In AAAI Fall Symposium Series: Sensory Aspects of Robotic Intelligence. Monterey, CA; Nov 15-17, 1991

Intelligent Agent Design Issues:

Internal Agent State and Incomplete Perception

Lonnie Chrisman
Computer Science
chrisman@cs.cmu.edu

Rich Caruana
Computer Science
caruana@cs.cmu.edu

Wayne Carriker
Electrical and Computer Engineering
carriker@ece.cmu.edu

Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Given the present diversity of architectures for intelligent agents it is important that we start identifying those environmental characteristics that impact architecture design. This would benefit architecture designers and those trying to understand and evaluate existing architectures. Unfortunately. the “Turing Tarpit” discourages precise characterizations. Nonetheless, we attempt here to make a precise characterization. We explore the situation where immediate perception of *all* pertinent aspects of the world is incomplete. Under this ubiquitous condition, we claim that internal state is necessary to compensate for sensory limitations. We identify types of internal state and examine the relationship between the characteristics of the environment and the appropriateness of each type of internal state for an agent.

1 Introduction

The number of cognitive architectures in Artificial Intelligence is large and rapidly growing. The experience of the community as a whole seems to indicate that a single universal architecture is not near on the horizon. Rather, each architecture appears to be tailored for particular environmental and task assumptions and excels under those conditions. Unfortunately, a clear elucidation of environment characteristics and a clear mapping of these to architecture design does not exist.

Elucidating the relationship between domain characteristics and architecture design would be valuable to designers. Such an elucidation could provide a basis for characterizing diverse architectures, understanding why (and where) they are effective, identifying their strengths and weaknesses, and enabling meaningful comparisons between competing methodologies. It could also be used to suggest and justify critical design decisions.

While the utility of this enterprise is obvious, few guidelines currently exist. One major obstacle to elucidating the relationship between domain characteristics and architecture features is the “Turing Tarpit”. This term, coined by Tom Mitchell, expresses the fact that all sufficiently powerful architectures are essentially Turing equivalent. Thus any behavior could. in theory, be implemented in any existing architecture. This discredits attempts to make hard claims about the inadequacy or superiority of any given approach. In other words, any argument such as “it is not possible to do X with architecture Y” can usually be shot down. The result is that, any arguments supporting one design choice over another must rest upon judgments of implementation complexity, naturalness, or other subjective measures. In this paper, we concentrate on a single aspect of intelligent agent design, the types of internal state employed by the agent. The goal of the paper is to identify how environment and task dimensions suggest the appropriate form of internal state that an agent should ideally use. We restrict the discussion to the situation of incomplete perception and claim that internal state is essential if the agent is to remain effective when perception is incomplete. The discussion, therefore, is centered around the use of internal state to compensate for incomplete sensing.

2 Complete and Incomplete Perception

Is it true that "the world is its own best model"?

Stated simply, perception is incomplete if at any instant all relevant features of the world cannot he observed. Features are relevant when the appropriate choice of action for a situation depends upon those features, and an action choice is appropriate if that choice leads to behavior comparable to or better than that resulting from the other possible choices.

Proponents of the reactive and situated action paradigms often tout the slogan "the world is its own best model" [Brooks, 1991]. Even if this is true, it ceases to be operational once the world can no longer be perceived completely at every instant in time. Where perception is always complete, the slogan may be justified, but with limited perception. the immediate observation stream cannot be sufficient for determining the next action. This leads to the following two conjectures:

Conjecture 1 *If perception is complete, then an agent whose actions are a function of percepts only (i.e. a purely reactive agent) can potentially be effective in all environments.*

Conjecture 2 *An agent with incomplete perception must maintain internal state information to remain effective.*

A few qualifications are warranted here. We use the term "purely reactive agent" to refer to an agent whose actions are a function of *current percepts only.* Note that most implemented "reactive" systems actually do choose their actions on the basis of some internal state in addition to their immediate percepts (cf. [Brooks, 1986], [Brooks, 1990], [Kaelbling & Rosenschein, 1990]). We justify our use of this term by asserting that the essence of reactivity accepts as its ideal what we have termed the purely reactive agent. A second point is that Conjecture 1 ignores aspects of implementation complexity which could justify (but not mandate) something more than a purely reactive agent. And finally, while internal state is necessary to remain effective, Conjecture 2 makes no claim that it is necessarily *possible* to remain effective when perception is incomplete. In this paper we only consider worlds where it is possible to compensate for perceptual incompleteness.

For agents such as mobile robots, the incompleteness of perception is nothing short of ubiquitous. There are a variety of reasons for this:

1. Sensor limitations:

   (a) Field of view.
   (b) Sensing at a distance and resolution.
   (c) Unreliable sensor readings.

2. Physical Obstructions: Walls or other opaque objects obstruct sensory capabilities.

3. Monetary cost of sensors: When building an agent, this may limit the ability to mount every desired sensor on the robot.

4. Computational (or energy) costs of sensing: When sensing operations are expensive or the number of sensing operations is large, an agent may have to intelligently forgo potentially relevant observations [Chrisman & Simmons, 1991], [Dean, 1990].

5. Mutually exclusive sensing: In some domains (e.g. X-rays or destructive sensing), certain sensing operations may prohibit the use of other relevant observations.

A proponent of reactivity might object to our ubiquity claim and instead propose that one ought to endow the agent with additional sensors until perceptual incompleteness is no longer significant. This objection, however, cannot possibly be sustained since the last four points above cannot be reasonably overcome in this manner. 'We conclude that perceptual incompleteness is a pervasive property of most autonomous agents in most domains.

2.1 Example: Tool Fetching

The following simple example should clarify the problem of incomplete perception and illustrate the various solutions discussed in this paper. Figure 1 depicts a broken object in the middle of a room near a toolshed. An agent, whose job is to repair broken objects, enters the room and discovers the object. To repair the object the agent must go to the toolshed, select an appropriate tool, and return to the object.

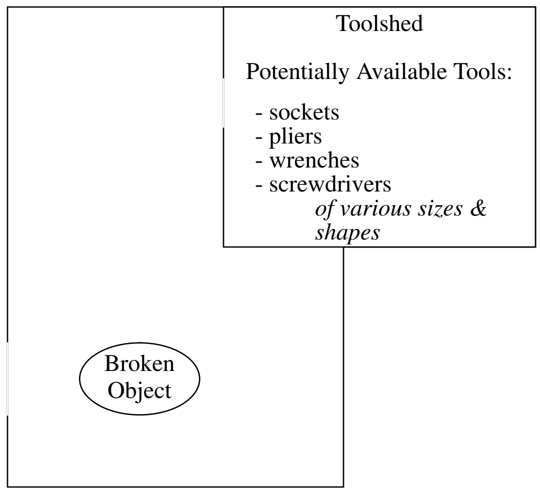

Figure 1: Example Domain: Tool Fetching

Perceptual incompleteness arises in this example in several places. First, when the agent turns its back to the object to head for the toolshed the object will move out of the agent's field of view and no longer be perceivable. An agent that does not maintain internal state would forget about the need to fetch the tool as soon as the object. was no longer visible! A less surmountable case occurs when the agent enters the toolshed and the walls completely obstruct the view of the broken object. At this point the robot must choose a tool appropriate for repairing the object, yet it no longer perceives the object nor the kind of damage that must be repaired.

We can identify three solutions for dealing with this particular incomplete perception problem, each corresponding to one of three general types of internal state that an agent might use.

The first consists of recording current world state. In this example, the agent records a description of the broken object. When the agent is inside the toolshed, it can select a tool based upon its internal description of the no-longer-visible object.

The second involves planning about the future. In this case, the agent might commit to the plan of entering the toolshed, grabbing a small phillips screwdriver, and returning to the object. The significant point here is that the agent commits to specific expectations about future world states (e.g., the availability of a small phillips screwdriver in the toolshed). Once in the middle of a plan it is not necessary for the agent to perceive the justification for that plan so the plan can carry the agent over perceptual voids.

The third method is to record that if the agent is ever in a situation where a phillips screwdriver is in view, it should pick it up. This approach differs from the second approach in that the agent makes no commitment to any future course of execution. It can be viewed as making all potential future decisions at the time the relevant percepts are available rather than at the time of encountering the choice point.

We conjecture that the three techniques above span the space of possible internal state types for compensating for sensory limitations. Furthermore, each approach can be characterized by the form of the internal state utilized. Finally, we claim that none of the above approaches is unconditionally better or worse than the other two; instead, the appropriateness of each depends upon various environmental parameters (which were unspecified in the example). The remainder of the paper is aimed at understanding this dependence.

3 Types of Internal State

In this section we explore three different types of internal state. We examine each type in broad terms - many instantiations for each type exist. Each type of internal state contains characteristics that separate it from the other two; moreover, all three types are mutually compatible such that an agent might make use of more than one. We conjecture that together, the three forms cover the space of useful internal states. The appropriateness of each form for dealing with perceptual limitations depends upon the particular characteristics of the task and environment.

We begin by examining each type of internal state individually. For each type of internal state we identify a representative agent that utilizes just that form of internal state. These agents are termed the *Model Reactive Agent*, the *Expectation Utilizing Agent*, and the *Contingency Anticipating Agent*. Examining each type of internal state separately aids in identifying the particular strength of each, thus giving us the necessary tools for examining hybrid cases.

3.1 Model Reactive Agent

The first type of internal state is remembering the current world state. A *pure* Model Reactive agent records only information that represents some aspect of the current world state; none of the internal state is used to represent future world states, future expectations, or decisions. As depicted in Figure 2, such an agent maintains an internal world model and chooses its next action based solely on this internal model. Note that this does not rule out sensing based reactivity because the internal model might faithfully reproduce all current sensing as well. Note also that a Model Reactive agent needs some predictive capability to track those parts of the world that it cannot directly perceive. Typically this predictive capability would be represented as implicit or explicit frame axioms that capture things like the fact that objects persist in time and do not become repaired automatically, etc.

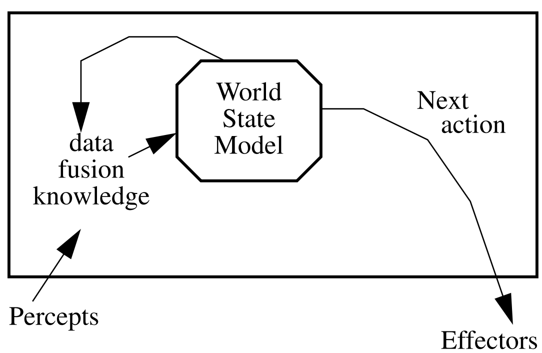

Figure 2: A Model Reactive Agent.

In the toolshed example, a Model Reactive agent memorizes the condition of the broken object either by an internal "image" of the broken object or via a simplified description). Subsequent behavior is a direct result of this stored information. Later when the agent has to decide what tool(s) to retrieve from the shed, it can refer to this internal model to determine which of the available tools seem most appropriate. The Model Reactive agent has an advantage over the other two pure agents when there are a large number of possible tools, when there is uncertainty about the tools that will be present in the toolshed, and when it is easy to remember the condition of the broken object (either because the object is simple or because there are a small number of possible broken states for it to be in). [Kaelbling & Rosenschein, 1990] view "[internal] state in the agent as carrying information content by virtue of its objective correlation with the environment," and thus their work provides an example of a primarily Model Reactive architecture.

3.2 Expectation Utilizing Agent

Agents that use only expectations about future world states to compensate for incomplete sensing are termed *Expectation Utilizing Agents*. Their distinguishing feature is that they use "plans" - courses of action selected in the present as a way of acting through spans of incomplete perception in the future. The agent selects plans (i.e. procedures) based on its current sensing, any plan that it might currently be executing, and its expectations about the future. By relying on plans, the agent need not base its present actions on just the perceivable current world state; all or part of the agent's actions may be the result of expectations deduced from previous sensing.

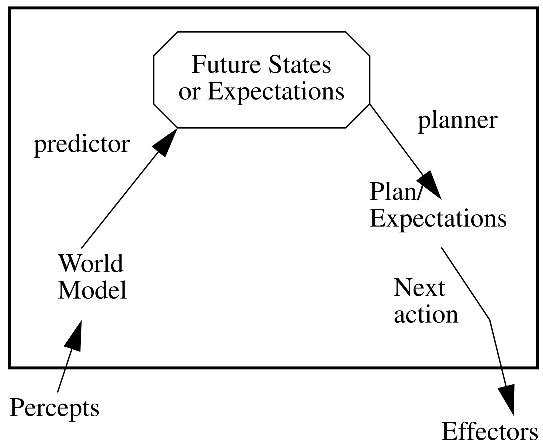

Figure 3: An Expectation Utilizing Agent.

Expectation Utilizing Agents can be thought of as "Plan Using Agents" if the reader accepts a more general definition of "planning" than that usually implied in AI, such as what [Agre & Chapman, 1987a] and [Agre & Chapman, 1987b] call "small-p planning". Expectation planning does not require that agents formulate new plans. It is possible that the agent retrieves plans from a precompiled library, as exemplified by PRS [Georgeff & Ingrand, 1989]. Expectation usage can be more general than a strict procedure following description implies. For example, the agent might select a schedule of behaviors rather than a detailed procedure. It sets up the initial behavior (e.g. "head for toolshed") and an expectation condition that, when met, signals a switch to a second set of reactive rules (cf. ERE *strategies* [Drummond, 1991]). Such an agent performs reactively while following a scheduled sequence of expectations.

In the toolshed example, the Expectation Utilizing Agent, retrieves a plan to: a. go to the toolshed, b. pick up a small phillips screwdriver, and then c. return to the broken object. From the point of view of compensating for perceptual gaps, there is no need for this agent to reason about why it is picking up a phillips screwdriver while in the toolshed. In fact, this reasoning would probably not be possible for the pure agent since it would not maintain any internal model of the current world state. This agent will have advantages over the other two when its expectations are reliable and the representation for the current world state is extremely complex.

3.3 Contingency Anticipating Agent

The third form of internal state cannot be directly interpreted as representing any part of the past, present, or future world state. This form of internal state, termed *contingency state*, changes only as a result of current percepts, and has the effect of triggering particular decisions in the future when specific conditions arise. An agent that uses only this form of internal state is called a *Contingency Anticipating Agent*. This agent differs from the Expectation Utilizing Agent in that it does not rely on expectations about the future. A plan-using agent considers only a single execution course into the future, while the anticipatory agent considers (implicitly or explicitly) all possible future choice points.

When an agent needs to select an action (encounters a choice point), the appropriate action may be a function of perceptual information observed only in the past. An agent can make an action choice either when it is faced with performing the action; or earlier when it is observing the percepts. The pure contingency anticipating agent always makes its decisions earlier when the relevant percepts are available. Thus, contingency state should be viewed as a store for all potential future decisions. At the moment an important piece of data is observed, the agent makes all potential future decisions that depend upon that observation. Firby's posting of a RAP (reactive action package) [Firby, 1987] is, perhaps, one of the best examples of setting contingency state to be implemented so far.

In the toolshed example, when the broken object is discovered, the contingency anticipating agent records that if it is ever in a situation with a small phillips screwdriver in view, then it should pick up that screwdriver. It also may record other decisions such as if the agent ever sees a toolshed, it should enter it. This agent has an advantage over the other two when the number of possible future choice points is small, the world state is highly complex, and expectations about the world are not reliable. In the toolshed example, this corresponds to the case where the relevant portions of the object are intricate and numerous, there are a small number of possible tools in the toolshed, and (as with the Model Reactive agent) the availability of any particular set of tools is unpredictable.

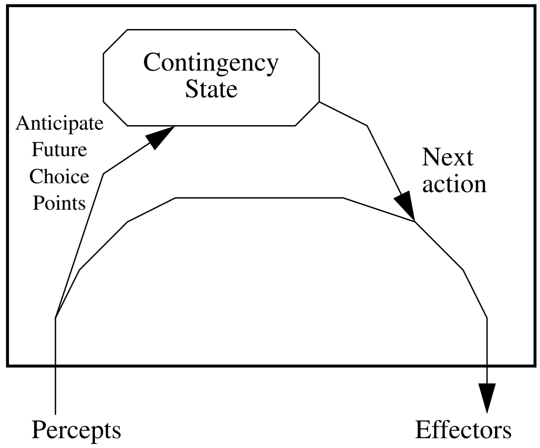

Figure 4: A Contingency Anticipating Agent.

4 Environmental Dimensions

Our survey of the three different types of internal state is now complete. We next devise environmental and task dimensions that suggest the appropriate types of internal states.

The first dimension is the complexity of the world state. When the enumeration of all pertinent aspects of the world state is long, we can view the complexity of the world state as being high. More generally, this dimension can be thought of as something similar to Kolmogorov complexity, i.e. the length of the shortest *program* (in a given formalism) capable of enumerating a given world state description[1].

The second dimension is the predictability of the environment. When an agent's expectations about the future are reliable, then predictability is high. Note that this makes predictability a function of both the environment and agent's ability, but this does not seem to be avoidable. Nonetheless, environment dynamics have no impact on predictability unless world state trajectory is unpredictable. We consider only environments that have some, possibly very small, level of predictability because internal state cannot be used to fill perceptual voids in completely unpredictable environments.

The third and final dimension measures the density of critical choice points in the space of all possible future execution paths. We are only concerned here with "critical" choice points, that is those choices where the actual decision impacts the effectiveness of the agent. When an agent is in a "decision-rich" domain, then the choice point density is high. When there are very few potential decisions, or the choices one makes is not significant. then the density is low.

'The reader should note that all three of the above dimensions are continuous and mutually orthogonal in that any point in this three dimensional space can in theory, be represented by some task domain. We assert that these three dimensions determine the justification for the use of each of the three previous forms of internal state. We now turn to identifying this dependence.

5 Trade-offs in Choice of Internal State Type

Figure 5 shows a 3-d view of the environment dimensions from the previous section. Each point in the space corresponds to environments with differing "degrees" of each dimension. The three arrows, one for each of the three

1. This characterization was suggested by Jeff Jackson.

pure agents discussed above, point in the direction in which that pure agent excels relative to the other agents.

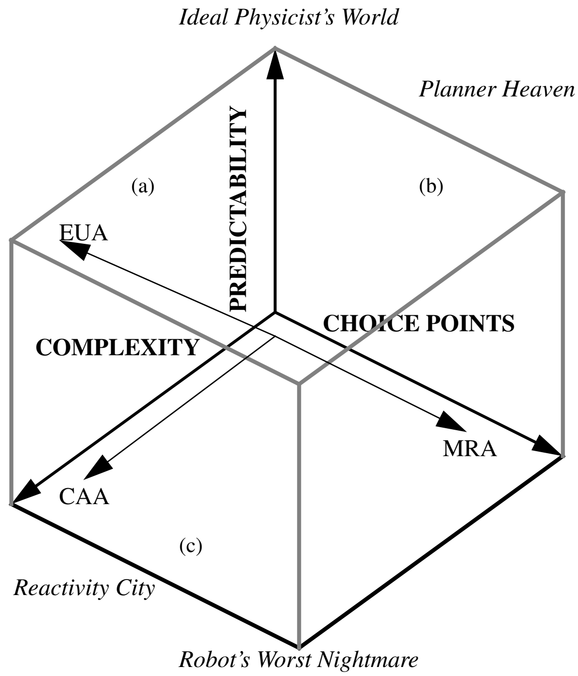

Figure 5: Environmental Dimensions.

The Model Reactive Agent is adversely sensitive to world state complexity since, as the world becomes more complex, it must maintain lengthier descriptions. Its arrow, therefore, rests low along the dimension of world state complexity. The arrow points towards increased choice point density since, unlike the other two pure agents, the Model Reactive Agent is indifferent to this complexity. Finally, since it does not take advantage of the predictability of the future, the arrow points away from that dimension. Note that the agent does rely on some predictability in order to update the non-visible portions of the world; however, as was discussed in the previous section, all agents that use internal state must rely on some minimal level of predictability.

The expectation utilizing (or plan using) agent does take advantage of the predictability of the future and is not adversely affected by world state complexity[1]; therefore, its arrow points positively in both these dimensions. As the density of choice points increases, this approach becomes less feasible due to the large number of possible futures that must be considered when retrieving or generating plans or expectations. Thus, the arrow is low along this dimension.

The Contingency Anticipating Agent is particularly sensitive to the density of choice points because it must. in essence, make all potential decisions at the time relevant observations are available. Since it does not maintain a model of the current world state, it is indifferent. to world state complexity. Also, it does not rely on expectations about the world and does not take advantage of predictability and, therefore, does not excel along that dimension.

The position of the three pure agents in the environment characterization space gives useful design information that relates the type of internal state to the kinds of environments for which it is most appropriate. Given an non-hybrid agent, the arrows point to the environments where that agent is likely to be most effective. Conversely, given an environment similar to one of the three regions containing pure agents, the arrow suggests the most efficacious form of internal state for an agent in that environment. The diagram also yields useful design information for hybrid agents and for hybrid environments.

We surmise that the closer an environment's characterization is to a pure agent's arrow, the more an agent in that

1. It is likely that the complexity of making a decision increases with the complexity of the world state. But this seems to impact all three architectures similarly so we do not account for this factor here.

environment should emulate that pure agent's method of handling incomplete perception. Consider, for example, an environment characterized by the point (a) in Figure 5. The world has a medium level of complexity, predictions are reliable, and choice point density is low. The ideal agent for this environment takes advantage of predictability, but also uses a partial model of the current world state to avoid over committing to expectations when they are no longer appropriate. However, given the relative magnitudes of world state complexity and predictability at (a), the ideal agent would rely more heavily on using expectations than upon modeling the current world state. In this region, agents that dynamically construct plans (as opposed to simply retrieving them from plan libraries) start to become justified.

As a second example, the effective agent for environment (b) uses both a current world state model and expectations or plans, but relies more heavily upon its world state model. The ideal agent in environment (c) uses primarily contingency anticipation state, but also to a lesser extent, makes use of expectations. And so on...

A few other interesting observations can be made about the diagram. First, as we move from the top right to the bottom left of the diagram, the ideal agent becomes increasingly reactive. Thus the diagram indicates that reactive architectures are more appropriate in environments with low predictability, high environmental complexity, and moderate to high density of decision points. Similarly, the diagram suggests that non-reactive architectures are more appropriate in environments where the future is relatively predictable, where the world can be adequately modeled internally, and where the density of decision points is not too high. Another observation is that the top of the diagram represents the "Ideal Physicist's World". Here the world is highly amenable to modeling and does not suffer from either source of complexity. Directly opposite the Ideal Physicist's World is the "Robot's Worst Nightmare" where the future is uncertain, the world is complex, and there are many decision points. It is interesting to consider how an agent in this environment might attempt to cope. Reasonably enough, the diagram suggests that the agent use a hybrid approach, i.e. that the agent must do everything possible to try to remain effective. This agent would probably model and track portions of the world that were not too complex and attempt to anticipate (but not plan for) important contingencies that critically depend upon the complex parts of the world.

Figure 5 serves as a heuristic design guide. It can be used to justify architectural design decisions concerning the use of particular forms of internal state. It can also be used to classify and predict the most appropriate environment for a given agent. The diagram cannot be expected to be perfect, however. It is proposed simply as a heuristic guideline, but even so, its inaccuracies seem far overshadowed by the insights it provides. A particular deficiency is that the diagram does not characterize sensitivity to increased perceptual incompleteness.

Guidelines like Figure 5 for intelligent agent design currently do not exist, but as we discussed in the introduction, the need for them is high. We hope that Figure 5 will provide a prototype for the development of further comparative scales, as well as improvement of the current characterization.

6 Conclusion

With the current diversity of architectures for intelligent agents, identifying similarities, differences, and relative advantages between them is far from trivial. Good performance is usually reported for one or two limited domains, but critics of an architecture are quick to find examples ill-suited to it. Clearly the appropriateness of a particular architecture or architectural feature is a function of the characteristics of the environment. Identifying the relevant environmental dimensions and their relationship to specific properties of intelligent architectures would aid comparison and evaluation of research results. It may also help create clearer design methodologies for the designers of intelligent agents.

This paper attempts to make progress towards this goal. In it we consider the specific case of incomplete perception and conjecture that an agent *must* use internal state to compensate for perceptual voids. We identify three forms of internal state and propose three environmental dimensions that serve to characterize the trade-offs between each form. These trade-offs are summarized in Figure 5. While Figure 5 is not perfect, we believe it contributes considerable insight into the pertinent trade-offs. As far as we know, this is the first attempt to concisely explicate a relationship between agent architecture and environment. We hope that additional analysis of the relationship between other architectural features and environment dimensions will serve to improve the coherence of the field as a whole.

7 Acknowledgments

The ideas in this paper were developed as part of the spring 1991 Intelligent Agent. Architecture course at CMU. We thank all the participants in the IA course for the lively discussions that shaped the ideas contained here. The course was organized by Reid Simmons, Tom Mitchell, and Chuck Thorpe. Additional thanks go to Reid Simmons and Keith Swanson for comments on an initial draft of this paper. Lonnie Chrisman is funded by NASA contract NAGW-1175. Rich Caruana. is funded by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), US Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-80-C-1465, Arpa Order No. 7597. Wayne Carriker's funding comes from an Engineering Creativity Award from the National Science Foundation, grant number ENG-8811580. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

References

[Agre & Chapman, 1987a] Phil Agre and David Chapman, "Pengi: An Implementation of a Theory of Activity." In Sixth National Conference on Artificial Intelligence; 1987.

[Agre & Chapman, 1987b] Phil Agre and David Chapman, "From Reaction to Participation." In DARPA Planning Workshop: Panel on Interleaving Planning and Execution; Santa Cruz, CA October,1987.

[Brooks, 1986] Rodney Brooks, "A Robust Layered Control System For a Mobile Robot." IEEE Journal of Robotics and Automation, Vol. RA-2, No. 1; March 1986.

[Brooks, 1990] Rodney Brooks, "Elephants Don't Play Chess." Robotics and Autonomous Systems, Vol. 6, No. 1-2, pg. 3-15; 1990.

[Brooks, 1991] Rodney Brooks, "Intelligence Without Reason." International Joint Conference on Artificial Intelligence; August 1991.

[Chrisman & Simmons, 1991] Lonnie Chrisman and Reid Simmons, "Sensible Planning: Focusing Perceptual Attention." In Ninth National Conference on Artificial Intelligence; 1991.

[Dean, 1990] Thomas Dean, "Planning Under Uncertainty and Time Pressure." In Proceedings of the Workshop On Innovative Approaches to Planning, Scheduling, and Control, pg. 390-395; Nov, 1990.

[Drummond, 1991] Mark Drummond, John Bresina, and Smadar Kedar, "The Entropy Reduction Engine: Integrating Planning, Scheduling, and Control." AAAI Spring Symposium Series; 1991.

[Firby, 1987] R. James Firby, "An Investigation into Reactive Planning in Complex Domains." In Sixth National Conference on Artificial Intelligence; 1987.

[Georgeff & Ingrand, 1989] Michael P. Georgeff and Francois Felix Ingrand, "Decision-Making in an Embedded Reasoning System." In Eleventh International Joint Conference on Artificial Intelligence; Detroit, MI; 1989.

[Kaelbling & Rosenschein, 1990] Leslie Pack Kaelbling and Stanley J. Rosenschein, "Action and Planning in Embedded Agents." Robotics and Autonomous Systems, Vol. 6, No. 1-2, pg. 35- 48; 1990.