Notes for Meeting 2
Frameworks and Metaphors for Intelligent Systems

Computer Science as an Empirical Discipline

Newell and Simon (1976) make some important claims about our field:

- Computer science is an empirical discipline, rather than a branch of mathematics.
- It is a science of the artificial, in that it constructs artifacts of sufficient complexity that formal analysis is not tractable.
- Instead, it must study these computational artifacts as though they were natural systems.
- Thus, it formulates hypotheses and collects evidence that supports or detracts from them.

In this article, they propose two hypotheses based on their founding work in list processing and artificial intelligence.

Laws of Qualitative Structure

The authors introduce the notion of laws of qualitative structure, which are crucial for the development of any scientific field:
- the cell doctrine in biology
- plate tectonics in geology
- the germ theory of disease
- the atomic theory of matter

Computer science also has a need for such high-level principles.

They propose two such laws, one related to mental structures and another and the other to mental processes.

What is a Physical Symbol System?

Newell and Simon introduce the notion of a physical symbol system:

- Symbols are physical patterns that remain stable unless modified.
- Symbol structures or expressions are organized sets of symbols.
- A physical symbol system includes processes for creating, modifying, copying, and destroying symbol structures; these let it:
  - Maintain structures that designate other objects or processes.
  - Interpret expressions that designate such processes.

These ideas are agnostic about the nature of the physical patterns; they can arise in neurons, in silicon chips, or as marks on paper.

The Physical Symbol System Hypothesis

Newell and Simon also present an important claim about such systems:

- A physical symbol system has the necessary and sufficient means for general intelligent action.
- They emphasize general intelligence here, such as humans exhibit, rather than specialized abilities.
- Research on the sufficiency side is associated with the field of artificial intelligence.
- Research on the necessity side is associated with computational models of human cognition.

This is a theoretical claim that is subject to empirical tests, but evidence to date generally supports it.

Development of the Hypothesis

Four historical developments in the 20th Century led to the Physical Symbol System Hypothesis:

- Studies in formal logic and symbol manipulation
- Turing machines and digital computers
- The concept of a stored program
- List processing and languages like IPL and Lisp

Later work on computing systems that designated and interpreted symbol structures built on these advances.

What is List Processing?

Newell, Shaw, and Simon introduced list processing in 1956. This involved three key ideas:

- Dynamic memory structures and mechanisms to alter them
- Data types and operations for different types
- Designation and manipulation of symbol structures

These support abstraction of structures and processes beyond the specific hardware on which they are implemented.

They demonstrated these ideas in IPL, the first list-processing language, although Lisp soon became more widely used.

Why is List Processing Important?

This new framework was important to AI’s development because it could:

- Encode arbitrarily complex structural descriptions
- Create new structural descriptions dynamically
- Use such structures to designate other structures
- Intepret these structures to produce behavior

As we will see, each of these abilities plays a crucial role in the construction of intelligent systems.

Problem Solving and Symbols Systems

- Human intelligence includes the ability to solve novel problems.
- How can we find solutions when we do not already know the answer?
  - This is the Meno Paradox from Plato.
- We can separate the generator of candidate solutions from the test.
  - This division eliminates the apparent paradox.
- This requires the ability to represent candidate solutions and to search through the resulting problem space.

The Heuristic Search Hypothesis

Newell and Simon propose a second law of qualitative structure:

- A problem solver represents candidate solutions as symbol structures.
- Problem solving occurs by generating/modifying these structures.
- The problem solver tests each candidate to determine acceptability.
- This continues until the problem solver finds a solution.

They include the term "heuristic" because, practically, an intelligent system cannot afford to search most problem spaces exhaustively.

Representing and Interpreting Problem States

- Before a system can search a problem space, it must encode states or situations in that space.
- Physical symbol systems have the representational capabilities needed for this purpose.
- The system must also be able to generate new problem states from previous ones.
- The dynamic character and interpretability of physical symbol systems let them support this capacity as well.
- Interpretability also enables the ability to determine whether a problem state satisfies the test and thus solves the problem.

Mitigating the Combinatorial Explosion

- Although the symbol structures that define a problem space may be quite small, the space itself may be very large.
- In many cases, the number of states grows exponentially with the number of steps from the initial to solution state.
- One key facet of intelligence lies in the ability to search such spaces selectively rather than exhaustively.
- Intelligence lies in the ability to use heuristic knowledge about regularities in the space to limit or even eliminate search.
- The term "heuristic" is viewed with disdain by some researchers, but it is central to understanding the mind.

The Logic Theorist

- In 1956, Newell, Shaw, and Simon developed the Logic Theory Machine, the first running AI program.
- This system proved theorems in the propositional calculus, finding proofs that differed from those in Russell and Whitehead.
- The Logic Theorist demonstrated the use of list processing, heuristic search, and goal-directed reasoning.
- It would difficult to overestimate the system’s impact on the new field of artificial intelligence.

Representation Selection in Problem Solving

- Newell and Simon note that a key step in problem solving - problem formulation - involves creating the problem space.
- This requires the system to select a representation for problem states that make up the space.
- Decisions about problem representation can have a major impact on solution difficulty (e.g., the mutilated checkerboard).
- One can view selection and revision of problem spaces as a higher form of heuristic search.
- However, there has been very little research done on this topic.

Summary

Newell and Simon make broad claims about the nature of intelligence:

1. Computer science and AI are empirical disciplines that involve studying the behavior of computational artifacts.
2. Physical symbol systems are necessary and sufficient for general intelligent behavior (the physical symbol system hypothesis).
3. Problem solving involves heuristic search through a space of states encoded as symbolic structures (the heuristic search hypothesis).

They incorporated these ideas into the Logic Theorist, the first AI system, and they provide a foundation for all that has followed.

Assignments for Meeting 3
Symbolic Processing and Lisp

Read the course notes:

- Nau, D. (2008). Introduction to Lisp. Computer Science Department, University of Maryland, College Park.
- Hall, M. (1994). Basic Lisp overview. Computer Science Department, Johns Hopkins University, Baltimore, MD.

Lisp has been one of the most influential and widely used languages in the history of artificial intelligence.

The framework clarifies many of Newell and Simon’s claims about physical symbol systems.