Agents, Habitats, and Routine Behavior

Richard Alterman and Roland Zito-Wolf
Computer Science Department, Center for Complex Systems
Brandeis University
Waltham, MA, USA 02254

Abstract

This paper describes research that characterizes the development of routine behavior based on a model of the historic relation of the agent to his environment. The view developed is that the agent forms a 'habitat' outside of which his performance degrades. Routine behavior emerges from the history of the relations between the agent and his habitats in the service of recurring goals. Routines are customized to the agent's environment, but so constructed as to support future related activities and the adaptation to new circumstances that extend the agent's range of activity. In this paper we focus on examining quantitatively how this customization reduces the agent's workload.

Introduction

One characteristic of recent work in the goal directed behavior of agents has been a tightening of the relationship between an agent and his environment. This tighter coupling of agent and environment has taken several forms, including: the development of agents that interface with the world rather than manipulating internal representations (Brooks, 1991); the design of special purpose agents for particular goals and environments; and models of agency that rely more on reacting than on planning.

The focus of this paper is on two aspects of the relation of agent and environment: the development of routines, and the grounding of routines in *habitats.* By routine it is meant those activities that are regularly repeated. One advantage of customizing the agent's behavior to his routines is the reduction of planning effort. Because some routines regularly occur in the same place, *habitats* can begin to emerge. The advantage of habitats over other places of operation include: greater skill in focussing on relevant details and the increased accessibility from memory of relevant information at each point of a routine interaction.

The mechanism we present will maintain, given a goal, a history of the agent's interactions with a particular kind of artifact (device) within a given habitat. This paper will examine how the history of the relations between a situated agent and his individual environment in the service of his recurring goals result in the development of routine behavior. Previous episodes of interaction accumulate in memory, and are segmented and sorted so as to guide the agent in future behavior. With each subsequent activity, the agent further explores and learns about his habitats. This learning magnifies the agent's abilities for the domains and situations with which he is most familiar. Thus, for example, at any given point of a routine interaction, occurring at its usual place, retrieval of relevant information from memory can be cued by the features of the external world that in the past have become available to the agent at that point in the interaction.

The relation of the agent to his environment

Routines. Agre (forthcoming) refers to recurring patterns of activity as *routines.* It is a counting exercise to show that a day is made up of routines: getting up in the morning, driving to work, calling home from the office, loading the dishwasher, brushing one's teeth. For day-to-day routine activities it is more often the case that the circumstances vary slightly than that the goals of the planner have changed. To the extent that an agent's activities fit this model, the development of standard routines for those activities will potentially save a great deal of planning effort.

Our focus will be on the routines which are common practices. We define *common practices* as routines, goals, and situations that are common to more than one agent within a given community of agents; *scripts* (Schank & Abelson, 1977) were a method for internally representing common practices. Examples of common practices abound: procedures for driving vehicles, riding public transportation, operating a dishwasher, using a library, renting a video, getting change from a change machine, using a vending machine, attending a ballgame.

Economic forces, technological innovation, social trends or local variations between communities can all be sources of variance in common practice. De-

spite continuous change, because of the shared set of practices from which new practices are created (e.g. Minsky, 1975; Schank and Abelson, 1977; Rumelhart, 1980), there is constancy in the world. The cultural history of practices acts as a background from which new practices are revised into being (Cole, 1990). As a rational agent operating in such a world, in many cases the ability to function effectively depends on the ability to construct an interpretation, an 'understanding', of the novel aspects of the current situation in terms of shared concepts of culture, community, home, and workplace. This is a doubly effective strategy because in many cases the world has been arranged (presented) so as to communicate to an agent the action to be taken; the prevalence of instructions is one piece of evidence that this sort of communication is occurring. In our view, the comprehension processes that are at work to build an interpretation of variance in practice are the same that are at work in the reading of a narrative.

Habitats. Our dictionary (Webster's 1970) defines a *habitat* as "The place where a person or thing is ordinarily found." Likewise, a situated agent performs activities not in arbitrary places, but in those places he frequents. The advantages of this kind of familiarity are nicely summarized by the following advertisement for Cannon copiers "Cannon will assign you a repairman who not only knows your type of machine but also knows your machine."

A given routine in memory will over time come to reflect the details of the habitat(s) where that activity is usually performed. An agent does not make ratatouille in some abstract kitchen, but rather in a specific kitchen, and after a while his ability to make ratatouille becomes dependent in part on that particular environment. Moreover, one would expect the agent's performance to degrade outside of the agent's habitats. In somebody else's kitchen, such issues as finding the spices, pots, and pans, and divining the idiosyncrasies of the oven, require time and effort. What distinguishes my habitats from other situations is my relation to and use of the features presented by that situation:

Familiarity In a habitat an agent has greater access to the details of the situation.

Focus Customizing routines to habitats facilitates the selective noticing of relevant features.

Perceptual Effort For non-habitats there are difficulties in locating the relevant details of a situation.

Retrieval Accuracy and Effort Because the agent and his memory is tightly integrated with its habitat the information in memory is more readily accessible. At each step in the performance of a routine within a habitat reminding is facilitated by the cues readily accessible at that point in the interaction.

This paper shows how a routine can be used as a framework to organize the relevant features extracted from habitats, and how the accumulation of such details in turn enriches and extends the routine. With respect to the discussion above we will empirically test the hypothesis that the accumulation of features associated with a given habitat over time reduces the information-processing load and the amount of active looking on the part of the agent for subsequent routine-based interactions.

A Case-Based Model of Activity

The FLOABN system (Alterman, Zito-Wolf, and Carpenter 1991) is a project exploring a range of issues including, interaction and activity, memory and case-based reasoning, instruction usage, spatial reasoning, and the social and cultural conditions of action. The domain of FLOABN is the everyday usage of office and household devices such as photocopiers and telephones. FLOABN acquires and revises routines via adaptation and through the interpretation of instructions and messages read or received during interaction with its environment.

The basic functioning of FLOABN combines case-based reasoning (CBR: Kolodner & Simpson, 1989; Rissland & Ashley, 1986; Hammond, 1990) and comprehension and is based on the idea of adaptive planning (Alterman, 1988). The model of adaptive planning is that a situated agent retrieves from memory a case that he adapts as he proceeds with his interaction with the world. These adaptations occur as a result of interpreting the external world. Construing the world is an important resource in selecting action, especially in terms of adapting to changes in common practice, where it can directly lead to action. An example of where interpretation is sufficient to select action is in the usage of instruction.

The core of FLOABN is an adaptive planner called SCAVENGER (Zito-Wolf & Alterman 1992). FLOABN also includes an instruction reader (IIMP: Carpenter & Alterman, 1991; SPRITe: Carpenter & Alterman, 1993) for processing instructions and other messages received during an activity. FLOABN interacts with a world defined by a discrete-event simulator.

The Representation of Routines as History

In FLOABN, the agent's memory of previous problem-solving episodes is organized using a *multicase*.[1,2] The history of previous related episodes of interaction, as organized by the multicase, frames routine interaction, providing both the expectations of how an interaction normally proceeds and background against which events are interpreted. At the same time, the multicase

[1] A detailed comparison of multicase to other case-based models of episodic memory can be found in (Zito-Wolf & Alterman, 1993)

[2] For further details on the multicase, the associated model of episodic memory, and the adaptive planner SCAVENGER, see Roland Zito-Wolf's (1993) dissertation *Case-Based Representations for Procedural Knowledge.*

is constituted from cases derived from episodes. That is, each episode has the potential to contribute new components and detail to the history of interactions, a process called the *enrichment* of the multicase.

We define a *decision point* (DP) as any point in a routine requiring selection among alternatives, such as branch points, projections of events, and parameter bindings. A multicase organizes a set of decision points into a directed graph.

Each decision point contains the knowledge relevant to making a single decision, represented in terms of cases. A DP specifies the known *options* at that point (GET-COPY-CARD, CHECK-POWER), associating with each option a set of *descriptors* of situations in which it is considered appropriate. Each descriptor represents one segment of some particular problem-solving episode; it describes that segment of experience by specifying a set of features and value available to the agent at that particular point in the interaction. It is because the descriptors are tied to features of the environment that the routine becomes grounded in the agent's individual world and, over time, habitats can emerge.

Each pair of option and descriptor is called a *case*. A DP may contain multiple cases for each option. To make a decision, the features of the current situation are matched against the descriptors stored with each option. Shown below is the DP expressing the fact that the copier may need to be enabled, either by turning power on (upstairs copier) or by inserting a copy card (downstairs copier).

```
(GO-TO-COPIER NEXT-STEP ->
    (CHECK-POWER    ((LOCATION UPSTAIRS-OFFICE)
                     (COPIER COPIER-004) ...))
    (GET-COPY-CARD  ((LOCATION ROOM-1-002)
                     (COPIER COPIER-001) ...)))
```

The descriptor list associated with each alternative lists the values of roles and parameters that are thought to characterize the circumstances under which that alternative is relevant. Decision points may be used to conditionalize any feature of any structure in FLOABN's episodic memory, such as parameter bindings for steps. For example, we associate with the DESTINATION role of the **goto-copier** step the set of copiers which the system has encountered thus far.

The decision points serve both to segment and sort the experience of the agent. Each decision point collected together all the related segment of episodes. Thus in Figure 1 the node having to do with lifting the cover of the photocopier has attached to it various differing episodes in the lifetime of the agent regarding lifting photocopier covers.

This organization serves several purposes. First, it segments each problem-solving episode into cases that are individually stored and indexed. This facilitates access to relevant parts of episodes for transfer to new situations. Second, the descriptors for the episodes both couple the routine to the agent's habitats and distinguish options at decision points. Lastly, the saved information allows episodes to be reconstructed from their descriptors.

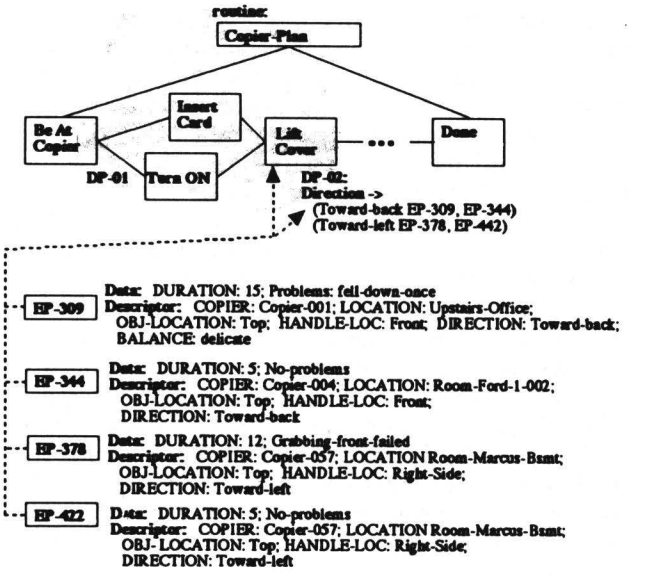

Figure 1: Decision Points Collect Together Related Episodes

Routine Behavior

An algorithm for acting using a multicase extends the adaptive planning algorithm described in (Alterman, 1988). SCAVENGER acts by selecting from the menu of options provided by the multicase at each step (or other decision) the element most appropriate to the current situation. This is done by selecting the option associated with the most similar previous situation. Options are evaluated by comparing their case descriptors to the current situation; the option associated with the best-matched descriptor is selected. Since each DP lists only the options relevant at its point in the routine the cost of this process is limited.

Performance of a step has several parts. First its preconditions are checked and any missing preconditions are repaired adaptively. Then the step's actions – its substeps, if any, followed by its "primitive component" (a code fragment) – are performed. If a step fails, SCAVENGER will try to adapt; if that fails, it will either press on or propagate the failure to the encompassing routine. The expected duration of steps is monitored; steps that are judged to be overdue are treated as having failed. Next, the step's expected outcomes are verified and repaired if necessary. A next step is selected and the cycle repeats. Routine performance is complete when performance of the root node completes and its expected outcomes are achieved.

Adaptation methods that are used include observa-

tion of alternatives available in the situation, adaptation of the problematic condition, adaptation of the step, or insertion of additional steps. In general sources of adaptation can involve a grab bag of techniques, such as experimentation, weak-method search, causal analysis, and memory. The current implementation relies on within-category search, adaptation based on relative similarity, subgoaling, instruction interpretation, passive observation (for unexpected events) and active observation (for role filling). Overall the primary focus is on the use of comprehension-based adaptations. The value of comprehension is especially clear in FLOABN's use of instructions, labels (a label on my stereo identifies the volume knob), iconographs (the 1/0 logo identifying the ON/OFF switch on many copiers), or affordances (the volume knob affords sliding left and right motions) that might be available. Another form of comprehension (not currently implemented) would involve copying the actions of some other agent.

Enrichment of Routines

The system begins with a skeleton routine such as an agent might acquire by having the task explained to it or seeing it performed. Each additional detail arises from some specific experience. Some experiences add new paths (e.g. running out of paper), some add detail to existing paths (e.g. observing lighting and movement as copies are made), and some modify existing steps or decision criteria (e.g., learning where to look for a power switch). Most paths through the multicase access elements contributed by a number of distinct experiences.

Multicases allow detail to be acquired incrementally through the overlay of old episodes with newer ones, resulting in a gradual enrichment of the multicase. As a given multicase is applied to a new situation SCAVENGER enriches that multicase. Whenever a choice is made, the choice (if new) can either result in the addition of a new decision point or it can be added to the list of options for an existing decision point. During the test sequence we show in the empirical section of the paper, 50 new decision points are acquired and several hundred options.

This constant case acquisition has two important effects: performance requires less effort over time, and the routine becomes customized to the details of specific habitats. By "becoming a routine" we mean that one's increased familiarity with a given habitat reduces certain specific costs (e.g., effort and time) of the activity; in the next section we will describe specific quantitative consequences of this process. Once the telephone-call multicase is extended to include inserting a dime in a pay phone, that modification does not need to be made again, nor will it be necessary to spend time receiving and interpreting messages about inserting a coin. Once the office copier is identified, it is remembered as an individual and, in the future, effort spent in identifying it is virtually eliminated.

An example of this customization to the normal places of engagement is found in the telephone domain. We find that as time goes on we differentiate habitats by place-relevant properties. For example, we learn that the phone in an office is usually on the desk (wherever that may be) except perhaps in Roy's office, where its often fallen on the floor or buried under paper. We learn that if you need a phone number in Rick's office, you check the wall. Pay-phones are found in halls rather than in offices. There can be no "axioms of phones" that state these facts, for they are merely regularities of the agent's experience. Nevertheless, they are reliable within the agent's habitats and form a rational basis for behavior.

Quantifying Routine Behavior

We examined the evolution of routine behavior over a span of episodes by presenting FLOABN with a sequence of examples including telephoning, copying, and vending machine transactions and observing the evolution of each multicase in response. FLOABN was provided initially with three skeleton multicases, one for each type of task. It was then presented with a sequence containing 15 different situations of the three types, plus 15 variant situations (one of each of the first 15 situations), plus 20 repetitions of some prior scenario (e.g., calling home from the office) for a total of 50 episodes. There were on average 25 steps per episode, yielding in excess of 1200 episodic cases. Each run of the example sequence required approximately 8 hours on an 8Mbyte Macintosh IIx under Allegro Common Lisp. For each episode we collected over 50 items of data about the evolution of memory and routine performance, including such items as the size and composition of decision points, the composition of plan memory, information usage and flow, and dynamic information on the evolution of routine structure.[3] Here we examine the change in three measures of effort involved in activity across the example sequence: the number of features observed, the amount of searching required, and overall workload.

Feature Extraction Effort. Earlier we discussed characteristics of of habitat-based activity. One difference was in the number of situation features to which the agent had to attend. In non-routine situations the agent has to expend effort determining the values of the significant features of the situation and sorting through irrelevant features, having less experience to focus his search. In his habitats, however, the agent already knows the values of many of the important features, and has a good idea of where to look for what he does not know. Figure 2 shows the number of features

[3]This paper focuses on the measures of information flow. Memory usage and decision effort are discussed briefly; more detail can be found in Zito-Wolf & Alterman (1993).

attended to by FLOABN during each episode. This measure counts all features – object existence, object properties, and object relations – FLOABN accessed from the situation (as opposed to from memory) in the course of activity. We expect this number to be reduced by routinization because as FLOABN becomes familiar with a situation, (1) relevant situation features are more likely to be available in memory; and (2) the number of features needing to be examined at all is reduced because fewer judgements need to be made and attention is better focused. There is indeed a clear reduction in attentive effort.

Active Looking for Particular Objects. We measured the amount of active looking that FLOABN does in a given situation, that is, the number of times FLOABN had to attend to some specific object. Such attending generally occurs in the process of locating a suitable object to fill some plan role, such as COPIER, COIN-SLOT, or PHONE. This correlates somewhat with the number of features attended, since object location is typically followed by feature examination to determine how well it matches the desired object or object type. However, it differs in that object searching is a measure of how many times the agent had to select and discriminate objects rather than a count of the features extracted from a situation. An object attended in two different searches is counted twice. The familiarity afforded by habitats helps reduce the amount of active looking in the course of an activity (Figure 3).

Agent Workload. To measure the overall change in effort expended in performing a task in a given situation due to routinization, we created a general measure of workload that combines effort of several types. In addition to the above measures of features extracted and active looking, this measure takes into account the number of steps performed and the number of adaptations made. These different components have been weighted to reflect differences in the mental effort they require:

$$\begin{aligned}\text{Workload} = \\ &\text{StepsPerformed} + \text{FeaturesAttended} + 2 * \text{FeaturesStored} \\ &+5 * \text{ActiveLooks} + 10 * \text{Adaptations}\end{aligned}$$

Feature storage is more expensive than attending a feature because storage in FLOABN involves first attending to a feature and then either building a new DP or modifying an existing one. Active looking involves both attending and storage of features, and it also has the additional costs of visual search. Finally, adaptation typically involves a significantly larger amount of cognitive effort than the performance of a single step because, at a minimum, it involves state-space search and evaluation. We feel that this latter number is fairly conservative as some adaptations can be fairly lengthy, involving time consuming processes like experimentation and the interpretation of instructions. The results are given in Figure 4. Routinization of behavior clearly reduces the workload, and the reduced slope of the graph toward the right shows that this difference is increasing as the situations of activity become more and more routine.

Discussion. Our results suggest that adapting to one's habitats is effective in reducing agent effort in several distinct ways: active looking, features examined, memory load, and plan modification (adaptation) effort. Intuitively it is reasonable that familiarity with a situation and an activity reduces the effort of performing that activity. However, the study of skill acquisition in Artificial Intelligence and Cognitive Science has typically focused on the role of internal transformations of knowledge (e.g., proceduralization in ACT (Anderson, 1983) and chunking in SOAR) and has largely ignored the effects of improvements in the coupling of the agent with his environment. In our view, skill acquisition is a phenomenon arising from several sources. ACT and SOAR have demonstrated that power law behavior can result from the mental transformations of knowledge such as proceduralization and chunking; Agre and Shrager (1990) have shown that power-law behavior can arise from the accumulation of local optimizations to an activity. Our work shows that similar performance improvements can arise from improving the fit between the agent and his world.

Summary and Conclusions

The research described in this paper has explored the relation of agent and environment under four conditions: 1) an individual agent is part of a community of agents; 2) individual agents have routines; 3) routines are grounded in the habitats of the individual agent; 4) some of these routines are common practices. Under these assumptions, we have shown how the enrichment of routines with details of the agent's habitats acquired through problem-solving experience can reduce both problem-solving and perceptual effort.

References

[1] P. Agre and J. Shrager. Routine evolution as the microgenetic basis of skill acquisition. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pages 694–701, Hillsdale, NJ, 1990. Lawrence Erlbaum Associates.

[2] Richard Alterman. Adaptive planning. *Cognitive Science*, 12:393–421, 1988.

[3] Richard Alterman, Roland Zito-Wolf, and Tamitha Carpenter. Interaction, comprehension, and instruction usage. *Journal of the Learning Sciences*, 1(4), 1991.

[4] John R. Anderson. *The Architecture of Cognition*. Harvard University Press, Cambridge, MA, 1983.

[5] Rodney Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

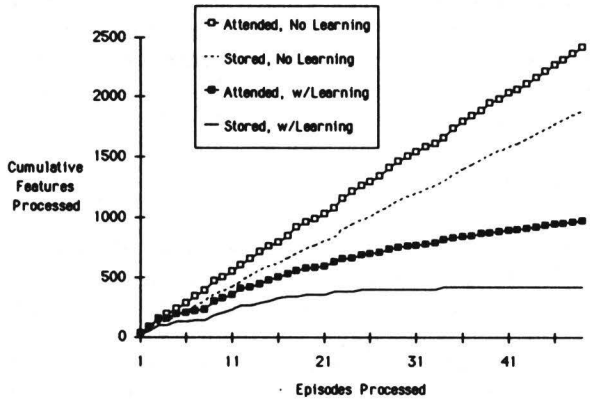

Figure 2: Features Observed per Episode

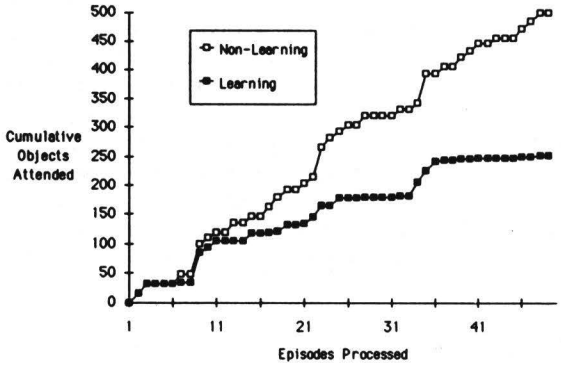

Figure 3: Objects Examined Per Episode

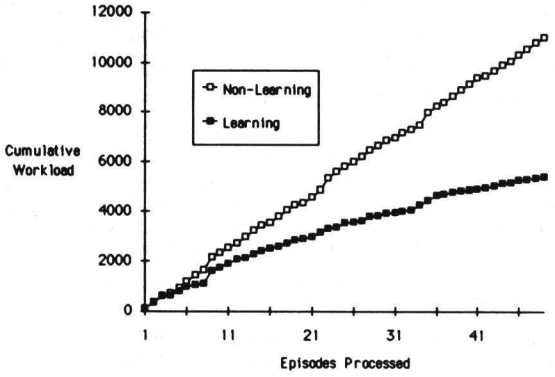

Figure 4: Effort Required per episode

[6] Tamitha Carpenter and Richard Alterman. Reading instructions. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ, 1991. Lawrence Erlbaum Associates.

[7] Tamitha Carpenter and Richard Alterman. Reading as a planned activity. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ, 1993. Lawrence Erlbaum Associates.

[8] Michael Cole. Cultural psychology: A once and future discipline? In J.J. Berman, editor, *Nebraska Symposium on Motivation, 1989: Cross-cultural perspectives, (Vol. 37)*. University of Nebraska Press, Lincoln, Nebraska, 1989.

[9] David B. Guralnik, editor. *Webster's New World Dictionary of the American Language, Second College Edition*. World Publishing Company, New York, 1970.

[10] Kristian J. Hammond. Case-based planning: A framework for planning from experience. *Cognitive Science*, 14:385–443, 1990.

[11] Janet L. Kolodner and Robert L. Simpson. The MEDIATOR: Analysis of an early case-based problem solver. *Cognitive Science*, 13:507–549, 1989.

[12] M. Minsky. A framework for representing knowledge. In P.W. Winston, editor, *The Psychology of Computer Vision*, pages 211–277. McGraw-Hill, New York, NY, 1975.

[13] Edwina Rissland and Kenneth Ashley. Hypotheticals as heuristic device. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 1986.

[14] D.E. Rumelhart. Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, and W. Brewer, editors, *Theoretical issues in reading comprehension*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1980.

[15] R. Schank and R. Abelson. *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1977.

[16] R. Wilensky. Meta-Planning: Representing and Using Knowledge About Planning in Problem Solving and Natural Language Understanding. *Cognitive Science Journal*, 5:197–233, 1981.

[17] Roland Zito-Wolf. Case-based representations for procedural knowledge. Unpublished doctoral dissertation, 1993.

[18] Roland Zito-Wolf and Richard Alterman. Multicases: A case-based representation for procedural knowledge. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 331–336. Lawrence Erlbaum Associates, 1992.

[19] Roland Zito-Wolf and Richard Alterman. A framework and an analysis of current proposals for the case-based representation of procedural knowledge. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*. Lawrence Erlbaum Associates, 1993. To appear.