Conventional AI: GPS

Jocelyn Paine

June 3, 1997

1 Introduction

Last week, I talked about heuristic search. We saw that a problem can be represented as a start state, a goal state, and a set of moves or *operators*. Now, heuristic searchers know *something* about operators; they have to, in order to change from one state to another. But their knowledge is very limited. Essentially, all they know is the operators that exist, and how, given an old state, each operator generates a new state, together perhaps with an evaluation function. But real problems demand more than this.

Other things being equal, if there are two operators that bring you nearer to your goal, and one brings you much closer than the other, it's worth trying that one first. An evaluation function will enable you to make that choice.

But it may be that you aren't quite in a condition to do so. I write this about an hour before I go to London for a talk. There's a choice of operators I can apply to change my current state (`here in Oxford`) to my desired state (`in London`). These include `walk` and `take the train`. The more powerful is `take the train`; but my state now is not quite right for that, because I'm in the Computing Service, and I can only take the train from the station. The answer is for me to walk to the station first. But a heuristic searcher would only find that answer after a lot of trial and error. What we want is a program that can *plan* intelligently — something that says to itself "Taking the train is a good choice; I'll accept that; now, I must set myself the subgoal of getting him from the Computing Service to the station". And that's what the General Problem Solver GPS does.

GPS was one of the first programs to show this kind of goal-directed behaviour. I think it worth studying for this, and other reasons:

- It exemplifies what AI people mean when they talk about formalising a problem.
- You may remember from the histories in the 1966 Scientific American article and Chapter 1 of Semantic Information Processing that there were two approaches to AI. One was modelling human behaviour; the other

was solving problems intelligently, without regard to how humans behave. GPS started off as a model of behaviour: Newell, Simon, and Shaw (amongst others in that group) were mainly concerned with cognitive modelling. And have been since: hence their later production-system models, which we do next week. In its later years though (1965 onwards ?), GPS was no longer regarded as a cognitive model. So GPS is a useful discussion point for AI as cognitive modelling vs AI as engineering.

- Wholly wonderful though GPS is, it's also the end of an era. It's the final program in the let's-find-a-generally-intelligent-algorithm approach, where everything could be regarded as a problem to be solved by the same method. Its limitations became obvious, and attention shifted from generally intelligent algorithms (dependent on embodying intelligence in the algorithm) to task-specific intelligence (intelligent use of task-specific knowledge). Expert systems are the best known example of the latter.

2 Essay and main reading

This week, I want you to write an essay on "How does the General Problem Solver work? What are its limitations, both as a model of human problem-solving, and as an engineering tool for solving problems in general?"

For the reading, start by working through sections 5.1 to 5.6 of the OU book. You can omit, or skim, section 5.5 (and-or trees) if you want; it's not relevant to how GPS works, though it does show how you can represent its behaviour.

You may also find useful the overview of GPS in *AI and Natural Man* by Boden, pp 354-356.

Unfortunately, the OU book doesn't give any examples of how to present problems to GPS. For this, read Winston Chapter 5, pages 146-156. Note on page 153 how GPS can be regarded as search. Note on pages 153-154 how it might be adapted to plan more efficiently by deferring detail. GPS itself was never used like this, though a sucessor called ABSTRIPS was.

These are adequate to tell you how GPS works. To put it in context, to find out why anyone should want to write such a program, read *GPS, A Program that Simulates Human Thought*, by A Newell and H A Simon, in *Computers and Thought* edited by E A Feigenbaum and J Feldman (RSL: per 19668 d 335; Psychology: KH:F.). Warning: what I called states, the authors call "objects".

As you read, ask yourself whether the protocol-taking technique really tells us what the subjects were doing? What does your knowledge of psychology have to say on this point?

For a critique of GPS, both as a cognitive model and as a tool for problem-solving, see Boden, *Computer Models of Mind*, pp 152-154.

If you want to find out more about GPS, you could have a look at *GPS - a case study in generality and problem solving*, by G W Ernst and A Newell

(RSL: 1878 d 142). You won't need it for the essay, or for Finals; I mention it just for interest. Note that parts of this are rather hard to read, because a lot of the general discussion is interspersed with terminology concerning GPS' internal workings. It does contain a lot of examples of how GPS was used, and the introduction discusses the motivation for writing GPS.

- The Introduction: for motivation.
- Chapter II, the section "An Approach to Generality": for its subsections on External Representation, Internal Representation, and Techniques. The first discusses how to present problems to a solver. The second is about what properties an internal representation should have. In the third, read the section on "Heuristic Search" for a summary of GPS and related methods. Note that the authors use "Heuristic Search" to mean search with operators, GPS-style; not evaluation-function search as the OU book does.
- Pages 91,92,249,250 for more comments on what a good representation should do for GPS. This may be useful to add to what we did in the Evans tutorial.
- The Summmary, pages 247-250.
- If you want examples of how problems can be formalised in terms of operators and states (something rare elsewhere in descriptions of GPS), then read the pages listed below. Because of the stuff about GPS internal workings, these may not be too easy to follow. At least try to see what operators are being used for each problem; how they reduce differences; how states are represented. Like the Computers and Thought paper, what we call states, the authors call "objects".
  - Page 65.
  - Pages 150-156: Hanoi.
  - Pages 181-184: Father and Son.
  - Pages 189-192: Monkey puzzle.
  - Pages 194-197: 3 coins.
  - Pages 200-206.
  - Pages 209-210: Bridges of Koenigsberg.
  - Pages 221-226: Water jugs.

3 Further reading

3.1 History

Machines Who Think (Psy) by McCorduck. Read the start of chapter 10, up to the end of the section on GPS. Note especially the footnote on page 212.

3.2 Planning

To get a more general idea of what planning is about, and where GPS fits in:

AI and Natural Man, the chapter on Problem Solving, and index references to GPS, Planning, Noah and Hacker.

Article on *Planning* in the *Encyclopaedia of AI* (open shelves RSL). You won't understand many of the technicalities; but try to see what kinds of problem planning research has tackled since GPS.

My tutorial notes on *Planning* (a copy is in the library).

Finally, come back to real life and everyday problems like catching a train, putting on a play, buying your food, or writing an essay. Imagine a robot doing them. How much would GPS and later planning work contribute? How far is there to go?