CPSC 625-600 Artificial Intelligence

- Instructor: Yoonsuck Choe
- TA: Randall Reams
- Grader: Anavil Tripathy
- See course web page for contact info and hours.
- All communications out of the class will be through email registered on email.tamu.edu, and the announcements on the web page, so regularly check out the web page.
- Class notes will be available on the web 24 hours prior to the class. It is **your responsibility** to print it out and bring it to the class. **http://courses.cs.tamu.edu/choe/16fall/625/lectures/**

Syllabus

http://courses.cs.tamu.edu/choe/16fall/625/625.pdf

See the syllabus for details on grading, course policy, etc.

Things You May Need

- Students with disability:
  Please contact the department office (HRBB 3rd floor) for assistance. See the syllabus for the full information.
- Computer (UNIX) accounts:
  If you don't have one, get one:
  `https://wiki.cse.tamu.edu/index.php/Getting_Started_Guide`

What is Intelligence

Textbook Definitions

- Thinking like humans
- Acting like humans
- Thinking rationally
- Acting rationally ←

However, it depends on the definition: **whatever the (intelligence) test tests.**

What is AI?

A folk (popular) view of AI

Universal studio's movie "Terminator" (bottom)

But Really, What is AI?

Diverse areas: http://www.aaai.org

- Problem solving
- Reasoning
- Theorem proving
- Learning
- Planning
- Knowledge representation
- Perception and Robotics
- Agents
- **and more**

Approaches

Two basic stances

- Strong AI:
  1. Build something that actually thinks intelligently.
  2. Simulation of thought would give rise to the pheonmenology of thought (i.e., how it feels like to think).
- Weak AI:
  1. Build something that behaves intelligently.
  2. Not worried about its feelings.

Problems

- Strong AI:
  Hard to determine if something is really consciously intelligent or not (the **other minds problem** in philosophy).
- Weak AI:
  Utility of the result is limited by the stated goal. Hard to achieve a **general usefulness** as in true intelligence.

How to do AI

Why not engineer AI, in the same way people engineered airplanes?

1. Flight
   goal is simple:
   - You know when a thing is flying.
2. Intelligence
   goal is complex and hard to define clearly:
   - Intelligence is a collection of many abilities.

There are many ways to meet a single clear goal (flight), but there can be only a small number of ways to simultaneously meet a huge number of unclearly defined goals (intelligence).

Why not Follow the Plane-Model?

Certain things may seem physically impossible (in terms of efficiency, etc.): e.g. the flight of flies.

- Flapping their wings cannot generate enough lift (for their body weight), but they do fly!
- Jet turbines cannot explain how the flies achieve such an impossible feat.
- Recent observation:
  Flies gyrate their wings to generate a vortex to create greater lift.

Moral: if you fail to build the impossible, study an existing solution.

How Flies Fly

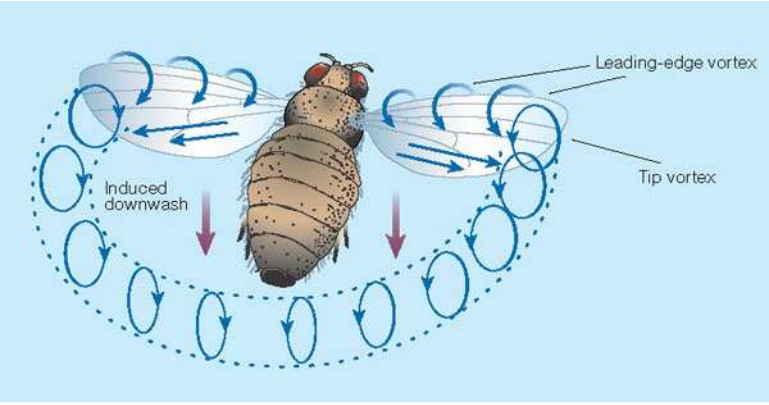

Source: Naturemagazine

Then, How to do AI?

Instructor's perspective:

- Importance of studying brain function.
- Influence of environmental regularities on brain development and function.
- Interaction of the brain with the environment through action.

We must think about the more fundamental issues from time to time when research seems to be at a dead-end.

Back to Reality

Let's be realistic. :-)

- Study strategies employed by humans in dealing with real-world problems.
- These include all the topics listed earlier.
- The background you learn in this course will enable you to appreciate the deepness of the problems, and to pursue further interest in AI, and in human and machine intelligence in general.

Overview

- Related academic disciplines
- History of AI
- Hard Problems
- Current Trends

Foundations of AI

- Philosophy
- Mathematics
- Psychology
- Cognitive Science
- Linguistics
- Neuroscience

Philosophy of Mind

The mind-body problem:

- Dualism: Mind and body are separate entities.
- Monism: Only mind or body exist, but not both
  1. Idealism: all things are mental
  2. Materialism: all things are material
- Epiphenomenalism: mental phenomena are just side-effects of physical change in the brain (i.e. they do not have causal power over behavior, like the smoke coming from a steam engine).

Too many variations to mention all.

Mathematics

- Algorithm (al-Khowarazmi)
- Boole
- Hilbert
- Gödel: Incompleteness theorem
- Turing: Halting problem
- Cook and Karp: P, NP, and the like

Representation/Interpretation, Symbol/Computing: the computer/software metahpore.

Psychology

- Behaviorism: stimulus-response and conditioning
- Functionalism: internal representations and processes. Implementation independent.
- Perceptual psychology: vision, audition, etc.
- Cognitive psychology: cognition as information processing.
- Holistic vs. localist debate: emergent vs. simple summation.

Linguistics

- WW II : machine translation.
- Phonetics, syntactic theory, semantics, discourse, etc.
- Innate vs. learned? : Chomsky
- Syntax: finite automata, context free grammar, etc.
- Semantics: semantic nets
- Sub-symbolic: self-organizing maps, episodic memory, recurrent neural nets, etc.

Cognitive Science

Interdisciplinary field studying human perception and cognition, ranging over:

- Neuroscience
- Behavioral science
- Social science
- Psychology
- Computational science
- Information theory
- Cultural studies

Neuroscience

- Staining: Golgi, Nissl
- Hubel and Wiesel: orderly structure of cat visual cortex
- PET scans and CAT scans: localizing functional modules
- fMRI imaging: cognitive and perceptual tasks
- Optical imaging: orderly structure
- TMS: zap and numb your brain
- Direct current stimulation

Connections

Scientific discoveries came from observing unexpected connections:

- Apple and gravity
- Cloud chamber and the discovery of subatomic particles
- Looms with punch-cards and modern computers

History of AI (I)

Gestation (1943–1956)

- McCulloch and Pitts: early neural nets
- Minsky and Papert: limitations of perceptron
- Newell and Simon: physical symbol system hypothesis
  - Logic Theorist
- Dartmouth Workshop (1956): AI was born
  It is 60 years old (2016)!
  `http://en.wikipedia.org/wiki/AI@50`

History of AI (II)

Early successes (1952–1969)

- General problem solver
- McCarthy: LISP
- Toy domains: ANALOGY, STUDENT (algebra).
- Widrow and Hoff: adalines
- Rosenblatt: perceptrons

History of AI (III)

The 60's and 70's

- ELIZA: pattern-matching-based NLP
- Genetic algorithms
- Knowledge-based systems: scientific domain, engineering domain, natural language
- Minky and Papert (1969): limitations of perceptron

The 80's:

- 5th generation AI – Prolog.
- Neural networks: Neocognitron, Convolutional Neural Networks, Back Propagation, etc.

History of AI (IV)

The 90's

- Probabilistic approaches. Bayesian networks. Hidden Markov Models.
- Support vector machines.
- Machine learning comes to central stage.
- Intelligent agents
- Statistical NLP

The 2000's – present

- Artificial General Intelligence (AGI)
- Big data, web, information retrieval
- Human-based computation: Recaptcha, Amazon Mechanical Turk
- Deep learning

History of AI (V)

50th anniversary in 2006: http://en.wikipedia.org/wiki/AI@50

- Some quotes from the 50th anniversary event (Rodney Brooks):
  - the social sophisitication of 10-year-old
  - the manual dexterity of a 6-year-old
  - the language ability of 4-year-old
  - the visual object recognition of a 2-year-old

Hard Problems (I)

- Physicalism, materialism, and naturalism: brain causes mind.
- Functionlism: if it functions in the same way, a silicon brain can also be conscious.
- Dualism and homunculus: the Cartesian theatre.
- Wide vs. narrow content: real correspondence with external world, or limited to experiential state?
- Intentionality (aboutness): how can we believe in things that do not exist, such as Poseidon.

Hard Problems (II)

- Semantic content and syntactic symbols: how can syntacic constructs possess intenionality?
- Symbol grounding: sensory devices produce grounded symbols, and composite symbols can be constructed on top of those.
- Problem of qualia: The feeling of perceiving, thinking, etc.
- Turing test and Searle's Chinese Room
  - system reply
  - robot reply

Hard Problems (III)

- However, the assumption that a collection of unconscious things are unconscious is invalid: think about organic vs. inorganic, life vsinanimate matter.
- Searle's point of view: mind is an emergent phenomena of the neural substrate (biological naturalism).

Current Trends

- Machine Learning
- Neural networks and statistical methods
- Genetic algorithms (Evolutionary computation)
- Embodied cognition, Dynamical systems approach
- Developmental robotics
- Computational Neuroscience
- Distributed Agents
- Connectomics
- Consciousness

Latest Achievements

- AI in computer games (1980's–current).
- Game playing (Chess, etc.): IBM's Deep Blue (1997), Google's AlphaGo (2015).
- Jeopardy: IBM's Watson (2011)
- Deep learning: Autonomous driving, Image recognition, Language translation, Speech recognition, etc. etc. Hinton, Schmidhuber, Le Cunn, Bengio, etc.

Remote Access

- SSH (secure)
- http://www.freessh.org/
- For windows, use PUTTY.EXE
  - use the SSH mode.
- On-campus:
  sun.cs.tamu.edu, etc.
- Off-campus:
  Only by using TAMU VPN
- Use TAMU vpn to access other unix hosts.

Little Bit of LISP

http://www.cs.tamu.edu/faculty/choe/courses/16fall/lisp-quickref.html

- CMUCL: Carnegie Mellon University Common LISP
- At the $*$ prompt, just type the expressions.

```
unix:~/> lisp
CMU Common Lisp CVS Head 2003-07-01 16:23:01, running on unix
With core: /usr/local/lib/cmucl/lib/lisp.core
Dumped on: Tue, 2003-07-01 16:01:00-05:00 on empic5
See <http://www.cons.org/cmucl/> for support information.
Loaded subsystems:
    Python 1.1, target UltraSparc/Solaris 7
    CLOS based on Gerd's PCL 2003/06/18 09:23:09
* (+ 10 20)

30
* (quit)
unix:~/>
```


Next Time

- LISP