Notes for Meeting 10
Reactive Control

Review Of Decision Making

In the previous session, we examined the task of decision making, which (in most formulations) involves:

- Comparing alternative choices on relevant dimensions

- Evaluating candidates in terms of these attributes

- Selecting one alternative in favor of others

However, intelligent agents must do more than make isolated choices.

They must operate in some environment over time, and thus make a SEQUENCE of decisions that let them reach their objectives.

The Task of Reactive Control

The most straightforward way to support extended behavior is to invoke decision making repeatedly.

We will refer to this as REACTIVE CONTROL, which we can state as:

- Given: A description of the current environmental situation;

- Given: A set of alternative actions that apply in the situation;

- Select: One of these actions for execution in the environment.

This formulation is very general but quite simple, which has made it attractive to researchers from a variety of disciplines.

Applications of Reactive Control

Despite their simplicity, reactive control methods have been applied successfully to well-defined tasks like:

- landing and launching airplanes

- driving autonomous ground vehicles

- playing Backgammon and other games

These domains share a close relationship between the environmental situation and appropriate responses.

Complex tasks may be decomposed manually into multiple reactive tasks.

Disciplines Adopting Reactive Control

The basic framework of reactive control has been studied intensively by at least three disciplines:

- Control theory

- Behaviorist psychology

- Artificial intelligence

These fields have approached the idea from different perspectives, but their responses have shown remarkable similarity.

In general, they all emphasize direct links between perception and action while downplaying high-level cognition.

Control Theory

One approach to reactive behavior - control theory - comes from applied mathematics.

This framework assumes that one controls a physical system by:

- measuring the values of sensor variables that describe the situation;

- using equations to calculate new values of control variables as a function of these sensor variables; and

- altering the values of control variables based on these calculations.

Work in this paradigm has emphasized guaranteed optimality of behavior.

This has led it to focus on linear and other restricted control schemes, although recently this has expanded to hybrid control methods.

Behaviorist Psychology

From 1940 to 1970, American psychology was dominated by the behaviorist paradigm. Its key characteristics included:

- rejection of hypothetical short-term mental structures (e.g., beliefs);

- emphasis on direct connections between stimuli (perceptions) and responses (actions);

- an abiding interest in motivation and learning from reward; and

- a reliance on rats, pigeons, and other non-primates as data sources.

These biases led behaviorism to uncover important regularities about learning, but also to ignore many aspects of high-level cognition.

Stimulus-Response Systems in AI

Another line of work, in AI, introduced reactive methods in response to perceived limits on planning systems.

The "stimulus-response systems" associated with this paradigm are characterized by:

- encoding knowledge as a set of stimulus-response patterns;

- operating in cycles, each of which involves:

  - retrieving a pattern that matches the current stimuli;

  - executing the associated response in the environment.

This approach is closely related to behaviorist psychology and forms the basis of most AI work on reinforcement learning.

Note that it assumes no short-term mental structures, such as internal beliefs or goals; one can argue these are not cognitive systems.

Limitations of Stimulus-Response Systems

What cognitive capacities do humans have that are lacking in stimulus-response systems? We know that people:

- Draw inferences about their current situation

- Have memories about past decisions and events

- Use knowledge about the order of activities

- Organize complex activities into subactivities

- Take goals into account when making decisions

One can demonstrate these capabilities on standard puzzles like the Tower of Hanoi.

Variations on Stimulus-Response Systems

Researchers have explored a number of extensions to the basic framework of stimulus-response systems that:

- carry out inference to enrich state descriptions

- introduce memory for past decisions/events into state descriptions

- organize knowledge into "behaviors" that influence each other

- organize knowledge into hierarchies that provide continuity

- explicitly encode goals and use them to guide decisions

These extensions provide capabilities that are much closer to those found in human cognition.

However, some communities continue to focus on purely reactive control, despite its inherent limitations.

Summary Remarks

In summary, reactive control is a simple and general framework for using decision making to produce sequential behavior.

This approach has been widely adopted in control theory and in some circles of psychology and artificial intelligence.

These efforts share a focus on sensori-motor abilities, typically to the exclusion of high-level cognition.

However, extensions to the basic framework support some higher-level functions and suggest a path to cognitive systems.

Assignments for Meeting 11
Hierarchical Control

Read the articles:

- Georgeff, M. P. and Lansky, A. L. (1987). Reactive reasoning and planning. Proceedings of the Sixth National Conference on Artificial Intelligence (pp. 677-682). Seattle, WA: AAAI Press. [required]

- Langley, P., & Choi, D. (2011). Icarus user’s manual (Technical Report). Institute for the Study of Learning and Expertise, Palo Alto, CA. (Read Section 3) [required]

- Examine the third exercise (due 11:59 PM on 3/2/2011) and bring questions about it to class.