The Origins of Artificial Intelligence and the Emergence of Collaborative Intelligent Systems

James L. Crowley
Professor Emeritus, Grenoble Institut Polytechnique
MIAI Chair on Collaborative Intelligent Systems
Univ. Grenoble-Alpes

Avant Propos: Science (T. Kuhn 1962)

Science:

The elaboration of **theories** and **models** that **predict** and **explain**.

Process: observe, document, hypothesize, predict, and test.
Truth is provided by experiments.

Science is a **social activity**.
Scientific research is conducted by **communities** that share concepts, problems and problem solutions (**paradigms**).

Scientific communities are born, grow, mature, age, decline, and die.

Part 1: The Origins of Artificial Intelligence

A scientific community devoted to Artificial Intelligence (AI) was created in the 1950s.

After a euphoric period in the 1980s, AI was declared “dead”.

Since 2010, the popular media increasingly claim that we are in an AI revolution.

What changed between 1980 and 2010?

Alan Turing: The Birth of Computer Science

Alan Turing (1912 – 1954)

1936: **The Turing Machine –** an abstract model for a universal computing machine. The foundation of Computer Science.

1940-45: The **Ultra Machine** (Bletchley Park). Decoded the unbreakable ENIGMA encoder of Germany.

1950: A practical definition for **Intelligence**

What do we mean by **Intelligence?**

The Turing Test: An imitation game

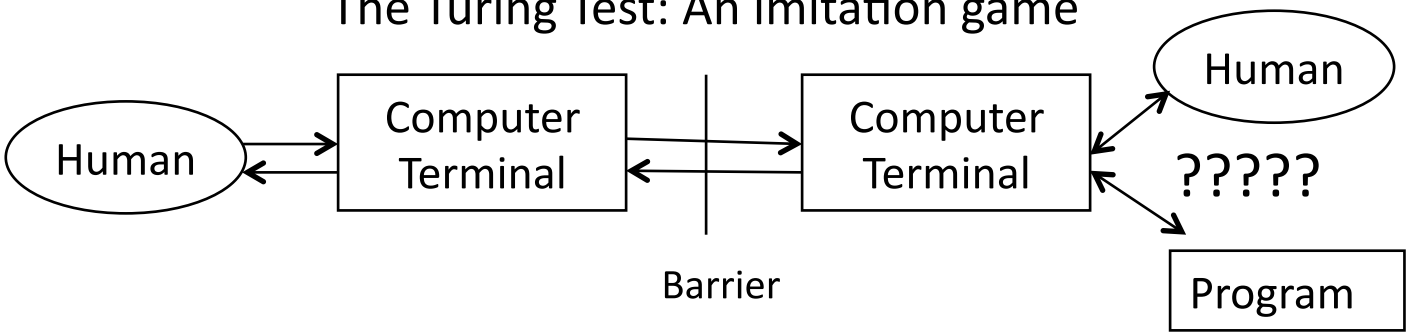

Intelligence according to Turing:
Human-level performance at (text-based) interaction.

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

What do we mean by **Intelligence?**

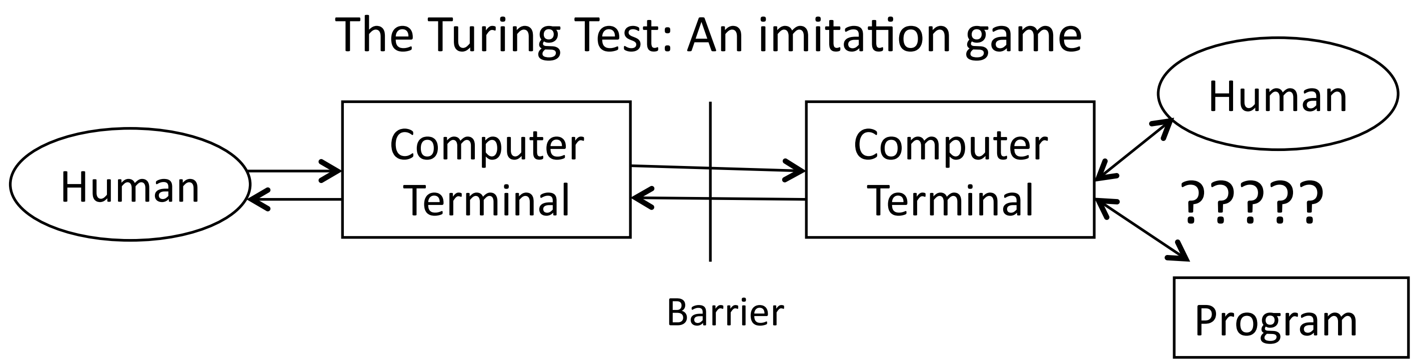

Modern technologies allow us to extend Turing's definition to tasks requiring perception, action, communication or interaction.

Intelligence: Human-level performance at tasks requiring perception, action, communication or interaction.

AI as a Modern Scientific Discipline

Arthur Samuel

John McCarthy

Marvin Minsky

Herbert Simon

Allen Newell

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.

Three schools dominated the field: MIT, Stanford and CMU

1965-1980: Three Approaches to AI


Marvin Minsky


John McCarthy

MIT: **Intelligence as Reasoning**
Universal Reasoning algorithm and a little bit of knowledge. Example: Logic Programming using Theorem Proving (Prolog)


Edward Feigenbaum


Bruce G. Buchanan

Stanford: **Intelligence as Knowledge**
Large knowledge base aided by weak reasoning. Example: Expert Systems (Mycin)


Herb Simon


Alan Newell

Carnegie-Mellon: **Intelligence as Cognition**
Study and imitate models of Human Intelligence
Example: Multi-agent Systems (Hearsay)

Expert System Design Process (1980)

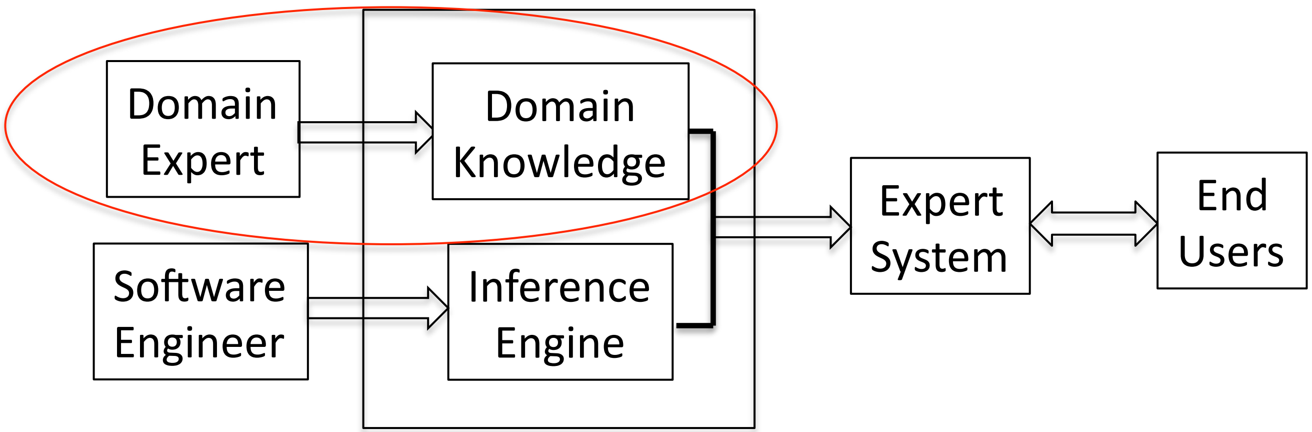

Example: MYCIN – Antibiotic Therapy Advisor (Feigenbaum et al 1980). Domain expert worked with Software Engineer to build system.

Fundamental Problem:

Prohibitive cost of generating domain knowledge.

Evolution of Artificial Intelligence

Dominant Paradigms for Artificial Intelligence:

- Pre-1960: Automata and Pattern Recognition
- 1960-1985: Planning, problem solving
- 1975-1990: Expert systems, symbolic reasoning
- 1985-2000: Logic programming, theorem proving
- 1995-2010: Bayesian methods, Semantic Web

Three Fundamental Barriers to AI:

(1) Insufficient Labeled Data for Supervised Learning.

(2) Insufficient Computing Power.

(3) Prohibitive Cost of Encoding Domain Knowledge.

The Hebbian Alternative: Spreading Activation


Donald Hebb (1904-1985)

information processing as propagation of energy.

Neurophysiological Postulate: When cell A repeatedly excites cell B a metabolic change takes that enables learning.

“Cells that fire together, wire together.”

D. O. Hebb, *The organization of behavior*. Wiley, New York, 1949

The Perceptron (1943)
An Alternative to the Turing Machine

Walter Pitts (1923 – 1969)

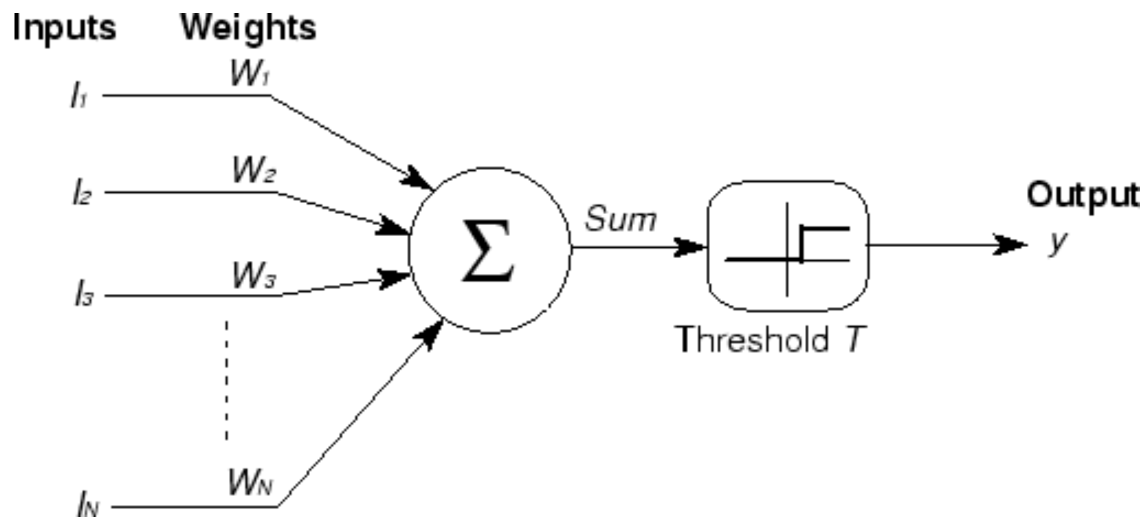

The Perceptron (W. McCulloch and W. Pitts 1943)
A machine that learns from its mistakes.

Warren McCulloch (1898– 1969)

McCulloch , W., and W. Pitts. "A logical Calculus of Ideas Immanent in Nervous Activity. Bull. Mathematical Biophysics, Vol. 5." (1943).

The Perceptron: A Machine That Learns

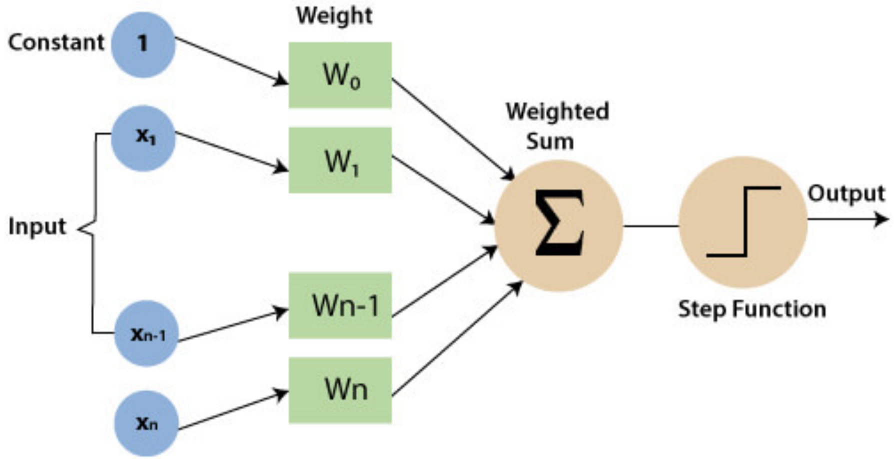

Perceptron: A supervised learning algorithm for a linear decision surface.

Problems: (1) Could only recognized pre-learned patterns
(2) Could only give a binary answer (yes/no)
(3) Required labeled training data for supervised learning.
(4) Required linearly separable properties for classes.

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

Frank Rosenblatt and the Perceptron (1957)

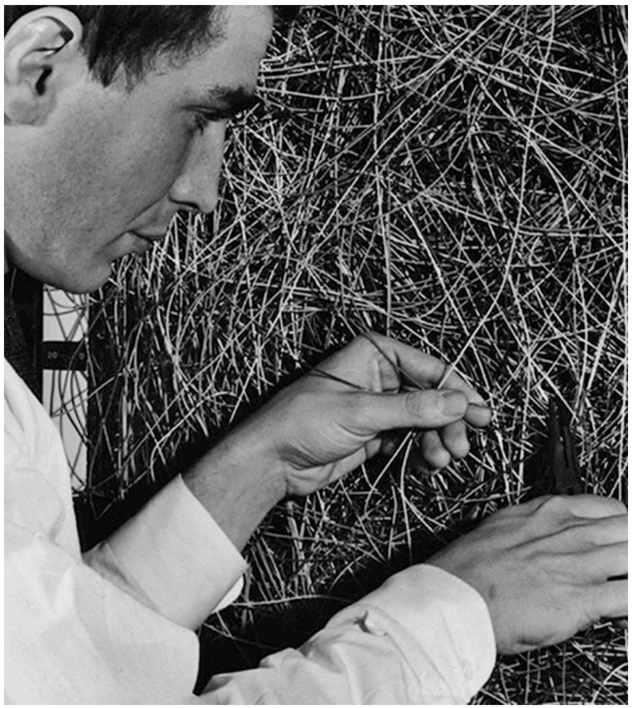

Frank Rosenblatt (1928 – 1971)

In 1957 Frank Rosenblatt at Cornell University, constructed a Perceptron and demonstrated it for the Press.

Rosenblatt claimed that the Perceptron was "the embryo of an electronic computer that will be able to walk, talk, see, write, reproduce itself and be conscious of its existence."

Journalists called it the electronic brain. Serious scientists were horrified.

Minsky and Papert: Perceptrons

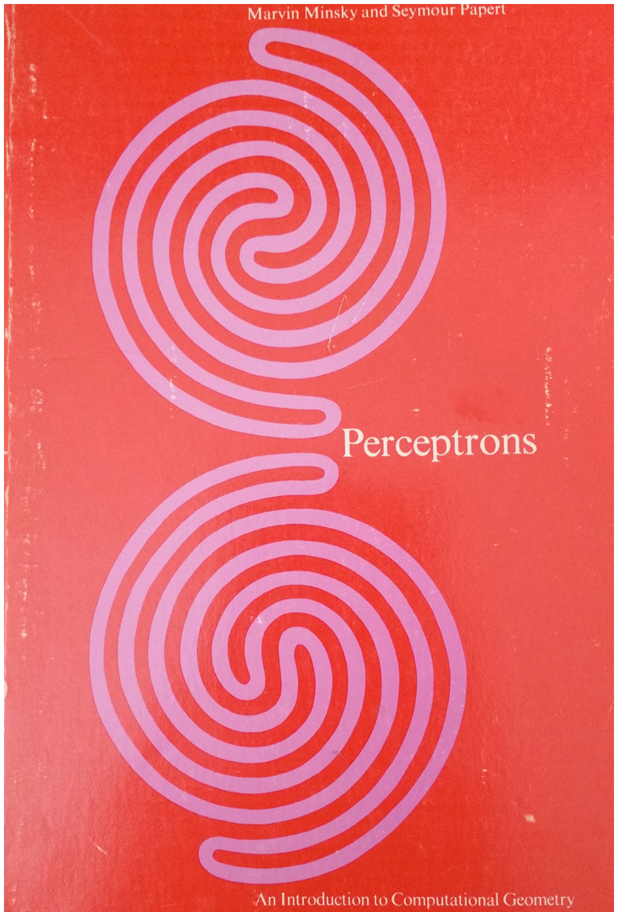

Marvin Minsky and Seymour Papert (MIT 1969)

In 1969 Minsky and Papert documented the many limitations of the Perceptron.
Scientific research on Perceptrons became disreputable.

M. Minsky and S. Papert. *Perceptrons. An Introduction to Computational Geometry*. MIT Press, Cambridge, Mass., 1969

Artificial Neural Networks (1975-1990)
A Hebbian Alternative to Turing Machines

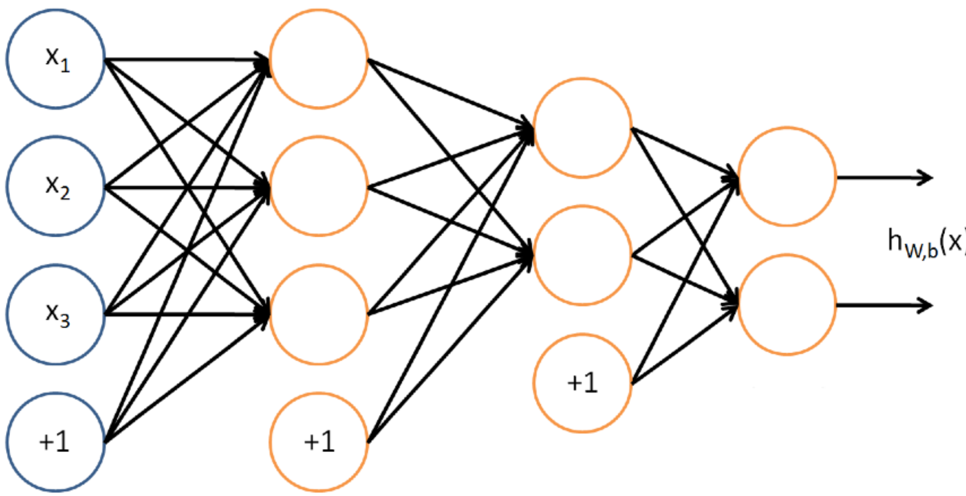

Hebbian Computation: Energy flow through a network.
Provided a simple alternative to symbolic computing

Artificial Neural Networks (1975-1990) – Two innovations

1) Multi-layer perceptrons with soft decision surface
2) Learning with Back-Propagation (distributed gradient descent).

The Hebbian Alternative: Spreading Activation

Geoffrey Hinton

David Rummelhart

Derived Back-propagation algorithm for distributed gradient descent, providing a scalable, universal learning algorithm for training networks of Perceptrons

Rumelhart, D.E., Hinton, G.E. and Williams, R.J., 1985. *Learning internal representations by error propagation*. California Univ San Diego La Jolla Inst for Cognitive Science.

Artificial Neural Networks

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

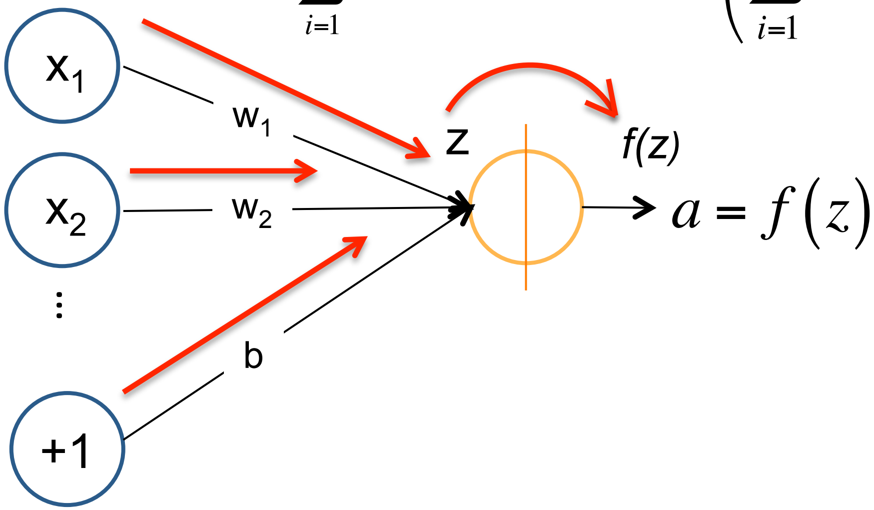

Decision: Sigmoid function

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

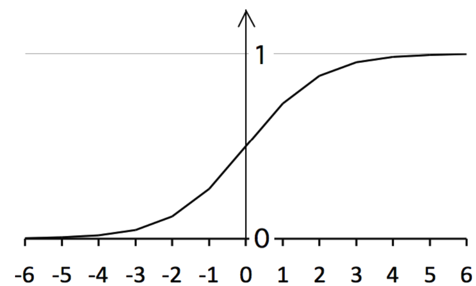

$$\frac{d\sigma(z)}{dz} = \sigma(z)(1 - \sigma(z))$$

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 $\{y_m\}$

$$C(\vec{X}_m, \vec{w}) = \frac{1}{2}(a_m^{out} - y_m)^2$$

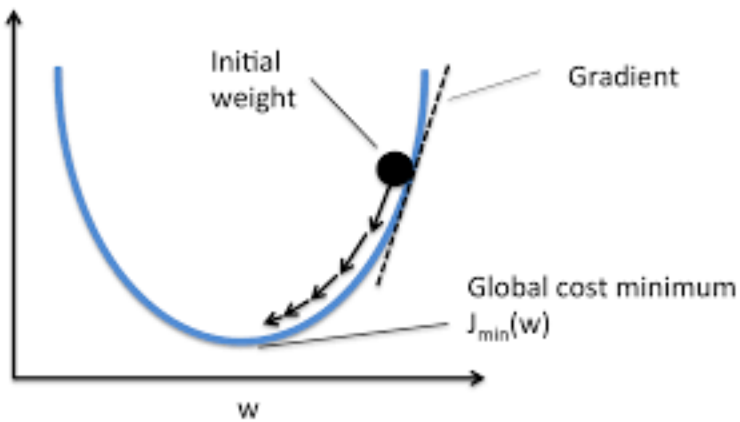

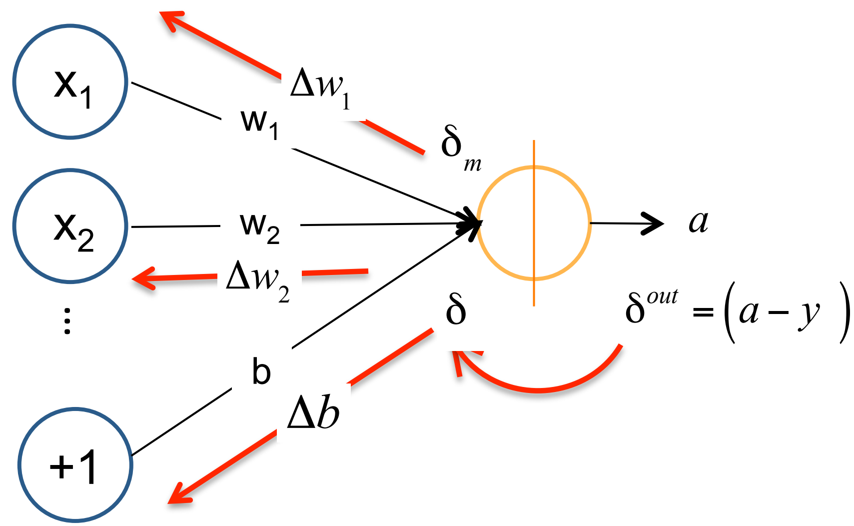

Back Propagation: A first-order iterative optimization algorithm for finding the minimum of a function.
Error flows back into the network to learn from mistakes

Generalized to Multi-Layer Networks

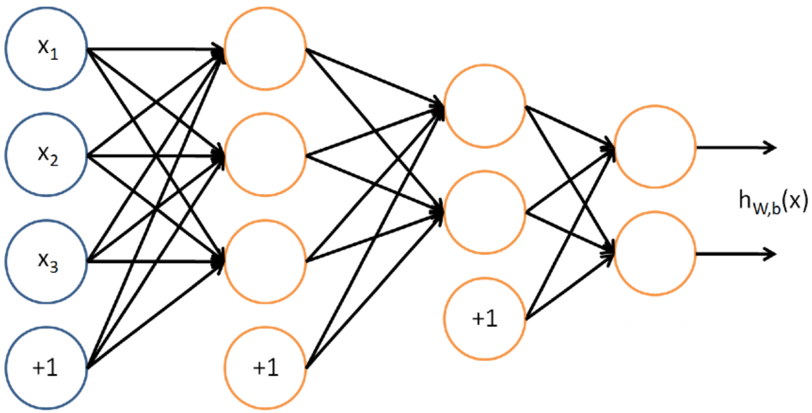

Recursive Feed-Forward calculation

$$\vec{a}^{(3)} = f(f(....f(w_{ij}^{(1)}\vec{X}_i + b_j^{(1)})))$$

Hebbian representation:
propagation of activation energy

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

A neural network is a distributed algorithm using **propagation of activation energy to compute.** enabling arbitrary scale networks using parallel computing.

Generalized to Multi-Layer Networks

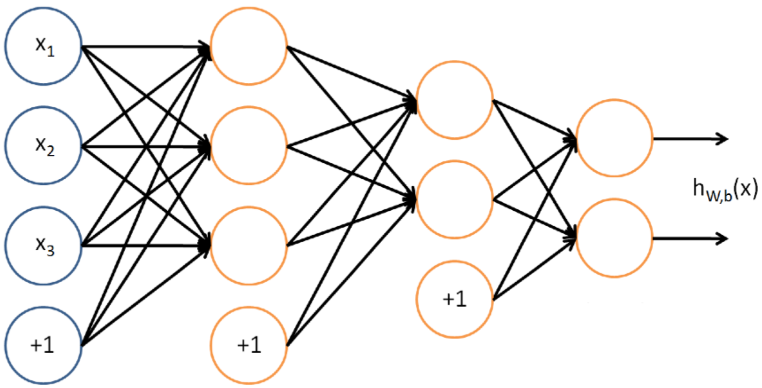

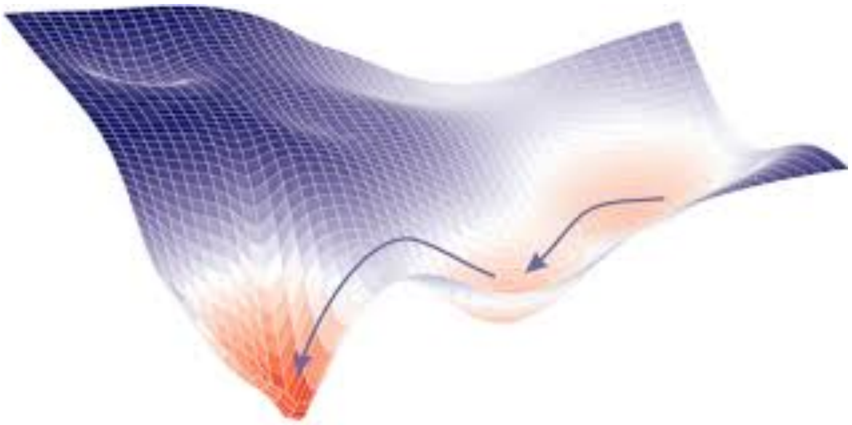

Training requires massive computing with massive data.

Difficulties: (1980's)

- Network has thousands (millions) of parameters
- Training data is very noisy.
- Loss function has local minima
- Results of learning difficult to explain or reproduce

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

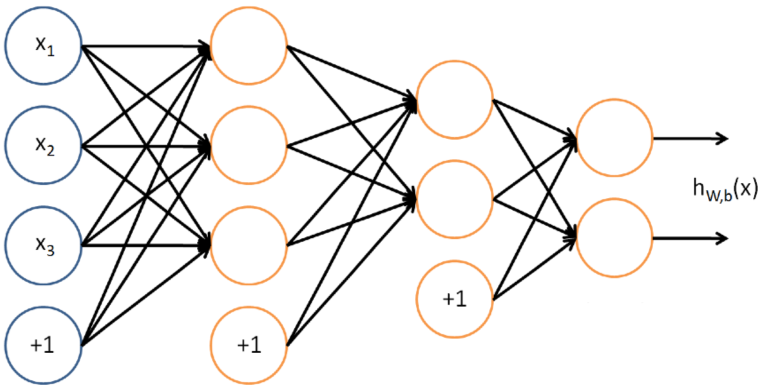

Problems:

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

Enabling Technologies arrived in the 2000's

Overcoming the three fundamental Barriers:

(1) Insufficient training data
- ⇒ Planetary scale data from the internet and the WWW
- ⇒ Data from realistic simulations

(2) Insufficient computing power
- ⇒ Moore's Law, GPUs, massively parallel computing

(3) Prohibitive cost of encoding knowledge
- ⇒ Self Supervised Learning

Le Net5 - 1994

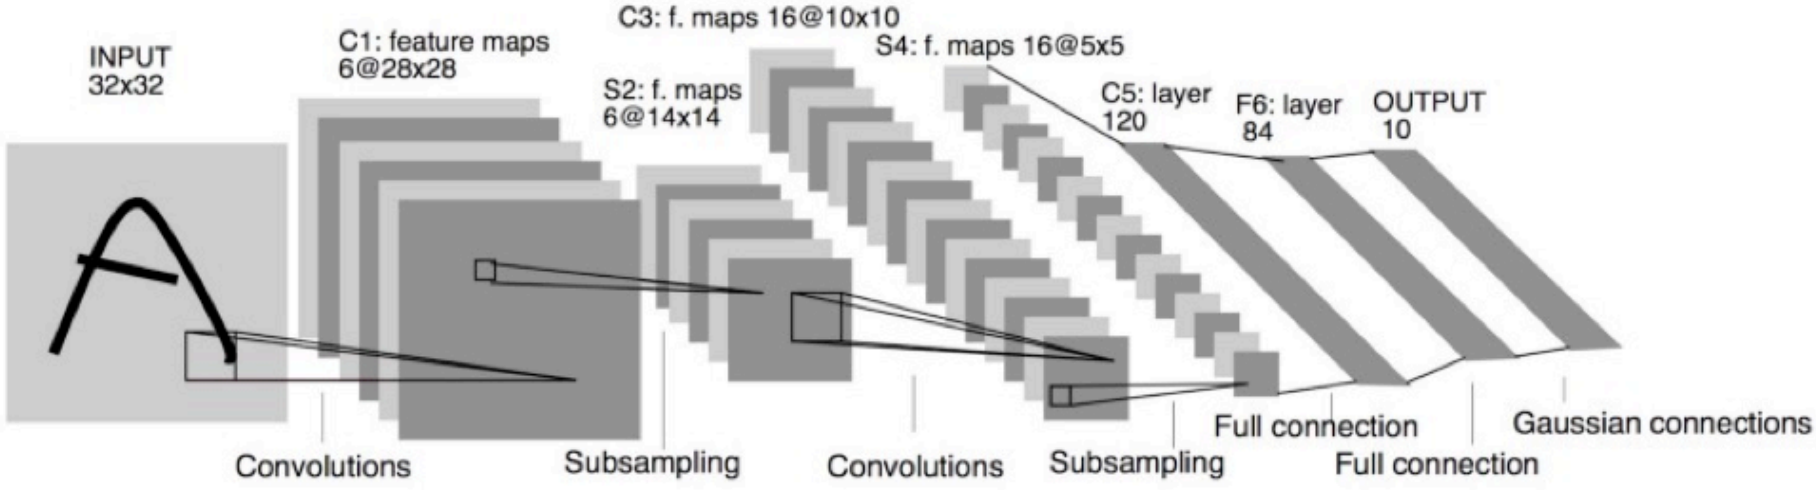

7-level convolutional network by Yann LeCun in 1998.
State of the art for recognizing hand-written numbers on checks.

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

AlexNet 2012

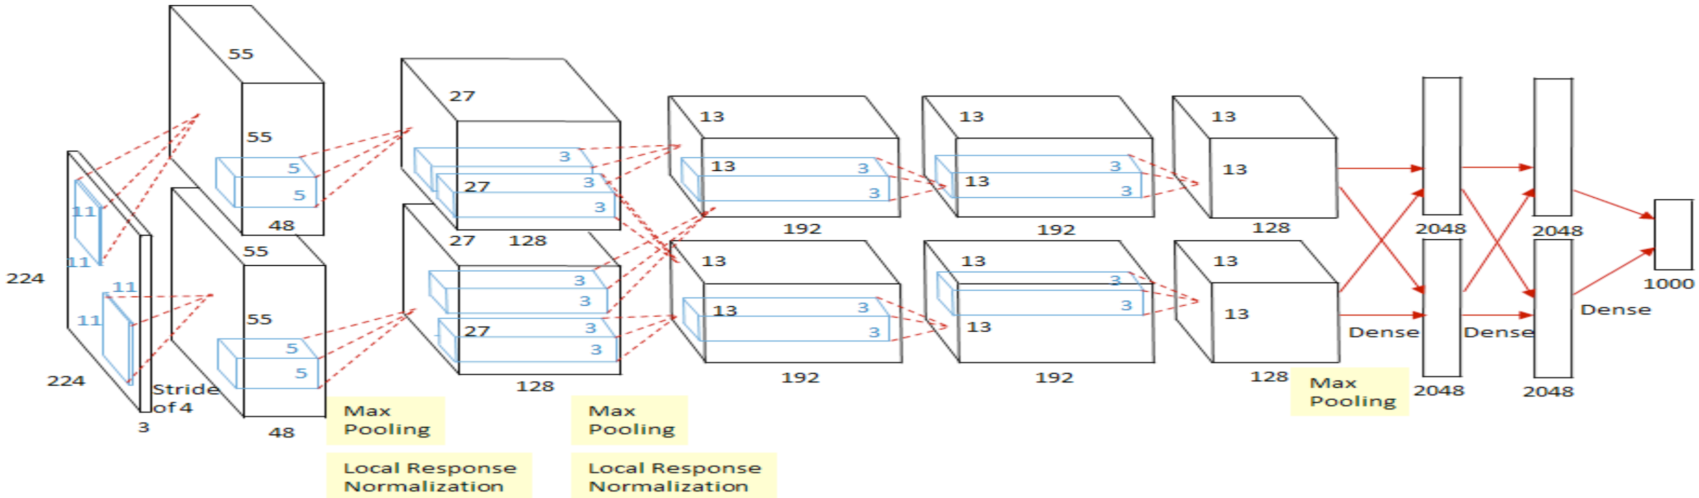

Created by Alex Krizhevsky and Geoff Hinton (based on LeNet)
Won the ImageNet Large Scale Visual Recognition Challenge in 2012 by a large margin with an error of around 15%

Triggered a paradigm shift for Computer Vision, Natural Language processing, 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

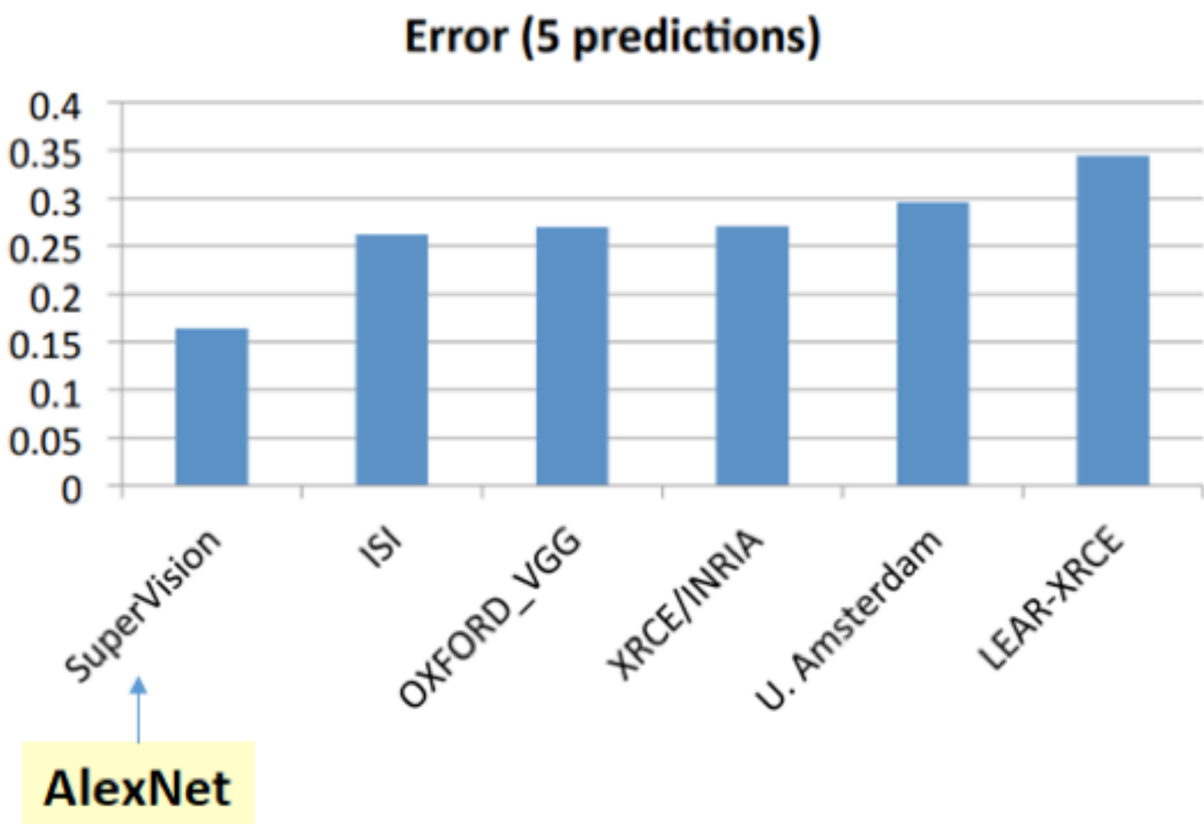

Artificial Neural Networks: Trainable Universal Function Approximation

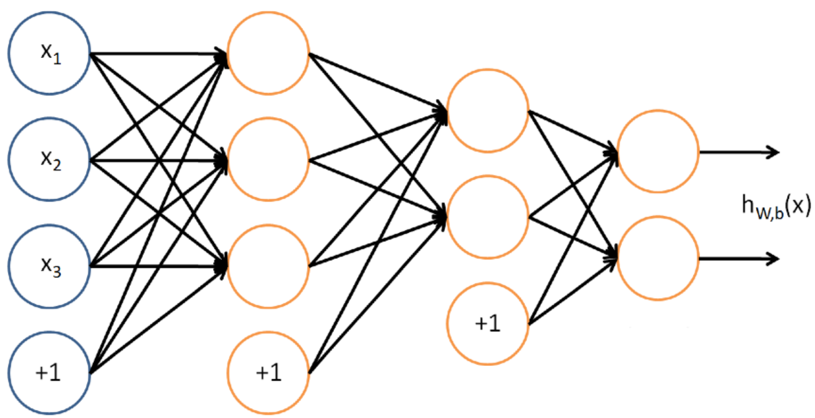

Recursive Feed-Forward calculation

$$\vec{a}^{(3)} = f(f(\dots f(w_{ij}^{(1)}\vec{X}_i + b_j^{(1)})))$$

Hebbian representation:
propagation of activation energy

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

A Neural Network is trainable **Universal Function Approximator.**

Generative and Discriminative Networks

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

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!)

Autoencoder

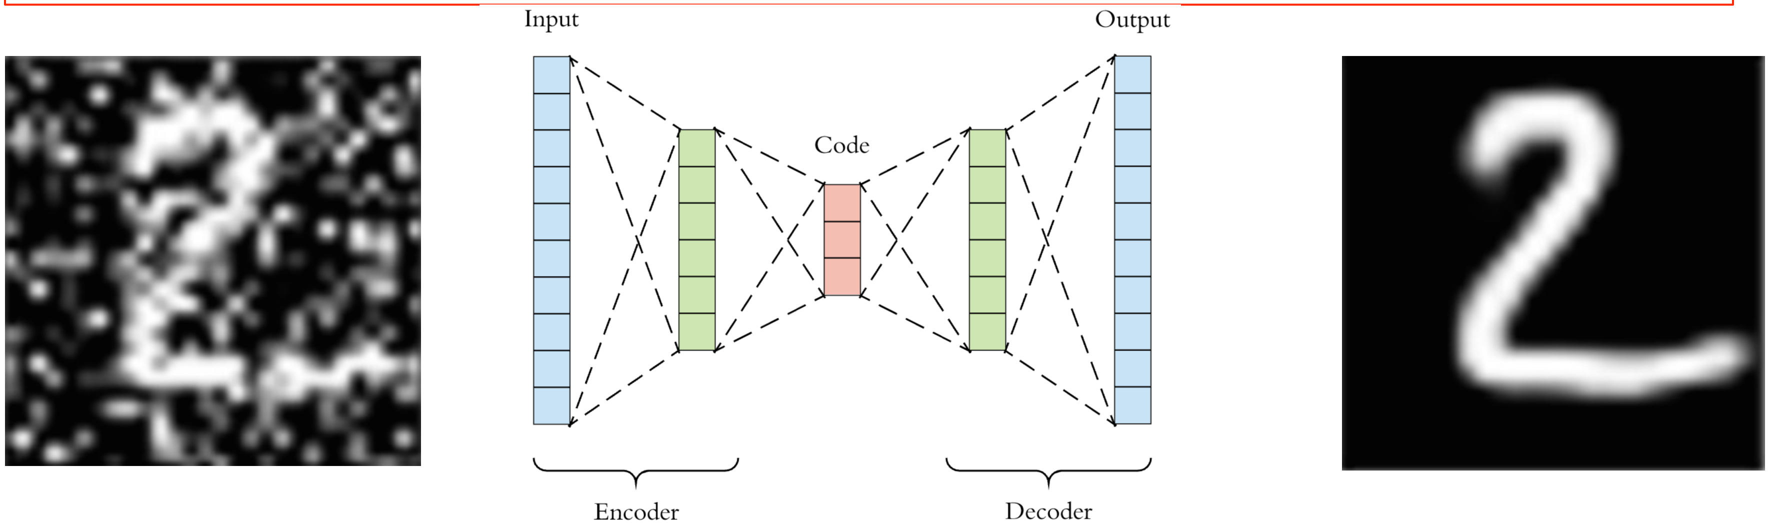

An Autocoder learns to reconstruct (generate) clean copies of data without noise.

Key concepts:

1) Training data is target. Error is difference between input and output
2) Scalable to any quantity of data using Back Propagation
3) Compresses the training data into a minimum number of
   independent hidden units (Code vector)

Used by Sejnowski and Hinton as a means to overcome lack of training data.

AutoEncoders

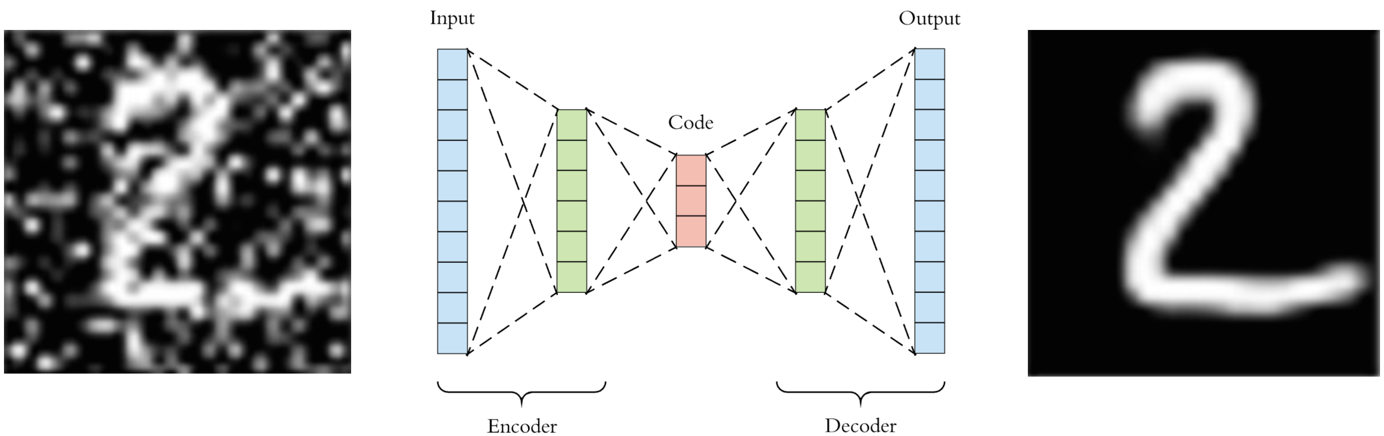

Autoencoders were originally used to compute **principal component analysis**, using least squares reconstruction error.

Adding a information theoretic “sparsity term” to the cost function provides **independent components analysis**, providing **unsupervised learning** of classes from data.

Variational Autoencoder

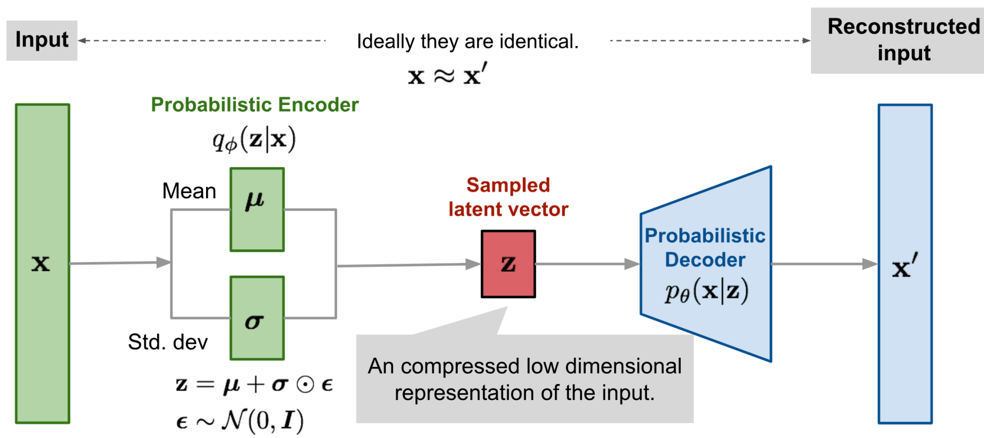

A VAE can be used to generate synthetic output.

Example:

1) Train VAE on dancers doing the same dance.
   => Code represents posture
2) Drive decoder of a dancer from encoder of another.

Auto-encoders Enable Self Supervised Learning

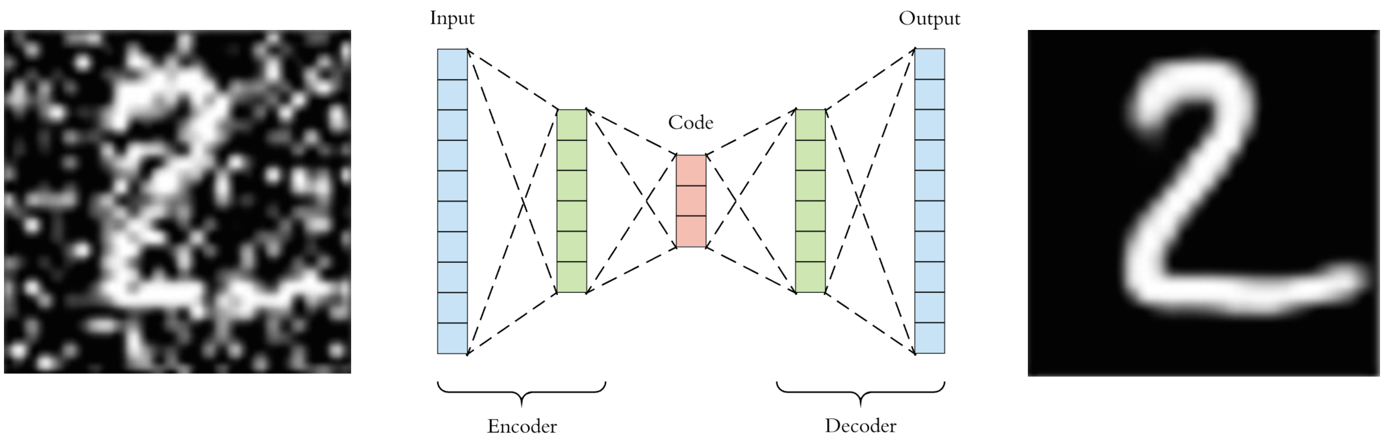

Self-supervised learning is a form of unsupervised learning.
The data its the ground truth.

The system learns to reconstruct missing parts in the data (**missing token replacement**), and to predict the adjacent data (**next token prediction**).

Self supervised learning unlocks all recorded human knowledge (and the internet) to machine learning.

Self Supervised Learning with Transformers

In 2017, a revolutionary paper by Vaswani et al from Google showed that the deep convolutional and recurrent networks using layers of could be completely replaced with stacked auto-encoders using self-attention.

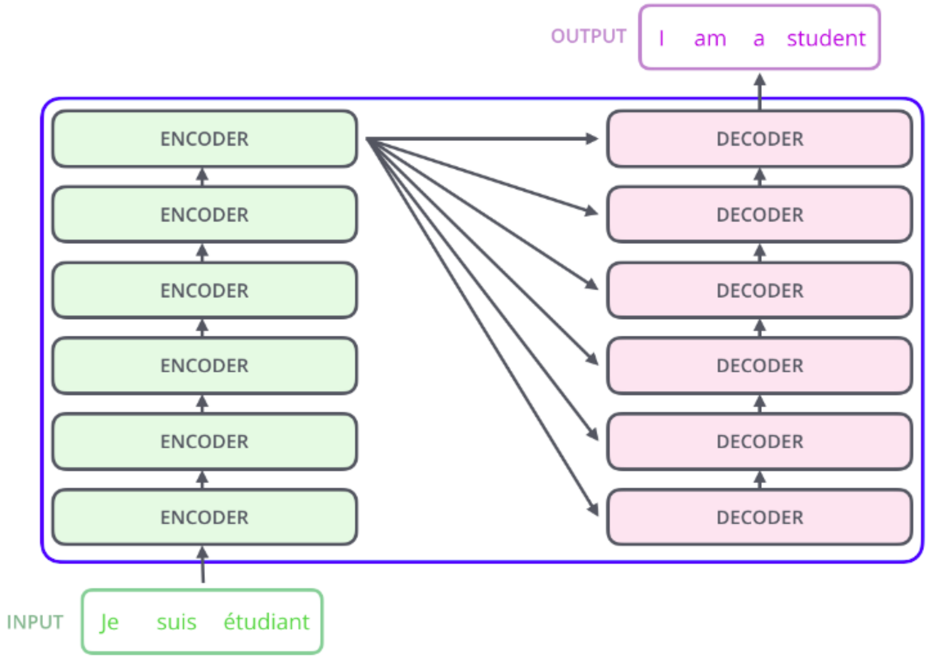

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. and Polosukhin, I. Attention is all you need. 2017

Transformers use attention to associate mutually relevant entities

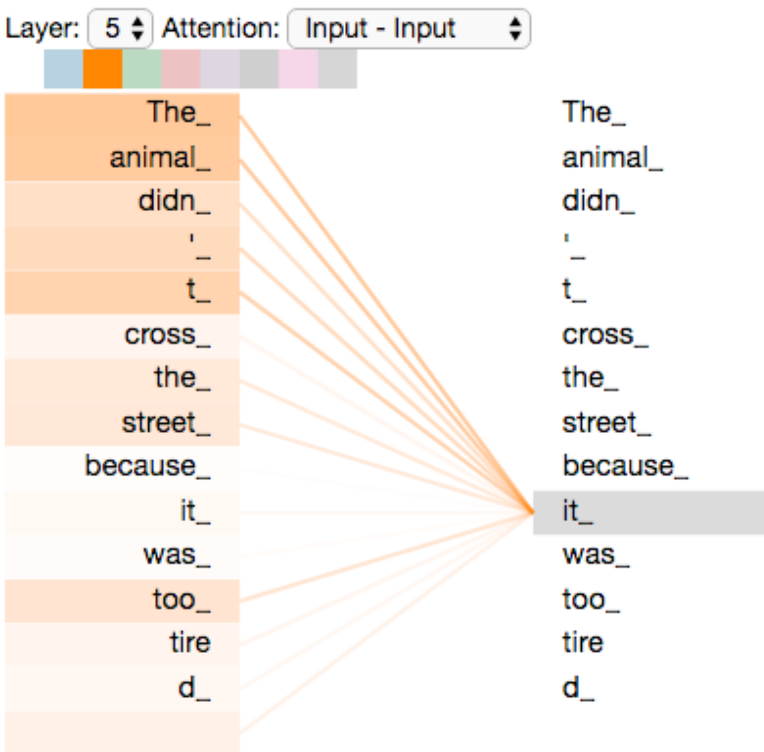

Transformers have become the dominant approach for natural language processing.

Self-attention associates words in a sentence or paragraph in order to provide context for a more abstract representation and establish meaning.

Jay Alammar, The Illustrated Transformer (http://jalammar.github.io/illustrated-transformer/)

Cross-Modal Self-Attention

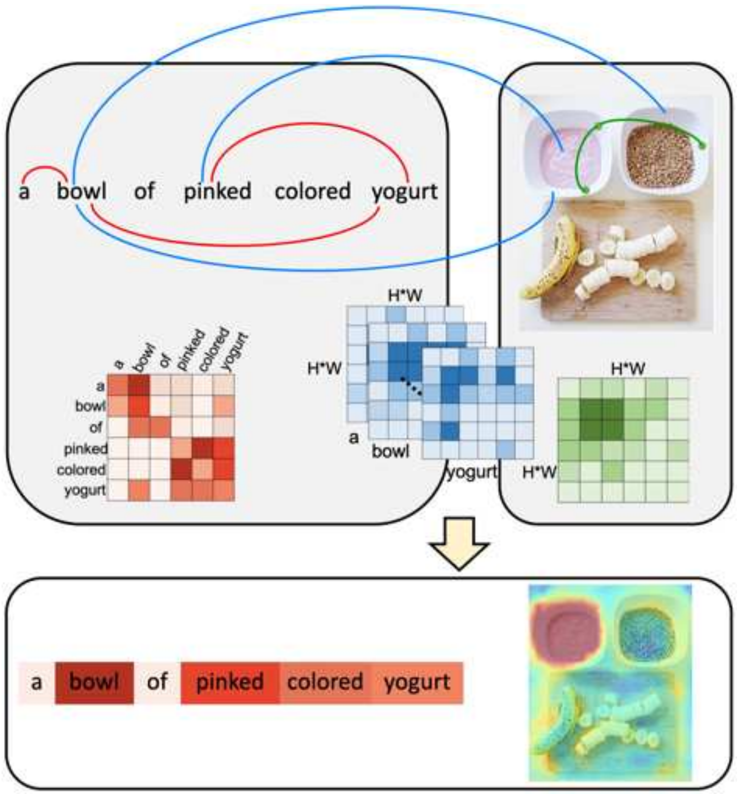

When used with multiple modalities, **self-attention** determines mutually relevant information.

Self-attention can be used to relate words to image patches as well as other words.

From: Ye et. Al. "Cross-Modal Self-Attention Network for Referring Image Segmentation", cvpr 2019, IEEE Conf. on Computer Vision and Pattern Recognition, June, 2019.

Extensions to Vision and Speech

Transformers are rapidly replacing Deep Recurrent Networks and Convolutional networks for **Speech Recognition** and **Computer Vision.**

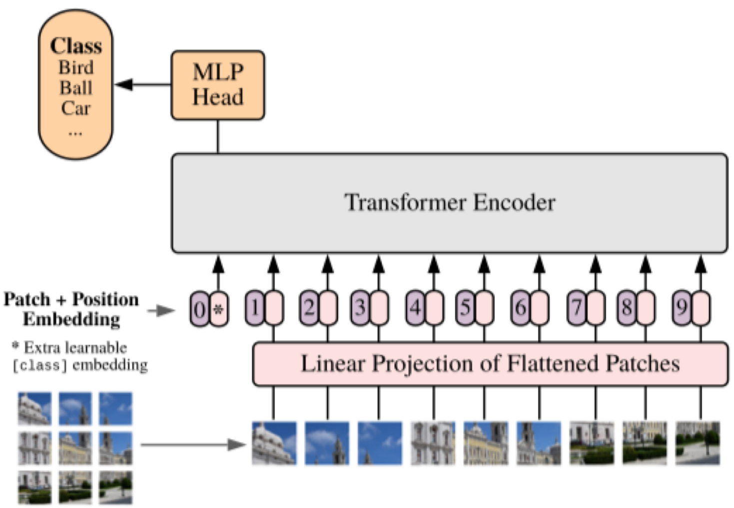

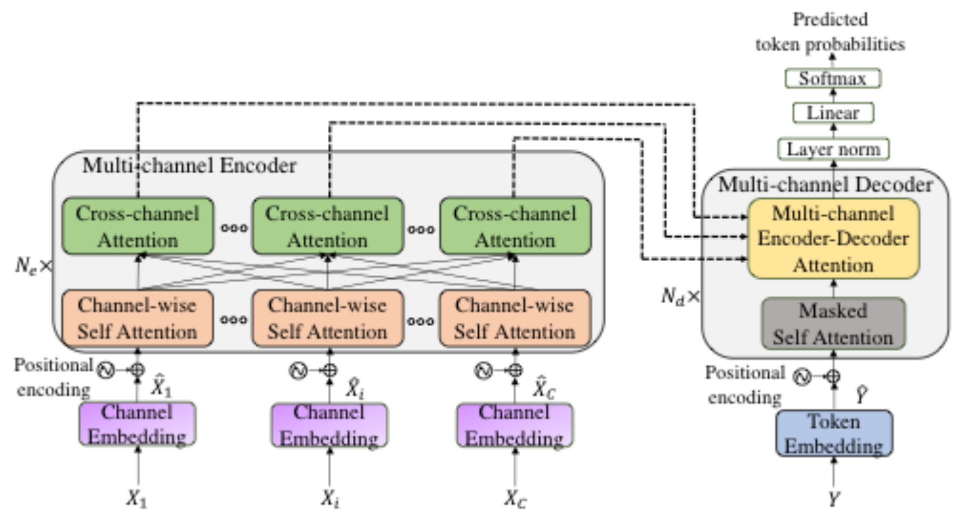

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S. and Uszkoreit, J. An image is worth 16x16 words: Transformers for image recognition at scale. ICLR, 2021

Chang, F. J., Radfar, M., Mouchtaris, A., King, B., & Kunzmann, S. (2021, June). End-to-End Multi-Channel Transformer for Speech Recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5884-5888). IEEE, 2021

Multimodal Perception with Transformers

Recent results indicate that Transformers are well adapted for
Multi-modal Perception, Robotics and Human-Computer Interaction

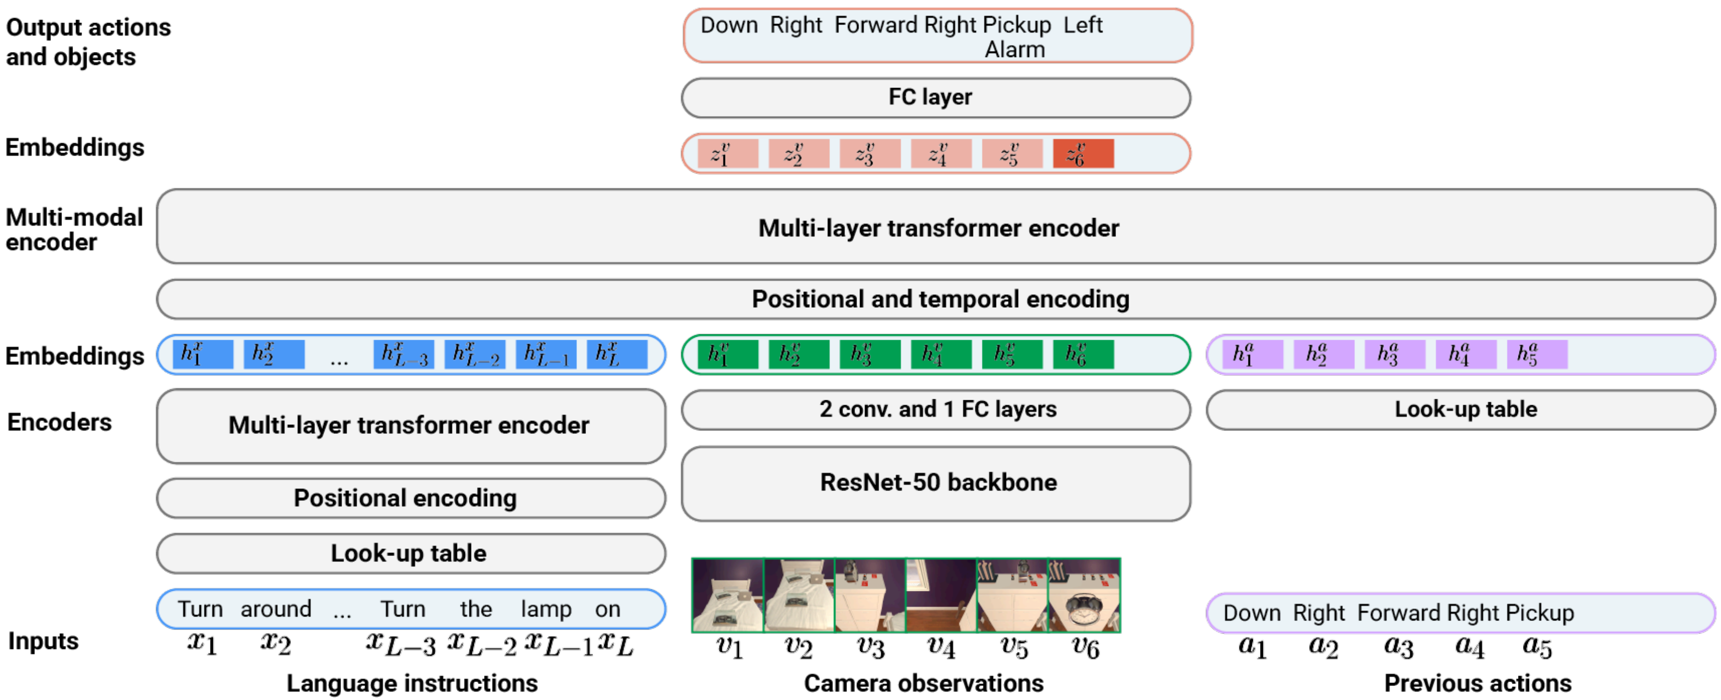

Pashevich, A., Schmid, C. and Sun, C., Episodic Transformer for Vision-and-Language Navigation, Int. Conf. on Computer Vision, ICCV 2021, Oct. 2021.

Part 2: The Emergence of Collaborative Intelligence

Outline:

- Multimodal Interaction
- Categories of Intelligent Systems
- A Hierarchical Framework For Collaborative AI

Multimodal Interaction: What is a modality?

Modality: A channel for perception or action

Examples of modalities:

- Perception: Vision, audition, olfaction, proprioception, …
- Action: Manipulation, Locomotion, Communication, …

Perception, Action, Cognition and Emotion

Perception: Interpretation of sensing through recognition, action, emotion, cognition

Action: Intentional movement to communicate, to sense, or to affect the environment.

Cognition: Abstract reasoning to predict phenomena and determine actions.

Emotion: Intuitive (somatic) reaction to a situation that provides rapid response.

P PERCEPTION

A ACTION

C COGNITION

E EMOTION

Emotion and cognition guide action and perception.

Hypothesis: Emotion guides cognition.

Emotion and Cognition:
Two Complementary Intentional Systems

<u>Emotion</u>:
Fast, Reactive, Predictable.

Enables rapid reaction to threats and opportunities.

Kahneman's System 1?

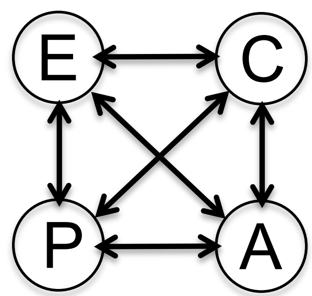

<u>Cognition</u>:
Slow, Deliberative, Creative

Enables planning, explanation, prediction and understanding.

Kahneman's System 2?

Kahneman D, Egan P. *Thinking, fast and slow*. New York: Farrar, Straus and Giroux; 2011

Multi-modal Interaction: Coupled Perception-Action

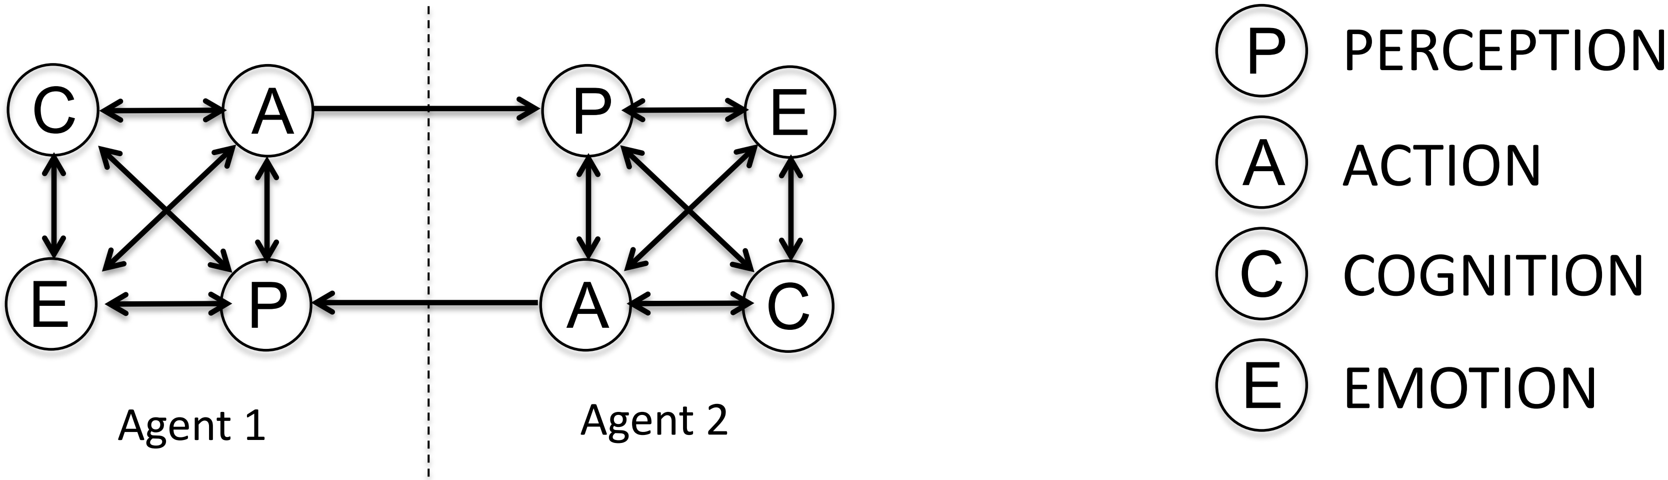

Multi-modal Interaction: Tightly coupled perception-action between two natural or artificial systems.

Categories of Interactive Systems

Tools

Affectors

Media

Advisors

Collaborators

Categories for interactive systems can be defined by the nature of interaction. (Crowley-Coutaz 2015).

Examples: Tools, Affectors, Media, Advisors, Collaborators, …

J. L. Crowley and J. Coutaz, "An Ecological View of Smart Home Technologies", 2015 European Conference on Ambient Intelligence, AMI 2015, Athens, Nov. 2015

Categories of Intelligent Services

Tools: A service used to achieve a goal. The behavior of a tool should be reliable, predictable and robust to environmental conditions.

Example: smart thermostat

Most current work on Human-Computer Interaction addresses the problem of interacting with tools.

Categories of Intelligent Services

Media: Extensions to human perception and experience, for entertainment, communications, and display of information. Can be interactive or simply peripheral, and ideally should provide a sense of immersion. Example: Augmented Reality.

Affectors: Services that inspire affection. Affectors can help compensate for a loss of social contact that can result from ageing or hospitalization.
Examples: Nabastag, Paro Affective Robot, Nao, Jibo,..

Categories of Intelligent Services

Advisors: Propose possible courses of actions. Should be completely obedient. Should not take initiatives or create unwanted distractions (nag-ware).
Examples: GPS Navigation system giving route advice

Collaborators: Intelligent systems that act (or interact) with people to achieve a common goal. Collaborators should share awareness of situations, operational plans, abilities, and practical knowledge (praxis).
Examples: Service Robots, Cobots, Sales Agents.

Beyond Advisors:
Collaborative Intelligent Systems

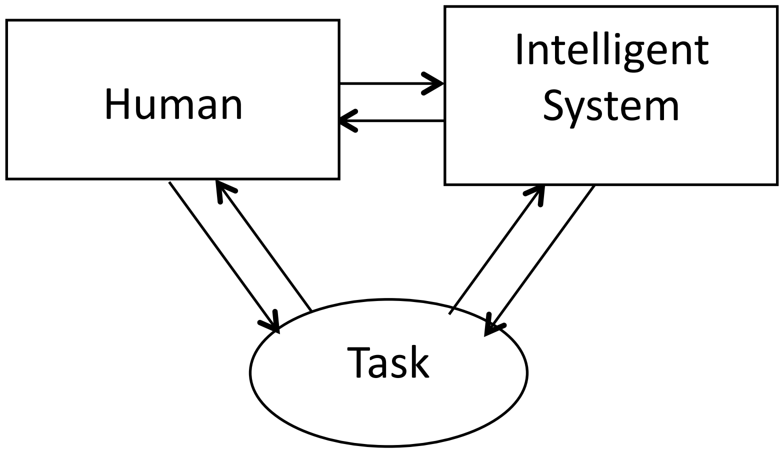

Collaboration is a process where two or more actors (agents) **work together** in order **to achieve** a shared **goal**.

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

A Hierarchical Framework for Collaborative AI

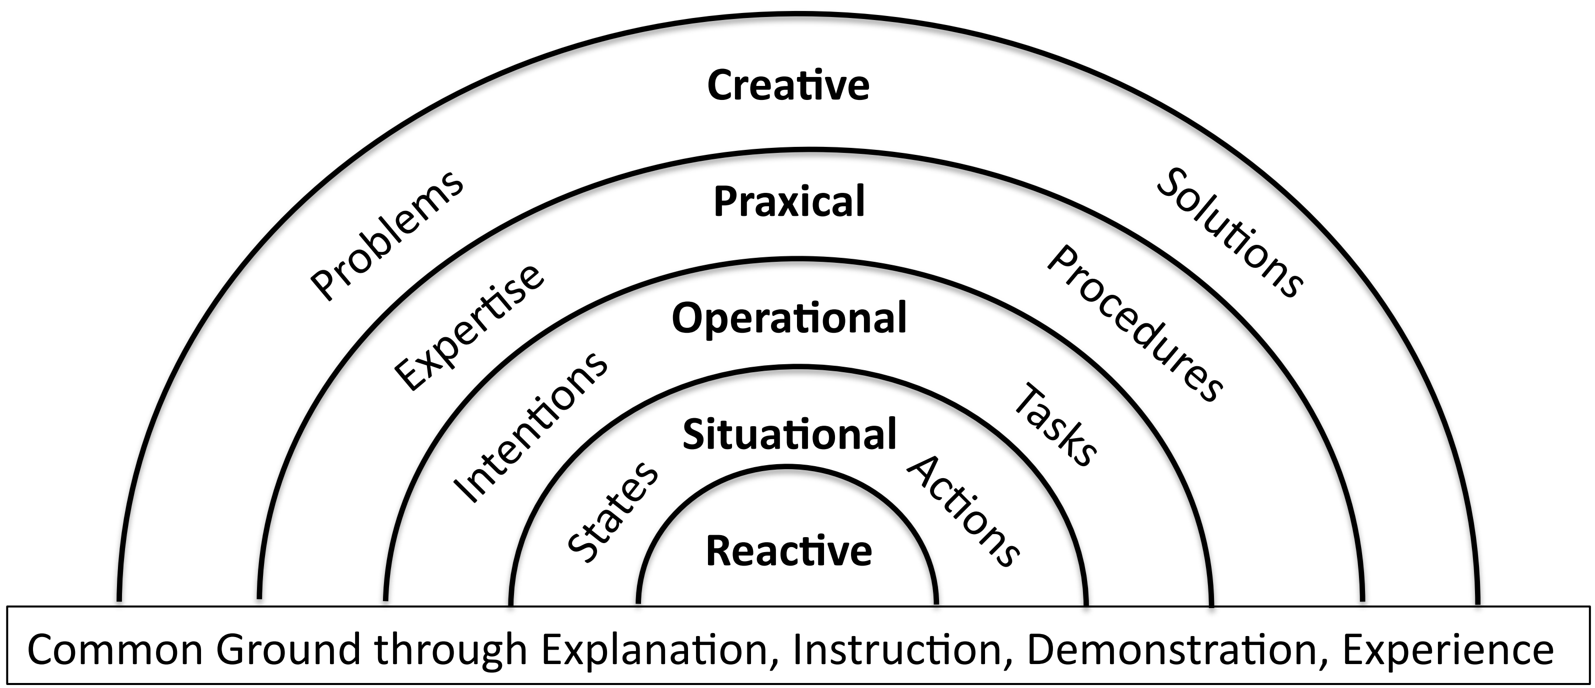

Reactive Collaboration

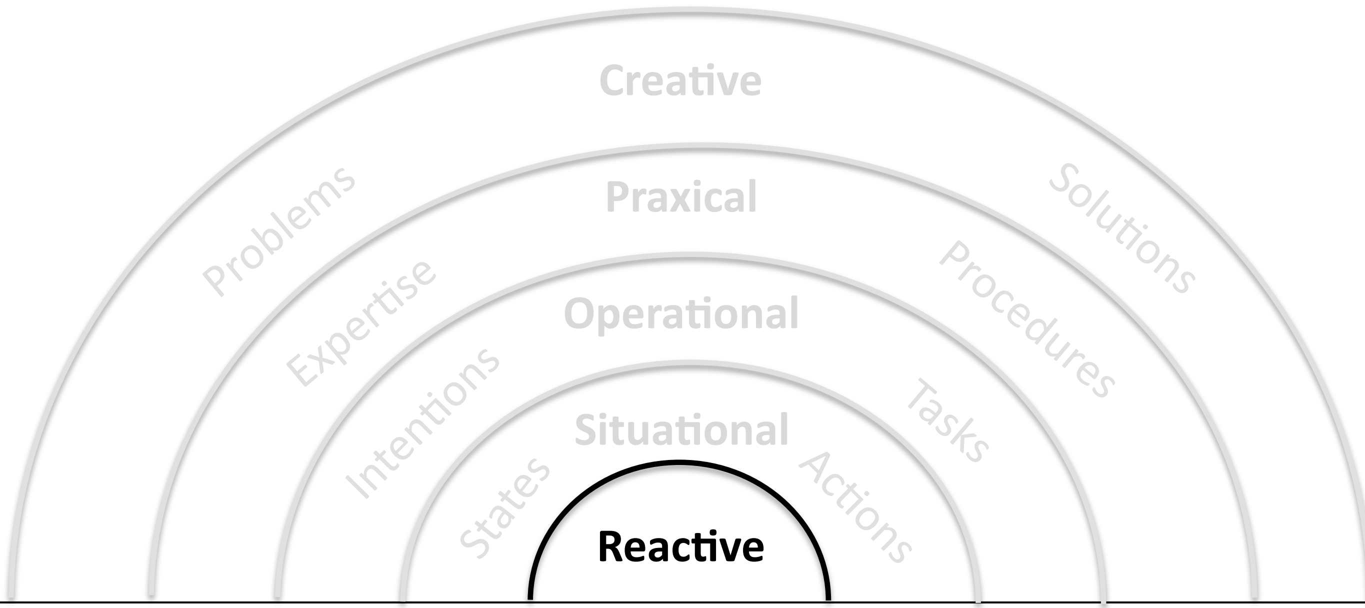

Common Ground through Explanation, Instruction, Demonstration, Experience

Reactive Collaboration

Tightly-coupled perception-action where actions of one agent are immediately sensed and interpreted as actions by the other

Example: Modern Automotive systems for lane following where observation of lane marking lines are translated to steering angle and communicated as control forces on the steering wheel.

Reactive Collaboration: Automotive Control

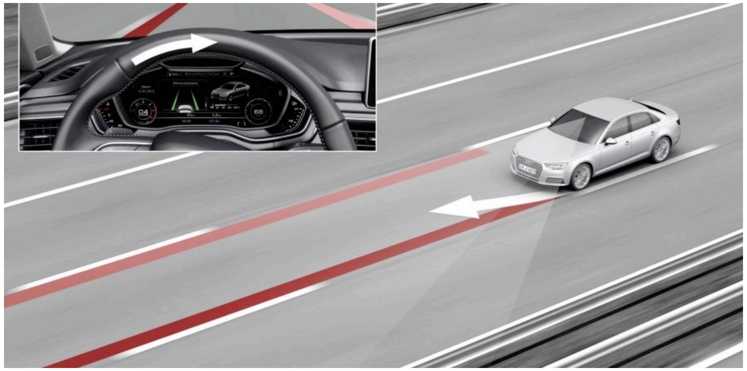


Intelligent Collaborative Automotive systems:

Lane Assist: Lane keeping generates auditory alarms and control forces on steering wheel

Blind spot detection: generates auditory alarms for vehicles that in your blind spot.

Brake Assist: Automatic adaptation of braking force based on time to contact

Reactive Social Interaction with AI

Reactive Social Interaction: Tightly-coupled perception-action where actions of one agent are immediately sensed and interpreted as actions of the other.

Example: Detection of engagement, emotional mirroring, emotional stimulation
Classic Technologies for Social IA

- Detection of engagement from posture, gaze and face expression
- Recognition of Ekman's basic emotions from Facial Action Units and prosody
- Sentiment analysis from language.

Example Research Challenges:

- Learning to recognize social signals from multimodal perception (vision, prosody, posture).
- Learning to imitate social interactions from multimodal perception (vision, prosody, posture) and display (graphical or mechanical animation)
- Learning to recognize and evoke emotion from multimodal perception and display.

Reactive Social Interaction: AI Companions

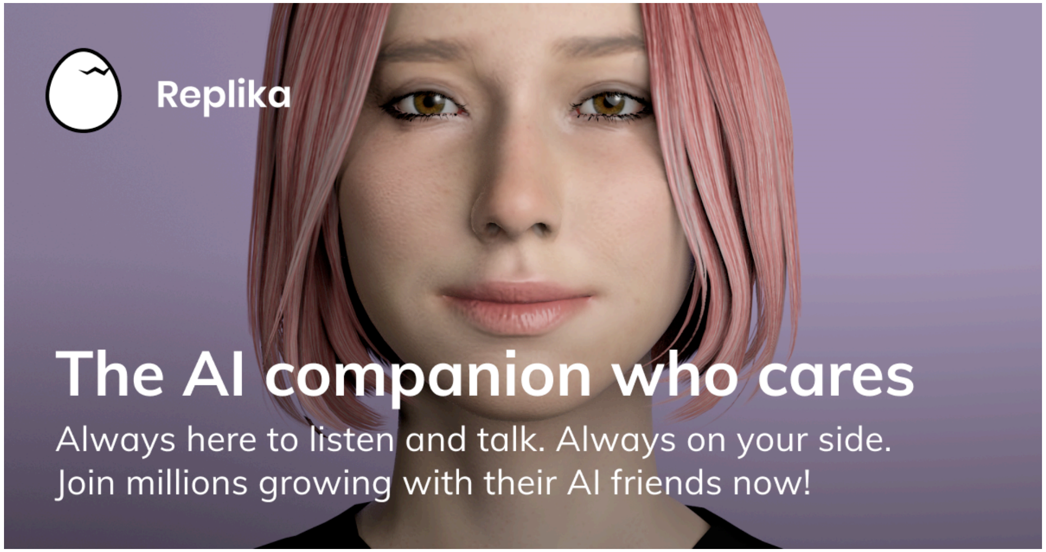

Replika. Based on OpenAI's GPT-3 transformer (trained on wikipedia)

Replika learned a realistic imitation of the language patterns by training on emails and text chat. The reactions are socially appropriate, but the interaction lacks context. There is no understanding of situation.

A Hierarchical Framework for Collaborative AI

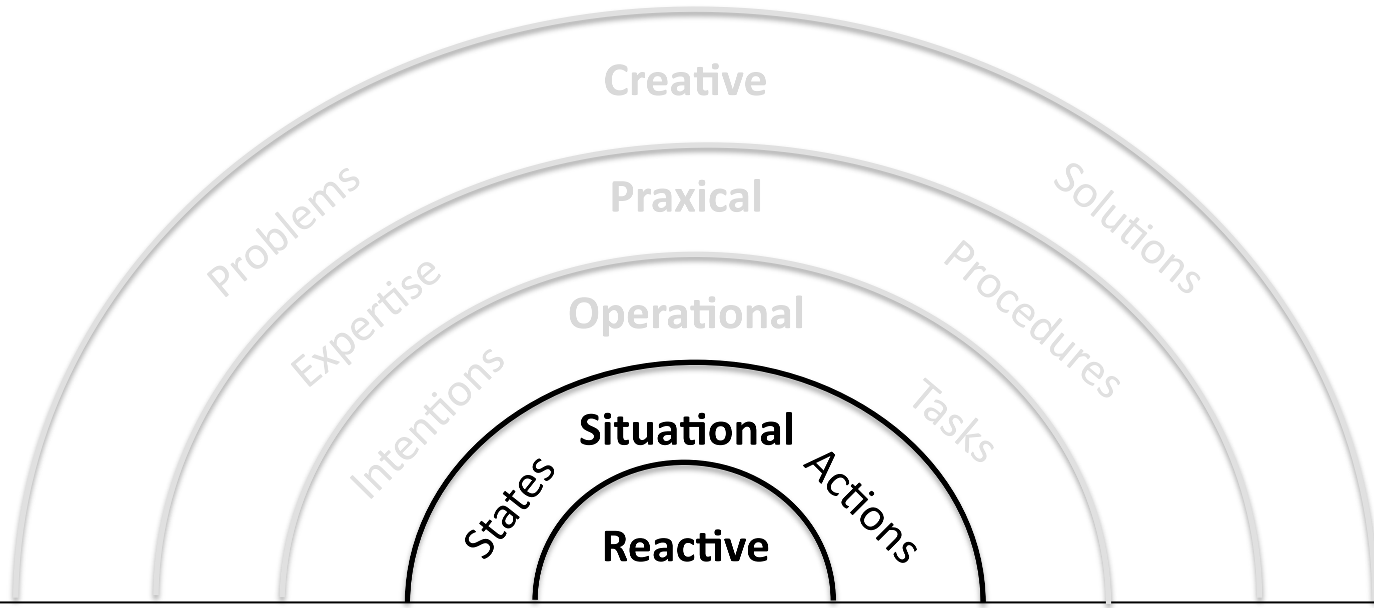

Common Ground through Explanation, Instruction, Demonstration, Experience

Situation-Aware Collaboration

Robby: Have you
Seen my Glasses?

Yes. They are on
the desk in your office.
Would you like me
to get them?

Situation Aware Interaction:
Perception, action and interaction
mediated by shared awareness of situation.

Situated Interaction Theory (Suchman 87)

Study of the interaction between an agent and its environment.

Core Concept: Mediation:

- Emphasizes the emergent, contingent nature of activity.
- Includes the environment as part of the cognitive process.
- Asserts that plans are artifacts of reasoning about actions (after the fact explanations, rather than deliberate procedures).

Situated interaction requires awareness

Suchman, L.A.. Plans and situated actions: The problem of human-machine communication. Cambridge university press , 1987

Situated Interaction Requires Awareness

Mica Endsley, Ph.D., P.E.
PhD USC 1990
editor-in-chief of the Journal of Cognitive Engineering and Decision Making
President: SA Technologies
Specialty: Cognitive Engineering
Application Domain: Aviation and critical systems.

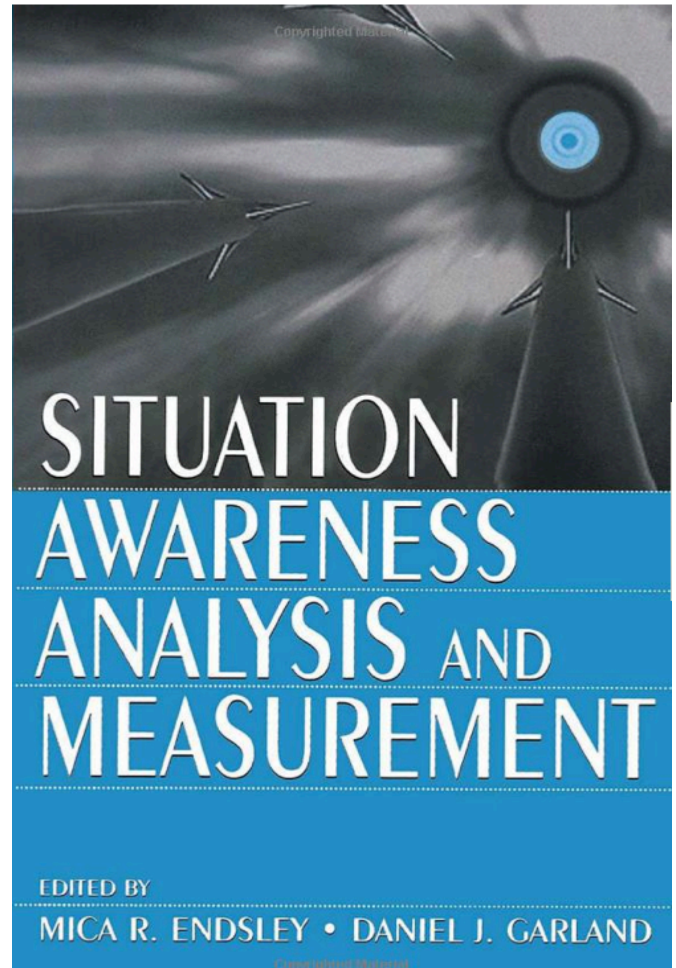

Situation Awareness

Situation Awareness : The Perception of [relevant] elements of the environment in a volume of space and time, the comprehension of their meaning and the projection of their status in the near future. (Endsley 2000)

Levels in Situation Awareness
1: Detection: Sensing of entities relevant to task
2: Assimilation: Associating perception with models that predict and explain.
3: Projection: Forecast events and dynamics of entities

M. Endsley, D. Garland, Situation Analysis and Awareness, Lawrence Erlbaum, 2000)

Situation Models: Philip Johnson-Laird

Philip N. Johnson-Laird

PhD Psychology, 1967, University College London
Stuart Professor of Psychology at Princeton Univ.
1971-1973: Inst. of Advanced Study, Princeton U.
1973-1989: Laboratory of Exp. Psychology, Univ of Sussex
1989- Applied Psychology Unit, Princeton Univ.

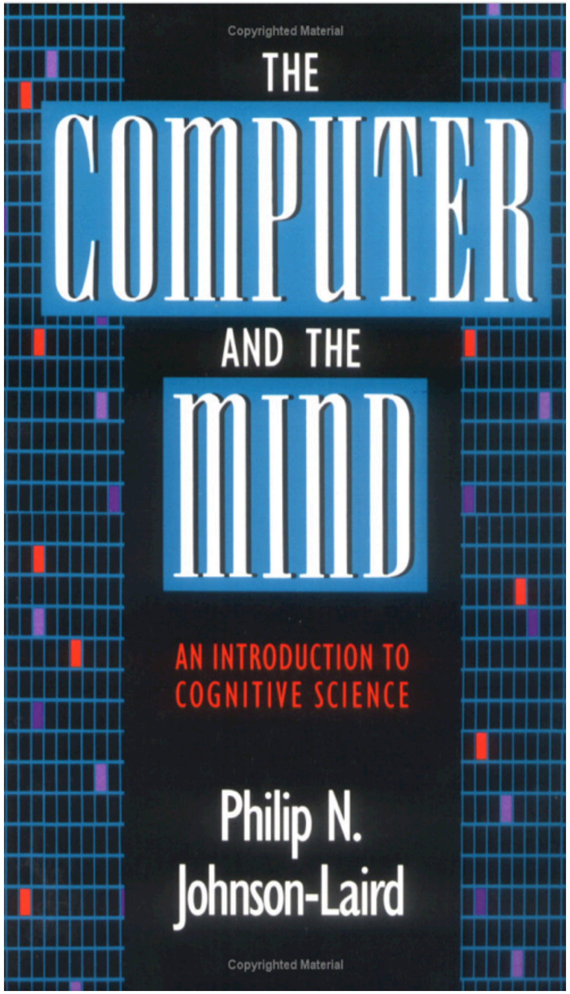

Situation Models:
a theory of mental models for natural language and inference.

Situation Models are widely used in Cognitive Psychology
to describe human abilities for

1) Providing context for story understanding
2) Interpreting ambiguous or misleading perceptions.
3) Reasoning with default information
4) Focusing attention for problem solving

Situation Models can provide a software framework for intelligent systems and services that interact with humans

Research Challenges for Situation-Aware Collaboration

Research Challenges

- How can we create technologies to permit humans and intelligent systems to share understanding of a situation.

- Can we use the latent variables from a multimodal transformer as a situation model for Situated Interaction?

A Hierarchical Framework for Collaborative AI

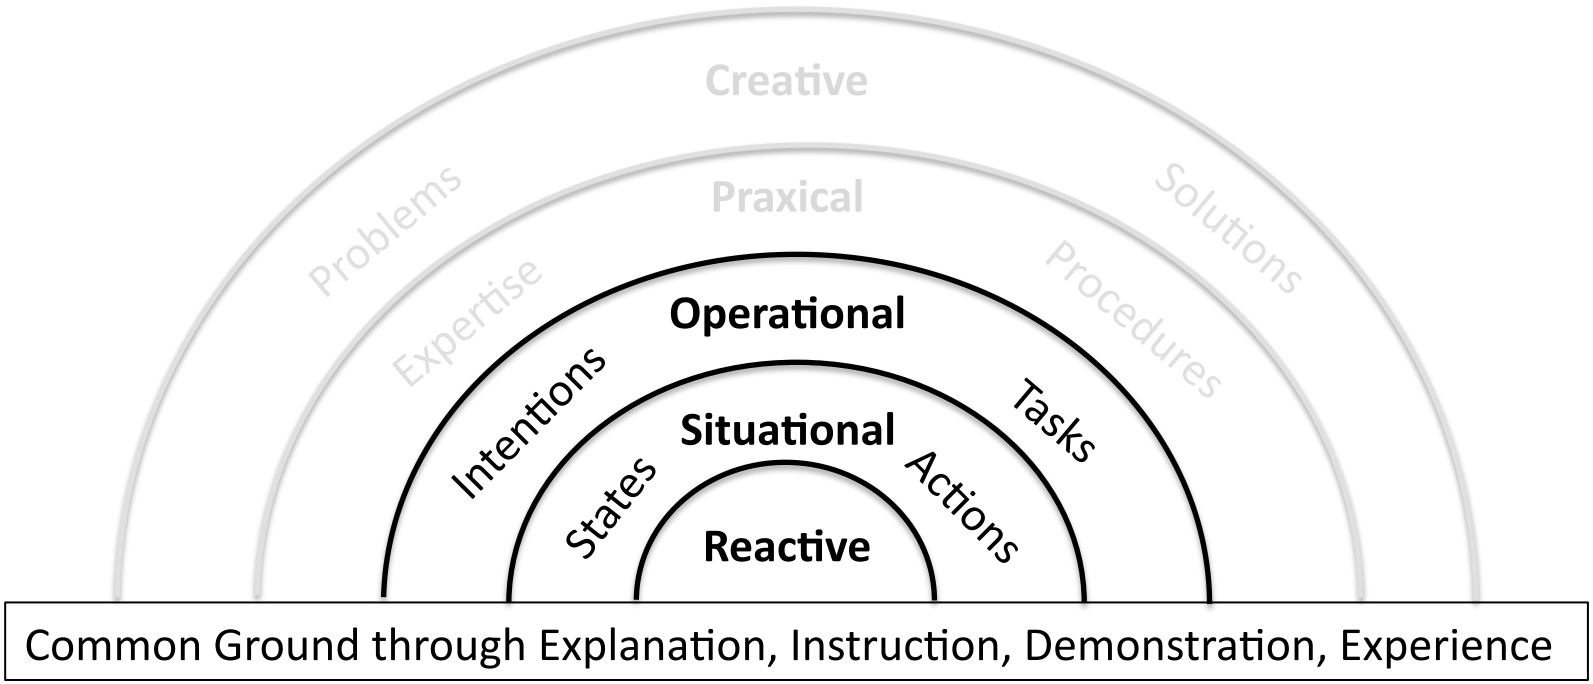

Operational Collaboration

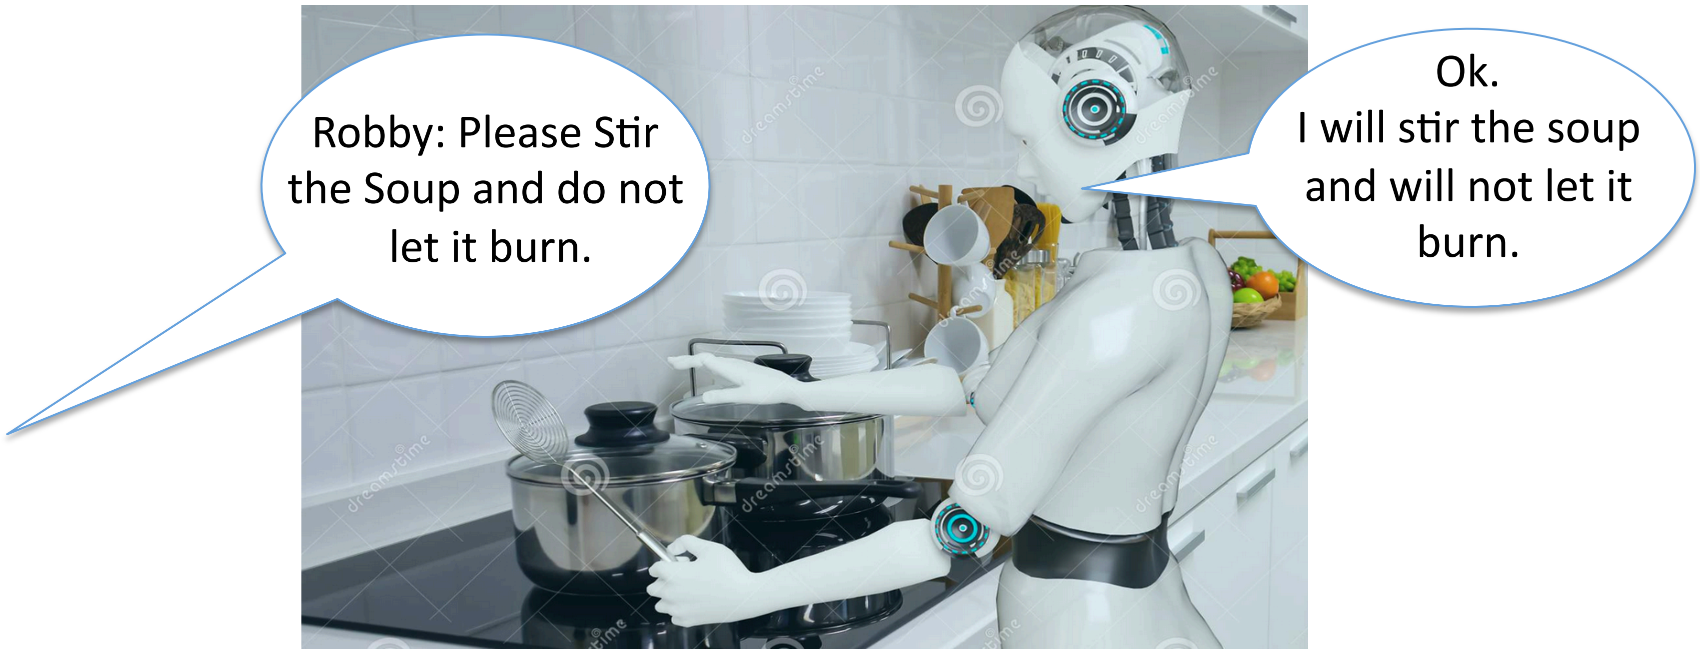

Operational Collaboration: shared authority over initiating, conducting, or terminating tasks and activities. (deciding who does what and when).
Operational Collaboration is typically mediated with **roles** and **protocols**.

Operational Collaboration

Shared authority over initiating, conducting, or terminating tasks and activities.

Examples:

- Meetings
- Games
- Talking on the phone
- Buying something in a shop

Roles and Protocols

Authority: Liberty to take actions.
Role: The behaviour expected of an individual who occupies a given social position or status.
Protocols: Rules and guidelines that govern actions and behaviors.

Social interactions are mediated by social protocols

Social interactions are mediated by shared protocols for polite interaction that governs perception, action and communication with roles.

Protocols simplify interactions by providing a script that prescribes a limited set of messages (greetings, communications, displays) that a participant should expect to receive, and a limited number of responses that should be communicated in response.

Technologies and Challenges for Operational Social Interaction with AI

Operational Social Interaction: Interaction mediated by social customs

Examples of human operational social interaction.

- Meetings
- Games
- Buying something in a shop
- Talking on the phone

Example Research Challenges

- Techniques to allow systems to learn social protocol from observation or explanation
- Technologies to permit humans and intelligent systems to negotiate and agree on interaction protocol

A Hierarchical Framework for Collaborative AI

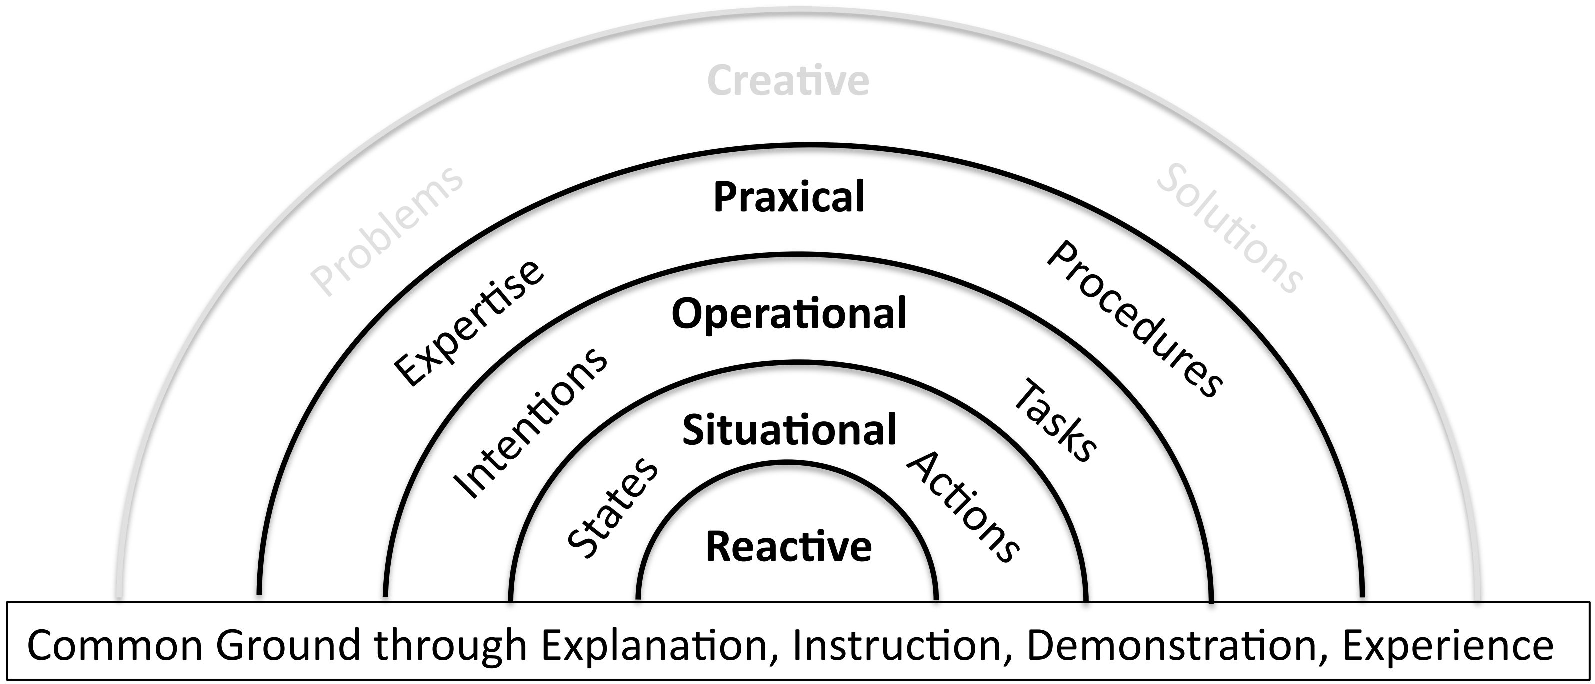

Technologies and Challenges for Praxical Collaboration

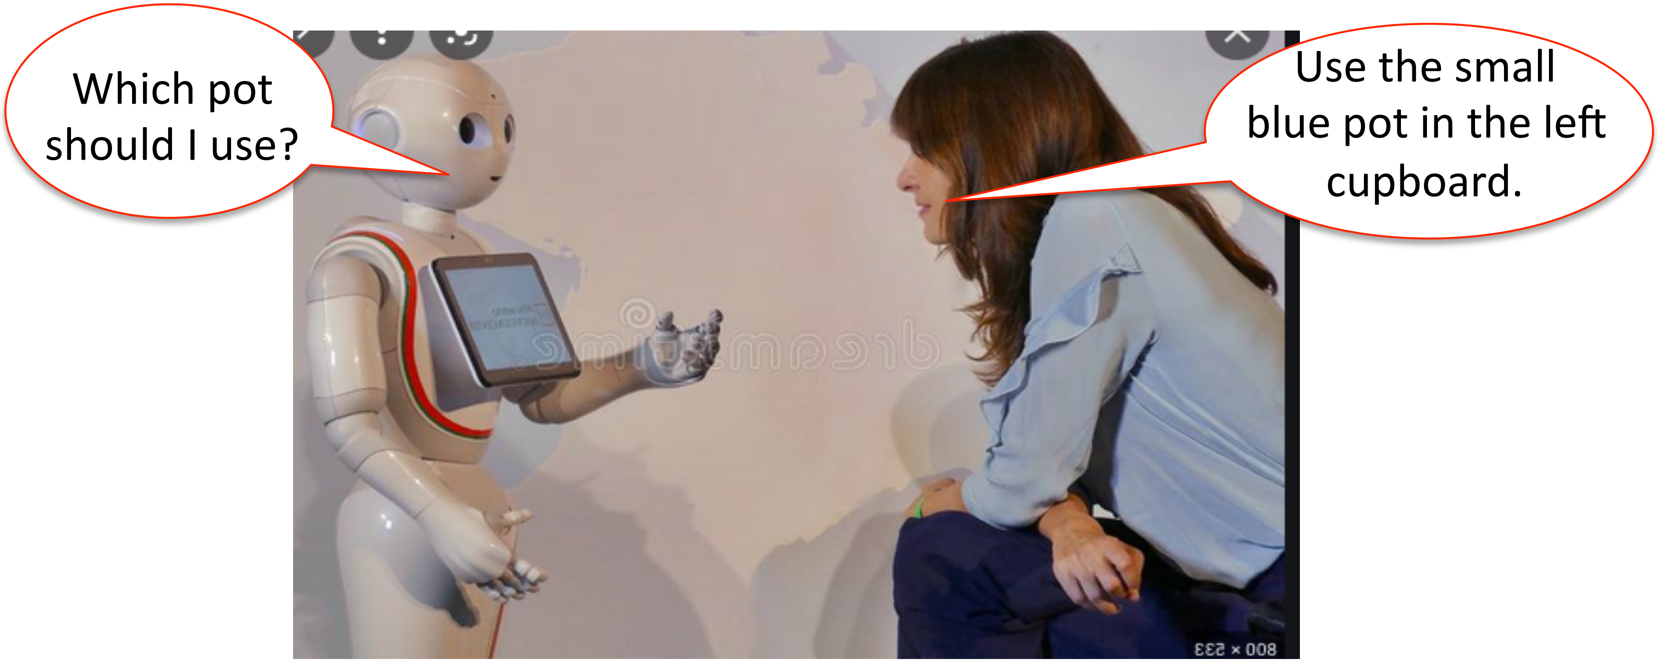

Praxical Collaboration: exchange of knowledge about how to attain goals and maximize value based on experience.

Praxical Collaboration:

Exchange of knowledge about how to attain goals and maximize value based on experience.

Examples:

- Sharing strategies for how to navigate in complex situations such as traffic related to sporting or political events.
- Explaining how to assemble or repair a device

Classic Technologies

- Rule based Expert Systems

Example Research Challenge

- Technologies to permit humans and intelligent systems to collaboratively plan and execute operations including contingencies.

A Hierarchical Framework for Collaborative AI

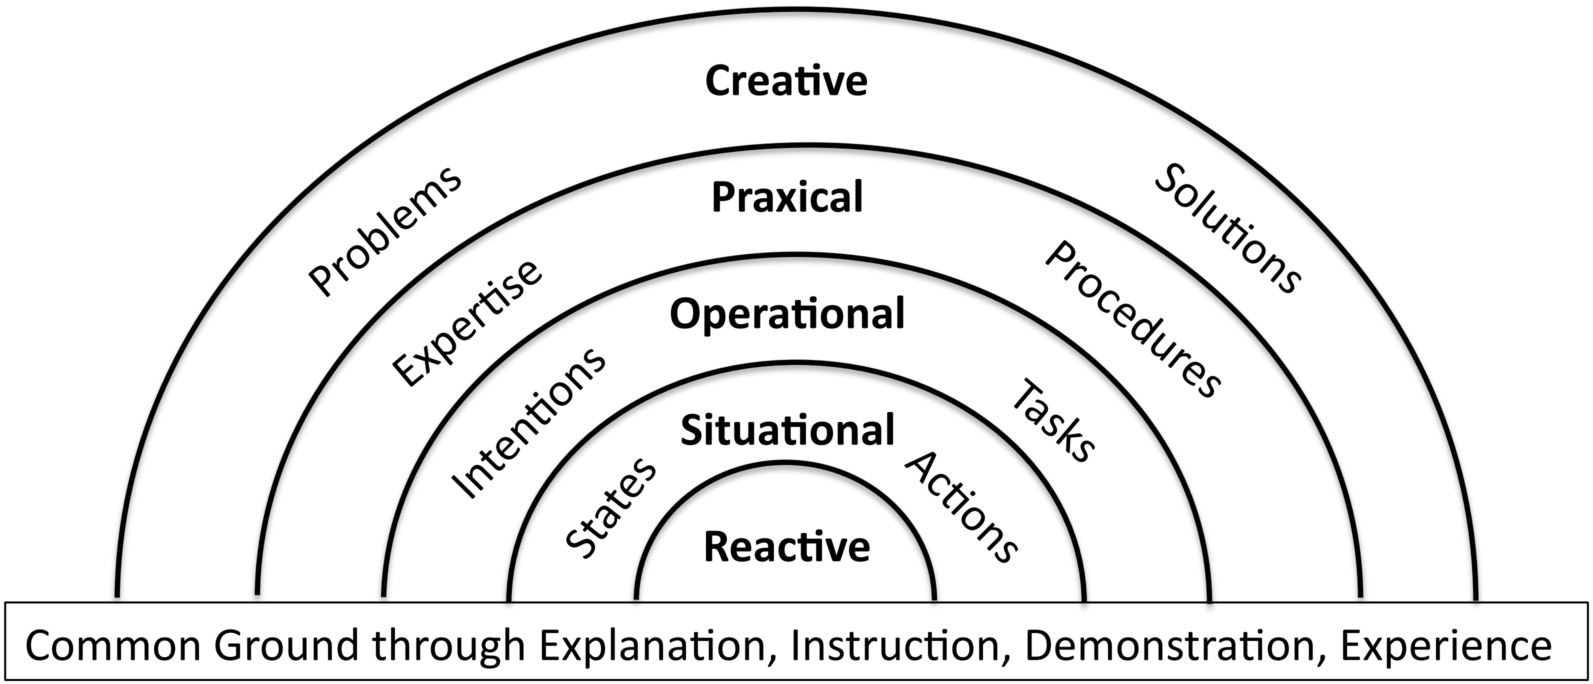

Creative Collaboration

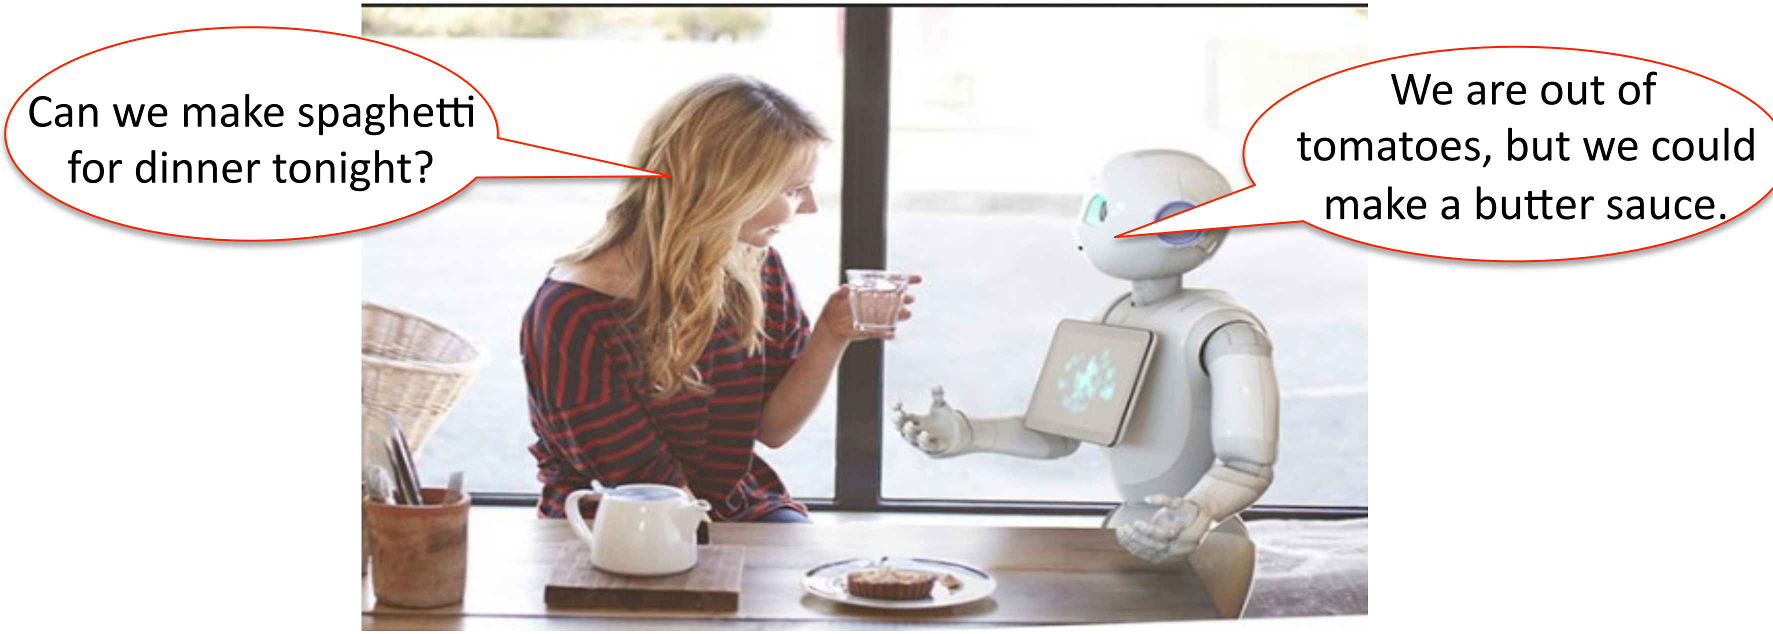

Two or more partners work together to solve a problem or create an original artifact.

Technologies and Challenges for Creative Collaboration

Creative Collaboration:

Partners work together to solve a problem or create an original artifact.

Classic Technologies

- EMYCIN (1978): Rule based expert system for antibiotoc therapy
- R1 (1980) Rule based Expert System for Configuring VAC computers

Modern Technologies

- Cognitive Computing (IBM Watson)
- Codex (automatic coder using GPT-3)
- DALL-EE (image generator based on GPT-3)
- LaMDA (conversational agent from Google)

So what happens next?

AI is a Transformative Technology, on the order of Fire or Electricity.

Technological transformations play out of several human generations. Gradually, humanity will become dependent on Intelligent Systems as we already are on Fire and Electricity.

Electric Appliances and their Penetration Rates

1% penetration: Appliance (years to 50% penetration)*

<u>Time Saving Appliances</u>
1890: Telephone (56 years)
1909: The Electric Iron (24 years)
1915: Vacuum Cleaners (40 years)
1916: Clothes Washers (20 years)
1934: Electric Kettle (33 years)
1948: Blender (22 years)
1950: Clothes Dryer (22 years)
1973: Microwave Oven (13 years)

<u>Quality of Life Appliances</u>
1913: Refrigerator (13 years)
1911: Air Conditioner (22 years)
1920: Radio (6 years)
1948: B&W Television (5 years)
1961: Color Television (6 years)
1969: VCR (9 years)

Quality of Life appliances achieve faster market penetration!

*S. Bowden and A. Offer, 1994.

So what happens next?

Artificial Intelligence is a Transformative Technology,
on the order of Fire or Electricity.

Will AI enslave Mankind? Probably not, but…..

Artificial Intelligence can be a very powerful technology for social control.

We see this already in China with ubiquitous control
We see this already with Deep Fake.
There are enormous dangers with social media.

The Origins of Artificial Intelligence and the Emergence of Collaborative Intelligent Systems

James L. Crowley
Professor Emeritus, Grenoble Institut Polytechnique
MIAI Chair on Collaborative Intelligent Systems
Univ. Grenoble-Alpes