

CSCE 625-600 Artificial Intelligence

- Instructor: Yoonsuck Choe
- 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/20spring/625/lectures/>

1


Syllabus

<http://courses.cs.tamu.edu/choe/20spring/625/625.pdf>

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

2

Textbooks



- Russell & Norvig (R&N) 3rd Edition: Main textbook.
- Brighton & Selina (B&S) : AI overview, including philosophical issues
 - Note: Introducing Artificial Intelligence: A Graphic Guide is the same book.

3

What is Intelligence

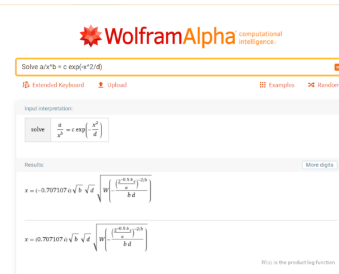

Textbook Definitions

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

“ We define AI as the study of agents that receive percepts from the environment and perform actions. ” – R&N

4

Who Is More Intelligent?



- Math Prof. vs. Kid vs. Mathematica (Wolfram)

5

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 recently: Deep Learning

7

AI in the Media



From <https://www.newworldai.com/>

6

Approaches

Two basic stances

- Strong AI:
 1. Build a machine capable of thought, consciousness and emotions, equalling that of humans.
 2. Humans are no more than elaborate computers.
 3. The model *is* the mind.
- Weak AI:
 1. Develop theories of human and animal intelligence, and then test these theories by building working models, usually in the form of computer programs or robots.
 2. The working models as a tool to aid in understanding.
 3. Does not claim that the machines themselves are capable of thought, consciousness, and emotions.

See B&S (2003), pp6–7.

8

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.

9

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).

10

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.

11

How Flies Fly



Source: *Nature*

12

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.
- Prediction, Exploiting the external medium
- Tool construction and tool use

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

13

Overview

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

15

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.

14

Foundations of AI

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

See B&S pp. 13–20

16

Philosophy of Mind

The mind-body problem (see B&S pp. 18–19):

- 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.

17

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.

19

Mathematics

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

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

18

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.

20

Cognitive Science

Interdisciplinary field studying human perception and cognition, ranging over:

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

21

Connections

Scientific discoveries came from observing unexpected connections:

- Apple and gravity
- Cloud chamber and the discovery of subatomic particles that led to the atomic bomb.
- Looms with punch-cards and modern computers

23

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

22

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 (in 2016)!

<http://en.wikipedia.org/wiki/AI@50>

24

History of AI (II)

Early successes (1952–1969)

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

25

History of AI (IV)

The 90's: Learn from Data

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

The 2000's – early 2010's: Learn from unstructured data

- Big data, web, information retrieval
- Human-based computation: Recaptcha, Amazon Mechanical Turk

27

History of AI (III)

The 60's and 70's: Hard-coded

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

The 80's: Symbolic Logic-based

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

26

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 sophistication 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

In 2020, it's 64 years old! Vision and language has greatly advanced since 2006. Manual dexterity somewhat improved. Social sophistication still a long way to go.

28

History of AI (VI)

Mid 2010's – present: Age of Deep Learning

- Deep learning
 - Vision : Convolutional Neural Networks
 - Speech : Recurrent Neural Networks
 - Natural language : Recurrent Neural Networks, Attention-based (e.g. Transformer)
 - Deep reinforcement learning: Complex games, Robot control, autonomous driving
- Artificial General Intelligence (AGI)
 - Reasoning, Common sense
 - Neurosymbolic AI

29

Summary and Outlook

31

Major Achievements

- AI in computer games (1980's–current).
- Game playing (Chess, etc.): IBM's Deep Blue (1997)
- Jeopardy: IBM's Watson (2011)
- Deep learning:
 - DeepMind's AlphaGo beats Go master 4-1 (2016)
 - DeepMind's AlphaStar beats top Starcraft gamers 10-1 (2019)
 - Hinton, Le Cunn, Bengio, wins the 2018 ACM Turing Award (2019).
 - Autonomous driving, visual recognition, Language translation, Speech recognition, etc. etc.

30

Symbolic AI: 1950's – mid 80's

- Symbolic Logic, Rule Based
- Theorem proving, Expert systems
- Issue: Can't deal with vague concepts ; Cannot learn from data

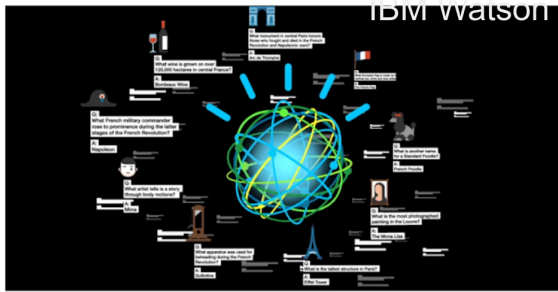


1. Train late (T) & no taxi at station ($\sim X$) John will be late for the meeting (L).	1. $T \wedge \sim X \rightarrow L$	=	$\sim T \vee X \vee L$
2. Train was late (T)	2. T		
3. John wasn't late for the meeting ($\sim L$).	3. $\sim L$		
4. [Question] Was there any taxi at the station?	4. X		

32

Machine Learning: mid 80's – mid 2010's

- Machine Learning), Probability based
- Unstructured Data,
- Some vision/language/speech tasks
- Decline of logic-based AI

Learn from Data



33

Machine Learning Primer

41	34	38	4	38	232	228	217	249	251
17	14	5	42	21	248	231	236	248	246
8	42	6	25	9	217	252	241	238	236
8	14	41	7	221	251	252	242	237	
23	40	27	28	11	235	231	237	244	231
28	15	30	27	28	237	232	232	255	245
1	31	36	18	8	249	237	245	252	237
16	11	40	5	18	231	244	235	246	231
23	9	15	35	19	243	231	236	242	239
17	12	17	16	17	241	233	235	234	237

Learn from data pair in the form of (x =PATTERN, y =LABEL).

Can recognize novel input x not seen during training

Data set 1 : Logic operation

x_1	x_2	y
0	0	0
0	1	1
1	0	1
1	1	1

* Deep learning is a form of machine learning

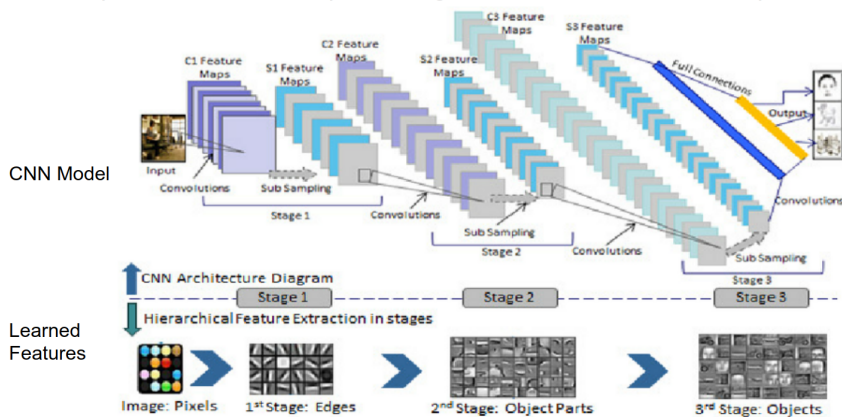
Data set 2 : Animal Recognition

x	y
	Dog
	Cat
	Dog
	Cat

34

Deep Learning: mid 2010's – present


- Rapid advances in Deep Learning: CNN, RNN, Attention, Deep RL



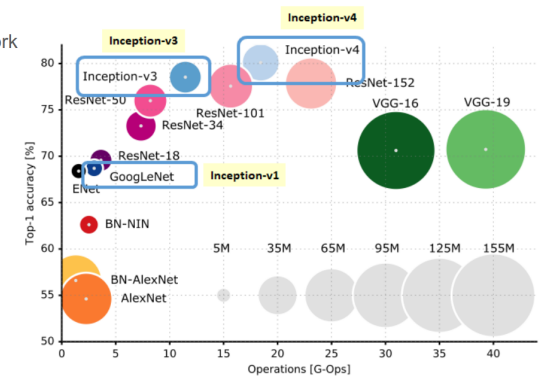
35

Advances in Deep Learning

- Fuel for DL Revolution
 - Algorithm: Very deep neural network
 - Big Data
 - Computing power (GPU-based)





CNN Model (Google's Inception)



36

Advances in Deep Learning

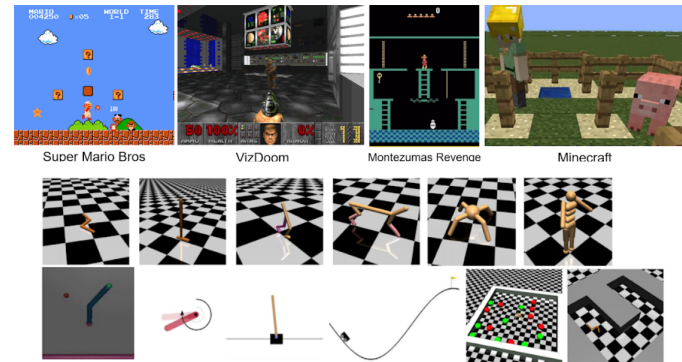
- Major advances in Vision, Speech, NLP - Superhuman in some cases
 - Beyond Classification: Detection, Segmentation (Semantic / Instance); Enhanced speech
 - Vision+Language (Visual Question Answering): Question answering based on images.



37

Advances in Deep Learning


- Beyond perception/language: Control of complex behavior based on perception.
 - Deep Reinforcement Learning (AlphaGo, AlphaStar) : Behave based on perception and learn from Reward.




38

Limitations of Deep Learning


- Limitations of deep learning
 - Need large data, Cannot do complex reasoning
 - Lack of common sense
 - Lack of interpretability, Susceptible to noise
 - Learn spurious features



Sensitive to small perturbations



What to rescue first in case of fire?






Learned snow field feature, not husky feature

39

Future of AI

- Combining symbolic AI and deep learning
 - Neurosymbolic AI
- New methods inspired by neuroscience
 - Neuroscience-inspired AI, Neuromorphic AI
- Nnets + evolution
 - Neuroevolution
- Philosophical issues
 - Consciousness
 - Epistemology
 - Post-humanism



40

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.

41

Next Time

- LISP

43

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:~/>
```

42