

Instructo-Soar: Learning from interactive natural language instructions (video abstract)*

Scott B. Huffman and John E. Laird

Artificial Intelligence Laboratory
The University of Michigan
Ann Arbor, Michigan 48109-2110
huffman@umich.edu

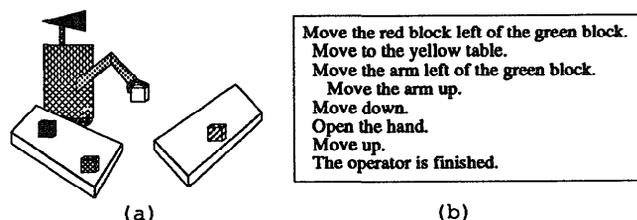


Figure 1: (a). An agent in an initial situation; (b) Instructions to teach a new procedure.

Despite its ubiquity in human learning, very little work has been done in artificial intelligence on learning from natural language instructions. In this video, we present a system, Instructo-Soar, that can both behave and learn from natural language instructions. The system is described in papers elsewhere [Huffman and Laird, 1993a; Huffman and Laird, 1993b]. The type of instruction we particularly address is *situated, interactive* instruction. Situated means that the student is within the task domain, attempting to perform tasks, when instruction is given. Interactive means that the student can request instruction as needed.

Instructo-Soar can learn completely new procedures from sequences of interactive instruction, and can also learn how to extend its knowledge of previously known procedures to new situations. The video demonstrates its application in a simple robotic domain. The system starts with a small set of primitive operators. Given instructions in the form of imperative natural language sentences, it is able to learn a hierarchy of complex operators. An example instruction scenario is shown in Figure 1.

Learning procedures from instructions involves more than simply memorization of instruction sequences. Acquiring a new procedure involves learning both the procedure's goal concept, and a general implementation for the procedure.

The instructed agent can learn the goal concept of a

new procedure after performing it (an inductive learning task). Instructo-Soar uses a simple difference-of-states heuristic to induce goal concepts; everything that has changed from the initial state to the final state during execution of the new procedure is considered part of the goal of the procedure. Recent versions of the system allow the instructor to give instructional feedback to alter the induced goal as needed.

To learn a general implementation for the procedure, the applicability conditions of each instruction in the implementation sequence must be determined. Instructo-Soar uses an explanation based approach for this: the agent attempts to explain to itself (via an internal forward simulation) how each instruction leads to achievement of the goal. This explanation process indicates which features of the situation and instruction are crucial for goal achievement.

Instructo-Soar exhibits a multiple execution learning process to learn a new procedure. Initial learning is rote and episodic in nature. After executing the new procedure the first time, the system can induce the goal concept of the procedure. During future executions, the system recalls the instructions it learned by rote initially, and explains how they contribute to reaching the procedure's goal, resulting in general learning. The learning curve that results closely matches the power law of practice.

This work represents first steps towards our long-term goal of building general, instructable autonomous agents.

References

- [Huffman and Laird, 1993a] Scott B. Huffman and John E. Laird. Learning from instruction: A knowledge-level capability within a unified theory of cognition. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, 1993.
- [Huffman and Laird, 1993b] Scott B. Huffman and John E. Laird. Learning procedures from interactive natural language instructions. In P. E. Utgoff, editor, *Machine Learning: Proceedings of the Tenth International Conference*, June 1993.

*This research was sponsored by NASA/ONR under contract NCC 2-517.