

Introduction to Artificial Intelligence for Computer Gaming

G. Michael Youngblood, Ph.D.
Computer Science
The University of North Carolina at Charlotte

What is Artificial Intelligence (AI)?

What is Artificial Intelligence (AI)?

“The exciting new effort to make computers think ... *machines with minds*, in the full and literal sense” (Haugeland, 1985)

“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning...” (Bellman, 1978)

“The study of mental faculties through the use of computational models”
(Charniak and McDermott, 1985)

“The study of the computations that make it possible to perceive, reason, and act”
(Winston, 1992)

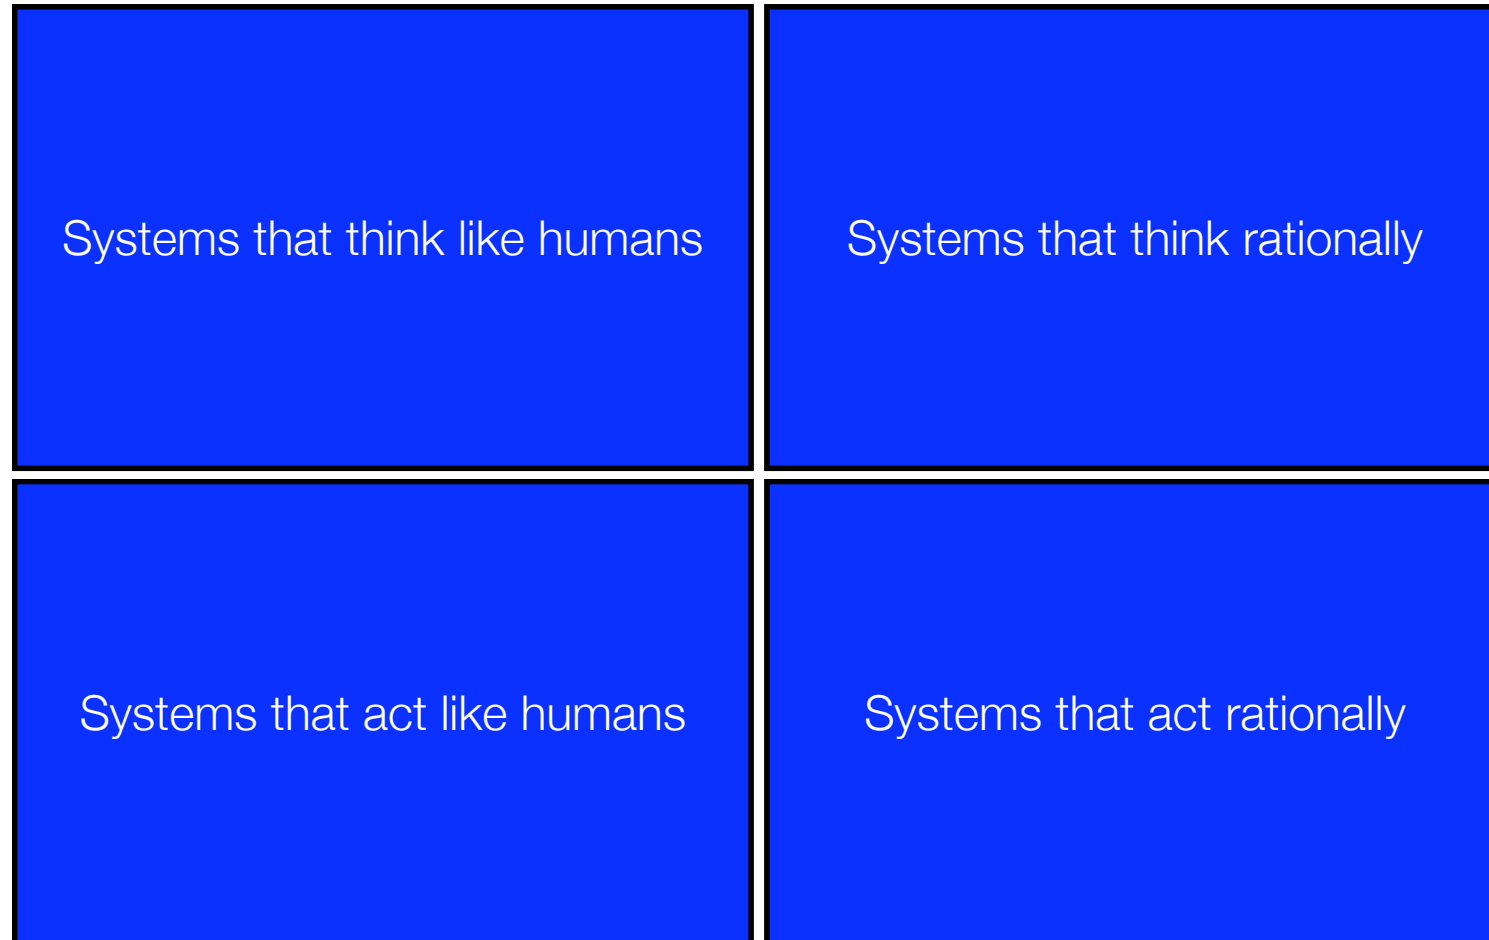
“The art of creating machines that perform functions that require intelligence when performed by people” (Kurzweil, 1990)

“The study of how to make computers do things at which, at the moment, people are better” (Rich and Knight, 1991)

“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes” (Schalkoff, 1990)

“The branch of computer science that is concerned with the automation of intelligent behavior”
(Luger and Stubblefield, 1993)

What is Artificial Intelligence (AI)?




Source: Russell and Norvig, p. 5

Acting humanly: The Turing test

Turing (1950) “Computing machinery and intelligence”:

- ◇ “Can machines think?” → “Can machines behave intelligently?”
- ◇ Operational test for intelligent behavior: the **Imitation Game**



- ◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- ◇ Anticipated all major arguments against AI in following 50 years
- ◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly: Cognitive Science

1960s “cognitive revolution”: information-processing psychology replaced prevailing orthodoxy of behaviorism

Requires scientific theories of internal activities of the brain

– What level of abstraction? “Knowledge” or “circuits”?

– How to validate? Requires

1) Predicting and testing behavior of human subjects (top-down)

or 2) Direct identification from neurological data (bottom-up)

Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI

Both share with AI the following characteristic:

**the available theories do not explain (or engender)
anything resembling human-level general intelligence**

Hence, all three fields share one principal direction!

Thinking rationally: Laws of Thought

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:

notation and **rules of derivation** for thoughts;
may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:

- 1) Not all intelligent behavior is mediated by logical deliberation
- 2) **What is the purpose of thinking?** What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

Rational agents

An **agent** is an entity that perceives and acts

This course is about designing **rational agents**

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: **computational limitations make perfect rationality unachievable**

→ design best **program** for given machine resources

A.L.I.C.E

The Chinese Room

- Imagine you are sitting in a room with a library of rule books, a bunch of blank exercise books, and a lot of writing utensils. Your only contact with the external world is through two slots in the wall labeled ``input" and ``output". Occasionally, pieces of paper with Chinese characters come into your room through the ``input" slot. Each time a piece of paper comes in through the input slot your task is to find the section in the rule books that matches the pattern of Chinese characters on the piece of paper. The rule book will tell you which pattern of characters to inscribe the appropriate pattern on a blank piece of paper. Once you have inscribed the appropriate pattern according to the rule book your task is simply to push it out the output slot.
- By the way, you don't understand Chinese, nor are you aware that the symbols that you are manipulating are Chinese symbols.
- In fact, the Chinese characters which you have been receiving as input have been questions about a story and the output you have been producing has been the appropriate, perhaps even "insightful," responses to the questions asked. Indeed, to the outside questioners your output has been so good that they are convinced that whoever (or whatever) has been producing the responses to their queries must be a native speaker of, or at least extremely fluent in, Chinese.

Searle's Chinese Room Argument

- The question that Searle asks is-Do you understand Chinese?
- Searle says NO
- **What do you think?**

Searle's Chinese Room Argument

- The question that Searle asks is-Do you understand Chinese?
- Searle says NO
- **What do you think?**
- Is this a refutation of the possibility of Artificial Intelligence?
- The Systems Reply
 - The individual is just part of the overall system, which does understand Chinese
- The Robot Reply
 - Put the same capabilities in a robot along with perceiving, walking, and any other relevant activities. Such a computer would seem to have genuine understanding and mental states.

The Loebner Prize

Jabberwacky

AI prehistory

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

Potted history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1952–69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: "AI Winter"
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth
"Nouvelle AI": ALife, GAs, soft computing
- 1995– Agents, agents, everywhere . . .
- 2003– Human-level AI back on the agenda

State of the art

Which of the following can be done at present?

- ◇ Play a decent game of table tennis
- ◇ Drive safely along a curving mountain road
- ◇ Drive safely along Tryon Street
- ◇ Buy a week's worth of groceries on the web
- ◇ Buy a week's worth of groceries at Harris Teeter
- ◇ Play a decent game of bridge
- ◇ Discover and prove a new mathematical theorem
- ◇ Design and execute a research program in molecular biology
- ◇ Write an intentionally funny story
- ◇ Give competent legal advice in a specialized area of law
- ◇ Translate spoken English into spoken Swedish in real time
- ◇ Converse successfully with another person for an hour
- ◇ Perform a complex surgical operation
- ◇ Unload any dishwasher and put everything away

What about AI for computer games?

- AI in games
- AI, Ai, and al

What is the goal of AI in computer games?

What is the goal of AI in computer games?

- Provide an opponent, companion, or game entity that
 - appears to imitate human components
 - acts intelligently
 - enhances the user experience
 - has complex behaviors
 - “react[s] autonomously to environmental stimuli” (Watt & Policarpo, p. 484)
 - “communicate[s] with other agents in the game” (Watt & Policarpo, p. 484)
- Commonly referred to as *agents*, *characters*, or “*bots*”

fin.