

A Simulated World for Modeling Learning and Development

Pat Langley, David Nicholas, David Klahr, and Greg Hood¹
Carnegie-Mellon University

1. Introduction

The creation of a complete intelligent system is a goal of many researchers in cognitive science and artificial intelligence. An elegant path to realizing that goal would be to let a system *develop* its intellect through interactions with the environment. Recent work in perceptual and motor skills has brought these hopes closer to fulfillment, but we are still many years from constructing systems that can adequately interact with the real world.

In this paper we describe a *simulated* world for developing and testing models of learning and development. Such a facility could play an important role in the emerging field of cognitive science, since it would encourage researchers to construct complete systems, and since it would provide a common ground on which competing theories might be compared. Below we summarize the aims of the project and the design criteria for the simulated environment. After this, we present an overview of the simulated world and of the sensory/effector interfaces through which model organisms may interact with it. We close with a brief discussion of the interface between the user and the environment.

2. Aims of the Research

The simulation system described here was conceived by, and for, cognitive scientists representing a variety of theoretical biases and paradigms. Its purpose is not to embody particular assumptions about complex information processing (human or otherwise), but rather to provide a medium for empirical research that can be shared by -- and tailored to the needs and goals of -- as wide a range of investigators as possible.

2.1. Focus on Learning and Development

The major goal of cognitive science is to understand the nature of intelligence. However, the knowledge and strategies responsible for intelligent behavior vary with time. Since any science searches for *invariant* regularities of behavior, this presents a difficult problem for our field. Langley & Simon (1981) have argued that we may find the desired invariants in a theory of learning and development, and preliminary steps towards a theory of the transition process have been taken by Klahr & Wallace (1976), among others.

Theories of learning and development can be characterized as having two problems to solve. First, they need to account for behavior at several different points in a developmental progression. Second, they must account for the mechanisms that enable the system to progress from state to state. The quality of the transitional theory is constrained by the quality of the theory that accounts for behavior at each state. One natural consequence of this argument is a focus on "expert" or "mature" performance: such models provide an ultimate target for the transitional processes, and would thus appear to have top priority in theory building efforts.

7312

However, we believe that this emphasis has resulted in an unfortunate limitation. The contact between current computer simulation models and an empirical base is almost exclusively via adult performance measures, and often skilled adults, at that. To the best of our knowledge, none of the currently proposed mechanisms for taking a system from state N to state $N+1$ could plausibly have brought the system to state N in the first place. That is, most models have little *developmental tractability*, although they may provide a reasonable account of the learning mechanism in an already well developed system.

¹The ideas presented here evolved out of the interactive efforts of the CMU World Modelers Group, in which Jaime Carbonell has played a major role and which has included Hans Berliner, Greg Harris, Marty Herman, and Glenn Iba, in addition to the present authors.

In order to remedy this situation, we need to ask questions about infant systems, about the rudimentary encodings and internal representations, and about the innate kernel of self-modifying processes. We think it is important to attempt to formulate developmental theories in which the ability to undergo self-modification in a plausibly supportive environment is the *primary* constraint on system design, with performance a problem to be solved within that constraint. (The converse is almost universally the case with modern theories of self-modification: performance models come first, with self-modification added on in ingenious ways.) The simulated world proposed here would facilitate such an effort. We envision our simulation system as a tool for investigating any of several domains of knowledge development, including form perception, object constancy, problem solving, or quantitative processes.

2.2. Constructing Complete Systems

A standard approach in science is to partition a phenomenon and study the pieces separately. In cognitive science, we find researchers specializing in language processing, problem solving, perception, and many other areas. And though this division of labor has clarified much about the components of intelligence, it has revealed little about their interaction. For example, researchers in language acquisition have separated that form of learning from concept formation and word learning, despite their strong interplay. Similarly, problem solving theorists have ignored perceptual and motor factors, though they may significantly influence the difficulty of a problem. We feel enough is understood of the components to initiate attempts to construct complete intelligent systems.

A second argument for creating complete systems comes from representational considerations. As long as one focuses on only an artificially bounded *subset* of behavior, the input of the system must be specified by the user. Thus, one might build a system with apparently general learning mechanisms, but which would learn only when presented with carefully hand-crafted data. The construction of a complete system should guard against such subtle kludges, since information would be obtained through direct interaction with the environment or by inferences the system made itself.

2.3. Simulated Worlds vs. Real Worlds

The ideal complete intelligence would be a robot, with sensory abilities for perceiving the real world and motor abilities for affecting it. Unfortunately, we are still far from understanding perceptual and motor behavior in the necessary detail. This leads us to propose the less impressive but more manageable option of devising a simulated environment. The notion of a simulated world has its own attractions, including its relative independence of hardware and its transportability between sites, making it an ideal tool for cognitive scientists to employ in their model construction. It also eliminates the computational constraint that cognitive processing be done in real time.

3. Criteria for a Simulated Environment

In order for the simulated world to be useful as an experimental tool in cognitive science studies of learning and development, the following points were considered central design criteria:

- **Independence and richness** -- The world model must be truly separate from the organisms and their internal "models" of the world. Moreover, the world must be sufficiently rich and unpredictable, so that no organism can internalize a complete model of the world in its lifetime. In particular, the environment should be much richer than either Becker's (1970) grid universe or Winograd's (1972) blocks world. This is crucial, since the real world elicits qualitatively different behavior than would a completely-internalizable world (where table-lookup and formula-evaluation would suffice for perfect behavior).
- **Extensibility and consistency** -- The world must be extensible, in terms of introducing arbitrary numbers of new objects and new organisms. However, physical laws must remain invariant.

- **Cross-sensory correlation** -- The world model should support multiple sensory media, and each organism can, in principle, be designed to take input from any subset of the sensory media. Cross-medium sensory correlations play a central role in theories of perceptual and cognitive development, and yet few computational studies of this phenomenon have appeared to date.
- **Multi-level sensory interfaces** -- In order to satisfy the objectives of different researchers exploiting the same world-model tool, interaction between the world model and the organisms should be mediated via sensory and effector interfaces, whose purpose is to provide the organism with data (or effector functions) specified at the desired level of abstraction.
- **Synchronous processing** -- The environment should not be forced to stop in order for the organisms to think at their leisure. Therefore, if the organisms are unable to process all incoming input, the focus-of-attention problem is introduced as an integral aspect of cognition.
- **Communication among organisms** -- The only communication possible is via the sensory and effector interfaces (by gestures, language, etc.), requiring organisms to pay attention in order for the communication to take place.

In formulating the above criteria we considered various issues we wished to research using the simulated-world environment. For instance, learning purposive action is a basic function in most higher-level organisms, including human "rational" thought. We wished to provide organisms with basic drives (e.g., hunger, curiosity, companionship of like organisms, etc.), basic actions, and learning mechanisms. Exactly what the starting point for learning processes ought to be is a matter of research and/or discretion according to the phenomenon one wishes to investigate.² For example, one type of purposive action is to learn subsumption goals (Wilemsky, 1978); e.g., secure more of a resource (such as food) than presently required to satisfy an internal drive -- if past experience has shown that the drive will recur and finding the resource may be an uncertain or costly operation. Another is to posit intrinsic satisfaction from the simple, repeated execution of activities: e.g., Piaget's "circular reactions."

4. An Overview of the Simulated Environment

The simulated three-dimensional environment contains *objects* of two types: *primitive* (the building blocks of the physical domain) and *complex* (hierarchical structures, aggregates of primitive objects). Every physical structure, including the manifestations of organisms in the simulated world, is either a primitive or a complex object.




Figure 1. Primitive Objects

4.1. Primitive Objects

Figure 1 illustrates the four types of primitive objects (polygon, circle, cylinder, sphere) and the spatial parameters that must be defined for each. A polygon is specified by an ordered list of

²One must be careful in researching learning by simulating autonomous organisms not to fall into the self-organizing-system fallacy. Early AI research in learning postulated extremely simple learning mechanisms and a virtual *tabula rasa* with respect to real world knowledge, with the expectation that such a system could organize itself into a thinking entity. The only parallel of such a "magical" feat in the real world is the evolutionary process, but this process required billions of years, billions of generations of organisms, and millions of individual organisms per generation. Our world model is designed with the objective of modeling developmental learning -- where a single organism can learn purposive action in a fraction of its lifetime. As such, it requires a non-zero starting point with some innate abilities, a high communication bandwidth with the external world, and built-in drives to focus its attention and guide its behavior (at least initially).

vertices, a circle by the coordinates of its center, a radius, and a normal vector, and so on. All primitive objects have an additional set of physical properties that must be specified: mass, center of mass, velocity, angular velocity, elasticity, static and dynamic coefficients of friction, temperature, taste, color, and texture.




Figure 2. Joint types used in complex objects

4.2. Complex Objects

The joining of primitive objects to one another, in any of the several ways shown in figure 2 (joints may be fixed, sliding, pivoting, or twisting), produces a complex object -- an organized structural hierarchy within which properties may be shared and which may be acted upon by the physical laws of the environment as a single unit. A joint may be broken, and the complex object separated into two (reorganized) hierarchies, by the application of force in excess of that joint's prespecified *stress limit*. A fifth, or "virtual," joint type may also be employed, when two objects are touching and not moving relative to one another (i.e., in static contact).

Consider the example of a table with a coffee cup resting on it. A typical table has four legs and a top: four primitive objects (long, thin cylinders) connected by fixed joints to a fifth (a circular or polygonal plane). The cup may be broken down into its body and a handle: the former a hollow cylinder closed at one end by a circle of equal radius, and the latter three more cylinders connected to resemble three sides of a square. A virtual joint connects the two complex objects, so that if a gentle force is applied to the table it will take the cup with it when it moves (depending, of course, on the exact magnitude and direction of the force, the frictional coefficients of the table top and the bottom of the cup, the inertia of the cup, etc.).

4.3. Physical Laws

The simulated environment must be "updated" at the start of each quantum of time. The positions of objects need to be adjusted, based on previous velocities and the application of forces; stresses on joints must be computed and, if necessary, objects fragmented. In short, changes that have resulted from the influences of objects (including the model organisms) upon one another -- or from some alteration introduced from without, by the user -- must be incorporated into the description of the current "state-of-the-world." This is accomplished by applying (recursively, to complex objects and their components) a set of simplified physical laws: rules for the modeling of relatively gross interactions in the environment.

5. Sensory and Effector Interfaces

Organisms must interact with the simulated environment in a manageable way that can be tailored to the needs of the researcher. As figure 3 illustrates, the sensory and effector interfaces (also called sensory-motor interfaces, or S-MI's) are independent of both the organisms and the external world model. The primary reason for this decision is that different researchers should be allowed to define the level of abstraction of the sensory information detected by each class of organism. Hence, a researcher interested in language acquisition can "plug in" a symbolic visual interface that provides the names and locations of objects in the visual field of the organism (in order to relate objects and actions to words and sentences). However, a researcher interested in perceptual or motor learning can substitute a lower-level interface that provides only collections of features for visible objects -- or yet lower-level visual data "decompiled" from the world model by the sensory interface. The same design principle holds for effector interfaces.

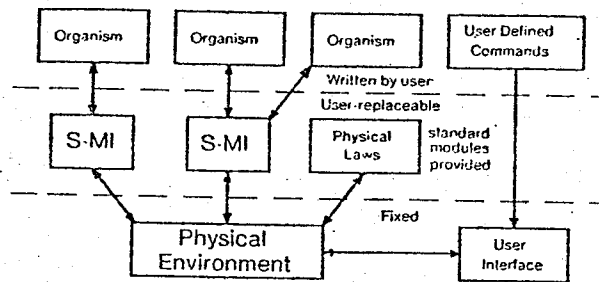


Figure 3. Modules of the System

We believe that the sensory and effector interfaces should themselves be subjects of research. As an example of the problems that arise in designing a sensory interface, consider the fact that human beings can tell with precision the location of nearby objects, but can only place approximate values on the absolute location of more distant objects. This phenomenon seems crucial for some forms of generalization. For example, it does not matter if a predator is 1001 or 1003 feet away -- the organism should flee; however, in reaching for an object that is 1 or 3 feet away, different actions are called for. Therefore, the perceptual mechanism that accentuates the latter difference, while glossing over the former, appears to be a desirable feature.

6. The User Interface

To exploit the advantages of a simulated world it is essential to have a powerful user interface. Of major importance is the ability to view the world as an observer within the world would see it. In our implementation, the user may create multiple windows, each of which is a perspective drawing of the world from a different viewpoint. These viewpoints may be fixed, or they may be bound to the "eyes" of particular organisms within the world so that the window accurately portrays the image seen by the organism. A graphical display also allows the user to easily construct new objects and organisms differing in physical characteristics.

In addition to having commands for controlling the windows, the user must also be able to control the motor and verbal behavior of certain organisms that function as teachers (of other organisms). The interface should allow the user to specify actions as high level commands instead of cumbersome primitives (such as forces to be applied at certain joints). Simulation of the world also enables a detailed record of events to be kept for later examination by the investigator. Assuming the cognitive states of the organisms are saved periodically, such a record could be backed up to a certain point in time and the simulation restarted, providing a valuable opportunity to explore alternative courses of action.

7. References

- Becker, Joseph D. An information-processing model of intermediate-level cognition. Stanford University Department of Computer Science AIM-119, May, 1970.
- Klahr, David & Wallace, J.G. Cognitive Development: An Information-Processing View. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Langley, Pat & Simon, Herbert A. The central role of learning in cognition. In J. R. Anderson (ed.), Cognitive Skills and Their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.
- Wilensky, Robert. Understanding goal-based stories. Yale University Department of Computer Science Research Report 140, September, 1978.
- Winograd, Terry. Understanding Natural Language. New York: Academic Press, 1972.