Artificial Intelligence

Text Book and Study Materials

- Suggested Book:

  Stuart Russell and Peter Norvig. 2009. Artificial Intelligence: A Modern Approach (3rd ed.). Prentice Hall Press, Upper Saddle River, NJ, USA.

Slides will be provided in SmartEdu/Google Classroom.

Lecture 1

Fundamentals of AI

What is AI??

AI: A Vision

- Could an intelligent agent living on your home computer **manage your email**, coordinate your **work and social activities**, help **plan your vacations**…… even **watch your house** while you take those well planned vacations?

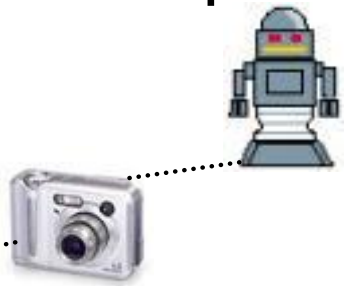

Main Goals of AI

- **Represent** and **store** knowledge
- **Retrieve** and **reason** about knowledge
- **Behave** intelligently in complex environments
- Develop interesting and useful **applications**
- **Interact** with people, agents, and the environment

Why AI?

- **Engineering:** To get machines to do a wider variety of useful things
  - e.g., understand spoken natural language, recognize individual people in visual scenes, find the best travel plan for your vacation, etc.
- **Cognitive Science:** As a way to understand how natural minds and mental phenomena work
  - e.g., visual perception, memory, learning, language, etc.
- **Philosophy:** As a way to explore some basic and interesting (and important) philosophical questions
  - e.g., the mind body problem, what is consciousness, etc.

Foundations of AI

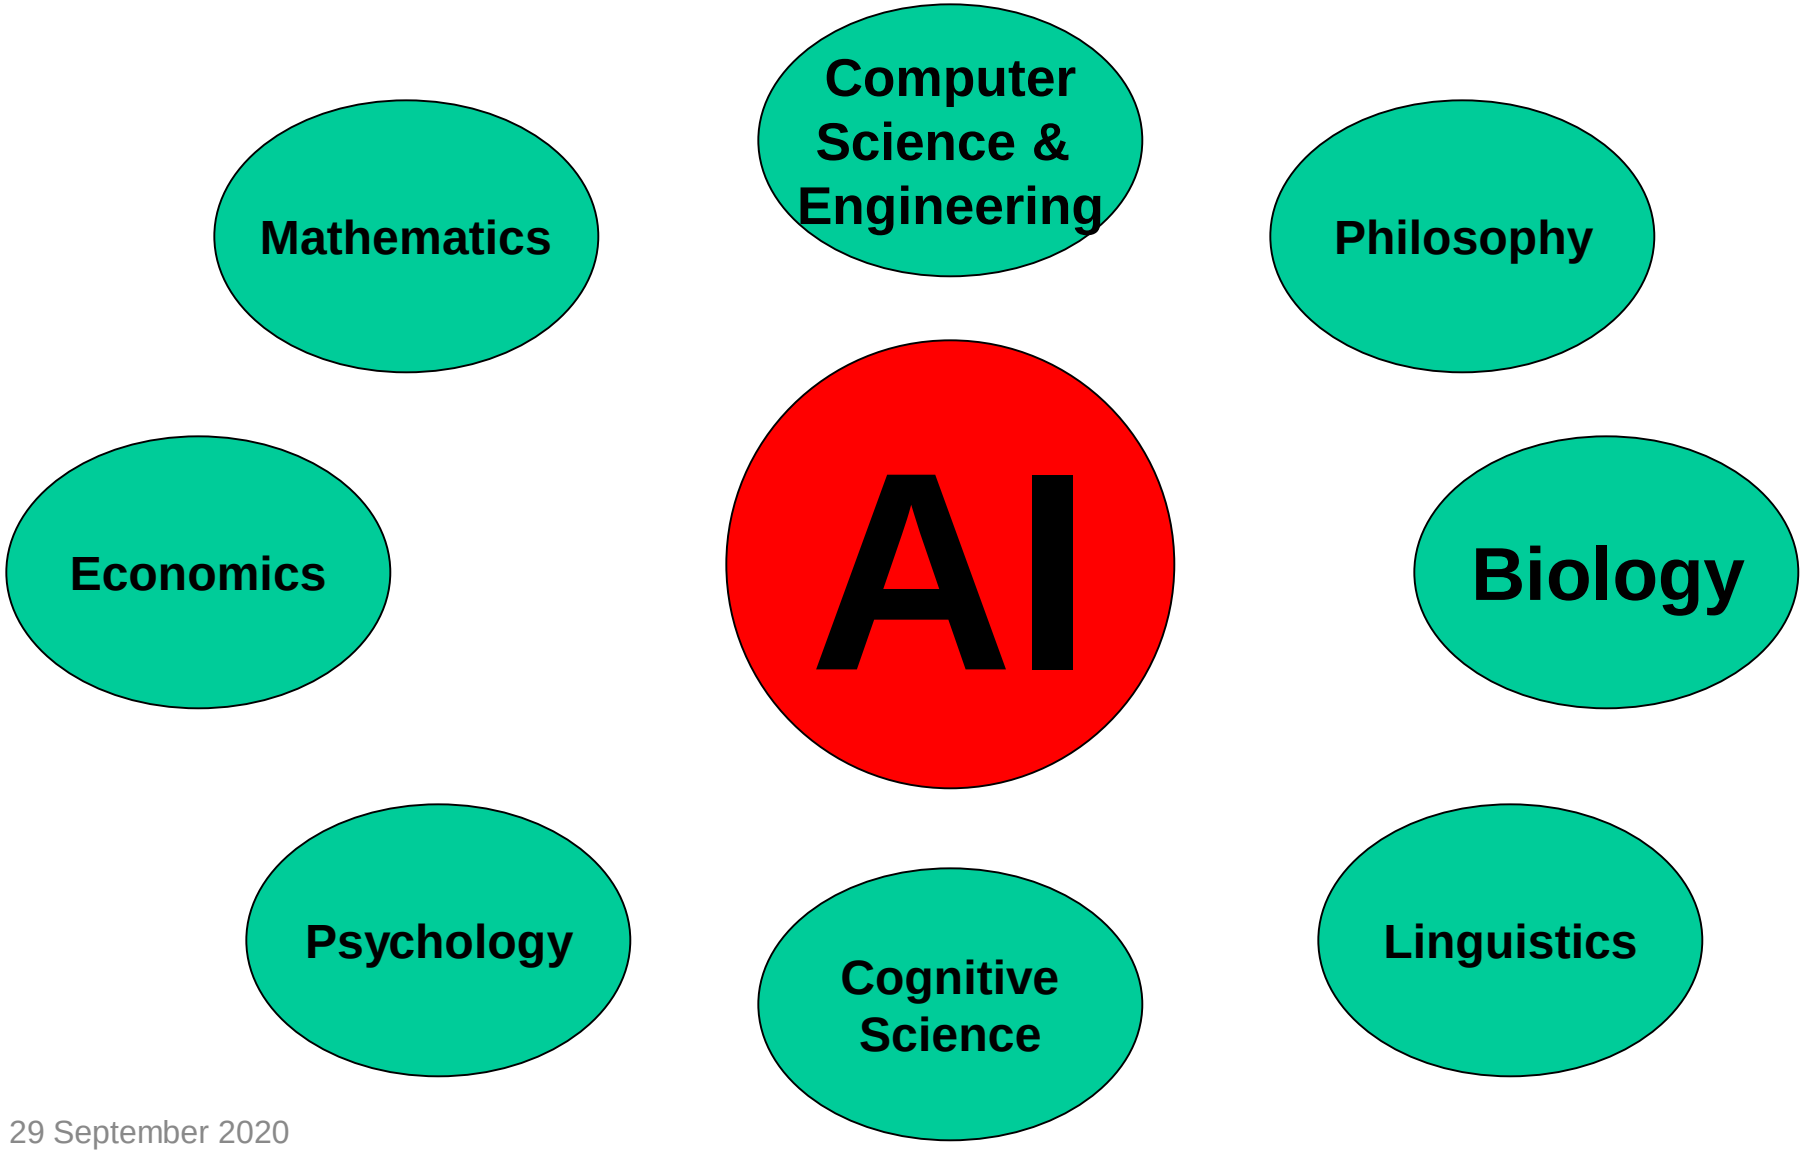

History

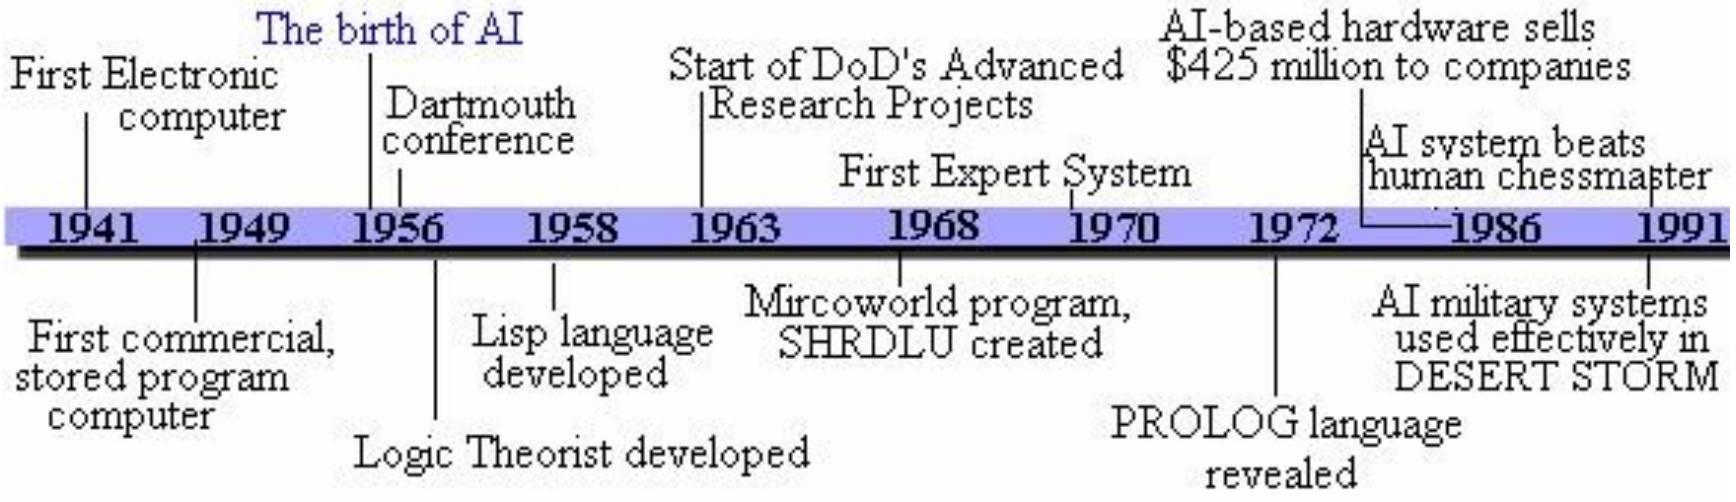

- **1997: Deep Blue beats Garry Kasparov (world champion)**
- 1998: Founding of Google
- **2000: Interactive robot pets**
- 2004: First DARPA Grand Challenge robot race
- **2004: Commercial recommender systems (TIVO, amazon.com)**
- **2007: Checkers is solved!**

Eliza

- ELIZA: A program that simulated a psychotherapist interacting with a patient and successfully passed the Turing Test.
- Coded at MIT during 1964-1966 by Joel Weizenbaum.
- First script was DOCTOR.
  - The script was a simple collection of syntactic patterns not unlike regular expressions
  - Each pattern had an associated reply which might include bits of the input (after simple transformations (my $\rightarrow$ your)
- Weizenbaum was shocked at reactions:
  - Psychiatrists thought it had potential.
  - People unequivocally anthropomorphized.
  - Many thought it solved the NL problem.

Eliza

- This is a transcript of a conversation with an Eliza-like program found in Emacs
- Try typing

Meta-x doctor

in your favorite Emacs.

- **I am the psychotherapist. Please, describe your problems. Each time you are finished talking, type RET twice.**
- **i'm having a hard time with my students**
- **Why do you say that?**
- **well, they keep asking me for money**
- **Maybe your plans have something to do with this.**
- **what plans?**
- **Is it because of your life that you say what plans?**
- **No, it's because I didn't know what you were talking about**
- **Is it because you didn't know what I was talking about that you came to me?**
- **no, it's because i wanted to see if you could pass the turing test**
- **Does it bother you that you wanted to see if I could really pass the turing test?**
- **…..**
- **How do you reconcile problems at school?**
- **This is a waste of time. Good bye**
- **My secretary will send you a bill.**

Big Questions

- Can machines think?
- If so, how?
- If not, why not?
- What does this say about human beings?
- What does this say about the mind?

What's Easy and What's Hard?

- It's been easier to mechanize many of the high-level tasks we usually associate with "intelligence" in people
  - e.g., symbolic integration, proving theorems, playing chess, medical diagnosis
- It's been very hard to mechanize tasks that lots of animals can do
  - walking around without running into things
  - catching prey and avoiding predators
  - interpreting complex sensory information (e.g., visual, aural, …)
  - modeling the internal states of other animals from their behavior
  - working as a team (e.g., with pack animals)
- Is there a fundamental difference between the two categories?

Turing Test

- Three rooms contain a person, a computer, and an interrogator.
- The interrogator can communicate with the other two by teleprinter.
- The interrogator tries to determine which is the person and which is the machine.
- The machine tries to fool the interrogator into believing that it is the person.
- If the machine succeeds, then we conclude that the machine can think.

What Can AI Systems Do?

- **Here are some example applications**
- **Computer vision:** face recognition from a large set
- **Robotics:** autonomous (mostly) automobile
- **Natural language processing:** simple machine translation
- **Expert systems:** medical diagnosis in a narrow domain
- **Spoken language systems:** ~1000 word continuous speech
- **Planning and scheduling:** Hubble Telescope experiments
- **Learning:** text categorization into ~1000 topics
- **User modeling:** Bayesian reasoning in Windows help (the infamous paper clip…)
- **Games:** Grand Master level in chess (world champion), perfect play in checkers, professional-level Go players

What Can't AI Systems Do Yet?

- Understand natural language robustly (e.g., read and understand articles in a newspaper)
- Surf the web
- Interpret an arbitrary visual scene
- Learn a natural language
- Play Go as well as the best human players
- Construct plans in dynamic real-time domains
- Refocus attention in complex environments
- Perform life-long learning

Exhibit true autonomy and intelligence!

Who Does AI?

- Academic researchers (perhaps the most Ph.D.-generating area of computer science in recent years)
  - Some of the top AI schools: CMU, Stanford, Berkeley, MIT, UIUC, UMd, U Alberta, UT Austin, ... (and, of course, UMBC!)
- Government and private research labs
  - NASA, NRL, NIST, IBM, AT&T, SRI, ISI, MERL, ...
- Lots of companies!
  - Google, Microsoft, Honeywell, Teknowledge, SAIC, MITRE, Fujitsu, Global InfoTek, BodyMedia, ...