What are routines good for?*

Henry H. Hexmoor

Department of Computer Science
SUNY at Buffalo, Buffalo NY 14260
hexmoor@cs.buffalo.edu

February 20, 1994

Abstract

Routines are patterns of interaction between an agent and its world. Getting in or out of a car, changing lane, and flipping pages of a book can be routines for an agent if the agent consistently engages in these activities in a similar way. I.e., a task for an agent is a routine if the agent that has choices about how to accomplish that task, nevertheless does it in the same way. Consistently putting on the left leg of pants before putting on the right leg would be a routine for an agent.

A routine is either imposed upon the agent (a plan at the conscious level to be followed), in which case it need not be discovered, or performed by the agent automatically, i.e., unconsciously. The latter may or may not ever be discovered, i.e., noticed and made conscious. However, the existence of such a routine may guide the agents actions. If it remains unconscious, it aids in choosing among competing actions unconsciously as an unexplained tendency or a preference. If it is noticed and made conscious, it can also be used as a concept and used in reasoning about actions and planning. We show how agents with routines a) use them to guide their everyday activity, b) use them to enrich their abstract concepts about acts.

1 Introduction

We see instances of routines as by products of research in planning, artificial life, and architectures for intelligent agency. Routines have not been direct objects of study. Colloquial use of the term "routine" dominates reference to this word in the relevant literature. In our work we try to understand and find ways in which routines can be formalized as well as their impact on learning about act, actions, and skills.

"Routines are patterns of interaction between an agent and its world" [AC87]. Getting in or out of a car, changing lanes, and flipping pages of a book can be routines for an agent if the agent consistently engages in these activities in a similar way. I.e., a task for an agent is a routine if the agent that has choices about how to accomplish that task, nevertheless does it in the same way. Consistently putting on the left leg of pants before putting on the right leg would be a routine for an agent.[1]

In [Agr88], a routine is described as a frequently repeated pattern of activity between an agent and the world, e.g., picking up a fork, making a bowl of cereal, putting on a heavy backpack, selecting a toll booth at the San Francisco Bay Bridge, putting a watch on, breaking an egg into a bowl, washing dishes in the kitchen sink after a large dinner, tossing a wad of paper into the trash can in the far corner of the office, and writing the word "the". Agre points out that he might not have any fixed way of pouring liquids, but in his morning stupor he might have set way of pouring his first cup of coffee. He goes on to emphasize the improvisational aspects of routines. However, he never gets around to suggesting that we learn from these routines. Agre does not examine the evolution of a routine from its

*This work was supported in part by Equipment Grant No. EDUD-US-932022 from SUN Microsystems Computer Corporation.

[1] We assume the agent has no medical or otherwise conditions that make it choose one leg over another.

constituent habitual preferences. Agre does not look at mechanisms for identifying and using skills by the agent. I sketch learning techniques that address learning skills from routine activities. Agents that have routines for tasks, use those routines to guide their activity.

Regardless of benefits of routines, they are either conscious, at KL, or unconscious, at PML. A conscious routine is a set of policies, a plan, or a script for doing something. An unconscious routine is a sequence of actions, as if the sequence is just a complex action.

A course of action or a set of tendencies for choosing among alternative actions can be a routine that is useful in many occasions and is part of many tasks. On the one hand, such a course of action can be considered to be a skill.[2] Since routines as skills are not plans, there can be no guarantee for success of a routine. Since we are concerned with every day activities, we are considering situated actions [Suc88], i.e., actions that an agent exhibits without reasoning or making conscious decisions, ergo unconscious behaviors.

On the other hand, a routine might be a plan, a set of policies, or a script to be followed.[3]

Following the research in situated activity [Suc88, AC87, Nil92], we would like an agent to operate with situated actions instead of planned actions. In the larger context of intelligent agency, we argue that situated activity is a particular form of interaction between an agent and its environment that is produced by an agent when it is involved in unconscious interaction [HLS92, HLCS93, HCBS93]. Elsewhere, in [HS93], we discuss that unconscious actions are part of an agent's skills as opposed to its knowledge. When routines are unconscious they can be considered to be skills. On the other hand, an agent may also engage in conscious, reasoned out routine.

We have developed an architecture that facilitates modeling conscious and unconscious interactions of an agent with the world. Our architecture is GLAIR¡ (Grounded Layered Architecture with Integrated Reasoning) [HN92, HLS92, HCBS93, HLCS93, HLS93, LHS93, HS93]. GLAIR is an architecture for intelligent autonomous agents that endows them with physical abilities that are "natural" for the agents, i.e., correspond to their physiology.

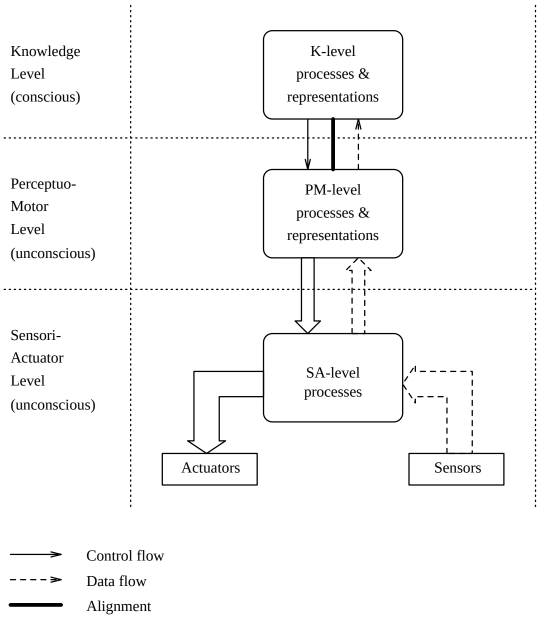

Figure 1: Schematic representation of the GLAIR architecture. Width of control and data paths suggests the amount of information passing through (bandwidth). Sensors include both world-sensors and proprio-sensors.

Figure 1 schematically presents GLAIR. GLAIR architecture is characterized by a three-level organization into a Knowledge Level (KL), a Perceptuo-Motor Level (PML), and a Sensori-Actuator Level (SAL). GLAIR is a general multi-level architecture for autonomous cognitive agents with integrated sensory and motor capabilities. GLAIR offers an unconscious layer for modeling tasks that exhibit a close affinity between sensing and acting, i.e., behavior based AI modules, and a conscious layer for modeling tasks that exhibit delays between sensing and acting. GLAIR provides learning mechanisms that allow for autonomous agents to learn emergent behaviors and add it to their repertoire of behaviors. GLAIR motivates the concept of *embodiment*.

In light of our distinction between conscious and unconscious interaction, we will specify the impact of routines on conscious and unconscious interactions with the environment. In general, we are interested in ways in which

[2] In computational terms this is like a macro. We consider skills to be routines at the unconscious.

[3] We consider plans/scripts to be routines at the conscious.

an agent discovers, learns, and uses routines. The following is our definition of a routine: "A conscious course of action or a set of unconscious preferences, that is habitually or by choice adhered to is a routine."

The above definition assumes an agent that lives in a world and performs a set of tasks that are part of its life. We assume the agent to start out with no apriori plan or preference for a course of action. By engaging in the tasks, the agent develops preferences for actions. Given a task, if the agent has determined preferences for all possible choice points in the space of possible actions, we say the agent has a routine for the task. The world of the agent may or may not be static or changing. If an agent develops a routine in a particular world and then the world changes, the routine maybe lost. For example if an agent develops a routine for finding an object in a cluttered but static room, it basically has learned the layout of the room.

If we now move objects in the room and it develops another routine for finding objects, we would want the new routine to work in the previous setting as well. I.e., we would like routines to endow the agent with ever richer preferences. Similarly, if the world becomes dynamic, i.e., the agent interacts with moving objects, new routines developed should cope with previously static situations. However, we don't expect the routines be so robust that we can change the agent's morphology, e.g., change its sensors and actuators, and expect the new routines to cope with the changed morphology.

Since a routine is a course of action or a set of preferences, within this course of action, there will be subsets of action courses that may remain the same even across morphology changes. For example, an agent that maneuvers around corners in a smoothed arc motion instead of a stop-turn-start sequence should be able to keep this ability over many task spaces and many minor morphological changes.

Many recent works in mobile robotics have addressed learning isolated actions, particularly through neural network approaches, e.g. [Neh92]. In our approach, we don't isolate actions for learning. Instead we focus on emeregnt sequences of actions that are parts of larger tasks and unspervised learning of these action sequences. Furhermore, sequences of actions learned by the agent become significant for the agent after being learned and the agent generates concepts corrseponding to these sequences to be used in subsequent interactions or reasoning.

In GLAIR we have developed a mechanism for unconscious behavior generation we call Perceptuo-Motor Automata, PMA. GLAIR and PMA are computational techniques designed to study improvisations and learning which is commonplace in many routine activities.

1.1 Uses of Routines

There are times that we notice our habitual behaviors, like tying a shoe lace, in a particular way. Why or when we form habits are beyond the scope of what we are interested. Here we seek ways to use and at times discover routines that are habitual.[4]

We expect discovered and learned routines to be used in one of three ways: a) as a mental notion to be reasoned about, b) as stand alone unconscious pieces of behavior, c) as tendencies for choosing certain commands over other applicable commands.

1.1.1 Routines Are Used to Guide Daily Activity

A routine is either imposed upon the agent (a plan at the conscious level to be followed), in which case it need not be discovered, or performed by the agent automatically, i.e., unconsciously. The latter may or may not ever be discovered, i.e., noticed and made conscious. However, the existence of such a routine may guide the agents actions. If it remains unconscious, it aids in choosing among competing actions unconsciously as an unexplained tendency or a preference. If it is noticed and made conscious, it can also be used as a concept in reasoning about actions and planning.

An important feature of GLAIR is that it allows knowledge to migrate from one level to another and in doing so changes representation in order to be consistent with the representation used at that level. The underlying assumption here is that parts of an agent's knowledge about its world have natural loci either in the conscious or the unconscious. As

[4] By discovering a routine, we mean making a conscious note of it and having a concept for that activity.

the agent interacts with the world, the knowledge may be gained at some level and later find its natural locus at a different level and migrate there. The migration of knowledge from conscious level to unconscious level often is a result of caching reasoning that takes place at the conscious level. So what is migrated to the unconscious is cached reasoning or just chunks of "what to do when". This is often known as chunking in psychological terms. Chunking allows the agent to quickly pick an action without deliberation. This lack of deliberation is a form of lessening the cognitive burden of deciding "what to do next". Naturally this lessening cognition aids in guiding the agent's actions.

We are interested in all routines including the ones that are not discovered but adopted. Reading a driver's manual we are told about passing another vehicle and although that may not be the only way to pass another vehicle we adopt it. Once our conscious knowledge of passing another vehicle is practiced it becomes part of our skills and in our framework of GLAIR [HLS92] we consider it to have become unconscious. We can consider this process of learning as a form of caching a plan into reactive plans as was done in [Sch87, Dru89].

Another way that the agent's actions is guided by routines is by refining skills at the unconscious level. There are many routines that need not have a concept associated with them at the conscious level, but nevertheless are useful in guiding actions. These are tendencies that the agent may fall into and stay with. Informally, we consider this use of routines as skill acquisition.

1.1.2 Routines are Used to Enrich an Agent's Knowledge and Skill with Acts

Beyond their capacity to guide an agent's activity, routines can enrich an agent's ability to interact with the world. This enrichment takes two forms. The first is to enrich the agent's internal model of its acts. This means that the agent learns more about its abilities and represents these abilities in its repertoire of acts. The second method is an enrichment that does not affect the agent's explicit model of the world. I.e., no modifications are made to what is represented. Instead, the agent improves its ability in choosing among alternate actions. With both methods, the agents learns more complex acts and behaviors.

2 Routines and Learning Techniques

It is hard to conceive of routines as things *to adopt* or *to reason about*. We consider routines to be unconsciously acquired through inductive learning.[5] Inductive learning is drawing inductive inferences from observed facts in the environment or by being told. A recent trend in inductive learning, especially as applied to planning, is reinforcement based learning, RBL. Reinforcement learning is a trial-and-error process where, upon executing an action, applicability of actions to situations are updated based on the reinforcement signal received in response to the action. An action elicited by a stimulus is known as a *respondent* [Hen69].[6] This is in contrast to actions controlled by their consequences [Hen69].[7]

For an RBL application, a set of states, a number of actions, and a reward system that determines rewards for being in a state are given. The connection between actions and states is not given or it is underspecified. This makes deciding what action to perform given a state nondeterministic. After running an RBL system, it learns a preferences for choosing actions given states that the system finds itself. This process reduces nondeterminacy of choosing actions.

In terms of GLAIR, we may have a situation where the agent's KL is rich and PML is lean. As the agent interacts with the world, it enriches its PML with decisions made at KL. However, some decisions may add nondeterminacy to PML. For instance if at PML for situation S there is a corresponding action A1 but the next time KL perceives situation S it decides on action A2, PML adds this nondeterminacy to its PMA. In the sense that conscious knowledge at KL migrates to unconscious in PML, we consider this as a form

[5] As opposed to analytic learning or deductive inference.

[6] Respondent conditioning is also known as instrumental learning. An example of respondent conditioning is laughter as a consequence of being amused.

[7] Operant conditioning is also known as classical conditioning. An example of operant conditioning is laughing to the boss's jokes.

of automaticity.[8] We have shown this type of learning in our implementation of a World War I style dogfight simulation game we call Air Battle Simulation (ABS). "Gabby," is the agent in this game which stands for "GLAIR air battler." Initially, Gabby has not acquired a PMA for the game yet, and so uses conscious level reasoning (i.e., SNePS behavioral rules [SR87]) to decide what move to make [HCBS93]. This is the case of a routine activity leading to becoming automatic. The routine is not being learned through automaticity. This form of learning does not encode or discover a routine but makes use of it.

What is developed with RBL is a connection between situations and actions. For example, given states S1 S2 S3 and Actions A1 A2, let's assume no preference for actions, e.g., given state S2, doing A1 or A2 are equally preferred. Let's assume that S3 is the most desirable state and S1 is the least desirable state. After RBL, we may learn that in state S1 action A1 is preferred and in state S2 action A2 is preferred. If the goal is to get to S3 and we always start at state S1, the system will repeat A1 by A2 and becomes deterministic. We can call A1 followed by A2 or visiting S1 and S2 a routine. OfCourse if we make the STRIPS assumption[9] and if wee assumed that the reward system stays the same, after RBL settles down and resolves all nondeterminacy, a routine can be detected by observing the recurring sequence of actions. In this example, A1 A2 as a repeating sequence is a routine.

Any system that changes goals will have a dynamic reward system. E.g., if the initial goal is to get to state S3, reward for getting to that state will be highest. Now if the goal changes to get to state S2, reaching state S2 will have the highest reward. Another agent maybe involved and the interaction of the agents determine the resulting state. E.g., after executing action A1 if the agent normally finds itself in state S2, due to another agent's interference it might find itself in state S1 or state S3. Let's extend our simple model to include goals. E.g., let's have G1 and G2 as two goals. Action sequences maybe recurrent under one goal or both. The agent may detect recurring sequences of actions common in multiple goals as routines.

Apart form multiple goals, there may not be rewards provided to make the system completely deterministic or the reward system dynamically changes. When rewards are dynamic or underspecified, the agent may detect recurring sequences of actions common under varying reward systems as routines.

An agent may have a rich PML and a thin KL. In this case, it is possible that the agent realized connections between its disparate pieces of knowledge. In GLAIR, knowledge level actions are decomposed in PML to more primitive behaviors and actuator commands. The agent relates a primitive action in KL with a behavior at PML. If PML takes over an agent's actions and the KL is inactive [10] the agent may notice changes of automatic behavior appropriate for the task. These changes correspond to switching between primitive actions in KL. At the KL, the sequence can be learned.

2.1 Summary of Learning Routines

- When goals in a system change, the corresponding rewards change. The agent may detect recurring sequences of actions common in multiple goals as routines. The learned routine will be in PML.

- When rewards are dynamic or underspecified, the agent may detect recurring sequences of actions common under varying reward systems as routines. The learned routine will be in PML.

- Learn connections between disparate pieces of KL corresponding to a rich PML. What is learned is a routine through migration of knowledge from PML to KL. The learned routine will be in KL.

[8] Automaticity is a psychological term. An agent experienced with an activity automatically (unconsciously) picks acts and execute them.

[9] Nothing changes unless the agent changes it

[10] This is like an skillsman working on something without reasoning.

3 Routines, Learning, and when Gerry Catches Garry

Gerry is a mobile robot.[11] It is 2.5' tall. Gerry is controlled by a 6.270 board.[12] Gerry has two independently driven wheels, an infrared transmitter, 4 infrared sensors, a front and a rear bumper, and four corner whiskers. Gerry is designed to have four primitive abilities of going straight, turning right, turning left, and stopping. When a Whisker or a bumper senses a contact with external objects, Gerry performs a reflex. There are reflexes uniquely corresponding to each whisker and bumper. Reflexes are implemented in SAL of Gerry's GLAIR. At PML of Gerry's GLAIR, behaviors are implemented using PMA.

Currently, Gerry has PMAs and some basic knowledge at KL for searching for another robot, Garry.[13]. Garry is 1' tall. Garry is controlled by a 6.270 board. Gerry has two independently driven wheels, an infrared transmitter, 2 infrared sensors, and three front bumpers/whiskers. When a Whisker or a bumper senses a contact with external objects, Garry performs a reflex.[14]

Figure /reffig:gerry-glair shows Gerry's KL knowledge to consist of three separate pieces of knowledge for spotting Garry, catching up with Garry, and grabbing Garry. At the KL, Gerry doesn't know that these are parts of a task for getting Garry. Perceptual reduction is implemented by having concepts at KL correspond to a pattern of stimulus at PML. For example, "no-garry" at KL corresponds to the pattern of no infrared being sensed at PML. At KL, no concepts for specific infrared signals as in PML exist. Each action in the KL is expanded into a PMA in KL. This is an implementation of action elaboration where actions at the PML are finer compared to KL actions.

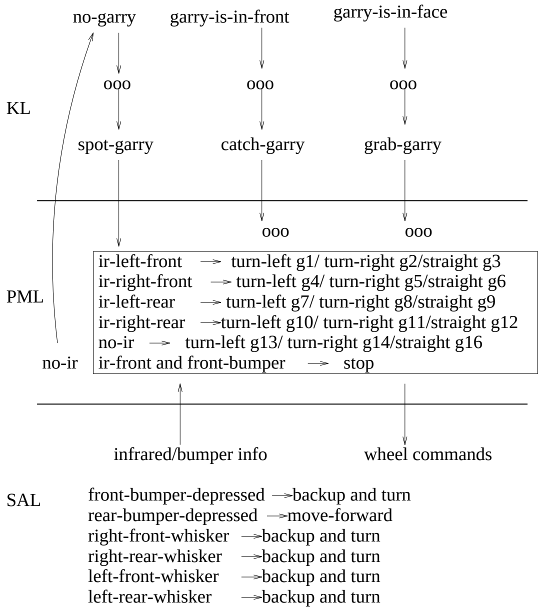

Figure 2: Schematic representation of the Gerry's GLAIR components. the box in PML shows a PMA for spotting Gerry, there are others for catching up with Garry and grabbing garry. gi in the PMA box next to each action represents goodness of action in the context of corresponding situation.

At SAL, Gerry has 6 reflexes for avoiding obstacles. Figure /reffig:gerry-glair shows pneumonic form of these reflexes. The actual implementation of actuator commands are in terms of wheel commands.

Concurrent behaviors at PML are allowed and also concurrent reflexes are possible. Behaviors and reflexes produce wheel commands. All wheel commands are given a priority level when they are generated. PML wheel commands have lower priority than the ones for reflexes. One exception is that behaviors have higher priority than reflexes if a reflex needs to be supressed. For example, to catch Garry, Gerry needs to have its front bumper touch Garry and sense high levels infrared in its front sensors. Normally touching an object with the front bumper trigger a reflex to backup and turn. However, we want to suppress this reflex. There is a module at SAL for arbitration and combination of all wheels commands. This module examines flags for suppression and if

[11] Gerry was a Omnibot 2000 toy robot before we converted into a research vehicle, gave it a brain, and called it Gerry. The first letter "G" stands for GLAIR and the two "r"s in the middle stand for roving robot.

[12] 6.270 boards are robot controller boards developed at MIT for use in their 6.270 class offering. 6.270 boards have a 6811 microprocessor with multitasking capabilities.

[13] Garry is a Lego robot. The first letter "G" stands for GLAIR and the two "r"s in the middle stand for roving robot.

[14] Details of Garry's implementation are irrelevant here.

reflexes are not suppressed, it sums the reflex generated wheel commands and passes it to the wheels. If no reflex is active, it takes all behavior generated wheel commands and sums and passes them to the wheels. If reflex suppression is turned on, it ignores reflexes and attends to the behavior generated commands.

We have programmed PMAs and the knowledge at KL so Gerry will look for Garry and will stop when it finds Garry. At the PML we are running reinforcement based learning to resolve nondeterminacy. Gerry's actions can be controlled either completely by KL or by PML.

We intend to experiment with Gerry learning a routine of learning by migration of knowledge from PML to KL. The routine we intend Gerry to learn is that when it sets on to spot Garry it follows through by catchig up and grabbing Garry. This is a realization for Gerry that its three actions at KL are part of the same routine. We want Gerry to produce a concept for the routine and that it consists of spotting, catching up and grabbing Garry.[15]

4 Conclusion

We outlined routines in general as arising from interaction between an agent and its world. Routines maybe conscious of unconscious. A routine may enrich an agent's conscious knowledge or increase its unconscious capacity for interacting with the world. Routines guide interaction with the world by lowering cognitive burden of deciding actions. We described how our architecture GLAIR facilitates demonstration of learning routines. We described ongoing research with Gerry, our mobile robot implemented using GLAIR.

References

[AC87] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of AAAI-87, Seattle, Wa.*, pages 268–272, July 1987.

[Agr88] Philip Agre. The dynamic structure of everyday life. Technical Report 1085, MIT Artificial Intelligence Laboratory, MIT, 1988.

[Dru89] Mark Drummond. Situated control rules. In *Proceedings Of the first conference on principles of knowledge representation*, pages 103–113. Morgan Kaufman, 1989.

[HCBS93] Henry Hexmoor, Guido Caicedo, Frank Bidwell, and Stuart Shapiro. Air battle simulation: An agent with conscious and unconscious layers. In *University of Buffalo Graduate Conference in Computer Science-93 (TR93-14)*. Dept. of Computer Science, SUNY at Buffalo, New York, 1993.

[Hen69] Derek P. Hendrey. *Conditioned Reinforcement*. The Dorsey Press, 1969.

[HLCS93] Henry Hexmoor, Johan Lammens, Guido Caicedo, and Stuart C. Shapiro. Behavior Based AI, Cognitive Processes, and Emergent Behaviors in Autonomous Agents. In *Proceedings of AI in Engineering*, 1993.

[HLS92] Henry Hexmoor, Joe Lammens, and Stuart Shapiro. An autonomous agent architecture for integrating perception and acting with grounded, embodied symbolic reasoning. Technical Report CS-92-21, Dept. of Computer Science, SUNY at Buffalo, New York, 1992.

[HLS93] Henry Hexmoor, Johan Lammens, and Stuart C. Shapiro. Embodiment in GLAIR: A Grounded Layered Architecture with Integrated Reasoning. In *Florida AI Research Symposium*, pages 325–329, 1993. also available as SUNYAB CS Dept. TR93-10.

[HN92] Henry Hexmoor and Donald Nute. Methods for deciding *what to do*

[15] Naturally, since we want this learning to be done by Gerry the name for the routine will be anonymous (a "gensym" in Lisp terms). Semantics of the routine for Gerry will be in terms of PMA.

next and learning. Technical Report AI-1992-01, AI Programs, The University of Georgia, Athens, Georgia, 1992. Also available from SUNY at Buffalo, CS Department TR-92-23.

[HS93] Henry Hexmoor and Stuart C. Shapiro. Examining the expert reveals expertise. In *3rd International Workshop on Human and Machine Cognition: Expertise in Context*, 1993. to appear.

[LHS93] Johan Lammens, Henry Hexmoor, and Stuart C. Shapiro. Of elephants and men. In *Proceedings of NATO-ASI*. Springer Verlag, 1993. to appear.

[Neh92] Ulrich Nehmzow. *Experiments in Comptence Acquisition for Autonomous Mobile Robots*. PhD thesis, University of Edinburgh, 1992.

[Nil92] Nils J. Nilsson. Toward agent programs with circuit semantics. Technical Report STAN-CS-92-1412, Stanford University, Stanford, CA., 1992.

[Sch87] Marcel J. Schoppers. Universal plans for unpredictable environments. In *Proceedings 10th IJCAI*, pages 1039–1046, 1987.

[SR87] S. C. Shapiro and W. J. Rapaport. SNePS considered as a fully intensional propositional semantic network. In N. Cercone and G. McGalla, editors, *The knowledge frontier: essays in the representation of knowledge*, pages 262–315. Springer, New York, 1987.

[Suc88] Lucy A. Suchman. *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge University Press, 1988.