Lecture 8: Reasoning Under Uncertainty

ICS 171, Summer 2000

Outline

- **Autonomous Agents**
  - need to be able to handle uncertainty
- **Probability as a tool for uncertainty**
  - basic principles
- **Decision-Marking and Uncertainty**
  - optimal decision-making
    - principle of maximum expected utility

Autonomous Agents

- **Consider an agent which is reasoning, planning, making decisions**

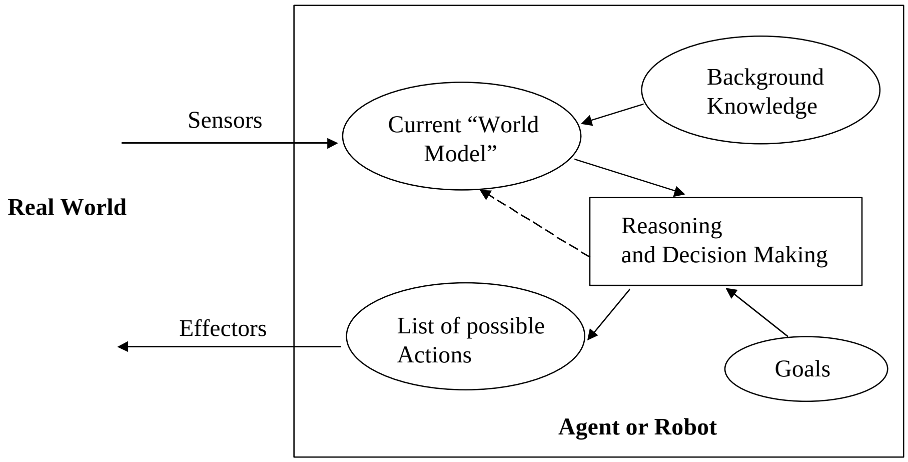

How an Agent Operates

- **Basic Cycle**
  - use sensors to sense the environment
  - update the world model
  - reason about the world (infer new facts)
  - update plan on how to reach goal
  - make decision on next action
  - use effectors to implement action

- **Basic cycle is repeated until goal is reached**

Example of an Autonomous Agent

- **A robot which drives a vehicle on the freeway**

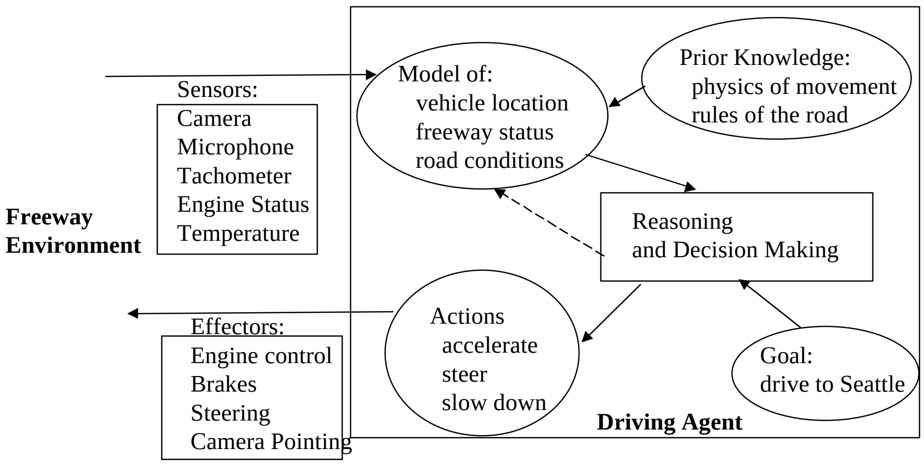

The Agent's World Model

- **World Model = internal representation of the external world**
  - combines
    - background knowledge
    - current inputs
- **Necessarily, the world model is a simplification**
  - e.g. in driving we cannot represent every detail
    - every pebble on the road?
    - details of every person in every other vehicle in sight?
- **A useful model is the State Space model - We used it in Search**
  - represent the world as a set of discrete states
    - e.g., variables = {Rainy, Windy, Temperature,.....}
      - state = {rain=T, windy=T, Temperature = cold, ...}
- **An agent must**
  - 1. figure out what state the world is in
  - 2. figure out how to get from the current state to the goal

Uncertainty in the World Model

- **The agent can never be completely certain about the external world state. i.e., there is ambiguity and uncertainty**
- **Why?**
  - sensors have limited precision
    - e.g., camera has only so many pixels to capture an image
  - sensors have limited accuracy
    - e.g., tachometer's estimate of velocity is approximate
  - there are hidden variables that sensors can't "see"
    - e.g., large truck behind vehicle
    - e.g., storm clouds approaching
  - the future is unknown, uncertain: i.e., we cannot foresee all possible future events which may happen
- **In general, our brain functions this way too:**
  - we have a **limited perception** of the real-world

Rules and Uncertainty

- Say we have a rule

  if toothache then problem = cavity

- But not all patients have toothaches because of cavities (although perhaps most do)

  So we could set up rules like

  if toothache and not(gum disease) and not(filling) and ......
  then problem = cavity

- This gets very complicated! a better method would be to say

  if toothache then problem = cavity with probability 0.8

  or p(cavity | toothache) = 0.8

Example of Uncertainty

- **Say we have a camera and vision system which can estimate the curvature of the road ahead:**

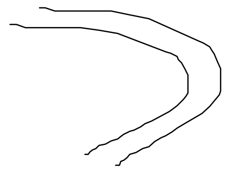

- **There is uncertainty about which way the road is curving**
  - limited pixel resolution, noise in image
  - algorithm for "road detection" is not perfect
- **We can represent this uncertainty with a simple probability model**
- **Probability of an event = a measure of agent's belief in the event given the evidence E**
  - e.g.,
    - p(road curves to left | E) = 0.6
    - p(road goes straight | E) = 0.3
    - p(road curves to right | E) = 0.1

Variables: Notation

- **Consider a variable, e.g., A,**
  - (usually in capitals)
  - assume A is discrete-valued
  - takes values in a **domain**
    - e.g. binary: domain = {true, false}
    - e.g., multivalued: domain = {clear, partly cloudy, all cloud}
  - variable takes one and only one value at a given time
    - i.e., values are mutually exclusive and exhaustive
  - The statement "A takes value a", or "A = a", is an event or proposition
    - this proposition can be true or false in real-world
  - An agent's uncertainty is represented by p(A = a)
    - this is the agent's belief that variable A takes value a (i.e., "world is in state a"), given no other information relating to A
  - Basic property: Σ p(a) = p(A=a1) + p(A=a2) + ... p(A=ak) = 1

Variables and Probability Distributions

- **Example:** Variable = Sky
  - takes values in {clear, partly cloudy, all cloud}
  - probabilities are p(clear), p(partly cloudy), p(all cloud), e.g:
    - » p(Sky = clear) = 0.6
    - » p(Sky = partly cloudy) = 0.3
    - » p(Sky = all cloud) = 0.1

- **Notation:**
  - we may use p(clear) as shorthand for p(Sky = clear)
  - If S is a variable, with taking values in {s1, s2, ... sk}
    - then s represents some value for S
    - i.e., if S is Sky, then p(s) means any of p(clear), p(all cloud), etc: this is a notational convenience
  - Probability distribution on S taking values in {s1, s2, ... sk}:
    - P(S) = the set of values {p(s1), p(s2), ... p(sk)}
    - If S takes k values, then P(S) is a set of k probabilities

Conjunctions of Events, Joint Variables

- **A = a is an event**
  - A is a variable, a is some value for A
- **We can generalize to speak of conjunctions of events**
  - A = a AND B = b (like propositional logic)
- **We can assign probabilities to these conjunctions**
  - p(A = a AND B = b)
  - This is called a *joint probability* on the event A=a AND B=b
  - Notation watch!
    - convention: use p(a, b) as shorthand for p(A = a AND B = b)
- **Joint Distributions**
  - let A, B, C be variables each taking k values
  - then P(A, B, C) is the joint distribution for A, B, and C
  - how many values are there in this joint distribution?

Summary of Notation and Conventions

- **Capitals denote variables, e.g., A, B, C...**
  - these are attributes in our world model
- **Lower-case denotes values of variables, e.g., a, b, c,**
  - these are possible states of the world
- **The statement A=a is an event (equivalent to a proposition)**
  - true or false in the real-world
- **We can generalize to conjunctions of events**
  - e.g., A=a, B=b, C = c (shorthand for AND(A=A, B=b, C=c)
- **lower case "p" denotes a single probability for a particular event**
  - e.g., p(A = a, B=b)
- **upper case "P" denotes a distribution for the full set of possible events (all possible variable-value pairs)**
  - e.g. P(A, B) = {p(a1,b1), p(a2,b1), p(a1,b2), p(a2,b2)}
  - often represented as a table

ICS-171:Lecture 8: 13

Axioms of Probability

- **What are the rules which govern the assignment of probabilities to events?**
- **Basic Axioms of Probability**
  - 1. 0 <= p(a) <= 1
    - probabilities are between 0 and 1
  - 2. p(T) = 1, p(F) = 0
    - if we believe something is absolutely true we give it probability 1
  - 3. p(not(a)) = 1 - p(a)
    - our belief in not(a) must be "one minus our belief in a"
  - 4. p(a or b) = p(a) + p(b) - p(a and b)
    - probability of 2 states is their sum minus their "intersection"
    - e.g., consider a = sunny and b = breezy
- **One can show that these rules are necessary if an agent is to behave rationally**

ICS-171:Lecture 8: 14

Conditional Probability

- **Define p(a | e) as the probability of a being true if we know that e is true, i.e., our belief in a is conditioned on e being true**
  - the symbol "|" is taken to mean that the event on the left is conditioned on the event on the right being true.
- **Conditional probabilities behave exactly like standard probabilities**
  - $0 <= p(a|e) <= 1$
    - conditional probabilities are between 0 and 1
  - $p(a_1|e) + p(a_2|e) + \ldots\ldots\ p(a_k|e) = 1$
    - i.e., conditional probabilities sum to 1.
    - here a_2, etc., are just specific values of A
  - we can have p(conjunction of events | e), e.g.,
    - p(a and b and c | e) is the agent's belief in the sentence "a and b and c" on the left conditioned on e being true.
- **Conditional probabilities are just a more general version of "standard" probabilities**

Actions and States

- **Let S be a discrete-valued variable**
  - i.e., V takes values in the set of states $\{s_1, \ldots. s_k\}$
    - values are mutually exclusive and exhaustive
  - represents a state of the world, e.g., road = {dry, wet}
- **Assume there exists a set of Possible Actions**
  - e.g., in driving
    - A = set of actions
      = {steer left, steer straight, steer right}
- **The Decision Problem**
  - what is the optimal action to take given our model of the world?
- **Rational Agent**
  - will want to take the best action given information about the states
  - e.g., given p(road straight), etc, decide on how to steer

Action-State Utility Matrix

- **u (A, s) = utility of action A when the world really is in state s**
  - u(A, s) = the utility to the agent which would result from that action if the world were really in state s
    - utility is usually measured in units of "negative cost"
    - e.g., u(write check for $1000, balance = $50) = -$10
  - important: the agent reasons "hypothetically" since it never really knows the state of the world exactly
  - for a set of actions A and a set of states S this gives **a utility matrix**

- **Example**
  - S = state_of_road, takes values in {l, s, r}
    - this the variable whose value is uncertain (=> probabilities)
  - A = actions, takes values in {SL, SS, SR, Halt}

  - u (SL, l) = 0
  - u (SR, l) = $20k
  - u (SL, r) = $1,000k

Example of an Action-State Utility Matrix

	STATE (Curvature of Road)		
ACTION	**Left**	**Straight**	**Right**
Steer Left (SL)	**0**	**-1000**	**-1000**
Steer Straight (SS)	**-20**	**0**	**-1000**
Steer Right (SR)	**-20**	**-20**	**0**
Halt (H)	**-500**	**-500**	**-500**

Note: you can think of utility as the negative cost for the agent
=> maximing utility is equivalent to minimizing cost

Expected Utilities

- **How can the driving agent choose the best action given probabilities about the state of the world?**
  - e.g., say p(l) = 0.2, p(s) = 0.7, p(r) = 0.1
- **Say we take action 1, i.e., steer left**
- **We can define the Expected Utility of this action by averaging over all possible states of the world, i.e.,**
  - expected utility (EU) =
    sum over states of {utility(Action, state) x p(state)}

EU(Steer Left) = u(SL | l) x p(l)
+ u (SL | s) x p(s)
+ u (SL | r) x p(r)

= - 0 x 0.2 + (- 20) x 0.7 + (- 1000) x 0.1
= - 114

Optimal Decision Making

- **Optimal Decision Making:**
  - **choose the action with maximum expected utility (MEU)**
  - **procedure**
    - calculate the expected utility for each action
    - choose among the actions which has maximum expected utility
- **The maximum expected utility strategy is the optimal strategy for an agent who must make decisions where there is uncertainty about the state of the world**
  - assumes that the probabilities are accurate
  - assumes that the utilities are accurate
  - is a “greedy” strategy: only optimizes 1-step ahead
- **Read Chapter 16, pages 471 to 479 for more background**

Example of Optimal Decision-Making

- **Use action-state utility matrix from before**
- **State Probabilities are p(l) = 0.2, p(s) = 0.7, p(r) = 0.1**

Expected_Utility(Steer Left) = 0 x 0.2 + (- 20) x 0.7 -1000 x 0.1
= -114

Expected_Utility(Steer Straight) = - 20 x 0.2 + 0 x 0.7 -1000 x 0.1
= -104

Expected_Utility(Steer Right) = - 20 x 0.2 + - 20 x 0.7 + 0 x 0.1
= -18

Expected_Utility(Halt) = - 500 x 0.2 + (- 500) x 0.7 + (-500) x 0.1
= -500

- **Maximum Utility Action = "Steer Right"**
  - note that this is the *least likely* state of the world!
  - but is the one which has maximum expected utility, i.e., it is the strategy which on average will minimize cost

Summary

- **Autonomous agents are involved in a cycle of**
  - sensing
  - estimating the state of the world
  - reasoning, planning
  - making decisions
  - taking actions
- **Probability allows the agent to represent uncertainty about the world**
  - agent can assign probabilities to states
  - agent can assign utilities to action-state pairs
- **Optimal Decision-Making = Maximum Expected Utility (MEU)Action**