Computational Intelligence

A Logical Approach

David Poole

Alan Mackworth

Randy Goebel

Oxford University Press

1998

What is Computational Intelligence?

The study of the design of intelligent agents.

An agent is something that acts in an environment.

An intelligent agent is an agent that acts intelligently:

- its actions are appropriate for its goals and circumstances
- it is flexible to changing environments and goals
- it learns from experience
- it makes appropriate choices given perceptual limitations and finite computation

Artificial or Computational Intelligence?

- The field is often called Artificial Intelligence.
- Scientific goal: to understand the principles that make intelligent behavior possible, in natural or artificial systems.
- Engineering goal: to specify methods for the design of useful, intelligent artifacts.
- Analogy between studying flying machines and thinking machines.

Central hypotheses of CI

Symbol-system hypothesis:

- Reasoning is symbol manipulation.

Church–Turing thesis:

- Any symbol manipulation can be carried out on a Turing machine.

Agents in the World

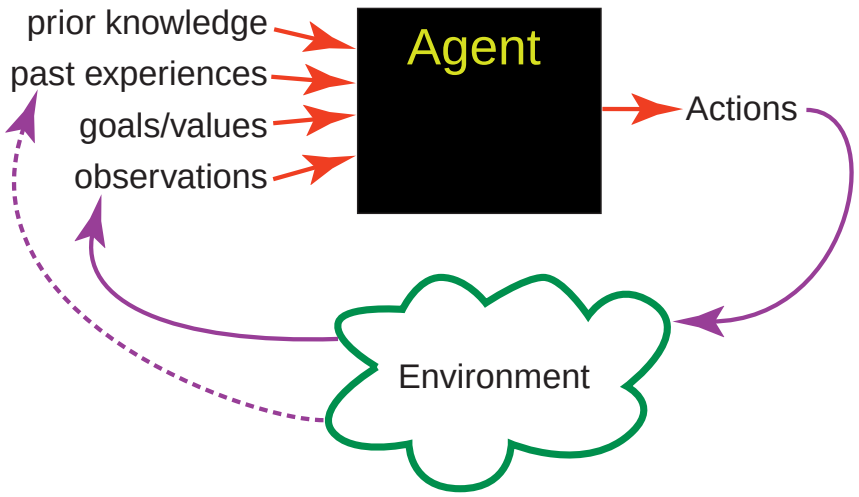

Representation and Reasoning

To use these inputs an agent needs a representation of them.

$\Longrightarrow$ knowledge

Most common sense tasks rely on a lot of knowledge.

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- Language to communicate with the computer.
- A way to assign meaning to the symbols.
- Procedures to compute answers or solve problems.

Example RRSs:

- Programming languages: Fortran, C++,…
- Natural Language

We want something between these extremes.

Example Application Domains

- Autonomous delivery robot roams around an office environment and delivers coffee, parcels,…
- Diagnostic assistant helps a human troubleshoot problems and suggests repairs or treatments. E.g., electrical problems, medical diagnosis.
- Infobot searches for information on a computer system or network.

Domain for Delivery Robot

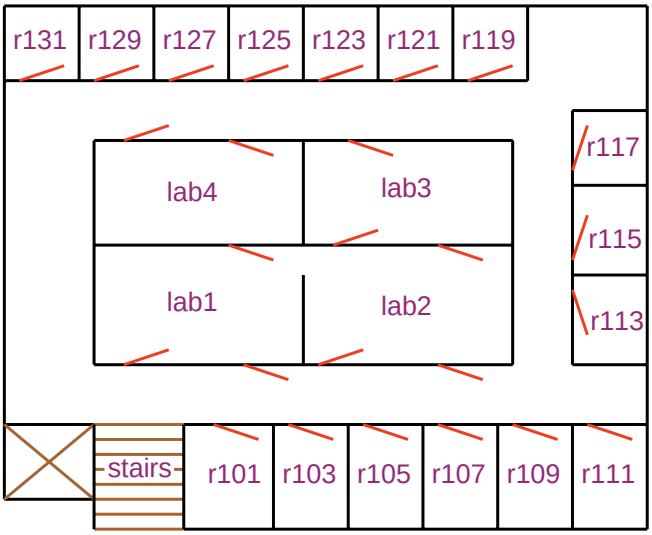

Autonomous Delivery Robot

Example inputs:

- Prior knowledge: its capabilities, objects it may encounter, maps.
- Past experience: which actions are useful and when, what objects are there, how its actions affect its position.
- Goals: what it needs to deliver and when, tradeoffs between acting quickly and acting safely.
- Observations: about its environment from cameras, sonar, sound, laser range finders, or keyboards.

What does the Delivery Robot need to do?

Determine where Craig's office is. Where coffee is…

Find a path between locations.

Plan how to carry out multiple tasks.

Make default assumptions about where Craig is.

Make tradeoffs under uncertainty: should it go near the stairs?

Learn from experience.

Sense the world, avoid obstacles, pickup and put down coffee.

Domain for Diagnostic Assistant

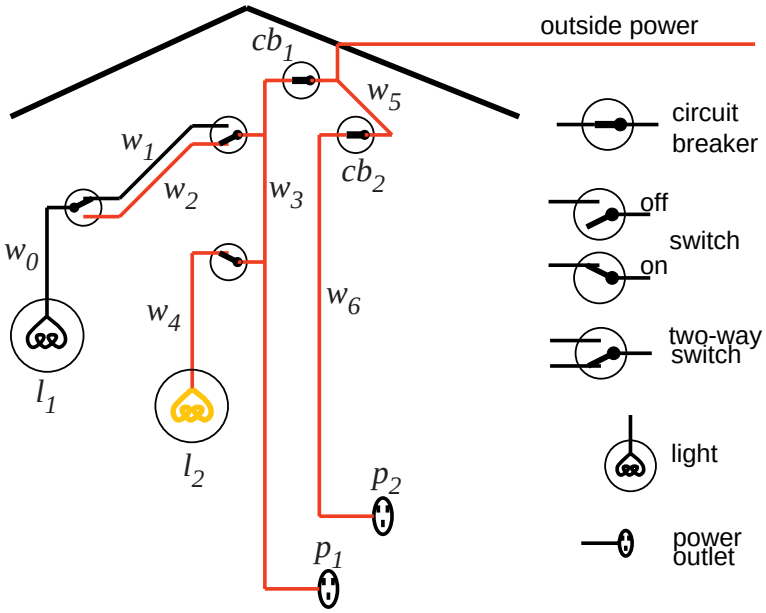

Diagnostic Assistant

Example inputs:

- Prior knowledge: how switches and lights work, how malfunctions manifest themselves, what information tests provide, the side effects of repairs.
- Past experience: the effects of repairs or treatments, the prevalence of faults or diseases.
- Goals: fixing the device and tradeoffs between fixing or replacing different components.
- Observations: symptoms of a device or patient.

Subtasks for the diagnostic assistant

Derive the effects of faults and interventions.

Search through the space of possible fault complexes.

Explain its reasoning to the human who is using it.

Derive possible causes for symptoms; rule out other causes.

Plan courses of tests and treatments to address the problems.

Reason about the uncertainties/ambiguities given symptoms.

Trade off alternate courses of action.

Learn about what symptoms are associated with the faults, the effects of treatments, and the accuracy of tests.

Infobot

Infobot interacts with an information environment:

- It takes in high-level, perhaps informal, queries.
- It finds relevant information.
- It presents the information in a meaningful way.

Infobot inputs

- Prior knowledge: the meaning of words, the types of information sources, and how to access information.
- Past experience: where information can be obtained, the relative speed of various servers, and information about the preferences of the user.
- Goals: the information it needs to find out; tradeoffs between the volume and quality of information and the expense involved.
- Observations: what information is at the current sites; what links are available; the load on various connections.

Example subtasks for the Infobot

Derive information that is only implicit in a knowledge base.

Interact in natural language.

Find good representations of knowledge.

Explain how an answer was derived and why some information was unavailable.

Make conclusions about the lack of knowledge or conflicting knowledge.

Make default inferences about where to find information.

Make tradeoffs between information quality and cost.

Learn the preferences of users.

Common Tasks of the Domains

- Modeling the environment Build models of the physical environment, patient, or information environment.
- Evidential reasoning or perception Given observations, determine what the world is like.
- Action Given a model of the world and a goal, determine what should be done.
- Learning from past experiences Learn about the specific case and the population of cases.

Our approach to teaching CI

- Our goal is to study these four tasks.
- We build the tools needed from the bottom up.
- We start with some restrictive simplifying assumptions and lift them as we get more sophisticated representations and more powerful reasoning strategies.
- The theory and practice are built from solid foundations.