

KNOWLEDGE PROGRAMMING IN LOOPS: Report on an Experimental Course

Mark Stefik, Daniel G. Bobrow,
Sanjay Mittal, and Lynn Conway¹

*Knowledge Systems Area
Xerox Palo Alto Research Center
Palo Alto, CA 94304*

Abstract

Early this year fifty people took an experimental course at Xerox PARC on knowledge programming in Loops. During the course, they extended and debugged small knowledge systems in a simulated economics domain called Truckin. Everyone learned how to use the Loops environment, formulated the knowledge for their own program, and represented it in Loops. At the end of the course a knowledge competition was run so that the strategies used in the different systems could be compared. The punchline to this story is that almost everyone learned enough about Loops to complete a small knowledge system in only three days. Although one must exercise caution in extrapolating from small experiments, the results suggest that there is substantial power in integrating multiple programming paradigms.

KNOWLEDGE PROGRAMMING is concerned with the techniques for representing knowledge in computer programs. It is important in many applications of AI, where the problems

are messy. As in many situations in life, pat solutions and simple mathematical models just aren't good enough. Things break. Information is missing. Assumptions fail. Situations are complicated. To cope with messiness, AI researchers have found that large amounts of problem-specific knowledge are usually needed. This places a premium on the use of powerful techniques for representing and testing knowledge in computer programs.

Very few people have been trained to build knowledge systems. This is a critical bottleneck that limits the scope and impact of knowledge engineering. It limits the number of things that can be tried, the number of good ideas that are propagated, and the number of successful applications that influence the way that others perceive the field.

A few numbers may serve to put this in perspective. About one computer science researcher in ten does some work in AI, and perhaps a fifth of those work in knowledge engineering. In 1980, approximately 265 people graduated with Ph.D.'s in Computer Science, according to the "Snowbird Report" (Denning, et al., 1981). Fewer than a half dozen doctoral theses appear each year on some aspect of building knowledge systems. An estimate in a brochure by Teknowledge, Inc., indicates that there are only about sixty people in the world with high level expertise in the design and development of knowledge systems. Although precise figures for these populations are difficult to obtain, all the evidence suggests that the community is tiny, indeed.

¹ Now with the Defense Advanced Research Projects Agency (DARPA).

Copyright © 1983 by Xerox Corporation

Thanks to Johan de Kleer, Richard Fikes and John McDermott for their reviews and comments on earlier drafts of this paper. We extend our special thanks to the course participants from Applied Expert Systems, Daisy Systems, ESL, Fairchild AI Lab, Lawrence-Livermore Laboratories, Schlumberger-Doll Research Laboratory, SRI International, Stanford University, Teknowledge, and Xerox Corporation. Their participation and feedback are vital to the ongoing experimental process for simplifying the techniques of knowledge programming. We enjoyed and will long remember their spirited involvement.

Training in knowledge engineering usually requires several years of study at one of a handful of universities. A group of us in the Knowledge Systems Area at Xerox PARC is trying to shorten this training time. Our goal is to increase the impact and scale of knowledge engineering by simplifying the methods of knowledge programming and making them more widely accessible. In doing this we have developed an experimental knowledge programming system called LOOPS (Bobrow & Stefik 1981; Stefik, et al., 1983a). Feedback about the adequacy of LOOPS is collected from beta-test sites which are using it to build knowledge systems. Feedback about the learnability of LOOPS is collected from participants in experimental courses.

Integration and Paradigms

An important principle of knowledge programming is that different paradigms are appropriate for different purposes. This contrasts with the use of a single programming paradigm for everything, be it logic programming as in Prolog (Clocksin & Mellish 1981), procedure-oriented programming as in Lisp (Winston & Horn 1978), object-oriented programming as in Smalltalk (Goldberg & Robson 1983), or whatever.

There are various metrics of cost for applying a programming paradigm across a spectrum of applications. Examples of metrics are the cost of learning, the cost of modifying, the cost of debugging, and the cost of running. These costs vary across paradigms and applications because different programming paradigms provide different ways of organizing information in programs. For a given metric and application, some programming paradigms can be more cost-effective than others. By allowing for choice and combina-




Figure 1.

The LOOPS Logo. Illustrating the different paradigms in the current version of LOOPS: procedure-oriented, object-oriented, access-oriented, and rule-oriented. The ring is intended to suggest that LOOPS integrates the paradigms. They are not just complementary, but are designed to be used together in building knowledge systems.

tion of paradigms, a knowledge programming system enables various costs to be lowered. For example, we attribute much of our success in the experimental courses to the low costs for learning and applying LOOPS. For each of the things that the course participants needed to represent in their knowledge systems, there was some paradigm in LOOPS in which the expression of the knowledge was concise and the learning cost was low. Although there is room for much more work on programming paradigms and their applications, the principle seems clear: it is expensive to use one simple programming paradigm for everything.

As indicated in the LOOPS logo in Figure 1, LOOPS currently integrates four programming paradigms:

Procedure-oriented programming: In this paradigm, large procedures are built from small ones by the use of subroutines. Data and programs are kept separate. Most computer languages are like this. The procedure-oriented part of LOOPS is INTERLISP-D (Teitelman 1978, Xerox 1982). INTERLISP-D is shown at the base of the LOOPS logo to suggest that it provides the solid foundation on which the rest of LOOPS is built.

Object-oriented programming: In this paradigm, information is organized in terms of objects, which combine both instructions and data. Large objects are built up from smaller objects. Objects communicate with each other by sending messages. The conventions for communicating with an object by using messages constitute message protocols. Standardized protocols enable different classes of objects to respond to the same kinds of messages. Inheritance in a class lattice enables the specialization of objects.

Access-oriented programming: This paradigm is useful for programs that monitor other programs. Its basic mechanism is a structure called an active value, which has procedures that are invoked when variables are accessed. A useful way to think of active values is as probes that can be placed on the object variables of a LOOPS program. These probes can trigger additional computations when data are changed or read. For example, they can drive gauges that display the values of variables graphically.

Rule-oriented programming: This paradigm is specialized for representing the decision-making knowledge in a program. In LOOPS, rules are organized into rulesets which specify the rules, a control structure, and other descriptions of the rules. Two key features of the rule language are that it provides techniques for factoring control information from the rules, and also dependency-trail facilities, which provide mechanisms for "explanation" and belief revision.

These different organizational methods determine the way that information is factored and shared. Each paradigm provides a vocabulary and a set of composition methods for organizing information in a program.

Procedure composition. The composition methods of INTERLISP-D are forms of familiar control statements for iteration, recursion, and procedure call.

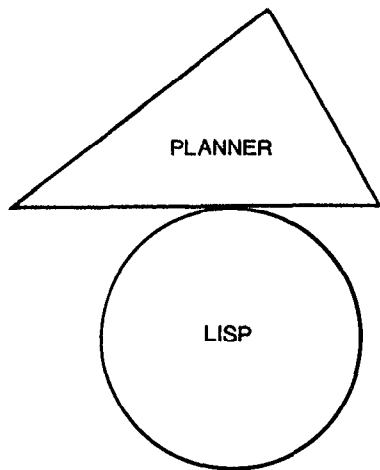


Figure 2.
Combining paradigms: The perch approach




Figure 3
Combining paradigms: The patch approach




Figure 4
Combining paradigms: The bridge approach.

Object composition This paradigm provides several composition methods (shown in figure 6). The simplest is the specialization of methods and variables of a superclass. Special classes called Mixins are used to impart a specific set of behaviors to a number of subclasses. The term "mixin" is borrowed from Flavors — (Weinreb & Moon 1981). Mixins exploit the multiple inheritance lattice by allowing inheritance to be factored. Composite objects extend the notion of objects to be recursive in structure so that multiple objects can be instantiated and linked together. Finally, perspectives in LOOPS are groupings of objects into a higher level object, such that each component is a view (or perspective) of the whole. Perspectives provide for the forwarding of messages to the appropriate view

Access composition Composition in this paradigm is done by nesting of active values. Analogous to the use of multiple probes in measuring a circuit, this composition assumes that the "probes" are for independent instruments and do not interfere with each other

Rule composition: The LOOPS rule-oriented paradigm provides for the sharing of rules among rulesets. It makes use of the other paradigms for organizing the interactions between the rules. Thus rules can call rulesets directly (using the procedural orientation), or invoke rulesets by sending messages (using the object orientation), or invoke rulesets by changing data (using the access orientation)

Integration has two major themes in LOOPS: integration to allow the paradigms to be used together in building a knowledge system; and integration of a programming environment for creating and debugging knowledge systems.

Some examples illustrate the integration of paradigms in LOOPS: the "workspace" of a ruleset is an object, rules are objects, and so are rulesets. Methods in classes can be either Lisp functions or rulesets. The procedures in active values can be LISP functions, rulesets, or calls on methods. The ring in the LOOPS logo reflects the fact that LOOPS not only contains the different paradigms, but integrates them. The paradigms are designed not only to complement each other, but also to be used together in combination.

Some examples from other systems illustrate the non-integration of programming paradigms. For example, Figure 2 shows the connection between PLANNER and LISP. PLANNER was implemented in LISP, but a programmer could not easily intermix PLANNER and LISP procedures. A simple mistake by a "naive" programmer could easily crash the whole system. Figure 3 shows the connection of list operations to PROLOG, reflecting the fact that list operations were added late to PROLOG, after the initial design. Figure 4 illustrates another approach, illustrated perhaps by the Spice Machine. In this example LISP and PASCAL communicate over a narrow bridge, making mutual use awkward and costly.

The second theme of integration is the integration with the programming environment. For example, LOOPS extends to other paradigms many of the facilities of INTERLISP-D,

such as the display-oriented break package, editors, and inspectors. In LOOPS, this integration has led to the same synergy that is exploited in using multiple paradigms for application programs. For example, the notion of “breaking” on access to a function is extended to breaking on access to a variable by using active values to invoke the break package; the notion of tracing is extended to the notion of having gauges that can monitor the values of variables.

Getting Ready for the First Course

On January 6, we began to plan the first LOOPS course that would be offered on January 31 to our beta-test sites. We made a preliminary course outline, but we knew that we needed some way to draw the participants into programming in LOOPS. The idea of a video game was suggested, say rocket ships with LOOPS programs controlling the thrust and phasers. This idea was rejected as being both too frivolous, and computationally too expensive. Another suggestion was a game for placing tiles. We knew from Malone (1980) that there were principles for making games motivating. Our course participants would be computing and other professionals drawn from research organizations and AI start-up firms, who were interested in using LOOPS for building expert systems. We needed something that they would find useful and appealing.

As brainstorming continued, some pedagogical principles began to emerge. The game should draw on the real world knowledge of our students. Rocket ships and tiles were wrong, because people didn't have experience with such things from their everyday lives. A board game like Monopoly was considered, and then our first concept of *Truckin'* emerged. It would be a board game with road stops (see Fig. 5). The players would drive trucks around buying and selling commodities. Their job would be to plan a route and make a profit. There would be various hazards along the way, places where goods and profits could be lost. Players would need to buy gas occasionally.

By mimicking real life, *Truckin'* would provide the kinds of difficulties that knowledge engineers encounter in building expert systems. We could create a rich and animated simulation environment for the “independent truckers.” The students would need to add knowledge to make their automated players more powerful. The simulation environment would draw on the student's real-world knowledge, and be rich enough to preclude a simple model. Much of the appeal of this was that the “common experience” character of *Truckin'* as a domain would enable us to side-step the usual knowledge acquisition bottleneck. The knowledge engineering experience would be accelerated by immediate feedback from the animated simulation. To help students get started, we would provide them with a simple expert system for playing *Truckin'*. We decided to teach students about knowledge programming in LOOPS by giving them a small knowledge system to extend.

At this point, we had less than a month to create the

course materials, lectures, and *Truckin'*. Sleep would become a rare and precious commodity. The *Truckin'* data base began to take shape. The players would start at Union Hall, and would try to be parked at Alice's Restaurant at the end of the game. There would be various kinds of hazards along the road. The player with the most cash at Alice's at the end of the game would win.

LOOPS was able to accommodate changes as our ideas evolved. Initially, we thought of the hazards as being road stops. This was probably a carry over of our childhood experiences with board games. Then we added the idea of “bandits” that could move around just like the independent truckers. Bandits were represented as an inheritance combination of players and consumers. We used active values on variables of the road stops to update the display for commodity prices and inventories. This meant that we did not need to find every place in the program where these things could potentially be changed, in order to update the display. The features of LOOPS worked for us, providing convenient techniques for factoring the program. We became experienced consumers of our own knowledge programming system as we raced to get ready for the course.

The simulation was designed to cause goal conflicts. A truck going quickly over a rough road would probably have its fragile merchandise damaged. A truck going quickly past a weigh station would probably get an extra fine, unless he was lucky or the weigh station was busy. On the other hand, a truck going slowly past a bandit would probably get intercepted. There would be perishable goods and fragile goods. We considered explosive goods and other such things, but removed them when they failed to add anything new to the game. Our pedagogical style was to leave some things out in order to keep it simple. A player could take only three kinds of actions: buy, sell, and move.

To facilitate the “suspension of disbelief” in watching the animated simulation, artistic attention had to be given to the appearances of things. Icons for the various commodities, hazards, and trucks were created. We experimented with different configurations of the gameboard, moving away from the outside edge configuration of most gameboards in order to pack enough road stops on the display screen. Highways were drawn next to the road stops, with a gray background and little dashed white lines in the center. People looked at intermediate versions of the gameboard and told us that the abrupt motion of the trucks was startling. We modified the code to simulate braking so that trucks would slow down as they arrived at their destinations. The visual appearance of *Truckin'* became seductive. People were drawn into it.

Prior to this, we had used a simple gauge in our demo to illustrate the application of active values. It was a crude looking gauge and had little generality. We decided to extend the collection of gauges so that people could use them for debugging and for monitoring their independent truckers during the simulation. A family of gauges was designed (see Figure 7). For further ideas on style, we collected some professional catalogs of gauges, and sought advice from

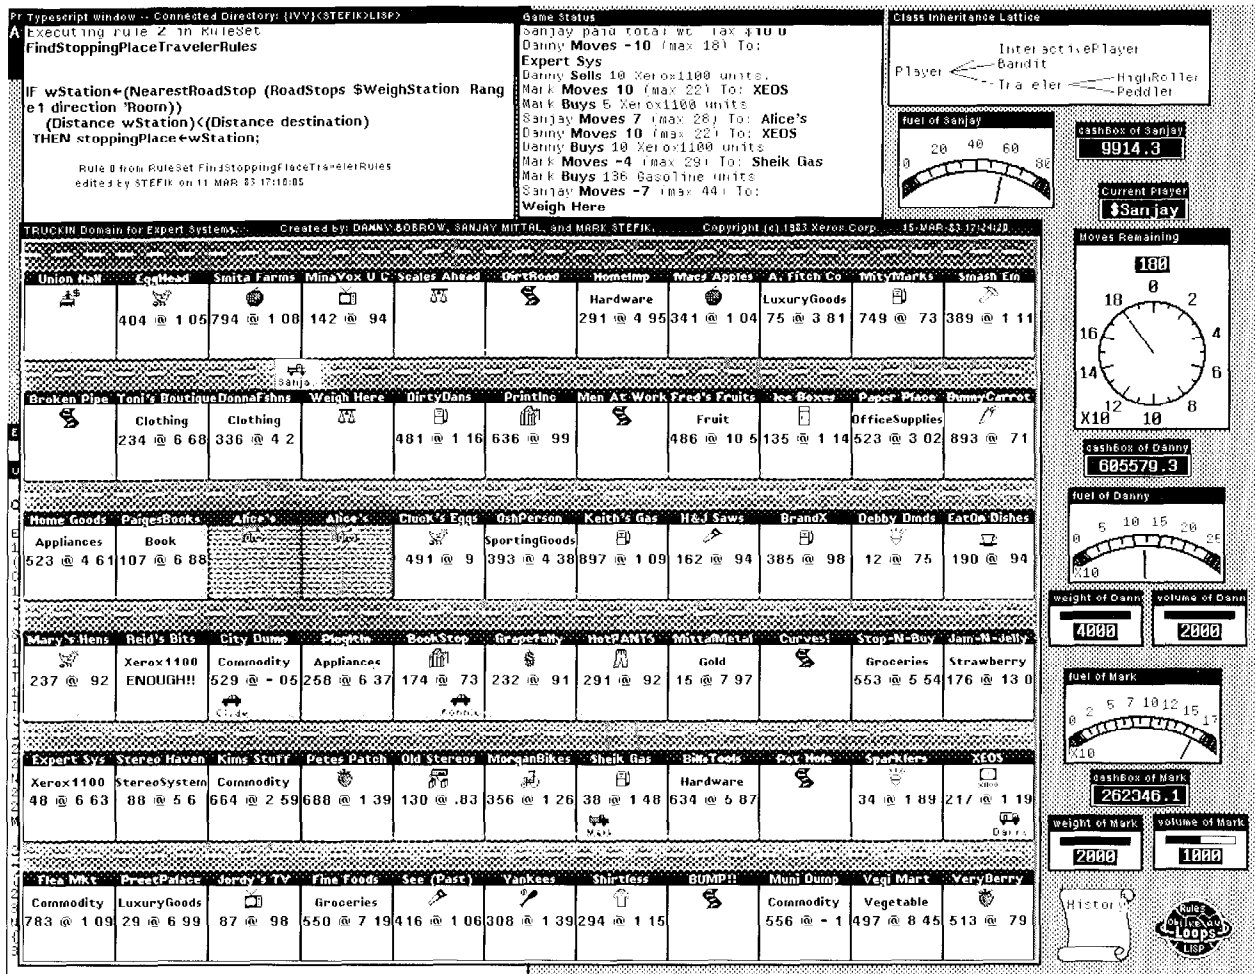


Figure 5 The Loops gameboard — for a game played by competing knowledge systems that emulate “independent truckers” The board’s squares are road stops, connected by the highway drawn above Roadstops can be producers, consumers, rough roads, weigh stations Roadstops with icons are producers, where players can buy Those with words (e.g., Clothing) are consumers, where players can sell The trucks for the players are shown parked or moving along the highway (e.g., Sanjay) To the right, a panoply of gauges monitors the status of various players In the upper left corner, a rule for one of the players is being traced.


Bruce Roberts on the Steamer project at BBN. The gauges went through several design reviews, to make the gauges simpler to use and modify Because of the extensive use of multiple inheritance and the interactions on the display between the parts of the hybrid gauges, a number of design issues surfaced. During these reviews, we created names for certain categories of design errors that we encountered For example, a *grainsize error* is a situation where the structural parts of an object (usually methods) are factored too coarsely for the fine control needed by its specializations. A *replication error* is a situation where almost the same structure is repeated in parallel classes, instead of factoring it in a way that would allow it to be shared. Such experiences gave us a deeper understanding of the programming issues that people would encounter in using the different paradigms

About two weeks before the course would begin, we sent







out notices to our beta-test sites inviting people to sign up for the course. We expected about a half dozen people We advertised that our course would provide hands-on experience in extending a “mini-expert system.” By word of mouth, the story spread Over fifty people called us, requesting to get on the list. We split the list in half and scheduled the second course for the end of February We didn’t send out any more advertisements. We had gone public and now we had to make it work

Suddenly it was the weekend before the course We made some guesses about the appropriate distribution of prices and penalties We created our first automated player, the Traveler, which would just travel along the board between Union Hall and Alice’s Restaurant As the Traveler cruised tirelessly around the game board, various bugs in the simulation surfaced. Meanwhile, we started work on a player to

Specialization



Mixins



Perspectives





Figure 6 Object composition in Loops. The inheritance lattice enables many forms of structural sharing in Loops. The simplest form is *specialization*, that is, creating a subclass that overrides and augments variables and methods inherited from the superclass. When multiple superclasses are used, the resulting subclass mixes together the attributes from the superclasses. *Perspectives* provide a way of grouping objects to act as views of a higher level object. *Perspectives* automatically forward messages to the appropriate object.

specialize in luxury goods called HighRoller. We didn't have time to debug it very well before the course started. We reasoned that the bugs were acceptable, since they would provide things for the course participants to fix. We were right, but in hindsight, we had a lot of gall.

The Courses as Experiments

We have now run two intensive knowledge programming courses, and also repeated the second course to a small group using videotape. By the time of publication of this article, the course will have been run for over 100 people. The courses are organized to alternate lectures and hands-on exercises (see Table 1). So far, everyone taking the course has learned enough about the LOOPS knowledge programming system to do some practice exercises (such as creating a new kind of gauge) and to build an extended (smarter) *Truckin'* player.

The most important aspect of the courses for our purposes is the opportunity that they provide for refining both LOOPS and the course materials. For us, the courses are

experiments, from which we are discovering how to make LOOPS and our teaching methods more effective. The basic structure of our experimental process is to run a course and take some measurements (for example, of the performance of students in terms of the problems that they complete, the questions that they ask, and the results of questionnaires that they return). We then change some parameters and take the measurements again during the next course. By examining how the observations and measurements differ, we can form hypotheses to guide subsequent iterations of the course.

- We substantially increased the emphasis on tools and techniques for debugging, and formulated explicit heuristics for programming in LOOPS. We taught students to use tools for understanding the behavior of a system. The second course led students to use gauges for monitoring the values of variables, explanation facilities (Fig. 8) for understanding which rule made a particular decision, and breaking and tracing facilities for discovering why some rules do not fire.
- In some cases, we introduced intermediate problems in the exercises, having hypothesized that some of the

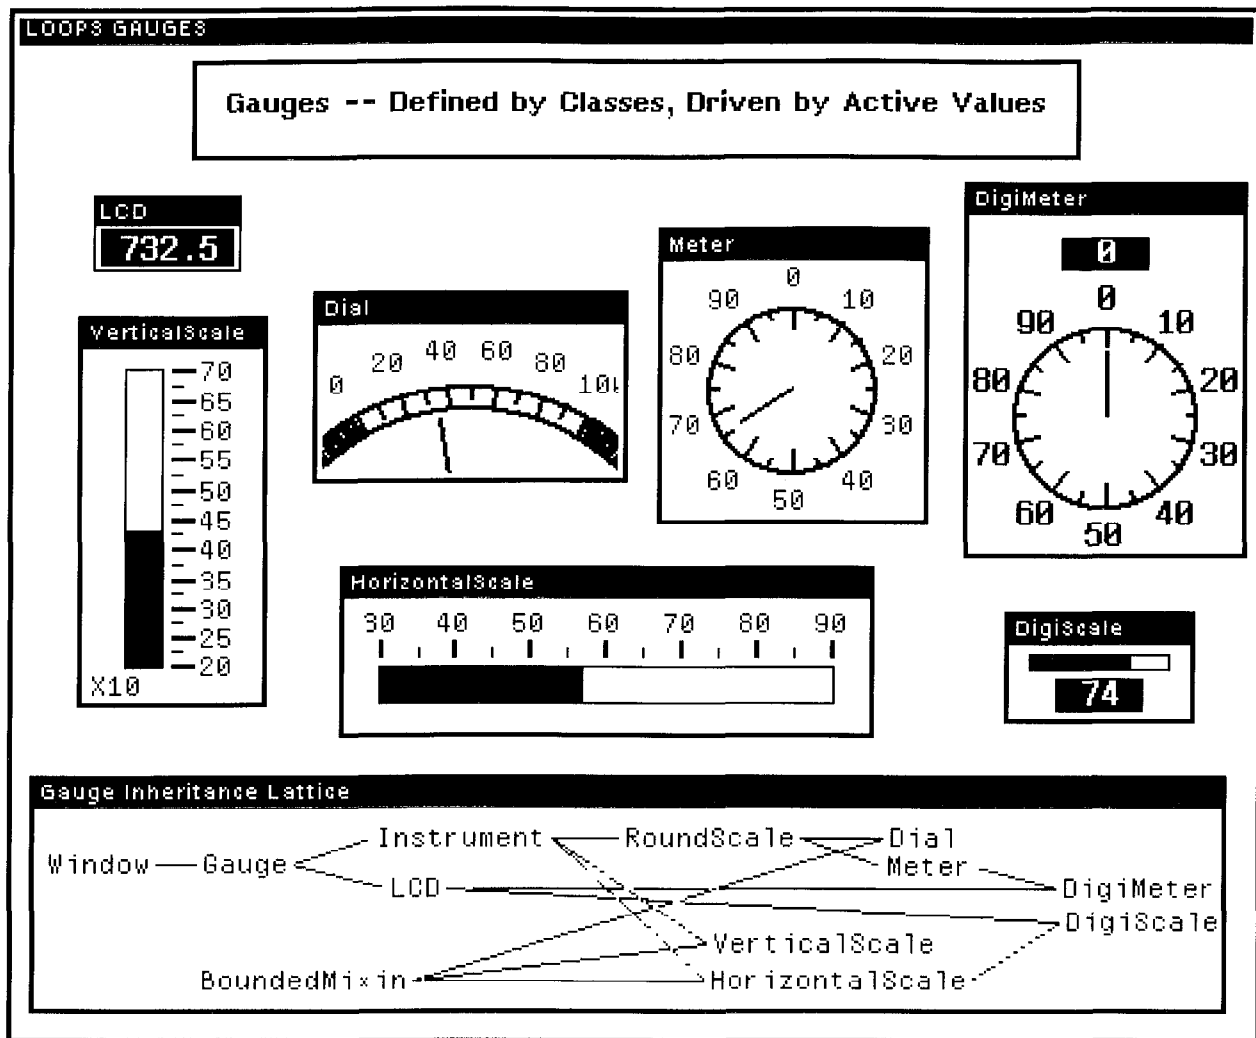


Figure 7 Loops gauges. Gauges are tools used to monitor the values of variables. They can be thought of as probes inserted onto the variables of an arbitrary Loops program. Gauges are defined in Loops as classes, and driven by active values – the mechanism behind access-oriented programming in Loops. A browser at the bottom of the figure illustrates the relationships between the classes of gauges. From this figure, we can see that the DigiMeter is a combination of a Meter and an LCD.

steps between exercises were too difficult to take all at once.

- We fashioned a new starting player for the second course, called the Peddler, which did a better job than HighRoller in factoring the concerns of an independent trucker. We hypothesized that restructuring HighRoller was too difficult to do in a three day course.
- We adjusted the commodity prices and risks to provide a greater reward and selection pressure for more sophisticated and knowledgeable truckers.
- We improved the browsers, that is, our interactive graphics for “browsing” information in a knowledge base (see Figure 9). We believed that we could reduce much of the cognitive load for restructuring objects and accessing information if we provided

more effective ways of making the right information visible.

- We fixed troublesome bugs in the rule compiler. During the first course, participants had to struggle with a compiler that did not reliably keep the generated LISP code in correspondence with the rules.

As a result of these changes, the participants in the second course were dramatically more successful than those in the first. In the first course, we had to slip the schedule for the knowledge competition by 90 minutes, in order to let people finish preparing their players. In the second course, people had players ready in about half of the allocated time, and spent the remaining time exploring other aspects of the system and tuning their strategies. Furthermore, the weakest player of the second course could easily dominate the best player from the first course.

Game Status
 at stereo haven
 Cash lost: \$83372 24
 >> #Mark intercepted!
 Sanjay Moves -4 (max 16) To:
Weigh Here
 Sanjay paid total Mt. Tax \$10 0
 Mark Moves -3 (max 18) To: HomeImp
 Mark Sells 40 San units
 Moves Remaining: 140
 Sanjay Moves -7 (max 46) To:
DirtyDans
 Sanjay Buys 32 Gasoline units
 Mark Moves 4 (max 4) To: Keith's Gas

Class Inheritance Lattice
 InteractivePlayer
 -Bandit
 Player
 -Traveler
 -HighRoller
 -Peddler

**IF gasStation+(FurthestRoadStop (RoadStops \$GasStation Range 1 direction 'Room'))
 (Distance gasStation)<(Distance destination)
 THEN stoppingPlace+gasStation;**
 Rule 0 from RuleSet FindStoppingPlaceTravelerRules
 edited by STEFIK on 11 MAR 83 17:18:05

Rule Exec (stopgap version)
 re: stoppingPlace
 #&(KeithGas (5 260921))
 re: why stoppingPlace
 (* Don't run out of gas.)
 IF goal = "Sit Tight"
 truck: fuel < 25 * truck::MaxFuel
 truck: cashBox > 0
 gasStation+(NearestRoadStop (RoadStops \$GasStation Range 1 NIL 'Room'))
 THEN stoppingPlace+gasStation;
 Rule 0 from RuleSet FindStoppingPlacePeddlerRules
 edited by STEFIK on 11 MAR 83 17:11:25

Gameboard Items:
 Union Hall: 514 @ .98
 Art's Art: 821 @ 1 05
 Broken Pipe: 567 @ 8
 Smirc's Farms: 63 @ 11 6
 DirtyDans: 471 @ 3 88
 Xerox 1100: 94 @ 1.04
 Keith's Gas: 442 @ 7 74
 HomeImp: 370 @ 1 18
 Pez Note: 780 @ 88
 Bunny Carrot: 500 @ 78
 Xerox 1100: 114 @ 6 86
 Commodity: 772 @ 2 92
 Lane Closed: 794 @ 1.23
 Flea Mkt: 41 @ 88
 40'er Miles: 172 @ 9 97
 Men At Work: 162 @ 4.75
 Weigh Here: 310 @ 14
 Fish Scales: 40
 Sheik Gas: 83 @ 1.18
 Luxury Goods: 48 @ 5 71
 Mongan Bikes: 420 @ 1.55
 Cheap Gas: 592 @ 94
 Gold: 14 @ 9 78
 Zears Stereo: 330 @ 2 27

Player Resources:
 fuel of Sanjay: 0 to 80
 cashBox of Sanjay: 6127.6
 Current Player: Mark
 Moves Remaining: 139
 fuel of Danny: 0 to 25
 cashBox of Danny: 4,799.00
 fuel of Mark: 0 to 17
 cashBox of Mark: 429649.0
 weight of Danny: 100
 volume of Danny: 100
 weight of Mark: 0
 volume of Mark: 0

Figure 8. Seeing the knowledge behind a decision. In this figure the game is interrupted, causing the Rule Exec window to pop up over the gameboard. The user has asked why his truck picked a particular stopping place, and Loops has displayed the rule that made the decision.

People asked far fewer questions in the second course, and were able to complete many more of the exercises. In addition, the questionnaires from the second course came back with radically different advice from those from the first course. The general response from the first course was "give us less on rules" and many people indicated substantial concern with many of the fundamental aspects of that paradigm. In the second course, the responses turned completely around. They said "give us more on rules and debugging."

We believe that in the first course the combination of a faulty rule compiler and lack of information on how to debug programs in this paradigm undermined confidence. During the second course, two members of a team were observed staring at a display. One of them said, "Why is it buying tomatoes?" and the other one elbowed him saying "Ask why! Ask why!" -- goading him into action at the LOOPS keyboard. They had learned their lessons well.

This process of simplifying methods and tuning the course in order to enhance learnability and propagatability

reflects our interest in the engineering of knowledge (Conway 1981, Stefik & Conway 1982). In this case we are engineering languages and techniques for knowledge programming. The courses provide a source of feedback on the effects of changes to the course materials, paradigms and programming environments. In time, we would like to extend our work to provide a framework that would simplify the process of creating higher level organizations in expert systems (Stefik *et al.*, 1982).

The Knowledge Competition

A very enjoyable thing about the LOOPS course is the electric excitement that erupts during a knowledge competition. People seem to project themselves into the players that they have created. They have put their player through many simulations and many playing conditions. In a sense, they have taught it everything. But during the competition there is a moment of truth. The rules cannot be changed. Success

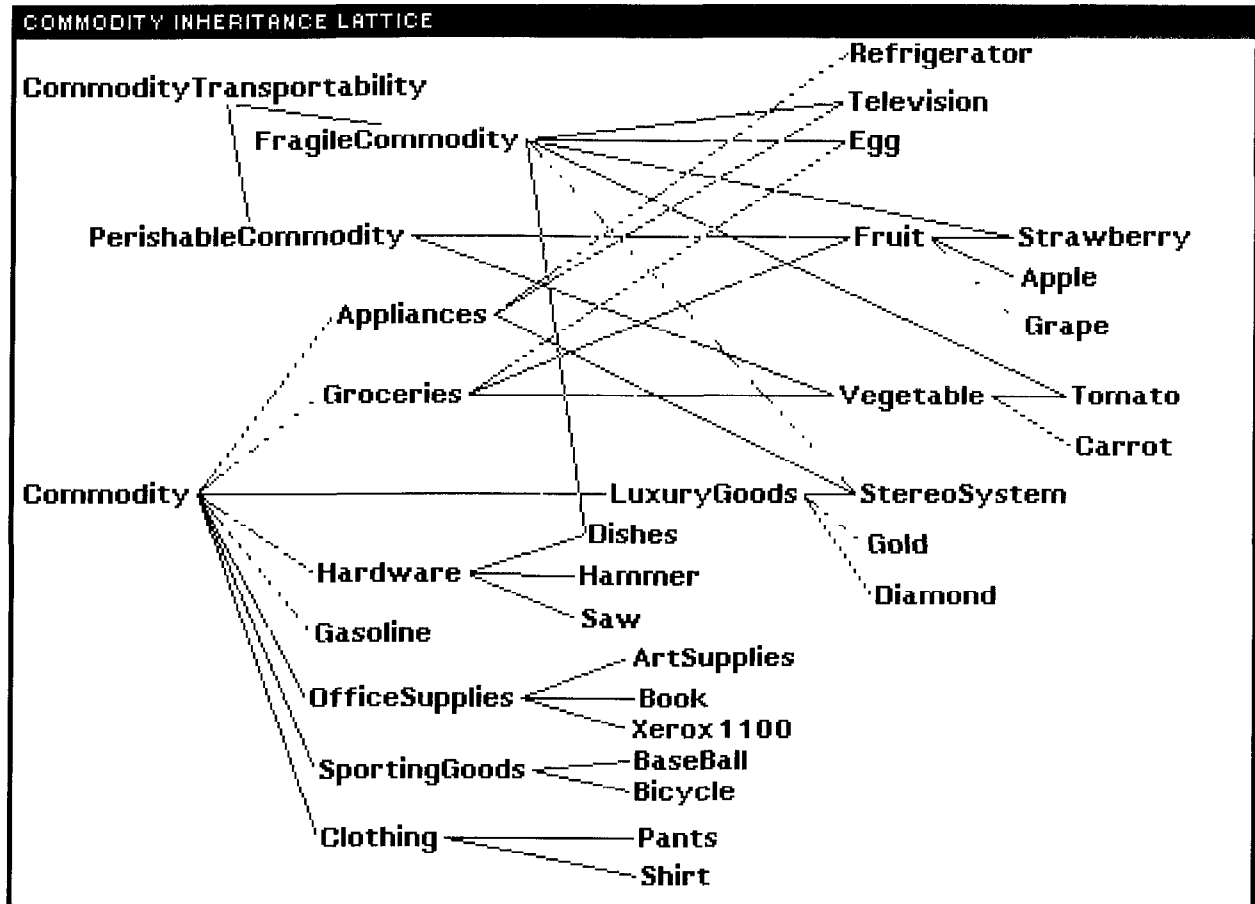


Figure 9. Class browser on commodities. Browsers are interactive programs used to browse through a knowledge base. The lines seen in a class browser indicate superclass relationships. For example, in this figure, a StereoSystem is a LuxuryGood, an Appliance, and a FragileCommodity. Browsers can be created to show other relationships too, and by selecting nodes in a browser, a user can access such further information.

in the short run is affected by chance, but on average, the most knowledgeable players will win.

The randomly generated game board comes up. As the simulation begins, there is a great deal of commentary and jibes as people compare their players. Who's ahead? Who just got robbed? In *Truckin'* the silliness of the ill-fated move is something that all the observers appreciate almost immediately. For example,

- A player may be racing to Alice's Restaurant. One move before the game ends it is unable to resist a business "opportunity" and doesn't make it to Alice's.
- A player may go to the closest place to sell some goods, even if it happens to be the City Dump, which unfortunately pays a "negative price."
- A player may become focused on a tight producer/consumer loop, making money faster than any other player on the board. If it is programmed to only buy fuel from stations along its route, but there is no gas station in the tight loop, the team will watch anxiously as the fuel gauge drops lower and lower.

- A player may try to park next to Alice's Restaurant near the end of the game, even if that happens to be the Union Hall, which confiscates all goods and cash.

In our experience so far, these oversights happen in the best of players. They provide a source of merriment during the competition, and an illustration of just how much knowledge is really needed to be powerful, even in an artificial environment.

The knowledge competition also serves as a source of examples and metaphors about the nature of knowledge. One example drawn from the first LOOPS course illustrates the interplay between knowledge and environment. For the first knowledge competition, two teams prepared players by simply fixing some of the bugs in the HighRoller. They had a private playoff just before the competition, and discovered that when both players were in the same game, the inventory of luxury goods on the game board became exhausted before the end of play. Neither player was able to cope with this situation. One of the heuristics that we now offer to teams

preparing for a knowledge competition is to test rules with many copies of the same player competing at once.

This interplay between knowledge and environment brings to mind the example of the ant on the beach (Simon 1981), in which the apparently complex movement of the ant is attributed to the complexity of the beach environment rather than the complexity of the ant. In *Truckin'*, the "ants" are mechanical and programmable. We have observed that even the complexity of the *Truckin'* environment creates a substantial selection pressure for resourceful and knowledgeable players. To win, the designers of the players must pit knowledge against complexity. Knowledge provides the adaptability needed for mastering the situations in the game

The name "knowledge competition" was inspired by the observation that it is truly the knowledge of players that is competing, and the most adaptable player wins. Recently in connection with the interest in fifth generation computers, Feigenbaum and McCorduck (1983) have characterized knowledge as the new "wealth of nations." In the knowledge competition and *Truckin'*, the competitive advantages of knowledge in a player is concrete and observable in short experiments

The success of the knowledge competition in motivating participation has led us to speculate on ways of alleviating the knowledge acquisition bottleneck. One idea is for a community of experts to interact through knowledge servers, which accept knowledge over a computer network and make themselves available for solving problems. Here again there would be a "competition" between different bodies of knowledge from the experts, competing to solve the problems that are posed

Implications

Sometimes the effects of a technological change can be surprising and widespread. Although our research and experimentation with LOOPS has not run its full course, there have been a few expert systems started at our beta-test sites: three systems that perform parts of VLSI design, a program for playing Bridge, an investment advisor, a program for expressing specifications of parallel programs, a tester for LOOPS

We sense that a technological change is emerging from such research on knowledge programming, a change in the infrastructure for building knowledge systems. The shift will have leveraging power in two ways: (1) the freeing of existing knowledge engineers from spending a year or two building the bottom of their knowledge representation systems, and (2) a measurable acceleration in the progress of the field if the simplified methods trigger an increase in numbers of practitioners from 100 to 1000 or more. Knowledge engineering can then begin to have a noticeable effect in many areas of our lives

LOOPS COURSE OUTLINE

FIRST DAY:

9:00-9:15	Introduction
9:15-10:15	Object-Oriented Programming: Classes - Objects Variables - Methods - Inheritance - Documentation
10:15-11:00	LOOPS Environment (Demonstration): Defining Methods - Editing - Printing - Inspecting Browsing - Gauges.
11:00-12:00	Exercise 1- Introductory hands-on session: Sending Messages - Browsing - Editing - Inspecting
12:00-1:00	Lunch.
1:00-2:00	Access-Oriented Programming: Active Values - FirstFetch - NamedObjects - AtCreation - Nested Active Values - The LOOPS Break Package
2:00-4:00	Exercise 2 - Gauges hands-on session: Specializing Classes - Instantiation - Using Gauges
4:00-4:30	The <i>Truckin'</i> mini-Expert System
4:30-5:00	Discussion

SECOND DAY:

9:00-9:15	Introduction.
9:15-10:15	Rule-Oriented Programming: RuleSets - Control Structures - Recording Rule Invocations.
10:15-12:00	Exercise 3 - Rules hands-on session: Editing RuleSets - Debugging RuleSets.
12:00-1:00	Lunch.
1:00-2:00	Knowledge Representation Examples from <i>Truckin'</i> .
2:00-4:30	Exercise 4 - Knowledge programming hands-on session: Rule-Oriented Programming - multiple paradigm programming.
4:30-5:00	Discussion.

THIRD DAY:

9:00-9:15	Introduction.
9:15-11:00	Initial Development of your player: hands-on session
11:00-12:00	Advanced LOOPS Features: Composite Objects Perspectives vs. Mixins - Meta Classes - System Mixins - Knowledge Bases.
12:00-1:00	Lunch
1:00-3:00	Final tuning of your player: hands-on Session
3:00-4:00	The <i>Truckin'</i> Knowledge Competition
4:00-4:30	LOOPS Environment: LOOPS Tester - Facilities for Bug Reporting
4:30-5:00	Wrap-up: LOOPS support - User Packages - Future Directions

Table 1.

References

- Bobrow, D G & Stefik, M (1981) *The LOOPS Manual* Tech. Rep KB-VLSI-81-13, Knowledge Systems Area, Xerox Palo Alto Research Center (PARC).
- Clocksinn, W F & Mellish, C S (1981) *Programmung in Prolog* Berlin: Springer-Verlag.
- Conway, J. (1981) The MPC adventures: Experiences with the generation of VLSI design and implementation methodologies *Proc of the Second Caltech Conference on Very Large Scale Integration*, 5-28
- Denning, P J., Feigenbaum, E , Gilmore, P , Hearn, A., Ritchie, R W , & Traub, J (1981) The Snowbird Report: A discipline in crisis. *Communications of the ACM*, 24:370-374
- Feigenbaum, E , & McCorduck, P (1983) *The fifth generation Artificial Intelligence and Japan's Challenge to the World* Reading, MA: Addison-Wesley
- Goldberg, A , & Robson, D. (1983) *Smalltalk-80 The language and its implementation* Reading, MA: Addison-Wesley
- Malone, T W. (1980) *What makes things fun to learn? A study of intrinsically motivating computer games* Technical Report CIS-7 (SSL-80-11), Xerox PARC
- Simon, H. A. (1981) *The sciences of the artificial* Cambridge, MA: The MIT Press
- Stefik, M., Bobrow, D , & Mittal, S. (1983) *Knowledge programming in LOOPS Highlights from an experimental course* Video Report KSA-83-1, Xerox PARC
- Stefik, M., Bell, A G , & Bobrow, D. G (1983) *Rule-oriented programming in LOOPS* Tech. Rep KB-VLSI-82-22, Knowledge Systems Area, Xerox PARC.
- Stefik, M., & Conway, L (1982) The principled engineering of knowledge *AI Magazine* 3(3):4-16.
- Stefik, M , Aikins, J., Balzer, R., Benoit, J., Birnbaum, L., Hayes-Roth, F., & Sacerdoti, E (1982) The organization of expert systems: A tutorial. *Artificial Intelligence* 18:135-173
- Teitelman, W (1978) *Interlisp Reference Manual* Technical Report, Xerox PARC
- Weinreb, D & Moon, D (1981) *Lisp Machine Manual*. Cambridge, MA: MIT Artificial Intelligence Laboratory
- Winston, P & Horn, B (1981) *Lisp* Reading, MA: Addison-Wesley
- Xerox Corporation (1982) *INTERLISP-D users guide* Pasadena, CA: Xerox Special Information Systems.

Note

LOOPS is available to selected Xerox customers designated as beta-test sites. The Knowledge Systems Area at Xerox PARC offers the intensive LOOPS course to selected applicants periodically for its research purposes.

Exciting new books from Harper & Row...

"This is the finest introduction to LISP ever written"
— Daniel L. Weinreb, Symbolics, Inc

David S. Touretzky LISP

A Gentle Introduction to Symbolic Computation

CONTENTS: Getting Acquainted. Functions and Data Lists. EVAL Notation Meet the Computer. Conditionals Global Variables and Side Effects List Data Structures Applicative Operators Recursion Elementary Input/Output. Iteration Property Lists. Appendix A—Recommended Further Reading Appendix B—Dialects of LISP. Appendix C—Extensions to LISP Appendix D—Answers to Exercises.

Marc Eisenstadt & Tim O'Shea ARTIFICIAL INTELLIGENCE

Tools, Techniques, and Applications

CONTENTS: TOOLS: An Introduction to Prolog, by William F Clocksin An Introduction to LISP, by Tony Hasemer Advanced LISP Programming, by Joachim Laubusch. A New Software Environment for List-Processing and Prolog Programming, by Steve Hardy
TECHNIQUES: How to Get a Ph.D. in AI, by Alan Bundy, Benedict du Boulay, Jim Howe, and Gordon Plotkin Cognitive Science Research, by Jon Slack Robot Control Systems, by Steve Hardy Kinematic and Geometric Structure in Robot Systems, by Joe Rooney Implementing Natural Language Parsers, by Henry Thompson and Graeme Ritchie. APPLICATIONS: Computer Vision, by John Mayhew and Henry Thompson Industrial Robotics, by William F Clocksin and Peter Davey Text Processing, by Paul Lefrere Planning and Operations Research, by Lesley Daniel.

ORDER TODAY.

Send this coupon to M Gonsky, Suite 5D, Harper & Row, 10 East 53d Street, New York, NY 10022



Harper & Row

Please indicate number of copies desired

- ___ LISP: A Gentle Introduction to Symbolic Computation @ \$17.95
- ___ ARTIFICIAL INTELLIGENCE: Tools, Techniques, and Applications @ \$21.50

Postage and handling: (Please include \$1.50 for the first book 50¢ for each additional copy) _____

Applicable sales tax: _____ Total: _____

Enclosed is my check/money order

Please charge my VISA MasterCard American Express

Exp. date: _____ Card # _____

Signature _____

Name _____

Address _____

City/State/Zip _____