

Artificial Intelligence For Human-Computer Interaction

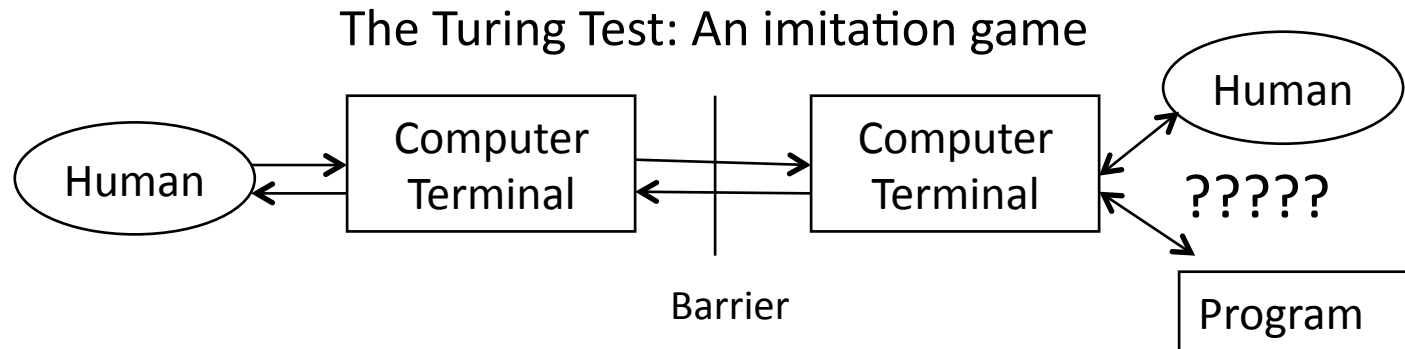
James L. Crowley
Professor Grenoble INP
Grenoble Informatics Laboratory (LIG)
INRIA Grenoble Rhone-Alpes
Univ. Grenoble Alpes

Artificial Intelligence For Human Computer Interaction

Outline:

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


Artificial Intelligence (AI)



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)



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

Intelligence: Human-level performance at interaction.
(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.

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

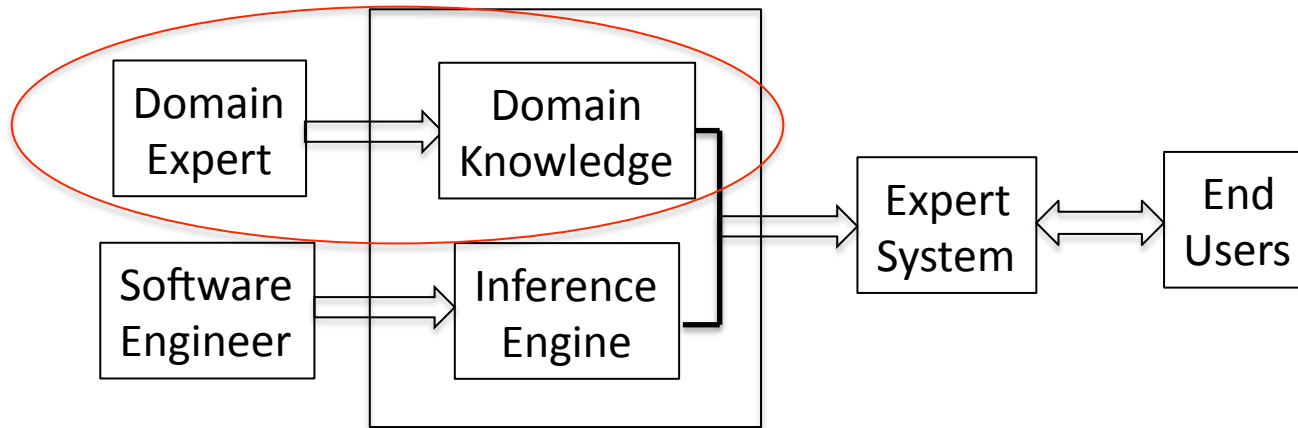
Dominant Paradigms for Artificial Intelligence:

- 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 (Gradient Descent).


Provided a simple alternative to symbolic computing

Feed Forward Network



Decision: Sigmoid function

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





$$\frac{df(z)}{dz} = f(z)(1 - f(z))$$

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




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

Effective training requires massive computing with massive data.

Artificial Neural Networks (1975-1990)

Multi-layer Perceptrons with Back Propagation Learning



Problems:

- 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

⇒ Planetary scale data from the internet and the WWW


(2) Insufficient computing power

⇒ Moore's Law, GPUs, grid computing

(3) Prohibitive cost of encoding knowledge

⇒ 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.

AlexNet 2012




Created by Alex Krizhevsky and Geoff Hinton

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.

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.

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



Generative and Discriminative Networks



Discriminative Networks:
Does data X contain class y ?


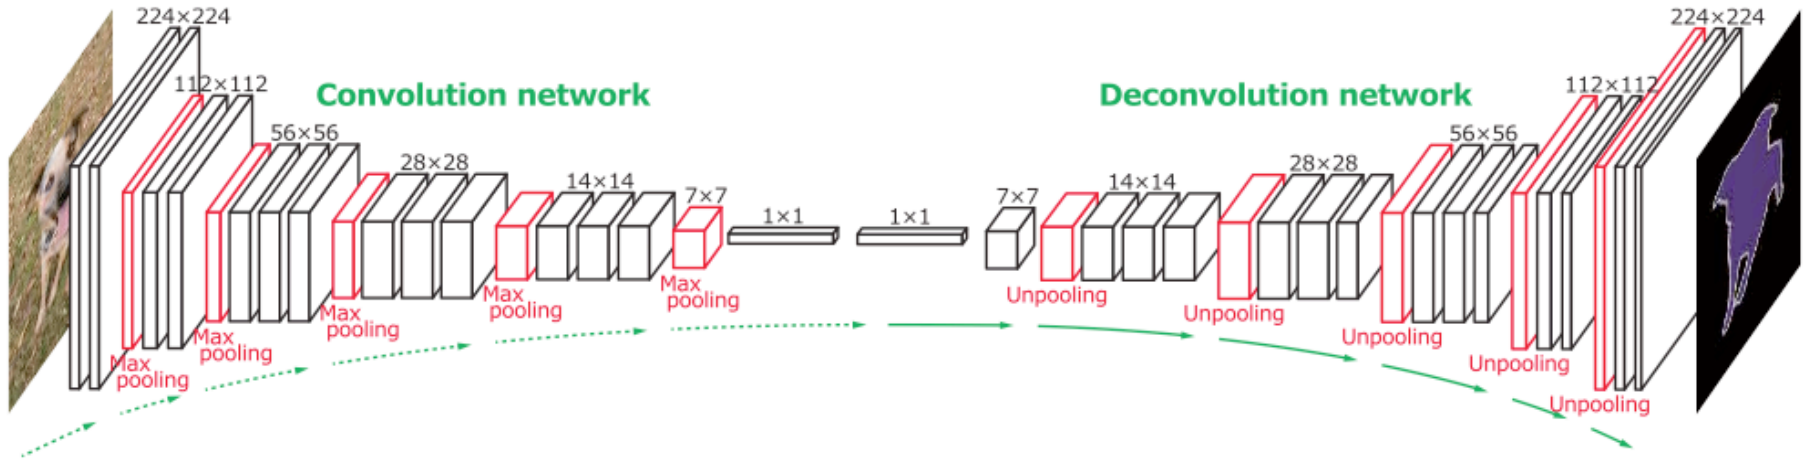
Generative Networks:
Generate pattern X for class y

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 Autoencoder 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.

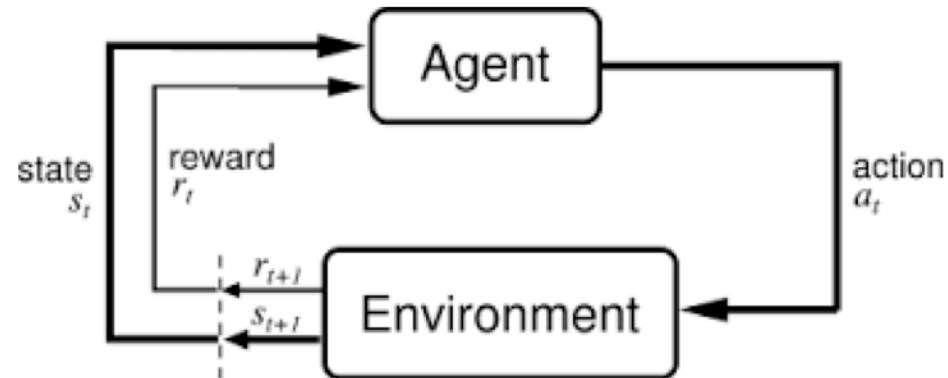
Generative Adversarial Networks



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-inforcement 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

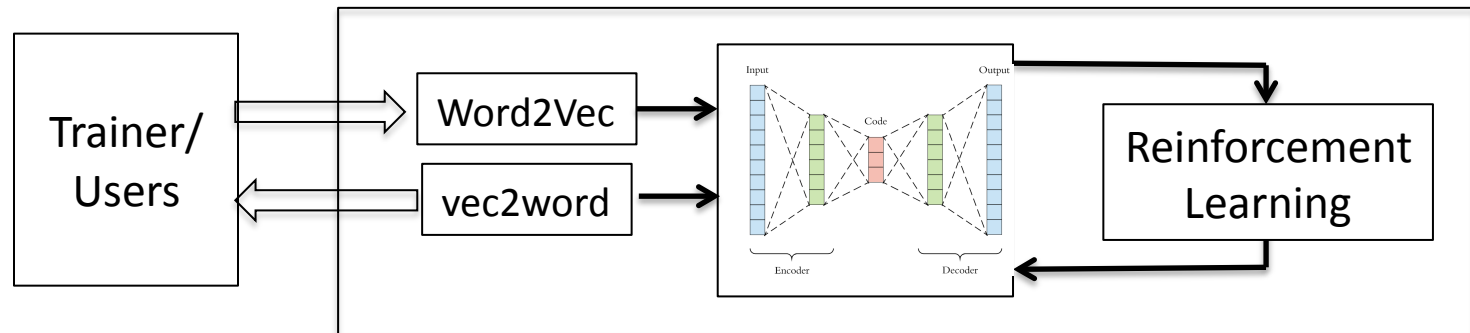
Word2Vec – Linguistic Encoding



- 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



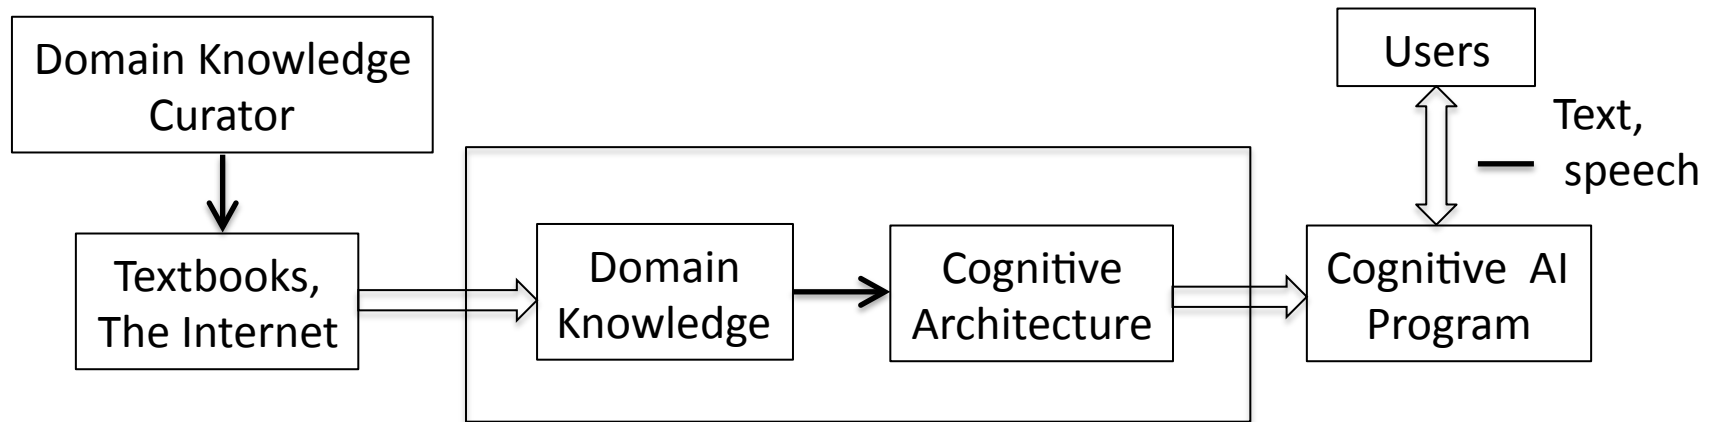
Research Challenge: Learn through interaction with a trainer.

Example: MicroSoft Tay AI twitterbot (2016)

Designed to imitate speech of 19 year old female and learn vocabulary and language by interaction.

Attacked by trolls and trained to generate insults.

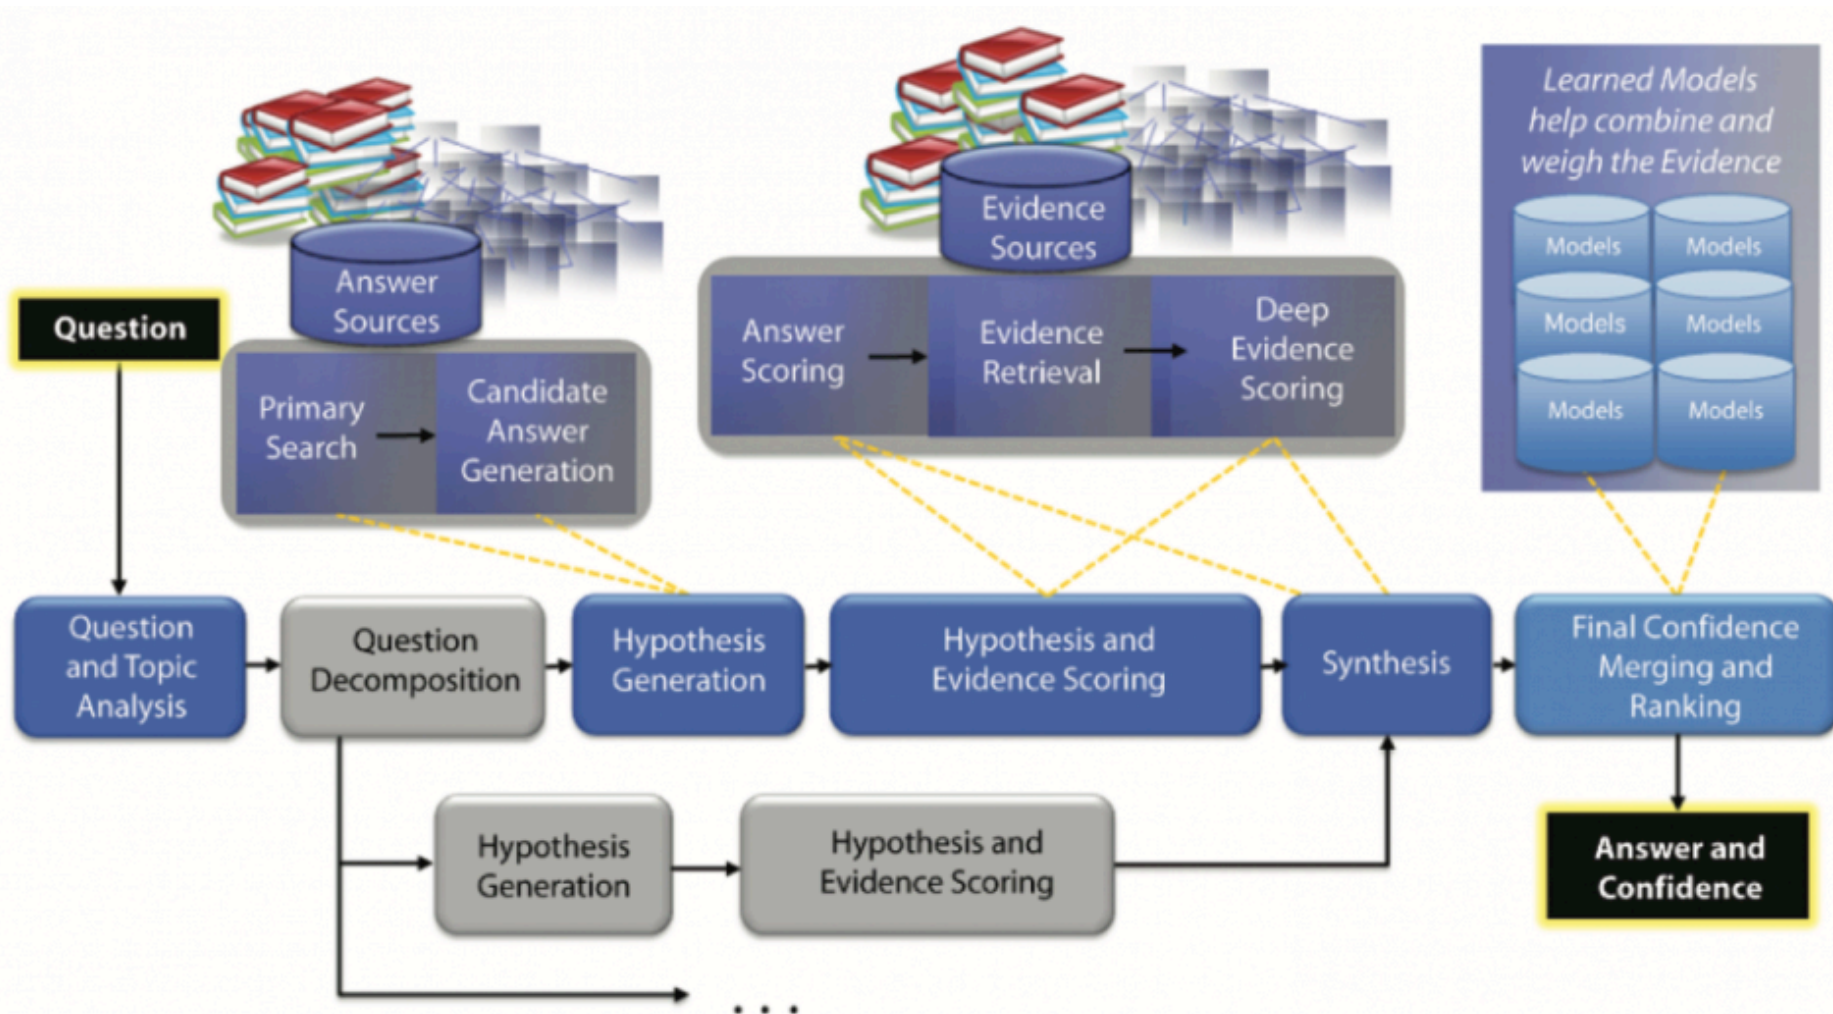
Cognitive Computing




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




Pepper – Animated by Watson



Pepper (Softbank Robotics) and Watson at CES 2016

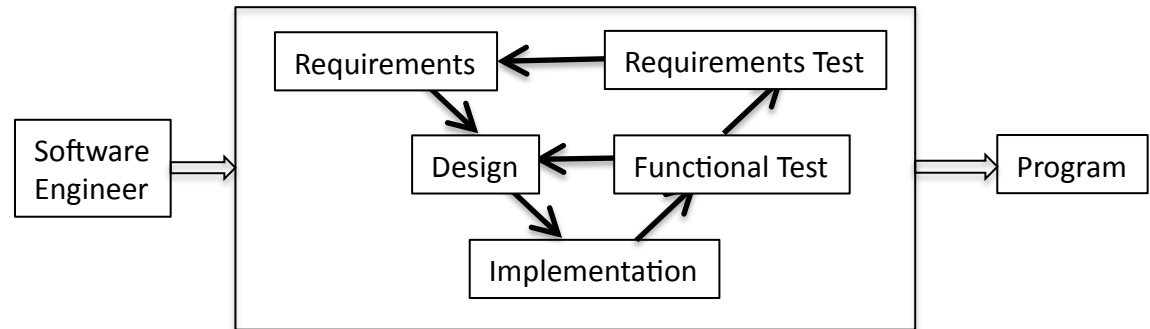
AI as a Service on the Cloud




Machine Learning :

A new paradigm for building systems

Classic Software Engineering:



Machine Learning:



With Machine Learning, the data becomes the code.


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

Potential Innovations from AI

What problem domains are most suitable for (vulnerable to) rupture from AI technologies?

What new paradigms are possible for HCI?

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 People:

=> Commerce, Education, Entertainment, Well Being

Interaction with the Physical World:

=> Robotics, Transportation, Manufacturing, ...

Interaction with Systems:

=> Virtual Personal Assistant, Smart Buildings...

Interaction with the People

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.

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.

Collaborative Intelligent Assistant

An artificial collaborative system that

- Assists an user at complex, critical tasks.
- Provides assistance and advice on demand
- Monitors actions, attention and awareness
- Anticipates problems, Warns of errors
- Can perform tasks on request


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

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

Interaction at Multiple levels

Intelligent collaboration requires interaction at multiple levels

Intentional: Shared goals and objectives
Operational: Dynamically selected actions
Spatial: Situation awareness
Sensory-motor: Sensor and motor signals.



Example: Intelligent Drivers Assistant



- Intentional: Shared navigation objectives.
(Where are we going?)
- Operational: Current route and vehicle state.
(How can we get there? Constraints?)
- Spatial: Spatio-temporal Situation Awareness
(nearby vehicles, routes, obstacles)
- Sensori-motor: Steering, acceleration, scan path (visual attention).

Collaborative Intelligent Systems

Collaboration requires mutual understanding

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: AI systems that provide solutions that can be understood by people.

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 a cognitive AI system.
- 3) Explain why one network architecture is better than another.

Explainable AI – Personal View

Explanation: a reason for deviation from expectation.

Humans provide “narratives” as explanations.

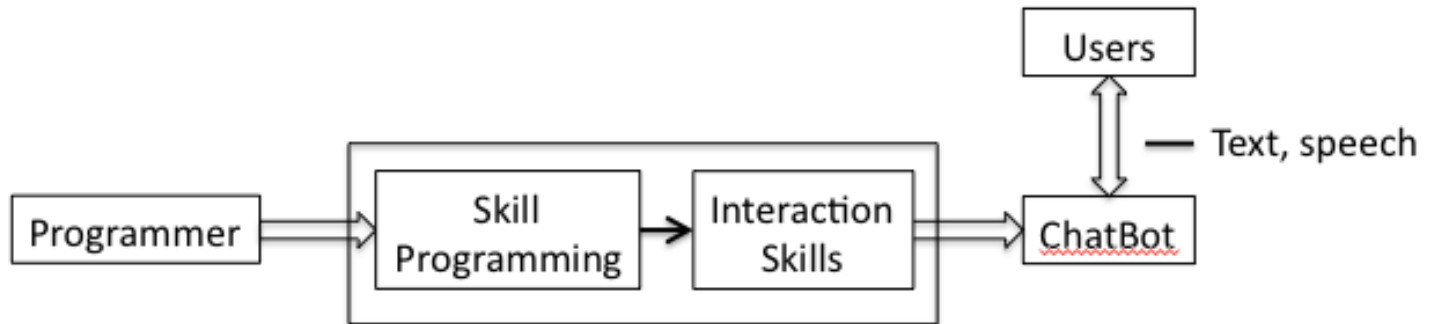
A narrative interprets a sequence of events as a story, placing the events in a context.

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

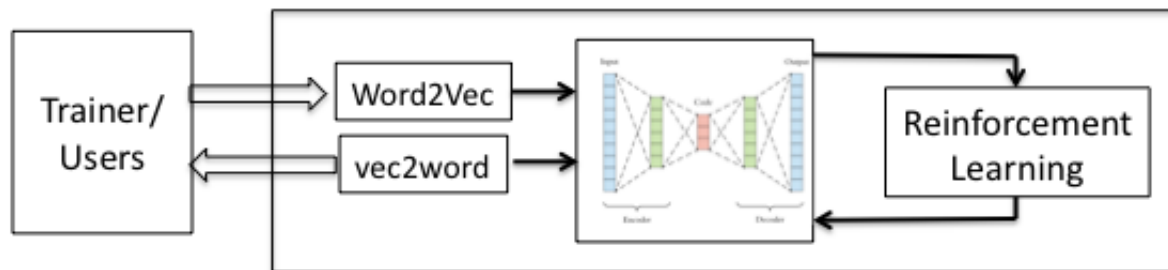
Explainable AI will require the ability to generate credible narratives to explain actions, decisions and consequences of systems.

Explainable AI – Two Approaches

Preprogrammed Explanation - (Amazon Skills, etc)




Training Explanations through interaction




Certi fiable 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

Pop Quiz

Q1: Artificial Intelligence is?

- a) A marketing slogan invented by IBM
- b) Science Fiction
- c) Human level performance at Interaction

Pop Quiz

Q2: Deep Learning is?

- a) A deep-sea diving experience
- b) An advanced degree program
- c) A rupture technology for innovation made possible by high-performance computing and planetary scale data.

Pop Quiz

Q3: Back Propagation is?

- a) A forest-fire control technique.
- b) A breeding technique for genetic selection.
- c) Distributed gradient descent for estimation of neural network parameters.

Pop Quiz

Q4: Innovation with AI requires?

- a) Planetary scale data
- b) High-performance parallel computing
- c) Machine Learning Architectures
- d) All of the above

Pop Quiz

Q5: AI Systems built with Deep Neural Networks...

- a) can explain their decisions.
- b) are easy to certify.
- c) Are trustworthy.
- d) None of the above

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, systems and the world.
- Explainable, Certifiable and Trustworthy AI systems are open challenges.

Artificial Intelligence For Human-Computer Interaction

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