Intelligent Systems: Introduction

- Poole, chap. 1, 2
- Negnevitsky, chap. 1

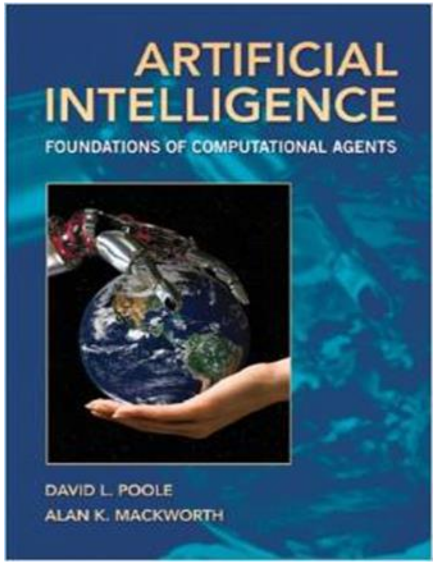

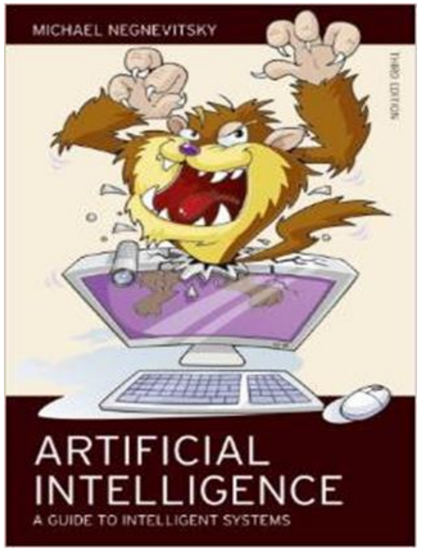

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?
- Agent acting in an environment
  - A test for intelligence
- Reducing Complexity in AI
- Dimensions of complexity
- Agents and Control
- What do we do in Intelligent Systems?
- Programming Languages
- Appendixes

What is Artificial Intelligence

- Artificial Intelligence is the synthesis and analysis of computational agents that act intelligently.
- An agent is something that acts in an environment.
- An agent acts intelligently if:
  - its actions are appropriate for its goals and circumstances
  - it is flexible to changing environments and goals
  - it learns from experience
  - it makes appropriate choices given perceptual and computational limitations

What is Intelligence, Anyways?

- are you intelligent if you:
  - can do complex arithmetic quickly?
    - humans are bad at it
    - … but computers are good at it
  - can recognize a face in a picture?
    - humans are good at it
    - … but hard to automate in a computer
  - hold a 15 min. conversation?
    - humans are good at it
    - … but really hard to automate in a computer

What is Intelligence?

- intellectual vs physical capabilities
  - a dog has a more acute sense of smell…
  - a bat can see at night…
- reflex vs planned/reasoned action
  - when the female wasp returns to her burrow with food, she first deposits it on the threshold, checks for intruders inside, and only then, if the coast is clear, carries her food inside.
  - but that's **instinctual** behavior
  - if the food is moved a few inches away while she is inside: on emerging, she will repeat the whole procedure as often as the food is displaced.
- awareness of existence (consciousness of itself)
  - if a system passes a test for intelligence but is not aware of it, is it intelligent?
  - but the only way to *really* know if a machine is thinking is to *be* the machine…

source of picture http://upload.wikimedia.org/wikipedia/en/8/85/Cicada_Killer_Wasp.jpg

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?

- Agent acting in an environment
  - A test for intelligence
- Reducing Complexity in AI
- Dimensions of complexity
- Agents and Control
- What do we do in Intelligent Systems?
- Programming Languages
- Appendixes

Goals of Artificial intelligence

- Scientific goal: to understand the principles that make intelligent behavior possible in natural or artificial systems.
  - analyze natural and artificial agents
  - formulate and test hypotheses about what it takes to construct intelligent agents
  - design, build, and experiment with computational systems that perform tasks that require intelligence
- Engineering goal: design useful, intelligent artifacts.
- Analogy between studying flying machines and thinking machines.

Agents acting in an environment

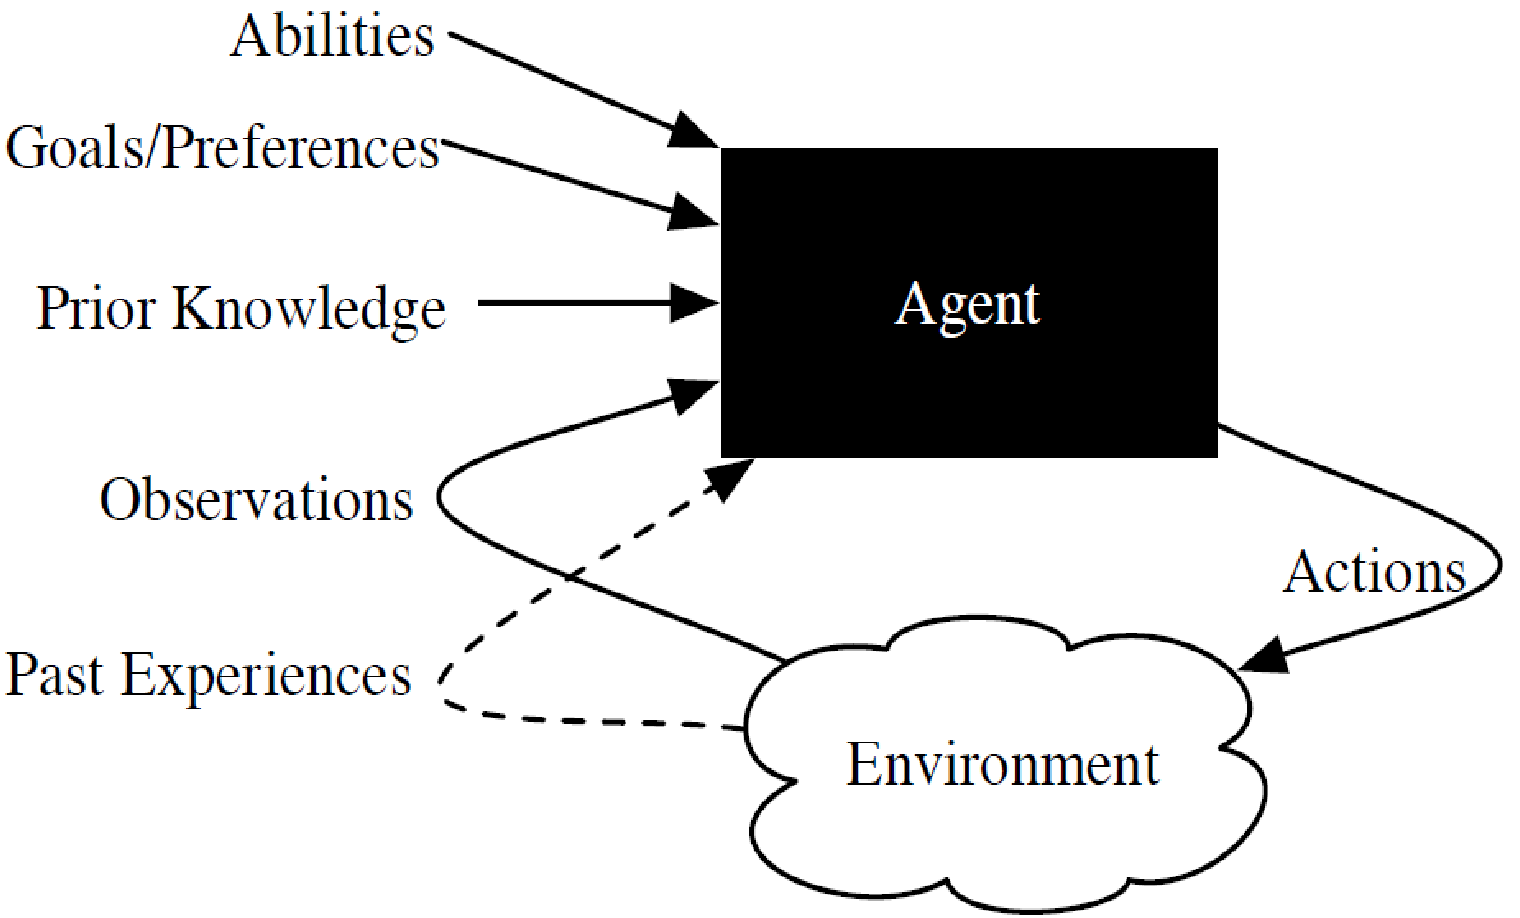

Inputs to an agent

- Abilities — the set of things it can do
- Goals/Preferences — what it wants, its desires, its values,...
- Prior Knowledge — what it comes into being knowing, what it doesn't get from experience,...
- History of observations (percepts, stimuli) of the environment
  - (current) observations — what it observes now
  - past experiences — what it has observed in the past

Examples of Agents

- Organisations Microsoft, Al Qaeda, Government of Canada, ...
- People teachers, physicians, stock traders, engineers, researchers, travel agents, farmers, waiters...
- Computers/devices thermostats, user interfaces, airplane controllers, network controllers, games, advising systems, tutoring systems, diagnostic assistants, robots, Google car, Mars rover...
- Animals dogs, mice, birds, insects, worms, bacteria...

Example agents

- Robot

  - abilities: movement, grippers, speech, facial expressions,. . .
  - goals: deliver food, rescue people, score goals, explore,. . .
  - prior knowledge: what is important feature, categories of objects, what a sensor tell us,. . .
  - observations: vision, sonar, sound, speech recognition, gesture recognition,. . .
  - past experiences: effect of steering, slipperiness, how people move,. . .

- Teacher

  - abilities: present new concept, drill, give test, explain concept,. . .
  - goals: particular knowledge, skills, inquisitiveness, social skills,. . .
  - prior knowledge: subject material, teaching strategies,. . .
  - observations: test results, facial expressions, errors, focus,. . .
  - past experiences: prior test results, effects of teaching strategies, . . .

How about these agents

- Medical doctor
- Autonomous car
- Apple Inc.
- ....

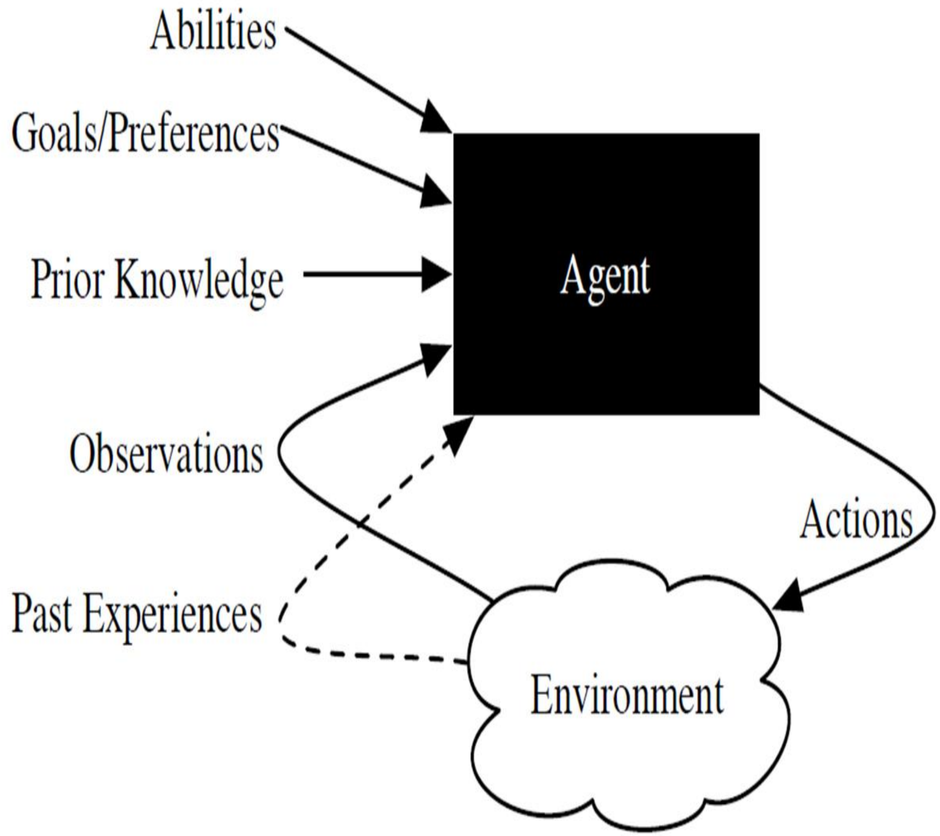

Is Deep Blue Intelligent?

- In 1996 and 1997 IBM's Deep Blue beat the human chess champion Kasparov in a six-games match.
- But Deep Blue uses:
  - plain brute force technique
  - on a massively parallel supercomputer
  - can explore 200,000,000 positions per second (Kasparov can examine 3/sec)
- Today, emphasis on more *intelligent* chess programs
- in Nov. 2006, Deep Fritz vs. Kramnik, ran on an ordinary Intel Core 2 Duo CPUs

source of image:
http://upload.wikimedia.org/wikipedia/en/c/c6/P11_kasparov_breakout.jpg

Is Chess Playing Intelligent?

"Chess is far easier than innumerable tasks performed by an infant, such as understanding a simple story, recognizing objects and their relationships, understanding speech, and so forth. For these and nearly all realistic AI problems, the brute force methods in Deep Blue are hopelessly inadequate."

– David Stork

Is Watson Intelligent?

- In 2011, IBM's Watson competed on *Jeopardy!*
- Watson beat Brad Rutter, the biggest all-time money winner on *Jeopardy!*, and Ken Jennings, the record holder for the longest championship streak
- Watson received the first prize of $1 million
- Watson is a question answering system... "an application of advanced Natural Language Processing, Information Retrieval, Knowledge Representation and Reasoning, and Machine Learning technologies to the field of open domain question answering"

source: http://en.wikipedia.org/wiki/Watson_%28artificial_intelligence_software%29

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?
- Agent acting in an environment
  - A test for intelligence

- Reducing Complexity in AI
- Dimensions of complexity
- Agents and Control
- What do we do in Intelligent Systems?
- Programming Languages
- Appendixes

A Test for Intelligence…

- The Turing Test
  - The "imitation game"
  - Proposed by Alan Turing in 1950
  - If a human interrogator cannot tell the computer and human apart, then the computer is intelligent
  - Measures the intelligence of a computer vs. a human
  - Turing predicted that by 2000, a machine might have a 30% chance of fooling a person for 5 minutes

The Turing Test

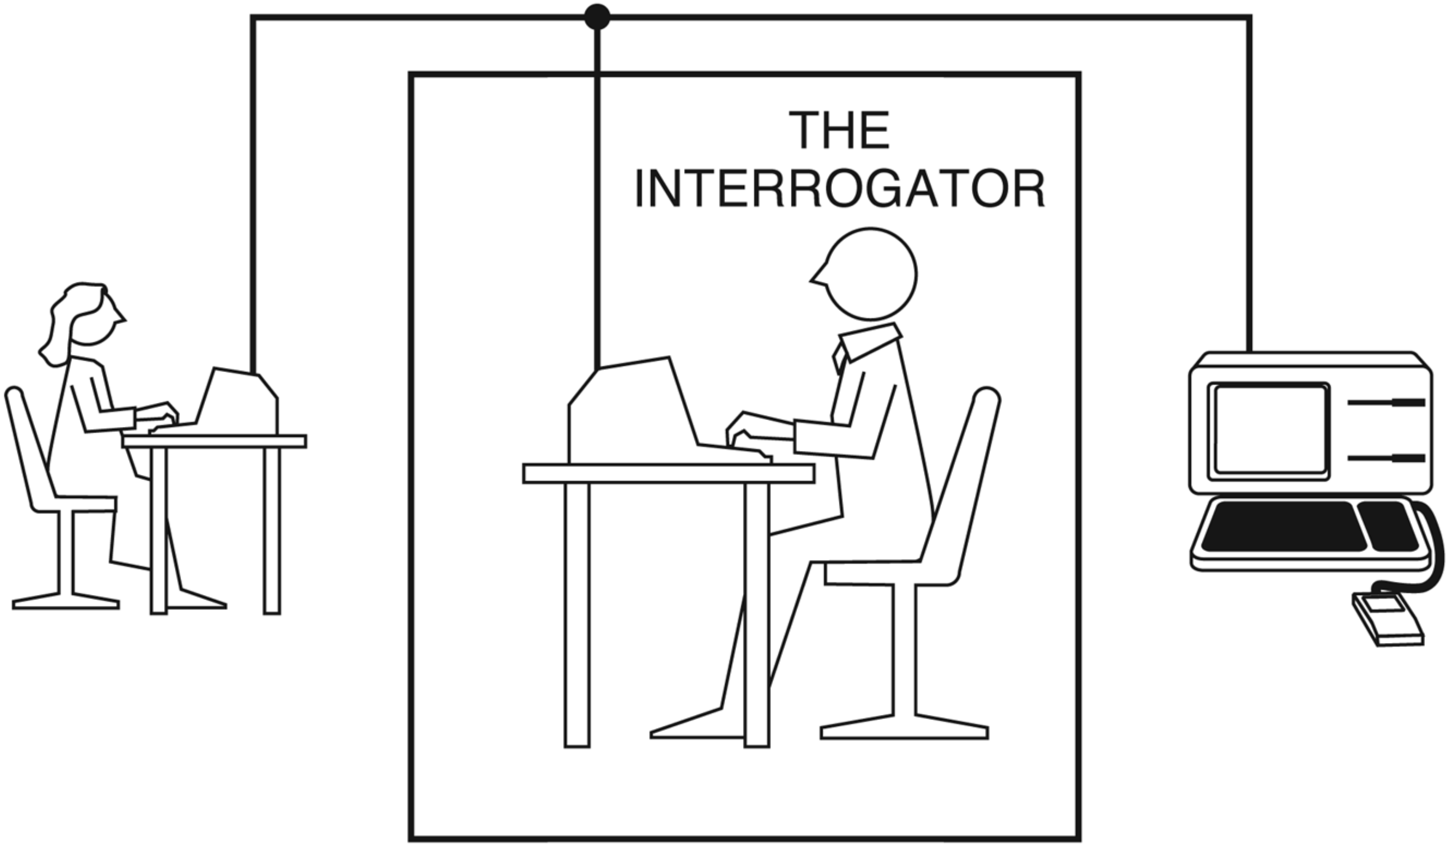

- A human mediates between the interrogator and the machine

The Turing Test

- Some capabilities required to pass the Turing test:
  - Natural Language Processing (NLP) to communicate
  - Knowledge Representation to store knowledge
  - Automated Reasoning to infer new knowledge
  - Machine Learning
  - ...

Arguments For the Turing Test

- Objective notion of intelligence
- Prevents us from arguments about the computer's consciousness
- Eliminates bias in favor of humans
- …

Current Turing Test

- No computer has passed Turing Test to date.
- The Long Bets Foundation has $20,000 bet between
  - Mitchell Kapor, founder of Lotus Development, and
  - Ray Kurzweil, inventor
  - Kapor bets that "*By 2029 no computer - or "machine intelligence" - will have passed the Turing Test. "*

Current Turing Test

- The Loebner Price
  - Goal: fool the judge into thinking that a program is a human.
  - started 1990
  - $100,000 for the 1st computer whose responses were indistinguishable from a human's… never awarded
  - In 2012, 4000$ will be given to the most human-like system
  - 25 minutes interaction
  - major criticism: does not advance AI.
    - ex: interaction is through text… so systems simulate *human* keyboard typing (various speeds, spelling mistakes and corrections,…)

Current Turing Test

CAPTCHA:

- **C**ompletely **A**utomated **P**ublic **T**uring test to tell **C**omputers and **H**umans **A**part
- the system asks a user to complete a test which the computer is able to generate and grade, but not able to solve.
- Because computers are unable to solve the CAPTCHA, any user entering a correct solution is presumed to be human.
- also known as **reverse Turing test**, because it is:
  - given by a machine and targeted to a human
  - in contrast to the Turing test that is given by a human and targeted to a machine.

source of image:
http://upload.wikimedia.org/wikipedia/commons/b/b6/Modern-captcha.jpg

How intelligent are these systems?

- SIRI: intelligent personal assistant and knowledge navigator. The application uses a natural language user interface to answer questions, make recommendations, and perform actions by delegating requests to a set of Web services.
- reCAPTCHA: asks users to enter words seen in distorted text images onscreen, which helps digitize the text of books.
- Google Self-Driving Car (Google Chauffeur)
- Watson supercomputer physician: is about to begin work evaluating evidence-based cancer treatment options that can be delivered to the physician in a matter of seconds for assessment.
- Coursera: online course offering
- Space Agents: *for Orbits* exploration mission of the European Space Agency and NASA

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?
- Agent acting in an environment
  - A test for intelligence
- Reducing Complexity in AI

- Dimensions of complexity
- Agents and Control
- What do we do in Intelligent Systems?
- Programming Languages
- Appendixes

Reducing Complexity in AI

- characterize simplifying assumptions made in building AI systems
- determine what simplifying assumptions particular AI systems are making
- suggest what assumptions to lift to build a more intelligent system than an existing one

Dimensions

- Research proceeds by making simplifying assumptions, and gradually reducing them.
- Each simplifying assumption gives a dimension of complexity
  - multiple values in a dimension: from simple to complex
  - simplifying assumptions can be relaxed in various combinations

Dimensions of Complexity

- Flat or modular or hierarchical
- Explicit states or features or individuals and relations
- Static or finite stage or indefinite stage or infinite stage
- Fully observable or partially observable
- Deterministic or stochastic dynamics
- Goals or complex preferences
- Single-agent or multiple agents
- Knowledge is given or knowledge is learned from experience
- Perfect rationality or bounded rationality

Modularity

- Model at one level of abstraction: flat
- Model with interacting modules that can be understood separately: modular
- Model with modules that are (recursively) decomposed into modules: hierarchical
- Example: Planning a trip from here to a see the Mona Lisa in Paris.
- Flat representations are adequate for simple systems.
- Complex biological systems, computer systems, organizations are all hierarchical
- A flat description is either continuous or discrete. Hierarchical reasoning is often a hybrid of continuous and discrete.

Succinctness and Expressiveness

Much of modern AI is about finding compact representations and exploiting the compactness for computational gains.
A agent can reason in terms of:

- Explicit states — a state is one way the world could be
- Features or propositions.
  - States can be described using features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.
- Individuals and relations
  - There is a feature for each relationship on each tuple of individuals.
  - Often an agent can reason without knowing the individuals or when there are infinitely many individuals.

Planning horizon

...how far the agent looks into the future when deciding what to do.

- Static: world does not change
- Finite stage: agent reasons about a fixed finite number of time steps
- Indefinite stage: agent reasons about a finite, but not predetermined, number of time steps
- Infinite stage: the agent plans for going on forever (process oriented)

Uncertainty

There are two dimensions for uncertainty. In each dimension an agent can have

- No uncertainty: the agent knows which world is true
- Disjunctive uncertainty: there is a set of worlds that are possible
- Probabilistic uncertainty: a probability distribution over the worlds.

Why Probability?

- Agents need to act even if they are uncertain.
- Predictions are needed to decide what to do:
  - definitive predictions: you will be run over tomorrow
  - disjunctions: be careful or you will be run over
  - point probabilities: probability you will be run over tomorrow is 0.002 if you are careful and 0.05 if you are not careful
  - probability ranges: you will be run over with probability in range [0.001,0.34]
- Acting is gambling: agents who don't use probabilities will lose to those who do.
- Probabilities can be learned from data and prior knowledge.

Uncertain dynamics

If an agent knew the initial state and its action, could it predict the resulting state?

The dynamics can be:

- Deterministic: the resulting state is determined from the action and the state
- Stochastic: there is uncertainty about the resulting state.

Sensing Uncertainty

Whether an agent can determine the state from its observations:

- Fully-observable: the agent can observe the state of the world.
- Partially-observable: there can be a number states that are possible given the agent's observations.

Goals or Complex Preferences

- achievement goal is a goal to achieve. This can be a complex logical formula.
- complex preferences may involve tradeoffs between various desiderata, perhaps at different times.
  - ordinal only the order matters
  - cardinal absolute values also matter
- Examples: coffee delivery robot, medical doctor

Single agent or multiple agents

- Single agent reasoning is where an agent assumes that any other agents are part of the environment.
- Multiple agent reasoning is when an agent reasons strategically about the reasoning of other agents.

Agents can have their own goals: cooperative, competitive, or goals can be independent of each other

Learning from experience

Whether the model is fully specified a priori:

- Knowledge is given.
- Knowledge is learned from data or past experience.

Perfect or bounded rationality

- Perfect rationality: the agent can determine the best course of action, without taking into account its limited computational resources.
- Bounded rationality: the agent must make good decisions based on its perceptual, computational and memory limitations.

Dimensions of complexity

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single-agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

State-space Search

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Classical Planning

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Decision Networks

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Markov Decision Processes (MDPs)

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Decision-theoretic Planning

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Reinforcement Learning

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Classical Game Theory

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Humans

- flat or modular or hierarchical
- explicit states or features or individuals and relations
- static or finite stage or indefinite stage or infinite stage
- fully observable or partially observable
- deterministic or stochastic dynamics
- goals or complex preferences
- single agent or multiple agents
- knowledge is given or knowledge is learned
- perfect rationality or bounded rationality

Dimensions Interact in Complex Ways

- Partial observability makes multi-agent and indefinite horizon reasoning more complex
- Modularity interacts with uncertainty and succinctness: some levels may be fully observable, some may be partially observable
- Three values of dimensions promise to make reasoning simpler for the agent:
  - Hierarchical reasoning
  - Individuals and relations
  - Bounded rationality

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?
- Agent acting in an environment
  - A test for intelligence
- Reducing Complexity in AI
- Dimensions of complexity
- Agents and Control

- What do we do in Intelligent Systems?
- Programming Languages
- Appendixes

Agents and control

Agents and Robots

A situated agent perceives, reasons, and acts in time in an environment.

- An agent is something that acts in the world.
- A purposive agent prefers some states of the world to other states, and acts to try to achieve worlds they prefer.
- Agents interact with the environment with a body.
- An embodied agent has a physical body.
- A robot is an artificial purposive embodied agent.

Agent Systems

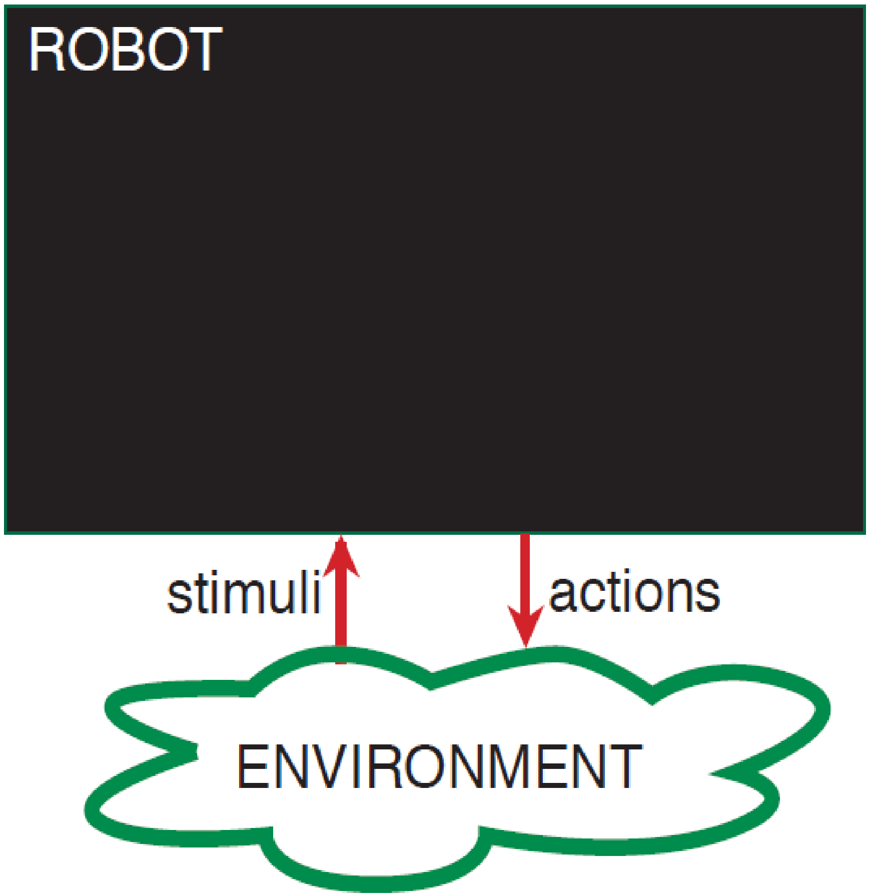

A agent system is made up of a agent and an environment.

- An agent receives stimuli from the environment
- An agent carries out actions in the environment.

Agent System Architecture

An agent is made up of a body and a controller.

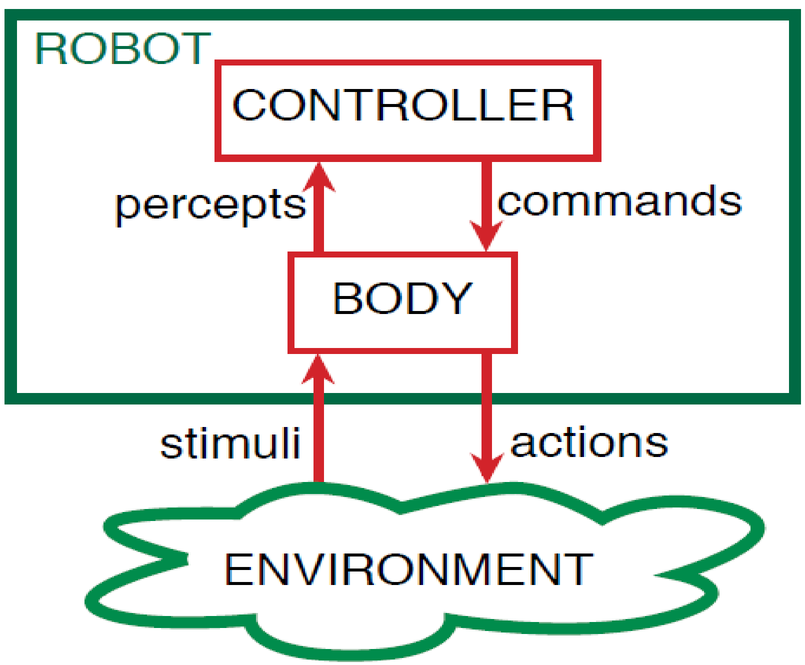

- An agent interacts with the environment through its body.
- The body is made up of:
  - sensors that interpret stimuli
  - actuators that carry out actions
- The controller receives percepts from the body.
- The controller sends commands to the body.
- The body can also have reactions that are not controlled.

Implementing a controller

- A controller is the brains of the agent.
- Agents are situated in time, they receive sensory data in time, and do actions in time.
- Controllers have (limited) memory and (limited) computational capabilities.
- The controller specifies the command at every time.
- The command at any time can depend on the current and previous percepts.

The agent Functions

- Let T be the set of time points.
- A percept trace is a sequence of all past, present, and future percepts received by the controller.
- A command trace is a sequence of all past, present, and future commands output by the controller.
- A transduction is a function from percept traces into command traces.
- A transduction is causal if the command trace up to time t depends only on percepts up to t.
- A controller is an implementation of a causal transduction.
- A causal transduction specifies a function from an agent's history at time t into its action at time t.

Belief States

- An agent doesn't have access to its entire history. It only has access to what it has remembered.
- The memory or belief state of an agent at time t encodes all of the agent's history that it has access to.
- The memory of an agent encapsulates the information about its past that it can use for current and future actions.
- At every time a controller has to decide on:
  - What should it do?
  - What should it remember?
    (How should it update its memory?)

— as a function of its percepts and its memory.

Controller

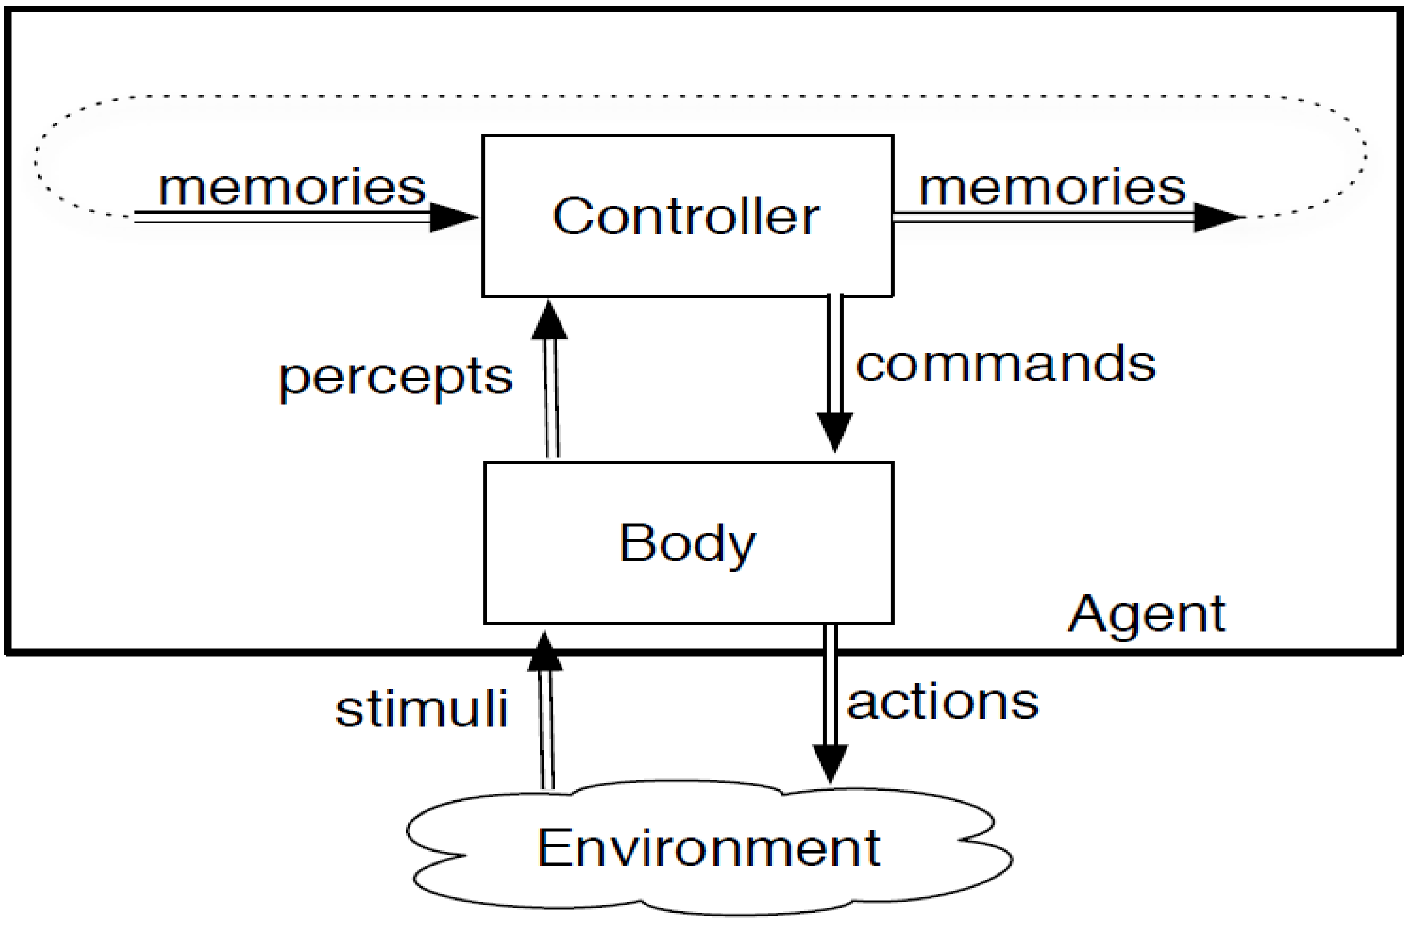

memories
Controller
memories
percepts
commands
Body
Agent
stimuli
actions
Environment

Functions implemented in Controller

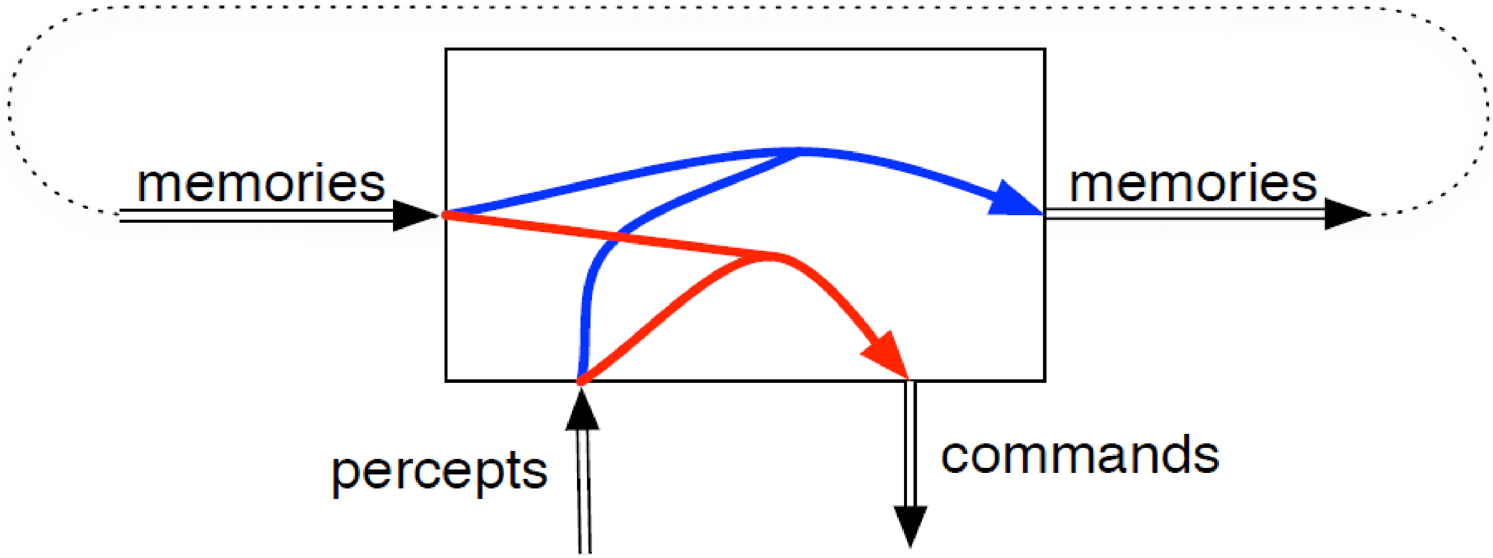

For discrete time, a controller implements:

- memory function *remember*(*memory*, *percept*), returns the next memory.
- command function *do*(*memory*, *percept*) returns the command for the agent.

Today

- Important questions
  1. What is artificial intelligence?
     - What is intelligence?
  2. What are the Goals of AI?
- Agent acting in an environment
  - A test for intelligence
- Reducing Complexity in AI
- Dimensions of complexity
- Agents and Control
- What do we do in Intelligent Systems?

- Programming Languages
- Appendixes

What do we do in Intelligent Systems?

- Knowledge representation
- Search, especially heuristic search (puzzles, games)
- Planning
- Reasoning under uncertainty, including probabilistic reasoning
- Learning
- Agent architectures
- Robotics and perception
- Natural language processing
- ...

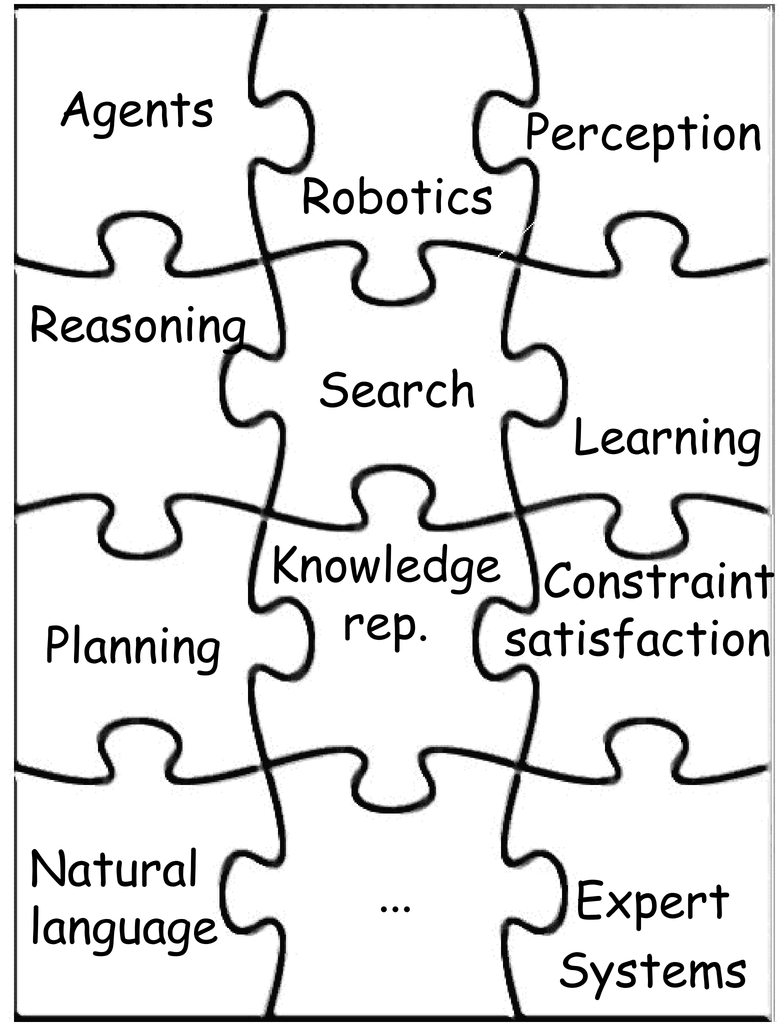

Programming Languages (1/2)

- Why do we need them?
  - Higher level of abstraction
  - Symbolic computation rather then numeric computation
    - `cold`, `hot` rather than `25.5°C`
    - `onTop(red)` rather than `position[1,0,0] = 50cm`
- Lisp
  - Functional language (based on recursive functions)
  - McCarthy in the 1950's
  - Gave rise to SCHEME, ML, …
- Prolog
  - Logical Language (based on first-order predicate logic)
  - Colmerauer in the 1970's for NL understanding
  - Based on Robinson's resolution method for theorem proving

Programming Languages (2/2)

- Approach
  - For cost effective reasoning and problems solving symbolic computation is used (e.g. Prolog, Lisp…, etc.)
  - For AI systems implementation general-purpose programming languages are typically used: C, C++, Java, …

Appendixes

1. What is artificial intelligence?

- No standard definition of AI among those working in the field
- AI has even been defined as:

"… the collection of problems and methodologies studied by artificial intelligence researchers."
- Luger and StubbleField

Other Definitions

Machines that **think like humans**

- *The exciting new effort to make computers think… machines with minds, in the full and literal sense* (Haugeland, 1985)
- *The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning…* (Bellmann 1978)

Machines that **act like humans**

- *The art of creating machines that perform functions that require intelligence when performed by people* (Kurzweil, 1990)
- *The study of how to make computers do things at which, at the moment, people are better.* (Rich and Knight, 1991)

Machines that **think intelligently**

- *The study of mental faculties through the use of computational models* (Charniak and McDermott, 1985)
- *The study of the computations that make it possible to perceive, reason, and act* (Winston, 1992)

Machines that **act intelligently**

- *A field of study that seeks to explain and emulate intelligent behavior in terms in terms of computational processes (Schalkoff, 1990)*
- *The branch of computer science that is concerned with the automation of intelligent behavior* (Luger & Stubblefield, 1993)

Approaches to AI

1. Engineering VS Cognitive Approach

- Engineering Approach:
  - Tries to find optimal solutions
  - No matter how (not necessarily what human do)
- Cognitive Approach:
  - Tries to understand the process
  - Tries to reproduce human behavior (even if wrong result)

- 4 points of view: Systems that can…

	Cognitive approach	**Engineering / Rational approach**
Behavior	act like humans	act intelligently
Reasoning	think like humans	think intelligently

Approaches to AI

2. Symbolic VS Sub-symbolic Approach

- Symbolic:
  - Reason from concepts/ideas/symbols
  - Associates intelligence with symbol manipulation
  - ex: Expert systems
  - ex: Many reasoning and planning systems
- Sub-symbolic:
  - Reason from the *underlying mechanism*
  - Reproduces intelligence through a lower level of abstraction
  - ex: neural networks -- same biological architecture ex: genetic algorithms -- manipulation of small features (genes)

Approaches to AI

3. Weak VS Strong AI

- Weak AI :
  - A system whose capabilities are **not** intended to match or exceed the capabilities of human beings.
  - A system *can* demonstrate intelligence, but does not necessarily have a mind, mental states or consciousness.
- Strong AI:
  - typically used in science fiction
  - A system that matches or exceeds human intelligence.
  - A system that could have: consciousness, self-awareness, the ability to *feel* sentiments, ...

Some Objections to AI

- Theological objection:
  - Thinking is part of humans' souls, and so animals/machines can't think.
- Head-in-the-sand objection:
  - Consequences of thinking machines are dreadful, so let's hope it's not possible.
- Machines will never be able to do X:
  - X = {be kind, friendly, have sense of humor, fall in love, etc.}

Why Does AI have a Bad Reputation?

- Unrealistic predictions
  - In 1950, Turing predicted that 50 years later (in 2000)
    - it will be possible to program a computer with ~100 Mb memory to pass the Turing Test 30% of the time, with 5 minute conversations.
    - It will be natural to speak of computers 'thinking'.
    - --> we still can't do that

  - Machine Translation:
    - In the 1950s, after World War II, we could translate automatically a few sentences from Russian to English.
    - Prediction: "Within three to five years, machine translation will be a solved problem."
    - --> we still can't do that

These Problems are Harder Than we Thought…

- MT has been the target of much [false?] criticism since the 50s
- ex: a common urban legend
  - in the 60s, the Department of Defense fed the phrase:
    "*the spirit is willing but the flesh is weak*"
  - into an English to Russian translator,
  - then fed the result back into a Russian to English system. the result came out:
    "*the vodka is strong but the meat is rotten*".
- MT is hard…
  - *e. g: les enfants et les femmes enceintes*
  - *--> pregnant children and women*

Some Successes of AI…

- A few years ago, all these were considered AI problems… now, no one thinks of them as AI
  - OCR – Optical Character Recognition
  - Speech Recognition
  - Information Retrieval
  - Spell checker and Grammar checker
  - Expert Systems
  - Data Mining
  - Word prediction
  - …
- A more pragmatic definition of AI today :

"AI research is that which computing scientists do not know how to do cost-effectively today."

History of AI

- Gestation period (1940-1956)
  - 1943: McCulloch & Pitts:
    - developed a boolean circuit model of the human brain
    - early work in artificial neural networks... but just a theory, no real implementation
  - 1950: Alan Turing
    - wrote *Computing Machinery and Intelligence*
    - fundamental paper in AI -- describes the Turing test
  - 1950: Claude Shannon
    - father of information theory
    - developed a Chess-Playing Machine
  - 1956: The Darmouth meeting
    - small get-together of the big guys
    - McCarthy, Minsky, Shannon & others
    - **the term "Artificial Intelligence" is first adopted**

History of AI

- The rise of AI (~1956 - late 70s)
  - Great enthusiasm, great ideas… but little success
  - 1958: John McCarthy develops LISP
  - Focus on heuristic search & learning techniques (ex. game playing)
  - Much work in early machine translation
  - Lots of work on knowledge representation
    - Minsky develops frames for knowledge representation
  - Lots of work on general methods for problem solving (weak methods)
    - Newell & Simon: development of GPS
    - 1965: Robinson develops a mechanical proof procedure, the Resolution Method
  - Lots of work on toy-world applications (that do not scale up…)
    - 1965: Weizenbaum develops ELIZA to simulate a dialogue between a psychotherapist and a patient
    - 1973: Winograd developsSHRDLU to understand English sentences in a restricted world and carried out instructions typed in English with a robot arm.
    - 1971-74: Sacerdoti et al. develops two of the first planning programs, STRIPS and ABSTRIPS
  - 1972: Colmeraurer develops Prolog

Example of a Toy-World

- toy-world = micro-word
- 1973: Winograd developed SHRDLU to understand English sentences in a restricted world and carried out instructions typed in English with a robot arm.

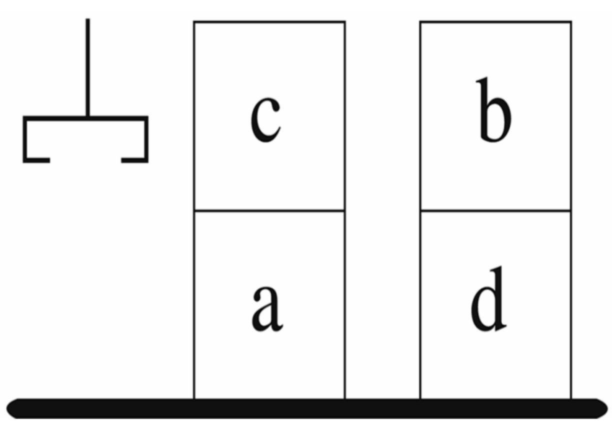

History of AI

- Reality hits (late 60s - early 70s)
  - 1966: the ALPAC report kills work in machine translation (and NLP in general)
  - People realized that scaling up from micro-worlds (toy-worlds) to reality is not just a manner of faster machines and larger memories…
  - Minsky & Papert's paper on the limits of perceptrons (cannot learn just any function…) kills work in neural networks
  - in 1971, the British government stops funding research in AI due to no significant results
  - it's the *AI Winter*…

History of AI

- Knowledge-based systems (70s - mid 80s)
  - Much work on expert systems
  - people realized that general-purpose problem solving (weak methods) do not work for practical applications
  - systems need specific domain-dependent knowledge (strong methods)
  - development of knowledge-intensive, rule-based techniques
  - 2 major expert systems
    - DENDRAL (1969): expert system to analyse chemical compounds for the molecular structure of Mars's soil
    - MYCIN (1972): expert system to diagnose infectious blood diseases
  - In the industry (1980s): First expert system shells and commercial applications.

History of AI (the end of *Good Old Fashioned AI*)

- Machine Learning is back (mid 80s-today)
  - More powerful machines -> usable implementation of neural networks
  - Huge developments in Neural Networks
  - Development of generic algorithms
  - Use of probability theory (Hidden Markov Models) for Natural Language Processing
- AI adopts the Scientific Method (mid 80s-today)
  - A good idea without a formal evaluation, just isn't enough anymore
  - Systems and theories are now rigorously evaluated on benchmarks and shared repositories
  - Regular "competition-style" conferences to evaluate approaches and compare them

History of AI

- Availability of very large datasets (2000-today)
  - Huge data sets are now available
    - document repositories in NLP,
    - datasets in ML,
    - billions on images for image retrieval,
    - billions of genomic sequences
    - …
  - So more emphasis in data-driven approaches, as opposed to algorithm-driven approaches
  - ie, use of generic machine learning algorithms on new datasets as opposed to developing sets of rules by hand.