

Artificial Intelligence Programming

Introduction

Chris Brooks

Department of Computer Science
University of San Francisco

Department of Computer Science — University of San Francisco — p.1/71

1-2: Course Mechanics

- Requirements: CS 245 or equivalent.
 - You should be comfortable writing medium-sized programs
 - We'll primarily program in Python.
 - We'll also use pre-existing tools written in Java
- Grading:
 - Homework/lab most weeks (35%)
 - 3 larger, more in-depth projects (30%)
 - Two midterms, plus a final (30%)
 - The dreaded "class participation" component (5%)
- Attendance is required.

Department of Computer Science — University of San Francisco — p.2/71

1-3: Course Mechanics

- Text: AI/MA, 2nd edition
- Late Policy
- Collaboration

Department of Computer Science — University of San Francisco — p.3/71

1-4: Topics

- Topics we'll cover:
 - Uninformed Search
 - Heuristic Search
 - Genetic Algorithms and Simulated Annealing
 - Knowledge Representation
 - Logic and Inference
 - Ontologies
 - Probabilistic Reasoning
 - Learning
 - Support Vector Machines

Department of Computer Science — University of San Francisco — p.4/71

1-5: non-AI topics

- Along the way, you'll also get exposed to some non-AI topics:
 - HTTP
 - HTML, XML
 - Semantic Web
 - Text classification
 - Python
 - and more!

Department of Computer Science — University of San Francisco — p.5/71

1-6: Topics we won't get to ...

- Stuff we probably won't have time for:
 - Robotics
 - Vision
 - Natural Language

Department of Computer Science — University of San Francisco — p.6/71

1-7: AI as a research field vs. AI in the popular press

- Ordinary folks typically think that AI is all about building the Terminator.
- AI research is typically **much** more narrowly focused.
 - Why is that?
- AI also has a perception of being a “failed research area”
 - Again, this hinges on how you define the field.
- It's important to carefully define the problems you want to solve.

1-8: What is AI?

- 'concerned with intelligent behavior in artifacts' - Nilsson
- 'study of the design of intelligent agents' - Poole
- 'Creating machines that perform functions that require intelligence when performed by people.' - Kurzweil
- 'The automation of activities we associate with human thinking, such as decision-making, problem-solving, learning' - Bellman
- 'The study of the computations that make it possible to perceive, reason, and act' - Winston

1-9: Reducing this ...

- Can machines think?
- Let's parse this -
 - *Can* - is it possible to build a machine that thinks?
 - Searle: Thought is a byproduct of human makeup - nothing less complex than a human can actually think.
 - “Brains cause minds” - we are intelligent because we have evolved to function in a very complex, spatially rich world.
 - Being completely embedded in the world is necessary for intelligence.
 - Is true machine consciousness possible (or desirable)?

1-10: Reducing this ...

- *Can* machines think?
- Newell and Simon: Physical Symbol System Hypothesis
- A physical symbol system has the *necessary* and *sufficient* means for general intelligent action.
- Symbols: letters, numbers, words, variables, etc.
- A symbol system is a collection of symbols, plus the rules or means to combine them into new symbols.
 - e.g. A CPU or a Turing machine
- Symbols are used to construct a *language* or *representation*

1-11: PSS Hypothesis

- necessary: symbol manipulation is required for intelligence.
 - This is controversial; some folks feel that intelligence can be achieved purely through subsymbolic manipulation of signals.
 - For example, recognizing someone's face is not a symbolic task.
- Sufficient - no other system or capability is required for intelligence if one has a PSS.
 - Also controversial - some people argue that the agent must be grounded in a sufficiently rich world.

1-12: Reducing this ...

- Can machines think?
- *machines* - What is a machine?
- We typically look at Turing-equivalent machines.
- Perhaps other types of machines can be intelligent.
- Cells can be programmed using DNA
- We can argue that the human body is just a very complex machine.

1-13: Reducing this ...

- Can machines think?
- *think* - This is the hardest one yet.
 - What does it mean for a machine (or a person) to think?
 - How can we as scientists say that some external thing is thinking?
 - How can we say that it is intelligent?

1-14: Thinking

- “A system is intelligent if it acts like a human”
- I assume people are intelligent because they act appropriately.
- Can we apply this metric to a computer?
- Turing test
- Response: Searle’s Chinese Room experiment

1-15: Problems with the Turing Test

- Sufficient, but not necessary.
- Only measures language and conversation.
 - What about learning, problem solving, vision, etc ...
- Test is subjective; depends on the questions asked.
- Do we really want an intelligent machine to be indistinguishable from a human?
 - Can’t do square roots, gets tired, can’t remember lists, etc.
- What about complicated mimicry programs? Do they count?

1-16: Thinking

- “A system is intelligent if it *thinks* like a human”
- If a system uses the same reasoning processes as a human, then it is intelligent.
- How do humans think? This is an ongoing research problem.
- Researchers in cognitive science and psychology construct computational models of human problem solving.
 - Gives a way to ‘look inside’ a model of the human brain
 - Goal: better understand the human brain/mind.
- Again, do we really want a machine to think ‘like a human’? Make logical errors, misestimate distances and probabilities, draw faulty conclusions, etc.

1-17: Thinking

- “A system is intelligent if it thinks *rationally*”
- An intelligent system is one that follows sound reasoning processes that always lead to correct outcomes.
- This leads to the study of logic and formal reasoning
- This used to be the dominant approach in AI.
- However, logic has its problems.
 - Formalizing common-sense knowledge
 - Dealing with uncertainty
 - Computational issues

1-18: Thinking

- “A system is intelligent if it *acts* rationally”
- Currently the most popular definition.
- Sidesteps all the sticky questions.
- A rational agent is one that acts so as to achieve the best possible (expected) outcome, given its knowledge and ability.
- Not omnipotent or omniscient
- This lets us work on smaller problems that require ‘intelligence’
- We can build agents that do certain tasks intelligently without having human-level intelligence.

1-19: AI as an engineering problem

- We'll focus on the engineering and design problems in AI
 - How do we build a program that can learn to do a task?
 - What sorts of knowledge is needed for a program to solve a word problem?
- Less emphasis on cognitive and philosophical issues.

1-20: Successes of AI

- A common criticism of AI is that it's a failure.
- Early AI researchers were overly optimistic.
 - Solutions to simple problems didn't scale up.
- "As soon as something works, it's no longer AI."
 - Voice recognition
 - Face recognition
 - Optimal scheduling
 - Automated translation
 - Factory robots
 - Object-oriented programming

1-21: Successes of AI

- Some successful AI programs:
 - Deep Blue
 - Chinook
 - TD-Gammon
 - Remote Agent
 - Expert systems (Dendral, MYCIN)
 - Alvin
 - Soar
 - Logistics planning

1-22: Current AI Applications

- Information Processing
- Expert Systems
- Space exploration
- Automated factories
- Robotic search-and-rescue
- Games and immersive environments
- Electronic Commerce
- Tutoring systems
- Others? ...