Tufts

Class #01:
Introduction to Machine Learning

Machine Learning (COMP 135): M. Allen, 15 Jan. 20

1

What is Artificial Intelligence?

- Historical definition (Dartmouth Workshop on AI, 1956):

> "The study of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

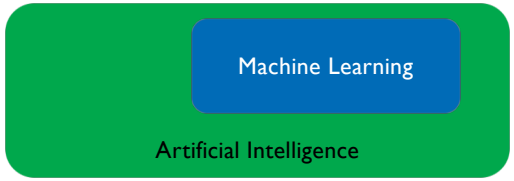

2

Modern AI: An Engineering Enterprise

- Building (partially) autonomous machines for a variety of tasks
  - Construction, transportation, search-and-rescue, exploration…
- Automating intelligence and formalizing knowledge
  - Internet search, expert systems, data mining, …
- Using computational models to understand complex behavior
  - Automated planning, large-scale crowd simulation, traffic analysis, …
- Using computers to discover new information
  - Medical image analysis, intrusion detection, stock market trading, …
- Allowing computers to work better with people
  - Reactive tutoring, automated assistants, "sensitive" GPS systems, …

3

Some Success: Space Exploration

- 1999: NASA allowed the Deep Space I vehicle to be piloted for two days by Remote Agent AI program
- Completely controlled craft operations, over 60 Million miles from Earth

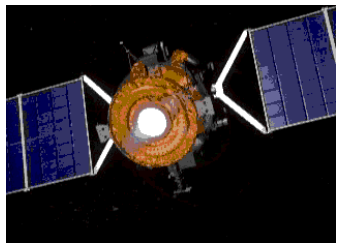

http://ti.arc.nasa.gov/tech/asr/planning-and-scheduling/remote-agent/

4

More Success: Natural Language Processing

- Modern tools have made lots of progress; free commercial software can now translate many highly ambiguous and complex phrases easily

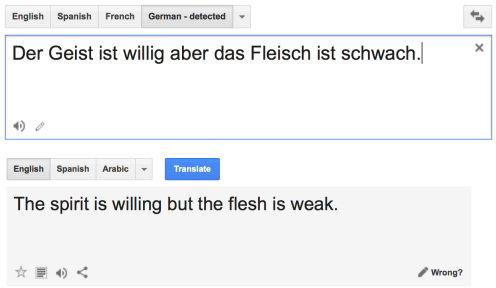

http://translate.google.com

5

Even More Success: Robotic Vehicles

- In the mid-2000's the CMU ALVINN system drove on its own from Washington, DC to San Diego, CA
  - Managed all but 52 of the over-2800 miles
  - Averaged 63 Miles per hour in day, night, bad weather

- In 2007, CMU's Boss system won the DARPA Urban Challenge
  - 60 Miles of urban driving
  - Merging with human traffic
  - Obeying all traffic laws and posted signs
- This work is now part of Google self-driving car system

6

How Do We Define Intelligence?

- It is not clear how "intelligence" should be understood (let alone how to get a machine to behave that way)
- How a ***human being*** might act?

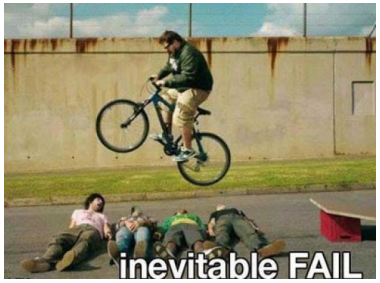

- Or is it some sort of ***ideal rationality***?

7

Turing Test: Intelligence = ***Acting Humanly***

- Alan Turing (1950) "Computing Machinery and Intelligence"
  - Proposed an imitation game
  - Predicted that by 2000, machines could fool average person for 5 minutes, 30% of the time
- One problem: not everyone agrees on the standard proposed by the test, and whether it is meaningful
- In any case, we still haven't got there yet…
  - Loebner prize for convincing bots would award up to $100,000 (and a gold medal) for a truly convincing interactive agent
  - No such agent has ever really been approached

8

What Should an Intelligent System Do?

- Following Turing, we take an operational approach:

> Intelligence is defined by some means of measuring performance in a set task.

- An intelligent system is one that optimizes some measure
- How much it changes things so that it gets closer towards the goals that have been set for it
  - The word-count of error-free text translated
  - Customer satisfaction for automated dialogue systems
  - Hours of accident free, real-time driving
  - Amount of data collected by an autonomous space-vehicle
  - …

9

Defining a Learning Problem

- Suppose we have three basic components:
  1. Set of tasks, *T*
  2. A performance measure, *P*
  3. Data describing some experience, *E*

> A computer program ***learns*** if its performance at tasks in *T*, as measured by *P*, improves based on *E*.
>
> From: Tom M. Mitchell, *Machine Learning* (1997)

10

Next Week

- Linear regression, polynomial regression, gradient methods, and over-fitting
- Readings: *TBA*
- Office Hours: 237 Halligan
  - Monday, 10:30 AM – Noon
  - Wednesday, 10:30 AM – Noon

11