Object descriptions add value to plans

Will Fitzgerald
R. James Firby
I/NET, Inc.
AAAI Fall Symposium 1996

July 2001 Printing

Descriptions Add Value

Traditional AI planning systems were concerned more with the ordering of plan operators than with the objects these operators acted upon. These systems assumed (usually without stating it) that the plan execution system could easily map from an internal, symbolic representation to the object represented in the real world. For example, if a planning system created the plan step `(pickup block-a)`, a one-to-one mapping existed between the symbol `block-a` and some block in the real world. It also assumed that the plan execution system could immediately sense and act that that block.

This is clearly an oversimplification; the real world does not provide an immediate connection from objects to a planning system's internal representations. This has been recognized for some time. For example, Agree and Chapman (1990), argue for a radically situated planning system in which object representations are strictly deictic. In their terms, this means an active, functionally motivated causal relationship between an agent and an object in the world, such as `the-bee-i-am-following`. Another approach is to distinguish between internal representations and sensor names (Firby, 1989). As an agent moves around in the world, it can use the information it receives from its sensors to disambiguate objects it senses into its internal representations.

What is missing from these approaches is the value of giving planning agents access to descriptions of objects in the world in addition to representations of objects in the world. Plans typically describe how to act on objects that are currently believed to have some properties. We suggest that it is often powerful to write plans to act on objects that meet a description of those properties.

For example, an agent might have a plan to paint all red fuel drums blue. If the agent acts only on what it currently believes about objects in the world, then it may miss some fuel drums, because it does not know enough about them to justify action. If, however, an agent carries with it the description of the properties, then it can act to find out more about objects in the world. For example, an agent might know that a particular object was a fuel drum, but not its current color; or it might know that it is red, but not what kind of object it is. Without access to a description of the objects upon which an agent is to act, an agent cannot tell whether it needs to get more information, act anyway, avoid action, etc.

An object description carries with it two types of information. First, it describes the kind of object that is being described (for example, that the object is a member of the set of fuel drums). Second, it describes the qualities of the object being described (for example, that the color is red). The information about qualities can be missing, of course (for example, an agent could look for any fuel drum). Similarly, the information about kind can be missing (for example, an agent can look for objects that are red in color). In our current syntax, this information is defined using describes and slot-value propositions. So the description underlying ”red fuel drum” is:

```
(describes object-description-n fuel-drum)
(slot-value object-description-n color red)
```

In conjunction with propositions of the sort:

```
(instance-of object-m fuel-drum)
```

```
(color object-m red)
```

different plans can be written which depend on either the object description or beliefs about the object. Plans can then depend on combinations of the object descriptions and beliefs about objects. For example, it is now possible to express, ”if you’re asked to paint fuel drums red, and the object is already red, don’t paint it.”

Object descriptions have three more useful qualities. First, recording object descriptions allows an agent to recognize that it is being requested to act on the same type of thing as it was before (for example, ”you just asked me for a red fuel drum, so I’m not going to comply with the request”). Second, they can mimic higher order propositional logical propositions within a weaker, frame-based logic, so one can assert, for example, John believes *p*, without being forced to assert *p*. Third, they map closely onto natural language descriptions of objects. When an agent is asked (for example) to paint, in natural language, the red fuel drums, what is being asked maps closely to the object descriptions we have described, not just propositions about objects. This is especially true when linguistic phenomena such as definite/indefinite reference are taken into consideration.

What we are arguing for is the use of information from three sources. First, there are sensing data from the real world such as whether the vision system senses a blue object. Second, there are stored beliefs about objects such as whether the memory system has recorded this object as being blue in color. Third, there are what we believe to be a distinct information source, object descriptions, such as whether the memory system has recorded that what is wanted is a blue object.

Bibliography

Agree, Philip E. and David Chapman (1990). What are plans for? Robotics and Autonomous Systems 6:17-34.

Firby, R. James (1989). Adaptive Execution in Complex Dynamic Worlds. Unpublished Ph.D. Thesis, Yale University.