Notes for Meeting 9
Decision Making and Choice

Two Aspects of Intelligent Behavior

Our recent sessions have focused on one important aspect of intelligent behavior that:

- interprets observations
- understands the environment
- makes inferences about situations

However, intelligent systems also have a generative side that:

- makes decisions
- carries out activities
- produces plans for the future

The most basic is decision making, which underlies many other activities.

The Need for Decision Making

Research on decision making has typically formulated this task as:

Given: Some aim one desires to achieve.
Given: A set of alternatives that may achieve that aim.
Find: The best of the alternatives, which involves:
- Evaluating the candidates
- Comparing them on relevant dimensions
- Selecting the most appealing candidate

This seems straightforward, but there are different accounts of how such decisions are (and how they should be) made.

Decision Theory

Decision theory is a formal framework that specifies how to choose among alternatives by:

- Listing the set of alternative actions or choices
- Listing the possible outcomes of each alternative
- Estimating the value of each possible outcome
- Estimating the probability of each possible outcome
- Multiplying the probability and value of each outcome
- Selecting the action or choice with the highest expected value.

Some variants also take the COST of each alternative into account.

Many view decision theory as the "proper" way to make choices; thus, it is a normative framework for decision making.

A Decision-Theoretic Example

Suppose that you need to buy a car, but you have limited options:

- There are two used cars - A and B - available for the same price.
- Either car may be in good working order or it may be a lemon.
- Car A would be more fun to drive (10) than car B (8), but if either is a lemon it would be little fun (0).
- The probability that A is a not a lemon is 0.7, while the same probability for B is 0.9.
- The expected value for A is 0.7 x 10 + 0.3 x 0 = 0.7, while the expected value for B is 0.9 x 8 + 0.1 x 0 = 0.72.
- Because car B has the highest score, decision theory states that you should prefer it over A.

This approach generalizes to situations involving N alternatives.

Other Decision-Making Examples

Now let us consider some examples of real-world decision-making tasks.

- You are diagnosed with a disease, but the operation is risky. How do you decide whether to have the procedure?
- You need to buy detergent for home. There are ten brands on the shelf in the supermarket. Which one do you pick?
- You are playing chess in a competition. How do you decide on your next move?
- You want to get married and raise a family. How do you decide on your partner?

These examples demonstrate that people do not typically use decision theory to make choices.

Does this mean they are irrational? Or could it mean there are other rational ways to decide?

The Incompleteness of Decision Theory

Simon (1993) notes that decision making is a more complex activity than normally assumed; it involves:

1. Selecting problems on which to focus attention;
2. Generating alternatives from which to choose; and
3. Evaluating and selecting among the generated alternatives.

Decision theorists have emphasized the third step over the others, which they typically ignore.

A full account of intelligent behavior must address all three issues.

Simon, H. A. (1993). Decision making: Rational, nonrational, and irrational. Educational Administration Quarterly, 29, 392-411.

Optimality and Satisficing

Some fields like economics assume that people make optimal decisions, but Simon (1993) notes that:

- Given the complexity of the world, the meaning of this claim is far from clear; what does optimal mean in such situations?
- Even in constrained contexts, there is clear evidence that people SATISFICE, i.e., they select acceptable alternatives.
- Simon refers to this as the theory of BOUNDED RATIONALITY; he views optimizing as prescriptive and satisficing as descriptive.
- The very notion of aspiration levels suggests that people are satisfied with situations that are good enough rather than best.

Could there be good reasons why humans operate in this manner? If so, should we build AI systems that work in the same way?

Satisficing and Heuristics

How do people make satisficing decisions with their bounded resources? Simon claims that they:

- They draw upon HEURISTICS or rules of thumb that generally produce acceptable results with limited effort.

- Many heuristics are linked to CHUNKS, which are long-term knowledge elements that encode useful patterns.

- They access these chunks through symbolic pattern matching that rapidly retrieves structures relevant to the current situation.

- Experts rely on chunks to behave intuitively, but can fall back on more costly methods (like search) if they encounter novel situations.

Taken together, these methods let humans survive in a complex world and accomplish many goals, despite their resources.

Misconceptions About Heuristics

Gigerenzer (2008) takes Simon's arguments further and identifies some erroneous beliefs about heuristics:

1. Optimization methods always produce better results than heuristics.

2. People rely on heuristics only because of cognitive limitations.

3. Humans resort to heuristics only on routine tasks of little import.

4. Only less cognitively sophisticated people rely on heuristics.

5. More information and computation always leads to better decisions.

He corrects each of these assumptions, some with empirical evidence and others with computational studies.

Gigerenzer concludes that heuristics will be central to any intelligent system that operates in a complex, uncertain environment.

Advantages of Heuristics

Gigerenzer also notes some general benefits of using heuristics over more "sophisticated" techniques; they are:

- Computationally tractable, in that they sidestep the combinatorial complexity of many tasks;

- Robust, in that they reduce the chance of overfitting which plagues more complex schemes; and

- Adaptive, in that they reflect which methods are effective in the physical and social environment.

Again, these arguments suggest that people do not use heuristics because they are cognitively impaired, but because they WORK.

Some General Heuristics

In addition, Gigerenzer lists some generic heuristics that people use:

- Recognition. Infer that a recognized alternative is better.
- Fluency. Assume that a more rapidly recognized choice is better.
- Take the best. Inspect attributes in order of importance and select the first candidate that dominates on an attribute.
- Tallying. Select the candidate with the most positive cues.
- Equality. Allocate resources equally to each of N alternatives.
- Default. If there is a default choice, then select it.
- Imitate majority. Do the same as most others in your peer group.
- Imitate success. Take a successful person as your role model.

These all have broad applicability in many decision-making contexts.

Summary Remarks

An intelligent agent must do more than understand its situation; it make also make choices about how to act.

Decision theory is a commonly adopted framework for determining the "correct" decision in a situation, but it is:

- Incomplete: it does not select problems or generate alternative;

- Intractable: it does not scale to combinatorial environments;

- Fragile: it relies on information that does not generalize well.

There is strong evidence that humans satisfice rather than optimize.

There are also powerful reasons to use heuristics and satisficing in AI systems.

Assignments for Meeting 10
Reactive Control

Read the articles:

- Horswill, I. (2008). Lightweight procedural animation with believable physical interaction. Proceedings of the Fourth Conference on Artificial Intelligence and Interactive Digital Entertainment. Stanford CA: AAAI Press. [required]

- Agre, P.E., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. Proceedings of the Sixth National Conference on Artificial Intelligence (pp. 268-272). Seattle: AAAI Press. [optional]

- Complete the third exercise (due 11:59 PM on 2/21/2011).