Artificial Intelligence:
A Rupture Technology for Scientific Research?

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

Is Artificial Intelligence A Rupture Technology for Scientific Research?

Artificial Intelligence (AI) has emerged as a Rupture Technology for Industry, Commerce and Social Innovation.

AI is increasingly recognized as a rupture technology for Informatics.

Can AI provide a rupture technology for scientific research?

Is Artificial Intelligence
A Rupture Technology for Scientific Research?

Outline:

- Definition for Artificial Intelligence
- History of Paradigms for AI
- Barriers and Enabling Technologies
- Potential for Innovations

Artificial Intelligence (AI)

The Turing Test: An imitation game

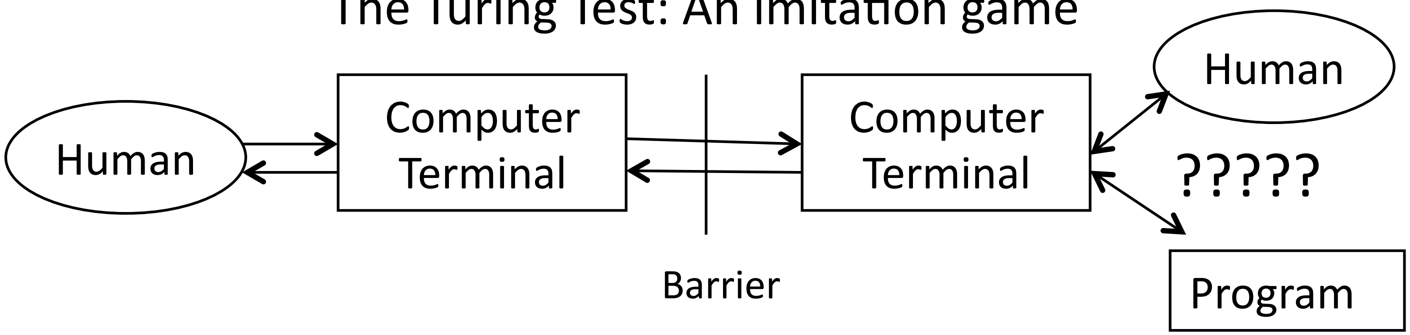

Intelligence according to Turing:

Human-level performance at (text-based) interaction.

The Turing Test: If a human cannot reliably discriminate between a human and a machine using text-based interaction then the machine is said to to be intelligent.

Artificial Intelligence (AI)

The Turing Test: An imitation game

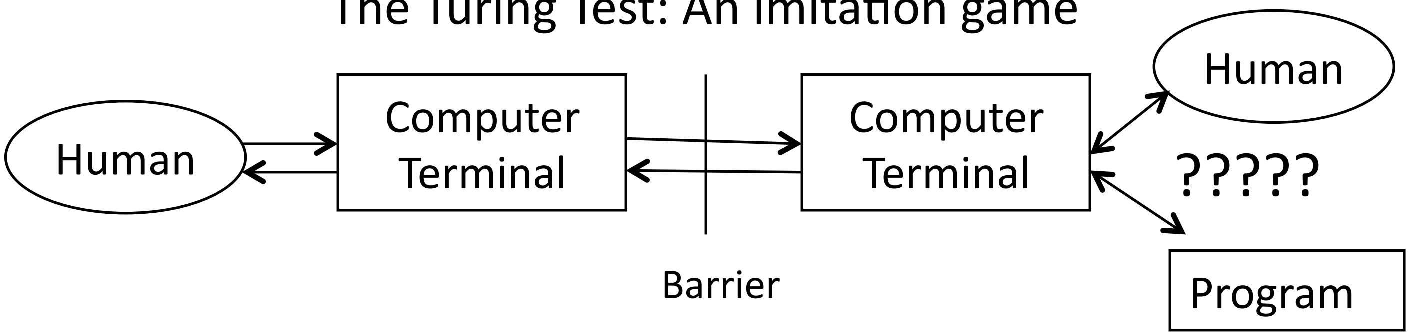

Modern technologies allow us to extend Turing's definition to tasks requiring perception, action, communication or interaction.

Intelligence: Human-level performance at Tasks.
(requiring perception, action, communication or interaction)

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.

Rosenblatt's Perceptron (1958)

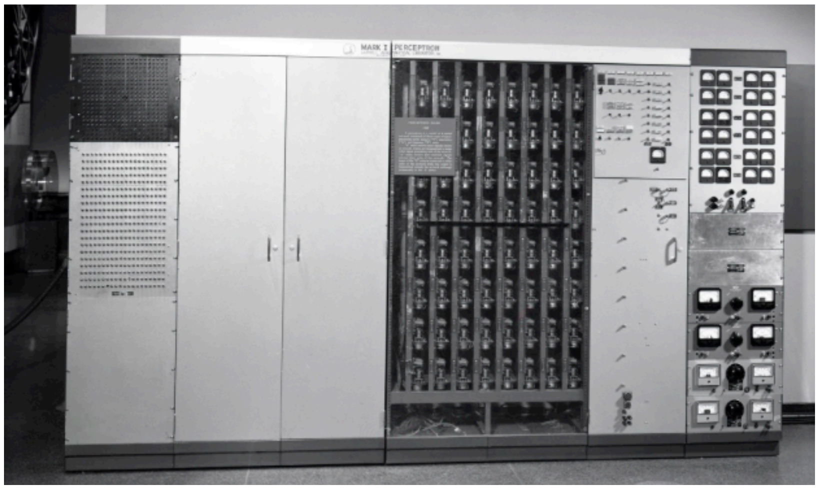

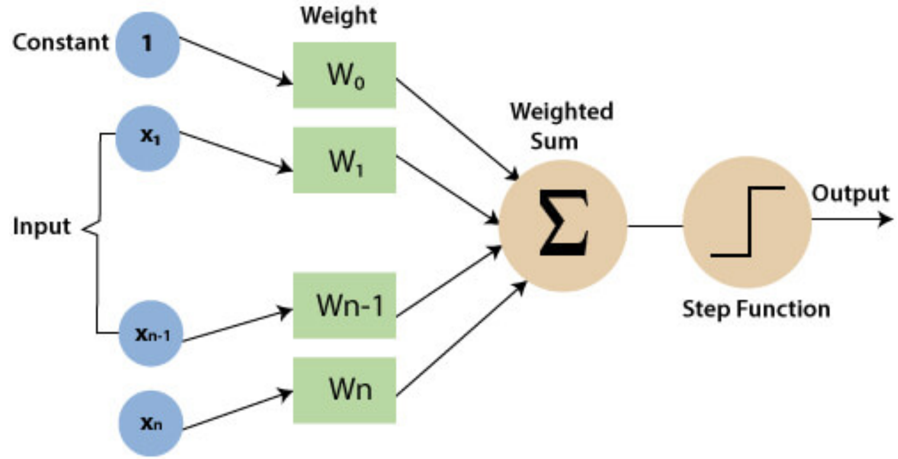

Perceptron: Learning algorithm for a linear decision surface.

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

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)

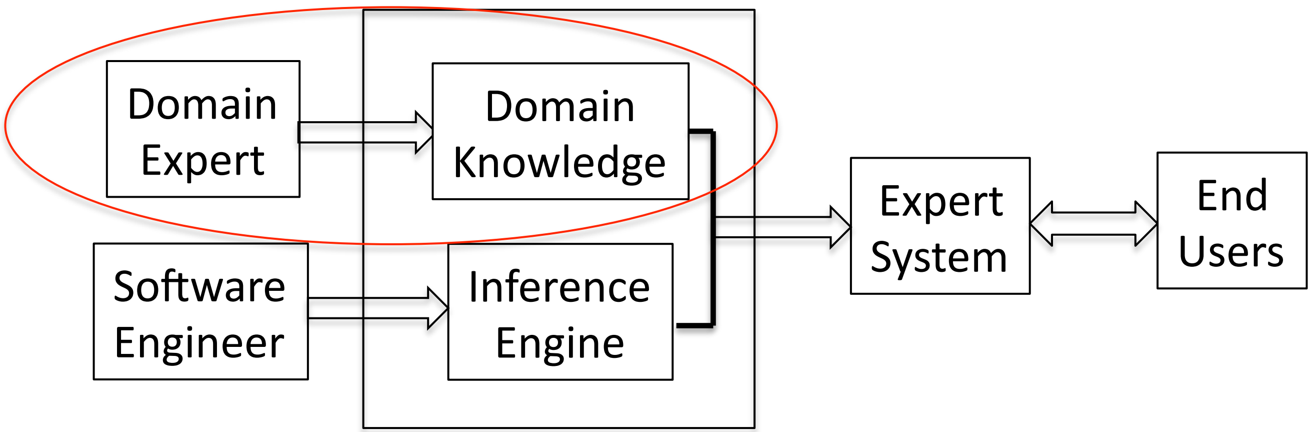

Example: MYCIN – Antibiotic Therapy Advisor (Feigenbaum et al 1980).
Domain expert worked with Software Engineer to build system.

Fundamental Problem:

Prohibitive cost of generating Domain Knowledge.

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

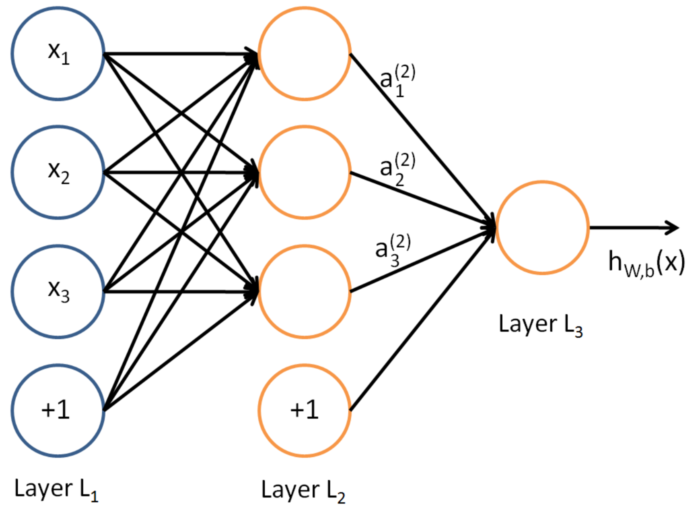

Artificial Neural Networks (1975-1990) – Two innovations

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

Provided a simple alternative to symbolic computing

Feed Forward Network

$$z = \sum_{i=1}^{N} w_i x_i + b \qquad a = f\left(\sum_{i=1}^{N} w_i x_i + b\right)$$

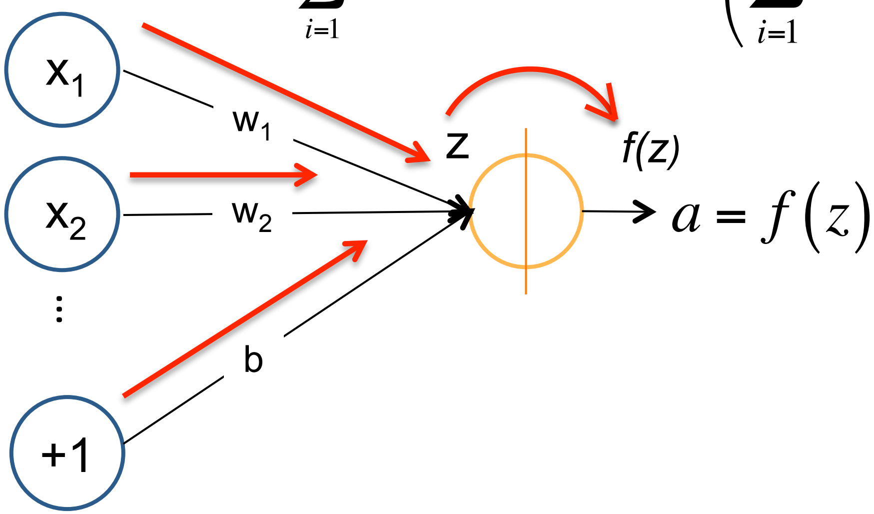

Decision: Sigmoid function

$$f(z) = \sigma(z) = \frac{1}{1+e^{-z}}$$

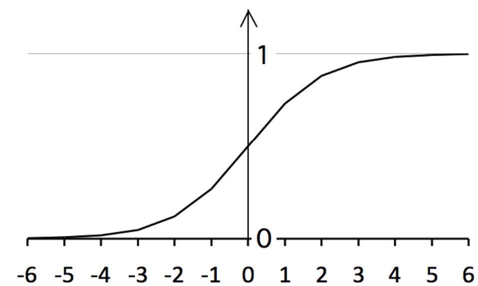

$$\frac{d\sigma(z)}{dz} = \sigma(z)(1-\sigma(z))$$

Important Innovation in the 1970's: Soft decision function.
A soft (differentiable) decision function makes it possible to learn from errors using Gradient Descent.

Back-propagation is Gradient Descent

Training Data: M samples $\{\vec{X}_m\}$ labeled with indicator Variables $\{y_m\}$

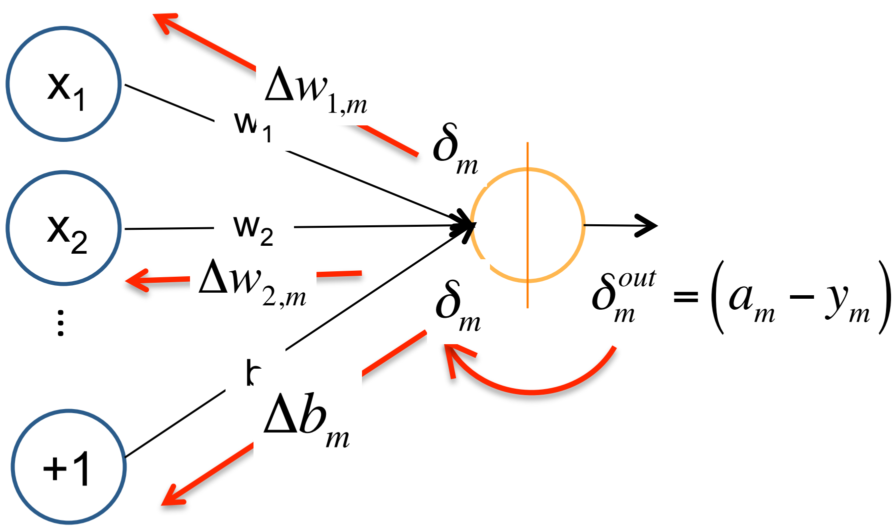

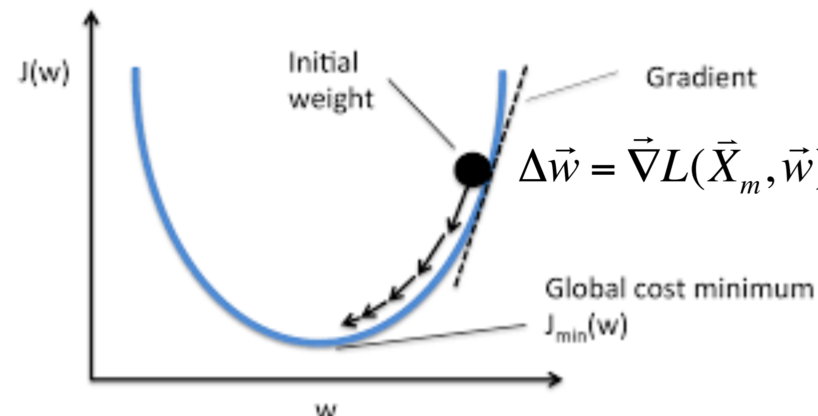

$$L(\vec{X}_m, \vec{w}) = \frac{1}{2M}\sum_{m=1}^{M}(\delta_m^{out})^2$$

Gradient descent: A first-order iterative optimization algorithm for finding the minimum of a function.
Used to determine the best weights and bias

Generalized to Multi-Layer Networks

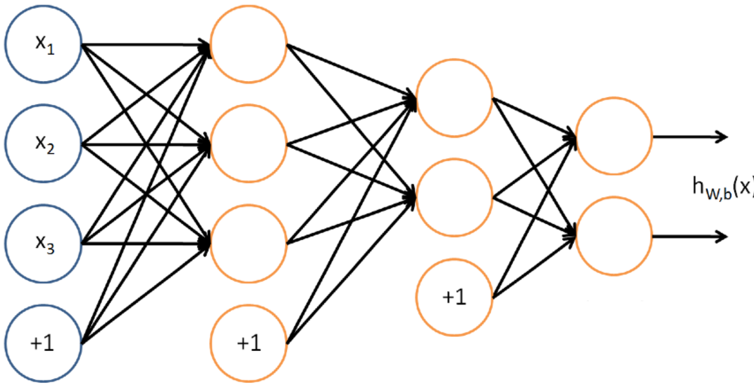

$$\delta_m^{(out)} = \left(a_m^{(L)} - y_m\right)$$

Activation of j^{th} unit of level *(l)*

$$a_j^{(l)} = f\left(\sum_{i=1}^{N^{(l-1)}} w_{ij}^{(l)} a_i^{(l-1)} + b_j^{(l)}\right)$$

Back-prop error for unit *j* of layer *(l)*

$$\delta_{j,m}^{(l)} = \frac{\partial f(z_j^{(l)})}{\partial z_j^{(l)}} \sum_{k=1}^{N^{(l+1)}} w_{jk}^{(l+1)} \delta_{k,m}^{(l+1)}$$

Correction of weight i of unit *j* at layer *(l)*

$$\Delta w_{ij,m}^{(l)} = a_i^{(l-1)} \delta_{j,m}^{(l)}$$

Correction for bias of unity *j* at layer *(l)*

$$\Delta b_{j,m}^{(l)} = \delta_{j,m}^{(l)}$$

Generalized to Multi-Layer Networks

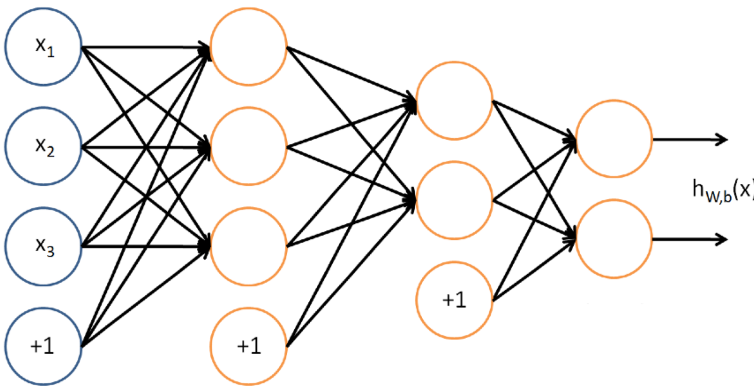

$$\vec{\delta}_m^{(out)} = \left(\vec{a}_m^{(3)} - \vec{y}_m\right)$$

Activation at Level 3

$$a_{j,m}^{(3)} = f\left(\sum_{i=1}^{N^{(2)}} w_{ij}^{(3)} f\left(\sum_{i=1}^{N^{(1)}} w_{ij}^{(2)} f\left(\sum_{i=1}^{D} w_{ij}^{(1)} \vec{X}_{i,m} + b_j^{(1)}\right) + b_j^{(2)}\right) + b_j^{(3)}\right)$$

Back-prop error at Level 1

$$\delta_{j,m}^{(1)} = \frac{\partial f(z_j^{(2)})}{\partial z_j^{(2)}} \sum_{k=1}^{N^{(2)}} w_{jk}^{(2)} \left(\frac{\partial f(z_j^{(3)})}{\partial z_j^{(3)}} \sum_{k=1}^{N^{(3)}} w_{jk}^{(3)} \left(\frac{\partial f(z_j^{(3)})}{\partial z_j^{(3)}} \delta_{k,m}^{(out)}\right)\right)$$

Generalized to Multi-Layer Networks

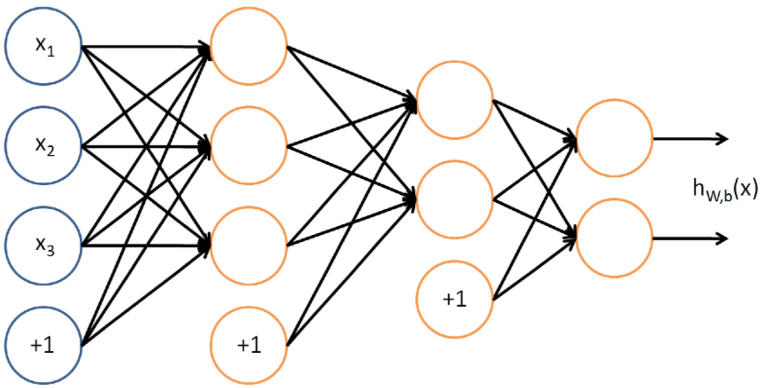

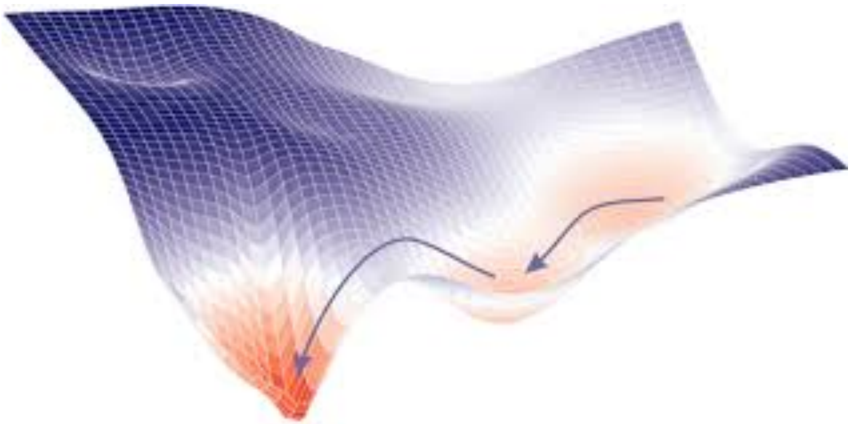

Training requires massive computing with massive data.

Difficulties:
- Network has millions of parameters
- Training data is very noisy.
- Loss function has local minima

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

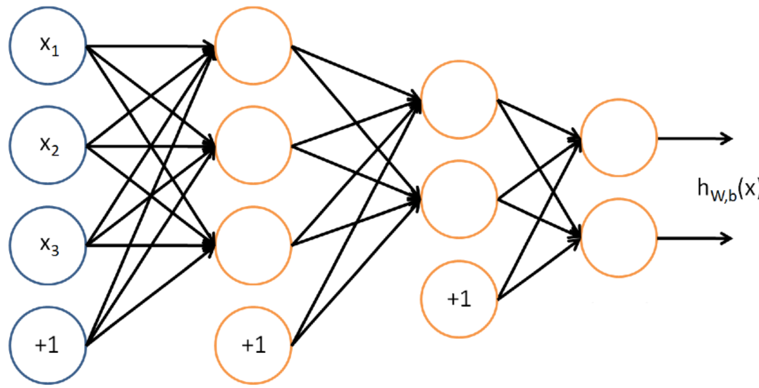

<u>Problems</u>:

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

Neural networks were (mostly) abandoned in the 1990s in favor of mathematically sound Bayesian machine learning.

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, grid computing

(3) Prohibitive cost of encoding knowledge

$\Rightarrow$ Generalized Deep Learning

Le Net5 - 1994

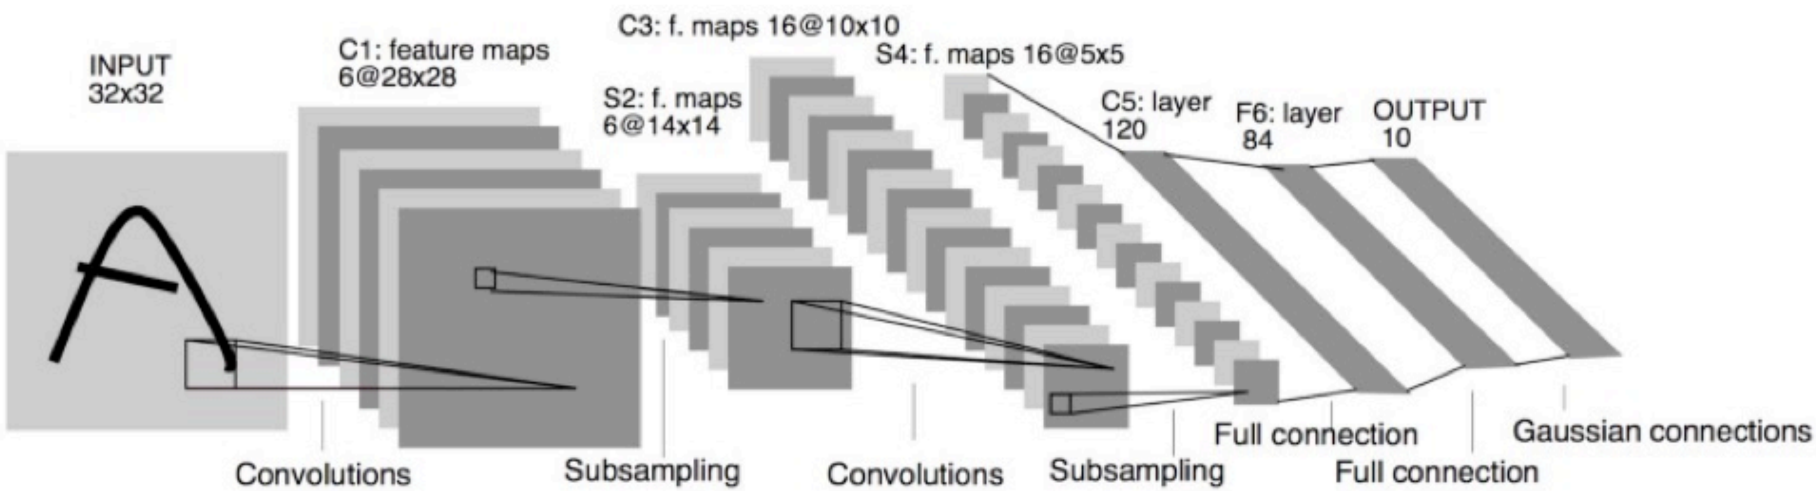

7-level convolutional network by Yann LeCun in 1998.
State of the art for recognizing hand-written numbers on checks.

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

Convolutional Neural Networks

Single Receptive Field per layer

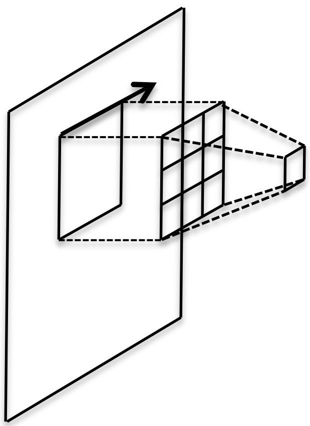

$$z_{l+1}(x,y) = \sum_{u,v} a_l(u-x, v-y) w_l(u,v)$$

Multiple Receptive Fields per layer

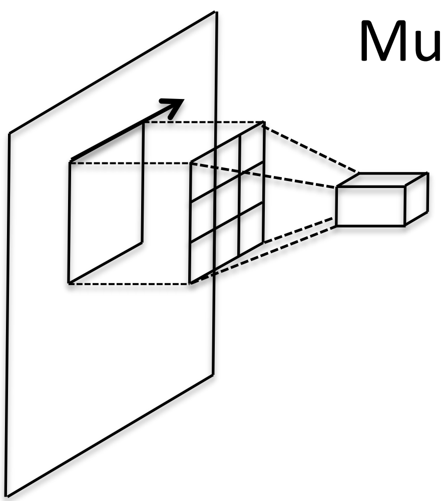

$$\vec{z}_{l+1}(x,y) = \sum_{u,v} \vec{a}_l(u-x, v-y) w_l(u,v)$$

Pooling (max Sampling)

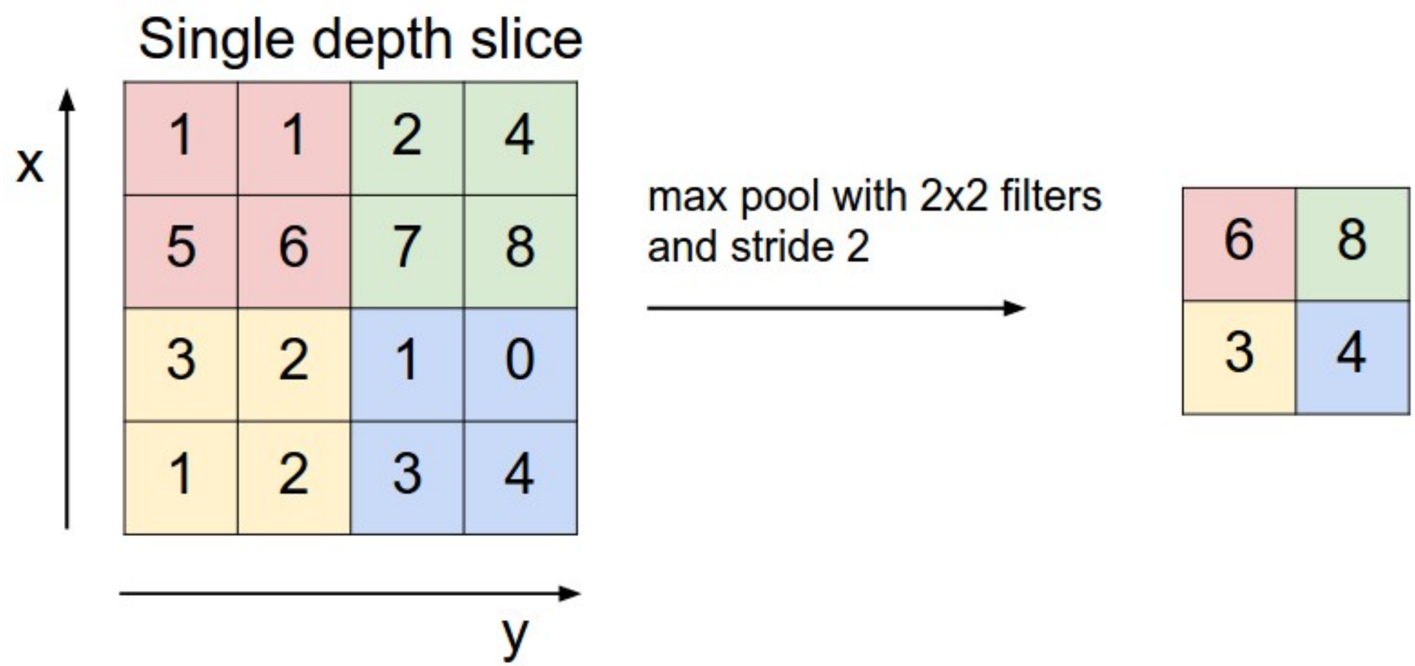

AlexNet 2012

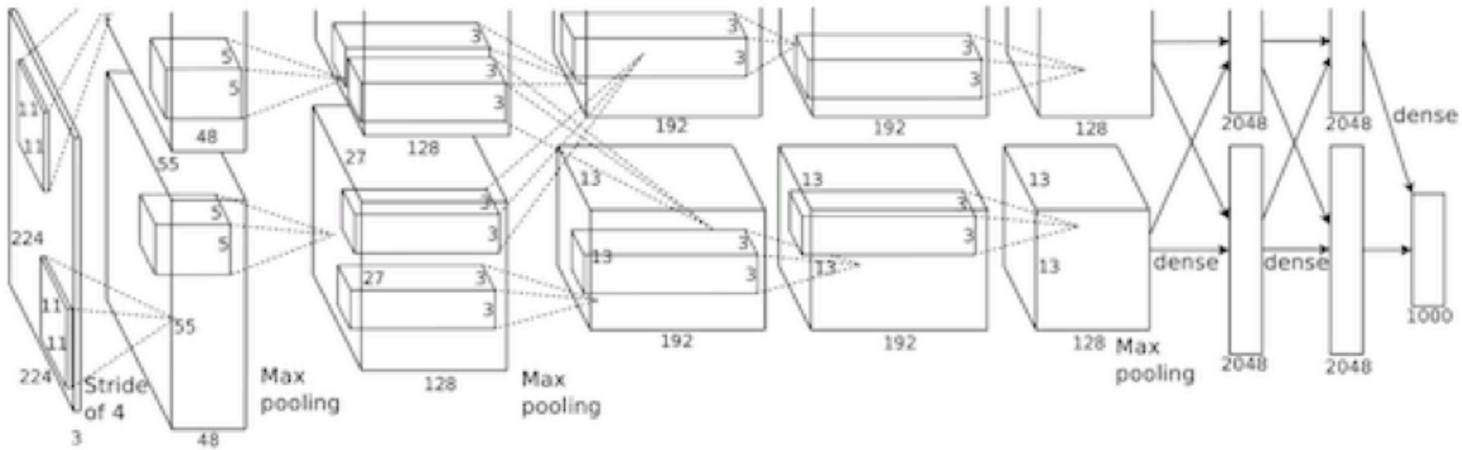

Created by Alex Krizhevsky and Geoff Hinton (based on LeNet)
Won the ImageNet Large Scale Visual Recognition Challenge in 2012 by a large margin with an error of around 15%

Triggered a paradigm shift for Computer Vision, Speech Recognition, Machine Learning and (more recently) Artificial Intelligence.

ImageNet Large Scale Visual Recognition Challenge in 2012

Ranking of the best results from each team

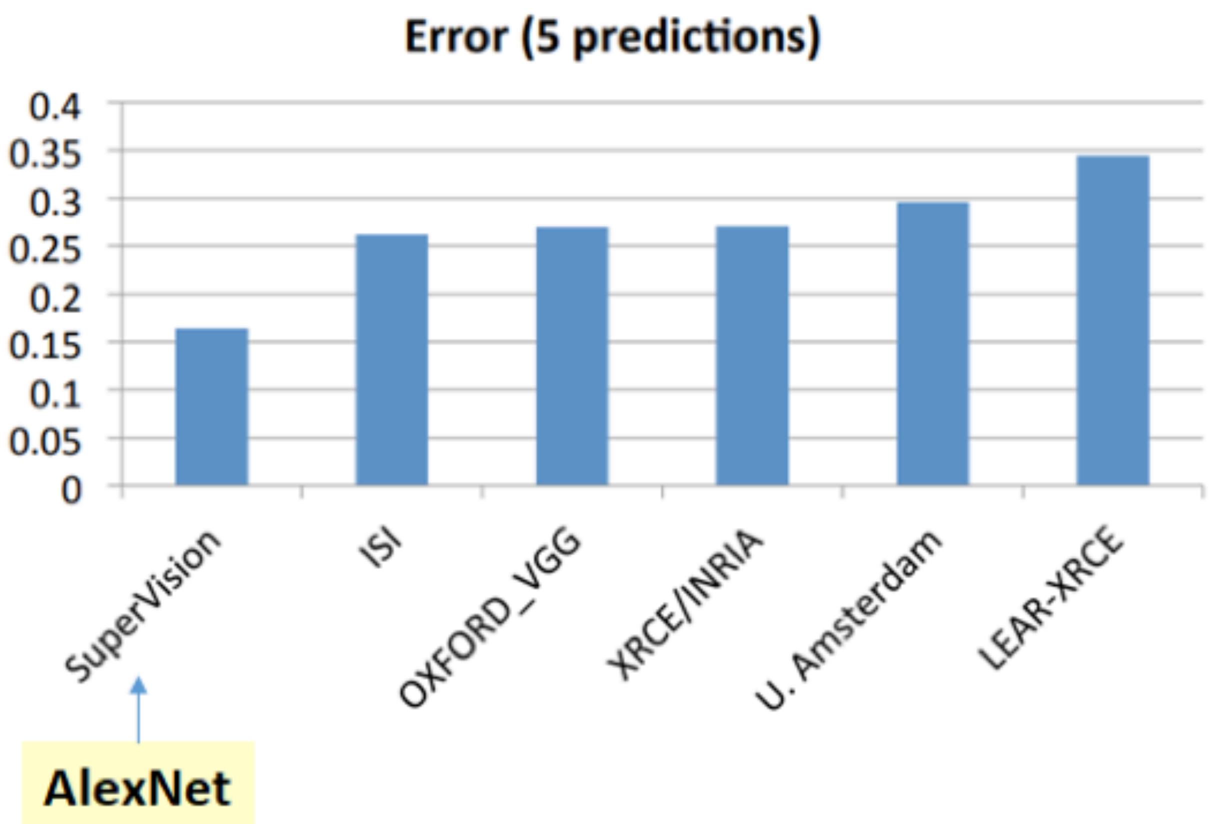

VGG 2015

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

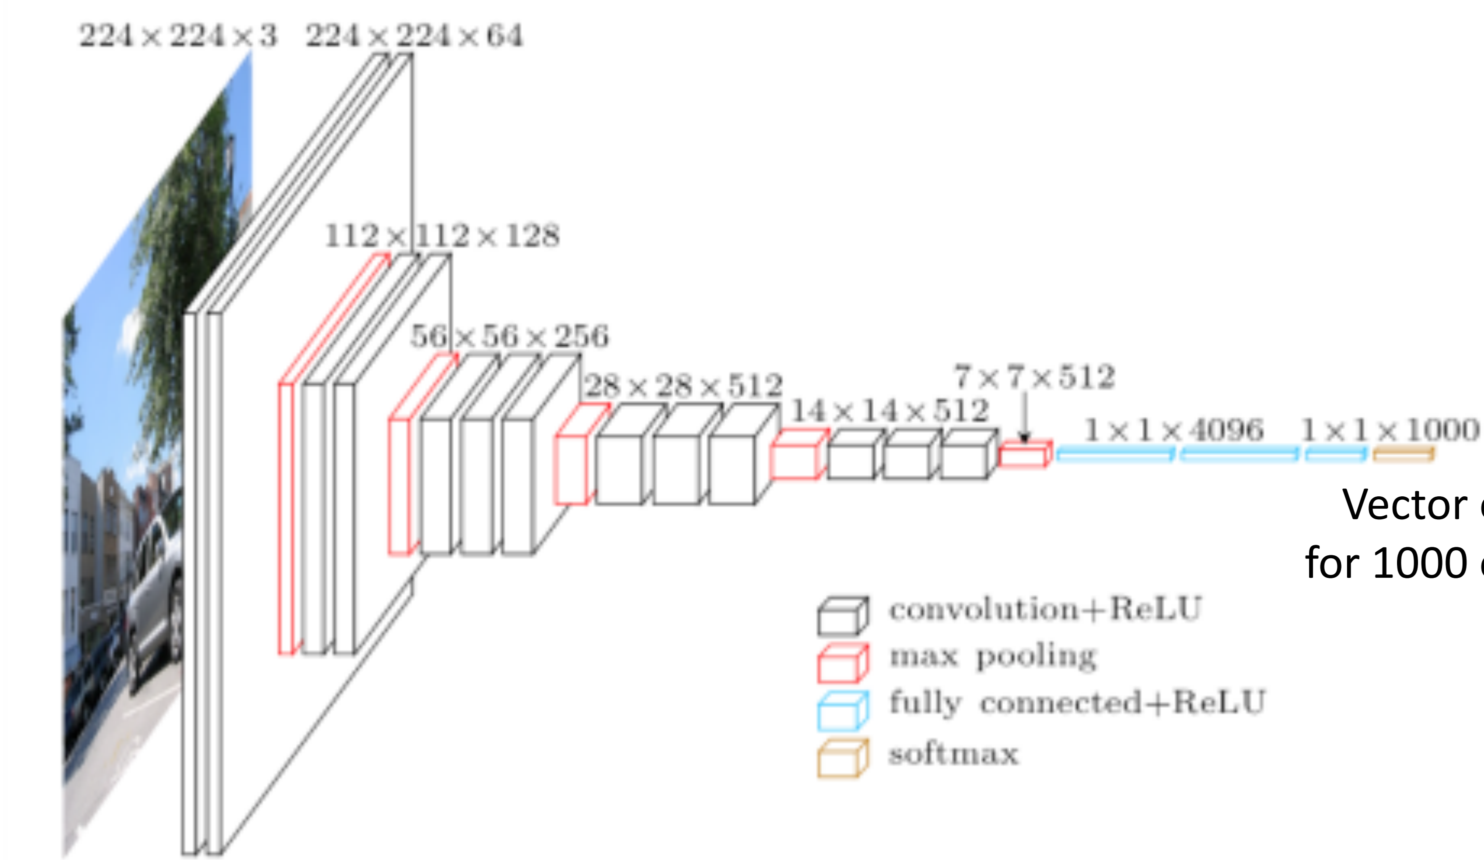

Output:
Vector of Likelihoods
for 1000 object classes

Generative and Discriminative Networks

$$\vec{X} \longrightarrow \boxed{\quad D(\vec{X}) \quad} \longrightarrow \hat{y} \qquad\qquad y \longrightarrow \boxed{\quad G(y) \quad} \longrightarrow \vec{X}$$

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

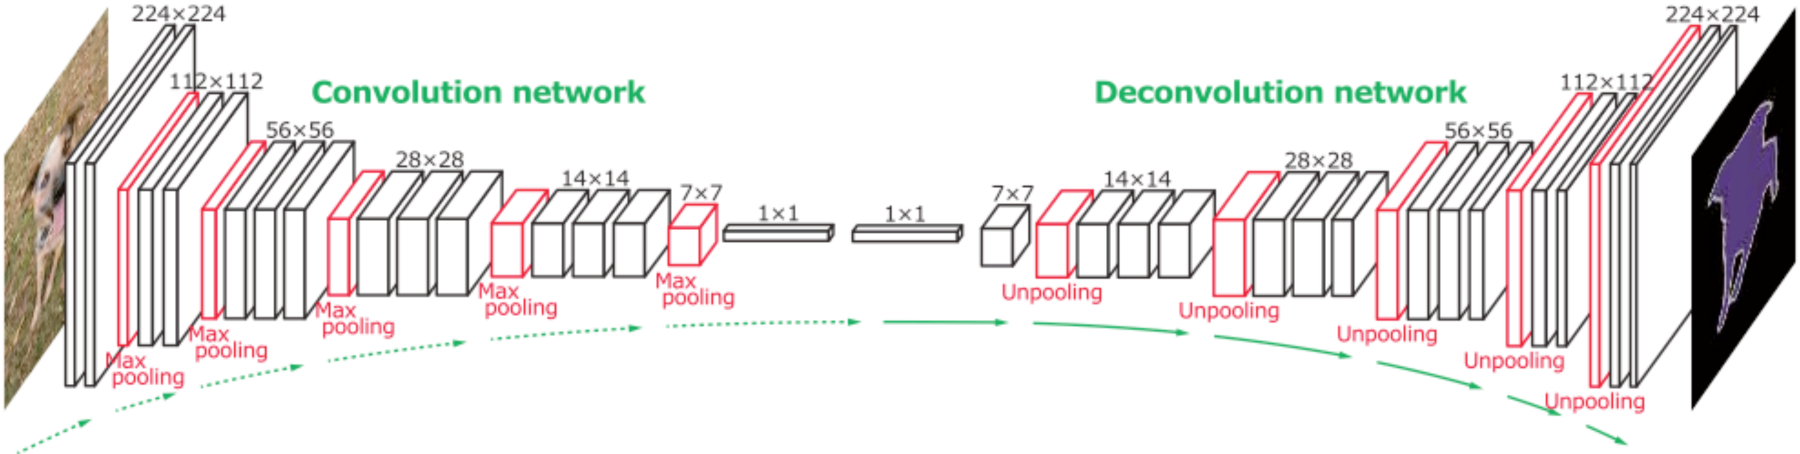

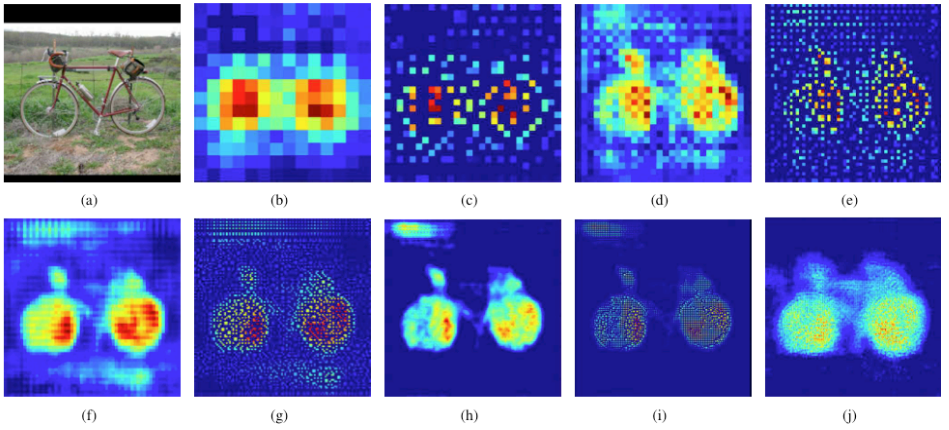

Autoencoder

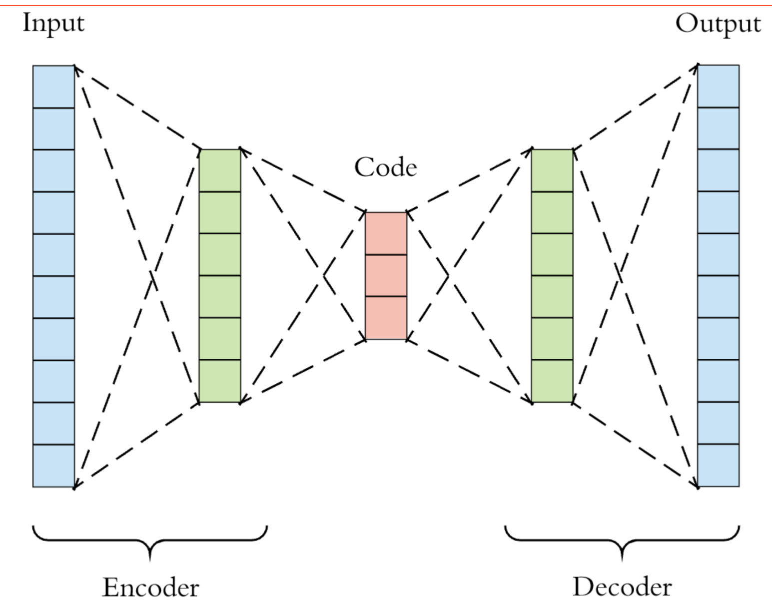

An Autocoder learns to reconstruct (generate) clean copies of data without noise. Key concepts:

1) Training data is target. Error is difference between input and output
2) addition of a “sparsity term” to loss function (modified back-prop)
3) Minimum number of independent hidden units (Code vector)

Problem: Output space is not continuous

Variational Autoencoder

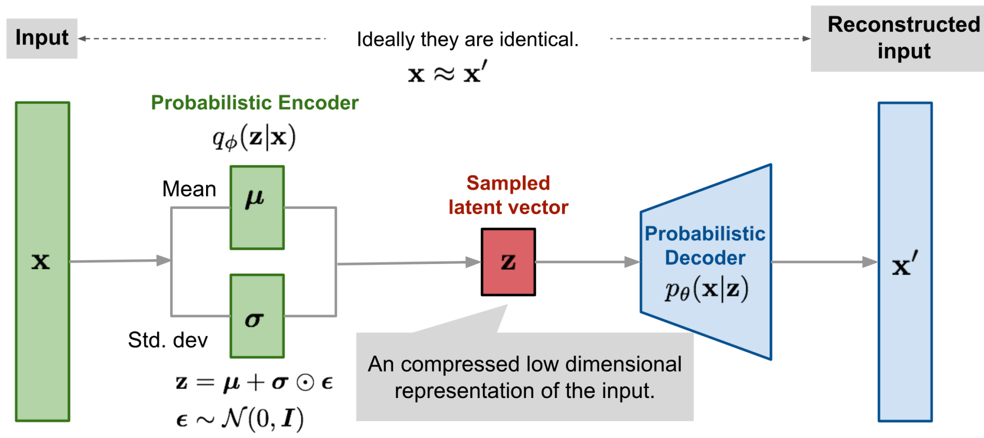

A VAE can be used to generate synthetic output.

Example:

1) Train VAE on dancers doing the same dance.

=> Code represents posture

2) Drive decoder of a dancer from encoder of another.

Generative Adversarial Networks

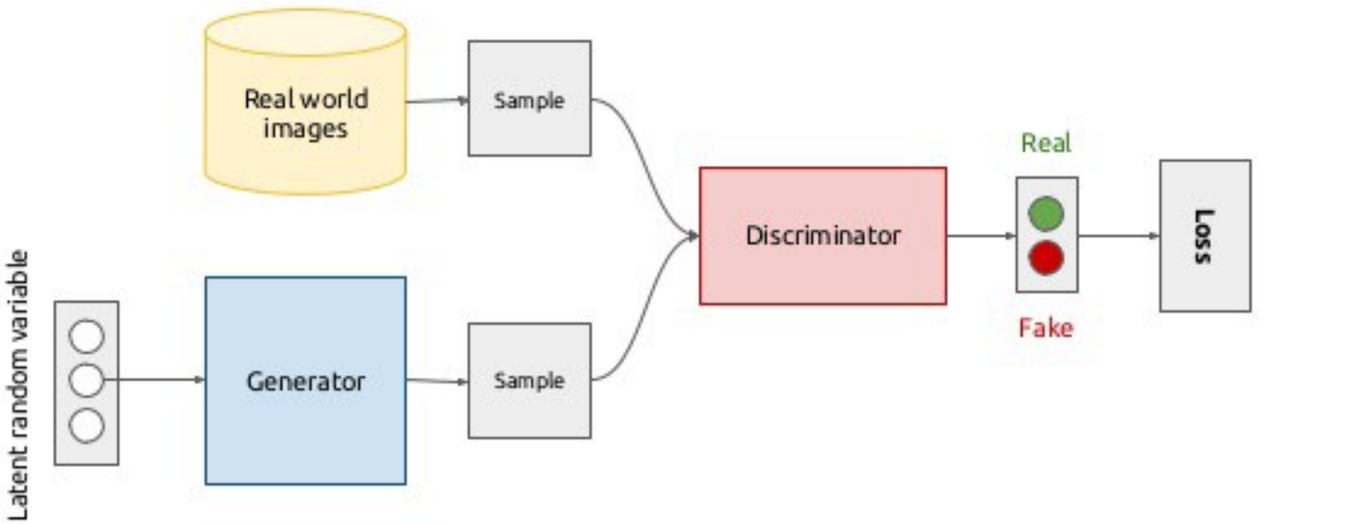

Unsupervised competitive learning between a Generative and a Discriminative network

Can be used to generate DeepFake, Realistic Speech synthesis, photo Realistic images (Hot topic at the IJCAI 2018)

Deep Reinforcement Learning

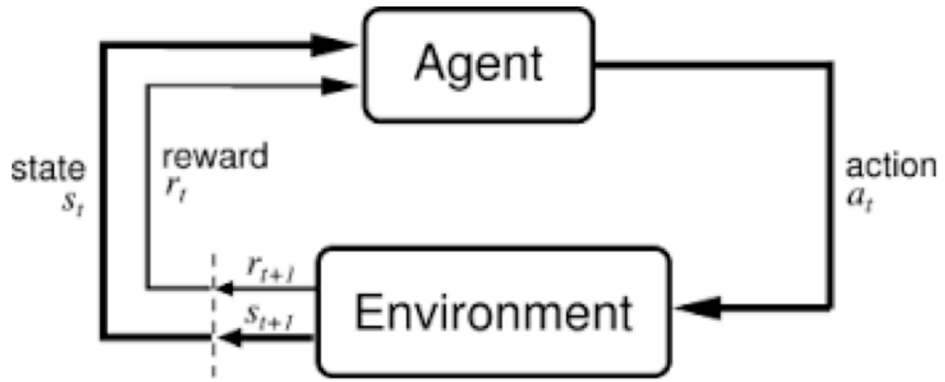

*Credit: Sutton & Barto

Re-enforcement learning with deep networks. Can be used with games or with simulated environments for self-supervised learning.

Enabled AlphaGO to beat world champion at Go in 2015

Chatbots: Current Technology

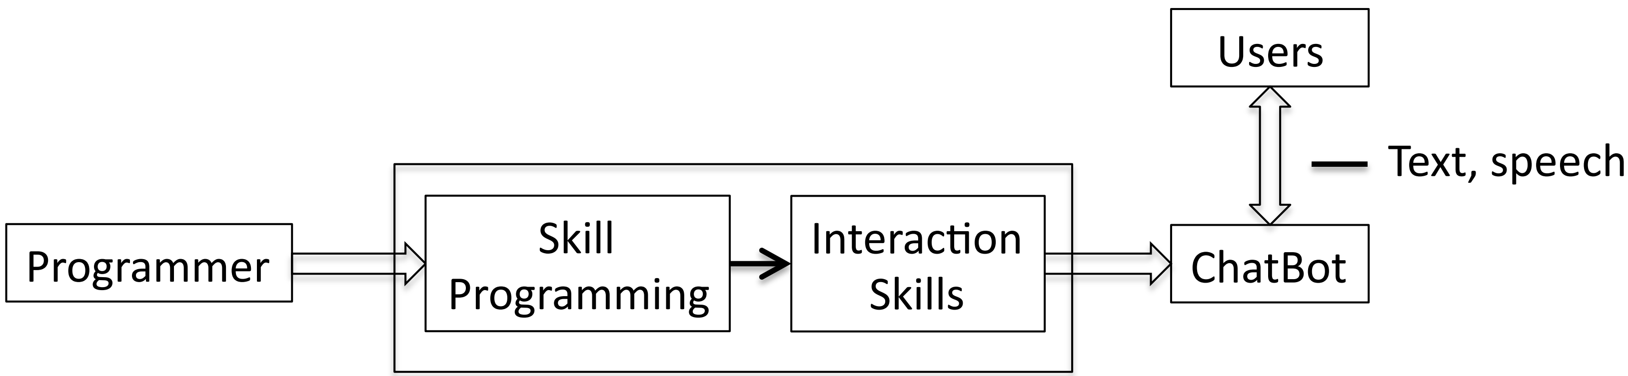

Chatbots “skills” are programmed using linguistic technologies

Chatbots do not “understand” the conversation, they respond with pre-programmed rules (Stimulus – Response)

Word2Vec – Linguistic Encoding

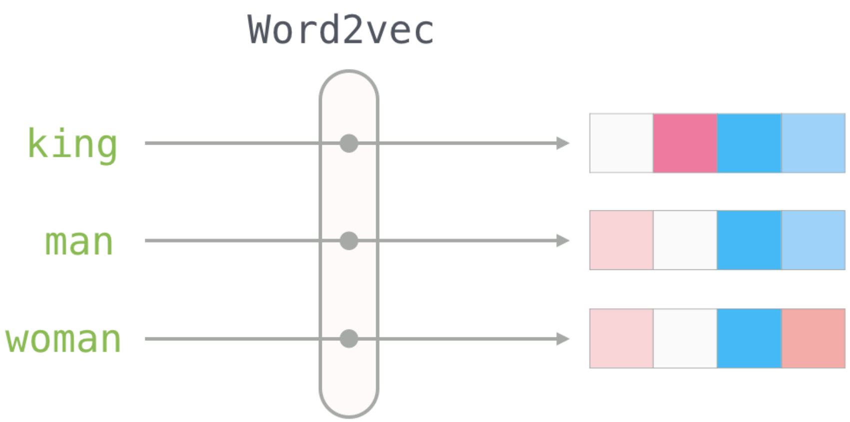

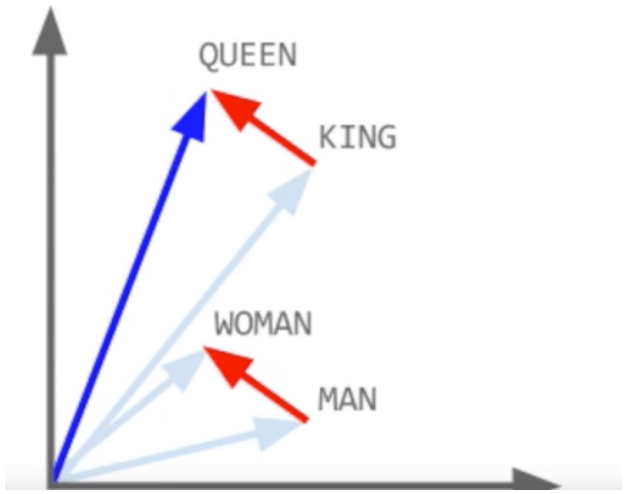

- Any document is a context, represented by word frequencies.
- Word2Vec projects each word onto a Context vector.
- Context vectors can be compared with cosine angles
- Context vectors can be used with Deep Learning!

Applications: Chat-bots, Data mining on the internet

Chatbots Training by Interaction

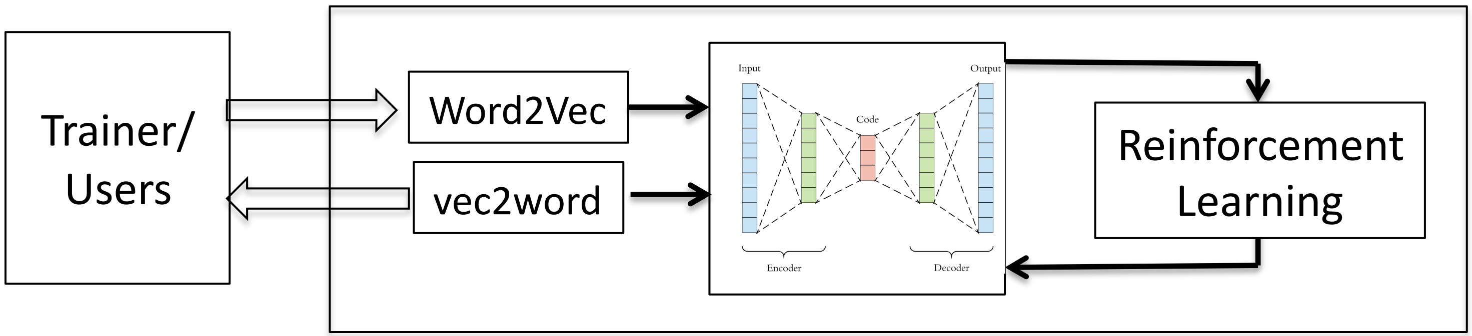

AI Chatbots: Learn through interaction with users.

Example: MicroSoft Tay AI twitterbot (2016)
Designed to imitate speech of 19 year old female and learn vocabulary and language by interaction.
Attacked by internet trolls and trained to generate insults.

Cognitive Computing

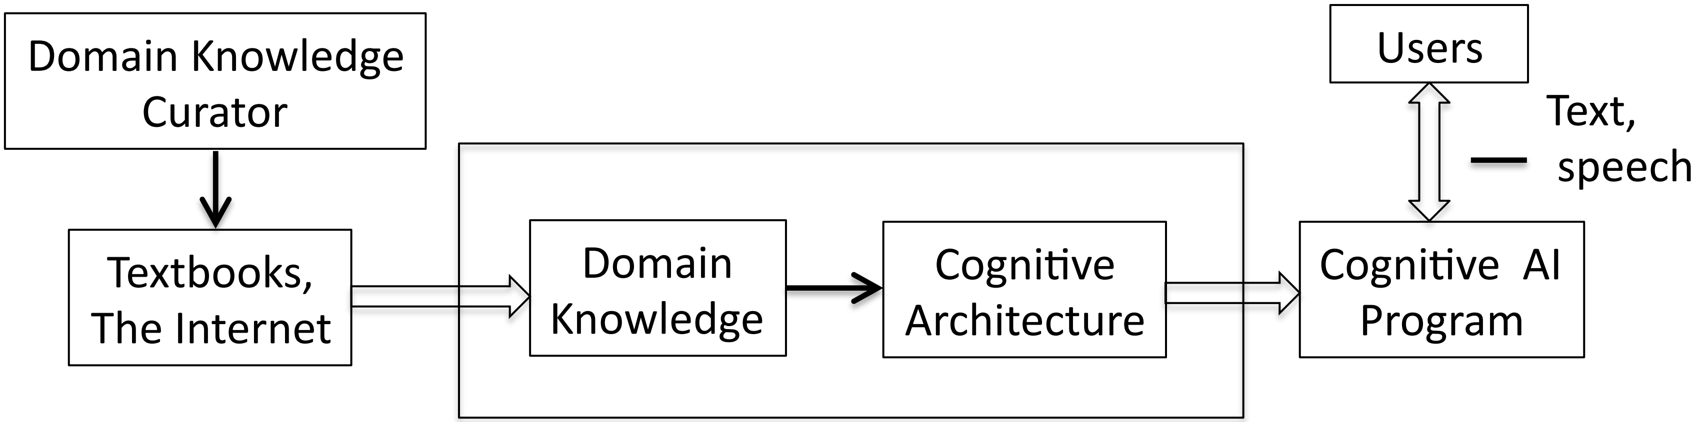

Cognitive Computing encodes knowledge from any written source (textbooks, literature, the internet) to generate a domain expert advisor program (an expert system!)

Example domains: Medical, Legal, Financial…

Cognitive Computing Example:
IBM Watson

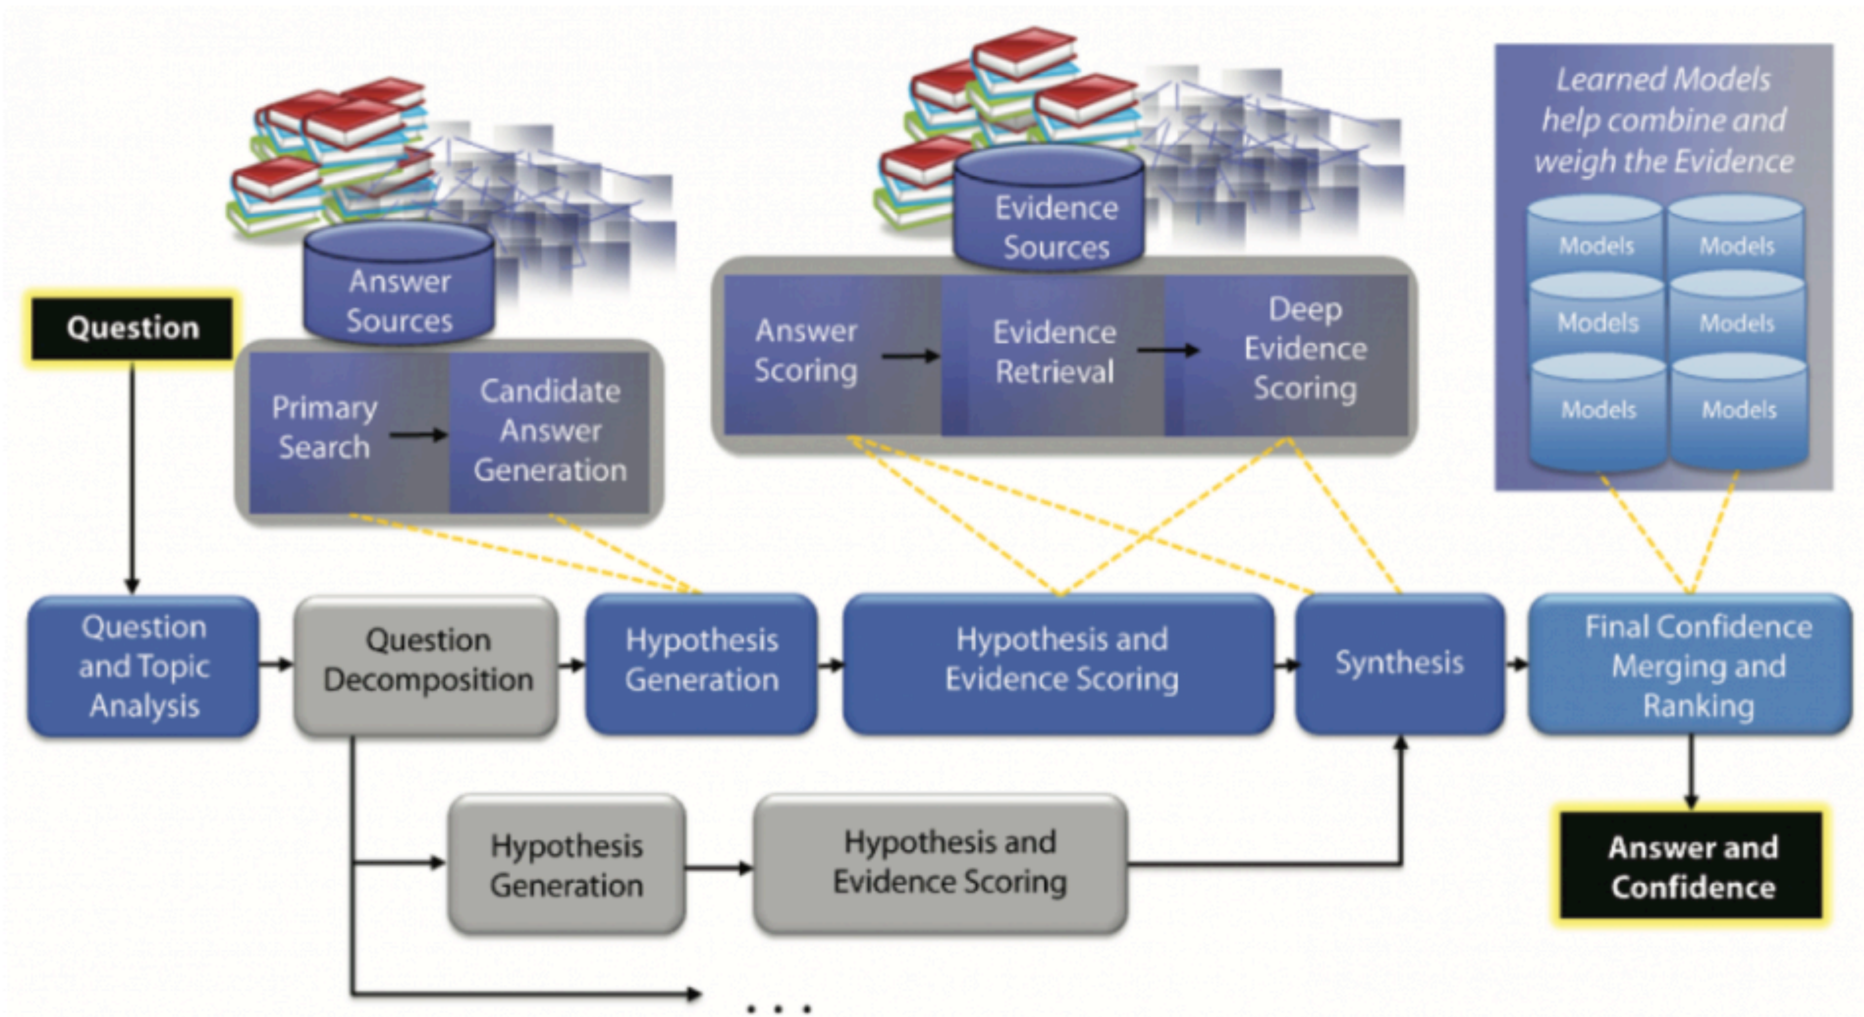

Source: IBM

Pepper – Animated by Watson

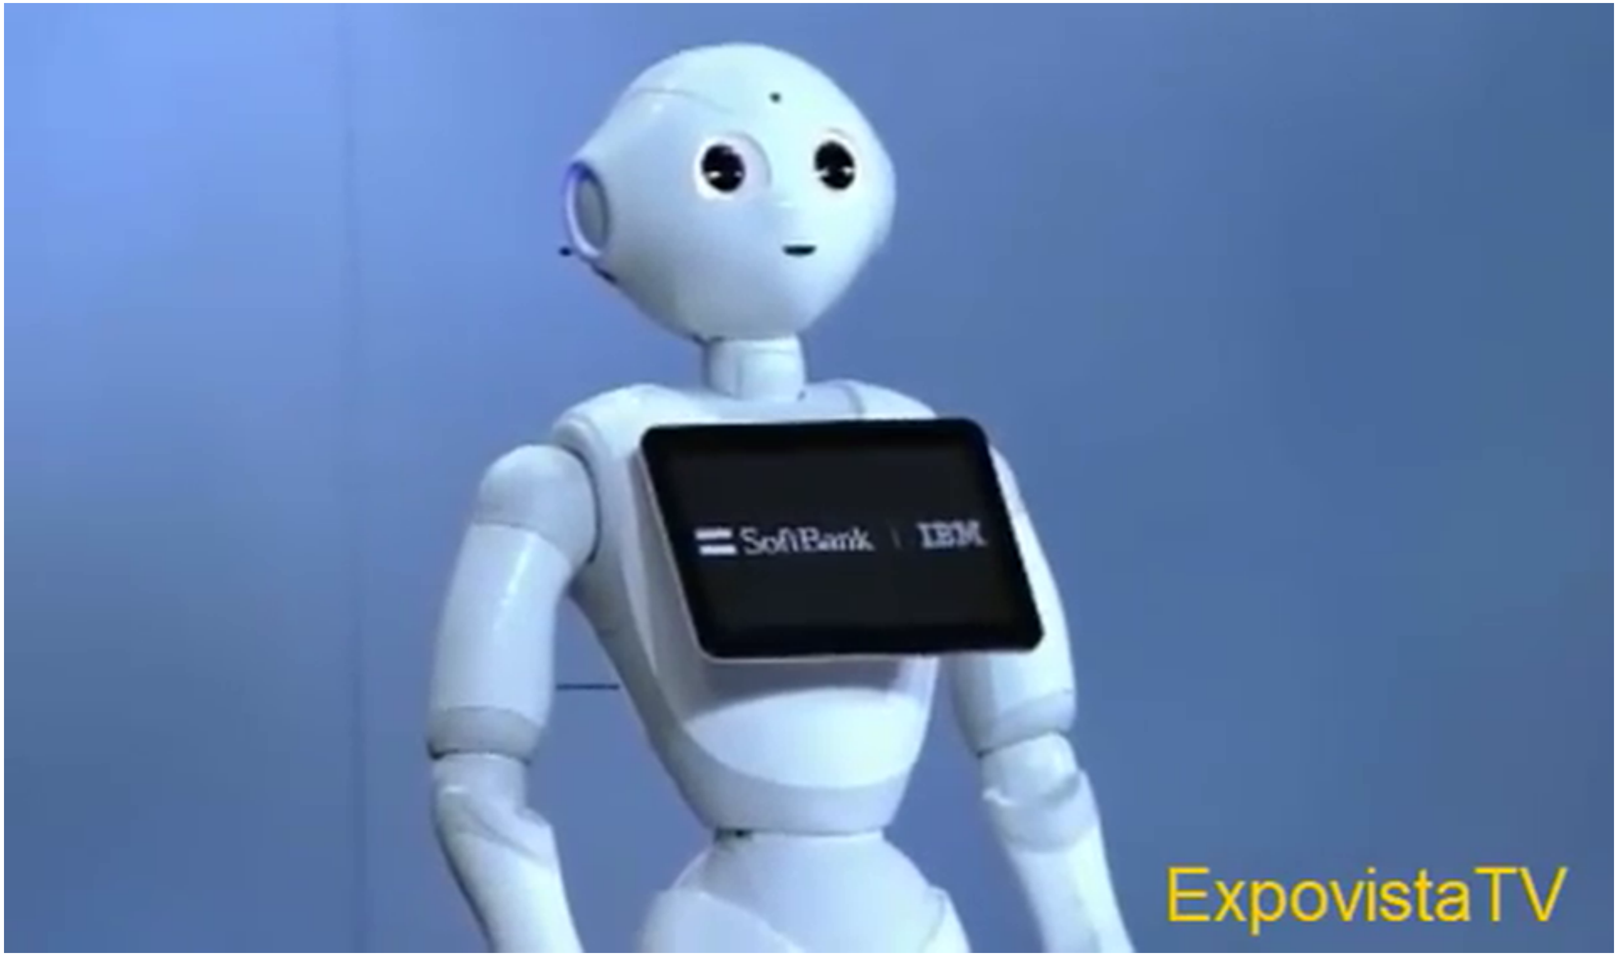

Pepper (Softbank Robotics) and Watson at CES 2016

AI as a Service on the Cloud

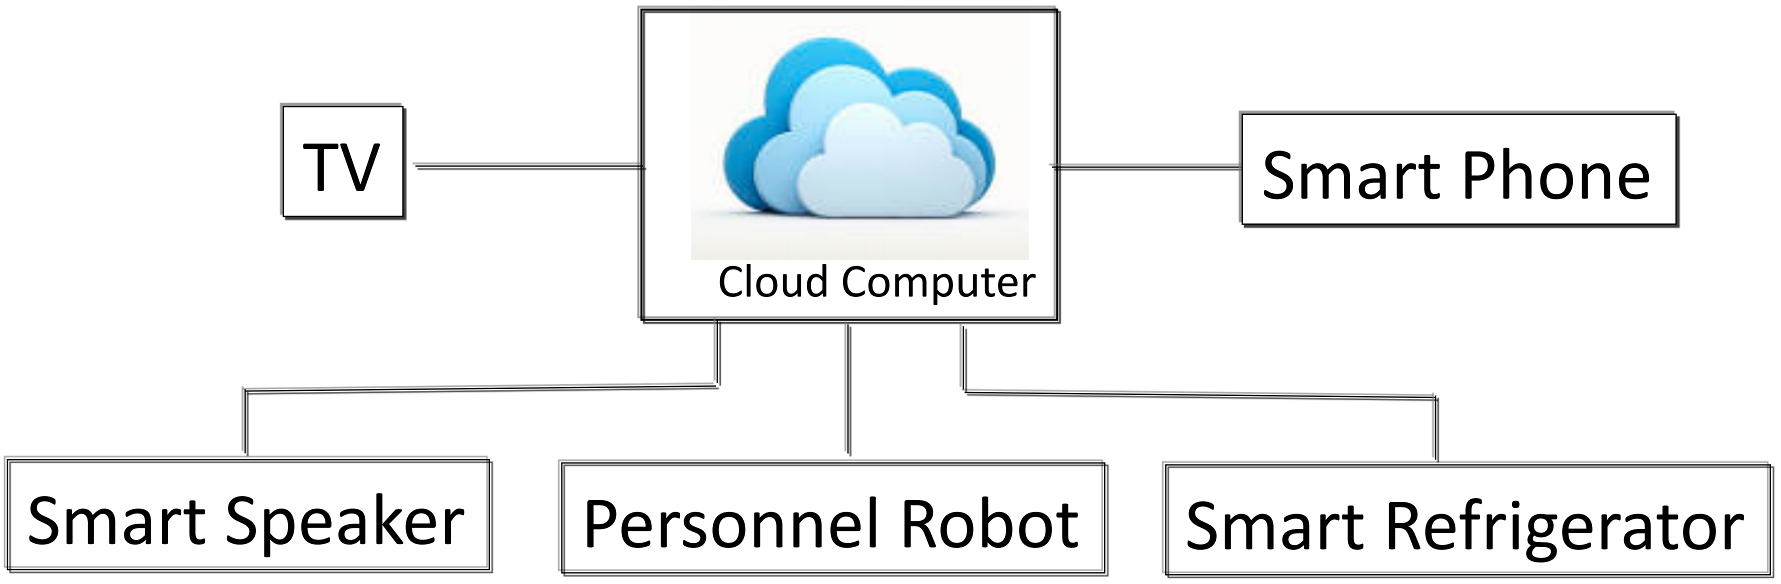

Machine Learning :
A new paradigm for building systems

Machine Learning from Data:

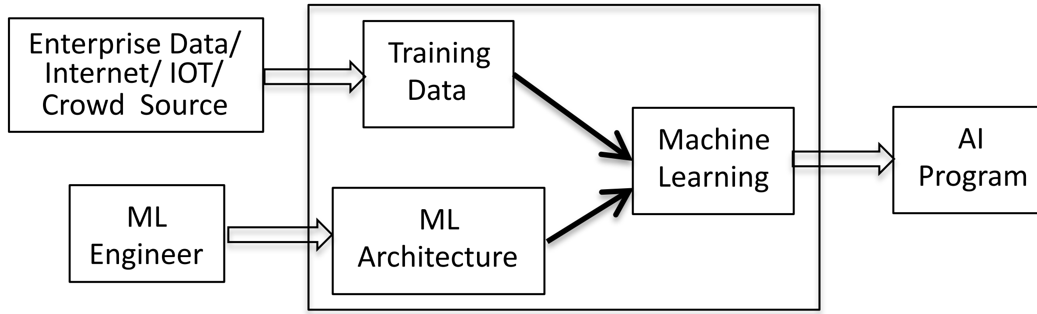

Machine Learning from a model: (fast approximate estimation)

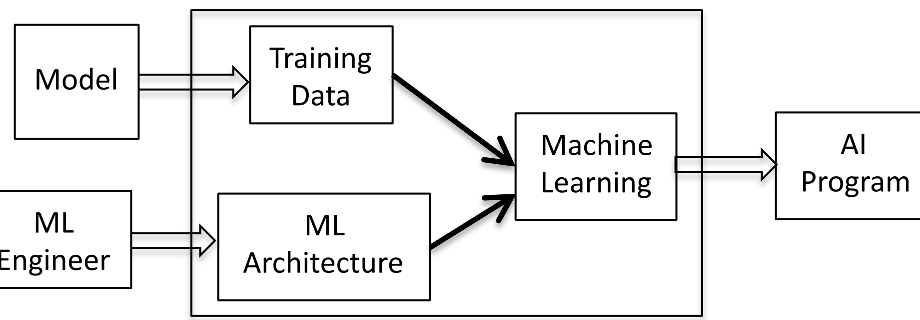

With Machine Learning, the data becomes the code.

Problems: Verification: How can you **certify** correct behavior?
Explanation: How can you **explain** output from network?

Can we certify AI for Critical Systems?

Classic Software Process

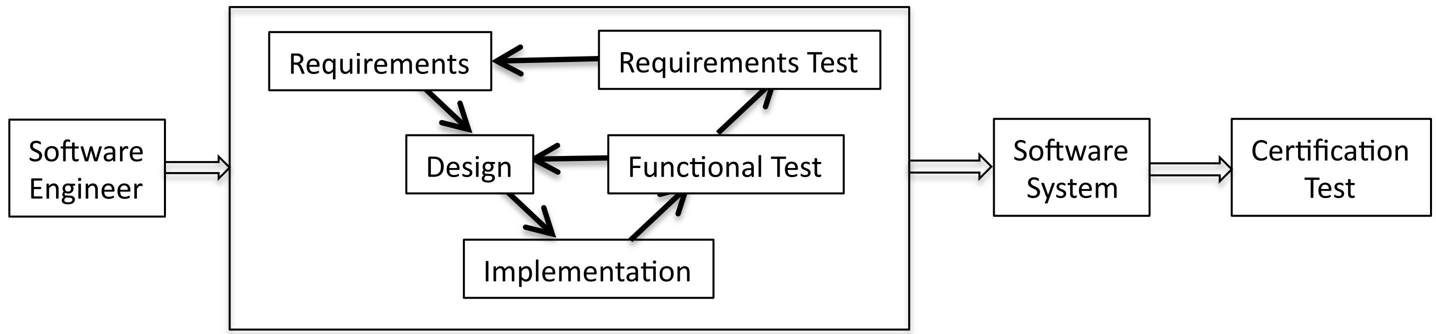

Classic Process: Certification guarantees compliance with specifications.
In case of Failure, the engineer is responsible.

Machine Learning

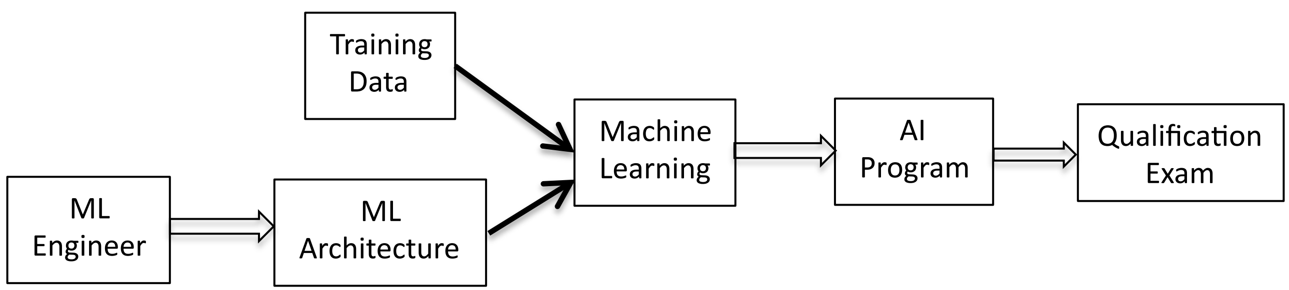

The AI Program Is a black box. What are the specifications? What do we test?
Who is responsible if the system fails?
Can we "qualify" AI systems as we do with programmers drivers, Airline pilots?

Explainable AI

AI systems that provide results that can be understood by people.

There are several different AI explainability problems.

1) Interpretability: Explain the output from a deep network.
2) Transparency: Explain/justify the decision given by an AI system.
3) Predictability: Explain why a system behaved in a certain manner.

(We currently lack a theory for deep learning.
Current Architectures are the result of empirical research.)

Chat Bot Approach to Explanation

Preprogrammed Explanation using linguistic tools:

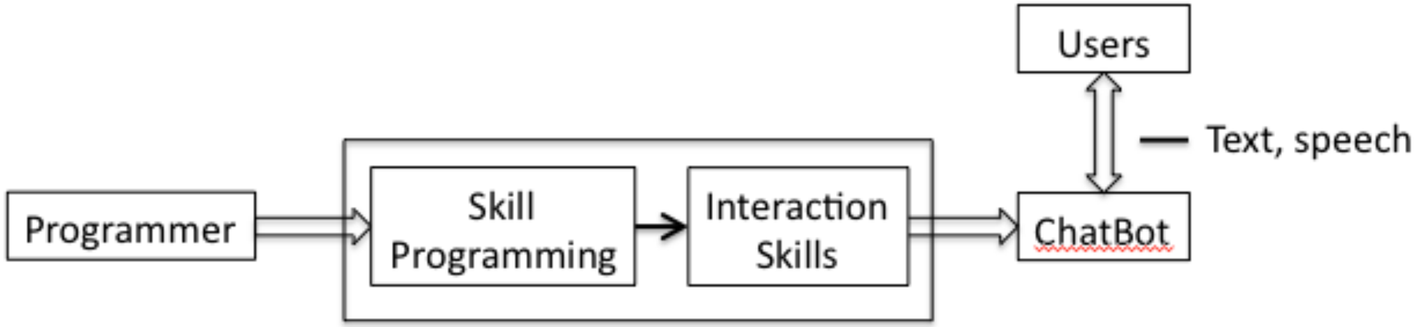

Training Explanations through interaction.

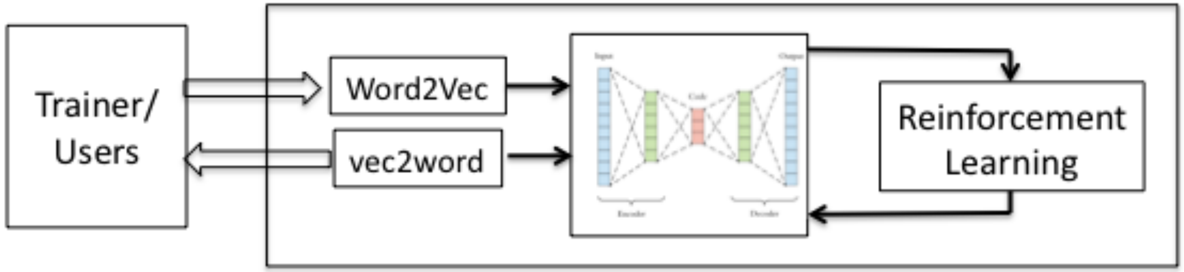

Explainable AI – Personal View

Explanation: a reason for 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.

Trustworthy AI

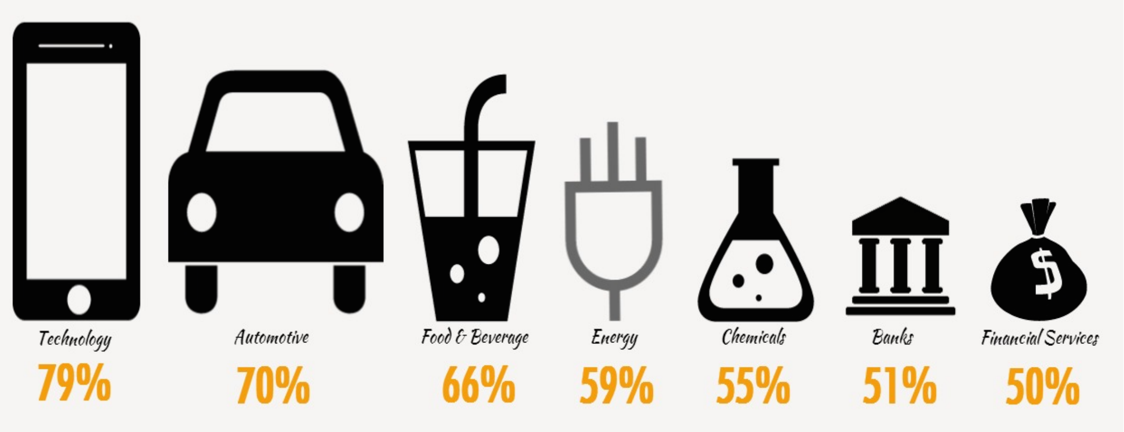

Trust: The ability to inspire confidence that a system is secure, available, and private.

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

Potential Innovations from AI

What domains are most suitable for rupture from AI technologies?

AI is the fire. Data is the fuel.

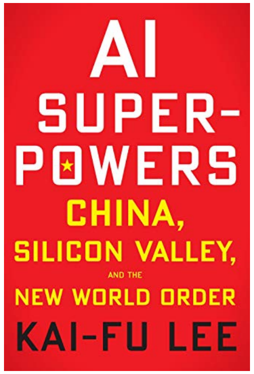

To predict AI innovation, look for the data (Kai Fu Lee).

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.

Potential Innovations from AI

AI: Human level ability at interaction

Interaction with <u>People</u>:
=> Commerce, Education, Entertainment, Well Being
Interaction with the <u>Physical World</u>:
=> Robotics, Transportation, Manufacturing, …
Interaction with <u>Systems</u>:
=> Smart Buildings, Smart City, …
Interaction with <u>Information</u>:
=> Virtual Personal Assistant, travel planning, …

New Categories of interactive Systems

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)

Advisors

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

Affectors

Affectors: Objects that interact to inspire affection

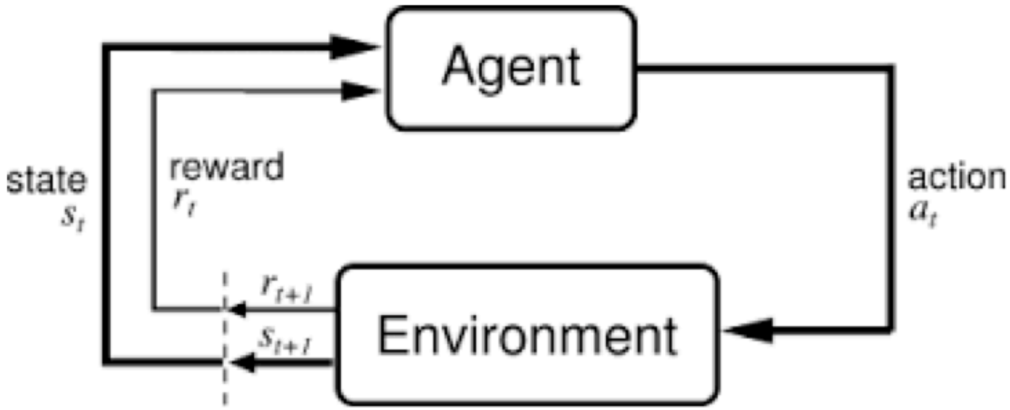

*Credit: Sutton & Barto

Multimodal perception of affect and **Deep Reinforcement Learning** can be used to learn actions for stimulating Affection.

Used with affective computing, can be adapted to any interaction.

Media: Augmented Reality

Current AR technology: Displayed information is “pre-programmed”

With AI and Machine learning:

- Computer vision can learn to recognize new phenomena
- Multi-modal interaction and Reinforcement Learning can be used to learn the appropriate information to display

Advisors: Cognitive Computing

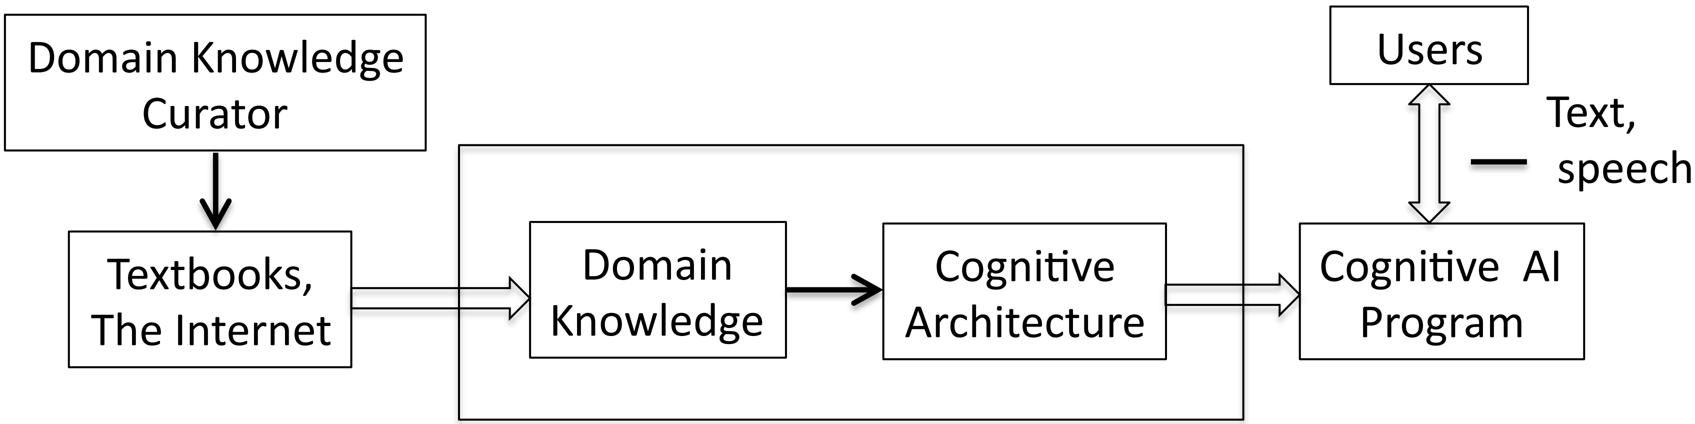

Cognitive Computing can be combined with perceptual user interfaces to provide expert advice.

Potential Innovations from AI

How can we use AI as a tool for scientific research?

Cognitive Computing as a tool for Monitoring the Scientific Literature

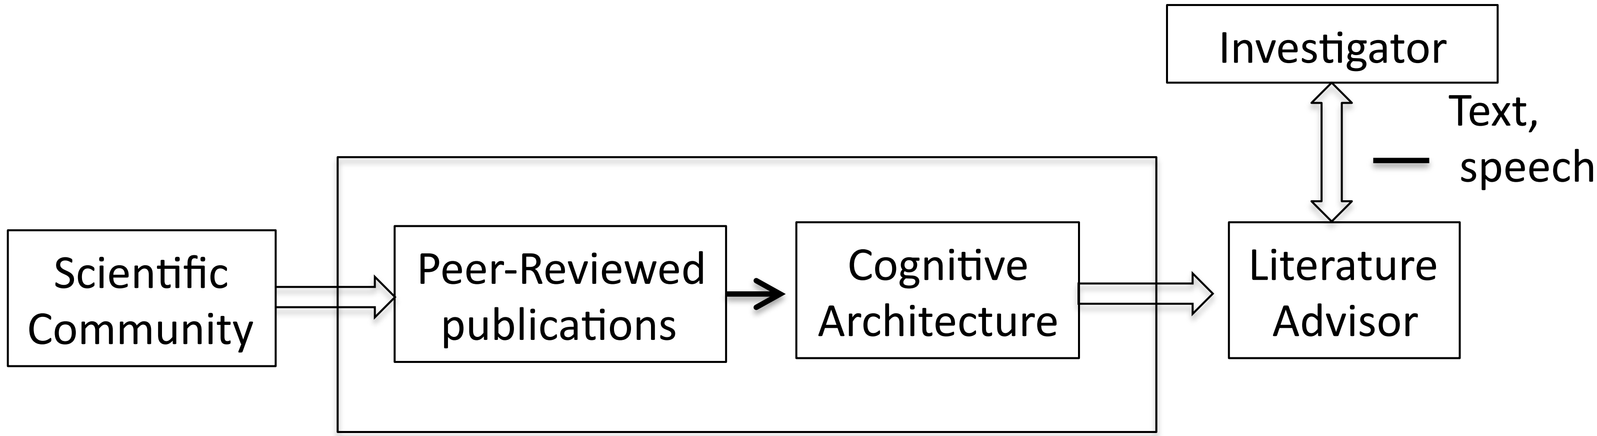

Cognitive computing can potentially provide a tool to provide advice and guidance in monitoring the scientific literature.

Cognitive computing will NOT discover new concepts, but it CAN provide a tool to augment human intelligence.

Beyond Advisors:
Collaborative Intelligent Systems

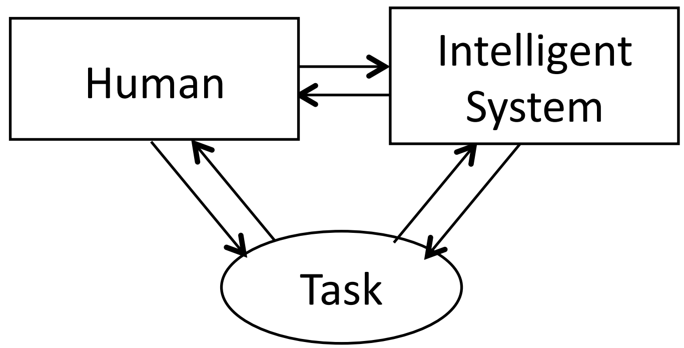

Collaboration is a process where two or more actors (agents) **work together** in order **to achieve** some shared **goals**.

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.

Empowering humans with AI.
Collaborative Intelligent Systems

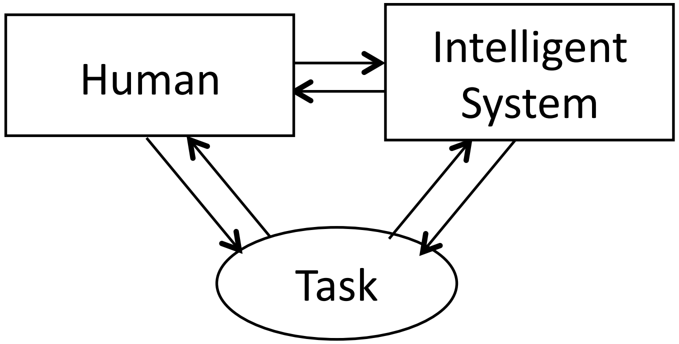

Challenge: Build collaborative intelligent systems that work synergistically with human user. Human and system each play roles suited for their abilities:

System: Collect, analyze and display massive volumes of data
Detect anomalous phenomena

Human: Determine challenges where theories and models fail.
Generate conceptual understanding.

Conclusions:

- Intelligence is Human level performance at Interaction
- Deep Learning is a rupture technology for AI.
- Deep Learning is made possible by planetary scale data, and massive computing.
- For Innovation: if AI is the fire, data is the fuel.
- Opportunities for interaction with people, environments and information.
- Explainable, Certifiable and Trustworthy AI systems are open challenges.

But is AI a rupture technology for Science?

- Deep learning systems can not explain their reasoning.
- We can train systems to explain human knowledge, however trained system can not invent new concepts.

Challenge:
Build collaborative systems that synergistically empower individuals to create new concepts, theories and models that predict and explain.

Is Artificial Intelligence A Rupture Technology for Scientific Research?

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