

The Soar Papers: Research on Integrated Intelligence

Paul S. Rosenbloom, John E. Laird, & Allen Newell (Editors)

A Book Proposal

August 8, 1990

The Soar project is an attempt to develop and apply a unified theory of natural and artificial intelligence. At the core of this effort is an investigation into the *architecture* — the fixed base of tightly-coupled mechanisms — underlying intelligent behavior. This architecture then forms the basis for wide-ranging investigations into basic intelligent capabilities — such as problem solving, planning, learning, knowledge representation, natural language, perception, and robotics — as well as applications in the areas of expert systems, and psychological modeling.

Since the project's official inception in 1983, it has grown from a three-man effort to a distributed, interdisciplinary community of more than seventy researchers. During this period, the Soar community has produced over one hundred articles and books on aspects of Soar and intelligence. The principal motivation for this new book is to bring together a nearly-comprehensive set of readings on Soar in a form that can be quickly, cheaply, easily and widely disseminated. It is hoped that such a book would not only be of interest to the Soar community, but also to significant segments of the broader context of artificial intelligence — particularly in such subareas as machine learning and intelligent agents — cognitive science, and cognitive psychology.

Towards this end we have selected out a set of sixty three articles for inclusion. This set includes articles previously published in journals, conferences, workshops, and books, as well as some technical reports and unpublished articles. There is thus a mixture in terms of size, quality — though we have eliminated the low quality articles — and accessibility. Included in the total are six articles that, though they are not about Soar, are direct precursors to it (and are by project researchers). Several of these are landmark papers that have languished in relative obscurity because of where they were originally published.

Due to the nature of the material covered by this book, it is structured in an manner that is atypical for a collection of readings. The standard approach is to divide the articles into chapters according to each article's principal contribution. The index, or discursive front matter, is then used to deal with hidden secondary contributions. When attempts were made to impose such a structure on this book, they failed. The

problems were that: (1) because Soar covers so many aspects of intelligence within a single system, too many of the articles make principal contributions in more than one area (in addition to secondary contributions); and (2) because the articles are all about a single system, there is a logical flow of the material that gets lost under such a partitioning. The book is therefore organized around chronological chapters – within years, the ordering is alphabetic by author – with a preceding topic hierarchy. Each major heading in the topic hierarchy corresponds to a chapter-level topic in a more traditional organization; however, here there is an opportunity to provide multiple organizations over the material, each corresponding to a viewpoint on Soar and intelligence. Each terminal element in the topic hierarchy is annotated with pointers – of the form *number:year* – to a logical sequence of articles directly relevant to the topic. The intent is that readers interested in the entire picture can read the articles in publication order, while readers interested only in a specific topic can go directly to a logical sequence of papers to read on that topic.

In addition to the topic hierarchy and the chronological readings, the book will include a detailed index, a complete bibliography of Soar writings – the topic hierarchy will be augmented to include all of these writings, with included articles being structurally distinguished from others – and an introductory essay of approximately 10-20 pages that sets the appropriate context for the project and system as a whole and, in somewhat more detail, for the individual topics in the hierarchy.

A rough page count for the book gives approximately 520 small-format pages and 730 large-format pages for the actual articles. If the small-format pages are published two per physical page, this yields approximately 885 pages, not counting front matter and indices. We are open to suggestions about shrinking this total, but prefer to shrink it as little as possible.

In selecting the set of sixty three articles we have attempted to eliminate any article that does not make a principal contribution on its own, independent of the other articles. This still leaves some overlap among the articles, but it is primarily limited to the introductory descriptions of Soar that, by necessity, occur in nearly all of the articles. Many of these descriptions actually provide distinct ways of viewing Soar – for example, as a hierarchy of cognitive levels, or as a general goal-oriented system – and thus maintain some degree of independent utility, while the remainder tend to be short. We plan on explicitly dealing with this via a note at the front of the book.

This new book should complement the two existing books on Soar. The first book – *Universal Subgoalting and Chunking: The Automatic Generation and Learning of Goal Hierarchies*, published by Kluwer Academic Publishers in 1986 – consisted of the two theses out of which the Soar project grew, plus one short article on their synthesis. The

two theses are not included in this new book, and predate most of the work to be included. The first book sold out two printings, and is now out of print. The second book — *Unified Theories of Cognition*, to be published by Harvard University Press in 1990 — is a focused treatment of Soar as a model of human cognition.

The remainder of this proposal contains the topic hierarchy, as currently developed, and the chronological list of articles to be included. Both of these may undergo further refinement before final publication.

Topic Hierarchy

1. Direct precursors
 - a. Weak problem solving methods [1:69]
 - b. Problem spaces [1:80]
 - c. Production systems [4:87]
 - d. Knowledge level [1:82]
 - e. Chunking [1:81, 4:87, 4:86]
2. Foundations
 - a. Cognitive architecture [7:89]
 - b. The Soar architecture [2:87, 2:91]
 - c. The knowledge level [7:90, 10:90]
 - d. The problem space computational model [7:90]
3. Levels of cognition
 - a. Knowledge level [1:82, 7:90, 10:90, 3:87, 10:89, 6:88, 8:90]
 - b. Problem space level [1:80, 7:90, 11:88, 15:89]
 - c. Symbol level [7:89, 9:89]
 - d. Neural level [9:89]
4. Mechanisms
 - a. Overview [2:87]
 - b. Problem spaces [1:80, 7:90]
 - c. Goals
 - d. Production systems
 1. Match algorithms [5:86, 4:88]
 2. Parallel implementation [9:88, 2:88]
 3. Alternative memory models [12:90, 13:90, 2:88, 9:89]
5. Capabilities
 - a. Problem solving
 1. Weak methods [1:69, 1:83]
 2. Abstraction [7:87, 13:89, 14:90, 9:90]
 3. Generic tasks [4:89]
 4. Planning [3:90, 9:90]
 5. Analogy and case-based reasoning [1:87, 12:89, 14:89]
 6. Time-limited problem solving [1:90]
 7. Mental models [5:88, 8:89, 5:89]
 8. Loop detection [1:90]
 - b. Learning
 1. Overview [1:86, 6:87]
 2. Skill acquisition [1:81, 4:87, 4:86, 1:84, 10:89, 2:89]
 3. Explanation-based learning [3:86]
 4. Knowledge acquisition [3:87, 10:89, 6:88, 10:90, 8:90, 6:90, 1:90]
 5. Learning from external guidance [1:87, 2:90]
 6. Inductive generalization [8:90, 6:90, 13:89, 5:90]
 7. Task acquisition [11:88, 5:89]
 8. Incorrect knowledge [2:86, 3:88, 2:90, 1:90]
 9. Expensive chunks [12:90, 13:90]

- 10. Integrated learning [6:87, 10:89, 6:88, 8:90]
- c. External interaction
 - 1. External guidance [1:87, 2:90]
 - 2. Instruction taking [11:88, 5:89]
 - 3. Natural language [4:90, 5:89, 5:90]
 - 4. Perception [15:90]
 - 5. Robotics [2:90, 3:90]
 - 6. Interruption [3:90]
- 6. Domains
 - a. Engineering design [6:89]
 - b. Algorithm and software design [5:87, 8:88, 12:89]
 - c. Computer configuration [1:85, 2:87, 7:87, 13:89, 15:89]
 - d. Medical diagnosis [10:88]
 - e. Factory scheduling [3:89]
 - f. Tool integration [11:90]
 - g. Human-computer interaction
 - 1. Browsing [5:90]
 - 2. Text editing [16:90]
 - h. Immediate reasoning
 - 1. Categorical syllogisms [5:88, 8:89]
 - 2. Relational reasoning [8:89]
 - 3. Wason selection task [8:89]
 - 4. Conditional reasoning [8:89]
 - i. Mathematics
 - 1. Algebra [1:87]
 - 2. Subtraction [2:91]
 - j. Puzzles and games
 - 1. Eight puzzle [2:87, 1:86, 1:84, 14:90]
 - 2. Tic-Tac-Toe [1:84]
 - 3. Block stacking [2:90, 3:90]
 - 4. Missionaries and Cannibals [3:88]
 - 5. Grid searching [12:90, 13:90]
 - k. Music [14:89]
 - l. Navigation [13:89, 14:90, 3:90]
- 7. Psychological modeling
 - a. Overview [5:90]
 - b. Immediate reasoning [5:88, 8:89, 5:89]
 - c. Strategy change [11:89]
 - d. Human-computer interaction [5:90, 16:90]
 - e. User modeling [16:90]
 - f. Perception [15:90]
 - g. Series completion [5:90]
 - h. Practice [1:81, 4:87, 4:86, 10:89]
 - i. Lexical acquisition [6:90]
 - j. Prototypes [6:90]
 - k. Verbal learning [3:87, 10:89, 10:90]

1. Unified theory of cognition [5:90]
8. Perspectives
 - a. Hierarchy of cognitive levels [2:91]
 - b. Knowledge level system [10:90, 7:90]
 - c. Meta-level (reflective) system [7:88]
 - d. General goal-oriented system [9:89]
 - e. Hybrid planning system [9:90]
 - f. Problem space system [7:90]
 - g. Symbol system [7:89, 2:87]
 - h. Neural network (connectionist system) [9:89]
 - i. Explanation-based learner [3:86]
9. Using Soar [15:89]
10. Comparisons [2:89]
11. Commentaries [1:91, 1:89]
12. Soar Group Sociology [1:88]

Included Articles

Prior to 1983

- 1:69. Newell, A., "Heuristic programming: Ill-structured problems", in *Progress in Operations Research, III*, Aronofsky, J., ed., Wiley, New York, 1969, pp. 360-414.
- 1:80. Newell, A., "Reasoning, problem solving and decision processes: The problem space as a fundamental category", in *Attention and Performance VIII*, R. Nickerson, ed., Erlbaum, Hillsdale, N.J., 1980.
- 1:81. Newell, A. & Rosenbloom, P. S., "Mechanisms of skill acquisition and the law of practice", in *Cognitive Skills and their Acquisition*, J. R. Anderson, ed., Erlbaum, Hillsdale, NJ, 1981, pp. 1-55.
- 1:82. Newell, A., "The knowledge level", *Artificial Intelligence*, Vol. 18, 1982, pp. 87-127.

1983 - 1985

- 1:83. Laird, J. E. & Newell, A., "A Universal Weak Method", Tech. report #83-141, Carnegie-Mellon University Computer Science Department, June 1983.
- 1:84. Laird, J. E., Rosenbloom, P. S., & Newell, A., "Towards chunking as a general learning mechanism", *Proceedings of the National Conference on Artificial Intelligence*, AAAI, Austin, 1984, pp. 188-192.
- 1:85. Rosenbloom, P. S., Laird, J. E., McDermott, J., Newell, A. & Orciuch, E., "R1-Soar: An experiment in knowledge-intensive programming in a problem-solving architecture", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 7, 1985, pp. 561-569.

1986

- 1:86. Laird, J. E., Rosenbloom, P. S., & Newell, A., "Chunking in Soar: The anatomy of a general learning mechanism", *Machine Learning*, Vol. 1, 1986, pp. 11-46.
- 2:86. Laird, J. E., Rosenbloom, P. S., & Newell, A., "Overgeneralization during knowledge compilation in Soar", *Proceedings of the Workshop on Knowledge Compilation*, T. G. Dietterich, ed., AAAI/Oregon State U., Otter Crest, 1986, pp. 46-57.
- 3:86. Rosenbloom, P. S., & Laird, J. E., "Mapping explanation-based generalization onto Soar", *Proceedings of the Fifth National Conference on Artificial Intelligence*, AAAI, Philadelphia, 1986, pp. 561-567.
- 4:86. Rosenbloom, P. S. & Newell, A., "The chunking of goal hierarchies: A generalized model of practice", in *Machine Learning: An Artificial Intelligence Approach, Volume II*, R. S. Michalski, J. G. Carbonell, &

T. M. Mitchell, eds., Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1986, pp. 247-288.

- 5:86. Scales, D. J., "Efficient Matching Algorithms for the Soar/Ops5 Production System", Tech. report KSL-86-47, Knowledge Systems Laboratory, Department of Computer Science, Stanford University, 1986.

1987

- 1:87. Golding, A. R., Rosenbloom, P. S., & Laird, J. E., "Learning general search control from outside guidance", *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, IJCAI, Milan, 1987, pp. 334-337.
- 2:87. Laird, J. E., Newell, A., & Rosenbloom, P. S., "Soar: An architecture for general intelligence", *Artificial Intelligence*, Vol. 33, 1987, pp. 1-64.
- 3:87. Rosenbloom, P. S., Laird, J. E., & Newell, A., "Knowledge level learning in Soar", *Proceedings of Sixth National Conference on Artificial Intelligence*, AAAI, Seattle, 1987, pp. 499-504.
- 4:87. Rosenbloom, P. S. & Newell, A., "Learning by chunking: A production-system model of practice", in *Production System Models of Learning and Development*, D. Klahr, P. Langley, R. Neches, eds., Bradford Books/MIT Press, Cambridge, MA, 1987, pp. 221-286.
- 5:87. Steier, D., "Cypress-Soar: A case study in search and learning in algorithm design", *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, IJCAI, Milan, 1987, pp. 327-330.
- 6:87. Steier, D. M., Laird, J. E., Newell, A., Rosenbloom, P. S., Flynn, R., Golding, A., Polk, T. A., Shivers, O. G., Unruh, A., & Yost, G. R., "Varieties of Learning in Soar: 1987", *Proceedings of the Fourth International Workshop on Machine Learning*, P. Langley, ed., Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1987, pp. 300-311.
- 7:87. Unruh, A., Rosenbloom, P. S., & Laird, J. E., "Dynamic abstraction problem solving in Soar", *Proceedings of the Third Annual Aerospace Applications of Artificial Intelligence Conference*, Dayton, OH, 1987, pp. 245-256.

1988

- 1:88. Carley, K. and Wendt, K., "Electronic Mail and the Diffusion of Scientific Information: The study of Soar and its dominant users", Department of Social and Decision Sciences, Carnegie Mellon University, July, 1988. Unpublished..
- 2:88. Flynn, R., "Placing Soar on the Connection Machine", Digital Equipment Corporation, September, 1988. An extended abstract of this paper was distributed at the AAAI Mini-Symposium, *How Can Slow Components Think So Fast?*, Spring, 1988..

- 3:88. Laird, J. E., "Recovery from incorrect knowledge in Soar", *Proceedings of the Seventh National Conference on Artificial Intelligence*, AAAI, St. Paul, 1988, pp. 618-623.
- 4:88. Nayak, P., & Gupta, A., & Rosenbloom, P. S., "Comparison of the Rete and Treat production matchers for Soar (a summary)", *Proceedings of the Seventh National Conference on Artificial Intelligence*, AAAI, St. Paul, 1988, pp. 693-698.
- 5:88. Polk, T. A. & Newell, A., "Modeling human syllogistic reasoning in Soar", *Proceedings of the 10th Annual Conference of the Cognitive Science Society*, Montreal, 1988, pp. 181-187.
- 6:88. Rosenbloom, P. S., "Beyond generalization as search: Towards a unified framework for the acquisition of new knowledge", *Proceedings of the AAAI Symposium on Explanation-Based Learning*, G. F. DeJong, ed., AAAI, Stanford, CA, 1988, pp. 17-21.
- 7:88. Rosenbloom, P. S., Laird, J. E., & Newell, A., "Meta-levels in Soar", in *Meta-Level Architectures and Reflection*, P. Maes & D. Nardi, eds., North Holland, Amsterdam, 1988, pp. 227-240.
- 8:88. Steier, D. M. & Newell, A., "Integrating multiple sources of knowledge into Designer-Soar an automatic algorithm designer", *Proceedings of the Seventh National Conference on Artificial Intelligence*, AAAI, St. Paul, 1988, pp. 8-13.
- 9:88. Tambe, M., Kalp, D., Gupta, A., Forgy, C.L., Milnes, B.G., and Newell, A., "Soar/PSM-E: Investigating match parallelism in a learning production system", *Proceedings of the ACM/SIGPLAN Symposium on Parallel Programming: Experience with Applications, Languages, and Systems*, July 1988, pp. 146-160.
- 10:88. Washington, R. & Rosenbloom, P. S., "Applying Problem Solving and Learning to Diagnosis".
- 11:88. Yost, G. R. & Newell, A., "Learning new tasks in Soar", Unpublished

1989

- 1:89. Bösser, T., "A discussion of 'The chunking of skill and knowledge' by Paul S. Rosenbloom, John E. Laird & Allen Newell", in *Working Models of Human Perception*, B. A. G. Elsendoorn & H. Bouma, eds., Academic Press, London, 1989, pp. 411-418.
- 2:89. Etzioni, O., & Mitchell, T. M., "A comparative analysis of chunking and decision-analytic control", *Proceedings of the 1989 AAAI Spring Symposium on AI and Limited Rationality*, Stanford, 1989.
- 3:89. Hsu, W., Prietula, M., and Steier, D., "Merl-Soar: Scheduling within a general architecture for intelligence", *Proceedings of the Third International*

- Conference on Expert Systems and the Leading Edge in Production and Operations Management*, May 1989, pp. 467-481.
- 4:89. Johnson, T. R., Smith, J. W. Jr., & Chandrasekaran, B., "Generic Tasks and Soar", *Working Notes of the AAAI Spring Symposium on Knowledge System Development Tools and Languages*, Stanford, CA, 1989, pp. 25-28.
- 5:89. Lewis, R. L., Newell, A., & Polk, T. A., "Toward a Soar theory of taking instructions for immediate reasoning tasks", *Proceedings of the 11th Annual Conference of the Cognitive Science Society*, Ann Arbor, MI, 1989, pp. 514-521.
- 6:89. Modi, A.K. and Westerberg, A.W., "Integrating learning and problem solving within a chemical process designer", Presented at the Annual Meeting of the American Institute of Chemical Engineers, November, 1989..
- 7:89. Newell, A., Rosenbloom, P. S., & Laird, J. E., "Symbolic architectures for cognition", in *Foundations of Cognitive Science*, M. I. Posner, ed., Bradford Books/MIT Press, Cambridge, MA, 1989, pp. 93-131.
- 8:89. Polk, T. A., Newell, A., & Lewis, R. L., "Toward a unified theory of immediate reasoning in Soar", *Proceedings of the 11th Annual Conference of the Cognitive Science Society*, Ann Arbor, MI, 1989, pp. 506-513.
- 9:89. Rosenbloom, P. S., "A symbolic goal-oriented perspective on connectionism and Soar", in *Connectionism in Perspective*, R. Pfeifer, Z. Schreter, F. Fogelman-Soulie, & L. Steels, eds., Elsevier (North-Holland), Amsterdam, 1989, pp. 245-263.
- 10:89. Rosenbloom, P. S., Laird, J. E., & Newell, A., "The chunking of skill and knowledge", in *Working Models of Human Perception*, B. A. G. Elsendoorn & H. Bouma, eds., Academic Press, London, 1989, pp. 391-410.
- 11:89. Ruiz, D. and Newell, A., "Tower-noticing triggers strategy-change in the Tower of Hanoi: A Soar model", *Proceedings of the Annual Conference of the Cognitive Science Society*, 1989, pp. 522-529.
- 12:89. Steier, D., "'But How Did You Know To Do That?": What a theory of algorithm design process can tell us", Engineering Design Research Center, Carnegie Mellon University, August, 1989, Unpublished..
- 13:89. Unruh, A. & Rosenbloom, P. S., "Abstraction in Problem Solving and Learning", *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, IJCAI, Detroit, 1989, pp. 681-687.
- 14:89. Vicinanza, S. and Prietula, M.J., "A computational model of musical creativity (Extended abstract)", Prepared for the AI and Music Workshop, Eleventh International Joint Conference on Artificial Intelligence, August 20-25, 1989..
- 15:89. Yost, G. R. & Newell, A., "A problem space approach to expert system specification", *Proceedings of Eleventh International Joint Conference on*

Artificial Intelligence, IJCAI, Detroit, 1989, pp. 621-627.

1990

- 1:90. Congdon, C.B., "Learning Control Knowledge in an Unsupervised Planning Domain", Artificial Intelligence Laboratory, The University of Michigan, February, 1990. Unpublished.
- 2:90. Laird, J.E., Hucka, M., Yager, E.S., and Tuck, C.M., "Correcting and extending domain knowledge using outside guidance", *Proceedings of the Seventh International Conference on Machine Learning*, 1990, pp. 235-243.
- 3:90. Laird, J. E., & Rosenbloom, P. S., "Integrating execution, planning, and learning in Soar for external environments", *Proceedings of the Eighth National Conference on Artificial Intelligence*, AAAI, Boston, 1990, pp. 1022-1029.
- 4:90. Lehman, J. F., Lewis, R. L., Newell, A., Tessler, S. G., & Huffman, S. B.(?), "NL-Soar: Natural language understanding in Soar (?)", Tech. report CMU-CS-90-???, Carnegie Mellon University, School of Computer Science, 1990, In preparation
- 5:90. Lewis, R. L., Huffman, S. B., John, B. E., Laird, J. E., Lehman, J. F., Newell, A., Rosenbloom, P. S., Simon, T., & Tessler, S. G., "Soar as a unified theory of cognition: Spring 1990", *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, Cambridge, MA, 1990.
- 6:90. Miller, C.S. and Laird, J.E., "A Simple, Symbolic Model for Associative Learning and Retrieval", Artificial Intelligence Laboratory, The University of Michigan, March, 1990. Unpublished.
- 7:90. Newell, A., Yost, G. R., Laird, J. E., Rosenbloom, P. S., & Altmann, E., "Formulating the problem space computational model", in *Carnegie Mellon Computer Science: A 25 Year Commemorative*, R. F. Rashid, ed., ACM Press/Addison-Wesley, 1990, In press
- 8:90. Rosenbloom, P. S. & Aasman, J., "Knowledge level and inductive uses of chunking (EBL)", *Proceedings of the Eighth National Conference on Artificial Intelligence*, AAAI, Boston, 1990, pp. 821-827.
- 9:90. Rosenbloom, P. S., Lee, S., & Unruh, A., "Responding to impasses in memory-driven behavior: A framework for planning", Submitted to the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control
- 10:90. Rosenbloom, P. S., Newell, A., & Laird, J. E., "Towards the knowledge level in Soar: The role of the architecture in the use of knowledge", in *Architectures for Intelligence*, K. VanLehn, ed., Erlbaum, Hillsdale, NJ, 1990, In press
- 11:90. Steier, D.M., "Intelligent architectures for integration", *Proceedings of the IEEE Conference on Systems Integration*, August 1990, A shorter version of

this paper will appear in the proceedings

- 12:90. Tambe, M., Newell, A., & Rosenbloom, P. S., "The problem of expensive chunks and its solution by restricting expressiveness", *Machine Learning*, 1990, In press
- 13:90. Tambe, M. & Rosenbloom, P. S., "A framework for investigating production system formulations with polynomially bounded match", *Proceedings of the Eighth National Conference on Artificial Intelligence*, AAAI, Boston, 1990, pp. 693-700.
- 14:90. Unruh, A., & Rosenbloom, P. S., "Two new weak method increments for abstraction", *Working Notes of the AAAI-90 Workshop on Automatic Generation of Approximations and Abstractions*, T. Ellman, ed., AAAI, Boston, 1990, pp. 78-86.
- 15:90. Wiesmeyer, M. and Laird, J., "A computer model of 2D visual attention and search", Tech. report, Department of Electrical Engineering and Computer Science, University of Michigan, 1990, A portion of this article appeared under the title "A computer model of 2D visual attention, in *Proceedings of the Annual Conference of the Cognitive Science Society*, 1990
- 16:90. Young, R.M. and Whittington, J., "Using a knowledge analysis to predict conceptual errors in text-editor usage", *Proceedings of the Conference on Human Factors and Computing Systems*, April 1990, pp. 91-97.

1991

- 1:91. Norman, D. A., "Approaches to the study of intelligence", *Artificial Intelligence*, 1991, In press
- 2:91. Rosenbloom, P. S., Laird, J. E., Newell, A., & McCarl, R., "A preliminary analysis of the Soar architecture as a basis for general intelligence", *Artificial Intelligence*, 1991, In press