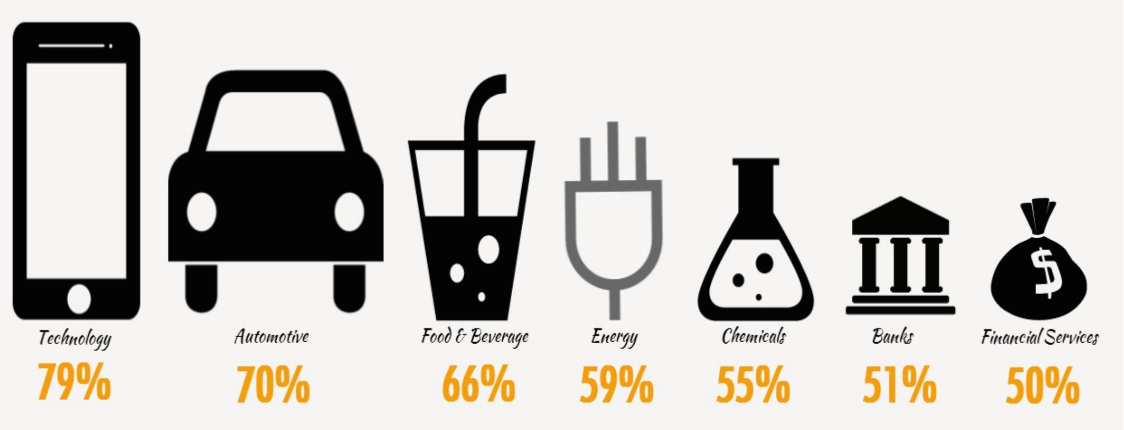

Trust: The ability to inspire confidence that a system is secure, available, and private.

Trusted systems may be insecure or not private
Secure private systems are not always trusted.

Trustworthy AI

According to the guidelines of the AI High-Level Expert Group, Trustworthy AI should be:

(1) lawful - respect all applicable laws and regulations

(2) ethical - respect ethical principles and values

(3) robust - both from a technical perspective while taking into account its social environment

Requirements: Controllable, Safe, Robust, Private, Transparent, and Accountable

Artificial Intelligence
a Rupture Technology for Innovation

Outline:

- Turing's Definition for Intelligence
- History of Paradigms for AI
- Barriers and Enabling Technologies
- Potential for Innovations
- Current open problems

Artificial Intelligence
a Rupture Technology for Innovation

James L. Crowley
Professor Grenoble INP
Grenoble Informatics Laboratory (LIG)
INRIA Grenoble Rhone-Alpes
Univ. Grenoble Alpes