

Chapter 8

Performance demonstration II: Use of Soar/MT components by others

While the environment is integrated, its components have been developed separately. As each component became available, it was spun off for use by others performing subsets of the tasks involved in model testing. The number of users of each tool, their comments, or both, provided feedback on how the various tools help perform (Tesler, 1983) specific tasks of model testing. Together they provides an estimate of the current and potential impact of the whole environment.

Spa-mode has had no use outside of this thesis. As noted earlier, the total environment, but for the displays, was used by V. Peck to perform two episodes of the Browser-Soar reanalysis. The underlying Dismal spreadsheet has had three to four additional users. It still has many problems, so a survey probably will not point out inadequacies not already known.

A survey was conducted of Soar users to find the strengths and weaknesses of the Developmental Soar Interface (DSI).

The other pieces of software either are not used by enough users (Spa-mode, Dismal), or they are so widely used that undertaking a survey is a more serious proposition (S-mode) than can be undertaken as part of this work. Portions of the DSI should no longer be considered pieces of developmental software, for out of the 60 Soar users responding to the survey, two-thirds now use some portion of it every time they use Soar.

8.1 Usage of the Developmental Soar Interface to develop Soar models

The three modules of the DSI (Soar-mode, Taql-mode, and the SX graphic display), have been through several releases. How to obtain them is explained in Appendix I. One or more of the modules are installed at each of the four principle Soar sites in the US, and at sites in Germany and the Netherlands, with over 40 researchers using one or more of the modules.

In the Fall of 1992, a survey (included as an appendix to this chapter) was sent to members of the Soar community identified through the Soar project's mailing lists, workshop attendance lists, and presenters at workshops, as most likely to use Soar in a routine way. In addition to the users directly targeted, an announcement of the survey was emailed to the general Soar mailing list, and an announcement was made at the Soar XI workshop in October, 1992.

Out of the 69 potential users identified, 63 returned a survey (a 92% response rate). The three people who never actually used Soar were dropped from later analyses. If users that were personally known did not fill in an item, or misidentified a portion of the DSI, this was corrected. Of the people responding, 50 are current members of the Soar community, and 13 are former members.

Table 8-33 shows a listing of the usage patterns. The columns list the components used, with each row representing a single user. The rows are grouped by the sets of components used. The primary tool used is Soar-mode, with 37 of the 60 users reporting using it. The SX graphic display has only been used as a routine tool for debugging by its developer and two other users, but 14 people have used it to create pictures of Soar models and to give demonstrations of their models. Taql-mode has been used and put aside by several people as they became more familiar with the TAQL grammar.

In users' responses of why they did not use additional modules, the largest number of responses (14) was that they did not use TAQL, so they did not need Taql-mode. (This would not necessarily translate into 14 users if they used TAQL.) The next largest concern (12) noted problems with installation and not knowing how to use the tools. Speed (5) was also a concern, and this concern was not limited just to the graphic display, a few users thought that Soar-mode and Taql-mode were slow to load. Most potential users of the SX graphic display were put off by how much it slowed down the system, and while only half the users reported dissatisfaction with its speed, this does not mean that

Table 8-33: Survey responses categorized by usage pattern.
Each row represents a user. Totals do not include "tried" users.

<u>Components used</u>	<u>Frequency of usage</u>			<u>Taql-mode</u>	<u>Totals</u>
	<u>Soar</u>	<u>Soar-mode</u>	<u>SX</u>		
<u>EVERYTHING</u>	daily	daily	Weekly	daily	7
	daily	daily	special	weekly	
	daily	daily	special	daily	
	daily	daily	daily	special	
	weekly	weekly	weekly	monthly	
	weekly	weekly	special	special	
	weekly	weekly	special	daily	
<u>SX & SOAR-MODE</u>	daily	daily	weekly		8
	daily	daily	special	tried	
	daily	daily	special		
	daily	daily	special		
	daily	daily	special		
	weekly	weekly	special		
	weekly	weekly	weekly	tried	
<u>TAQL-MODE</u>	daily	tried		daily	1
<u>SOAR- & TAQL-MODE</u>	daily	daily		daily	6
	daily	daily	tried	daily	
	daily	daily		daily	
	daily	daily		daily	
	weekly	weekly		weekly	
<u>SOAR-MODE</u>	monthly	monthly		monthly	17
	daily	weekly	tried		
	daily	daily		tried	
	daily	daily			
	daily	daily			
	daily	daily			
	daily	daily			
	daily	daily			
	daily	daily			
	weekly	weekly		tried	
	weekly	weekly			
	weekly	weekly			
	weekly	weekly			
	weekly	weekly			
	monthly	monthly			
monthly	monthly	tried	tried		
<u>NOTHING</u>	daily	tried		tried	21
	daily	tried			
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	daily				
	weekly				
	weekly				
	weekly				
monthly	tried		tried		
monthly					
quarterly					
na					
na					
<u>Totals</u>	60	38	15	14	60

they were satisfied with it. There were no underlying problems reported with the metaphor, representations, and manipulation of the problem space level objects.

Other users had problems with the underlying systems that the tools were built on. Several users (4) reported that they did not have a machine that could run the X window system, and some users (2) did not know or want to learn Emacs. A few users, perhaps four or five, use a Macintosh exclusively, or nearly exclusively, and the current environment is unavailable to them.

While only two respondents had not heard of all the software, a few were misinformed. One user did not know that they were using Soar-mode (but loaded it in their startup files), and one did not know that they were using Taql-mode (but when reporting useful Soar-mode features included a feature only in Taql-mode).

Use in video productions. The SX graphic display has been used to make three videos of Soar and Soar models that have been shown outside of CMU. A 20 minute tape of NTD-Soar was shown at a NASA contractors' meeting and as part of a research talk at Queen Mary & Westerfield College, both in the Spring of 1992. A 2 minute video showing the basic interaction method with the DSI and how Soar uses the Garnet toolkit has been shown four times: at the CHI '91 Garnet Special interest group meeting, at the CHI '92 Doctoral Consortium, May 1992, and as part of research talks at the Applied Psychology Unit in Cambridge, England and at Queen Mary & Westerfield College in the Spring of 1992.

Work is underway to create an introductory video explaining Soar (Newell, P., et al., forthcoming). This video is a demonstration of what will be a general capability to take a graphic description of Soar models and create high quality graphic output suitable for commercial broadcast. The initial depictions of the Soar model were created with the SX graphic display and then sent to a commercial computer graphics company for visual enhancement. The project is expected to be completed in the Spring of 1993.

Impact of the DSI on the next release of Soar software: Soar6. In the next release of the Soar software, called Soar 6, several of the features of the DSI have been incorporated or have encouraged the Soar 6 developers to include similar features. These include a very customizable trace, hooks for interacting with Soar-mode, and a better command line interpreter. Soar 6 is still under development; given time, we hope to migrate additional features to Soar 6, such as the ability to display the match set continuously, and the ability to provide a display of which productions will fire on the next decision cycle.

8.2 Usage of S-mode to create functions in S

S-mode has been distributed through three sources that make its total usage hard to compute. It appears, however, to be one of the dominant ways of interacting with S. It was first placed in 1991 in the GNU-Emacs archives at The Ohio State University. This makes it available via anonymous FTP. S-mode has also been distributed via anonymous FTP from the authors' machines. The number of users who picked it up in these two ways cannot be known.

The second mode of distribution, through a statistics software mail server, allows an approximation of a lower bound. Statlib, run by Dr. Michael Meyer at CMU, is a system for distributing statistical software and datasets by electronic mail. The system keeps track of the mail requests for each piece of software and can provide a listing of who requested each piece. Since S-mode was first placed in the Statlib server, there has been 1,043 requests for it, including requests for updated versions (personal communication, M. Meyer, October, 1992).

The exact size of its distribution is confounded further by the nature of GNU-Emacs' copy protection and the nature of S-mode's installation. GNU-Emacs and S-mode are copylefted, which means that users are entitled to (and indeed legally obligated to) provide others with copies upon request, although a copying fee can be charged. How many sites have passed S-mode on would be impossible to

compute. GNU-Emacs and its extensions are installed primarily on multi-user machines and distributed file systems. Once installed, many users can use the same piece of software although on different machines. For example, S-mode has been installed with GNU-Emacs on the Andrew system at CMU (and I don't even know who installed it). Any of the approximately 5,000 Andrew users at CMU can use S-mode.

Appendix to Chapter 8: Survey distributed to Soar users

Survey on the Developmental Soar Interface

Frank Ritter

12-Oct-92

I'm writing up my thesis and would like to get a better headcount of how many people use the DSI, and how they use it. Your comments will also be used to improve the current interface and serve as background for future versions.

* How often do you use Soar?

Daily Weekly Monthly Quarterly Other (describe)

* Which of the following have you heard of and which have you used?

	Heard of		Have used	
	Y	N	Y	N
SX graphic display (triangle thingy)	Y	N	Y	N
Soar-mode	Y	N	Y	N
Taql-mode	Y	N	Y	N

* For items you've heard of, but never used, have you considered using any? Any specific reasons why you have not used them?

* Are there any features that you would like to see added to the Soar interface for programming, editing, or understanding Soar models ?

If you have not used any items, you can quit here. Thank you.

SX graphic display (triangle thingy)

* How often you use it ? (tick one and/or write in a modifying number)

Daily	Weekly	Monthly	Quarterly
Tried once or twice	Never		
Special purpose (e.g., demos, making figures; please explain)			

* If you don't use the SX graphic display, why don't you use it?

* How do you use it? (you may tick more than one)

I've only tried it.	
I use it for special debugging.	I use it for routine development.
I use it for demos.	I use it to make presentation diagrams

* How long have you used it (e.g., 3/91 to present) ?

* What are the most valuable features ?

* What are the worst problems/bugs/factors stopping you from using the SX graphic display more often?

Soar-mode, extensions to gnu-emacs for editing productions.

* How often you use Soar-mode ? (tick one and/or write in a modifying number)

Daily	Weekly	Monthly	Quarterly
Tried once or twice	Never		
Special purpose (e.g., demos, making figures; please explain)			

* How do you use Soar-mode? (you may tick more than one)

I use it for special debugging.	I use it for routine development.
I use it for demos.	I've only tried it.

* How long have you used soar-mode (e.g., 3/91 to present) ?

* What are the most valuable features of Soar-mode?

* What are the worst problems/bugs/factors stopping you from using Soar-mode more often?

* If you don't use Soar-mode, why don't you use it?

Taq1-mode (extensions to Emacs for editing TAQL constructs)

* How often do you use taql-mode ? (tick one and/or write in a modifying number)

Daily	Weekly	Monthly	Quarterly
Tried once or twice	Never		
Special purpose (e.g., demos, making figures; please explain)			

* How do you use taql-mode? (you may tick more than one)

I've only tried it.	I use it for demos.
I use it for special debugging.	I use it for routine development.

* How long have you used taql-mode (e.g., 3/91 to present) ?

* What are the most valuable features of taql-mode?

* What are the worst problems/bugs/factors stopping you from using taql-mode more often?

* If you don't use taql-mode, why don't you use it?

Additional on-line & hardcopy copies available from Frank Ritter@cs.cmu.edu

Please return surveys by email or hardcopy to Frank Ritter@cs.cmu.edu