## Intelligent Agents (Ch. 2)

- examples of agents
  - webbots, ticket purchasing, electronic assistant, *Siri*, news filtering, *autonomous vehicles*, printer/copier monitor, Robocup soccer, NPCs in Quake, Halo, Call of Duty...
- agents are a unifying theme for Al
  - use search and knowledge, planning, learning...
  - focus on <u>decision-making</u>
  - must deal with uncertainty, other actors in environment

## Characteristics of Agents

- essential characteristics
  - agents are <u>situated</u>: can sense and manipulate an environment that changes over time
  - 2. agents are goal-oriented
  - 3. agents are <u>autonomous</u>
- other common (but not universal) aspects of agents:
  - 1. adaptive (learns from experience)
  - 2. optimizing (rational)
  - 3. social (i.e. cooperative, teamwork, coordination)
  - 4. life-like (e.g. in games, interactions with humans)



- policy mapping of states (or histories) to actions
  - π(s)=a
  - $\pi(s_1,...s_t)=a_t$
- Performance measures:
  - utility function, rewards, costs, goals
  - mapping of *states* (or *states*×*actions*) into R,
    S |-> ℜ or S,A |-> ℜ

## Rational behavior (rationality)

- rationality: "for each possible percept sequence, a rational agent should <u>select an action</u> <u>that is expected to maximize</u> <u>its performance measure</u>, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has"
- colloquially, being rational means "doing the right thing"



## Rationality

- <u>select an action that is expected to maximize its</u> <u>performance measure</u>
- consider a set of possible outcomes, {*o<sub>i</sub>*}
- select the action *i* that leads to the outcome with the highest payoff/reward, *argmax<sub>i</sub> payoff(o<sub>i</sub>)*
- in uncertain (stochastic) environments, if an action could lead to several outcome, <u>take the average</u> <u>outcome</u>, weighted by probability

remember Expectiminimax?

7/27/2024

 $Expectiminimax(s) = \begin{bmatrix} u_1(s) & \text{if is a terminal node} \\ max{Expectiminimax(s') | s' \in \text{succ}(s)} & \text{if max node} \\ min{Expectiminimax(s') | s' \in \text{succ}(s)} & \text{if min node} \\ \sum_{s' \in \text{succ}(s)} P(s') \cdot Expectiminimax(s') & \text{if chance node} \end{bmatrix}$ 

take action that leads to highest average mm score over children

### Task Environments

• The architecture or design of an agent is strongly influenced by characteristics of the environment

Discrete	Continuous	
Static	Dynamic	(read the
Deterministic	Stochastic	definitions and examples in
Episodic	Sequential	the textbook)
Fully Observable	Partially Observable	
Single-Agent	Multi-Agent	

- Reactive/Reflex Agents
  - stimulus-response
  - condition-action <u>lookup table</u>
  - efficient
  - goals are implicit
  - sense-decide-act loop
  - OODA loop (observe-orient-decide-act)





State	Action
S1	A1
S2	A2
S3	A3
•••	





return action

(This is specifically a Rule-based Reflex Agent. Could also use a Lookup Table to selection *action*.)

- Rule-based Reactive Agents
  - condition-action trigger rules
    - if carInFrontIsBraking then InitiateBraking
  - more compact than table
  - issue: how to choose which rule to fire?
    - must prioritize rules, if more than one rule can fire
- implementations
  - if-then-else cascades
  - CLIPS; JESS Java Expert System
  - Subsumption Architecture (Rodney Brooks, MIT)
    - hierarchical design <u>behaviors</u> in layers
    - e.g. obstacle avoidance overrides moving toward goal

- Model-based Agents
  - use local variables to represent and remember the state of the world and infer unobservable aspects



 $rule \leftarrow RULE-MATCH (state, rules)$ 

action ← RULE-ACTION [rule]

*state* ← <u>UPDATE-STATE</u> (*state, action*) // predict, remember <sup>7/27/2024</sup> **return** *action* 

- Knowledge-based Agents
  - knowledge base containing logical rules for:
    - inferring unobservable aspects of state
    - inferring effects of actions
    - inferring what is likely to happen
  - use inference algorithm to decide what to do next, given state and goals
    - use forward/backward chaining, natural deduction, resolution...
    - prove: Percepts $\cup$ KB $\cup$ Goals |= do( $\alpha_i$ ) for some action  $\alpha_i$

- Proactive agents
  - (in contrast to Reactive agents)
  - more than just stimulus-response reason about what is going to happen
  - proactive agents don't just take actions for immediate payoff - they think several steps ahead, because achieving some goals/rewards takes multiple actions
  - typically, proactive agents are knowledge-based, because they have to do more complex reasoning
  - we will talk about 2 specific architectures for designing proactive agents: Planning Agents, and Utility-based Agents (MDPs and RL)...

- Goal-based Agents (Planning Agents)
  - search for plan (sequence of actions) that will transform  $S_{init}$  into  $S_{goal}$
  - state-space search (forward from S<sub>init</sub>, e.g. using A\*)
  - goal-regression (backward from S<sub>goal</sub>)
    - reason about effects of actions
  - SATplan, GraphPlan,
    PartialOrderPlan (POP)...



#### Goal-based agents



note: plans must be maintained on an <u>agenda</u> and carried out over time - these are <u>intentions</u>

- Utility-based Agents
  - utility function: maps states to real values, quantifies "goodness" of states,  $u(s) \rightarrow \Re$
  - agents select actions to maximize utility
    - sometimes payoffs are immediate (think "reactive")
    - othertimes payoffs are delayed:
      - <u>Sequential Decision Problems</u>
      - maximize long-term reward

#### Markov Decision Problems (MDPs)

- transition function:  $T(s,a) \rightarrow S$ 
  - outcomes of actions
  - could be probabilistic (distribution over successors states)
- reward/cost function:  $R(s,a) \rightarrow \Re$
- "plans" are encoded in *policies* 
  - mappings from states to actions:  $\pi: S \rightarrow A$
  - Markov property: probabilities only depend on current state
- the goal: maximize reward over time
  - long-term discounted reward  $\sum_{t=0}^{\infty} \gamma^t R_{a_t}(s_t, s_{t+1})$



7/27/2024

## Multi-Agent Systems

- Collaborative Agents
  - competition (Minimax) vs. collaboration
  - collaboration: is there a way agents can work together so they mutually benefit?
  - "open" agent environment: assume all agents are self-interested (have their own utility function)

# Market-based methods for Multi-Agent Systems

- mechanisms to incentivize collaboration
  - <u>contract networks</u> agents make bids to do tasks for each other, negotiate price, make commitments
  - <u>auctions</u> agents bid on resources
    - first-price, second-price, open vs sealed bid, asc vs descending
    - strategy to maximize utility?
    - bidding on *combinations* of resources is more complicated
  - <u>consensus</u> algorithms voting (weight choices by utility)
- important issues:
  - do these mechanisms <u>incentivize agents to be rational</u> and bid their true values?; are they free of exploitation/manipulation?
  - efficiency: do these mechanisms <u>maximize social benefit</u>? (sum of utility of outcomes over all agents)

## Methods for Collaborative Agents

- Agent Teamwork
  - shared goals, joint intentions
    - assume teammates are not just self-interested
    - teammates can compensate for each other if a team goal is at risk
  - well-defined roles, responsibilities
  - communication among teammates is key
- BDI *modal logic* for representing **B**eliefs, **D**esires (goals), and Intentions (actions) of other agents
  - Bel(self,empty(ammo))
     ∧Bel(teammate,¬empty(ammo))
     ∧Goal(teammate,shoot(gun))
     → Tell(teammate,empty(ammo))
  - *intentions* are actions that we select and commit to, which means we plan to do them (or keep trying till we succeed)
  - modal operators go beyond FOL: Bel(<agt>,<sentence>)

7/27/2024