

Artificial Intelligence

Overview

Instructor

- **Dr. Changhe Yuan**
- **Email: changhe.yuan@qc.cuny.edu**
- **Phone: 3484**
- **Office: A342**
- **Office Hours: M 2:30-3:30pm**

Lectures

- **M 10:30am – 12:00pm**
- **SB D141**

Course website

- <http://url.cs.qc.cuny.edu/teaching/CSCI363Fall2017/syllabus.html>

Prerequisite

- **CSCI 323 Design and Analysis of Algorithms**
- **MATH 241 Introduction to Probability and Mathematical Statistics**
- **MATH 231/232 Linear Algebra**
- **All with a grade of C or better**
- **Programming skills important!**

Class Resources

- **Artificial Intelligence: Modern Approach (3rd Edition)** by Stuart J. Russell, Peter Norvig
- **Lecture notes**

Grading

- **Participation** **5%**
- **Assignments** **30%**
- **Mid-term exams** **45%**
- **Semester project** **20%**

What is Artificial Intelligence?

- Intelligence in daily life?

Recommendation systems
Nest - thermostat
Siri

What is Artificial Intelligence?


- **The science and engineering of making intelligent machines [John McCarthy].**
- **The scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines [AAAI].**
- **The study and design of intelligent agents, where an intelligent agent is a system that perceives its environment and takes actions that maximize its chances of success [Russell & Norvig].**

AI Topics

- **Agents**
- **Applications/Expert Systems**
- **Cognitive Science**
- **Education**
- **Ethical & Social**
- **Games & Puzzles**
- **History**
- **Interfaces**
- **Machine Learning**
- **Natural Language**
- **Philosophy**
- **Reasoning**
- **Representation**
- **Robots**
- **Science Fiction**
- **Speech**
- **Systems & Languages**
- **Turing Test**
- **Vision**

Search Overview

- On holiday in Romania; currently in Arad.
- Flight leaves tomorrow from Bucharest
- **Formulate goal:**
 - be in Bucharest
- **Formulate problem:**
 - **states:** various cities
 - **actions:** drive between cities
- **Find solution:**
 - sequence of cities, e.g.,
 - Arad, Sibiu, Fagaras, Bucharest



A* search


- **Main idea:** expand paths that seem most promising and avoid paths that are already expensive
- Use the following evaluation function to guide the search: $f(n) = g(n) + h(n)$
 - $g(n)$ = cost so far to reach n
 - $h(n)$ = estimated cost from n to goal
 - $f(n)$ = estimated total cost of path through n to goal

A* Path Finding Algorithm

- <https://www.youtube.com/watch?v=19h1g22hby8>
(1:06)

Game Search

- E.g., 2-ply game:



Endgame: Challenging the Masters

- Deep Blue was a chess-playing computer developed by IBM. On May 11, 1997, the machine won a six-game match by two wins to one with three draws against world champion Garry Kasparov.
- <http://www.youtube.com/watch?v=NJarxpYyoFI&feature=related> (6:07)

Alpha Go

- Google's 'AlphaGo' and the world's top ranked Go player go head-to-head in a battle to decide whether or not an AI can be programmed to win a game as complicated as Go.
- <https://www.youtube.com/watch?v=KsbQHfX6Pg>
(4:07)

Uncertainty Reasoning Overview

- **Uncertainty:** The lack of certainty, a state of having limited knowledge where it is impossible to exactly describe existing state or future outcome, more than one possible outcome.
- **Probability:** A measurement of uncertainty; A set of possible states or outcomes where probabilities are assigned to each possible state or outcome – this also includes the application of a probability density function to continuous variables [Wikipedia]

Example

- **Problem:** A screening test has a 90% chance of registering breast cancer if it exists, as well as a 20% chance of falsely registering cancer when it does not exist. About one in one hundred women requesting the screening test end up diagnosed with breast cancer.
- Ms. X has just been told that her screening test was positive. What is the probability that she has breast cancer?
- **Goal:** develop methods and algorithms for modeling and reasoning under uncertainty

Uncertainty Reasoning

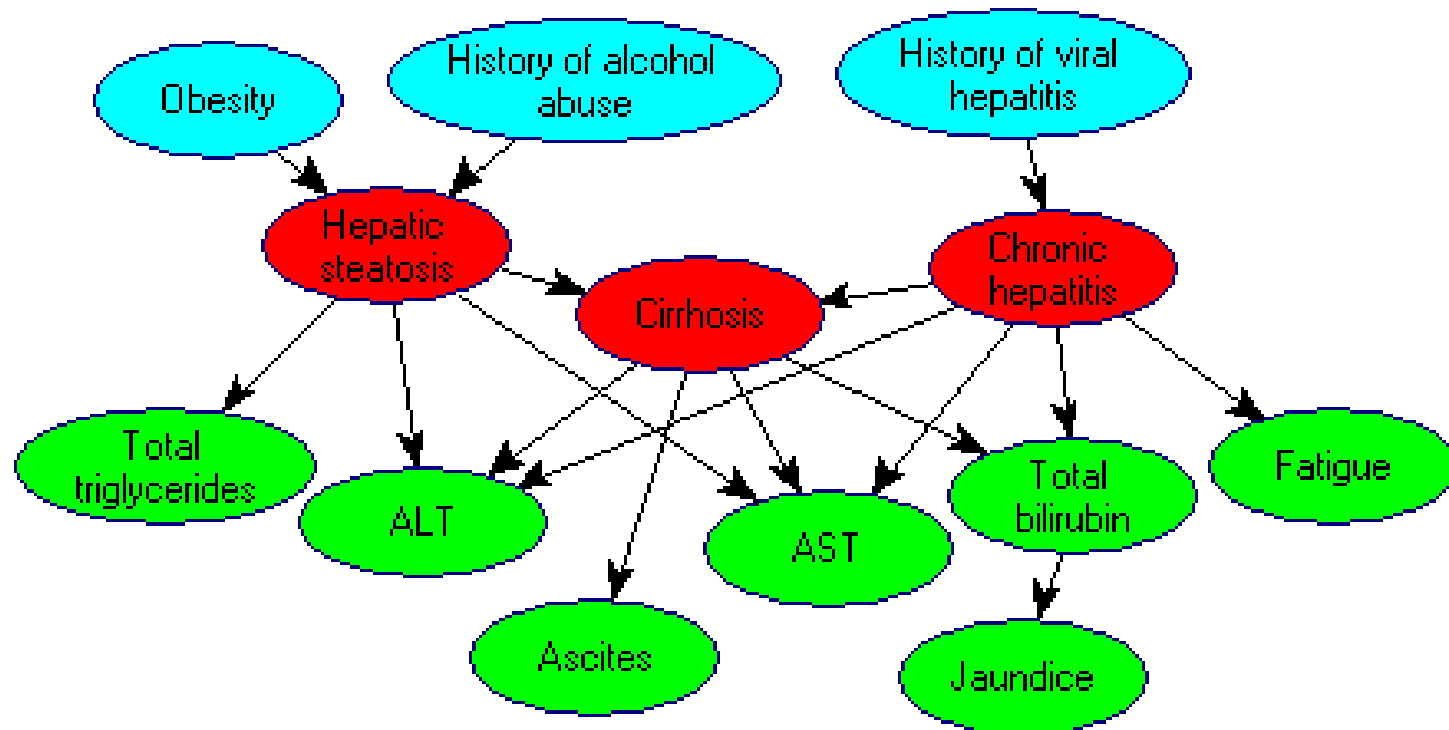


- **Bayes' rule:** $P(a | b) = \frac{P(b | a)P(a)}{P(b)}$

- Useful for assessing **diagnostic** probability from **causal** probability:

$$P(\text{cause} | \text{effect}) = \frac{P(\text{effect} | \text{cause})P(\text{cause})}{P(\text{effect})}$$

Bayesian networks




True Skill

- <http://www.youtube.com/watch?v=b7142mL9-f8>
(4:08) [Microsoft]

Hidden Markov Models

Two components:

- A Markov chain of hidden states H_1, \dots, H_n
- Observations X_1, \dots, X_n



Speech Recognition

- **Dragon:**
<http://www.youtube.com/watch?v=W3DhnpLIKCQ&feature=related> (2:03)
- **Microsoft:**
http://www.youtube.com/watch?v=2Y_Jp6PxsSQ
(1:34)

Machine Learning Overview

- **Machine learning is a scientific discipline that is concerned with the design and development of algorithms that allow computers to evolve behaviors based on empirical data, such as from sensor data or databases. [Wikipedia]**

Machine Learning

Problem: decide whether to wait for a table at a restaurant, based on the following attributes:

1. **Alternate:** is there an alternative restaurant nearby?
2. **Bar:** is there a comfortable bar area to wait in?
3. **Fri/Sat:** is today Friday or Saturday?
4. **Hungry:** are we hungry?
5. **Patrons:** number of people in the restaurant (None, Some, Full)
6. **Price:** price range (\$, \$\$, \$\$\$)
7. **Raining:** is it raining outside?
8. **Reservation:** have we made a reservation?
9. **Type:** kind of restaurant (French, Italian, Thai, Burger)
10. **WaitEstimate:** estimated waiting time (0-10, 10-30, 30-60, >60)

Goal: Build a machine learning model based on your prior behavior (**data**) to predict your future decision

Learning

- Simplest form: learn a function from examples

f is the **target function**

An **example** is a pair $(x, f(x))$

Problem: find a **hypothesis** h

such that $h \approx f$

given a **training set** of examples

Machine Learning Problem Types

- **Supervised learning** generates a function that maps inputs to desired outputs.
 - For example, in a **classification** problem, the learner approximates a function mapping a vector into classes by looking at input-output examples of the function.
- **Unsupervised learning** models a set of inputs
 - For example, **clustering** is the assignment of a set of observations into subsets (called *clusters*) so that observations in the same cluster are similar in some sense.
- **Semi-supervised learning** combines both labeled and unlabeled examples to generate an appropriate function or classifier.
- **Reinforcement learning** learns how to act given an observation of the world. Every action has some impact in the environment, and the environment provides feedback in the form of rewards that guides the learning algorithm. [Wikipedia]

Classification

- **Mind Reading**

http://www.youtube.com/watch?v=jq_nAm4rOu0
(1:36)

Clustering

The screenshot shows the Pandora internet radio website. At the top, the Pandora logo is displayed with the tagline "internet radio". To the right of the logo, the user's email "charlieyuan@yah..." is shown, along with links for "account", "sign out", and "upgrade". Below the header is a navigation bar with buttons for "Your Profile", "About the Music", "Mobile", and "Help".

The main content area features a music player interface. On the left, there is a "Create a New Station..." button and a list of "Your Stations" including "Smooth Jazz Radio", "Today's Adult Hits ...", "Club / Dance Radio", "80s Pop Radio", "Today's Country Ra...", and "Love Songs Radio" (which is highlighted). Below the stations is a "QuickMix" button. On the right, the music player shows the song "Heaven" by Bryan Adams from the album "So Far So Go...". The album cover art is visible, and there are "share", "play/pause", and "volume" controls at the top of the player. A "buy" button is also present.

Below the music player, a bio for Bryan Adams is displayed. It includes a small image of Bryan Adams playing a guitar, with the text "BRYAN ADAMS 11" overlaid. The bio text reads: "Bryan Adams negotiated the shifting tides of the '80s so well that it never seemed like he was changing his music to fit the times. A veteran of the '70s studio arena rock game, Adams struck out on his own in the early '80s, turning into a star in his native Canada and making headway with his 1983 album, Cuts...". Below the bio is a link for "full bio..." and three small dots indicating more options.

Natural language processing

- **Natural language processing (NLP) is a field of computer science, artificial intelligence, and computational linguistics concerned with the interactions between computers and human (natural) languages. [Wikipedia]**
- **Many challenges in NLP involve natural language understanding, that is, enabling computers to derive meaning from human or natural language input, and others involve natural language generation.**

Information Retrieval

- **IBM Watson**

<https://www.youtube.com/watch?v=DywO4zksfXw> (6:42)

<http://www.youtube.com/watch?v=Y2wQQ-xSE4s&feature=related>
(1:50)

- **Google search engine**

<http://www.youtube.com/watch?v=BNHR6IQJGZs&feature=related>
(3:15)

Concluding remarks

- **Artificial intelligence is a fun topic**
 - A lot of practical applications
 - Highly interdisciplinary
 - Lots of AI jobs
- **Also, they are not as difficult as you may have thought!**