

## Preliminary Studies in Agent Design in Simulated Environments

Scott B. Hunter

Department of Computer Science  
Cornell University  
Ithaca, NY 14853  
E-mail: hunter@cs.cornell.edu

It is known that, in general, the point along the purely-reactive/classical-planning axis of the controller spectrum that is most appropriate for a particular environment/task (E/T) pair will be determined by characteristics of the environment, the agent's perceptual and effectual capabilities, and the task. Instead of proposing another hybrid architecture, we want to determine criteria for determining which architectural compromise is best suited for a given E/T. Our goal is to understand relationships between E/T pairs and the agent architecture, so that we can predict the performance of the architecture under parametric variations of the environment and/or the architecture. This is a first step toward constructing methods for automatic synthesis of agents as in (Ros89).

Our example of a domain where the choice of architectural basis is not so clear is the game of XChomp (programmed by Jerry J. Shekhel), a close relative of the commercial game PacMan. This domain allows for easy change of parameters to simulate a number of discrete combinatorial problem domains.

Interesting characteristics of the game that make it different from the E/Ts considered in (AC87), (Bro86), (Cha91), (Sch87) among others, are:

*There are non-local tasks.* By non-local, we refer to not only spatial and temporal extents, but also universal quantification of parameters of the task.

*Hostile aspects of the environment may be temporarily made not only benign, but positive concrete goals; these conversions are under the control of the player.* Classification of objects changes over time, complicating the decision of how to respond to such objects, as these are based on projections of possible futures.

*There is not much flexibility with respect to movement.* Not only is movement within this environment restricted to the four cardinal directions, but it is a maze, so that in most locations, only two of those four may be used. When the cost of making mistakes is high, the extra effort to get it right the first time is (possibly) justified.

*There are multiple conflicting objectives.* While having multiple objectives is not particularly novel, those in this environment have a nasty habit of pulling their acquisitions at cross purposes.

Any designer must answer the following questions:

On what informational basis does an agent make its action selection choice? What aspects of the world does it perform forward projection on, what aspects does it sense, and what is the map from external and internal state of the agent to an action (or sequence) that maximizes the objective function of the agent? We need to be able to construct a solution and justify it using methods other than pointing to the constructed solution as an existence proof. This follows in the spirit of work done in (Hor93) for a mobile robot.

To enable us to study these questions, we isolate a range of E/T combinations based on the XChomp game. We implement a range of controllers that exploit the information needed for "optimal" play and test their task performance experimentally. We then vary the performance requirements and environmental specifications in a form of perturbation analysis to determine how robust the agents are; and how to modify them to be effective in new situations.

In addition, starting from a very simplified version of this E/T, we are developing a theoretical basis upon which to justify the agents we develop. This basis is expected to not only be used in determining how an agent should behave, but also what a designer should not be concerned about.

### References

- P. Agre and D. Chapman. Pengi: an implementation of a theory of activity. In *Proceedings of AAAI-87*. Morgan Kaufmann, 1987.
- R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1), 1986.
- D. Chapman. *Vision, Instruction and Action*. PhD thesis, MIT AI Lab, 1991.
- I. Horswill. Poly: A vision-based artificial agent. In *Proceedings of AAAI-93*. Morgan Kaufmann/MIT Press, 1993.
- S.J. Rosenschein. Synthesizing information-tracking automata from environment descriptions. In *Proceedings of KR-89*. Morgan Kaufmann, 1989.
- M.J. Schoppers. Universal plans for reactive robots in unpredictable domains. In *Proceedings of IJCAI-87*. Morgan Kaufmann, 1987.