CS 2710 Foundations of AI
Lecture 1

Course overview

Milos Hauskrecht

milos@pitt.edu

5329 Sennott Square

Course administrivia

Instructor: Prof. Milos Hauskrecht

5329 Sennott Square

milos@pitt.edu

TA: TBA

Course related information is now in Canvas

Temporary course web page:

http://people.cs.pitt.edu/~milos/courses/cs2710-Fall2020/

Course textbook

Course textbook:

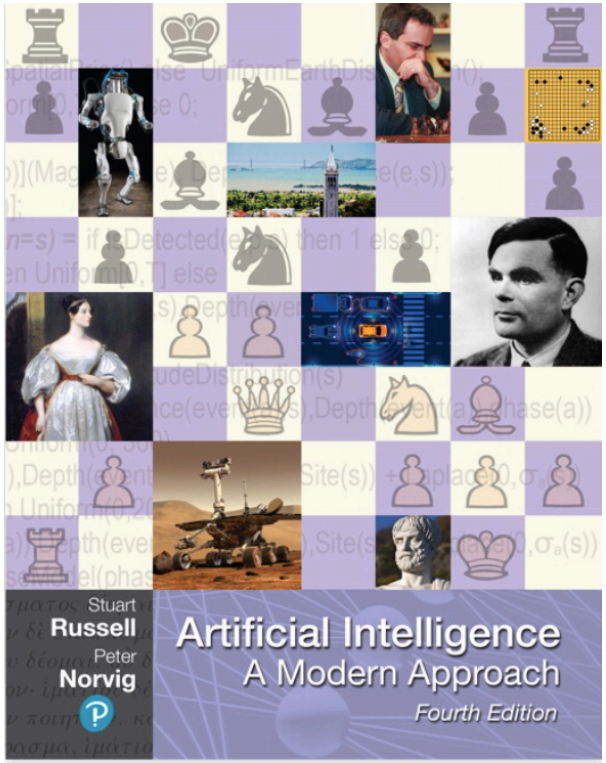

Stuart Russell, Peter Norvig.

Artificial Intelligence: A modern approach.

4th edition, *Pearson*, **2020**

Grading

- **Homework assignments 50%**
- **Exams: Midterm, Final 45%**
- **Lectures/short quizzes 5%**

Lectures

- **5 % of the grade**
- **Short quizzes + Activity**
  - Random
  - Short question(s) covering previous lectures

Homework assignements

- **Homework assignments:**
  - **50 % of the grade**
  - Weekly assignments: **Thursdays** (3 Thursdays cycle)
  - **Due at 4:15pm** (15 min before the lecture)
  - A mix of theoretical and programming problems
  - No extensions. Homework due dates are strict.
- **Collaborations:**
  - **No collaborations** on homework assignments
- **Programming part:**
  - Python
  - Python 3.6
- **Submissions** via **Canvas: report + programs**

Exams

- **Midterm**
  - around mid October
- **Final**
  - Cumulative (whole semester)

Academic honesty

- All the work in this course should be **done independently.**
- **Collaborations on homework assignments, quizzes and exams are not permitted.**
- Cheating and any other anti-intellectual behavior, including giving your work to someone else, will be dealt with severely.

- **Academic Integrity Policy for School of Computing and Information (SCI) :**

  http://sci.pitt.edu/current-students/policies/academic-integrity-policy/

Artificial Intelligence

- The field of **Artificial intelligence**:
  - The design and study of computer systems that behave intelligently
- **AI programs:**
  - Go beyond numerical computations and manipulations
  - Focus on problems that require reasoning (intelligence)
- **Why is AI research important?**
  - Engineering aspect
    - solving of hard problems
  - Cognitive aspect
    - Understanding the nature of human intelligence

Intelligence and machines

Can we make machines intelligent?

- **Symbolic system hypothesis** (Newell and Simon)
  - Intelligence is substrate neutral;
  - A physical symbol system has necessary and sufficient means for general intelligent action.
- **Biological substrate only** (John Searle, philosopher).
  - Intelligence is substrate dependent. The material humans are made of is fundamental for our intelligence. Thinking is possible only in special machines – living ones made of proteins.
- Other researchers believe that **sub-symbolic processing** (signal processing) may be needed to replicate intelligence.

Turing test

How to test the "intelligence" of a program?

- **Behavioral view of intelligence**
  - A behavior capable of fooling a human interrogator

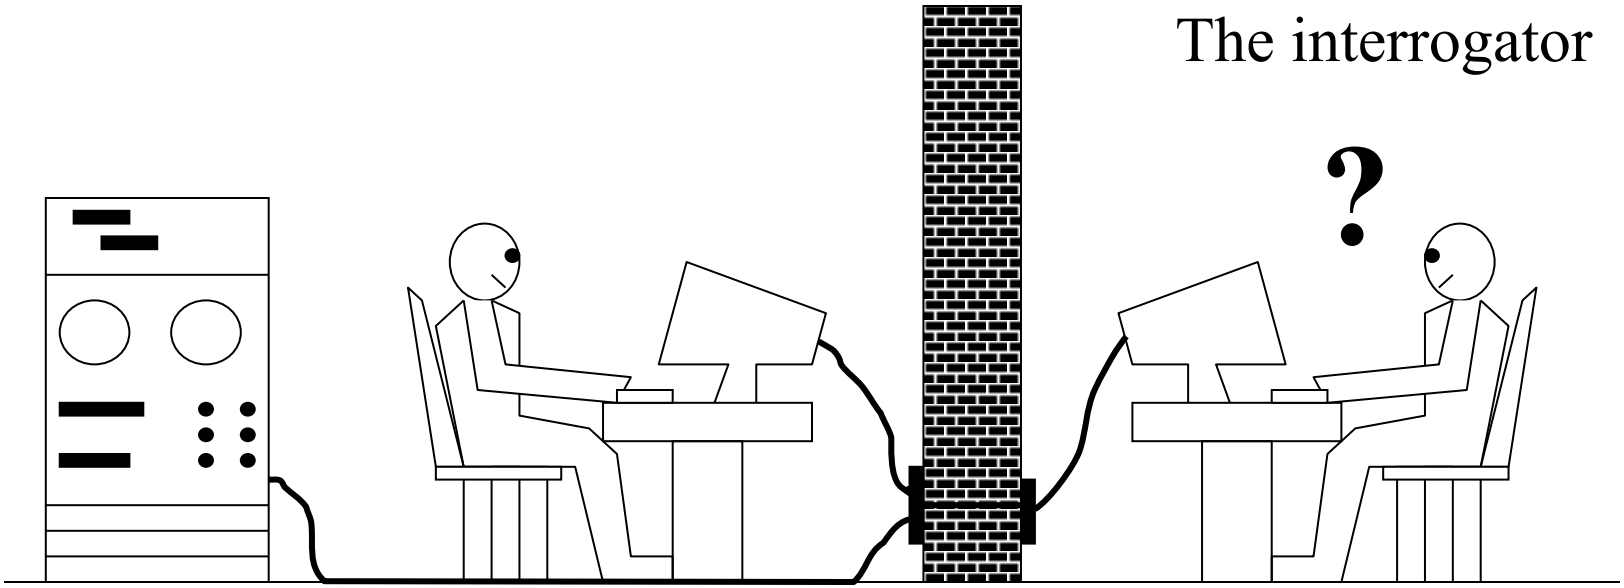

- Acting humanly is sufficient to pass the test.
- Is it necessary the AI system also thinks humanly?

Eliza program

- **Eliza, a computer therapist**, created by **Joseph Weizenbaum in 1966 at MIT**, is one of the most famous early AI programs.
- Eliza functions by "twisting the statements of her 'patients' back at them in the classic manner of a non-directive psychotherapist.“
- The fact that it understands the subject's statements is an illusion. Surprisingly, many users were taking its performance quite seriously.
- **You can run Eliza at:**
  - **http://www-ai.ijs.si/eliza/eliza.html**
- **Info:**
  - **http://web.mit.edu/STS001/www/Team7/eliza.html**

Eliza program. Sample dialogue.

Eliza: Hello. I am ELIZA. How can I help you?

You: I am frustrated by the traffic I have to face every morning.

Eliza: How long have you been frustrated by the traffic you have to face every morning?

You: For about a year now.

Eliza: Please go on.

You: Why should I go on? Give me some advice.

Eliza: Can you elaborate on that?

You: Why? Give me some solution.

Eliza: Can you elaborate on that?

You: It is obvious you cannot pass the Turing test

Eliza: Oh, i cant pass the turing test.

You: That is correct.

Eliza: Please go on.

…

Try it yourself at: http://www-ai.ijs.si/eliza/eliza.html

Passing the Turing test

More recent news on Turing Test: June 7, 2014

- A computer program that pretends to be a 13-year-old Ukrainian boy called Eugene Goostman passed a Turing test at the Royal Society in London on June 6, 2014 by convincing 33 percent of the judges that it was human during a five-minute typed conversation.

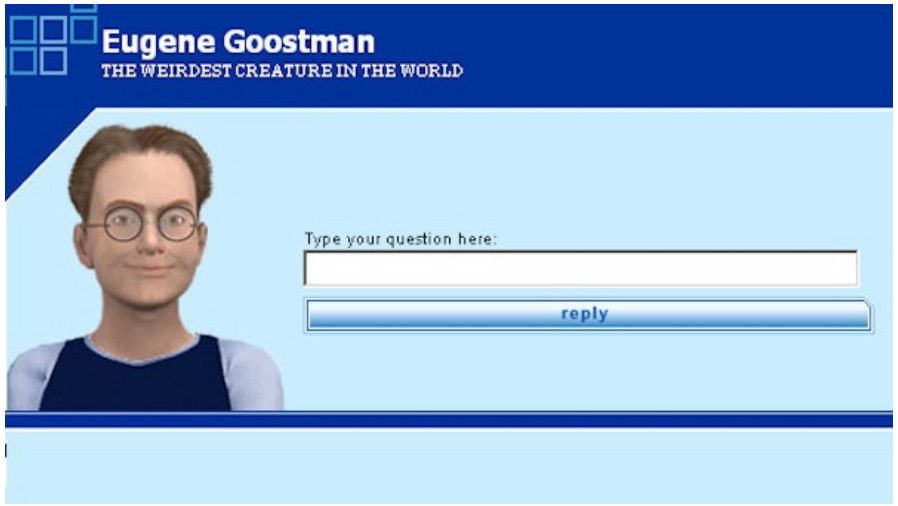

Evaluation of an AI system

Depends on what matters more.

- **Reasoning vs. Behavior**

  - the **computational process** or the **end-product** matters

- **Human performance vs. Rationality**
  - Compare against a human model (with its weaknesses) or a normative "ideal" model (rational system)

RN textbook

- The textbook adopts the **rational agent perspective**

  Focus on behavior and rational (normative) models

- **Agent:** an entity that perceives and acts
  - On abstract level the agent maps percepts to actions

$$f : Percepts \rightarrow Actions$$

- **Design goal:** for any given environment find the agent that performs the best with respect to some normative model

- **Caveat:** The design may be limited by resources: memory, time
  - Find agents with best resource-performance trade-off

History of AI

- **Artificial Intelligence** – name adopted at Dartmouth conference in 1956

- "Contemporary" AI starts in 20th century (1940s),
  But the origins go back many years.

Origins of AI:

- **Artificial people**.
  - Beings or devices capable of substituting or replacing humans in various activities.
- **Mathematical models of reasoning**.
  - Formal models of thought and reasoning.

Artificial people

- Beings or devices capable of substituting or replacing humans in various activities
- **Legends, stories:**
  - **Androids** (artificial people):
    - Android constructed by Albert the Great (13-th century)
    - Golem: made from clay, household chores (14-th century)
  - **Homunkulus** – a human-like being created in other than natural way (Paracelcus, 16-th century)
- **Mechanical people** capable of writing, drawing, playing instruments (18-th century)
- **Kempelen's chess machine** (18-th century).
- **Robots** Drama R.U.R. by K. Capek (early 20th century)

Mathematical models of reasoning.

- Philosophers and mathematicians worked on **models of reasoning and thought**.

- **Aristotle** (384-322 B.C), ancient Greece, philosopher
  - Tried to explain and codify certain types of **deductive reasoning** he called syllogisms.
- **George Boole** (1854)
  - Foundations of **propositional logic**
    - Formal language for making logical inferences.
- **Gottlieb Frege** (end of 19-th century).
  - **First order logic**

The beginnings of AI (40s-50s).

Two streams:

- **Neural network approach** (McCulloch and Pitts 1943).
  - Models of a human brain

- **Computer programs** capable of **simple reasoning tasks**:
  - chess programs (Shannon 1950, Newell, Shaw & Simon 1958)
  - checkers (Samuel 1959)
  - Theorem prover in geometry (Gelernter 1959)
  - Logic Theorist (Newell, Shaw & Simon 1957). Used propositional logic to prove theorems.

- **Dartmouth meeting (1956),** the name **Artificial Intelligence** adopted (proposed by John McCarthy)

60s.

Developments in the two streams:

- **Neural network models for learning patterns and pattern recognition**
  - **Objective:** replicate self-organization and subsequently phenomenon intelligence
  - Build on McCulloch and Pitts' work (1943)
  - **Adaline networks** (Widrow, Hoff 1960)
  - **Perceptrons** (Rosenblatt 1961)
  - Minsky and Papert (1969) – strong critique of perceptrons, it killed the area for a decade
- **Symbolic problem solvers**:
  - **General problem solver** (Newell, Simon) – think humanly
  - **LISP –** AI-specific programming language
  - **Micro-worlds –** focus on problem-solving in restricted worlds (e.g. blocks world)

70s. Knowledge-based system era.

- Early AI systems did not scale-up well to large applications
- The need for background knowledge

Edward Feigenbaum: **"knowledge is the power"**

Power of the system derived from the knowledge it uses

- **Expert systems:** obtain the knowledge from experts in the field, and replicate their problem-solving

Examples of KB/expert systems:

- **Dendral** system (Buchanan et al.). Molecular structure elicitation from mass spectrometer readings.
- **Mycin.** Diagnosis of bacterial infections.
- **Internist** (Pople, Myers, Miller). Medical diagnosis.

80s. AI goes commercial.

AI becomes an industry

- Many tools for the design of KB systems were developed

Revival of neural network (connectionist) approach.

- **Multi-layer neural networks**
  - Modeling and learning of non-linear functions.
  - Back-propagation algorithm (learning)

Failure of AI in 80s

- High expectations in very short time
- Computational complexity: some problems are intrinsically hard
- Modeling uncertainty
- Separation of connectionist - logic approaches.

90s. Recovery. Moving ahead

Modeling uncertainty (a breakthrough in late 80s)

- Bayesian belief networks, probabilistic graphical models.

Subcommunities/subareas covered originally by AI mature and develop:

- **Machine learning and data mining**
  - Analysis of large volumes of data
  - Finding patterns in data
  - Learning to predict, act
- **Image analysis and vision**
- **Natural language processing**
- **Autonomous agents** with intelligence:
  - Software agents
  - Robots

AI: this century

AI is more rigorous and depends strongly on: applied math, statistics, probability, control and decision theories

Advances:

- Machine Learning and Data mining
- Image analysis and vision
- Natural language processing
- Optimization
- Robotics

Applications:

- **Achieve partial intelligence** (not all human capabilities)
- Systems with components of intelligence in a specific application area;

AI: recent development

- **Success in solving many non-trivial problems**
- **AI effects everyday life**

Advances:

- Deep learning:
  - New solutions for many tasks in image/vision, speech recognition, game playing problems
- Mining large scientific and commercial datasets
  - Big data methods and analytics
  - Network analysis
- Robotics
  - New generation of autonomous systems

AI applications: Software systems

- **Diagnosis of:** software, technical components

- **Adaptive systems**
  - **Adapt systems to user needs**
  - **Adapt systems to specific tasks**

- **Examples:**
  - **Intelligent interfaces**
  - **Intelligent helper applications**
  - **Collaborative filtering**
  - **Target advertising**

Search and information retrieval

Web search engines

- Improve the quality of search
- Rely on methods/algorithms developed in AI
- Add inferences and knowledge to search queries

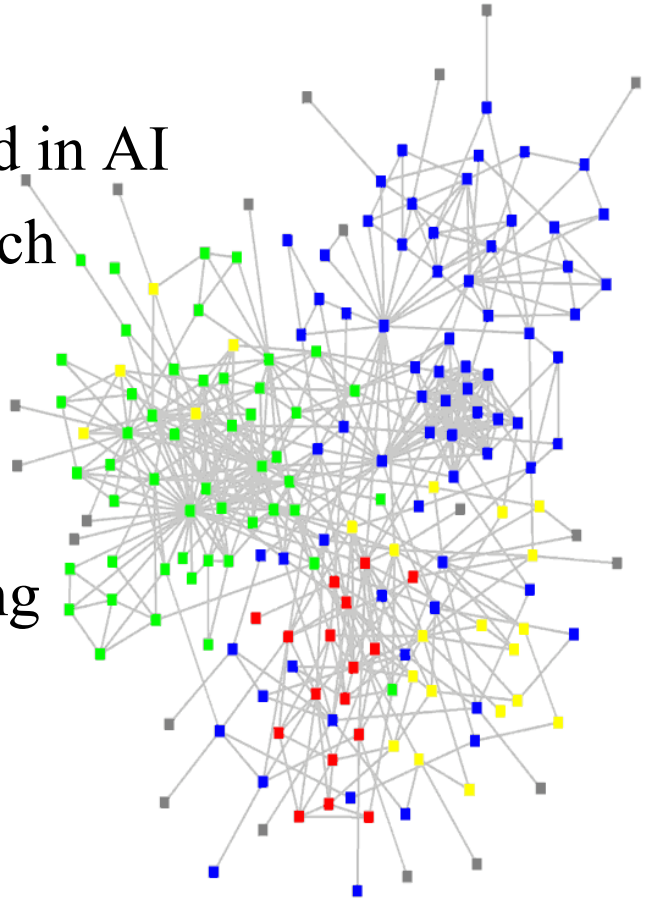

Semantic web (or web 2):

- From information to knowledge sharing
- Ontology languages

Speech recognition

Speech recognition systems:

- Systems based on statistical models,
- Hidden Markov models

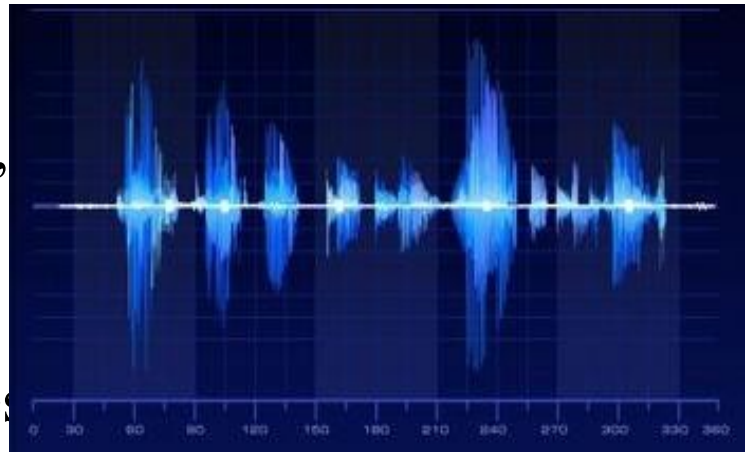

Multi-user speech recognition

- Voice command/voice activated devices
  - No training – works for many users

Adaptive speech systems

- Adapt to the user (training)
- continuous speech

Speech recognition powered devices:

Space exploration

Autonomous rovers,
intelligent probes

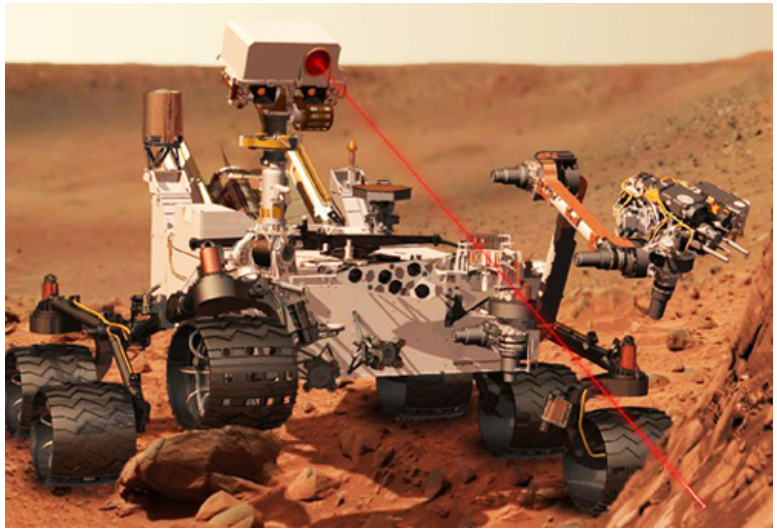

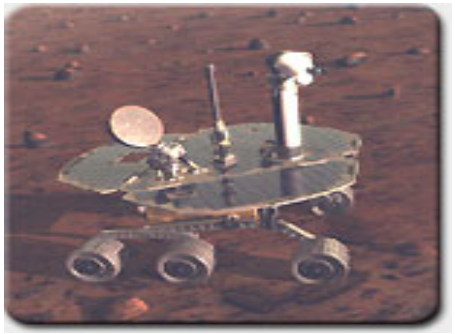

Analysis of sky
Survey data

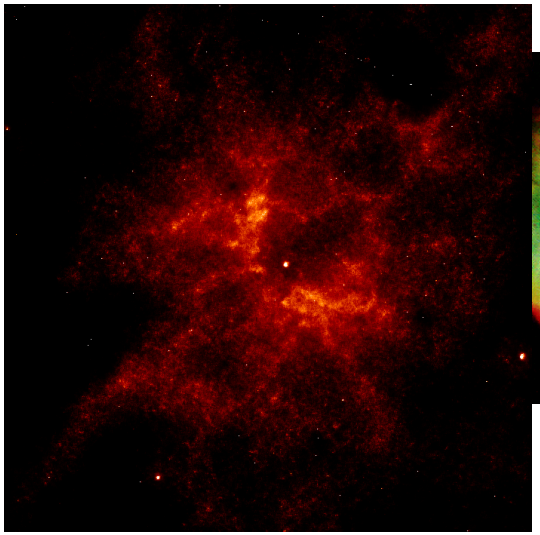

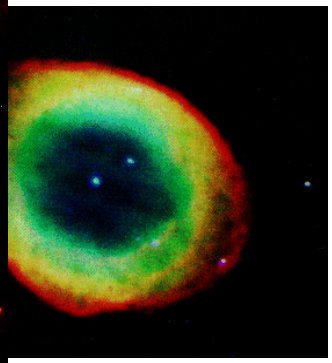

AI applications: Medicine

- **Medical diagnosis**
- **Patient Monitoring and Alerting:**
  - Decision support
- **Medical imaging**
  - Classification of body structures and visualization
- **Robotic surgeries**

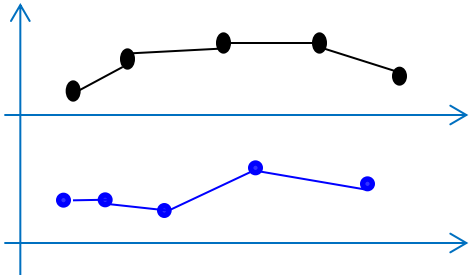

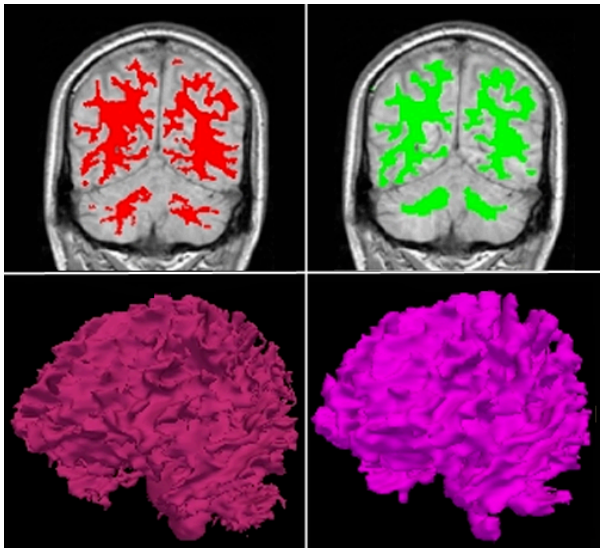

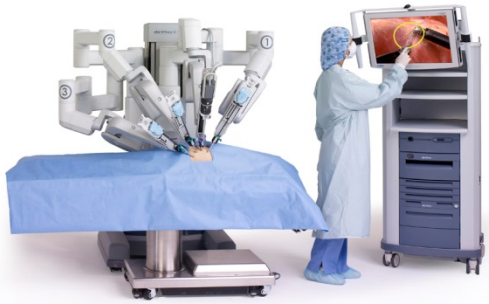

AI applications: Bioinformatics

- **Omics: Genomics, Proteomics, Metalobomics**
  - RNA sequence analysis
  - Prediction of gene regions on DNA
  - Analysis gene expressions and proteomic profiles: find genes, proteins (peptides) that characterize a specific disease
  - Regulatory networks

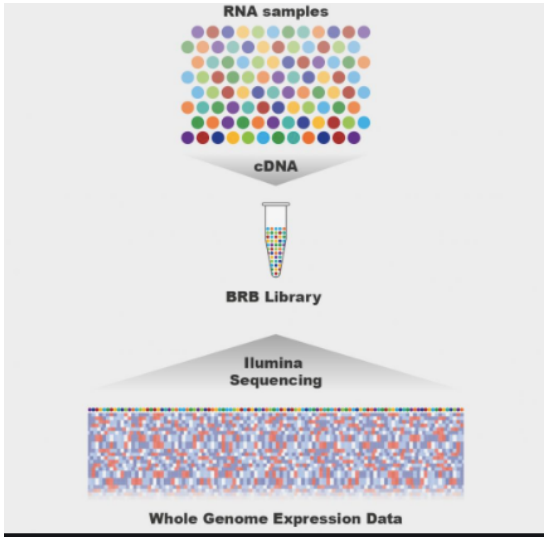

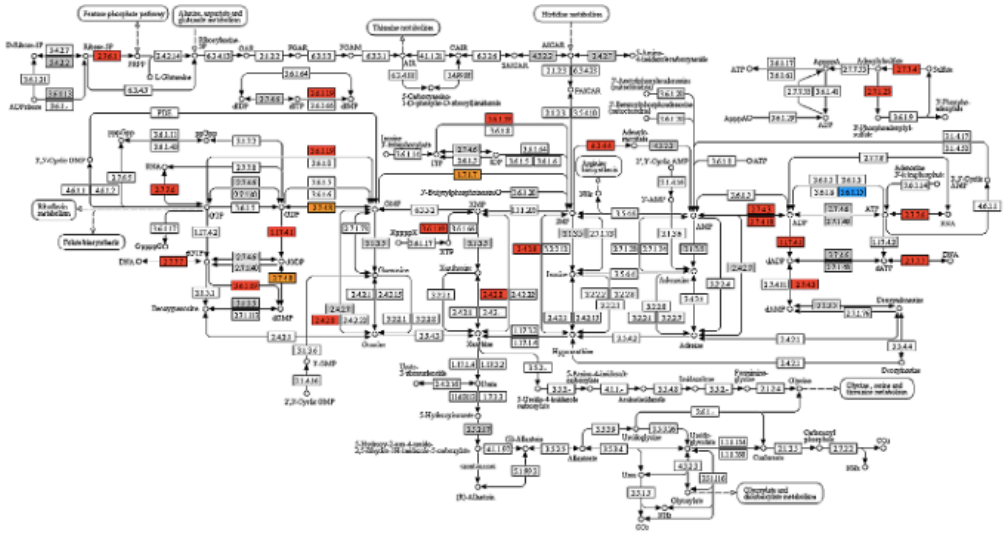

AI applications: Transportation

Autonomous vehicle control:

- ALVINN (CMU, Pomerleau 1993)
- Series of DARPA challenges (http://www.darpa.mil/grandchallenge/)
  - 2004, 2005 Drive across Mojave desert
  - 2007 - DARPA Urban Challenge

Now:

- autonomous vehicles are close to reality:
  - Uber, Tesla, ArgoAI

Other applications

- **Traffic monitoring**
- **Navigation/route optimizations**

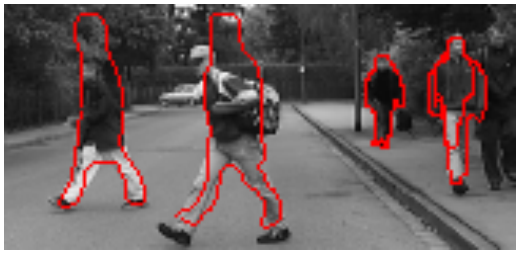

Image recognition/annotation

ImageNet

Game playing

- **Backgammon**
  - TD-backgammon: A reinforcement learning program that learned to play backgammon at the championship level
- **Chess**
  - **Deep blue (IBM) program**

    defeated Kasparov in 1997
- **Game of Go**
  - **AlphaGo (DeepMind) program**
  - defeated LeeSedol in 2016
- **Card playing programs**
  - **Bridge, Poker**

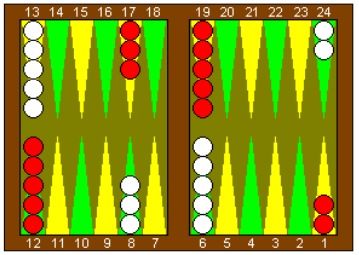

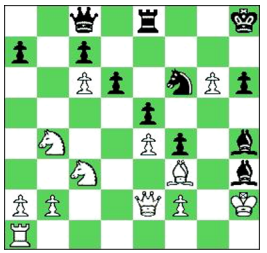

Natural language processing

Understanding/annotation/translation of free text

- **Document analysis:**
  - Automatic classification of articles
  - Content extraction/inference
  - Email SPAM detection

- **Language translation**

Knowledge extraction

Knowledge extraction

- **Knowledge Graph by Google** and its services to enhance its search engine's results with information gathered from a variety of sources

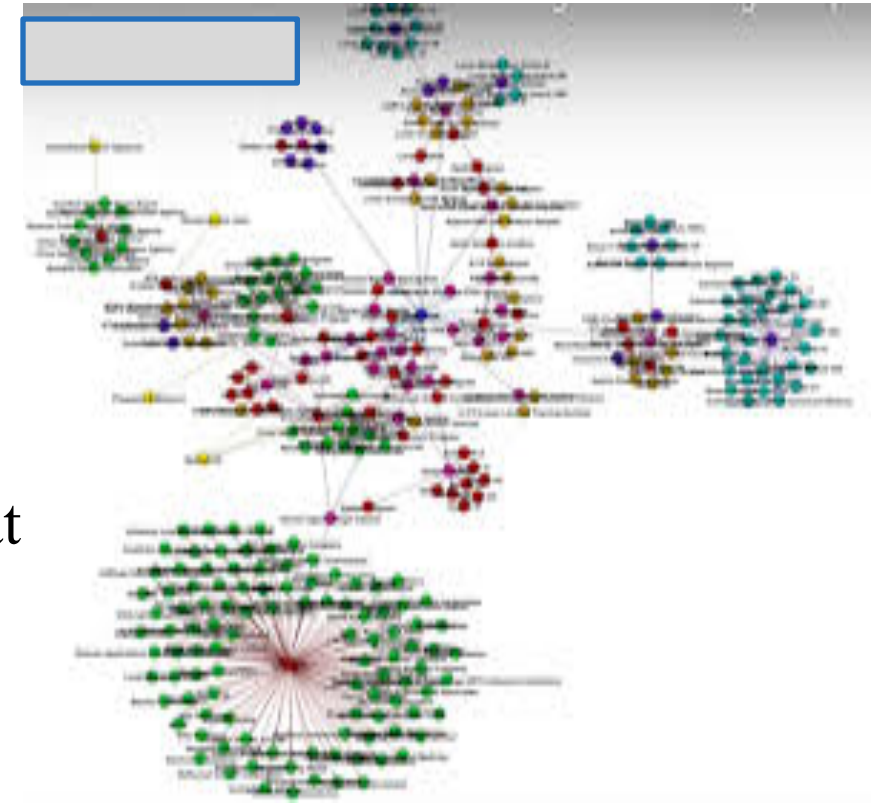

- **Kbpedia** (http://kbpedia.org): an open-source knowledge graph that combines seven leading public knowledge bases into an integrated and computable structure.

Robots

- **Robotic toys**
  - Sony's Aibo
- **Robotic tools**
  - Vacuum cleaners
- **Humanoid robots**
- **Robotics limbs**
- **Autonomous robotic systems**




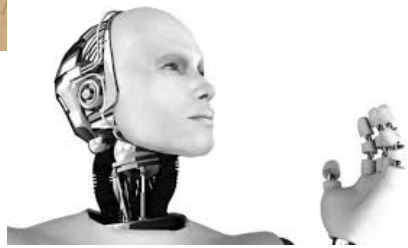
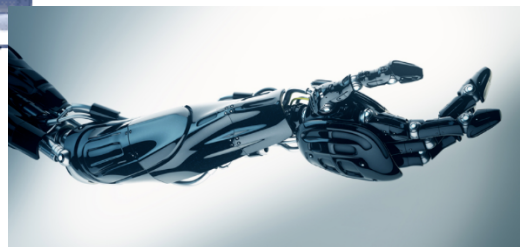



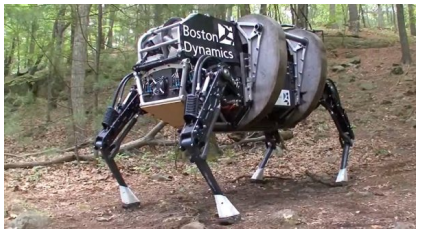

Other application areas

- **Handwriting analysis/ detection**
- **Human face detection**
- **Video stream annotation**
- **Object tracking**
- **Music composition, picture drawing**
- **…**

Topics

- **Problem solving and search**
  - Formulating a search problem, Search methods, Combinatorial and Parametric Optimization.
- **Knowledge representation, reasoning, planning**
  - Propositional and First order logic, inference
  - STRIPS, Partial-order planners
- **Modeling uncertainty, inference, decision making**
  - Probabilities, Bayesian belief networks, Inference in BBNs, Decision making in the presence of uncertainty.
- **Machine Learning**
  - Supervised learning, Unsupervised learning, Selected Machine learning topics