Module 5 Notes

ACTING UNDER UNCERTAINTY

Agents often encounter uncertainty due to **partial observability or nondeterminism**. Strategies like maintaining a **belief state** or **generating contingency plans** help handle this uncertainty. However, such approaches have limitations:

1. **Logical agents** must consider all possible explanations for observations, leading to overly complex belief-state representations.
2. **Contingency plans** can become excessively large and must account for improbable events.
3. Sometimes, there's no guaranteed plan for achieving a goal, yet the agent must act.

For instance, an automated taxi aiming to get a passenger to the airport on time may face uncertainty regarding various factors like traffic or accidents. Despite uncertainty, the agent must make decisions that maximize its expected performance, considering its knowledge about the environment. The right decision depends on balancing goals and their likelihood of achievement.

This highlights the importance of **rational decision-making under uncertainty**, where agents choose actions that are expected to optimize their performance measures relative to their knowledge about the environment.

Summarizing Uncertainty

Consider the scenario of diagnosing a dental patient's toothache, which exemplifies uncertain reasoning. Diagnosis, whether in medicine or other fields, typically involves **uncertainty**. Attempting to devise rules for dental diagnosis using propositional logic reveals the breakdown of the logical approach.

For instance, a simple rule like

- **"Toothache ⇒ Cavity"**

is inadequate because not all patients with toothaches have cavities; some may have gum disease, abscesses, or other issues.

- **Toothache ⇒ Cavity ∨ GumProblem ∨ Abscess ...**

To rectify this, one would need to enumerate an extensive list of possible problems, making the rule impractical.

Even transforming the rule into a causal one like:
"Cavity ⇒ Toothache"
proves inaccurate because not all cavities cause pain.

The limitations of using logic for medical diagnosis stem from three main factors:

1. **Laziness:** It's cumbersome to list all antecedents or consequents required for a rule, and using such exhaustive rules is challenging.
2. **Theoretical ignorance**: Medical science lacks a complete theory for the domain.
3. **Practical ignorance**: Even with knowledge of rules, uncertainty remains due to incomplete testing or unavailable tests.

In domains like medicine, logical consequences between phenomena aren't straightforward. Instead, the agent's knowledge provides only a degree of belief in relevant statements. Probability theory addresses uncertainty arising from laziness and ignorance, thus resolving the qualification problem.

Probability allows summarizing uncertainty, indicating, for example, an 80% chance that a patient with a **toothache has a cavity**. This belief could derive from statistical data or general dental knowledge.

Probability statements reflect knowledge states rather than **absolute truths** about the **real world**.

For instance, stating "**The probability of a cavity, given a toothache, is 0.8**" reflects a certain knowledge state. As new information emerges, such as a history of **gum disease**, the probability statement may change accordingly without contradicting previous assessments. Each statement represents a distinct assertion about a specific knowledge state.

Uncertainty and rational decisions

Let's reconsider the A90 plan for reaching the airport. Suppose it gives us a 97% chance of catching our flight. Does this necessarily make it the rational choice? Not always. There may be other plans, like A180, with higher probabilities of success. If avoiding missing the flight is crucial, it might be worth enduring a longer wait at the airport.

Now, what about A1440, a plan involving leaving home 24 hours in advance? In most cases, this isn't a good choice. Although it nearly guarantees arriving on time, it entails an intolerable wait, not to mention potentially unpleasant airport food.

In making such decisions, an agent must establish preferences among the potential outcomes of various plans. Each outcome encompasses factors like punctuality and airport wait time. We utilize utility theory to represent and reason with these preferences. Here, utility refers to the usefulness or desirability of a state to an agent.

Preferences, expressed through utilities, are combined with probabilities in the general theory of rational decisions called decision theory:

Decision theory=probability theory + utility theory

The fundamental idea of decision theory is that an agent is rational if and only if it chooses the action that yields the highest expected utility, averaged over all the possible outcomes of the action. This is known as the principle of **maximum expected utility (MEU).** Expected here refers to the statistical mean of outcomes, weighted by the probability of each outcome. This principle is a fundamental concept across various decision-making scenarios.

```
function DT-AGENT(percept) returns an action
  persistent: belief_state, probabilistic beliefs about the current state of the world
              action, the agent's action

  update belief_state based on action and percept
  calculate outcome probabilities for actions,
     given action descriptions and current belief_state
  select action with highest expected utility
     given probabilities of outcomes and utility information
  return action
```

Figure 13.1 A decision-theoretic agent that selects rational actions.

Function Name: DT-AGENT(percept)

Purpose: To make decisions in an uncertain environment based on observed precepts.

Persistent Variables:

- belief state: Probabilistic beliefs regarding the current state of the environment.

- action: Represents the action chosen by the agent.

Functionality:

- Updates the belief state based on the agent's action and the received percept.
- Calculates the probabilities of different outcomes for each potential action, considering action descriptions and the current belief state.
- Selects the action with the highest expected utility, taking into account the probabilities of outcomes and utility information associated with each outcome.

- **Returns the selected action**.

Figure 13.1 illustrates the structure of an agent using decision theory to select actions. While the agent's abstract structure resembles those described in previous chapters, the primary difference lies in the inclusion of probabilities alongside possibilities for world states in the agent's belief state. With this belief state, the agent can predict action outcomes probabilistically and select the action with the highest expected utility.

The function DT-AGENT(percept) is a decision-making function for an agent operating in an uncertain environment. Here's a breakdown of its components and how it works:

Inputs: The function takes a percept as input. A percept is the agent's observation of the current state of the environment. It could include sensor readings, information from the environment, or any other data relevant to the agent's decision-making process.

Persistent Variables:Persistent variables are those that retain their values across multiple iterations or invocations of a function or program.

- **belief state**: This variable represents the agent's probabilistic beliefs about the current state of the world. It encapsulates the agent's understanding of the environment based on past observations, actions, and any prior knowledge.
- **action:** This variable holds the action selected by the agent.

Update Belief State: The function updates the agent's belief state based on the action taken and the current percept received. This update incorporates new information obtained from the environment and adjusts the agent's understanding of the world accordingly.

Calculate Outcome Probabilities: The function calculates the probabilities of different outcomes for each possible action. It uses action descriptions and the current belief state to estimate the likelihood of various results or consequences that may occur as a result of taking each action.

Select Action with Highest Expected Utility: Based on the calculated probabilities of outcomes and utility information associated with each outcome, the function selects the action with the highest expected utility. Expected utility is a measure that takes into account both the probability of an outcome and its utility (i.e., how desirable or beneficial it is to the agent).

Return Action: Finally, the function returns the selected action, which the agent will then execute in the environment.

In summary, the DT-AGENT function enables an agent to make decisions in an uncertain environment by updating its beliefs, estimating the probabilities of different outcomes for potential actions, and selecting the action that maximizes its expected utility based on these probabilities and utility information.

The upcoming chapters focus on representing and computing with probabilistic information in general, as well as specific tasks such as representing and updating belief states over time and predicting the environment. Further chapters delve deeper into utility theory and develop algorithms for planning sequences of actions in uncertain environments.