


CS 6511: Artificial Intelligence

Introduction

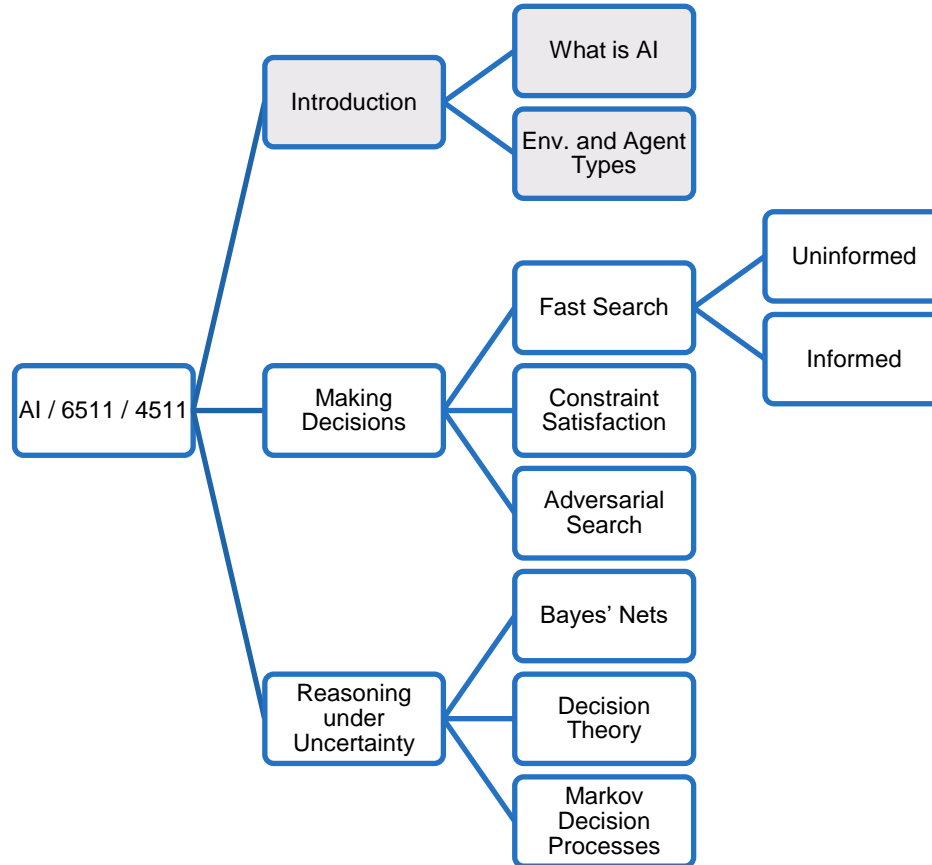


Instructor: Amrinder Arora
The George Washington University

Course Outline

Other AI Topics (Not included in this class):


- Machine Learning
- Big Data
- AR/VR
- Speech Synthesis
- *Robotics*
- *NLP*



Today

- What is artificial intelligence?
- What can AI do?
- What is this course?

Different Perspectives of AI



VR/Sensory

Neural Networks

Machine Learning

NLP

Algorithms

Big Data

Rational Decisions

We use the term **rational** in a very specific, technical way:


- Rational: maximally achieving pre-defined goals
- Rationality only concerns what decisions are made (not the thought process behind them)
- Goals are expressed in terms of the **utility** of outcomes
- Being rational means **maximizing your expected utility**

A better title (but perhaps not so marketable) for this course would be:


Rational Agent Design

Designing Rational Agents


- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- Characteristics of the **percepts, environment, and action space** dictate techniques for selecting rational actions
- **This course** is about:
 - General AI techniques for a variety of problem types
 - Learning to recognize when and how a new problem can be solved with an existing technique



Pac-Man™ as an Agent



* Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes



Agent and Task Environment Interaction

Letter	Attribute	Description
P	Performance Measures	What we are trying to achieve
E	Environment	How our environment behaves
A	Actuators	What we can use to interact with and influence our task environment
S	Sensors	What we can use to sense and observe the environment

Characterizing Task Environment Types

THERE ARE 7 KEY ATTRIBUTES

Environment Type: Observability

- Observable?
- Unobservable?
- Partially observable?

Environment Type: Single vs. Multi Agent

- Single Agent?
- Or multiple?

- Competitive, or collaborative?

Environment Type: Deterministic vs. Stochastic

- Deterministic: Effect defined by action
- Stochastic: Effect depends on chance (in addition to the action)

Environment Type: Episodic vs. Sequential

- Episodic: Atomic incidents (episodes)
- Sequential: Previous decision affects next decisions

Environment Type: Static vs. Dynamic

- Static: Does not change while agent is responding
- Dynamic: Environment keeps changing
- Semi-dynamic: Environment itself is not changing, but the agent's reward is changing

Environment Type: Discrete vs. Continuous

- Discrete: Discrete set of actions and observations
- Continuous: Continuous set of actions/observations

Environment Type: Known vs. Unknown

- Known: We know what the actions do (the laws of physics are known)
- Unknown: We don't know what affect the actions will have
- (This is different from observable or unobservable)

Characterizing Task Environment: Quick Summary

1. Fully observable or Partially? (Or unobservable)
2. Number of Agents – 1 or many?
3. Deterministic, or Stochastic (Probabilistic)
4. Episodic (independent tasks) or Sequential (look ahead required)
5. Static vs. Dynamic (Observe once, or again and again?)
6. Discrete or Continuous
7. Known vs. Unknown

Looking into the Agents

AGENT TYPES

Agent Types

- Simple reflex agent: Current observation + Predefined rules
- Model based reflex agent: Have an internal model to interpret the environment
- Goal based: Choose the action that can help achieve the goal
- Utility based: Use the concept of utility, which they measure and choose the action on the basis of the utility (represents the agent's assessment of the performance measure)
- Learning agents: To deal with unknown environments

Simple Reflex Agent

- Great for full observable environment

Model Based Reflex Agent

- Great for full observable environment and can work with partially observable environments

Goal Based Agent

Utility Based Agent

- Great for full observable environment

Learning Agents

- To deal with unknown environments

States

- Atomic: Each state is treated as a black box
- Factored: States are defined by a set of features
- Structured: States are expressed in form of objects and relations between them. (Also called facts and/or relations.)

Quick Summary of PEAS

- PEAS
 - Performance Measures
 - Environment
 - Actuators
 - Sensors

Wrap Up

- Concepts
 - “Agent”
 - “Utility”
 - “Rational Agent”
 - Environment
- This class = Rational Agent Design
- Agent Types
 - Agent States
- Next Topics
 - Search Space
 - DFS/BFS as tools to explore search space