04.11.20 13:39

CSCI 101
Connecting with Computer Science
Artificial Intelligence I

Jetic Gū

2020 Fall Semester (S3)

Overview

- Focus: Artificial Intelligence
- Readings: -
- Core Ideas:
    1. What is Artificial Intelligence?
    2. History
    3. Challenges
    4. Discussion

What is Artificial Intelligence?

It’s not very intelligence though

P1
What?

What is Artificial Intelligence?

- Robotics
- Virtual Assistant
- Autonomous Vehicles
- …
- It’s more!

What?

P1
What?

What is Artificial Intelligence?

- Core: Statistics and Biology
  - Perception: interpret input
  - Take Actions: execution through commands and motor functions
  - Maximise The Goal: accomplish objective

Concept

P1
What?

Artificial Intelligence Areas

- Computer vision: Image processing
- Natural Language Processing: Processing Human Language
- Signal Processing: Audio, etc.
- Knowledge base, Reasoning: Logical representation of knowledge, Logical Reasoning
- Social Intelligence, Creativity

Concept

P1
What?

Modern AI Research

- Infancy, 1 month old child
  - Model: Statistics models, focusing on pattern recognition
  - We cannot yet do reasoning
  - We cannot yet process human language well
  - We cannot yet even recognise objects in a picture
  - Our models are way too simple

Concept

History of AI Research

Let’s go back

History of AI Research

- Alan Turing: Imitation Game[1] (a.k.a Turing test)
  - A human asks the same questions to a Machine and another Human
  - Can he/she tell which responses are from a human, which are from the machine?
  - What is the key? **Language**!

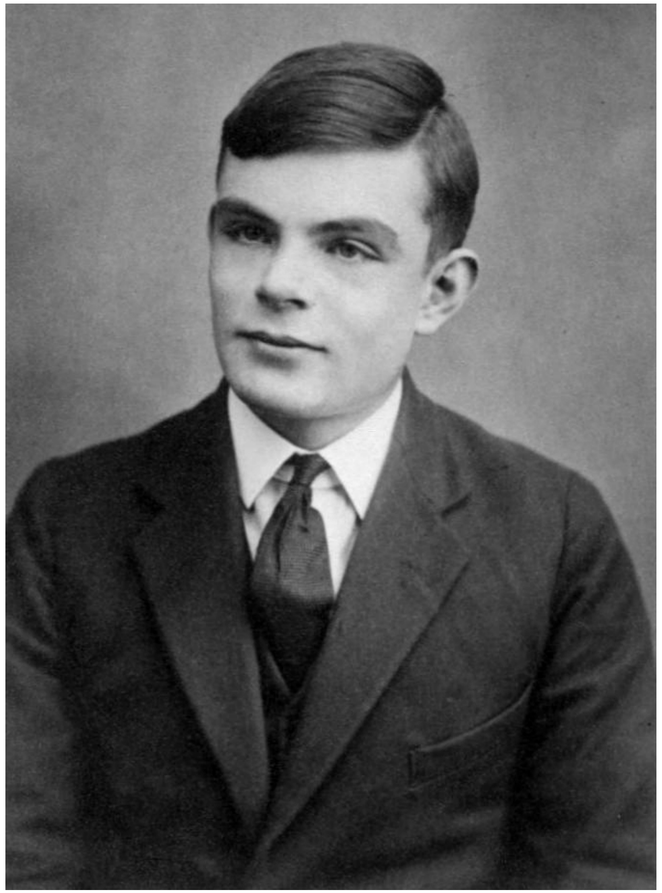

1. Turing, Alan M. "Computing machinery and intelligence." Parsing the turing test. Springer, Dordrecht, 2009. 23-65.

History of AI Research

- Direction 1: Linguistics and Logic
  - Universal Grammar
  - Formal Logic Inference
- Direction 2: Biology
  - Machine Learning and Pattern Recognition
  - Computer Vision

P2
History

History of AI Research: Ling

- Noam Chomsky: Universal Grammar[1]
  - Chomsky is a linguist
  - *There should be structural rules to languages, applicable to all human languages*
  - If such rules exist, we can convert all human languages to logical representations (Formal Logic)!

History

1. Hauser, Marc D., Noam Chomsky, and W. Tecumseh Fitch. "The faculty of language: what is it, who has it, and how did it evolve?." science 298.5598 (2002): 1569-1579.

P2
History

History of AI Research: Ling

If I give her my heart, she will love me.
I give her my heart.
She loves me.

- Formal logic
  - Assume p_1: I give her my heart, p_2: She loves me
  - Known
    - $p_1 \rightarrow p_2$
    - p_1
  - Inference: p_2, she loves me!

Example

1. Example here is first-order logic

History of AI Research: Ling

- Problem with Universal Grammar and Formal Logic

  1. Human language is too complex, Universal Grammar cannot be achieved

  2. Formal Logic (first-order or higher-order) cannot establish link to physical world (Naming problem)

     - Wittgenstein
       whereof one cannot speak, thereof one must be silent[1]

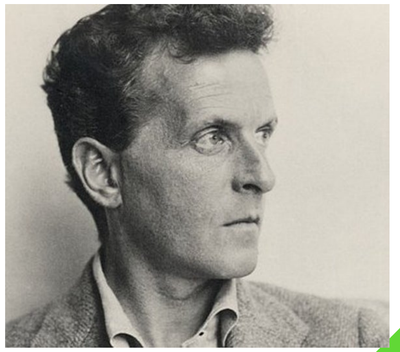

1. Wittgenstein, Ludwig. "Tractatus Logico-Philosophicus (trans. Pears and McGuinness)." (1961).

History of AI Research: ML

- Machine Learning
  - Not AI research, but based in Statistics
  - "When it is sunny, people are 90% more likely to buy ice cream"
    - Statistics: Empirical observation of facts
    - Fact: it is sunny, Model prediction: buy ice cream

History of AI Research: ML

- Machine Learning
  - Gathering evidence/features
    $F = x_0, x_1, \ldots, x_n$
  - How does an ML model tell if it is a dog?
    - x_0: "It is fluffy"
    - x_1: "It has oval eyes"
    - …

1. Similar to a Perceptron Model

History of AI Research: ML

- Machine Learning
  - Summing up features and learned weights
  - $P(L = \text{cat} \mid F) = \Sigma_i(w_i x_i)$
  - e.g.: picture on the right, $P(L = \text{cat} \mid F) = 0.98$

1. Similar to a Perceptron Model

History of AI Research: ML

- Machine Learning
  - Features: we define the features
  - Weight: learned using MATH and a large number of examples
    - This process is called machine learning
    - You need a lot of training data

F = (0,1,0,…….) L=cat
F = (1,1,0,…….) L=not cat
F = (0,1,0,…….) L=cat
F = (1,1,0,…….) L=not cat
F = (0,1,0,…….) L=cat
F = (1,1,0,…….) L=not cat
F = (0,1,0,…….) L=cat
F = (1,1,0,…….) L=not cat
…

History

1. Similar to a Perceptron Model

History of AI Research: ML

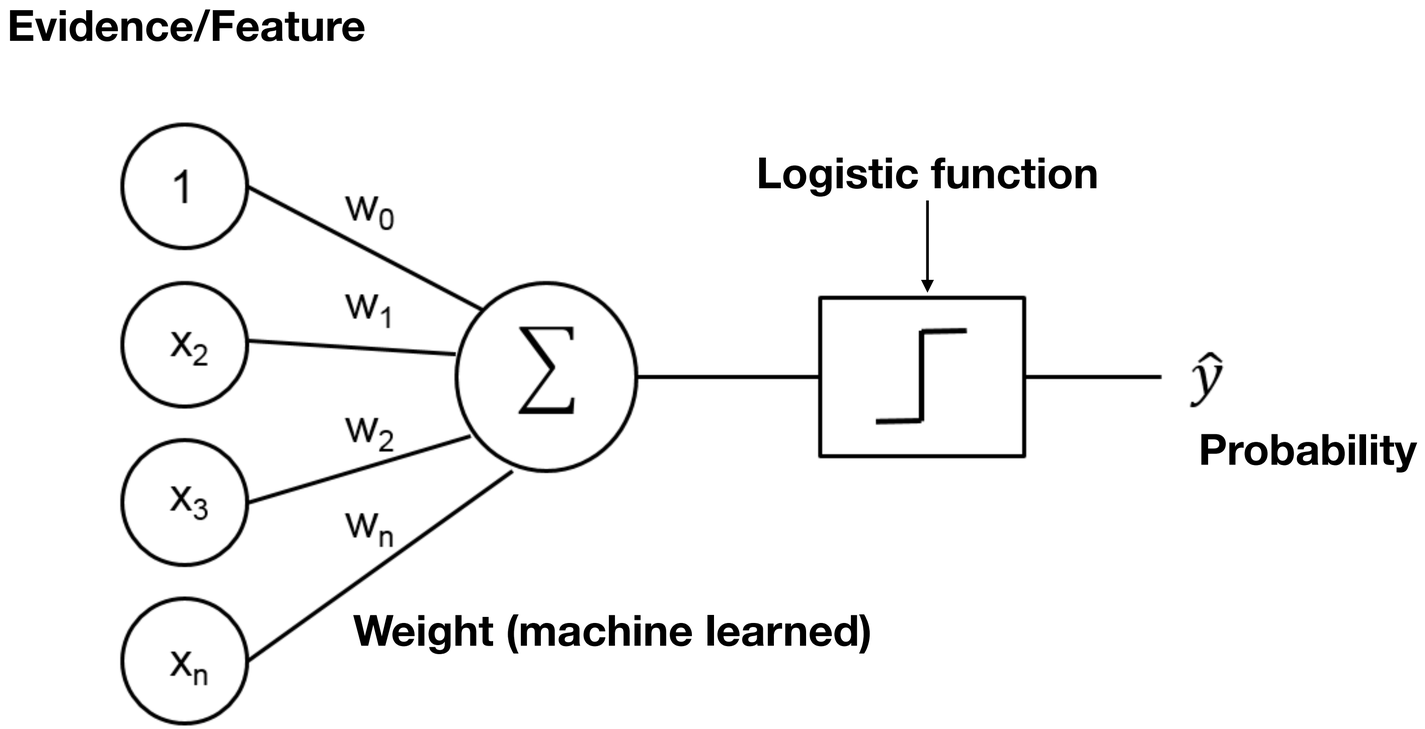

Any Machine Learning Model is a Probability Model!

1. Perceptron Model

History of AI Research: ML

- Probability Model
  - Recognise Handwritten Digit from Image X
    - $P(L = 9 \mid X) + P(L = 8 \mid X) + \ldots + P(L = 0 \mid X) = 1$
  - Sentiment Analysis for Sentence X
    - $P(L = \text{Positive} \mid X) + P(L = \text{Negative} \mid X) + P(L = \text{Neutral} \mid X) = 1$
  - Voice Recognition from audio clip X
    - $P(L = \text{Open my door} \mid X) + P(L = \text{Get weather report} \mid X) + \ldots$

Example

What are the challenges for AI?

P3
Challenges

All AI Models Now Are Probabilistic Models

- $P(L \mid X)$
  - The set L is finite, but the real world is not finite, possibilities are infinite
  - Heavily dependent on training data
    - $< x, l >$ is an observation of Input and Correct label
    - There is no way we can expose the model to all possible input/correct label set
    - Our best ML model cannot do 100% accuracy on all seen data

Concept

P3
Challenges

Our Probabilistic Models are WAAYYYY too Simple

- Neural Network
  - More complex than the Perceptron Model we’ve seen before
  - Best we can do: tens of thousands of "Artificial Neurons"
  - Human Brain: 86 billion
  - Rabbit: 0.5 billion
  - Snail: 11,000, our AI might be less complicated as Snails

Concept

P3
Challenges

We don’t know anything about the real world

- Naming and Identity: The Ship of Theseus
- What do I mean by "chair" in "bring me the chair next door"?
  - In the physical world, there is a chair next door.
  - In my mind, there is a chair next door.
  - In an infinite possible worlds, there is one in which there is a chair next door.

Concept

P3
Challenges

AI is a Philosophical Problem

- And we’ve just started

Concept

P4
Discussions

What are the questions you have about AI?

Summary

P4
Discussions

What I will discuss

- Wednesday
  - Brief introduction to Neural Networks
- Thursday
  - A list of modern problems in AI and their solutions
- I want to know what you want to know!

Concept