SESSION XI TEACHING APPLICATIONS OF MICROCOMPUTERS

N. John Castellan, *Presider*

Teaching students to model neural circuits and neural networks using an electronic spreadsheet simulator

THOMAS T. HEWETT

Drexel University, Philadelphia, Pennsylvania

An electronic spreadsheet simulator can be used to enable students to conduct simulated micro-electrode recording experiments. In addition, it can be used both to let students explore the operation of models of hypothetical neural networks and to let them design and develop their own neural models.

A recent paper (Smith et al., 1984) described the educational goals of student ownership of a personal microcomputer at Drexel University. In a later paper, Hewett and Perkey (1984) argued that many computer programs intended for instructional purposes are impractical or of little use when each student has a personal machine. However, required student access to a personal computer creates new opportunities for the use of microcomputer applications programs. One obvious, but nontrivial, example is the impact of the word processor on what is considered acceptable written work, from lab reports to term papers. When every student has a word processor, the instructor no longer need worry about creating time-consuming extra work (e.g., retyping) when demanding written work be revised until acceptable.

A less obvious example of a new use for a microcomputer application program involves the electronic spreadsheet simulator (e.g., Multiplan, SuperCalc, VisiCalc, etc.). This kind of program was originally designed to simulate and replace the large paper worksheet used in accounting and financial planning. It is, however, a more generally useful tool that can be used to study and explore functional relationships among a number of parameters. Consequently, it can be used in fields as diverse as ecology (e.g., Silvert, 1984), nutrition, physics, and psychology.

An electronic spreadsheet simulator is oriented around a matrix of rows and columns, into which the user enters text, data, formulas, or instructions. The program executes user-specified calculations either automatically, after a new value is entered anywhere in the matrix, or upon

demand. This calculation cycle often involves recalculation of new values for previously entered formulas that depend upon still other values that may have been changed during or since the last calculation cycle. The typical spreadsheet simulator also offers a number of programming language capabilities, including the ability to manipulate strings and to do iterative calculations. In addition, built-in functions can be combined to produce a limited plotting capability, thereby providing both a graphical and a numerical method for illustrating relational concepts. Although the spreadsheet program requires that the simulation designer be able to specify the relationships to be explored, it does not require that the end user be a programmer.

There are a number of areas of instruction in psychology in which this tool can be applied. For example, in dealing with sensation, perception, and pattern recognition, it is sometimes desirable for students to understand both the basic neurophysiology and the functional operation of various hypothetical neural circuits and networks. One approach to helping students build their understanding was illustrated by Bott and Munro (1977) and Lindsay and Norman (1977). This approach involves presenting basic neurophysiological data, describing the procedures by which the data have been gathered, and then abstracting fundamental concepts. These basic concepts then become the building blocks used in developing a series of increasingly more complex, and more adequate, explanatory models. A number of these theoretical models can be quickly and easily represented and simulated using an electronic spreadsheet program. By beginning with inadequate but intuitively simple models, one can gradually introduce the necessary complexities required to deal with a broad range of available phenomena.

The author's mailing address is: Department of Psychology and Sociology, Drexel University, Philadelphia, PA 19104.

SIMULATING MICROELECTRODE RECORDING

Consider the following heuristic simplification. A typical microelectrode recording experiment involves inserting a microelectrode into a nerve bundle in an attempt to pick up the neural impulses generated by a single cell. Once the investigator is recording from a single transmitter, different patterns of stimulation are introduced to the receptors. Then an attempt is made to discover meaningful relationships by comparing the patterns of stimulation with their effects on recorded cell output. A further simplification can be introduced by narrowing focus to a single, one-cell-wide slice of neural tissue, and concentrating on two sets of cells, a set of receptors, and a set of transmitters.

Now, assume that the output of a single transmitter is being recorded. With no receptor stimulation, this cell shows a steady background firing rate. By temporarily stimulating each individual receptor, observed changes in the firing pattern of the transmitter cell can be attributed to the effect of the stimulated receptor. A similar result can be obtained by moving a square wave of light from one end to the other over the set of receptors. Either procedure allows identification of the receptor cells that are part of the neural circuit, identification of whether the effect of a receptor is excitatory or inhibitory, and determination of how strong the effect is.

One way of simulating this situation is represented in Figure 1, which is the screen display from an Apple Macintosh running Microsoft Multiplan. The first column symbolizes the current input values for each of a set of receptor cells. Next, the display represents the one-cell-wide slice of neural tissue with two sets of neural cells—receptors and transmitters. The vertical orientation is not required, but does decrease computation time. It also makes possible some limited graphics and plotting. The grid lines that typically provide visual definition for individual locations in a spreadsheet matrix are not visible here. This is a result of the spreadsheet template's hav-

ing been protected so that a user can change only the underlined values.

Between the columns of receptors and transmitters, there is a column for the most recently calculated output of each receptor and a column of question marks. The question marks symbolize a lack of knowledge about the nature and number of connections between the two layers of cells. Finally, there is a single value representing the most recently calculated firing rate of the transmitter neuron that is being recorded. With no receptor stimulation, this transmitter shows a background rate of 100 impulses/sec. The user's problem is to discover which receptors affect the transmitter, whether the effects are excitatory or inhibitory, and how large each effect is.

What is not visible here is which transmitter cell is being recorded and the underlying formula that the spreadsheet program uses to calculate the transmitter's output. In that formula, the total output is the sum of the transmitter's background firing rate and values that represent the effects of receptors. Each value for a particular receptor is the current receptor firing rate multiplied by a gain factor between 1.0 and -1.0 . This gain factor indicates the nature and magnitude of the receptor's effect. Positive values represent excitatory connections, negative values, inhibitory ones. The absolute value of a gain factor determines the percentage of the receptor firing rate that is passed on to the transmitter being recorded. These firing rates and gain factors can be specified in the underlying formula in a variety of ways. They can be entered as constants, they can be referenced in the formula by identifying another location in the spreadsheet that the program is to access in order to find the needed value, or they can be the result of a calculation. In this particular example, the contribution of each receptor in the unknown circuit is the product of a constant and a value found in another spreadsheet location.

After appropriate preparation, a student can be asked to discover the relationships by modifying the levels of stimulation to the receptors. For example, Figure 2 represents a screen display after a user has begun to pass a square wave of light over the receptors. The receptor input values here are in the number of units of light required to produce 100 impulses/sec in the receptors. As indicated by the pull-down menu on the upper right, the spreadsheet is set for automatic calculation. This means that each time the user enters another value, the program automatically calculates or recalculates every other value. For those not familiar with the Macintosh, the pull-down menu is not visible to the user during calculation.

At this point, the user has just entered an input to the fifth receptor cell. Prior entries produced no changes in the transmitter firing rate. However, after the fifth entry, the transmitter firing rate has dropped to 90 impulses/sec. This is a decrease by 10% of the receptor cell's output and indicates two things. First, the fifth receptor has an inhibitory effect. Second, the gain is -0.1 for this connection. Continuing would allow the user to discover that

Figure 1. Screen display of a spreadsheet template simulating the recording of a single transmitter neuron.

Figure 2. Screen display of the spreadsheet after successive stimulation of several receptors.

stimulating the sixth receptor with 100 units of light increases the transmitter firing rate to 190, that the same stimulation of the seventh receptor decreases the transmitter rate to 180, and that similar changes for the eighth, ninth, and tenth receptors produce no effect on the transmitter. This information enables the student to draw a diagram of a neural circuit in which the fifth and seventh receptors with gains of -0.1 , have inhibitory effects and the sixth receptor, with a gain of 1.0 , has an excitatory effect.

A student who has discovered this simple circuit can now be provided with a series of increasingly more complex "unknowns." For example, by duplicating this spreadsheet template and then changing the underlying formula for the transmitter output, the instructor can easily create a variety of additional problems. Furthermore, as is the case here, the designer of the circuit problems can use the column hiding and protection capabilities available in some spreadsheet programs to conceal the formulas and gain factors so that the student cannot see or change them. As a fringe benefit of the fact that the spreadsheet simulator offers two modes of calculation, automatic and manual, a student has the freedom to follow either of two experimental procedures. In one case, the student can change the stimulation to one receptor at a time, with the transmitter output being recalculated after each entry. In the other case, the student can present a stimulus pattern to a set of receptors and then have the effects calculated. Given a neural circuit with unknown characteristics, the student who adopts the wrong procedure—changing more than one value at a time—will also eventually learn an important lesson in experimental methodology at relatively low cost.

LATERAL INHIBITION

Given the computational capabilities of an electronic spreadsheet, neural model simulations can involve the output of several transmitter cells, allowing the user to create neural networks composed of multiple circuits. Consider the process of lateral inhibition thought to be

responsible for visual experiences such as the contrast effect produced by putting patches of a constant shade of gray against various backgrounds differing only in brightness, or the gray spots that appear at the intersections of the Hering Grid, or the differential shading seen in Mach bands—the increasingly brighter series of adjacent gray bands that also appear to have differences in brightness between the edges and the central portion of each band.

Part of a spreadsheet simulation of this latter phenomenon is illustrated in Figure 3. On the left, a section of the spreadsheet is devoted to a "typical circuit." Here the user can enter gain factors for the connections between a set of receptors and a single transmitter cell. This creates a neural circuit that is then replicated in the network on the right. In this example, the central receptor in the typical circuit, $R(0)$, has a gain of 1.0 , and each receptor immediately on either side has a gain of -0.1 . Associated with the circuit is a column of OKs. This spreadsheet template has been programmed to allow only values between 1.0 and -1.0 as gain factors. The OK indicates a value within range.

Once the typical circuit has been designed, each transmitter in the network on the right side of Figure 3 bears that identical relationship to the receptor cells that affect it. Thus, the third transmitter, $T(3)$, for example, receives a 100% excitatory effect from the third receptor, $R(3)$, a -10% inhibitory effect from the second, $R(2)$, and a second -10% effect from the fourth, $R(4)$. Another way of looking at this is to see that each receptor cell has a strong excitatory effect upon its corresponding transmitter cell, for example, $R(3)$ on $T(3)$. In addition, each receptor cell has a slight inhibitory effect on each of the two transmitter cells on either side of its corresponding transmitter cell, for example, $R(3)$ on $T(2)$ and $T(4)$.

In the "Receptor Input" column in Figure 3, several values have already been entered. The first five each represent the number of units of light required to produce a receptor output of 100 impulses/sec, and the next five, 200 impulses/sec. Scrolling the window downward would reveal a total of five receptors with 300 as the input value and another five with 400. For purposes of this example, the receptors are assumed to extend indefinitely beyond

Figure 3. Screen display of a spreadsheet template simulating the operation of a neural network composed of several identical circuits.

those at either end of the set shown in the spreadsheet template. The light input for the last receptor at either end of the set shown is also assumed to extend indefinitely.

Comparing the pattern of transmitter outputs with the corresponding receptor outputs, a student can begin to see the results of the excitatory and inhibitory effects of the receptors. The output of 180 impulses/sec for the third transmitter, T(3), is the sum of its background rate, plus the excitatory effect of the third receptor, R(3), plus the inhibitory effects of the second and fourth receptors, R(2) and R(4). From this information, the background rate for the third transmitter, T(3), can easily be calculated to be 100 impulses/sec.

Of more interest are the effects of different levels of receptor stimulation upon the firing rates for the fifth and sixth transmitter cells, T(5) and T(6). The fifth transmitter, T(5), being inhibited by both the fourth and sixth receptors, R(4) and R(6), shows a firing rate lower than the rates for the first four transmitters. Since the input to the sixth receptor, R(6), is greater, it has a greater inhibitory effect than is produced by the corresponding receptors that inhibit the first four transmitters. Similarly, the sixth transmitter, T(6), shows a higher firing rate than the next three, T(7), T(8), and T(9). This again results from the lack of symmetry in the pattern of inputs for the receptors—R(5) with 100 units, R(7) with 200 units—that inhibit T(6). A similar pattern can be seen in comparing the outputs of T(9), T(10), T(11), and T(12).

To create a visual representation of this pattern of output values, several spreadsheet locations can each be assigned a formula that refers to a transmitter output value, and that utilizes the string-manipulation capability of the spreadsheet program. In Figure 4, the window on the spreadsheet has been scrolled to the right and the display shows a plot of the transmitter cell output values. To prevent the user from having to search all over the spreadsheet for plotted points, each transmitter output value has been divided by the maximum transmitter output and then scaled to fit the size of the display. The location of each asterisk is determined by a formula that fills each cell with the computed number of blank spaces concatenated with

Figure 4. Screen display of a plot of the output of several transmitter cells.

Figure 5. Screen display of the neural network with a new 'typical' circuit and new input values.

Figure 6. Screen display of the plot of the new transmitter output values.

the asterisk. Although the plotting is done on transformed scores and the axes of the graph have been rotated 90° from their conventional orientation, the user can see in this plot both a correspondence to the numbers on the left of the screen and a graphic analog corresponding to the visual experience reported by those who view the Mach bands.

The "typical" circuit in this example is quite simple, but the user can easily make the network more complex. Figure 5 shows a case in which a gain has been specified for each receptor cell in the typical circuit and all receptor inputs are either 100 or 200. The firing rate for each transmitter in the network has been calculated by a formula that includes a constant for that cell's background firing rate, a reference to each of the nine gain-factor locations in the typical circuit, and a reference to the locations of nine receptor-cell output values. This increases the flexibility of the simulation considerably. In addition to being able to change the level of stimulation for each receptor, the user can also design a wide variety of circuits simply by changing the gain factors. This includes deleting any receptor cell from the circuit by setting the gain factor to zero. Naturally, such changes also change the characteristics of the network. The plot of transmit-

ter output values resulting from the changes shown in Figure 5 is shown in Figure 6.

CONCLUSION

The neural model simulations described here offer an opportunity for active student exploration of neural circuits and networks in ways that cost the student very little except in time spent learning. This is especially true when the student already owns a personal microcomputer and the spreadsheet program. In addition, the simulations can be designed in such a way that the student is able to "experiment" with a simulation—experiment both in the sense of discovering the characteristics of an unknown model and in the sense of modifying various components of a known model to see how the simulation is affected.

In this way, a student is able to begin working with intuitively simple neural models and basic components in ways that allow for a natural generalization to the models and explanations required to deal with more complex phenomena (e.g., the operation of feature detectors, etc.). This active involvement in the process of discovery also creates an opportunity for a student to develop both a formal knowledge and an intuitive understanding of structures, processes, and relationships. Furthermore, the level of programming expertise required by a spreadsheet program is low enough that an interested student can and will go one step further to use the programming capabilities of the spreadsheet simulator in creating models of circuits and networks with characteristics other than those provided by the instructor, thereby actively engaging in the process of theoretical modeling.

The examples discussed here by no means exhaust the phenomena that can be modeled using a spreadsheet program. In addition, there are a number of complexities that could be introduced if desired. In both examples, it has been assumed that the input to the receptor cells is in units of stimulation required to produce the corresponding number of neural impulses per second. Other simplifications include representing only two layers of neural cells, and using only a single row of receptor cells in the receptive field. These are not necessary simplifications. An electronic spreadsheet program easily allows the use of a functional relationship other than identity between the level of stimulation and the response rate of the receptors. Furthermore, the typical electronic spreadsheet is powerful enough to deal with the additional complexity and detail involved in adding more rows and/or more layers of cells to the simulation.

To date, student reactions to these neural model simulations have generally been favorable. The reactions, however, have been collected in a context of formative evaluation. Consequently, there are no experimental data to suggest that the simulations have any particular worth. Nonetheless, student comments clearly indicate three things that also appear to be true for simulations in general. First, a preuse discussion and a postuse debriefing are required and do seem to enhance the effective-

ness of the simulations. Second, some limitations of the simulations can be compensated for by proper preuse preparation and postuse debriefing. Third, many students do feel they gain an enhanced depth and quality of understanding from the interactive learning and immediate feedback. (Some of them even report having fun "playing" with the network.) Those who did not feel they gained something from the simulations seem to be either those who felt they already understood the material or those who were confused by some aspect of the material or by the way it was represented in the simulation.

As a final note, several contrasts can be drawn with the more usual procedure of developing a special-purpose program to create each simulation. Most of the advantages and disadvantages of each approach are reflected in the words of one individual who read an earlier version of this paper:

I think I could write a good Pascal program for this purpose in a couple of days which would provide better icons for the particular functions, or "physiological" type stimuli, and provide outputs that might closer match those of a neurophysiologist (e.g., having the speaker give the pulsing rate); however, it might be a year before I could find the time.

Although the neural model simulations described here lack the important advantages offered by sophisticated graphics and sound effects, the spreadsheet program does provide some compensating advantages. First, the cost in development time and effort is relatively small. For someone already familiar with an electronic spreadsheet, both simulations described here can be set up with just a few hours of work. In fact, it takes longer to write a description of one than it does to create it. For someone not familiar with an electronic spreadsheet, the required time would be several hours greater, since a degree of mastery of the spreadsheet program would have to be developed first. However, that extra time would be well invested because that person would also have gained the use of a powerful, multipurpose tool.

This versatility is also the second major advantage of the spreadsheet program over a special-purpose program. The spreadsheet program is already written with a great deal of built-in flexibility. The simulation designer, whether instructor or student, can draft, test, and change a variety of models without having to worry about introducing bugs into the underlying program. To build a general-purpose neural modeling program that could provide this kind of flexibility would require a considerable investment in time. A second aspect to this flexibility is that the designer of the neural models is able to give end users varying degrees of access to the flexibility of the underlying spreadsheet program, while still controlling the entries that can be made and preventing certain kinds of mistakes. Again, this can be done without the considerable investment in time required to build the same kind of end user flexibility into a special-purpose program. Fi-

nally, neither the instructor nor the student needs to know a programming language in order to take advantage of the spreadsheet program's capabilities.

REFERENCES

- BOTT, R., & MUNRO, A. (1977). *Study guide for human information processing*. New York: Academic Press.
- HEWETT, T. T., & PERKEY, D. J. (1984). The mythical "mountain" of software. *Collegiate Microcomputer*, 2, 207-210.
- LINDSAY, P. H., & NORMAN, D. A. (1977). *Human information processing*. New York: Academic Press.
- SILVERT, W. (1984). Teaching ecological modeling with electronic spreadsheets. *Collegiate Microcomputer*, 2, 129-133.
- SMITH, A. L., HEWETT, T. T., PERKEY, D. J., RUSSELL, R. N., BURTON, J. S., EISENSTEIN, B. A., HERR, A. J., CARMONE, F. J., FAROUK, B., GORDON, W. J., & STRONG, G. W. (1984). On selecting a microcomputer for 1800 freshmen at Drexel. *Collegiate Microcomputer*, 2, 107-112.