Ritter, F. E., Jones, G., & Young, R. M. (1997). Report on Tutorial 1: Introduction to the Soar cognitive architecture. AI and Simulation of Behaviour Quarterly, 98, 50. AISB QUARTERLY

Introduction to the SOAR Cognitive Architecture

Report by Frank Rittter and Richard Young, update also by Gary Jones

1994 Report

We presented a one day tutorial at the AISB Spring Workshop series that covered the fundamentals of the Soar architecture. Soar is a cognitive architecture built around multiple problem spaces for representing knowledge and is implemented as a production system. It includes a simple built-in learning mechanism called chunking. Soar has been proposed as an exemplar and candidate unified theory of cognition. We presented the tutorial to eight people. Two lecturers came along to further their education, and the remainder were typically graduate students who were evaluating Soar to use in their thesis.

The tutorial was based on a mixture of lectures and practical exercises that required tutees to run and modify two Soar models. The morning was spent with a verbal introduction to Soar, and six exercises modifying a simple model that ate when it was hungry, drank when it was thirsty, and learned how to choose between them. In the afternoon, tutees solved several exercises modifying a much more complete Soar model for simple analogy recently presented at the CHI'94 conference.

Tutees performed the exercises in pairs on Macintosh Powerbooks. While we don't have any hard data, our feelings and the tutees' reports indicate that this worked well. We believe that pairs are important to reduce the cognitive load of simultaneously learning a new language and a new interface. The compact size, integrated keyboards and efficient display, facilitated movement about the room, made interaction less fussy, and allowed lectures to proceed directly and reference the screen naturally. We hope to continue using this format in the future, but Powerbooks are not yet sufficiently commonplace that tutees can routinely bring their own.

We believe that the hands-on portions of the tutorial are quite important. The largest lesson tutees learn is the ability to understand the process of cognitive modelling in Soar. In particular, the tutees walked away understanding the phrase 'what Soar will do with a task'. This does not mean that Soar can pick up a task and do it, but that an analyst/programmer will have to write productions (knowledge) that guide the processing steps built into the architecture, even if it's not possible to learn how to do this fluently in eight hours. This understanding is not available from a lecture.

Even so, tutees don't always think they have learned enough. Each time we have given this tutorial, the tutees have suggested that the tutorial last longer. We think that this is a good sign, for each time we have included more and remained coherent. Tutees have decried the inability to walk out, after four or eight hours, creating Soar models. Soar remains a complex cognitive modelling approach. If you don't need what it offers (e.g. ubiquitous and robust learning, the ability to incorporate up to a million rules), you are best choosing another modelling tool. If you do need the set of facilities that it provides, you have little choice. There are a variety of current research efforts to make Soar more easily usable.

Being part of a series of concurrent workshops did provide us with an opportunity to meet and interact with attendees from other workshops and viewpoints. This was useful in keeping cognitive architectures, and ours in particular, in perspective.

1996 Update

A few things have changed in the last two years, but the material, because it is an introduction, has remained basically the same. The Soar architecture has been re-released several times. The tutorial has been updated to run with the latest release (now release 7). The material has been presented several times, which we believe has improved it each time, including twice in classes at Nottingham.

The tutorial material has been translated into HTML documents and expanded from the outline format previously used. This allows the material to be more easily presented stand-alone, or at remote sites. Using HTML also allowed us to provide more contextual help, that is, the user can click on the material they are reading for further help or manual pages. A pointer to this is available from Frank Ritter.

In previous versions, the tutorial has been improved by polishing the explanations and attempting to create better examples or better sequences of examples. Although these areas are important and can certainly be improved, we believe that the most important improvements in teaching Soar will come from improved displays. We have created a graphical user interface for Soar, which we are currently using ourselves. The displays were introduced briefly at the end of the last University of Nottingham course, and were well received by the students. When the tutorial is offered next at the first European Workshop on Cognitive Modelling, we plan to use the displays extensively.

Frank Ritter and Gary Jones, Psychology Dept., Nottingham University, England, UK

Richard Young, MRC Applied Psychology Unit, Cambridge, England, UK