Kennedy, W.G. (2012) Long-Term Learning in Soar. In Encyclopedia of the Sciences of Learning. Norbert M. Seel (Ed.) 1st Edition 2012. ISBN 978-1-4419-1427-9. Springer.

LONG-TERM LEARNING IN SOAR

William G. Kennedy
Krasnow Institute for Advanced Study, George Mason University
Fairfax, Virginia
United States of America
wkennedy@gmu.edu

Synonyms

Expertise
Expertise development
Long-term expertise development

Definition

Long-term learning in Soar is the process of accumulating procedural knowledge throughout the existence of an intelligent agent. In Soar, such knowledge is in the form of IF-THEN productions, or rules, called "chunks". The learned chunks are new rules capturing the results of resolving obstacles in the reasoning process. This long-term knowledge is maintained by the system with the expectation that it will be useful during the existence of the intelligent agent. Declarative memory is not part of Soar's long-term memory.

Theoretical Background

Allen Newell, through his book, "Unified Theories of Cognition" (Newell 1990), proposed many partial theories of cognition, both abstract and human, and offered the Soar architecture as a candidate unified theory of cognition. He defined a unified theory of cognition as "a single set of mechanisms for all of cognitive behavior". Three parts of his theory are applicable to long-term learning in Soar: symbols, memory, and learning. First, he, with Herbert A. Simon, theorized that humans were symbol-processing systems and that symbols were representations of knowledge that are useful in the process of storing and retrieving that knowledge. Next, the consensus was and still is that there are two separable kinds of memory. The first is declarative memory, which is the collection of the facts including processing status used in problem solving. The collection of this information forms the working memory of the system and he proposed this knowledge was the system's short-term memory. The other accepted form of memory is procedural memory, which is made up of the IF-THEN rules or productions that use the declarative knowledge to develop additional declarative knowledge and to take actions to solve the overall problem. Then, as part of the theoretical basis of Soar, he proposed that long-term procedural memory could be considered as a single production system. Although he acknowledged that episodic and semantic memories were the next most commonly proposed separate memory structures, he proposed "trying to live with" the single representational system.

Newell also proposed one form of learning, called chunking. The implications of the chunking theory of learning were twofold. The Soar qualitative theory of learning included the premise that learning occurs at an approximately constant rate and that the productions, which make up the long-term memory, are maintained permanently. He proposed no mechanism to remove learned knowledge, i.e., forgetting. Therefore, Soar, as a unified theory of cognition, includes a monotonically increasing long-term memory associated with its constant learning process, when learning employed. However, many models using Soar do not employ learning and meet their performance goals using only productions that are initially loaded into the system.

Constantly increasing long-term memory results in performance problems, both theoretical and practical, and there are other important research questions. For practical purposes, the most common approach by Soar users is to pre-load the system with all the rules necessary to solve the intended scope of presented problems and to not have the agent learn at all.

Important Scientific Research and Open Questions

Long-term learning in Soar, as well as other symbolic learning and monotonically increasing systems, led to performance problems associated with the maintenance, retrieval, and use of the knowledge. This situation was identified as the "utility problem" (Minton 1990). Approaches to address this problem initially include restricting or stopping the learning process and improving the matching of current conditions to memory. They proved to be somewhat successful, but did not resolve the theoretical basis of the problem. Even restricting the expressiveness of the rule language was not successful because it necessitated a large increase in the number of rules necessary to represent the same amount of knowledge (Tambe, Newell, & Rosenbloom 1990). Another approach is to begin with learning but to later stop learning. The conditions for which this approach is appropriate have not been explored for the Soar system specifically. Improvements in the matching process have extended practical performance of Soar to beyond 100,000 rules (Doorenbos 1993), but because there is still growth in the match costs with additional chunks, it does not resolve the theoretical side of the problem.

Removal of previously learned knowledge, i.e., forgetting, appears to be necessary to address the theoretical basis of the utility problem. Forgetting based on the lack of recent use has been explored in Soar (Kennedy & De Jong 2003) and in both Soar and ACT-R, the other important cognitive architecture (Kennedy & Trafton 2007). Forgetting based on other characteristics of the long-term memory, the agent's goals, or environment, have yet to be explored.

The Soar system continues to be the basis of active research and to be a widely applicable real world problem. Recent activities have added episodic, semantic, and spatial memories to its architecture (Laird 2008). What the long-term effects are on the theoretical and practical performance of long-term learning in Soar with these new forms of long-term memory are open research questions.

Cross-References

- Soar - learning
- Chunking mechanisms and learning
- Learning by chunking
- Matching
- Long-term expertise development
- Expertise

References

Doorenbos, R. (1993). Matching 100,000 Learned Rules. *Proceedings of the Eleventh National Conference on Artificial Intelligence* (pp. 290-296). Menlo Park, CA: AAAI Press.

Kennedy, W.G. & De Jong, K.A. (2003) Characteristics of long-term learning in Soar and its application to the utility problem. *Proceedings of the Twentieth International Conference on Machine Learning.* (pp. 337-344). Menlo Park, CA: AAAI Press.

Kennedy, W.G., & Trafton, J.G. (2007) Long-term symbolic learning. *Cognitive Systems Research*, *8*, 237-247.

Laird, J.E. (2008) Extending the Soar cognitive architecture. In *Artificial General Intelligence 2008: Proceedings of the First AIG Conference*. ISO Press.

Minton, S. (1990). Quantitative results concerning the utility of explanation-based learning. *Artificial Intelligence*, *42*, 363–391.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Tambe, M., Newell, A., & Rosenbloom, P.S. (1990) The problem of expensive chunks and its solution by restricting expressiveness. *Machine Learning, 5,* 299-348.