AI and Problem Solving: Planning

Jocelyn Paine

June 3, 1997

In this tutorial, I shall try to provide a background to planning: to present some of the main ideas in conventional planning; to demonstrate the problems faced in the real world; to relate planning in AI systems to cognition; and to introduce some unconventional but promising approaches.

1 Planning — the early days

The point of this section is to show what planning is, and how it (some of it, anyway) proceeded from search and GPS. Apart from GPS, none of the planning methods described are intended as psychological models — they're all exercises in getting machines to plan, exercises which uncovered more and more difficulties in what originally seemed a simple problem. This lack of psychological relevance isn't a criticism, of course; AI is concerned not only with existing natural intelligences, but also with the structure of all possible intelligences, existing or not. But although none of these post-GPS planners are intended as cognitive models, some of them have undoubtedly taken inspiration from human problem solving. And they have enabled us to clarify and formalise problems (such as interacting subgoals) which all planners must solve, whether human or computer.

I want to introduce these ideas, for each of which I'll give some references.

1. What is planning? Some terminology.
2. The blocks world.
3. Planning as state-space search.
4. Means-end analysis.
5. Planning with abstraction hierarchies.
6. The linearity assumption and interaction between goals.

Here are the references. When reading, concentrate on the concepts, rather than on technicalities. Remember that in Finals, you will have at most an hour to talk about planning. You should be able, by comparing and contrasting different planners, to write an essay on the problems of planning and the way in which people have tried to solve them. It will be helpful to have a little example ready for each, to illustrate the way in which each planner might work with a problem; but you won't be expected (I hope) to go into tremendous detail. (Having said this, if you think that abstraction spaces, for instance, is a really beautiful concept, don't feel inhibited from taking it as far as you want.)

Planning and terminology. Essentially, a plan is a set of actions which, when executed, will (barring failures) enable some agent to achieve its goal or goals, given the current state of the world. Most simply, a goal is (as in state-space search and GPS) just a new state of the world. But there may be other types of goal. The planner needs to know how the agent's actions will affect the world; and it may have to deal with other changes or events.

Start by finding *Readings in Planning* (now in the Psychology library) edited by Allen, Hendler and Tate. Go to the first article — *Planning* — and read enough of sections 1, 2, 2.1 and 3, 3.1, 4, and 4.1 to acquaint yourself with the terms action, state, event, execution, planning, plan structure and subplans, atomic actions or plans, partial plan, goal, agent, and multi-agent planning. These are fairly well-defined terms, and papers on planning will often use them without further ado. Note the different types of goal that are described in section 3.

This and the next paper *A Review of AI Planning Techniques* form a general introduction to planning, and you may find some parts of both useful throughout this set of readings (though other parts will be too mathematical). Note the family tree on page 27.

The blocks world. Many early planners planned for an agent living in the blocks world. Make sure you know what is meant by this — what objects are found in a typical blocks world, and what actions can you perform on them? Almost any book on AI mentions the blocks world; *AI and Natural Man* is as good as any. What are the differences between the blocks world and real life?

Planning as state-space search. The simplest way to plan is to search all possible states of the world until you find the one you want. This is state-space search, which I'll assume you've done in the *Conventional AI: Search* tutorial. Consider a typical problem, Missionaries and Cannibals for example. What might a solution plan look like? What are the defects of planning this way?

Means-end analysis. After search came GPS. Again, I'll assume you've covered this in the *Conventional AI: GPS* tutorial. What were the problems with it?

Planning with abstraction hierarchies. One of the problems with GPS was that it couldn't defer details. If it set itself a subgoal to satisfy a particular operator, it had to plan how to fulfil all the preconditions of that operator before it could go any further. So it might keep going down and down, building ever more finicky subplans, before ever getting past that operator.

To avoid this problem, the idea of planning at various levels of abstraction was introduced, the first planner that used it being ABSTRIPS. Three references, in increasing order of detail: Winston page 154; *AI and Natural Man* pages 357–360 (end of ¶1); *Planning in a Hierarchy of Abstraction Spaces* in *Readings in Planning*, particularly pages 98–100.

The linearity assumption and interaction between goals. The methods above all worked on the linearity (or additivity) assumption: that subgoals could be tackled independently and the subplans for each pasted together to form a total plan. But that won't always work: sometimes interactions between goals make it impossible to consider each in isolation.

The first program to attempt interacting goals was Hacker. It, essentially, planned as though the goals were independent, and then tried to correct faults in the resulting plans. It could not deal with all interacting goal problems this way: for example, Sussman's problem (illustrated in the references below).

For this section, I shall concentrate on Hacker, and then refer you to a goal-interaction problem that it couldn't solve, and ask you to accept it as the motivation for further work.

Start with the *Introduction to Chapter 3* of *Readings in Planning*, pages 108–109. Then, for Hacker, read pages 286–297 of *AI and Natural Man*; if you want more, read *The Virtuous Nature of Bugs* in *Readings in Planning.* Finally, *AI and Natural Man* pages 360–366.

2 Planners versus production systems

At this point, you might think about the difference between planners and production systems. Apart from GPS, few planners have been used as models of human planning, whereas there are many production system models. Planners and production systems are both programs, but what is the difference between them?

First, the input — the specification of the problem to be solved — is in a rather different form. A planner is specialised for reading this specification as a list of actions and their effects (e.g. GPS' operator tables) plus a statement of the current world state and the goal. A production system (at least if used as a general model) must be much more general, because it may have to deal with many other kinds of input. In fact, to call the action and goal specifications input is a bit misleading, since they were probably created by the system in the first place. Indeed, it's possible that their creation and the act of planning might not be separable — are they in people?

Second, a planner contains a lot of specialised knowledge about *how to plan* — how to abstract, how to re-order goals, and so on. Where is this knowledge in a production system? Certainly not in the inference method.

3 Case-based planning

As I said above, few planners were intended as psychological models. Now I'll come to one that might form the basis for such a model. The idea is that when we plan, rather than building plans from scratch, we work by modifying a large memory of existing plans. This is called case-based planning.

There's an introduction and two references: the first is probably easier to follow. *Introduction to Part IV* of *Readings in Planning*, page 647. Then the article on *Chef*, pages 655–659. And *Case-based planning: a framework for planning from experience* by Hammond, in *Cognitive Science* volume 14 number 3 (July–September 1990). Don't read these articles if you're feeling hungry.

As above, don't go into too much detail. I don't suppose you'd ever need to write for more than twenty minutes on this one topic. So concentrate on the main points: the way in which Chef might model a person's experience-based planning, and a few example of how it mutates old plans into new.

4 Planning under limited knowledge or time

In this and the following sections, I'm working towards real robots or animals living in the real world, and I'm going to show just how big a leap this is from classical planning.

The previous two sections have concerned planning in some kind of ideal world, where an agent always has enough time and knowledge to predict the effect of a plan. That need not be so. One may lack time to form a good plan, so have to improvise with a bad; one may lack knowledge needed to choose one plan over another (or the time necessary to apply that knowledge). To get an idea of the approaches that have been tried, read *Introduction to Chapter 4* of *Readings in Planning*, pages 187–188; *Introduction to Artificial Intelligence* by Charniak and McDermott, pages 518–519; and if you want to go further, *Decision Theory and AI II* in *Readings in Planning*, especially pages 207–208 and 212–217.

5 Planning and acting

Though now considering resource limitations, we still haven't said much about executing plans. How do you fix matters up if one of your actions has an unexpected result, or if there's an interruption? Read Charniak and McDermott pages 524–527, and if you want to go further, skim *Planning and Acting* in *Readings in Planning*.

6 Opportunistic planning

As page 526 of Charniak and McDermott says, some have argued that we so often experience surprises, whether pleasant or unpleasant, that to think of planning as the systematic generation of plans is misleading; instead, planners must be able to take advantage of opportunities, whenever they occur. This is certainly a feature of human planning: see *A Cognitive Model of Planning* in *Readings in Planning*.

7 What else is needed?

Well, quite a lot actually. People from conventional AI often underestimate the material problems a robot or animal would face in the real world. These include motion control, perception, noise, and dealing with the opposition. Skim Charniak and McDermott pages 527–543 to get an idea of the problems faced by (a) assembly robots and (b) games players.

8 Unconventional approaches

Some researchers say that conventional AI has had such little success that we need a radically new approach. Their attitude, if one can generalise across the whole work, is that most of what's needed in a natural intelligence was placed there by evolution long before the human brain evolved (consider the two billion years life has existed with the history of the human race). Conventional AI has concentrated on rational intelligence, printing the postage stamp on top of the Empire State Building without worrying about the building itself. What we should be doing is to work our way up the evolutionary tree, making a complete creature at each step and then building the next on top of the last.

This work is new, and we can't yet say where it will lead. But it seems to have some tempting successes, and anyone who's going to take AI seriously should know of it. Start by finding the first issue of *Artificial Intelligence* for 1991 (numbers 1–3 in one volume), and read the article by Brooks and its criticism by Kirsh (incidentally, president of the AI Society here, 1982–83).

For more details, look at *AI at MIT, volume 2* edited by Winston and Shellard, pages 3–26; *Minimalist Mobile Robots* by Connell, chapter 1; and *Intelligence as Adaptive Behaviour* chapter 1. All these are now in the Psychology library.

Then note that not everybody agrees: read the contrary view in the *Forward* to *Readings in Planning.*

9 Planning and story-understanding

I'll end with an unexpected application, which suggests that we need planners to understand other people's plans as well as to create our own. Read chapter 26 of *Thinking — Readings in Cognitive Science*, edited by P.N.Johnson-Laird and P.C.Wason. It's in the Psychology library. You'll find more about this in Chapter 10 of Charniak and McDermott (which contains a bit on understanding characters' emotions).

If you're not familiar with scripts and frames, both referred to in this article, you may find the articles *Scripts* and *Frames* in the *Encyclopaedia of AI* (RSL) useful.

10 Must all intelligences plan in the same way?

To wind this up (and provide an antidote to the irrational view of aliens which TV will doubtless be portraying over the Christmas repeat season), read *Why intelligent aliens will be intelligible* by Minsky. It's photocopy M131 in the library.