

Artificial Intelligence

Chapter 1

(Some slides adapted from Stuart Russel, Dan Klein, and others. Thanks guys!)

1

Chapter Outline

- Motivations to study AI...
- What is AI anyway?
- A brief history of the field
- The state of the art

2

So you wanna do AI?

- Why are you interested in AI personally? Why are you here?
 - ✓ Because I think Artificial Insemination is key to our nation's agricultural success? (fail)
 - ✓ I need another class to graduate
 - ✓ ???
- Why you should be interested...
 - Intelligence *defines us*, as human, as a species...
- Opportunity...
 - AI will be huge...
 - AI vs Physics as a field...

What is Artificial Intelligence?

- We could start by looking at (supposed) areas of application...
- ... and there are plenty...

What is Artificial Intelligence?

- Or maybe see what sorts of companies are working on what...

100 STARTUPS USING ARTIFICIAL INTELLIGENCE TO TRANSFORM INDUSTRIES


The infographic lists 100 startups across various industries:

- CONVERSATIONAL AI/ BOTS:** Intellifit, Mokuks, Automot, etc.
- VISION:** Clarifai, @ronocom, etc.
- AUTO:** @nocom, drive.ai, etc.
- ROBOTICS:** @nocom, Rekid, etc.
- CYBERSECURITY:** Cylance, silt science, etc.
- BUSINESS INTELLIGENCE & ANALYTICS:** DataRobot, etc.
- AD, SALES, CRM:** TalkIQ, etc.
- CORE AI:** @nocom, etc.
- HEALTHCARE:** @nocom, etc.
- TEXT ANALYSIS/ GENERATION:** @nocom, etc.
- IOT/IOT:** @nocom, etc.
- COMMERCE:** @nocom, etc.
- FINTECH & INSURANCE:** @nocom, etc.
- OTHER:** @nocom, etc.

Credit: <http://www.techforkorea.com>

What is Artificial Intelligence?

- Hmm, or maybe we could boil this down to some key "topic areas"...

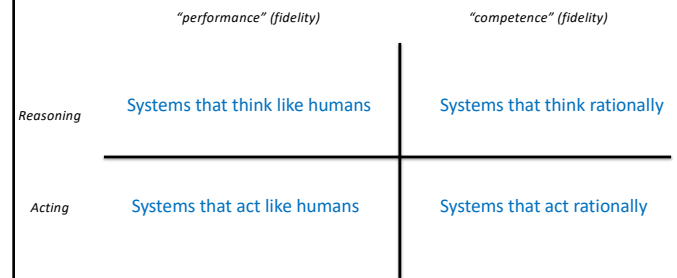


What is Artificial Intelligence?

- We're still not at the heart of it...
- **WHAT IS INTELLIGENCE REALLY?**
 ← Uhh... well... ya know... *thinking*. Too vague!
 ✓

What is AI?


- At the core, it's about software systems that...*behave in a certain way*.
 - Historically, we can discern four different perspectives
 - Similar functionally...but quite different philosophically



Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the **Imitation Game**



- Predicted that by 2000, 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 representation, reasoning, language understanding, learning

Problem: Turing test is not **reproducible, constructive**, or amenable to **mathematical analysis**

9

Thinking humanly: Cognitive Science

- 1960s "**cognitive revolution**": information-processing psychology replaced prevailing orthodoxy of **behaviorism**
- Requires scientific theories of internal activities of the brain
 - What level of abstraction? "Knowledge" or "circuits"?
 - How to validate? Requires either:
 - Predicting and testing behavior of human subjects (top-down)
 - Direct identification from neurological data (bottom-up)
 - ... plus, ideally, modeling findings in software!
- Both approaches (roughly, **Cognitive Science** and **Cognitive Neuroscience**) are now distinct from AI
- Both share with AI the following characteristic:
 - the available theories do not explain (or engender) anything resembling human-level general intelligence
- Hence, all three fields share one principal direction: focus on understanding intelligent behavior

10

Thinking rationally: Laws of Thought

- **Normative** (or **prescriptive**) rather than **descriptive** approach
- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of **logic**:
 - notation and rules of derivation for thought
 - Will see more when we look at reasoning agents.
- may or may not have proceeded the idea of mechanization
 - i.e. pure philosophy versus application orientation.
- Direct line through mathematics and philosophy to modern AI

Problems:

1. Not all intelligent behavior is mediated by logical deliberation
2. Not goal driven. What is the **purpose** of thinking? What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

11

Acting rationally: doing the "right" thing

- **Rational** behavior: doing "the right thing"
 - Don't worry about how humans perform it
 - Don't worry about logical truth
 - Focus on results:
- **The right thing:** that which is expected to maximize goal achievement, given the available information.
- Doesn't necessarily involve thinking
 - e.g., blinking reflex...but thinking should be in the service of rational action
- This course in AI is about **engineering**
 - About how to build it, how to make it happen
 - Not about philosophy...or even theory of cognition
- Thus: we will focus on this practical view of AI.
 - How can we make the machine intelligently solve problems?
 - Formally: Designing **agents** that act rationally

12

Rational agents

- An **agent** is an entity that perceives and acts
 - This course is about designing **rational agents**
- Defn: a software agent that acts to achieve best **expected** outcome modulo:
 - Available knowledge at that moment
 - Uncertainty of knowledge that it does have
- Or often, realistically: **Limited rationality**: take the most rational action *given some time limit to act*.
- Abstractly, an agent is a function from percept histories to actions:

$$f : P^* \rightarrow A$$
- For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: **computational limitations make perfect rationality unachievable**
 - design **best program for given machine/situational resources**

13

AI prehistory: Influences

Philosophy	logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof, concept of algorithms, computation, (un)decidability, (in)tractability, probability
Psychology	Adaptation, phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions (decision theory), Game theory, Max-Min strategies, Adversarial reasoning
Linguistics	knowledge representation, grammar (for NLP)
Neuroscience	Model of plastic physical substrate for mental activity
Control Theory	homeostatic systems, stability, simple optimal agent designs

14

Brief glance: history of AI

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1952-69	Look Ma, no hands! Early automatons...
1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Universal solver: Robinson's complete algorithm for logical reasoning
1966-74	AI discovers computational complexity Neural network research almost disappears
1969-79	Early development of knowledge-based systems
1980-1988	Expert systems industry booms
1988-93	Expert systems industry busts: "AI Winter"
1985-95	Neural networks concepts resuscitated...a new way forward
1988-	Resurgence of probability; general increase in technical depth "Nouvelle AI": ALife, GAs, soft computing
1995-	Agents, agents, everywhere . . .
2003-	Human-level AI back on the agenda (the next bubble?)

15

Quick quiz: How much do you know about AI?



Which of the following can be done at present?

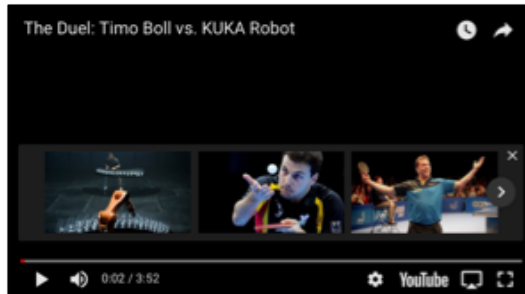
- Play a decent game of table tennis
- Drive safely along a curving mountain road
- Drive safely in rush downtown phoenix
- Buy a week's worth of groceries on the web
- Buy a week's worth of groceries at Bashas
- Play a decent game of bridge
- Beat world champions in GO
- Discover and prove a new mathematical theorem
- Design and execute a research program in molecular biology
- Give competent legal advice in a specialized area of law
- Converse successfully with another person for an hour
- Perform a complex surgical operation
- Unload any dishwasher and put everything away
- Write an intentionally funny story

16

State of the art

Which of the following can be done at present?

- Play a decent game of table tennis (Kuka Robotics, 2014)



- Truth: No, it can't really beat Timo Boll. But robots can play decent parlor ping-pong...and are getting better. See: <https://spectrum.ieee.org/automaton/robotics/industrial-robots/robots-playing-ping-pong-whats-real-and-whats-not>

State of the art

Which of the following can be done at present?

- Play a decent game of table tennis
- Drive safely along a curving mountain road
- Drive safely in rush downtown phoenix
- Buy a week's worth of groceries on the web
- Buy a week's worth of groceries at Bashas
- Play a decent game of bridge
- Beat world champions in GO
- Discover and prove a new mathematical theorem
- Design and execute a research program in molecular biology
- Give competent legal advice in a specialized area of law
- Converse successfully with another person for an hour
- Perform a complex surgical operation
- Unload any dishwasher and put everything away
- Write an intentionally funny story



18

Unintentionally funny stories: The best AI can do...

One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe threatened to hit Irving if he didn't tell him where some honey was. The End.

Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.

Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.

19

Natural Language Processing

Speech technologies

- Automatic speech recognition (ASR)
- Text-to-speech synthesis (TTS)
- Dialog systems




Language processing technologies

- Machine translation:
 - Aux dres de son président, la commission serait en mesure de le faire .
 - According to the president, the commission would be able to do so .
 - Il faut du sang dans les veines et du cran .
 - We must blood in the veins and the courage .
 - There is no backbone , and no teeth .
- Information extraction
- Information retrieval, question answering
- Text classification, spam filtering, etc...



Computer Vision



Images from Jitendra Malik

Robotics

- **Robotics**
 - Part mech. eng.
 - Part AI
 - Reality much harder than simulations!
- **Technologies**
 - Vehicles
 - Rescue
 - Soccer!
 - Lots of automation...
- **Social Disruptor?**
 - Manufacturing automation
 - Transportation
 - The end of blue collar?



Images from stanfordracing.org, CMU RoboCup, Honda ASIMO sites

Logic and Decision Making

- **Logical Systems**
 - Theorem Provers
 - NASA Fault diagnosis
 - Question answering (Prolog +++)
 - Ex: Proved Robbins Conjecture, unsolved for decades.
- **Solution Methods:**
 - Deduction Systems
 - Constraint satisfaction models
 - Satisfaction solvers (huge advances!)
- **Decision Support, Planning**
 - Scheduling (e.g. airline routing, military)
 - Ex: During 1991 Gulf War, US forces deployed AI logistics planning/scheduling for 50k vehicles, cargos, etc.
 - Route planning (Google Maps)
 - Medical Diagnosis (e.g. Pathfinder sys)
 - Automated help desks
 - Fraud detection



Game Playing

- **May 1997: IBM Deep Blue rules chess**
 - Beats Gary Kasparov, world champion
 - First match won against a world champion
 - Integrates "intelligent creative play"
 - Examines 2M boards...per second.
 - Humans understood 99.9% of Deep Blue's moves.
 - Can do about the same with decent cluster today...
- **2005: Machines can't play GO**
 - Several orders of magnitude harder than chess computationally
 - Human experts scoff at machines chances. Very weak play.
- **2016: IBMs GO spans world champions**
 - Secretly entered in online tournament
 - Beats the world champion...four times in a row.
- **Open questions:**
 - How can human compete at all? How does "slow" human cognition do it?



The new Difficult Ethics of AI

- AI will pose some hard questions for us all ... soon...
 - Who is liable if a robot driver has an accident?
 - "Self-Driving Tesla Was Involved in Fatal Crash, U.S. Says", NY-Times, 6/30/2016
 - Will you like the decisions that machines make? Is rational always right?
 - "Self-driving cars programmed to decide who dies in a crash", USA Today, 11/23/17
 - Societal effects of...well... "human obsolescence"?
 - What will large people do when the work is done by machines?
 - How will it re-distribute social power structures?
 - Colossal restructuring of human life. Whole model is built around work...
 - Will machines surpass human intelligence? What will we do with superintelligent machines?
 - the "singularity"
 - Make Terminators? Save the human race?
 - Would intelligent machines have conscious existence? Legal rights?
 - Do human rights attach to biology... or cognition?

Thought experiment: Rationality wins? Always?

It's a bright, sunny day and you're zooming along alone in your spanking new self-driving vehicle. You're looking at the window at the great scenery...because you can!

As you approach a rise in the road, heading south, a school bus appears, driving north, one driven by a human, and it veers sharply toward you. There is no time to stop safely, and no time for you to take control of the car. Does the car:

- Swerve sharply into the trees, possibly killing you but possibly saving the bus and its occupants?
- Perform a sharp evasive maneuver around the bus and into the oncoming lane, possibly saving you, but sending the bus and its driver swerving into the trees, killing her and some of the children on board?
- Hit the bus, possibly killing you as well as the driver and kids on the bus?

The existential question:

Who dies when the car is forced into a no win situation?

It's unaddressed...even as legislation to allow masses of autonomous vehicles onto the road is moving through Congress.

Fun quote:

Daimler: "our its autonomous vehicles would prioritize the lives of its passengers over anyone outside the car." Market cars based on their selfish life preservation algos?!

Adapted from: <https://www.usatoday.com/story/money/cars/2017/11/23/self-driving-cars-programmed-decide-who-dies-crash/891493001/>

Summary: So where are we headed in the future?

- For sure: Better task-specific agents
 - Focus on one fairly narrow aspect of AI puzzle:
 - Speech recognition (e.g. Siri, Alexa, etc.)
 - Handwriting recognition (tons of note-taking apps)
 - Search engines (Google, Amazon, etc.)
 - "Smart" data aggregators (content-based news aggregators)
 - Vision systems (Kinect, Facial recognition, security)
 - Games (self-adapting, automatic generation. The holodeck...)
 - Continues clear trend of application-specific improvement in two decades
- Human level AI (HLAI) ??
 - Many wizened AI gurus (McCarthy, Minsky, Winston, etc.) have complained
 - "these dog-n-pony tricks are not "real" AI
 - Need to return to broad goal of integrated, flexible, *thinking* machines
 - Ex: Siri
 - ✓ Speech recognition
 - ✓ Intelligent search
 - ✓ Good speech production
 - ❖ Terrible "intelligence"!!!

