




## Introduction

CIS\*3700 (Winter 2007)





## Contact Information

- Instructor:
  - Fei Song (Thornbrough 1388, ext. 58067)
  - Email: [fsong@cis.uoguelph.ca](mailto:fsong@cis.uoguelph.ca)
  - Office Hours: 2-3pm, Tuesday and Thursday
- Teaching Assistants:
  - Robert Collier and Wei Wang
  - Email: [ta3700@cis.uoguelph.ca](mailto:ta3700@cis.uoguelph.ca)
  - Office Hour: TBA



## Overview



- Fundamental concepts and techniques:
  - Agents and environments
  - Search
  - Knowledge representation and reasoning
  - Learning
- Specific subfields: game search, expert systems, decision tree learning, statistical natural language processing.



## Evaluation



• Assignments (4 x 10%)	40%
• Midterm	20%
• Final exam	40%

- Late policy: late submissions are acceptable, but there will be a reduction of 10% for one day late, 25% for two days late, and 50% for three days late. After that, no marks will be given.





## Textbook and References

- Stuart Russell and Peter Norvig. "Artificial Intelligence: A Modern Approach". Second Edition. Prentice-Hall, 2003.
- References:
  - Thomas Dean, James Allen, and Yiannis Aloimonos. "Artificial Intelligence: Theory and Practice". Benjamin/Cummings Publishing, 1995.
  - Matt Ginsberg. "Essentials of Artificial Intelligence". Morgan Kaufmann Publishers, 1993.



## Why is AI Exciting?

- Powerful: not just to understand but also to build intelligent entities.
- New: work started in mid-1940's and the name was coined in 1956 by John McCarthy.
- Broad: a wide range of subfields, such as game search, expert systems, learning, information extraction and retrieval, perception and robotics, neural networks, fuzzy logic, and genetic computing.



## What is AI?

- Four categories of definitions:

	Human-centered	Rationality-centered
Thought-centered	Systems that think like humans	Systems that think rationally
Behavior-centered	Systems that act like humans	Systems that act rationally

## Acting Humanly: the Turing Test Approach

- Short version: AI is the enterprise of constructing an intelligent artifact.
  - Intelligence: things people do well exhibit intelligence, but not all of them (e.g., bodily functions)
  - Turing test: a functional definition.
- Long version: AI is the enterprise of constructing a physical symbol system that can reliably pass the Turing test (Ginsberg, 1993).

## Thinking Humanly: the Cognitive Modeling Approach

- Require an understanding of the actual workings of human minds.
  - Cognitive science: construct precise and testable theories of the workings of human minds through computer models and psychological experiments.
- Is it necessary for intelligent entities to duplicate human thought process?
  - The distinction has helped both AI and Cognitive Science, especially in areas of vision and natural language.

## Thinking Rationally: the "Laws of Thought" Approach

- Build intelligent systems using logic:
  - Logic rules such as "syllogisms": Aristotle is a man; all men are mortal; so Aristotle is mortal.
  - In principle, any solvable problem can be solved by logic.
- Not always practical:
  - Not easy to express knowledge with 100% certainty in logical notation.
  - Intractable: quickly exhaust the current computing power and resources.

## Acting Rationally: the Rational Agent Approach

- An rational agent acts so as to achieve the best outcome, or when there is uncertainty, the best expected outcome.
  - No need to imitate human, either internally or externally.
  - Not necessarily based on logic.
  - Perfect rationality may be computationally too expensive, but a good starting point for limited rationality.