Dr. George Karraz, Ph. D.

Artificial Intelligence

Lecture I

Introduction

Dr. George Karraz, Ph. D.

Prerequisites

- Comfortable programming in language such as C (or C++) or Java
- Some knowledge of algorithmic concepts such as running times of algorithms
- Ideally, some familiarity with probability (we will go over this from the beginning but we will cover the basics only briefly)
- Not scared of mathematics; ideally, some background in discrete mathematics, able to do simple mathematical proofs

Grading

- Midterm exams: 20 %
- Final exam:50 %
- Participation: 5%

Syllabus

- Introduction (lecture I)
- Propositional Logic (Lecture II)
- First-Order Logic (Lecture III)
- Decision Trees (Lecture IV)
- Linear Regression (Lecture V)
- Logistic Regression (Lecture VI)
- Neural Networks (Lecture VII, VIII, IX)
- Support Vector Machine (Lecture X)
- Principal Components Analysis (Lecture XI)
- K-mean Clustering (Lecture XII)

Real AI

- 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 is really difficult to understand at all.
- Special-purpose AI is more doable (nontrivial)
  - E.g., chess playing programs, logistics planning, automated translation, voice recognition, web search, data mining, medical diagnosis, keeping a car on the road, … … … …

Definitions of AI

if our system can be more rational than humans in some cases, why not?

focus on action avoids philosophical issues such as "is the system conscious" etc.

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

- We will 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

"Chinese room" argument [Searle 1980]

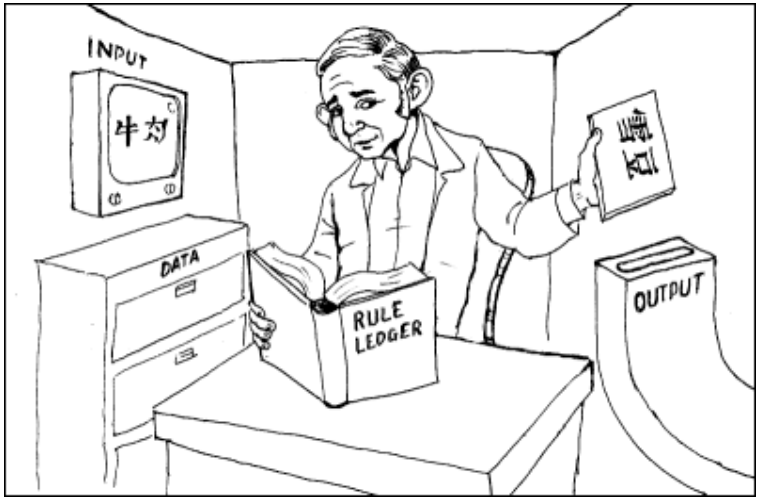

image from http://www.unc.edu/~prinz/pictures/c-room.gif

- 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!

Early history of AI

- 50s/60s: Early successes! AI can draw logical conclusions, prove some theorems, create simple plans… Some initial work on neural networks…
- 70s, 80s: Creation of expert systems (systems specialized for one particular task based on experts' knowledge), wide industry adoption

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

This course

Is focused on general AI techniques that have been useful in many applications