
Artificial Intelligence

Lecture 1 - Introduction

Brigitte Jaumard

Dept of Computer Science and Software Engineering

Concordia University

Montreal (Quebec) Canada

Outline

- What is AI?
- Foundations of AI
- Short history of AI
- Philosophical discussions

What is AI?

Views of AI fall into four categories:

Think like humans	Thinking rationally
Act like humans	Acting rationally

The textbook advocates "acting rationally"

Think like humans


- 1960s "cognitive revolution"
 - information-processing psychology: get inside the actual workings of human minds
 - **Introspection:** catch our own thoughts as they go by
 - **Psychological experiments:** observing a person in action
 - **Brain imaging:** observing the brain in action
- Scientific theories of internal activities of the brain
 - What level of abstraction? "Knowledge" or "circuits"?
 - **Cognitive science:** Predicting and testing behavior of human subjects (top-down)
 - A. Newell and H. Simon (GPS – General Problem Solver): concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problem
 - This approach now distinct from AI
 - Share with AI the following characteristic:
 - ~~The available theories do not explain anything resembling human-level general intelligence~~

Act like humans

- Turing (1950) "Computing machinery and intelligence"
 - **Natural language processing:** enable to communicate in, e.g., English or Vietnamese
 - **Knowledge representation:** store what it knows or hears
 - **Automated Reasoning:** store information to answer questions
 - **Machine Learning:** adapt to new circumstances, detect and extrapolate patterns
 - **Computer Vision:** perceive objects
 - **Robotics:** manipulate objects and move about
- "Can machines think?" → "Can machines behave intelligently?"

Act like humans (cont'd)

- Operational test for intelligent behavior: **the Imitation Game (Turing test)**
 - A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer
- Predicted that by 2,000, 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



Thinking rationally

- The “Laws of Thought” approach
 - What does it mean to “think rationally”?
 - Normative / prescriptive rather than descriptive
- Logician tradition
 - Logic: notation and rules of derivation for thoughts
 - Aristotle: what are correct arguments/thought processes?
 - E.g.: Socrat is a human, human cannot live forever → Socrat human cannot live forever
 - Direct line through mathematics, philosophy, to modern AI
- Problems:
 - Not all intelligent behavior is mediated by logical deliberation
 - What is the purpose of thinking? What thoughts should I have?
 - Logical systems tend to do the wrong thing in the presence of uncertainty

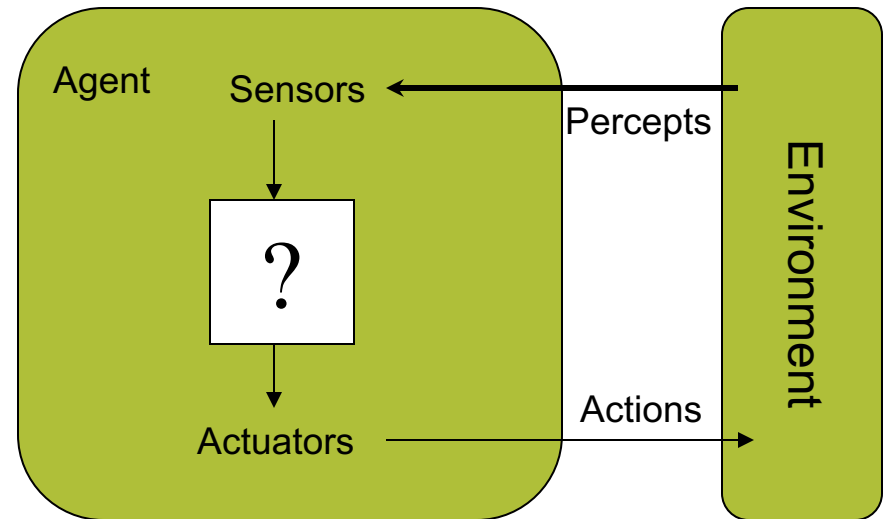
Acting rationally

- 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
 - Thinking can be in the service of rational action
 - Entirely dependent on goals!
 - Irrational \neq insane, irrationality is sub-optimal action
 - Rational \neq successful
- Our focus here: rational agents
 - Systems which make the best possible decisions given goals, evidences, and constraints
 - In the real world, usually lots of uncertainty... and lots of complexity
 - Usually, we're just approximating rationality
- “Computational rationality” a better title for this course

Rational agents

- An **agent** is an entity that perceives and acts
- An agent function maps from percept histories to actions:

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



- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Computational limitations make perfect rationality unachievable
- So we want the best program for given machine resources

Foundations of AI

- 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
- Economics utility, decision theory
- Neuroscience physical substrate for mental activity
- Psychology phenomena of perception and motor control, experimental techniques
- Computer engineering building fast computers
- Control theory design systems that maximize an objective function over time
- Linguistics knowledge representation, grammar

Short history of AI

- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's "Computing Machinery and Intelligence"
- 1950—70: Excitement: 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
 - 1964: ELIZA
 - 1965: Robinson's complete algorithm for logical reasoning
- 1970—88: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms
- 1988—93: Expert systems industry busts: "AI Winter"
- 1988—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - Agents, agents, everywhere... "AI Spring"?
- 2000—: Where are we now?

ELIZA

- Early natural language processing computer program created from 1964 to 1966 at the MIT Artificial Intelligence Laboratory by Joseph Weizenbaum
 - Created to demonstrate the superficiality of communication between humans and machines
 - Eliza simulated conversation by using a 'pattern matching' and substitution methodology that gave users an illusion of understanding on the part of the program, but had no built in framework for contextualizing events
-

ELIZA (cont'd)

Welcome to

```
EEEEEE LL      IIII ZZZZZZZ AAAAAA
EE      LL      II     ZZ     AA  AA
EEEEEE LL      II     ZZZ     AAAAAAA
EE      LL      II     ZZ     AA  AA
EEEEEE LLLLLL IIII ZZZZZZZ AA  AA
```

Eliza is a mock Rogerian psychotherapist.

The original program was described by Joseph Weizenbaum in 1966.

This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?

YOU: Men are all alike.

ELIZA: What is the connection, do you suppose ?

YOU: They're always bugging us about something or other.

ELIZA: Can you think of a specific example ?

YOU: Well, my boyfriend made me come here.

ELIZA: Is it important to you that your boyfriend made you come here ?

YOU: He says I'm depressed much of the time.

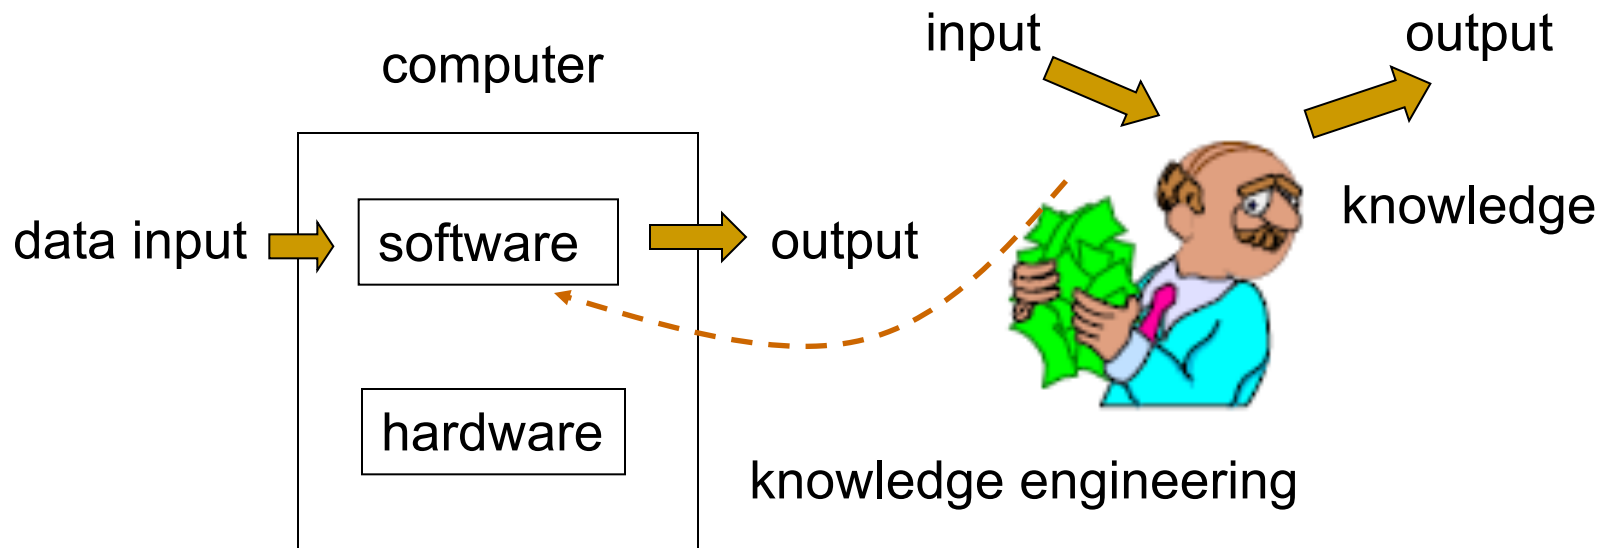
ELIZA: I am sorry to hear that you are depressed.

YOU: It's true. I am unhappy.

ELIZA: Can you explain what made you unhappy ?

YOU: █

Expert systems




Expert system = Human Expertise + Inference/Reasoning

Some examples: DENDRAL, MYCIN, PROSPECTOR,
MOLGEN, ICAD/ICAM

State of the art

- May, '97: Deep Blue vs. Kasparov
 - First match won against world-champion
 - “Intelligent creative” play
 - 200 million board positions per second!
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a big PC cluster
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- Proverb solves crossword puzzles better than most humans
- **AlphaGo** became the first computer Go program to beat a human professional Go player without handicaps on a full-sized 19×19 board



Philosophical discussions

What Can AI Do?

- Play a decent game of table tennis?
- Drive safely along a curving mountain road?
- Buy a week's worth of groceries on the web?
- Discover and prove a new mathematical theorem?
- Converse successfully with another person for an hour?
- Perform a complex surgical operation?
- Unload a dishwasher and put everything away?
- Translate spoken English into spoken Vietnamese in real time?
- Write an intentionally funny story?

Can machine think?

Ethical Questions

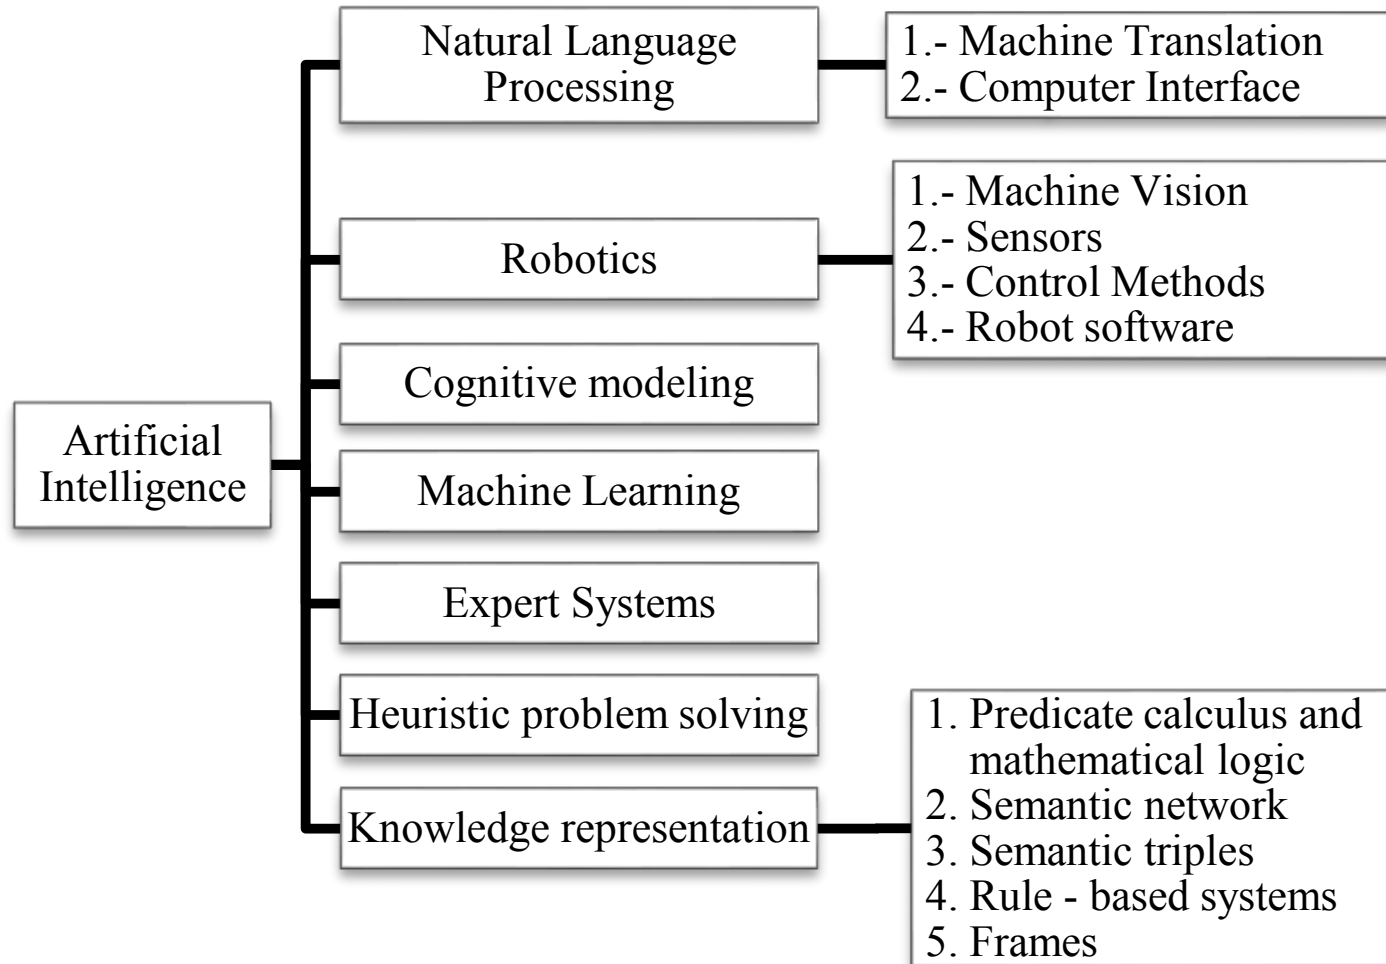
<https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial-intelligence/>

1. **Unemployment.** What happens after the end of jobs?
2. **Inequality.** How do we distribute the wealth created by machines?
3. **Humanity.** How do machines affect our behaviour and interaction?
4. **Artificial stupidity.** How can we guard against mistakes?
5. **Racist robots.** How do we eliminate AI bias?
6. **Security.** How do we keep AI safe from adversaries?
 - not only robots produced to replace human soldiers, or autonomous weapons, but AI systems (e.g., drones) that can cause damage if used maliciously.
7. **Evil genies.** How do we protect against unintended consequences?
8. **Singularity.** How do we stay in control of a complex intelligent system?
9. **Robot rights.** How do we define the humane treatment of AI?

Some problems with AI

- People might lose their jobs to automation.
- People might have too much (or too little) leisure time.
- People might lose their sense of being unique.
- People might lose some of their privacy rights.
- The use of AI systems might result in a loss of accountability.
- The success of AI might mean the end of the human race.


AI covers seven subfields of computer science



Benefits of the use of Artificial Intelligence (AI) to the businesses

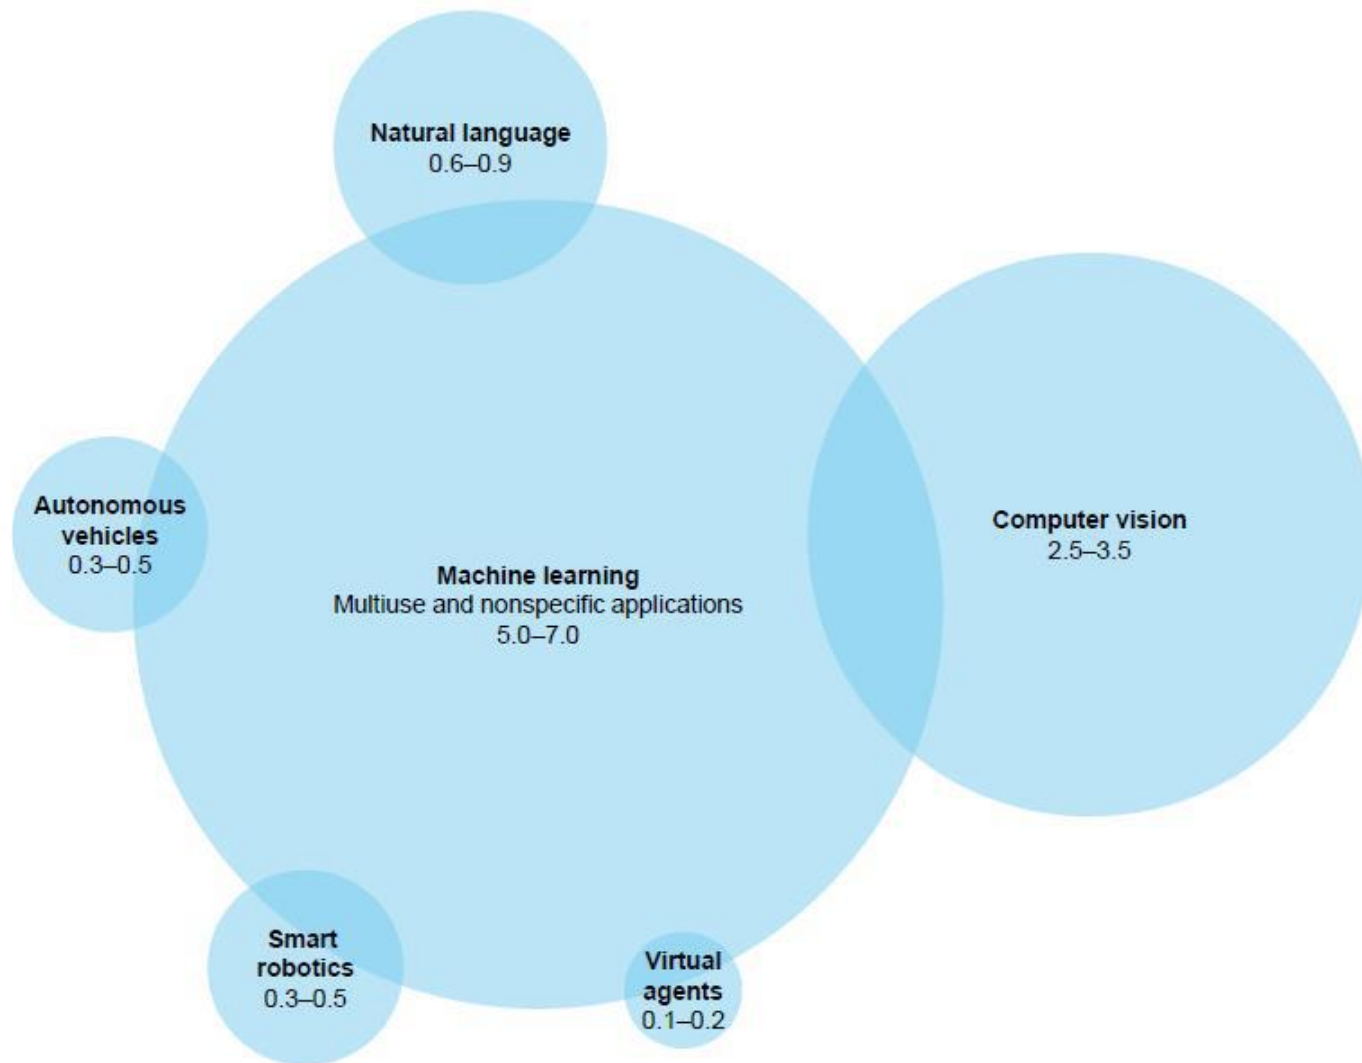
Enterprise artificial intelligence market revenue worldwide 2016-2025

Revenues from the artificial intelligence for enterprise applications market worldwide, from 2016 to 2025 (in million U.S. dollars)



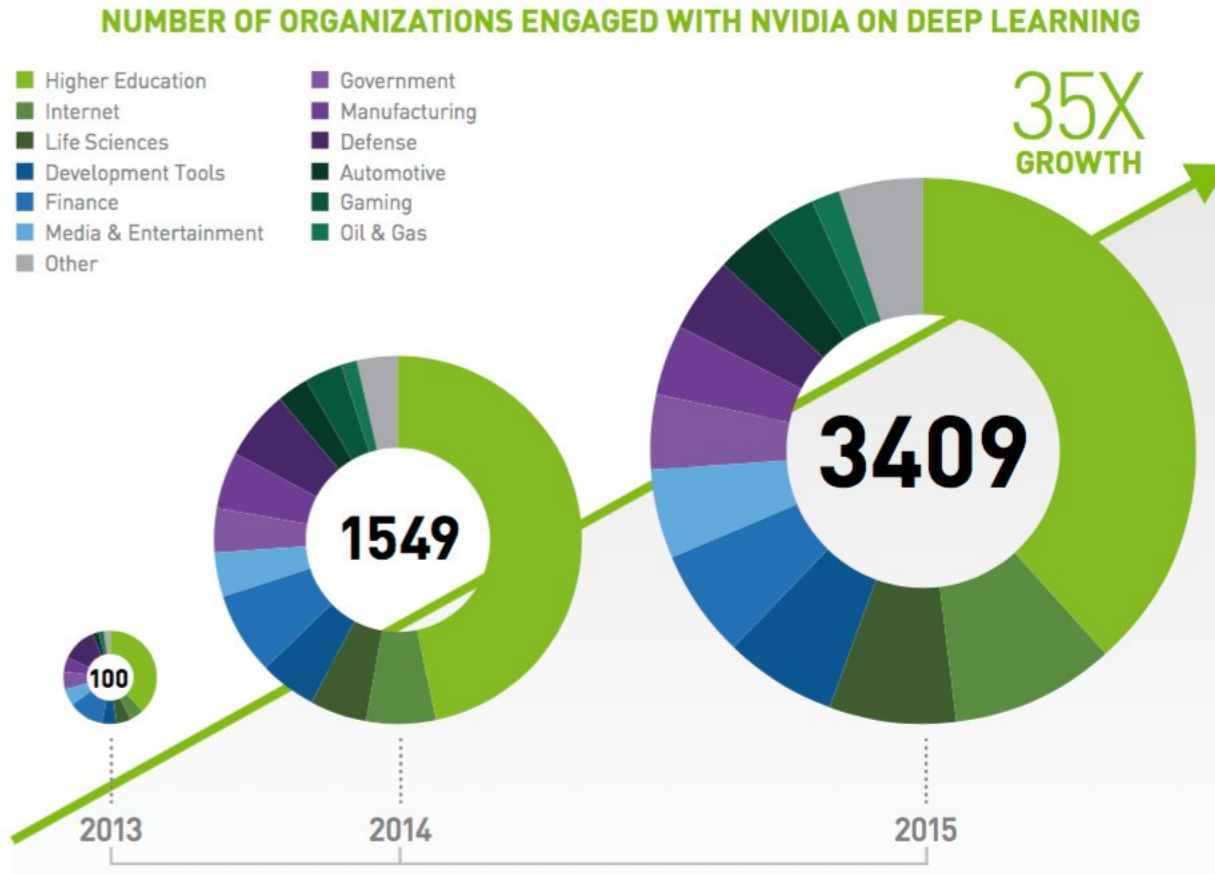
Machine learning received the most investment, although boundaries between technologies are not clear-cut

External investment in AI-focused companies by technology category, 2016¹
\$ billion



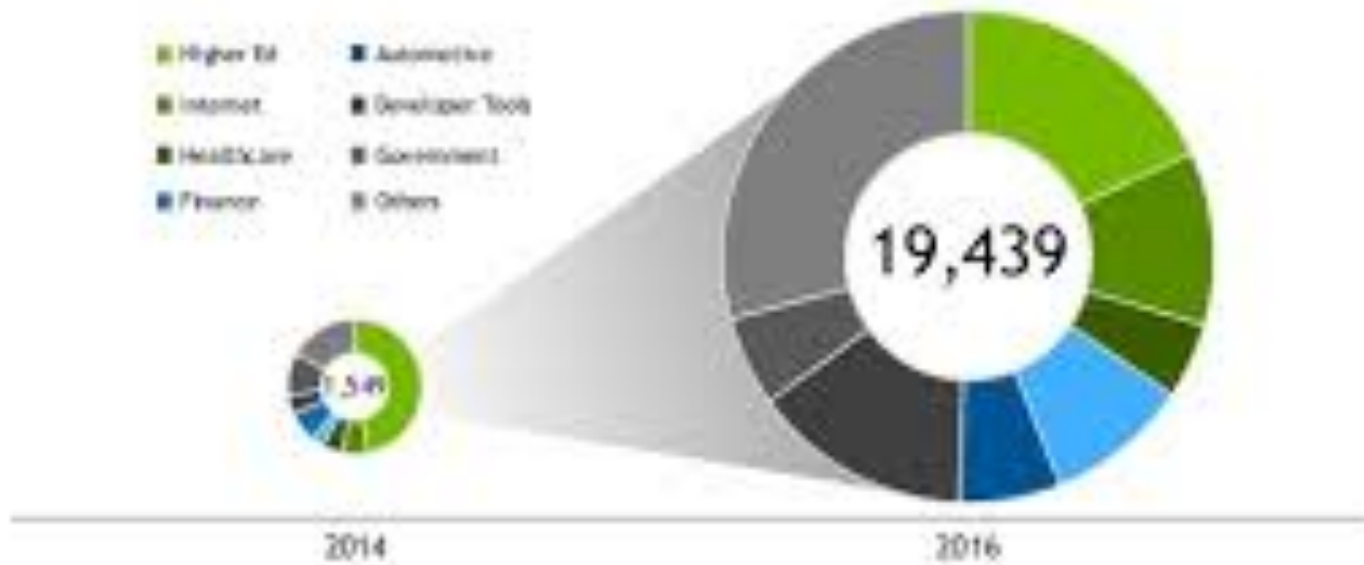
¹ Estimates consist of annual VC investment in AI-focused companies, PE investment in AI-related companies, and M&A by corporations. Includes only disclosed data available in databases, and assumes that all registered deals were completed within the year of transaction.

New GPU Computing Model for Artificial Intelligence



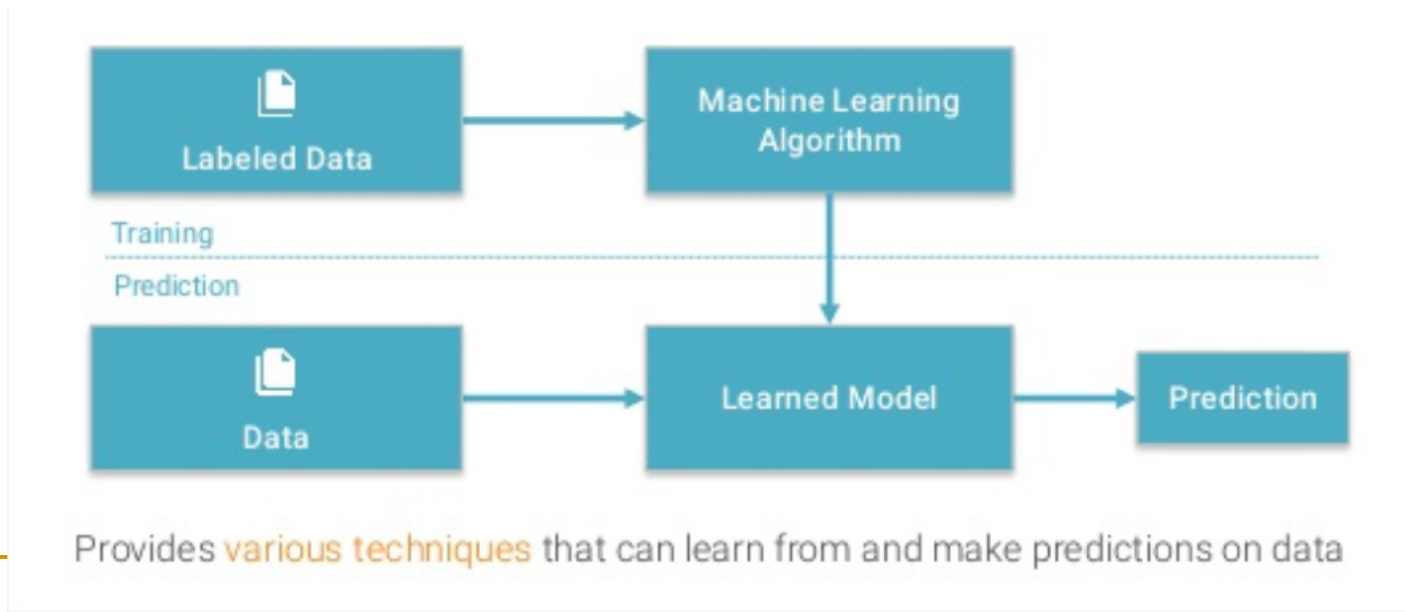
New GPU Computing Model for Artificial Intelligence

EVERY INDUSTRY HAS AWOKEN TO AI Organizations engaged with NVIDIA on Deep Learning




Machine Learning

- **Machine Learning** is a type of Artificial Intelligence that provides computers with the ability to **learn without being explicitly programmed**



Machine Learning: Hype of Reality?

- “Machine Learning is a core transformation way by which we are **rethinking everything** we are doing” – Sandar Pichai (CEO Google)



Learning Approaches



Supervised Learning: Learning with a **labeled training set**
Example: email spam detector with training set of already labeled emails



Unsupervised Learning: **Discovering patterns** in unlabeled data
Example: cluster similar documents based on the text content



Reinforcement Learning: learning based on **feedback** or reward
Example: learn to play chess by winning or losing

Deep Learning



Part of the machine learning field of learning representations of data. Exceptional effective at learning patterns.




Utilizes learning algorithms that derive meaning out of data by using a hierarchy of multiple layers that mimic the neural networks of our brain.




If you provide the system tons of information, it begins to understand it and respond in useful ways.

Deep Learning Basics

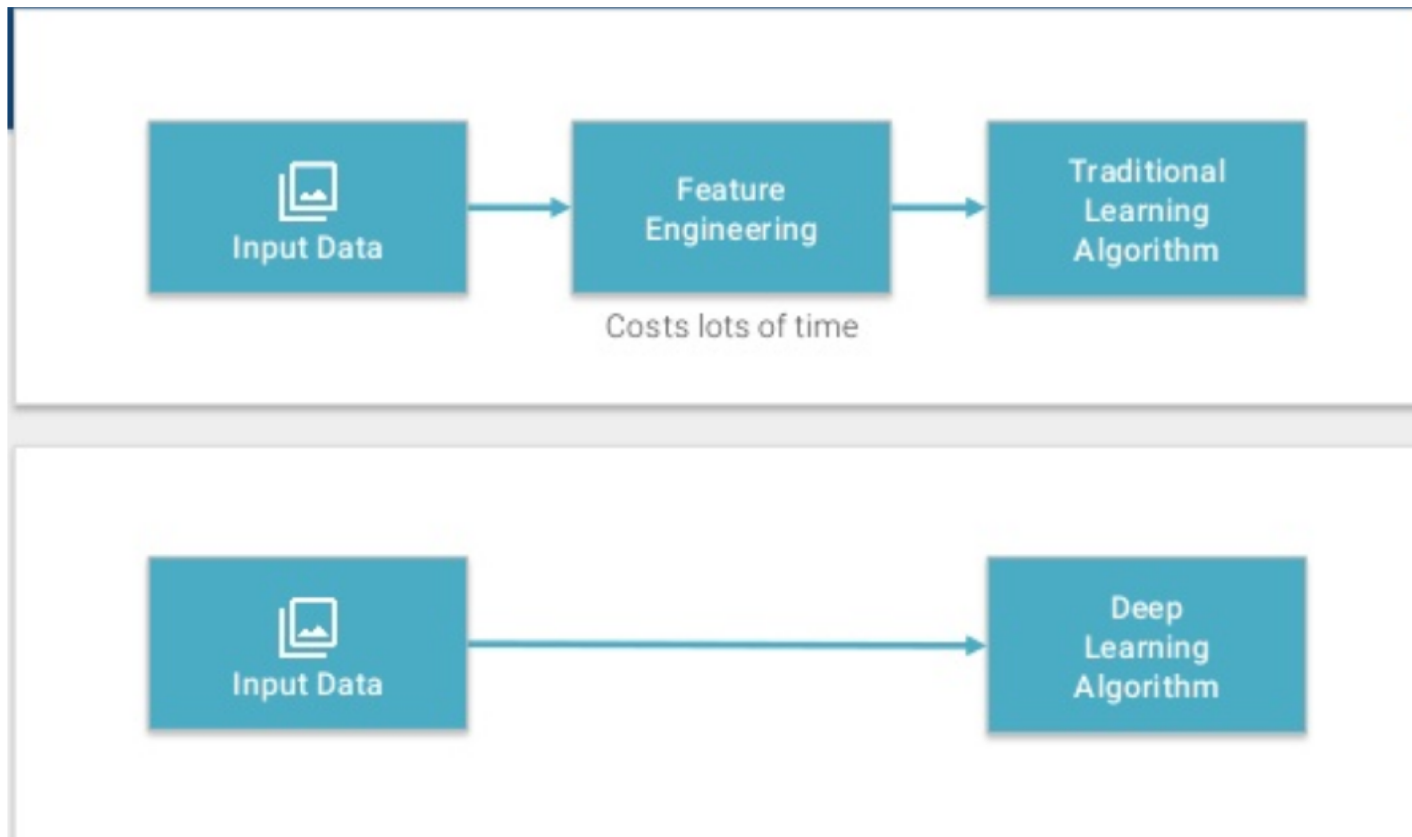


A deep neural network consists of a **hierarchy of layers**, whereby each layer **transforms the input data** into more abstract representations (e.g. edge -> nose -> face). The output layer combines those features to make predictions.

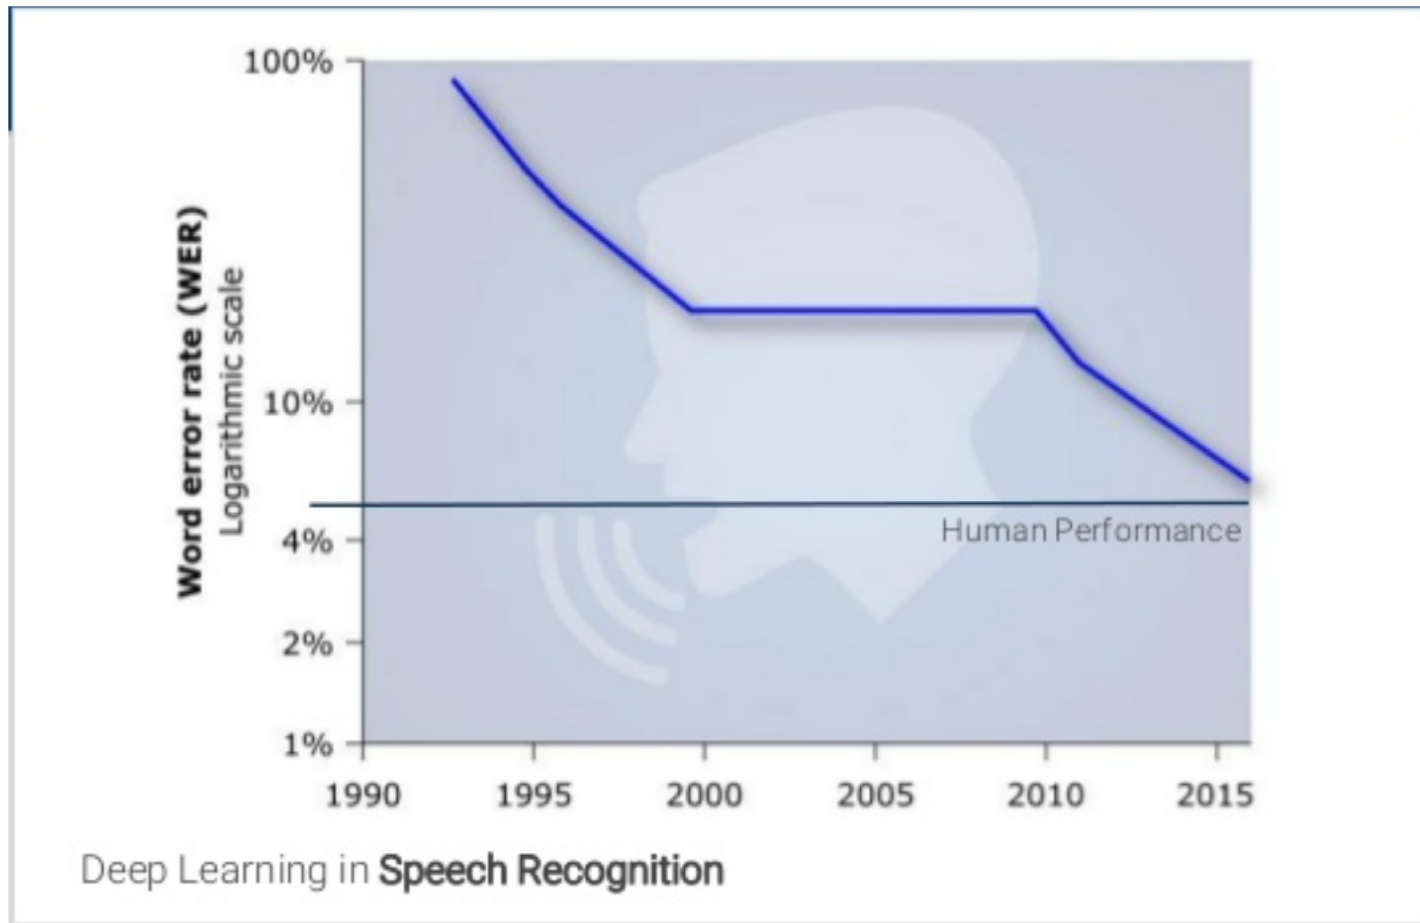
Deep Learning Basics: What did it learn?



Deep Learning: No More Feature Engineering



A Brief History: Big Bang aka “One net to rule them all”



Reading and Suggested Exercises

- Chapter 1
- Suggested Exercises: 1.2, 1.14, 1.15

Exercises to be submitted

- **Exercise 1.** Read those two papers, and discuss whether there are still up to date. For the debatable elements, provide your views. No more than one or two paragraphs per paper.
 - <https://homes.cs.washington.edu/~pedrod/papers/cacm12.pdf>
 - <http://www2.math.uu.se/~thulin/mm/breiman.pdf>

Exercises to be submitted (cont'd)

- **Exercise 2.** After going through exercise 1.15, can you find some 3 more recent contests in Artificial Intelligence?

Exercises to be submitted (cont'd)

■ Exercise 3.

- Examine the AI literature to comment (do you agree/disagree) on the following claims/criticisms of machine learning/deep learning? For each claim/criticism, provide 3-5 references and write 10-20 lines
 - *Claim 1.* Deep-learning networks are highly susceptible to the butterfly effect--small variations in the input data can lead to drastically different results, making them inherently unstable
 - *Claim 2.* It is difficult to understand how deep learning networks arrive at insights, therefore cannot be used in mission-critical applications like predictive maintenance or clinical decision-making
 - *Claim 3.* Deep learning networks require a lot of time for training, thereby making it very hard to quickly retrain models on the edge using newly available information