

Designing an Integrated Architecture: The PRODIGY View *

Jaime G. Carbonell, Yolanda Gil, Robert Joseph,
Craig A. Knoblock, Steve Minton[†], and Manuela M. Veloso

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

[†] Sterling Federal Systems
NASA Ames Research Center
AI Research Branch, Mail Stop: 244-17
Moffett Field, CA 94035

ABSTRACT

Artificial intelligence has progressed to the point where multiple cognitive capabilities are being integrated into computational architectures, such as SOAR, PRODIGY, THEO, and ICARUS. This paper reports on the PRODIGY architecture, describing its planning and problem solving capabilities and touching upon its multiple learning methods. Learning in PRODIGY occurs at all decision points and integration in PRODIGY is at the knowledge level; the learning and reasoning modules produce mutually interpretable knowledge structures. Issues in architectural design are discussed, providing a context to examine the underlying tenets of the PRODIGY architecture.

1 Introduction

A common dream for many AI researches, present authors included, is the construction of a general purpose learning and reasoning system that given basic axiomatic knowledge of a domain is capable of becoming an expert problem solver.

Our machine learning approach, implemented in PRODIGY [Carbonell *et al.*, 1990], starts with a general problem-solving engine based on a possibly incomplete domain theory. The problem solver improves its performance through experience by refining the initial domain knowledge and learning knowledge to control the search process. The paper is divided into two parts. The first part describes the basic architecture, including the problem solver and the various learning modules. The second part discusses the design issues in building an integrated architecture.

*This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976, Amendment 20, under contract number F33615-87-C-1499, monitored by the Air Force Avionics Laboratory, in part by the Office of Naval Research under contracts N00014-84-K-0345 (N91) and N00014-86-K-0678-N123, in part by NASA under contract NCC 2-463, in part by the Army Research Institute under contract MDA903-85-C-0324, and in part by small contributions from private institutions. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of DARPA, ONR, NASA, ARI, or the US Government. The fourth author was supported by an Air Force Laboratory Graduate Fellowship, and the third author was supported by an AT&T Bell Labs Ph.D. Scholarship.

2 The PRODIGY Architecture

2.1 The Problem Solver

PRODIGY's basic reasoning engine is a general-purpose problem solver and planner [Veloso, 1989, Minton *et al.*, 1989] that searches for sequences of operators (i.e., plans) to accomplish a set of goals from a specified initial state description. Search in PRODIGY is guided by a set of *control rules* that apply at each decision point. Search control rules may be general or domain specific, hand-coded or automatically acquired, and may consist of heuristic preferences or definitive selections. In the absence of any search control, PRODIGY defaults to depth-first means-ends analysis. But, with appropriate search control rules it can emulate other search disciplines, including breath-first search, depth-first iterative-deepening, best-first search, and knowledge-based plan instantiation.

2.1.1 Domains of application

PRODIGY has been applied to many different domains: robotic path planning, the blocks-world, an augmented version of the STRIPS domain, matrix algebra manipulation, discrete machine-shop planning and scheduling, computer configuration, logistics planning, and several others. In order to solve problems in a particular domain, PRODIGY must first be given a specification of that domain, consisting of a set of operators and inferences rules.

2.1.2 Knowledge representation

Each operator has a precondition expression that must be satisfied before the operator can be applied, and a list of effects that describe how the application of the operator changes the world. Precondition expressions are well-formed formulas in PDL, a form of predicate logic encompassing negation, conjunction, disjunction, and existential and universal quantification. The effects are atomic formulas that describe the conditions that are added or deleted from the current state when the operator is applied. Operators may also contain conditional effects, which represent changes to the world that are dependent on the state in which the operator is applied.

2.2 Problem definition

A problem consists of an initial state and a goal expression. To solve a problem, PRODIGY must find a sequence of operators that, if applied to the initial state, produces a final state satisfying the goal expression. The search tree initially starts out as a single node containing the initial state and goal expression. The tree is expanded by repeating the following two steps:

1. **Decision phase:** There are four types of decisions that PRODIGY makes during problem solving. First, it must decide what node in the search tree to expand next, defaulting to a depth-first expansion. Each node consists of a set of goals and a state describing the world. After a node has been selected, one of the node's goals must be

selected, and then an operator relevant to this goal must be chosen. Finally, a set of bindings for the parameters of that operator must be decided upon.

2. **Expansion phase:** If the instantiated operator's preconditions are satisfied, the operator is applied. Otherwise, PRODIGY subgoals on the unmatched preconditions. In either case, a new node is created with updated information about the state or the subgoals.

The search terminates after creating a node whose state satisfies the top-level goal expression.

2.3 Control Rules

As PRODIGY attempts to solve a problem, it must make decisions about which node to expand, which goal to work on, which operator to apply, and which objects to use. These decisions can be influenced by control rules for the following purposes:

1. **To increase the efficiency of the problem solver's search.** Control rules guide the problem solver down the correct path so that solutions are found faster.
2. **To improve the quality of the solutions that are found.** There is usually more than one solution to a problem, but only the first one that is found will be returned. By directing the problem solver's attention along a particular path, control rules can express preferences for solutions that are qualitatively better (e.g., more reliable, less costly to execute, etc.).
3. **To direct the problem solver along paths that it would not explore otherwise.** As with most planners, for efficiency PRODIGY normally explores only a small portion of the complete search space. However, control rules can be used to force PRODIGY to explore a path that would not be expanded by the default search.

PRODIGY's reliance on explicit control rules, which can be learned for specific domains, distinguishes it from most domain independent problem solvers. Instead of using a least-commitment search strategy, for example, PRODIGY expects that any important decisions will be guided by the presence of appropriate control knowledge. If no control rules are relevant to a decision, then PRODIGY makes a quick, arbitrary choice. If in fact the wrong choice is made, and costly backtracking proves necessary, an attempt will be made to learn the control knowledge that must be missing. The rationale for PRODIGY's *casual commitment* strategy is that for any decision with significant ramifications, control rules should be present; if they are not, the problem solver should not attempt to be clever without knowledge, rather, the cleverness should come about as a result of learning. Thus, our emphasis is on an elegant and simple problem solving architecture which can produce sophisticated behavior by learning the appropriate, domain-specific control knowledge.

Control rules can be employed to guide the four decisions described in Section 2.2. Each control rule has a left-hand side condition testing applicability and a right-hand side indicating whether to SELECT, REJECT, or PREFER a particular candidate. To make a control decision, given a default set of candidates (nodes, goals, operators, or bindings, depending on the decision), PRODIGY first applies the applicable selection rules to select a subset of the

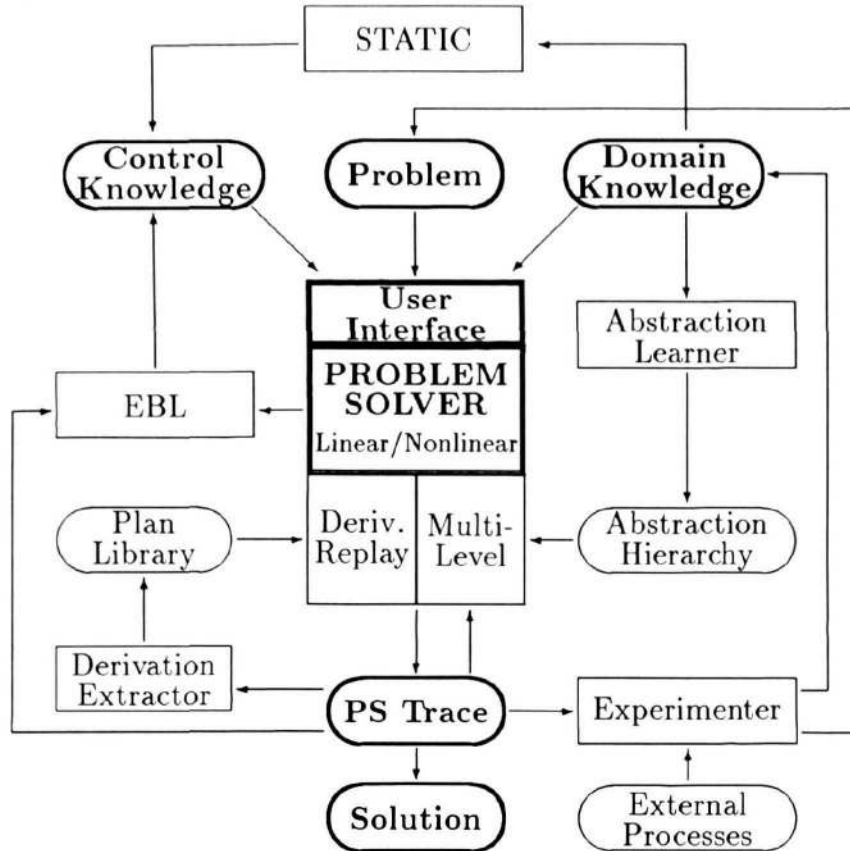


Figure 1: The PRODIGY Architecture

candidates. (If no selection rules are applicable, all the candidates are included.) Next rejection rules further filter this set by explicit elimination of particular remaining candidates, and finally preference rules are used to find the most preferred alternative. If backtracking is necessary, the next most preferred candidate is attempted, and so on, until a global solution is found, or until all selected and non-rejected candidates are exhausted.

2.4 The Learning Modules

PRODIGY's general problem solver is combined with several learning modules. The PRODIGY architecture was designed both as a unified testbed for different learning methods and as a general architecture to solve interesting problems in complex task domains. Let us now focus on the architecture itself, as diagrammed in Figure 1.

The problem solver produces a complete search tree, encapsulating all decisions – right ones and wrong ones – as well as the final solution. This information is used by each learning component in different ways. In addition to the central problem solver, PRODIGY has the following learning components:

APPRENTICE: A user interface that can participate in an apprentice-like dialogue [Joseph, 1989], enabling the user to evaluate and guide the system's problem solving and learning. The interface is graphic-based and tied directly to the problem solver, so

that it can both acquire domain knowledge or accept advice as it is solving a problem.

EBL: An explanation-based learning facility [Minton, 1988] for acquiring control rules from a problem-solving trace. Explanations are constructed from an axiomatized theory describing both the domain and relevant aspects of the problem solver's architecture. Then the resulting descriptions are expressed in control rule form, and control rules whose utility in search reduction outweighs their application overhead are retained.

STATIC: A method for learning control rules by analyzing **PRODIGY**'s domain descriptions prior to problem solving. The **STATIC** program [Etzioni, 1990] produces control rules without utilizing any training examples. **STATIC** matches **EBL**'s performance on some domains and exhibits one to two orders of magnitude faster learning rate. However, not all problem spaces permit purely static learning, requiring **EBL** to learn control rules dynamically.

ANALOGY: A derivational analogy engine [Carbonell and Veloso, 1988, Veloso and Carbonell, 1989] that uses similar previously solved problems to solve new problems. The problem solver records the justifications for each decision during its search process. These justifications are then used to guide the reconstruction of the solution for subsequent problem solving situations where equivalent justifications hold true. Both analogy and **EBL** are independent mechanisms to acquire domain-specific control knowledge.

ALPINE: An abstraction learning and planning module [Knoblock, 1990]. The axiomatized domain knowledge is divided into multiple abstraction levels based on an analysis of the domain. Then, during problem solving, **PRODIGY** first finds a solution in an abstract space and then uses the abstract solution to guide the search for solutions in more detailed problem spaces. This method is orthogonal to analogy and **EBL**, in that both can apply at each level of abstraction.

EXPERIMENT: A learning-by-experimentation module for refining domain knowledge that is incompletely or incorrectly specified. Experimentation is triggered when plan execution monitoring detects a divergence between internal expectations and external observations. The main focus of experimentation is to refine the factual domain knowledge, rather than the control knowledge.

3 Issues

Many controversial issues arise in the design and construction of an integrated architecture. These issues are resolved differently in **PRODIGY**, **SOAR** [Rosenbloom *et al.*, 1990], **THEO** [Mitchell *et al.*, 1990], **ICARUS** [Langley *et al.*, 1989], and other attempts at integrating multiple aspects of cognition and perception. **PRODIGY** takes specific positions in these design dimensions and derivative architectural attributes, which we elucidate below to promote discussion.

3.1 Types of Reasoning

Cognitive behaviors range from the emotive and intuitional to the rational and deliberative. Within the latter category, one can distinguish real-time decision making and long-term planning. PRODIGY models only deliberative planning and problem-solving, albeit in resource-limited domains.

Within planning and problem-solving task domains, several reasoning paradigms have been proposed, including least-commitment planning, “reactive planning”, and casual commitment planning. PRODIGY subscribes to the latter, as explained in Section 2.3, learning new control knowledge from correct and incorrect past commitments. “Reactive planning”, an archetypical oxymoron, refers to conditioned-reflex and the total absence of any higher-level deliberation or knowledge representation. We believe that whereas there is a role for such behavior, it is patently absurd to negate the need for higher level cognition as sometimes advocated by the proponents of “reactive planning” [Brooks, 1986, Agre and Chapman, 1987].

3.2 Modular vs. Monolithic Architecture

We do not know whether the human mind compartmentalizes distinct cognitive abilities, or distinct methods of achieving similar ends (such as our multiple learning methods). However, modularization at this level is a sound engineering principle, and therefore we have adhered to it. Integration is brought about by sharing both a uniform knowledge representation, and a common problem-solving and planning engine. Other systems, such as SOAR, take on a monolithic structure. There is only one learning mechanism, chunking, which can never be turned off or even modulated. Which is better? Clearly, we believe the former to be superior – but only from engineering principles rather than psychological ones.

3.3 Scaling Up

Any integrated architecture must address increasingly large tasks, whether its objective is to model human cognition or to build useful knowledge-based systems for complex tasks. Scalability can be calibrated in multiple ways, but all relate to efficient behavior with increasing complexity, as measured by:

- Size of the domain: total number of objects, attributes, relations, operators, inference rules, etc.
- Size of the problem: number of steps in the solution plan, number of conjuncts in the goal expression, size of the visited search space, etc.
- Variety: number of qualitatively different actions and object types in the domain.
- Perplexity: average fan-out at every decision point in the search space (with and without learned control knowledge).

In PRODIGY we seek to achieve a reasonable measure of scalability in all these dimensions. The learning techniques strive to reduce the visited search space in future problems with respect to the virtual (complete) search space.

3.4 Psychological Validity vs Cognitive Engineering

As mentioned earlier, the PRODIGY project strives to produce a useful, scalable, and maintainable reasoning and learning architecture. Where this matches human cognition, it is so by accident, by the limited imagination of the PRODIGY designers, or perhaps because the human mind has indeed optimized such aspects of cognition. In all other aspects, the goal is at reengineering cognition the way it ought to be, in order to be most useful in problem solving, planning, and learning. Here we enumerate a few additional ways in which PRODIGY differs from human thought and from other cognitive architectures:

- PRODIGY forgets when it chooses to do so. For instance, a control rule whose testing and application overhead is greater than the search reduction benefits accrued over time may be discarded. All other systems retain all acquired knowledge. Minton [Minton, 1988] demonstrated that the effectiveness of explanation-based learning is improved by measuring the utility of acquired knowledge and retaining only those rules with positive utility.
- PRODIGY deliberates on any and all decisions: which goal to work on next, which operator to apply, what objects to apply the operator to, where to backtrack given local failure, whether to remember newly acquired knowledge, whether to refine an operator that makes inaccurate predictions, and so on. It can introspect fully into its decision cycle and thus modify it at will. This is not consistent with the human mind, yet it is an extremely useful faculty for rapid learning.
- PRODIGY's knowledge acquired in one module is open to inspection and interpretation by other modules. Abstracted operators can be used to plan, to drive EBL, to analogize with past memory, and so forth. The compartmentalization is at the level of learning methods, and the sharing is at the level of all knowledge acquired. This is an architectural property different from those of other integrated architectures.

Acknowledgements

The authors gratefully acknowledge the contributions of the other members of the PRODIGY project: Daniel Borrajo, Oren Etzioni, Dan Kahn, Dan Kuokka, Michael Miller, Alicia Perez, William Reilly, Santiago Rementeria, Nobuyoshi Wada, and Xuemei Wang.

References

- [Agre and Chapman, 1987] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of the National Conference on Artificial Intelligence*. Seattle, WA, 1987.
- [Brooks, 1986] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1), 1986.

- [Carbonell and Veloso, 1988] J. G. Carbonell and M. M. Veloso. Integrating derivational analogy into a general problem solving architecture. In *Proceedings of the First Workshop on Case-Based Reasoning*, pages 104–124, Los Altos, CA, May 1988. Morgan Kaufmann.
- [Carbonell *et al.*, 1990] Jaime G. Carbonell, Craig A. Knoblock, and Steven Minton. Prodigy: An integrated architecture for planning and learning. In Kurt VanLehn, editor, *Architectures for Intelligence*. Erlbaum, Hillsdale, NJ, 1990.
- [Etzioni, 1990] Oren Etzioni. Why PRODIGY/EBL works. In *Proceedings of Eighth National Conference on Artificial Intelligence*, Boston, MA, 1990.
- [Joseph, 1989] Robert L. Joseph. Graphical knowledge acquisition. In *Proceedings of the 4th Knowledge Acquisition For Knowledge-Based Systems Workshop*, Banff, Canada, 1989.
- [Knoblock, 1990] Craig A. Knoblock. Learning effective abstraction hierarchies. In *Proceedings of Eighth National Conference on Artificial Intelligence*, Boston, MA, 1990.
- [Langley *et al.*, 1989] Pat Langley, Kevin Thompson, John A. Allen, Wayne F. Iba, and John Gennari. An integrated cognitive architecture for autonomous agents. Technical report, Department of Information and Computer Science, University of California, Irvine, 1989.
- [Minton *et al.*, 1989] Steven Minton, Craig A. Knoblock, Daniel R. Kuokka, Yolanda Gil, Robert L. Joseph, and Jaime G. Carbonell. PRODIGY 2.0: The manual and tutorial. Technical Report CMU-CS-89-146, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, May 1989.
- [Minton, 1988] Steven Minton. *Learning Effective Search Control Knowledge: An Explanation-Based Approach*. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, 1988.
- [Mitchell *et al.*, 1990] Tom M. Mitchell, John Allen, Prasad Chalasani, John Cheng, Oren Etzioni, Marc Ringuette, and Jeffrey C. Schlimmer. Theo: A framework for self-improving systems. In Kurt VanLehn, editor, *Architectures for Intelligence*. Erlbaum, Hillsdale, NJ, 1990.
- [Rosenbloom *et al.*, 1990] Paul S. Rosenbloom, Allen Newell, and John E. Laird. Towards the knowledge level in SOAR: The role of the architecture in the use of knowledge. In Kurt VanLehn, editor, *Architectures for Intelligence*. Erlbaum, Hillsdale, NJ, 1990.
- [Veloso and Carbonell, 1989] M. M. Veloso and J. G. Carbonell. Learning analogies by analogy — the closed loop of memory organization and problem solving. In *Proceedings of the Second Workshop on Case-Based Reasoning*, pages 153–159, Los Altos, CA, May 1989. Morgan Kaufmann.
- [Veloso, 1989] Manuela M. Veloso. Nonlinear problem solving using intelligent casual-commitment. Technical Report CMU-CS-89-210, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1989.