



COL333/671: Introduction to AI

Semester I, 2022-23

What is AI?

Rohan Paul


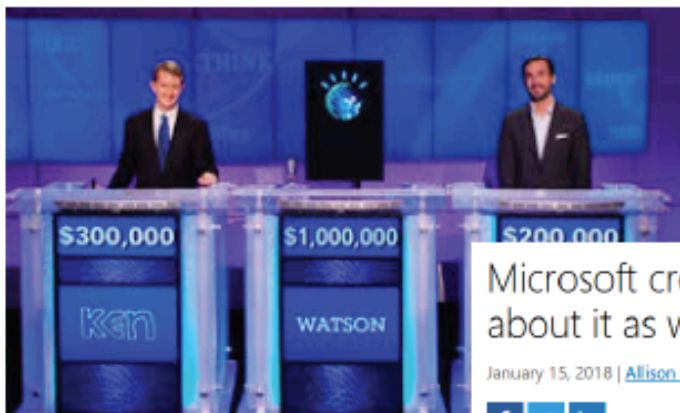


This Class

- Different theoretical foundations of AI.
- What does AI try to accomplish?
- Historical developments.
- AI problem solving techniques.
- Reference Material
 - AIMA Ch. 1 and Ch. 2 (2.1-2.3)

Acknowledgement

These slides are intended for teaching purposes only. Some material has been used/adapted from web sources and from slides by Dorsa Sadigh, Percy Liang, Mausam, Dan Klein, Nicholas Roy and others.



Microsoft creates AI that can read a document and answer questions about it as well as a person

January 15, 2018 | [Allison Linn](#)



Microsoft researchers achieve new conversational speech recognition milestone

August 20, 2017 | By [Xuedon](#)



June 24, 2014

DeepFace: Closing Performance in Face Recognition

Conference on Computer Vision and Pattern Recognition (CVPR)

By: [Yaniv Taigman](#), [Ming Yang](#), [Marc'Aurelio Ranzato](#), [Lior Wolf](#)

Abstract

In modern face recognition, the conventional pipeline co-classify. We revisit both the alignment step and the representation in order to apply a piecewise affine transformation layer deep neural network. This deep network involves more locally connected layers without weight sharing, rather than fully connected layers. We trained it on the largest facial dataset to-date, an identity dataset belonging to more than 4,000 identities.

If you think AI will never replace radiologists—you may want to think again

May 14, 2018 | [Michael Walter](#) | [Artificial Intelligence](#)



It's one of the most frequently discussed questions in radiology today: What kind of long-term impact will artificial intelligence (AI) have on radiologists?

Robert Schier, MD, a radiologist for RadNet, shared his own thoughts on the topic in a [new commentary](#) published by the *Journal of the American College of Radiology*—and he's not quite as optimistic as some of his colleagues throughout the industry.





It [AI] would take off on its own and redesign itself at an ever increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded.

— *Stephen Hawking* —

AZ QUOTES



Human-compatible AI

1. The robot's only objective is to maximize the realization of human values
2. The robot is initially uncertain about what those values are
3. Its behavior provides information about those values

AI: Informal Definition

- Models and algorithms that lead to intelligent behavior or solve problems that require human-like intelligence.

What is Intelligence?

- Dictionary.com: *capacity for learning, reasoning, understanding, and similar forms of mental activity*
- Ability to perceive and act in the world
- Reasoning: proving theorems, medical diagnosis
- Planning: take decisions
- Learning and Adaptation: recommend movies, learn traffic patterns
- Understanding: text, speech, visual scene

A machine can never do a task X. AI researchers respond by making a computer do X.

Perceptual Tasks: Natural Language

- Speech technologies (e.g. Siri, Alexa)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Language generation
- Language processing technologies
 - Question answering
 - Machine translation
 - Web search
 - Text classification, spam filtering, etc...



"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine de Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine. Le Chine sur ses gardes



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard



Visual Chatbot



Please wait for a few seconds while a caption is being generated for

Caption: a man is standing on a street with a skateboard

in front of a building

he is wearing a yellow shirt


where is the man standing

what is the colour of the shirt

Perceptual Tasks: Visual Recognition

- Object and face recognition
- Scene segmentation
- Image classification

<https://vision-explorer.allenai.org/>






Images from Erik Sudderth (left), wikipedia (right)

Adapted from D. Klein

Planning Tasks

- Predicting Structures
 - Given an amino acid sequence predict its structure.
 - How a protein folds -> functional characteristics
- Sequential Decision-making
 - Game Playing
 - Deep Blue
 - Alpha Go
 - Treatment recommendation



Logic and Reasoning

- Logical systems
 - Theorem provers
 - Fault diagnosis. Medical diagnosis.
 - Manufacturing planning
- Methods:
 - Deduction systems
 - Satisfiability solvers

An Autonomous Diagnostic and Prognostic Monitoring System for NASA's Deep Space Network

Mark L. James
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-8488
Mark.James@jpl.nasa.gov

Lydia P. Dubon
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 393-5288
Lydia.P.Dubon@jpl.nasa.gov

Abstract—Our objective is to provide a framework of automated tools and techniques for reducing operational and maintenance costs in the NASA's Deep Space Network (DSN). The focus of our technology application is fault diagnostics and prognostics for ground systems during DSN tracking operations. The domain chosen to demonstrate our capability is the new DSN Full Spectrum Processing Array configuration located at the Goldstone Deep Space Communications Complex (GDSCC) which is monitored by the Jet Propulsion Laboratory (JPL).

Communications Complex (GDSCC) which is monitored by the Jet Propulsion Laboratory (JPL).

To accomplish our goals we use two JPL developed tools: Beacon-based Exception Analysis for Multi-missions (BEAM) and Spacecraft Health Inference Engine (SHINE). BEAM is used as a highly advanced prognostic state estimator and SHINE is being used for hard real-time diagnostics and interpretation of the system state output by BEAM. These technologies provide new insights into system visibility that were not previously possible using channel-based diagnostics techniques thereby making near zero false alarms attainable. Raw sensor data and software-derived data are simultaneously fused in real-time to automatically abstract system physics and information invariants (constants). This methodology enables a system to be ultra-sensitive to degradation and changes, and to isolate significant events in both time and space to specific sensors.

This paper provides an overview of the synergistic approach to applying BEAM and SHINE technologies to DSN ground tracking systems, which maximizes the benefits from each of these technologies.

TABLE OF CONTENTS

1. ABSTRACT
2. INTRODUCTION
3. BACKGROUND
4. DEEP SPACE NETWORK
5. OBJECTIVE
6. DSN AUTONOMOUS FDI FRAMEWORK
7. DISCUSSION OF BEAM
8. DISCUSSION OF SHINE
9. BEAM & SHINE IN DSN OPERATIONS
10. CONCLUSIONS
11. ACKNOWLEDGEMENTS
12. REFERENCES
13. BIOGRAPHY

1. INTRODUCTION

Our objective is to provide a framework of automated tools and techniques for reducing operational and maintenance costs in the NASA's Deep Space Network (DSN). The focus of our technology application is fault diagnostics and prognostics for ground systems during DSN tracking operations. The domain chosen to demonstrate our capability is the new DSN Full Spectrum Processing Array configuration located at the Goldstone Deep Space

2. BACKGROUND

Automation of DSN downlink operations is a critical step in the advancement of NASA's communication link to future unmanned spacecraft. The forces behind the development of autonomous ground systems are both economic and technical. The era of NASA's New Millennium, Discovery and Mars Exploration programs will result in a series of "faster, better, cheaper" missions. These new mission series will approximately triple the mission load for DSN operations, thereby, increasing the demand for reliable and efficient ground tracking systems with minimum system failures and minimum downtime. An increase in operations staff would result in prohibitive costs



Robotics/Embodied AI

- Intelligent cars that can drive autonomously
- Intelligent manipulation tasks
- Unmanned exploration.
- Machine capable of walking



<https://www.youtube.com/watch?v=JTVJkJavU6g>

<https://www.youtube.com/watch?v=uWv-I7XMoB8>



When is a computer program or system displaying AI capabilities?

Various perspectives based on how we assess performance


- Human-centered
 - Comparison with human performance.
 - Empirical observations and hypothesis about human behaviour.
- Engineering viewpoint.
 - Comparison with a rational or right thing given what it knows (objective function).
 - Performance w.r.t. an objective.
- Thinking vs. Acting
 - Displaying input/output behaviour vs. Human-like way to arrive at the conclusion.

<p>Thinking Humanly</p> <p>“The exciting new effort to make computers think . . . <i>machines with minds</i>, in the full and literal sense.” (Haugeland, 1985)</p> <p>“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)</p>	<p>Thinking Rationally</p> <p>“The study of mental faculties through the use of computational models.” (Chamiak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)</p>
<p>Acting Humanly</p> <p>“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)</p>	<p>Acting Rationally</p> <p>“Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998)</p> <p>“AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)</p>
<p>Figure 1.1 Some definitions of artificial intelligence, organized into four categories.</p>	

Acting Humanly: The Turing Test

Turing (1950) "Computing machinery and intelligence":

- ◇ "Can machines think?" → "Can machines behave intelligently?"
- ◇ Operational test for intelligent behavior: the Imitation Game



Faculties


Natural Language Processing
Knowledge Representation
Automated Reasoning
Machine Learning
Computer Vision
Robotics



Alan Turing

Thinking Humanly

- Cognitive Science
 - How humans or animals perceive and act in the world.
 - Efforts to build testable models of human reasoning about tasks and problems.
- What are the tools?
 - Introspection. Psychological experiments.
- Aim
 - A program or an AI system is considered intelligent if its trace or reasoning process matches the steps that a cognitive model of human would take for a task/problem.
 - Back and forth. A computer program really well on a task. Then it should inform our understanding of human reasoning.
- Applications
 - Intelligent tutoring systems
 - Human computer interaction

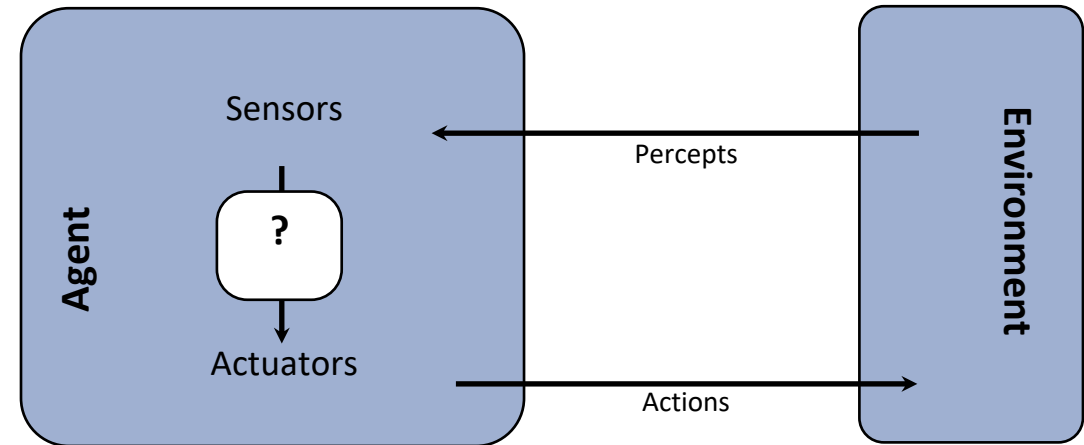


AI: Two views

- Thinking and Acting humanly
 - Thought processes and reasoning.
 - Leading to human-like behavior.
- **Thinking and Acting rationally**
 - System is rational if it does the “right thing” given what it knows.
 - Measuring against an ideal performance measure.
 - Engineering approach.

Agent View of AI

- What is an agent?
 - An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.
- Examples
 - Alexa
 - Robotic system
 - Refinery controller
 - Question answering system
 - Crossword puzzle solver
 -



Agent View of AI

- Agent Type
- Performance Measure
- Environment
- Actuators
- Sensors

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments, referrals	Keyboard entry of symptoms, findings, patient's answers
Satellite image analysis system	Correct image categorization	Downlink from orbiting satellite	Display of scene categorization	Color pixel arrays
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, operators	Valves, pumps, heaters, displays	Temperature, pressure, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, suggestions, corrections	Keyboard entry

Figure 2.5 Examples of agent types and their PEAS descriptions.



Domain Characteristics

- Fully or Partially observed
- Single or Multiple Agents
- Deterministic or Stochastic
- Episodic or Sequential
- Static or Dynamic
- Discrete or Continuous

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete


Acting Rationally: Maximizing Expected Utility

- An agent is an entity that perceives and acts.
- Rational agent
 - A rational agent selects actions that maximize its (expected) utility. The agent prefers those actions that take it closer to its objective.
 - Rationality implies: the agent must act to achieve the best outcome (deterministic case) or the best expected outcome (stochastic case).
 - Characteristics of the percepts, environment, and action space dictate techniques for selecting rational action.
- Rationality viewpoint lends itself to a mathematical formalism
 - Objective function and costs and algorithms that can maximize the agent's objective. Find the best agent for the architecture.



Thinking Rationally: Laws of Thought

- This perspective says that an AI system should display a logical thought process.
 - Aristotle
 - Logical way of deduction and reasoning.
- But sometimes we act even without deliberation
 - Reflex actions occur without deliberation.
 - E.g., reflex action if we touch something hot.



Strong AI vs. Weak AI Hypothesis

- Weak AI Hypothesis
 - Can machines act intelligently?
 - A system that passes the Turing test (appears to be acting humanly) may not be actually thinking. It may be only be a simulation of thinking.
- Strong AI Hypothesis
 - Can machines really think?
 - Machines that act intelligently not just by simulating but actually by thinking. Awareness of its mental states and process of arriving at a solution.

How to solve AI tasks?

Problem: planning routes in a city.



AI system: computational model that can provide you a route.

```
# This structure is for supporting on-line search
class PriorityQueue:
    def __init__(self):
        self.heap = []
        self.heap_size = 0
        self.priority = 0
        # Insert (state, cost, time, priority, (parent, child)) if
        # (state, cost, time, priority) is better than the existing
        # priority.
        # Return whether the priority queue was updated
        def update(self, state, cost, time, priority, (parent, child)):
            # Compare the new state with the existing state
            # If (state, cost, time, priority) is better than the existing
            # state, then update the priority queue.
            return True

        # Return (state, cost, time, priority, (parent, child))
        # or (None, None) if the priority queue is empty.
        def search(self):
            # If the heap is empty, return (None, None)
            if self.heap_size == 0:
                return (None, None)
            # Return the state with minimum priority, cost, time
            # or (None, None) if the priority queue is empty.
            # Return (state, cost, time, priority, (parent, child))
            # or (None, None) if the priority queue is empty.

# Example search problem on the number line.
# Start at 0, need to go to 10, cost 1 to move down, 2 to move up.
class NumberLineSearchProblem:
    def __init__(self, start, goal):
        self.start = start
        self.goal = goal
        self.current = start
        self.direction = 1
        self.steps = 0

    def is_goal(self):
        return self.current == self.goal

    def actions(self):
        actions = []
        if self.current < self.goal:
            actions.append('up')
        elif self.current > self.goal:
            actions.append('down')
        return actions

    def result(self, action):
        new_state = self.current
        if action == 'up':
            new_state += 1
        elif action == 'down':
            new_state -= 1
        return new_state


    def cost(self, state):
        return abs(state - self.start)

    def time(self, state):
        return abs(state - self.start)


    def priority(self, state):
        return 0
```

1. Modeling

- Routing problem as a graph.




Modeling




2. Inference

- Path finding algorithm that runs on the graph.




Inference




3. Learning

- We can make use of past data.
- Learning parameters





+data

Learning




We will explore a variety of models

- State-based Models
- Variable based Models
- Decision-making models
- Reflex Models




Models that reason with world states



Reflex models. Example, a neural network classifying images.

5	3			7			
6			1	9	5		
	9	8				6	
8			6				3
4			8	3			1
7			2				6
	6				2	8	
			4	1	9		5
			8			7	9

Constraint satisfaction models



Probabilistic Models

History – Gestation Phase (1943-1955)

- Can we simulate a neuron?
 - Artificial neurons that turn on/off states based on inputs.
 - Showing that a neural network could learn (Mc Culloch and Pitts)
 - Simple updating rule to learn (Hebbian learning by Donald Hebb 1949)
- Turing
 - *Proposed the Turing Test*
 - *“Computing Machinery and Intelligence”*

Donald Hebb (1949)



When an axon of cell 1 is near enough to excite a cell 2 and repeatedly and persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that 1's efficacy, as one of the cells firing 2, is increased.

History – Birth of the Field (1956)

- Workshop at Dartmouth College;
 - John McCarthy, Marvin Minsky, Claude Shannon, etc.
- Set an agenda for researchers:
 - Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it.
 - Attempt to build machines that will function autonomously in complex changing environments.
 - Proposed a separate field of AI to address this agenda.



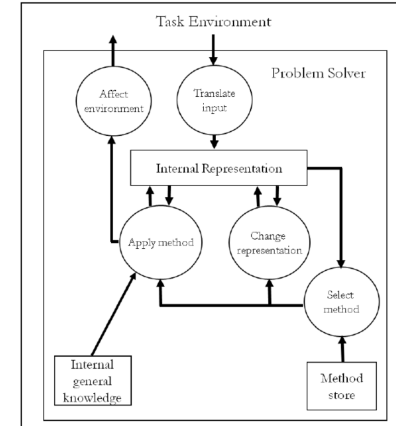
Dartmouth

History – Early Era (till 1970s)

- General Problem Solver (1955)
 - Newell & Simon's Logic Theorist: prove theorems in Principia Mathematica using search + heuristics, solving puzzles.
 - Accept new axioms on the fly.
- SHRLDU (1972)
 - Winograd. Micro-world to re-arrange blocks on the table. Language understanding
- Challenges
 - Limited Computation (could not manage state spaces)
 - Limited information (complexity of the real world)



Newell & Simon



SHRLDU:

<https://www.youtube.com/watch?v=bo4RvYJYOzI>

History – Expert Systems (1969 - 1979)

- Knowledge-based systems
 - Elicit specific domain knowledge from experts in form of rules:
 - if [premises] then [conclusion]
 - DENDRAL
 - Infer molecular structure from mass spectrometry data. Input is a molecular formula and how it gets fragmented during experiments.
 - The program deduced possible structures consistent with data.
- Limitation
 - No notion of uncertainty




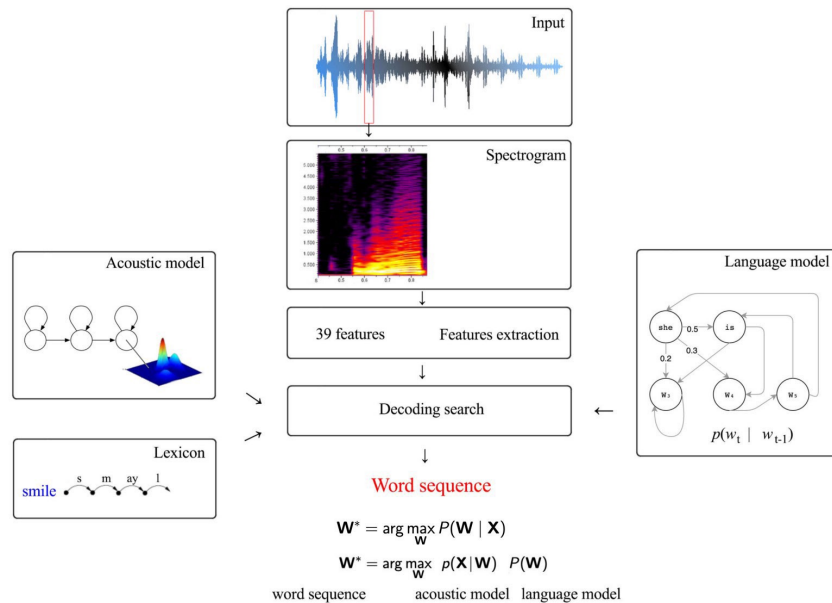
Fig. 2. McLafferty rearrangement. (Reproduced by permission. Source: C. W. A. Miles, ed., *Mass Spectrometry*, Wiley, 1971, p. 100.)

DENDRAL expert system


History – Embracing uncertainty (1987s onwards till 2000s)

- Decision making under uncertainty.
- Probabilistic models.

Speech recognition. HMMs.



Darpa Grand Challenge (2005) <https://www.youtube.com/watch?v=M2AcMnfzpNg>




Remote agent program on Deep Space I (1999) <https://ti.arc.nasa.gov/tech/asr/groups/planning-and-scheduling/remote-agent/>

History – Deep learning (2012 onwards)

- Confluence of three things:

- Large data sets
- Large amounts of computation
- Improved algorithms and optimization.



AlexNet (2012): huge gains in object recognition; transformed computer vision community overnight



AlphaGo (2016): deep reinforcement learning, defeat world champion Lee Sedol

- Vast improvement in certain tasks.

- Recognition.
- Game playing.

<https://github.com/features/copilot>

<https://affinelayer.com/pixsrv/>

<https://deepai.org/machine-learning-model/text2img>

<https://alphafold.ebi.ac.uk/>

System I and System II View


SYSTEM 1 VS. SYSTEM 2 COGNITION

2 systems (and categories of cognitive tasks):

System 1

- Intuitive, fast, **UNCONSCIOUS**, non-linguistic, habitual
- Current DL

Implicit knowledge

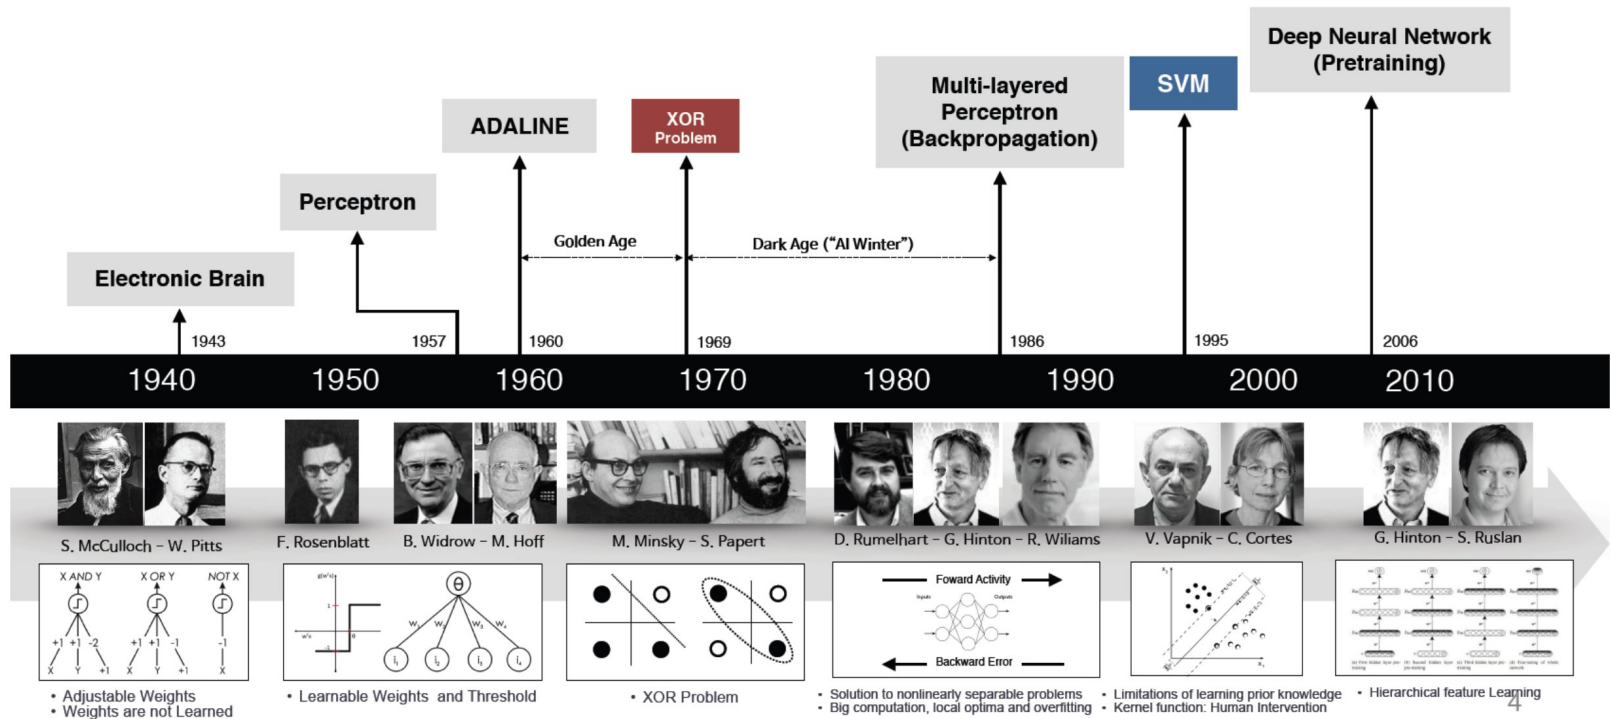


System 2

- Slow, logical, sequential, **CONSCIOUS**, linguistic, algorithmic, planning, reasoning
- Future DL

Explicit knowledge

Manipulates high-level / semantic concepts, which can be recombined combinatorially



There were two major winters in 1974–1980 and 1987–1993^[6] and several smaller episodes, including the following:

- 1966: failure of **machine translation**
- 1970: abandonment of **connectionism**
- Period of overlapping trends:
 - 1971–75: **DARPA's** frustration with the **Speech Understanding Research** program at **Carnegie Mellon University**
 - 1973: large decrease in AI research in the United Kingdom in response to the **Lighthill report**
 - 1973–74: **DARPA's** cutbacks to academic AI research in general
- 1987: collapse of the **LISP machine** market
- 1988: cancellation of new spending on AI by the **Strategic Computing Initiative**
- 1993: resistance to new **expert systems** deployment and maintenance
- 1990s: end of the **Fifth Generation computer** project's original goals

AI has been influenced by many disciplines

- Bayes rule (Bayes, 1763) from **probability**
- Least squares regression (Gauss, 1795) from **astronomy**
- First-order logic (Frege, 1893) from **logic**
- Maximum likelihood (Fisher, 1922) from **statistics**
- Artificial neural networks (McCulloch/Pitts, 1943) from **neuro-science**
- Minimax games (von Neumann, 1944) from **economics**
- Stochastic gradient descent (Robbins/Monro, 1951) from **optimization**
- Uniform cost search (Dijkstra, 1956) from **algorithms**
- Value iteration (Bellman, 1957) from **control theory**

Next Time

- This Class
 - What is AI?
- Next Class
 - Problem solving as search