IOT ONLINE COURSE

Fundamentals of Artificial Intelligence

F-AI-1: Introduction – What is AI?

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham
Université de Pau, France

Slides realized by J. Mantilla

IoT – from idea to reality

What is Artificial Intelligence?

- A machine that mimics cognitive functions.
- It is the science and engineering of making intellingent machines.
- In Computer Science, Artificial Intelligence (**AI**) research is defined as the study of « intelligent agents »

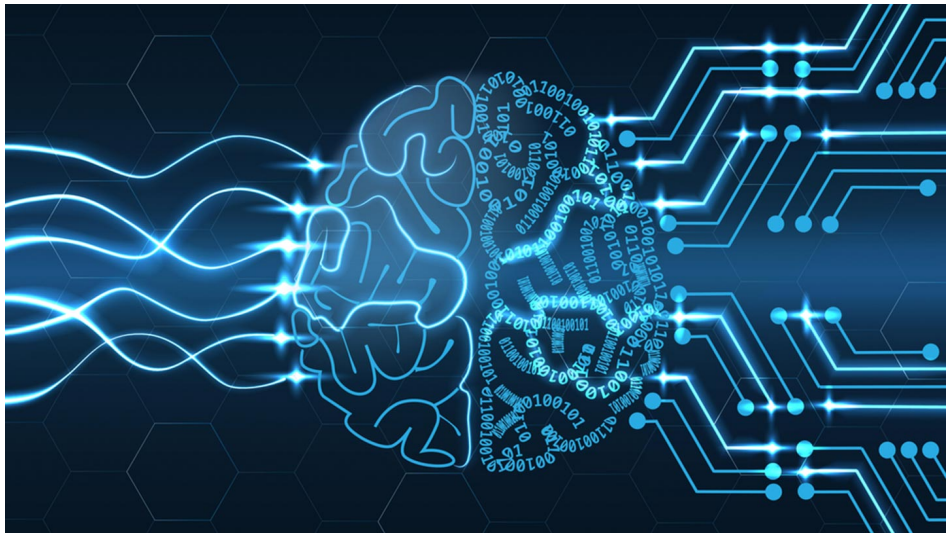

Definitions

- A serious science.
- General-purpose AI like the robots of science fiction is incredibly hard
  - Human brain appears to have lots of special and general functions, integrated in some amazing way that we really do not understand (yet)
- Special-purpose AI is more doable (nontrivial)
  - E.g., chess/poker playing programs, logistics planning, automated translation, speech and image recognition, web search, data mining, medical diagnosis, keeping a car on the road.

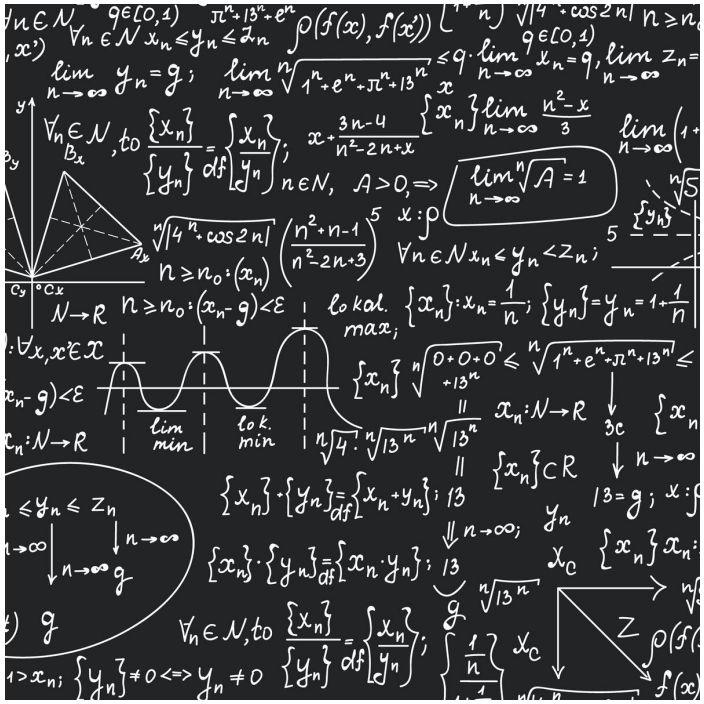

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

The Turing Test

- Proposed By Alan Turing in 1950
- To be called intelligent, a machine must produce responses that are indistinguishable from those of a human.
- Human judge communicates with a human and a machine over text-only channel.
- Both human and machine try to act like a human.
- Judge tries to tell which is which.
- Is Turing Test the right goal?

  "Aeronautical engineering texts do not define the goal of their field as making 'machines that fly so exactly like pigeons that they can fool even other pigeons.'" [Russell and Norvig]

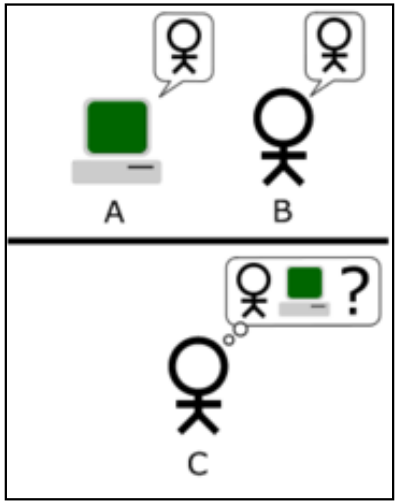

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Reflection

if AI can be more **rational** than humans in some cases, why not?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that **act rationally**

AI focus on **action.** Avoids philosophical issues such as "is the system conscious" etc.

Premise

- We will mostly follow "**act rationally**" approach
- Distinction may not be that important
- Acting rationally like a human presumably requires (some sort of) thinking rationally like a human
- Humans much more rational anyway in complex domains

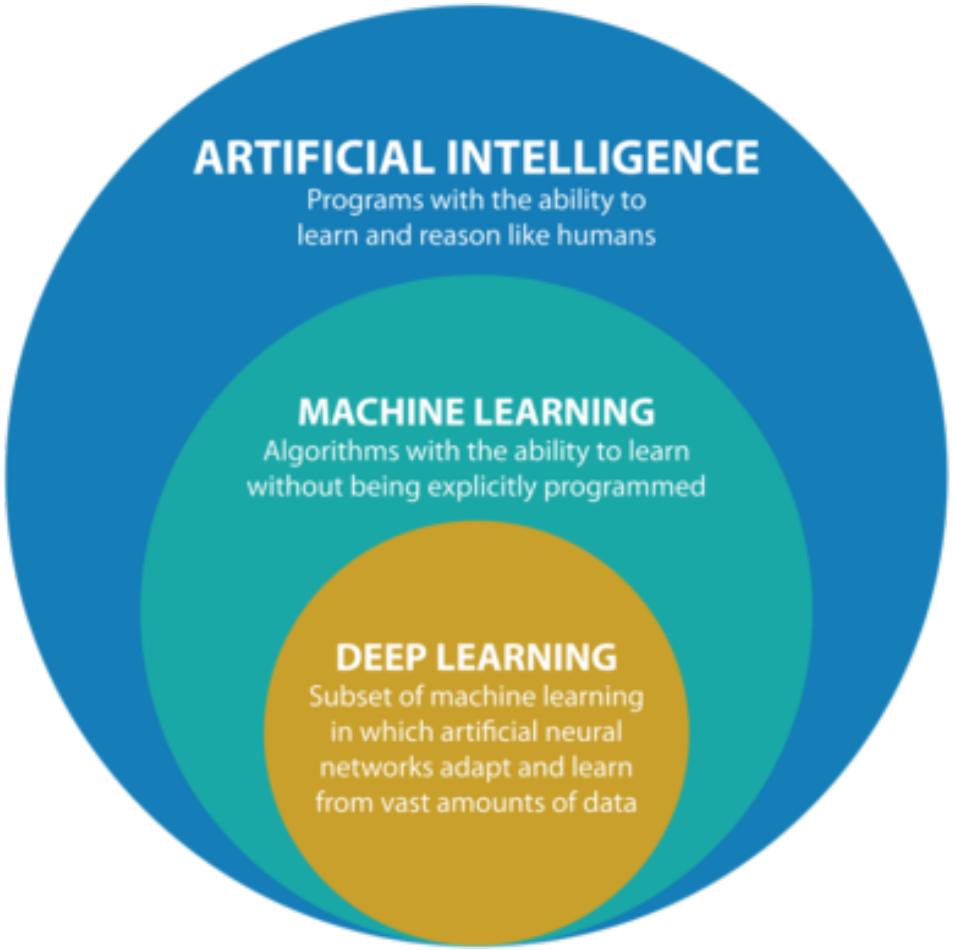

History

1956 Dartmouth Conference: The Founding Fathers of AI

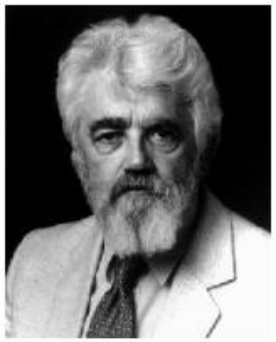

John McCarthy

Marvin Minsky

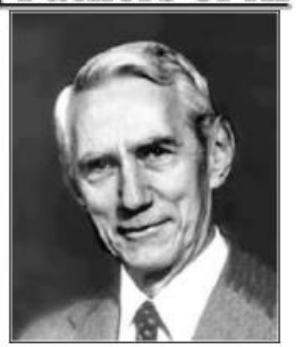

Claude Shannon

Ray Solomonoff

Alan Newell

Herbert Simon

Arthur Samuel

And three others...

Oliver Selfridge
(Pandemonium theory)

Nathaniel Rochester
(IBM, designed 701)

Trenchard More
(Natural Deduction)

1950s and 1960s

Early successes. AI can draw logical conclusions, prove some theorems, create simple plans... Some initial work on neural networks...

- Led to overhyping: researchers promised funding agencies spectacular progress, but started running into difficulties:
- Ambiguity: highly funded translation programs (Russian to English) were good at syntactic manipulation but bad at disambiguation

  "The spirit is willing but the flesh is weak" becomes "The vodka is good but the meat is rotten"
- Scalability/complexity: early examples were very small, programs could not scale to bigger instances
- Limitations of representations used

1960s, 1970s and 1980s

- The Logic Theorist

  Often considered as the first AI Program. It solved 38 of the first 52 theorems in Principia Mathematica. Written by Allen Newell, Herbert A. Simon and Cliff Shaw.

- Arthur Samuel

  Writes a self learning program for the game of checkers.

- Expert Systems

  During 70's and 80's. AI program that simulated the knowledge and analytical skills of human experts.

- AI Labs

  Set up at places like Stanford and MIT.

- AI Winters

  During 1974 and 1987. Funding cuts.

Timeline

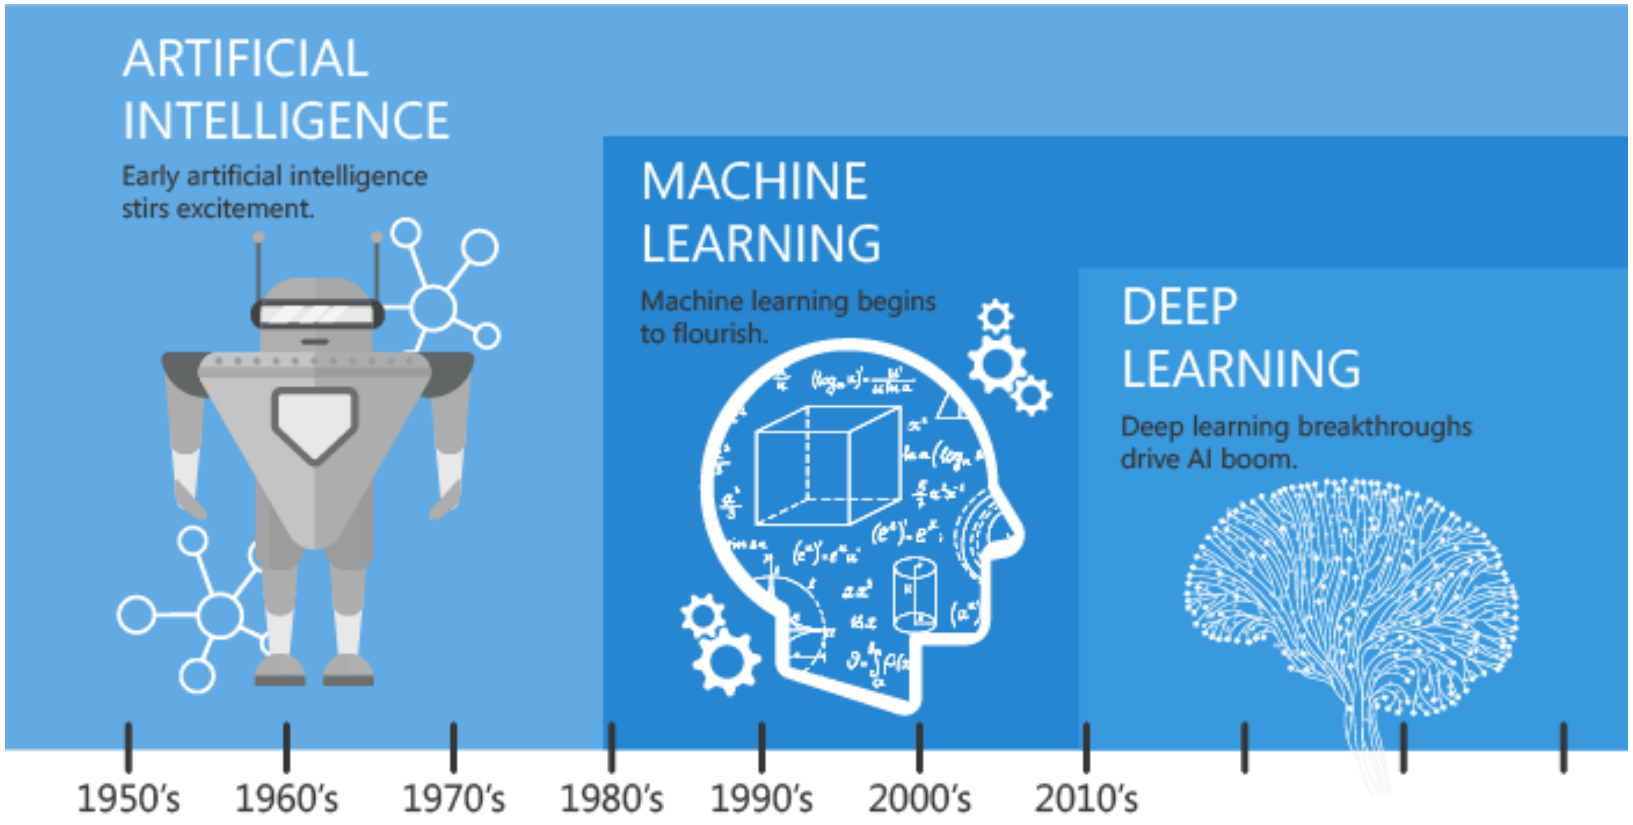

Since an early flush of optimism in the 1950's, smaller subsets of artificial intelligence - first machine learning, then deep learning, a subset of machine learning - have created ever larger disruptions.

“Chinese Room” Argument

- Proposed by John Searle in 1980
- Person who knows English but not Chinese sits in room
- Receives notes in Chinese
- Has systematic English rule book for how to write new Chinese characters based on input Chinese characters, returns his notes
  - Person=CPU, rule book=AI program, really also need lots of paper (storage)
  - Has no understanding of what they mean
  - But from the outside, the room gives perfectly reasonable answers in Chinese!
- Searle’s argument: the room has no intelligence in it!

Modern AI

- More rigorous, scientific, formal/mathematical
- Fewer grandiose promises
- Divided into many subareas interested in particular aspects
- More directly connected to “neighboring” disciplines
  - Theoretical computer science, statistics, economics, operations research, biology, psychology/neuroscience, …
  - Often leads to question “Is this really AI”?
- Some senior AI researchers are calling for re-integration of all these topics, return to more grandiose goals of AI
  - Somewhat risky proposition for graduate students and junior faculty…

Artificial General Intelligence

- Artificial General Intelligence (AGI):
- Also referred to as "Strong AI"
- AGI is a machine with general intelligence and, much like a human being, it can apply that intelligence to solve any problem.
- The kind of artificial intelligence we see in the movies, like the robots from *Westworld* or Data from *StarTrek: The Next Generation*.

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Narrow Artificial Intelligence

- Also referred to as "Weak AI"
- Operates within a limited context.
- Simulation of human intelligence.
- Focused on performing a single task extremely well.
- These machines may seem intelligent, they operate under far more constraints and limitations than even the most basic human intelligence.

Narrow AI Examples

- Smart assistants (like Siri and Alexa)
- Prediction tools
- Self driving cars
- Conversational bots for marketing and customer service
- Robo-advisors for stock trading
- Spam filters on email
- Social media monitoring tools for dangerous content or false news
- Song or TV show recommendations from Spotify and Netflix

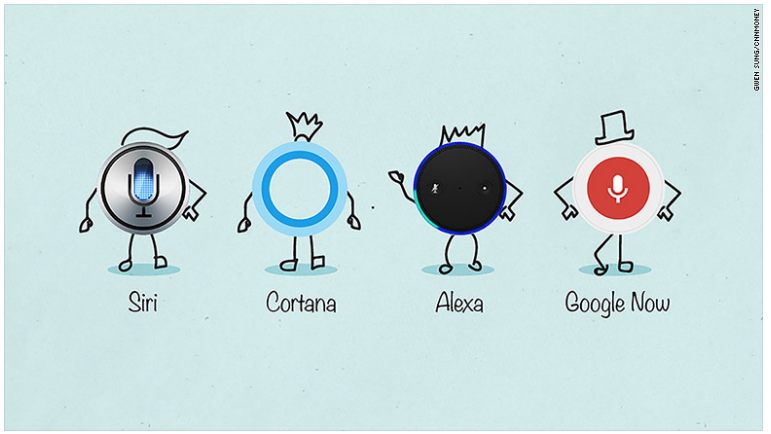

Prof. Congduc Pham
http://www.univ-pau.fr/~cpham

Symbolic Artificial Intelligence

- A type of Narrow Aritificial Intelligence
- Requires meticulous programming define the rules that specify the behavior of an intelligent system.
- Suitable for applications where the environment is predictable and the rules are clear-cut.
- Fallen from grace in the past years, but most of the applications we use today are rule-based systems.

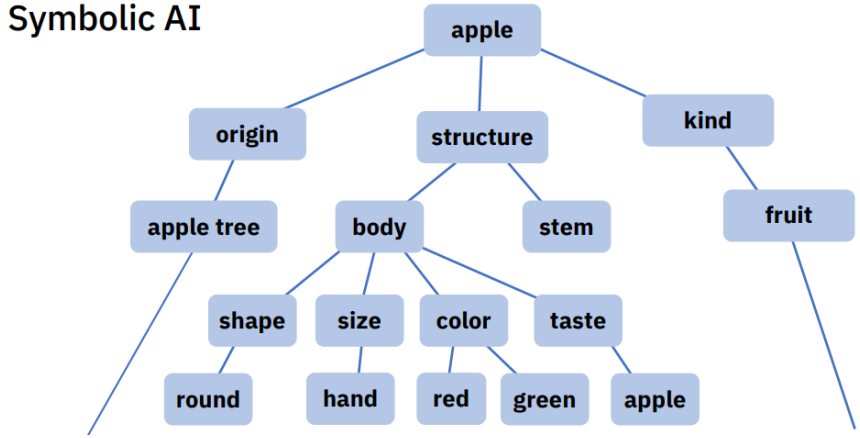

Machine Learning

- A type of Narrow Aritificial Intelligence
- Develops intelligent systems through examples.
- A developer creates a model and then “trains” it by providing it with many examples.
- The machine learning algorithm processes the examples and creates a mathematical representation of the data that can perform prediction and classification tasks.
- For instance, a machine-learning algorithm trained on thousands of bank transactions with their outcome (legitimate or fraudulent) will be able to predict if a new bank transaction is fraudulent or not.

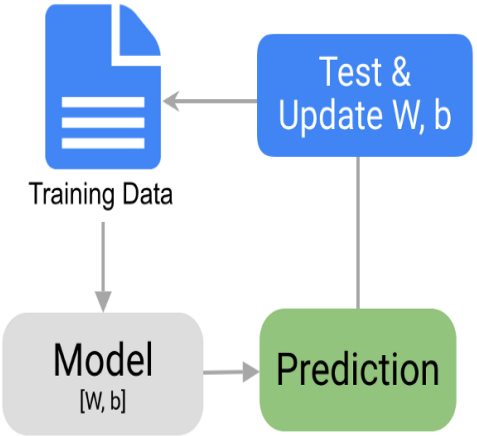

AI Technologies

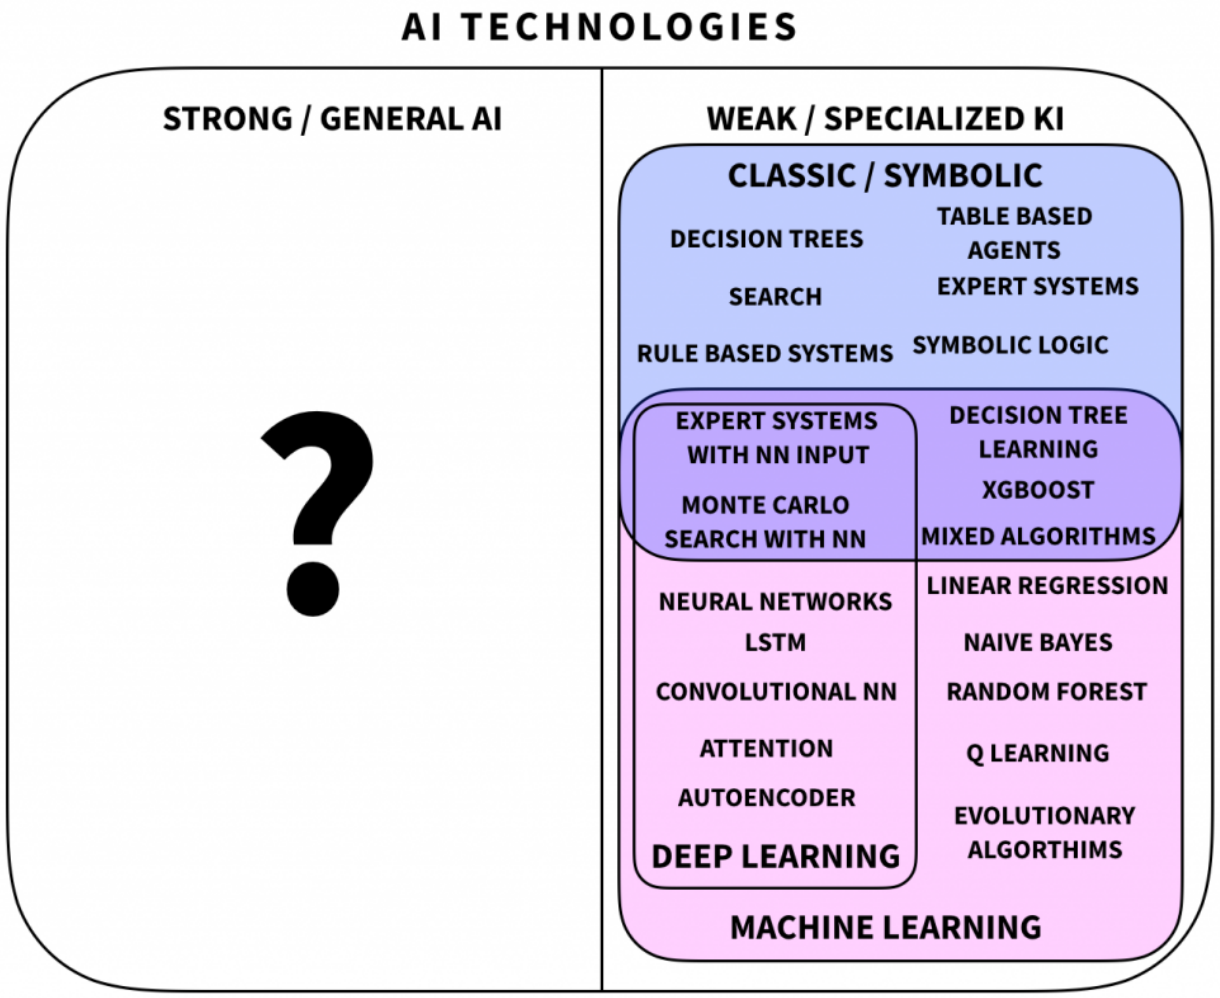

IOT Online Course

Fundamentals of Artificial Intelligence

Continue with
F-AI-2: Machine Learning

WAZIup
IoT – from idea to reality
WAZIhub