CS440/ECE448: Artificial Intelligence
Lecture 1: Course Intro

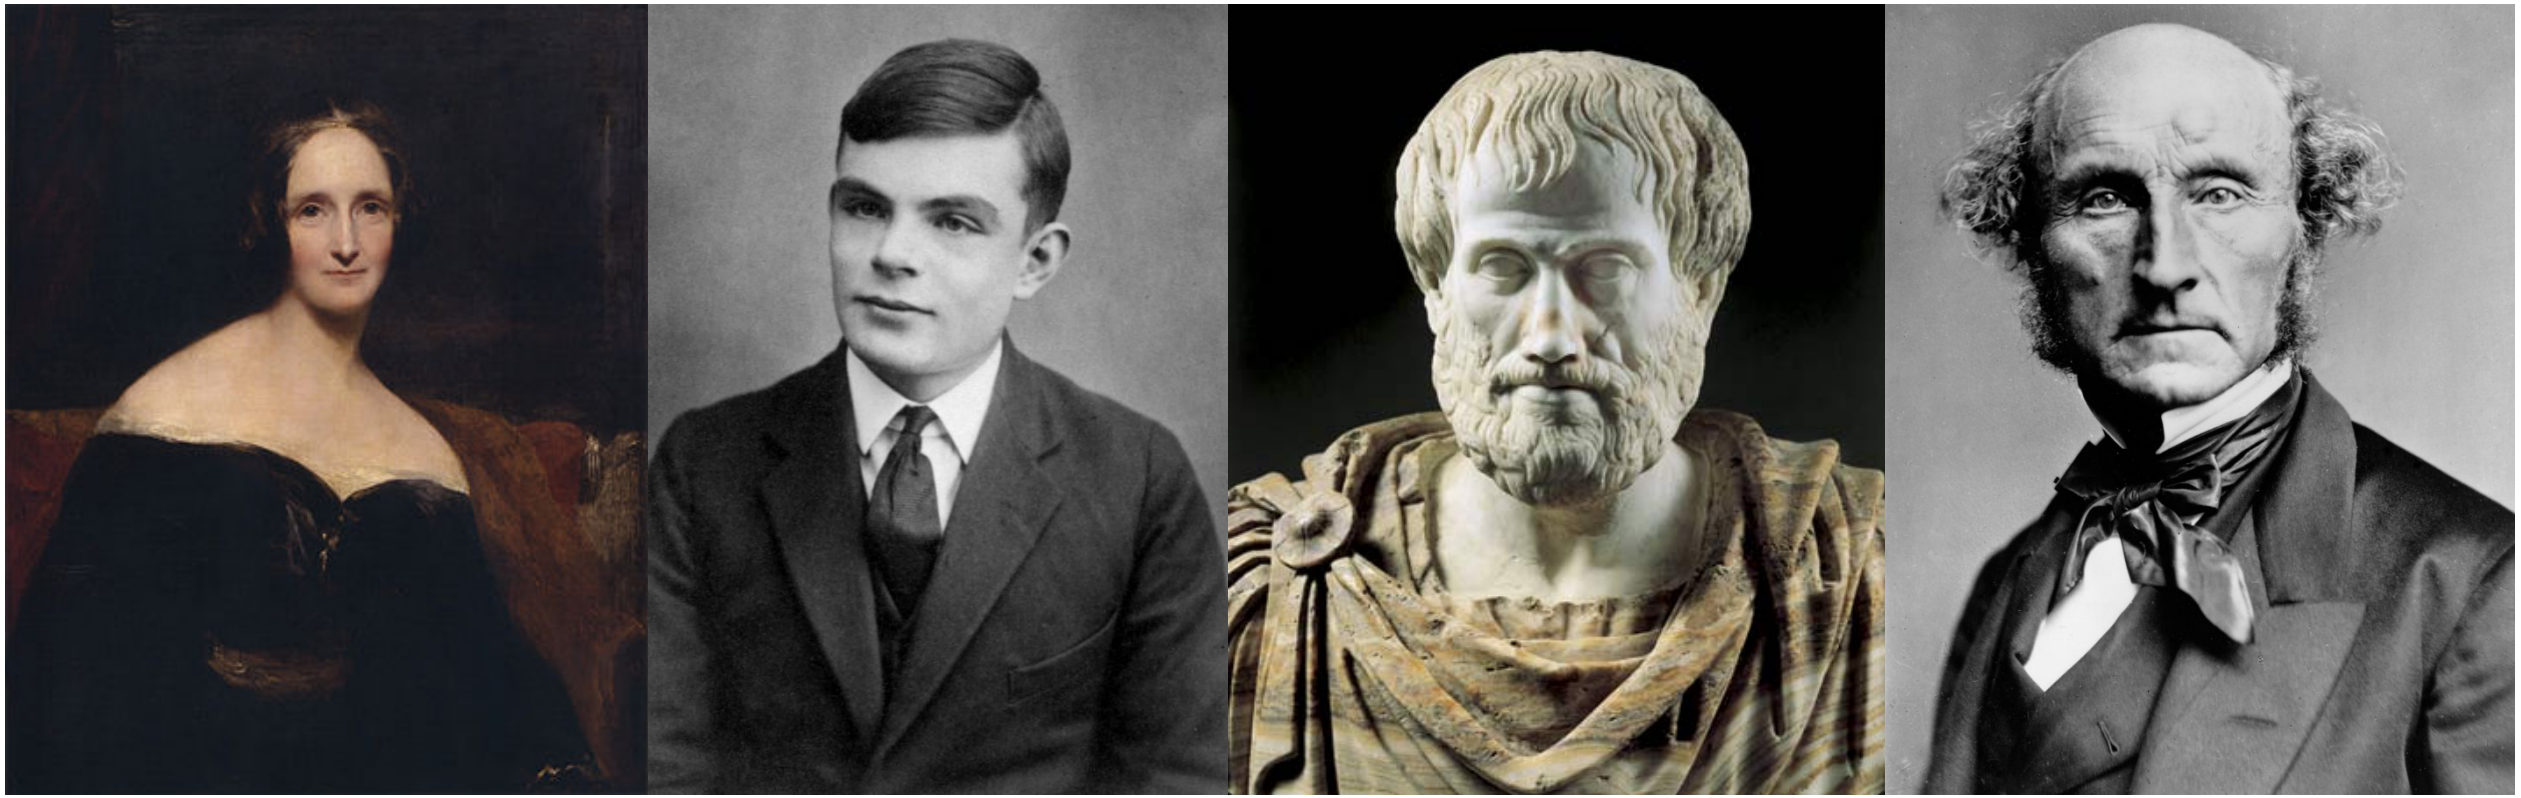

Course Intro: Syllabus

- Web page: https://courses.engr.illinois.edu/ece448/sp2020/
- Grading
- Homework
- Apps
- Textbook

Grading

- 3-credit and 4-credit sections will be graded on separate curves.
- 3-credit: 60% homework, 40% exams.
- 4-credit: 50% homework, 40% exams, 10% review papers.
- Grade cutoffs: 90% will get you at least an A-, 80% will get you at least a B-, 70% will get you at least a C-. Curves are likely to reduce the B- and C- thresholds.

Homework

- 7 machine problems, mostly autograded, at gradescope.com.
- MP1 (search): released Monday 1/22, due Monday 2/5 by 11:59pm.
- Most MPs are released two weeks before they are due.
- Plan in advance! Deadline extensions are not given for routine causes like being sick, having an on-site job interview, etc.

Apps

- All homework and exams graded and submitted at: https://www.gradescope.com/
- We will copy your grades to either https://learn.illinois.edu or https://compass2g.illinois.edu (not sure which, yet) so you can see where you stand relative to class average and standard deviation.
- Q/A forum: https://piazza.com
- Videos of lectures: https://echo360.org
- Searchable lecture videos? https://classtranscribe.ncsa.illinois.edu

Textbook

Artificial Intelligence, A Modern Approach: Third Edition by Russell & Norvig

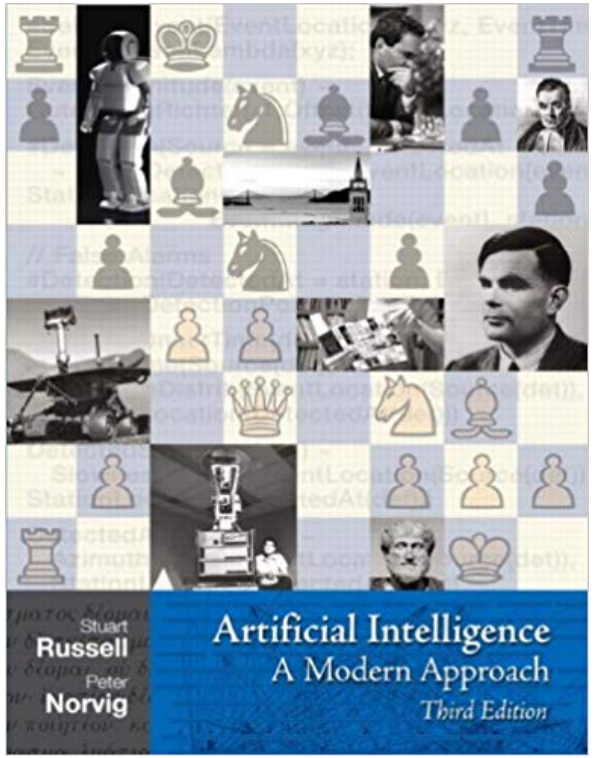

Hardcover

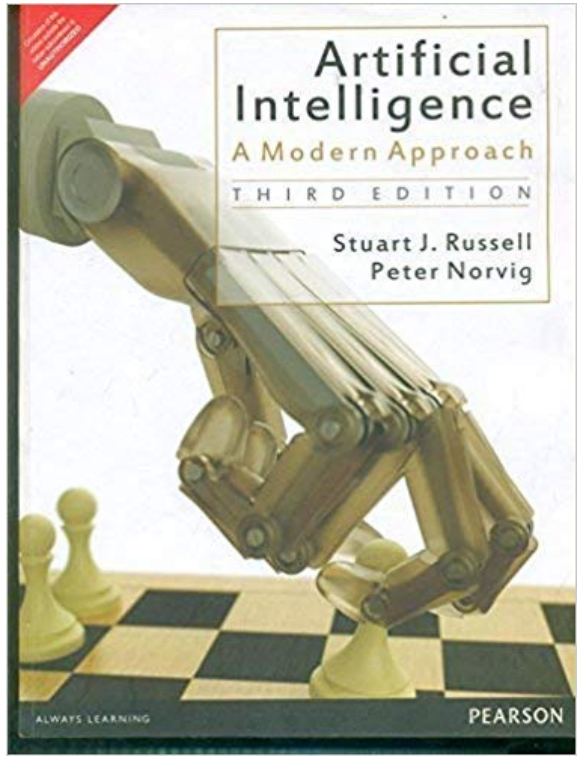

Paperback

- Pretty good listing of the topics covered in this course
- In-depth treatment of knowledge-based/expert-system AI; introduces probabilistic and learning-based methods
- Sample problems and readings will be specified when applicable

Outline, Remainder of Today's Lecture

- What is AI? What is Intelligence?
- Russell & Norvig's "four approaches to AI" (chapter 1)
- Brief history of AI (chapter 2)

What is Intelligence?

The word “intelligence” is surprisingly recent. Ancients used it to mean “the universal mind.” Early moderns (e.g., Bacon, Hobbes; 1500s) ridiculed it, and stopped using it. It was then repurposed to its current meaning by psychologists and eugenicists in the early 20th century.

What is Intelligence?

Charles Spearman popularized the modern definition in his paper "General intelligence objectively determined and measured," American Journal of Psychology 15(2):201-292.

- He showed that test scores are correlated across many subjects and proposed "general intelligence" as the faculty that unifies them.

Spearman's correlation matrix for six measures of school performance. All the correlations are positive, the *positive manifold* phenomenon. The bottom row shows the *g* loadings of each performance measure.[7]

	Classics	French	English	Math	Pitch	Music
Classics	–					
French	.83	–				
English	.78	.67	–			
Math	.70	.67	.64	–		
Pitch discrimination	.66	.65	.54	.45	–	
Music	.63	.57	.51	.51	.40	–
g	.958	.882	.803	.750	.673	.646

https://en.wikipedia.org/wiki/G_factor_(psychometrics)

What is “Artificial Intelligence”?

The term was invented in (John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” August 1955):

“We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer … An attempt will be made to find how to make machines

1. use language,
2. form abstractions and concepts,
3. solve kinds of problems now reserved for humans
4. improve themselves.”

What is Artificial Intelligence?

Russell & Norvig's "four approaches to AI" (chapter 1): Intelligence means…

1. Thinking like a Human	2. Acting like a Human
3. Thinking Rationally	4. Acting Rationally

1. Thinking like a Human

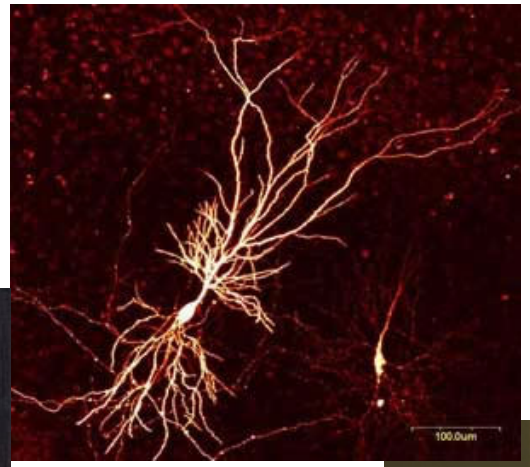

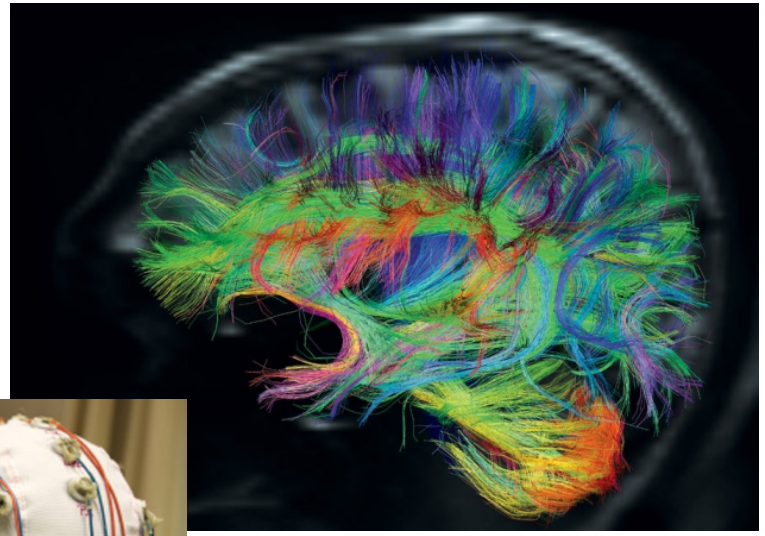

Mary Shelley, author of *Frankenstein: The Modern Prometheus*; Neuron, showing branching of the dendrites; EEG cap; Cortical connectivity map, computed using diffusion tensor MRI

Can we simulate a human brain?

How many binary computations per second can the brain perform?

- Spatial scale: there are 100 trillion neurons (10^14).
- Numerical precision: each neuron either generates an action potential or doesn't (binary!).
- Temporal scale: Other neurons are sensitive to timing with a resolution of perhaps roughly 1 millisecond (1000 bits/second).

Answer: if each neuron performs 1000 binary computations/second, then the brain performs up to (100 trillion)X(1000) = 10^17 binary computations/second (100 Peta-ops: about 100,000 GPUs)

Then why can't we simulate a human brain?

How many brain computations can we IMAGE?

- Temporal scale: no problem, EEG (electroencephalography) gets ~5000 samples/second
- Spatial scale is the problem:
  - EEG: 100 pixels/brain
  - fMRI and ECOG: 1mm scale (~10^5 voxels/brain)
  - Versus 10^14 neurons/brain

Then why can’t we simulate a human brain?

- The short answer: we can’t find out what computations a living human brain is performing, because any current imaging modality that precise would kill it.
- …and we are currently about 9 orders of magnitude (10^9) away from the necessary level of precision (volume).
- MRI improved by roughly 2 orders of magnitude per decade from 1970 to 2000, then slowed significantly, has improved perhaps 1 o.o.m. per two decades since then. So perhaps this approach will be possible in 180 years.

What is Artificial Intelligence?

Russell & Norvig's "four approaches to AI" (chapter 1): Intelligence means…

1. Thinking like a Human	2. Acting like a Human
3. Thinking Rationally	4. Acting Rationally

What is Artificial Intelligence?

Russell & Norvig's "four approaches to AI" (chapter 1):

1. Thinking humanly	**2. Acting humanly**
3. Thinking rationally	**4. Acting rationally**

2. Acting like a Human

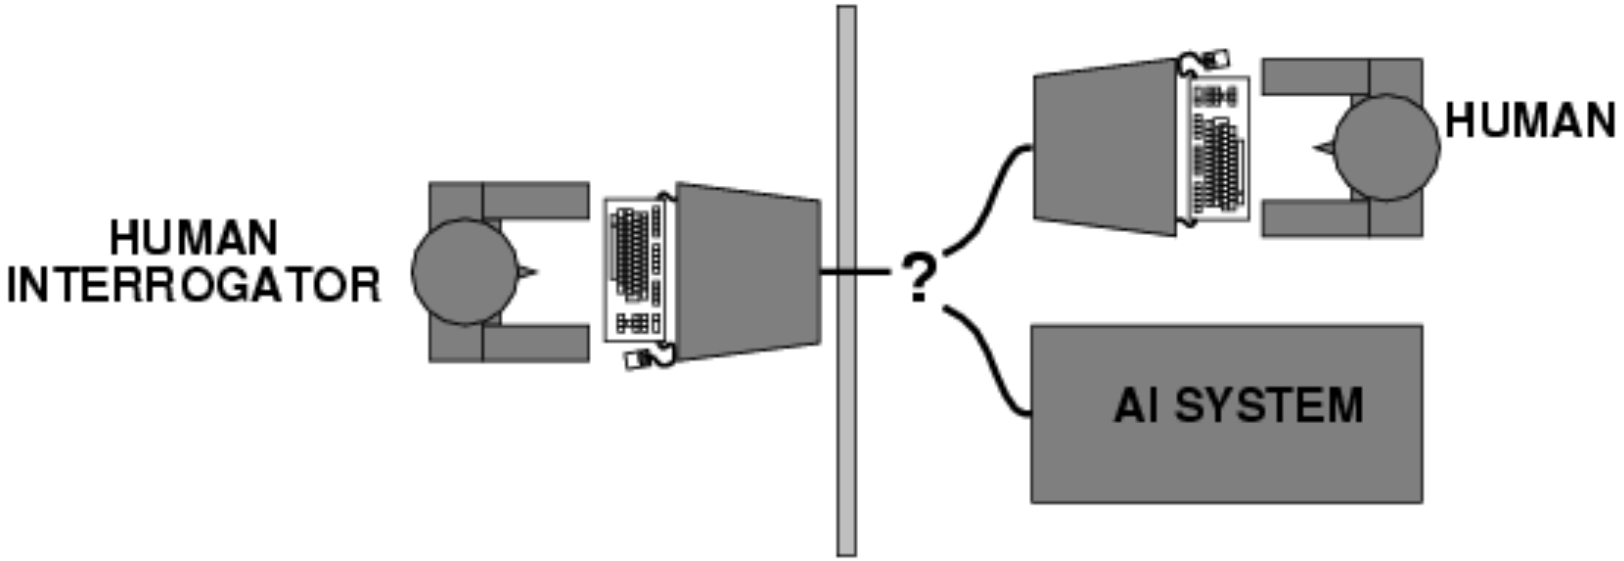

Schematic of the Turing test; Alan Turing

The Turing Test

- Alan Turing, “Intelligent Machinery,” 1947:

“Now get three men as subjects for the experiment. A, B and C. A and C are to be rather poor chess players, B is the operator who works the paper machine… a game is played between C and either A or the paper machine.

C may find it quite difficult to tell which he is playing...

These questions replace our original, ‘Can machines think?’”

Practical Problems with the Turing Test

- Can’t be automated (you need human judges).
- Human judges can be fooled by misdirection, e.g., by a chatbot that pretends to be a paranoid schizophrenic (https://en.wikipedia.org/wiki/PARRY) or a 13-year-old Ukrainian boy (https://en.wikipedia.org/wiki/Eugene_Goostman)

Winograd Schema

- Winograd schema (H. Levesque, *On our best behaviour*, IJCAI 2013) attempts to solve the practical problems with the Turing test
- Multiple choice questions that can be easily answered by people but cannot be answered by computers using “cheap tricks”
- Always arranged in pairs:

The trophy would not fit in the brown suitcase because it was so small. What was so small?

  - *The trophy*
  - *The brown suitcase*

Winograd Schema

- Winograd schema (H. Levesque, *On our best behaviour*, IJCAI 2013) attempts to solve the practical problems with the Turing test
- Multiple choice questions that can be easily answered by people but cannot be answered by computers using "cheap tricks"
- Always arranged in pairs:

*The trophy would not fit in the brown suitcase because it was so **large**.*
*What was so **large**?*

  - *The trophy*
  - *The brown suitcase*

A theoretical problem with the Turing test

Why is human behavior the standard?

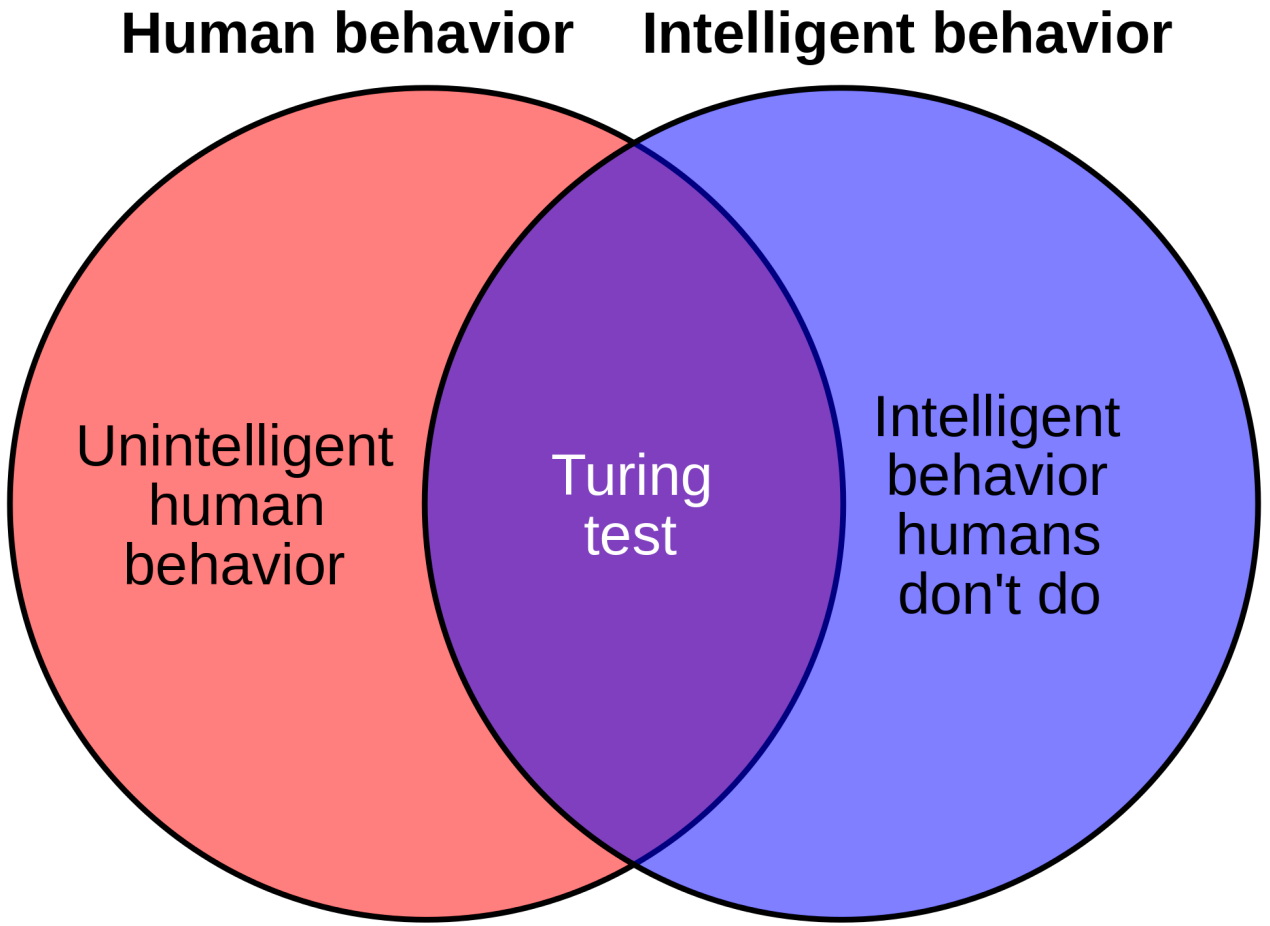

By en:User:CharlesGillingham, User:Stannered - en:Image:Weakness of Turing test 1.jpg, Public Domain, https://commons.wikimedia.org/w/index.php?curid=3457053

What is Artificial Intelligence?

Russell & Norvig's "four approaches to AI" (chapter 1): Intelligence means…

1. Thinking like a Human	2. Acting like a Human
3. Thinking Rationally	4. Acting Rationally

AI definition 3: Thinking rationally

Aristotle, 384-322 BC

AI definition 3: Thinking rationally

- Idealized or “right” way of thinking
- **Logic:** patterns of argument that always yield correct conclusions when supplied with correct premises
  - *“Socrates is a man; all men are mortal; therefore Socrates is mortal.”*
- **Logicist approach to AI:** describe problem in formal logical notation and apply general deduction procedures to solve it

Successes of Logicist Approach: Expert Systems

- Expert system = (knowledge base) + (logical rules)
  - Knowledge base = easy to collect from human judges and/or encyclopedia
  - Logical rules = easy to deduce from examples, and easy to verify by asking human judges
  - Combination of the two: able to analyze never-before-seen examples of complicated problems, and generate the correct answer

- Example: speed control system of the https://en.wikipedia.org/wiki/Sendai_Subway_Namboku_Line. “This system (developed by Hitachi) accounts for the relative smoothness of the starts and stops when compared to other trains, and is 10% more energy efficient than human-controlled acceleration.”

Failures of Logicist Approach: Robust AI

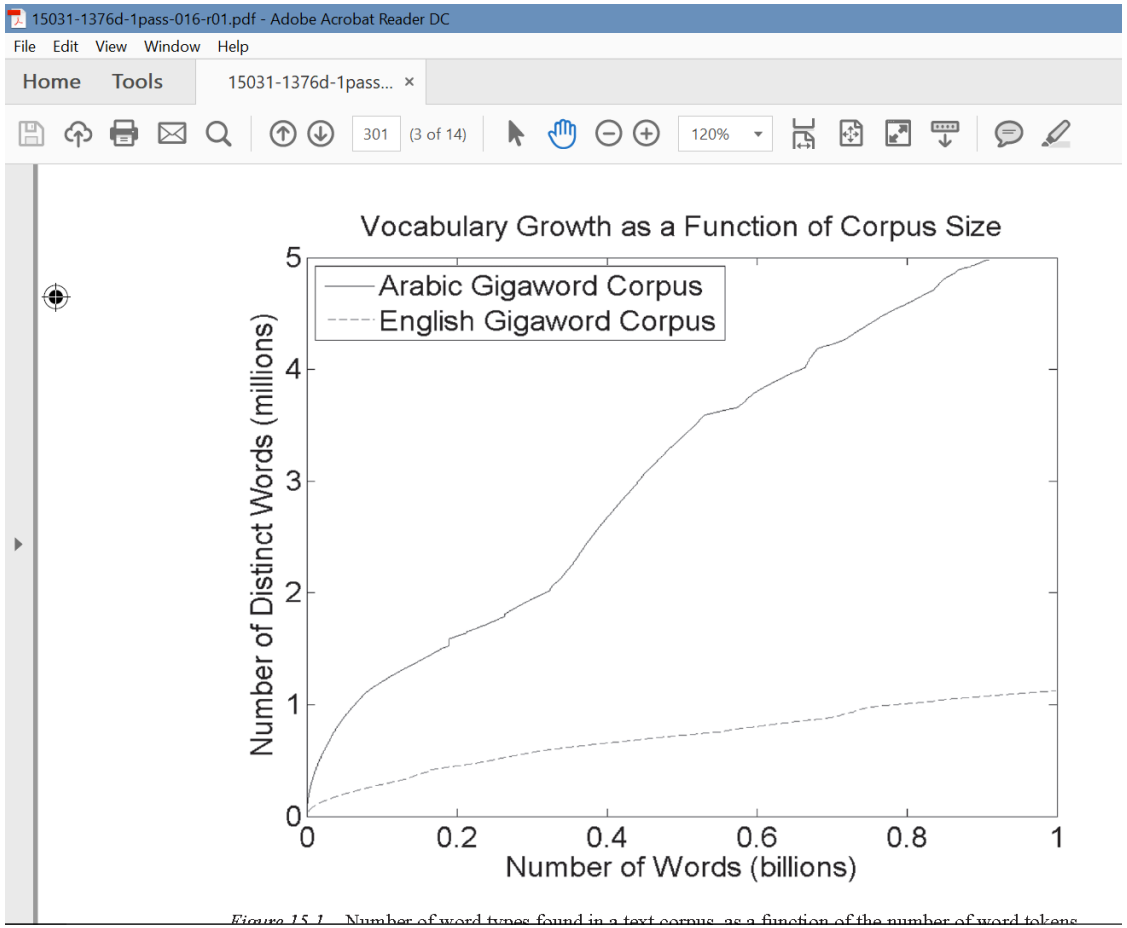

- Humans commonly believe that there are a finite number of facts that must be entered into a knowledge base. Evidence suggests that this is incorrect.
- Example (Hasegawa-Johnson, Elmahdy & Mustafawi, "Arabic Speech and Language Technology," 2017): the number of distinct words in any corpus of text is linearly proportional to the number of words. In English, a never-before-seen word occurs ~once/1000 words; in Arabic, ~once/180 words.

What is Artificial Intelligence?

Russell & Norvig's "four approaches to AI" (chapter 1): Intelligence means…

1. Thinking like a Human	2. Acting like a Human
3. Thinking Rationally	**4. Acting Rationally**

AI definition 4: Acting rationally

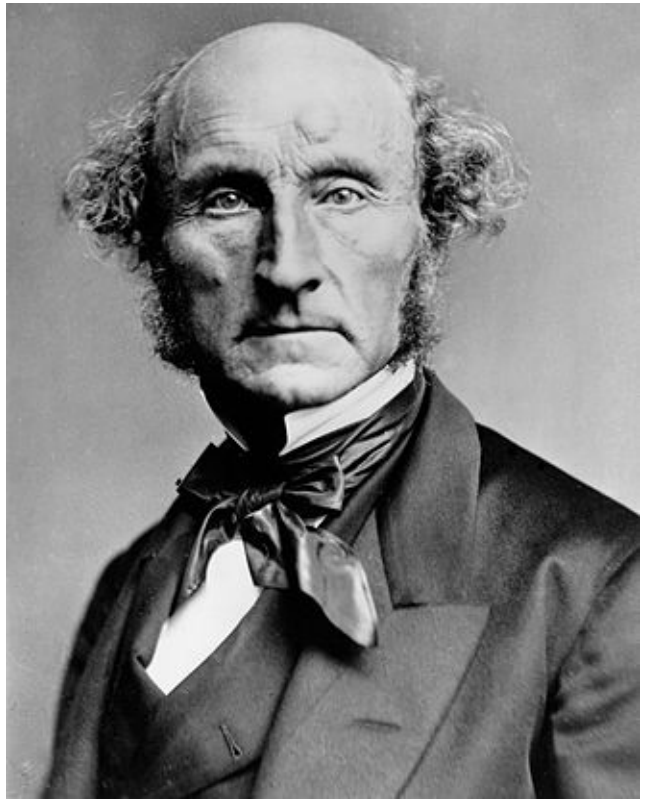

John Stuart Mill, 1806-1873

AI definition 4: Acting rationally

- **Rational agent =** acts to optimally achieve its goals
  - Goals are application-dependent and are expressed in terms of the **utility of outcomes**
  - Being rational means **maximizing your (expected) utility**
- This definition of rationality only concerns the decisions/actions that are made, not the cognitive process behind them
- An unexpected step: rational agent theory was originally developed in the field of economics
  - Norvig and Russell: “most people think Economists study money. Economists think that what they study is the behavior of rational actors seeking to maximize their own happiness.”

Russell & Norvig's "four approaches to AI"

Intelligence is…

1. Thinking like a human
   - Sometimes called "grounded AI" – create an AI with neurons like ours
2. Acting like a human
   - Turing's definition of AI; ignores the underlying thought process
   - Might include acting irrationally
3. Thinking rationally
   - Logicist AI: must use a rational/logical thought process
4. Acting rationally
   - Utilitarianism: act in order to maximize your own benefit, regardless of the thought process you use

Outline, Remainder of Today's Lecture

- What is AI? What is Intelligence?
- Russell & Norvig's "four approaches to AI" (chapter 1)
- Brief history of AI (chapter 2)

A very brief history of AI

Decade	Symbolic AI	Neural AI
1943		McCulloch-Pitts
1956	Logic Theorist	
1966	ALPAC Report	
1986		Back-propagation
1996	Deep Blue	
2016		AlphaGo

1956: Logic Theorist

- **<u>Reasoning as search</u>**:
  - Root of the search tree: initial hypothesis
  - Branch: a deduction based on the rules of logic
  - Goal state: the theorem to be proven
- Proved 38 of the first 52 theorems in chapter 2 of the *Principia Mathematica*.

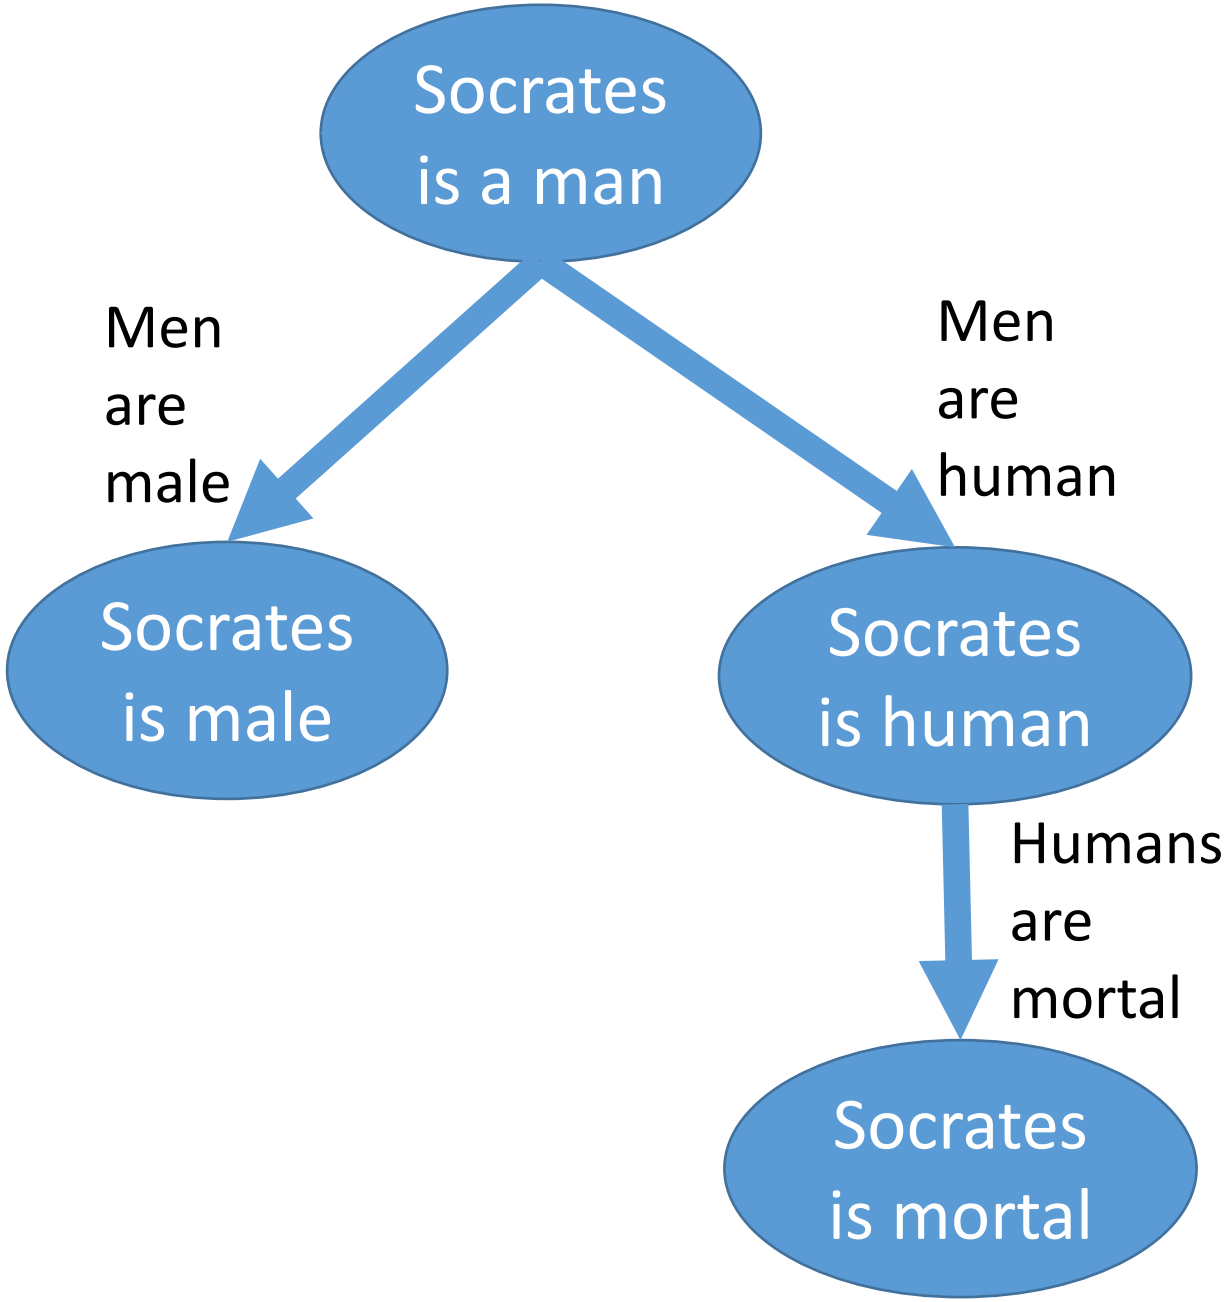

The ALPAC Report of 1966

“They concluded, in a famous 1966 report, that machine translation was more expensive, less accurate and slower than human translation.”

Photo: Eldon Lyttle, https://commons.wikimedia.org/wiki/File:Computer-translation_Briefing_for_Gerald_Ford.jpg

AI Winter

1943: McCulloch-Pitts Neuron Model

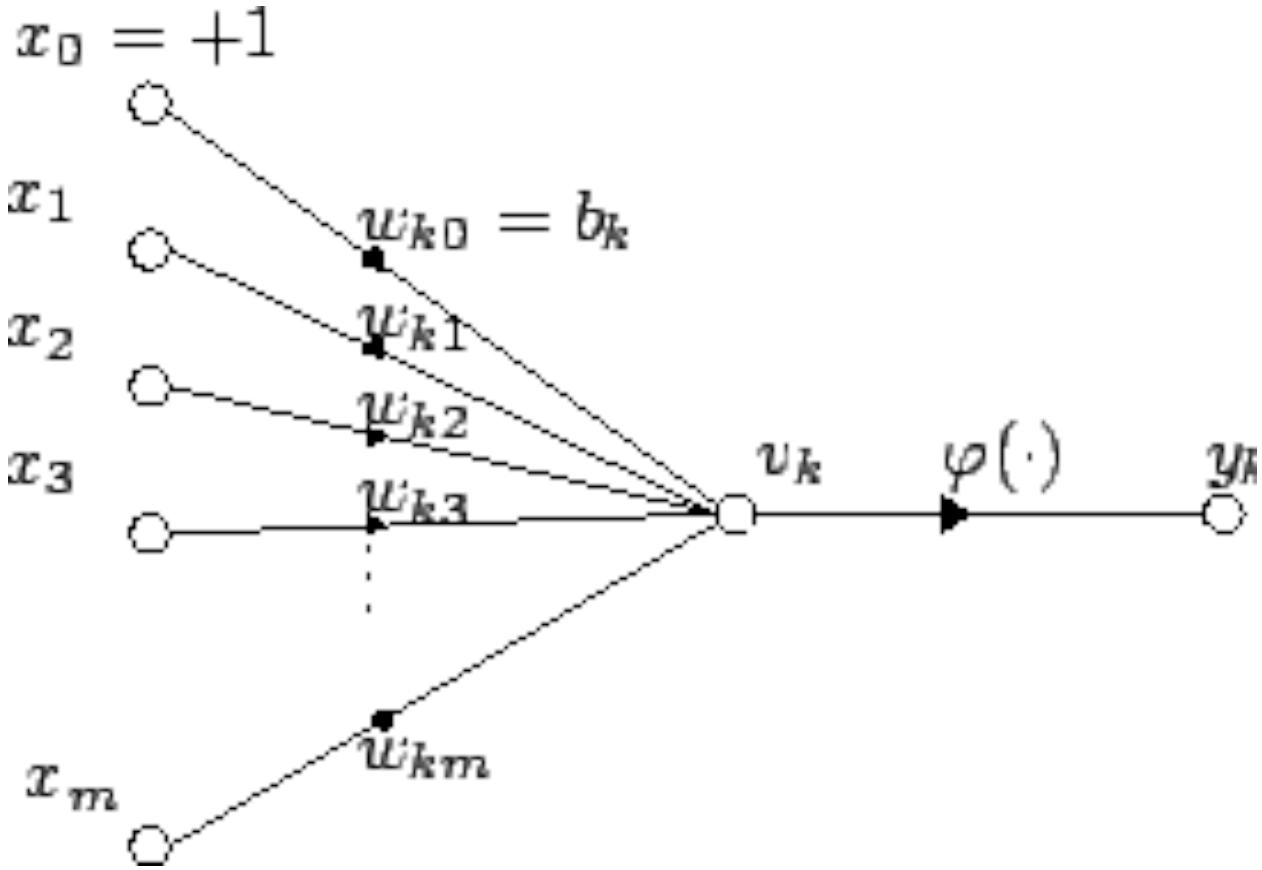

Attribution: Plarroy, https://commons.wikimedia.org/wiki/File:Artificial_neuron.png

1986: Invention of back-propagation makes multi-layer neural networks practical

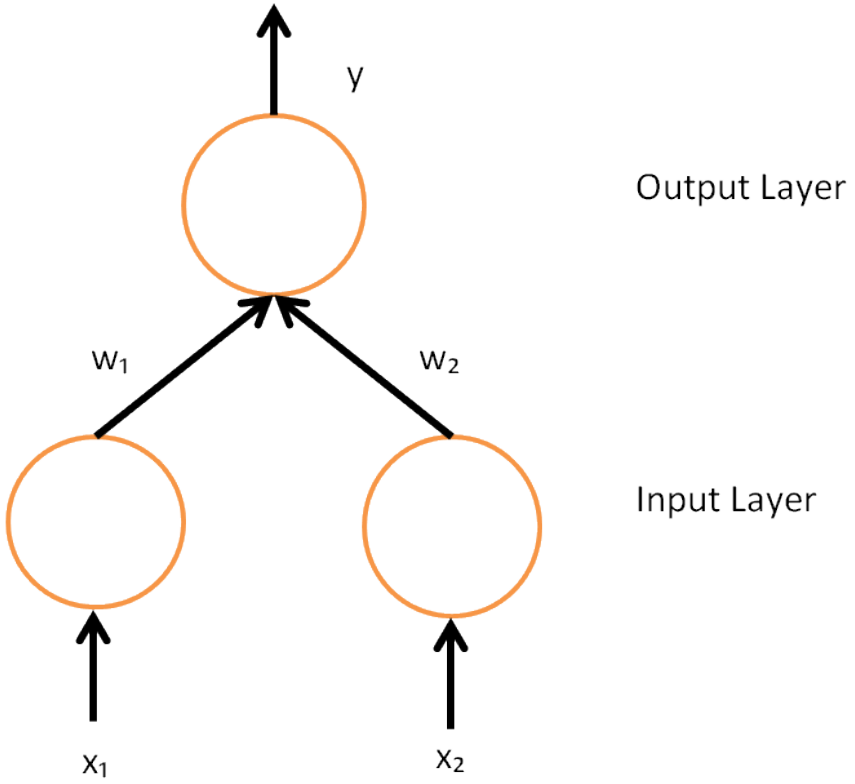

David E Rumelhart, Geoffrey E Hinton, Ronald J Williams, "Learning internal representations by error-propagation," in Parallel Distributed Processing: Explorations in the Microstructure of Cognition 1:318-362

1996: Deep Blue beats Gary Kasparov

By Kasparov_Magath_1985_Hamburg.jpg: GFHundderivative work: Hardy Linke (talk) - Kasparov_Magath_1985_Hamburg.jpg, CC BY 3.0, https://commons.wikimedia.org/w/index.php?curid=14640503

2016: AlphaGo

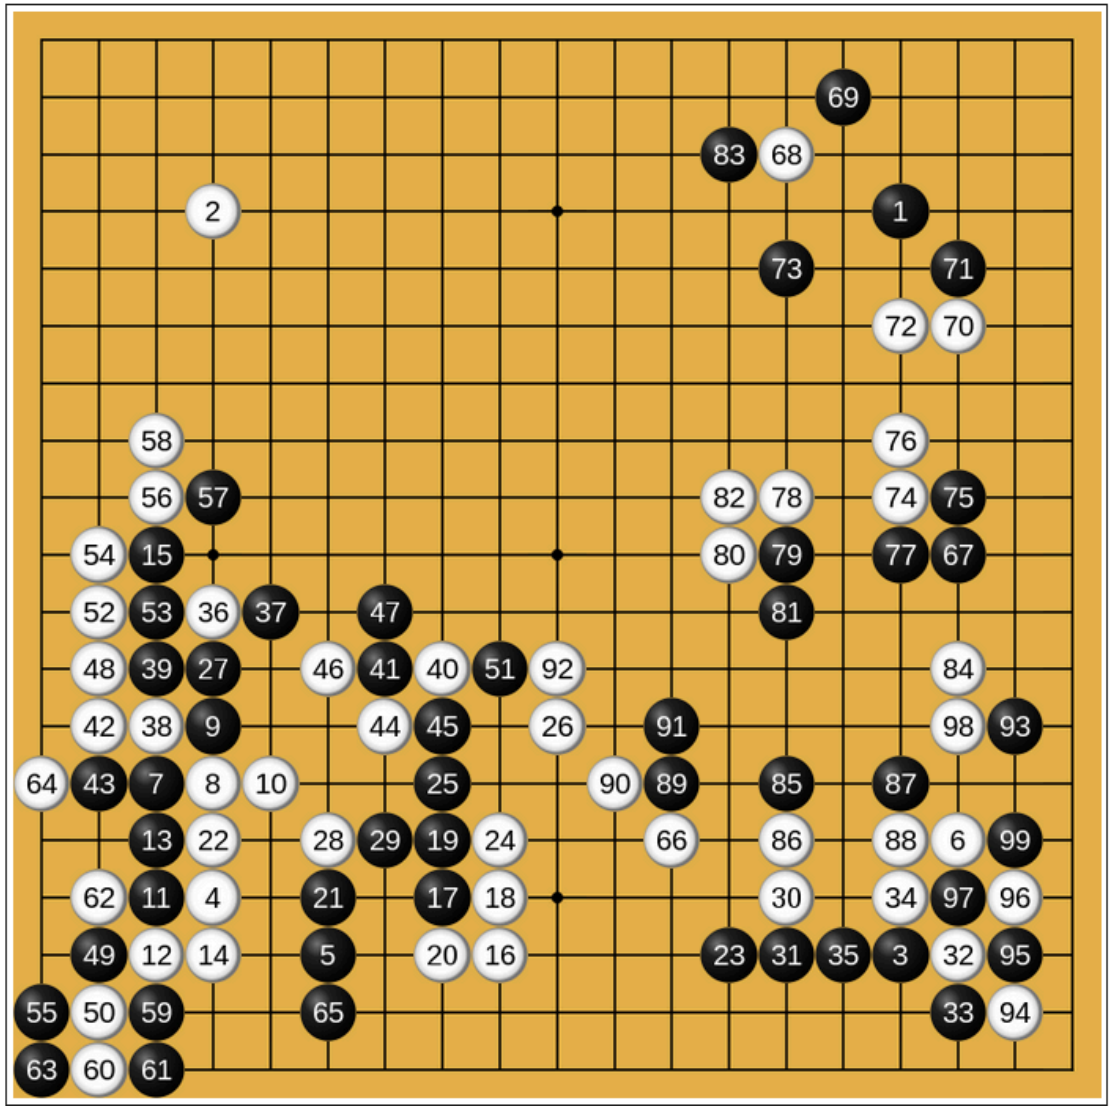

A very brief history of AI

Decade	Symbolic AI	Neural AI
1943		McCulloch-Pitts
1956	Logic Theorist	
1966	ALPAC Report	
1986		Back-propagation
1996	Deep Blue	
2016		AlphaGo

A modern research problem: how to best combine these two approaches in order to make really “thinking” machines.

Upcoming: course outline

1. Thinking rationally: search and planning (weeks 1-4)
2. Acting rationally: probability (weeks 5-8)
3. Thinking like a human: neural nets (weeks 9-11)
4. Acting like a human: reinforcement learning, games, and vector semantics (weeks 12-14)