SPATIAL LEARNING SIMULATION SYSTEMS*

Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson
Computer and Information Science Department
Graduate Research Center
University of Massachusetts
Amherst, MA 01003

We describe a highly interactive color graphics simulation system designed for studying spatial movement strategies and models of spatial learning. The system provides multiple representations of model activity through animated color displays of the structures simulated.

Introduction

The movement strategies of animals form a fascinating area of study, not only for what they tell us about animal behavior and its evolution but also for what they suggest to us about methods of solving problems in general. For example, the common bacterium Escherichia Coli responds to levels of attractants (e.g. molecules indicating the presence of food) by tending to swim in a straight line as long it senses increasing attractant concentration and by tending to change its direction (to a random new one) when the attractant level decreases [see Berg (4) or Koshland (7) for discussions of bacterial chemotaxis]. Selfridge (8), an expert in the theory of adaptive systems, argues that this strategy is one of the most important primitive adaptive strategies for introducing adaptiveness into a system, be it a natural or a man-made system. He calls it, using Berg's (4) terms, the "Run and Twiddle" strategy: If things are getting better, keep doing what you are doing; if things are getting worse, do something (anything!) else. Similarly, more sophisticated adaptive strategies, including strategies for learning, can be clearly understood when played out as controlling movement through space.

We have developed a highly interactive, color graphics based simulation system called EXPERIMENTER for creating, modifying, manipulating, and simulating model organisms and their two dimensional planar environments. The model world consists of up to a large number of independently moving model organisms, goal or landmark objects, and movement restricting barriers. A user provides the subroutines for the control of the model organisms, and EXPERIMENTER allows him to experiment with the capabilities of his control algorithm in a range of different environments, much as a psychologist would investigate the behavior of some experimental animal. The limitations on model organisms and environmental structures are particularly suited to the investigation of elementary forms of spatial learning, reasoning, and planning.

Our particular interest in this system has been with the investigation of a variety of learning algorithms that are in the form of rules for storing information in structures known as associative memory networks. These networks consist of a collection of simple processing units that behave as idealized neurons, altering their connectivity weights according to their input/output histories. Although it has been unfashionable since the late 1960's to develop systems based on networks of neuron-like adaptive elements, it is now being realized that this approach offers unique advantages for highly associative processes, especially for those that must be performed in real time such as in vision and robot control [see Hinton and Anderson (6), Albus (1)].

To provide animated graphical displays of the movement control networks used with EXPERIMENTER, we created a program called DESIGNNET for interactively creating color graphics displays of associative memory networks. When used with the EXPERIMENTER system, the network displays are animated to show the values of the network variables as they change over time. We are thus able to observe the internal behavior of the movement control network as well as its external spatial behavior. See Barto and Sutton (2), and Barto, Anderson, and Sutton (3) for descriptions of some of the results obtained using these programs.

Color Graphics as a Simulation Tool

As high quality color graphics systems become more available at modest cost, it is quickly being realized how useful color graphics can be for presenting simulation data. Whereas the most common application of color graphics in simulation is probably the production of color coded multidimensional graphs [see Graedel and McGill (5)], the use of color graphics is not limited to static presentation of results. In the case of EXPERIMENTER, simulations produce animated displays of spatial movement similar to what one might see, for example, by observing bacteria in a petri dish with a microscope. Color is used to distinguish various kinds of landmarks and to distinguish the trails left behind the various moving model organisms. The animated display of the controlling associative memory network shows it as a dynamic mechanism, with parts of its structure changing over time in color, shape, or size.

Our simulation system runs on a Digital Equipment Corporation VAX 11/780, coupled with a Grinnell GMR-27 color graphics display system and an Aydin monitor. The display has a resolution of 512 by 512 pixels with 16 bits per pixel. The primary method of interaction with the simulation system is through a joystick or trackball controlling a cursor, and some associated buttons.

Space as a Dynamical System

To a mobile organism, space can be viewed as a particular kind of dynamical system with which the organism interacts. In the simplest case of a spatial environment containing no moving objects, the state of this dynamical system is the organism's position and orientation. The organism's actions are inputs to its spatial environment which "computes" a new position and orientation from these inputs together with the old position and orientation. The spatial environment's outputs are stimuli generated by various landmarks. These outputs are a function of the environment's state.

We are thus able to view the organism environment interaction as a special case of the classical control loop where the "plant" to be controlled in this case is the spatial environment. EXPERIMENTER allows us to interactively specify the dynamics of this special kind of plant and to interactively conduct simulation experiments. EXPERIMENTER is only appropriate, however, where a two dimensional view of space is sufficient, such as for spatial planning and reasoning tasks. It is not appropriate for applications requiring a three-dimensional model environment or objects of complex structure.

---

*This research was supported by the Air Force Office of Scientific Research and the Avionics Laboratory (Air Force Wright Aeronautical Laboratories) through contract F33615-80-C-1088.

Editing the World

Most interaction with the program occurs through the joystick controlled cursor of the color display. Various menus of options appear on the CRT, and an option can be selected by moving the cursor into a box in the option list and pressing a button on the control console. Fig. 1 shows most of the hierarchically organized menu structure. We will describe the action of a few of these options to illustrate some of the interactive facilities.

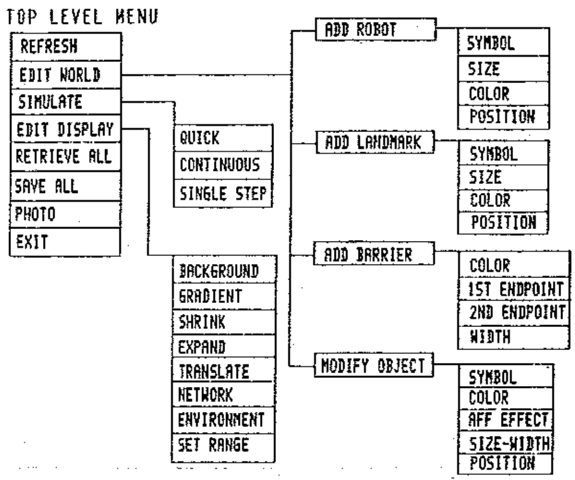

Figure 1. A portion of EXPERIMENTER's menu structure.

The EDIT WORLD option allows the user to change some aspect of the environment. When it is selected, another list of options appears as shown in Fig. 1. Selecting, for example, the ADD ROBOT option (where by "robot" we mean a model organism) from this new list causes another list of options to appear which ask for SYMBOL TYPE, SIZE, COLOR, and POSITION. Selecting SYMBOL TYPE causes several different symbols, e.g. square, circle, or box, to appear on the screen. Placing the cursor on the symbol desired to represent the new model organism and pressing a button causes that symbol to appear in the middle of the screen (and the other symbols to disappear). Selecting SIZE allows the user to shrink or enlarge the symbol by moving the cursor toward or away from the center of the screen. When COLOR is chosen, the two options ARBITRARY and BY MATCHING appear. Choosing ARBITRARY allows the user to use the cursor to select intensities for the primary colors red, green, and blue to create the desired color. Choosing BY MATCHING allows the user to give the new model organism the color of another object in the environment by placing the cursor near the object having the desired color and pressing a button. Finally, the POSITION option permits the user to position the model organism to any location in the environment by moving the cursor. The procedure for adding LANDMARKS and BARRIERS are essentially the same, and MODIFY OBJECT lets one alter the size, color, position, etc., of existing model organisms, landmarks, or barriers. The AFF EFFECT option is used to specify how the environment's state affects the input channels (or "afferents") of the movement controller.

Multiple Representations

The EDIT DISPLAY option allows the user to alter the view of the simulated system that is presented on the CRT. When option ENVIRONMENT is chosen, a bird's-eye-view of the spatial environment is presented. Option NETWORK causes a display of the controlling network to appear (we discuss network display below). The options SHRINK, EXPAND, and TRANSLATE allow the user to zoom and pan over the environment display. By being able to switch from one representation to another at any point in a simulation, and to resume simulation while viewing the new display, a user can quickly understand what the model is doing. One can also save the system state using the SAVE ALL option so that at a later time simulation can be restarted at that state (using RETRIEVE ALL) and the subsequent behavior viewed again in a different graphical form.

For the spatial learning experiments that we have performed using this system, we have found several other display representations useful. As a simulation proceeds, the model organisms can leave trails showing where they have been. Each trail has the color of its originating organism so that the paths of various organisms can be distinguised (Fig. 2). The GRADIENT option superimposes a vector field over the environment display that shows the direction the simulated organism would move, given its present internal memory state, if it were to be placed at each of many positions in the environment. Fig. 3 shows an example of this type of display. It is very easy to assess the progress of a learning algorithm by using this display option to show the overall behavioral consequences of the information it has accumulated.

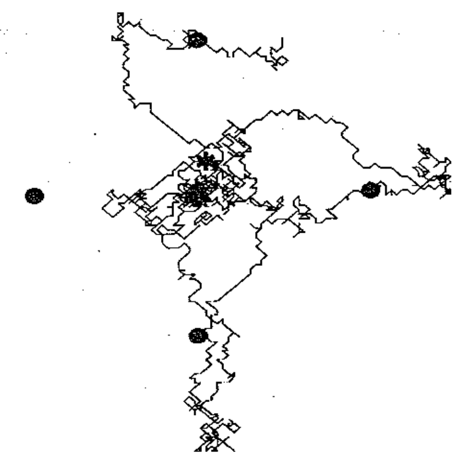

Figure 2. The trail of a model organism [from Barto and Sutton (2)].

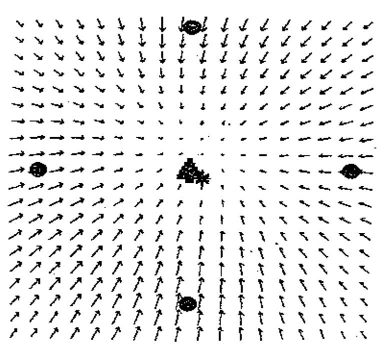

Figure 3. A vector field showing the model organism's movement tendencies [from Barto and Sutton (2)].

Network Displays

Although EXPERIMENTER can be used with any sort of control program, our experiments have been with control programs in the form of associative memory networks. We cannot completely describe associative memory networks here, nor can we give more than a brief description of the system DESIGNNET which we have constructed for designing displays of them. In essence, an associative memory network is a matrix of real numbers that is used to determine an output vector, or "recollection", from an input vector, or "key", by matrix multiplication [see Hinton and Anderson (6)]. Using DESIGNNET, one can very quickly create complex network displays by selecting and positioning modules which represent entire associative memory networks. Fig. 4 shows a display made using two modules, one whose outputs act as inputs to the second. This structure is animated during simulation and dynamically shows the values of network variables using a variety of graphical conventions. For example, the values of up to three variables can be displayed at each intersection of the input and output pathways. One variable's values is shown as the radius of a colored disk (whose color indicates the sign) as shown in Fig. 4.

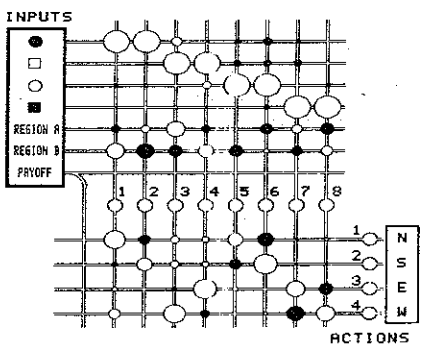

Figure 4. A display of an associative memory network [from Barto, Anderson, and Sutton (3)].

The other variables are shown as the radius of a white circle, and half the width of a white square. The values of other variables, such as input and output variables, are shown as the color intensity of their corresponding pathways. When viewed during simulation, the animated network display shows network activity as changes in color intensity, disk size, etc.

Conclusion

We have used some of the enormous potential of color graphics systems to create a graphical surrogate experimental system for investigating strategies for moving in space. Rather than producing traditional graphical displays of simulation results, such as plots of descriptive variable time trajectories, the system produces animated graphical representations of the structures simulated which show the time courses of the variables. By providing multiple representations of model behavior and extensive interactive means for manipulating the simulated spatial environment, the system makes it possible to test hypotheses very efficiently. Our experience with this system has convinced us that animated color graphics must continue to play a central role in our research program.

References

1. Albus, J.S.: Brains, behavior, and robotics. Peterborough, N.H.: McGraw-Hill 1981

2. Barto, A.G., Sutton, R.S.: Landmark learning: An illustration of associative search. Biol. Cybern. 42, 1-8 (1981)

3. Barto, A.G., Anderson, C.W., Sutton, R.S.: Synthesis of nonlinear control surfaces by a layered associative search network. Biol. Cybern. (to appear)

4. Berg, H.C.: How bacteria swim. Sci. Amer. 233, 36-44 (August 1975)

5. Graedel, T.E., McGill, R.: Graphical presentation of results from scientific computer models. Science 215, 1191-1198 (March 1982)

6. Hinton, G.E., Anderson, J.A.: Parallel models of associative memory. N.J.: Lawrence Erlbaum 1981

7. Koshland, D.E., Jr.: A model regulatory system: bacterial chemotaxis. Physiol. Rev. 59, 811-862 (1979)

8. Selfridge, O.G.: Tracking and trailing: adaptation in movement strategies. Unpublished manuscript (1978)