CPSC 322
Introduction to Artificial Intelligence

September 13, 2004

Highlights from last time

We're moving forward under some assumptions:

Whatever intelligence is, it results from some kind of computation and it's platform independent. It's not unique to brains.

Symbol manipulation is a type of computation that is sufficient to give rise to intelligent behavior.

Any symbol manipulation can be carried out on a Turing machine (and, by the way, computers are TMs).

But keep in mind that these assumptions may be wrong.

The Diagnostic Assistant Domain

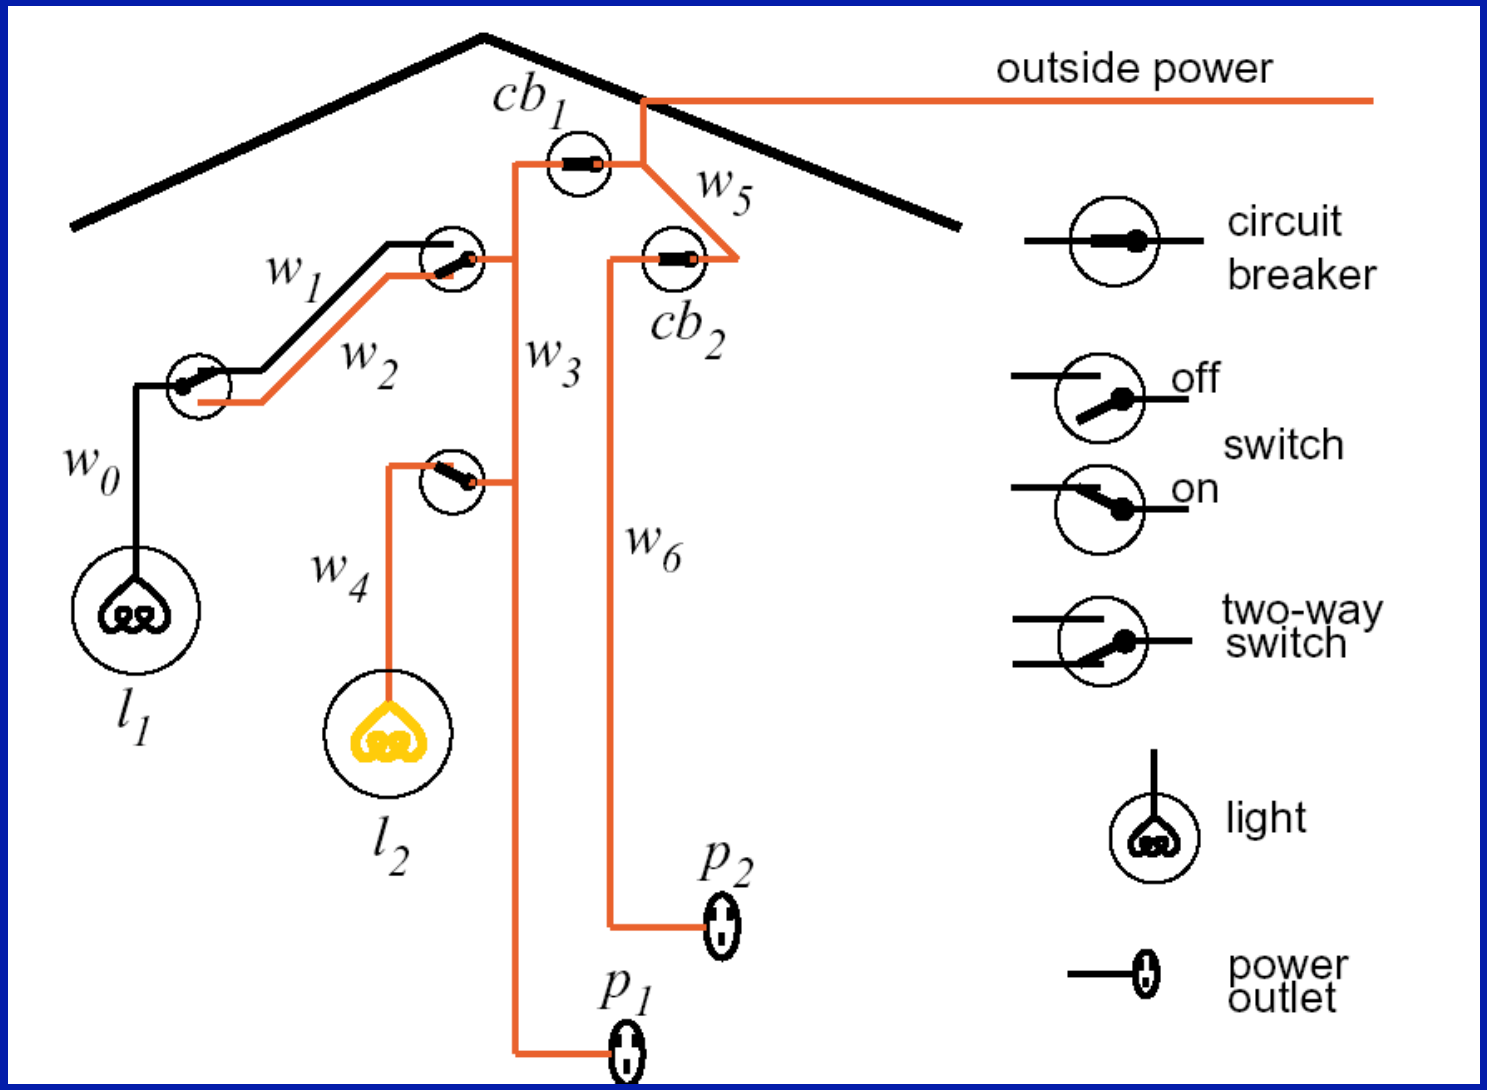

What the Ideal Assistant Should Do

Derive effects of faults

Search through the space of possible faults

Explain its reasoning to human users

Derive possible causes for symptoms; rule out other causes

Plan courses of tests and repairs to address problems

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

What Our Assistant Will Do*

Tell us whether one specific light bulb is on

* today...more to come later

The Intelligent Agent

An intelligent agent is a system that

- acts appropriately for its circumstances and its goal
- is flexible when faced with changing environments and changing goals
- learns from experience
- makes appropriate choices given perceptual limitations and finite computation

The Intelligent Agent as Black Box

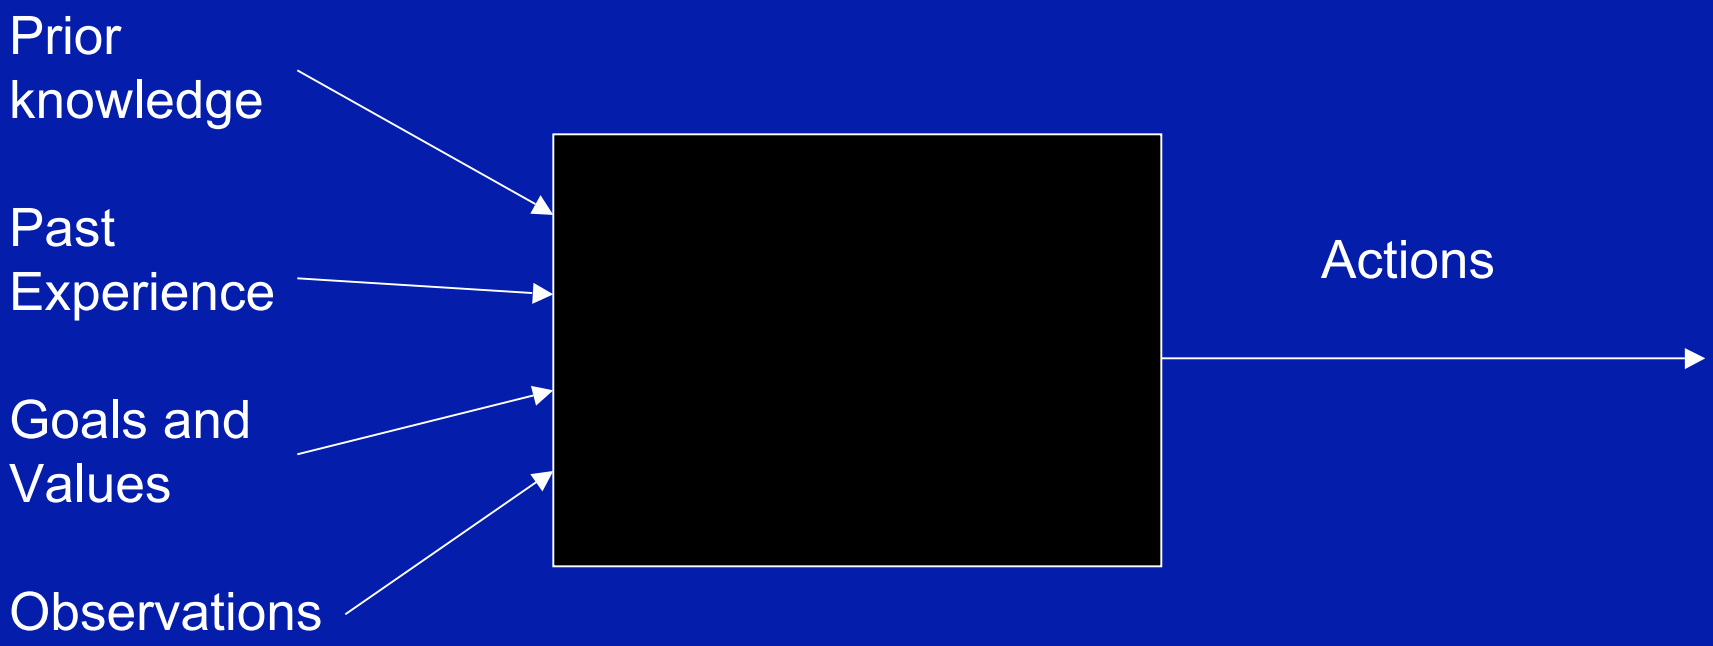

The Intelligent Agent as Black Box

prior knowledge - how switches and lights work, how malfunctions manifest themselves, what information do tests provide

past experiences - data from previous cases including effects of repairs, prevalence of faults, prevalence of symptoms for faults, accuracy of tests

goals and values - goals of fixing the device, tradeoffs between fixing or replacing components

observations - symptoms of a device in failure

The Intelligent Agent as Black Box

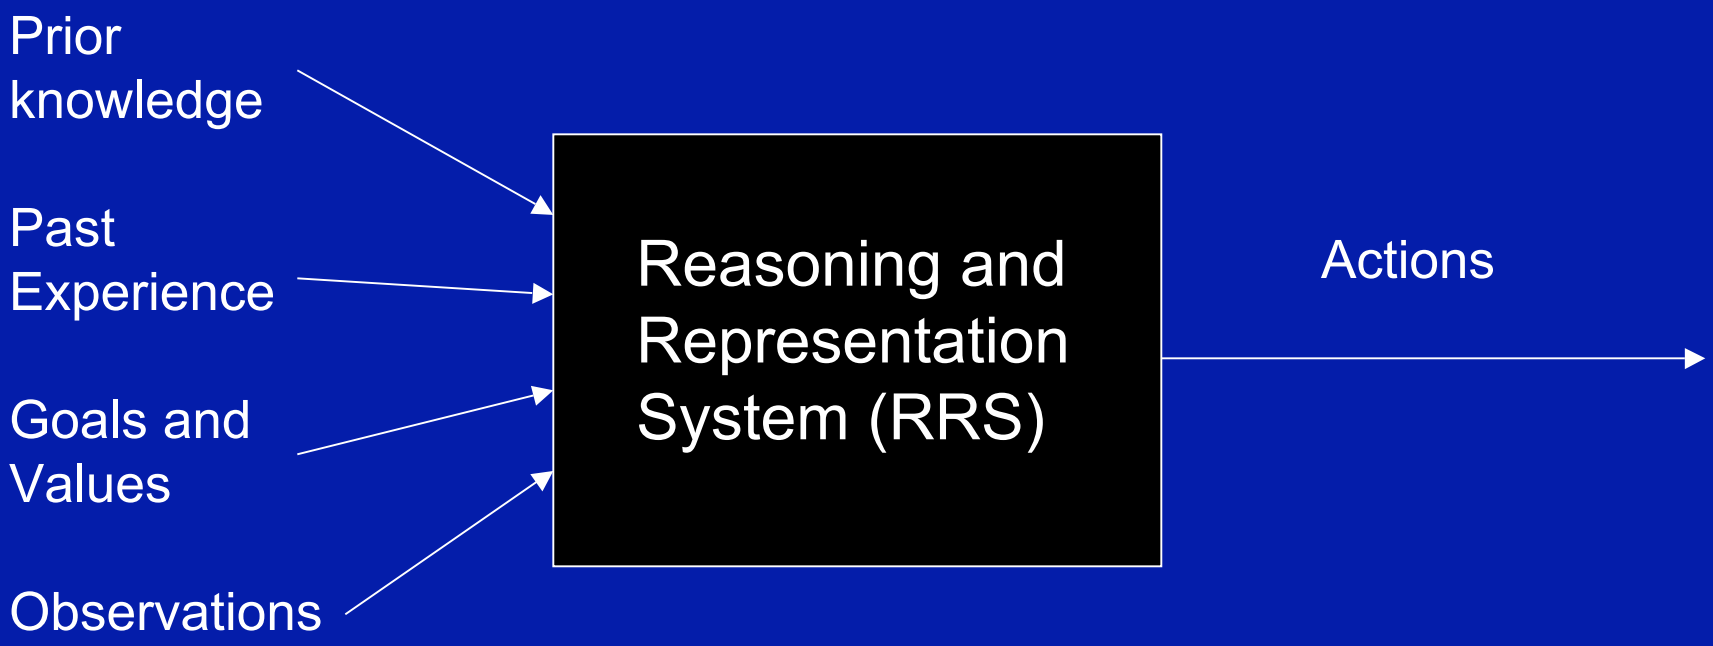

Reasoning and Representation System

A language for communication with the computer

A way to assign meaning to the language

Procedures to compute answers to problems given as input in the language

Where does the RRS come from? You!

Five Simple Steps to World Domination

(or how to build the black box)

1. Begin with a task domain that you want to characterize
2. Distinguish the things you want to talk about in the domain (the ontology)

What goes into the ontology?

Description of what exists in the domain of interest:

individuals/”things”
properties/attributes of individuals
relationships between individuals

Five Simple Steps to World Domination

(or how to build the black box)

1. Begin with a task domain that you want to characterize
2. Distinguish the things you want to talk about in the domain (the ontology)
3. Use symbols in the computer to represent objects and relations in the domain. (Symbols *denote* objects...they're not really the objects.)
4. Tell the computer the knowledge about the domain.
5. Ask the RRS a question which prompts the RRS to reason with its knowledge to solve a problem, produce answers, or generate actions

Whoa!

Dude! Reasoning?! The last step is reasoning? Isn't this like saying "and then a miracle happens"? Where does the reasoning come from? Why don't you tell us the 273 steps toward building the reasoning part?

Starting up CILOG

From any UBC CS undergraduate machines:

```
% ~cs322/cilog/cilog
```

Or from within SWI Prolog:

```
?- [cilog_swi].

CILOG Version 0.12. Copyright 1998, David Poole.
CILOG comes with absolutely no warranty.
All inputs end with a period. Type "help." for
help.

cilog:
```

Is the representation sufficient?

What was left out?

How about relationships like

connected(l2, w4)

Or

between(s3, l2, outside_power)

How do you know when you've got it right?

Starting on Wednesday...

We'll be talking about the CILOG RRS in much more detail

Things to do

Start reading Chapter 2.
Read it carefully.
Read it again.
Read it with CILOG running in front of you.
Run all the examples while you're reading.
Then go back and read it again.