

---

# Virtual Empirical Investigation: Concept Formation and Theory Justification

---

Dan Kalman

---

**1. INTRODUCTION.** The big picture is important in mathematics. If you have the big picture or the broad insight about what is really going on, you can usually reconstruct details that may have been forgotten. Unfortunately, this is an aspect of the subject that is largely unobserved by students, especially in the first two years of the college curriculum. All too often, they approach mathematics entirely at the molecular level. That is, they try to learn every fact, technique, concept, method, or result with equal emphasis, never seeing how these molecules combine to create larger structures. Little wonder, then, if they don't appreciate how a few key molecules can establish the framework of a much larger whole.

Although I know this, and keep it in mind when I prepare to teach my classes, my good intentions seem to have little impact on many of the students. For example, there is a clear, even compelling, picture that goes with Newton's method: slide down the tangent line until you hit the  $x$ -axis; go up to the curve; repeat. This is the big picture I have in mind when I plan instruction for this topic. Do my students learn this about Newton's method? Do they hang onto this one key picture as the important thing to remember? Alas, no.

This article describes an idea for helping students see the big picture. It involves computer activities that allow students to manipulate mathematical objects at an almost tactile level. There are two reasons for doing this. First, carefully constructed activities can draw a student's attention to the big picture in a way that paper and pencil textbook exercises can not. Second, I believe that these activities can play a role in the college mathematics curriculum akin to the use of physical manipulatives in the elementary grades.

I organize the presentation as follows. First, to give the flavor of the activities, one sample is described in detail. That is followed by a discussion of the pedagogical rationale for these activities and by descriptions of a few more sample activities. Next, there is a brief discussion of evidence supporting the use of computer activities. Finally, I explain how I create the activities.

**2. NEWTON'S METHOD.** The big picture for Newton's method, as described in the introduction, is dynamic and visual. The computer activity is supposed to highlight that big picture. Figure 1 shows the computer screen layout for the activity. The main feature is a window that displays the graph of a function. The expression that defines the function appears in a text box. Students can edit the expression to define any function they wish to consider. Another text box specifies the starting value of  $x$ . A push button triggers the execution of a single step of the Newton algorithm: in an animated display, the student sees a tangent line grow from the point on the curve down to the  $x$ -axis, sees this point marked with a small square, sees a vertical line grow up to the corresponding point of the graph. At the same time, a textual display shows the numerical value of  $x$  and  $f(x)$  for this new point, as well as the index  $n$  that counts the number of iterations as they are performed.

I provide the students with a set of instructions for using this activity, along with a series of specific questions to investigate. For a specified function and starting point,

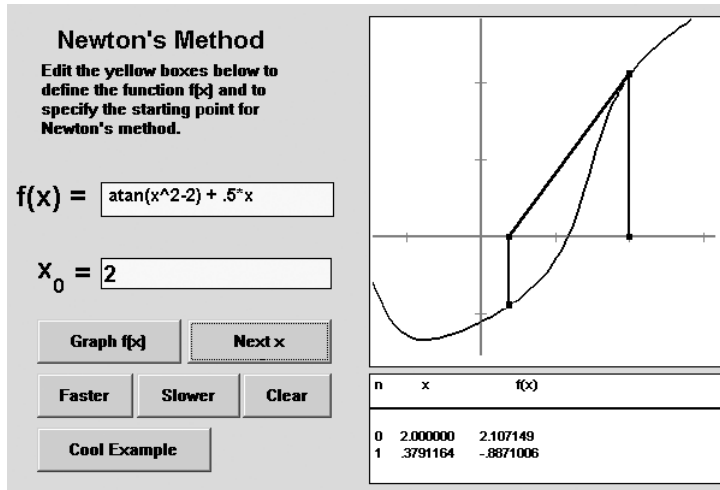


Figure 1. Newton's method.

they are asked to find how many iterations are required to find the root to four decimal place accuracy. Then to six decimal place accuracy. Then eight. I have the students change the initial value of  $x$  and report on how that affects the process. I give students an example with two roots and ask them to determine which initial points lead to each root and which initial points don't lead to any root. Finally, I give them a highly contrived example in which the starting value is very close to a point of period six (Figure 2). The iteration mimics the six-term cycle once, then settles down on one root. Changing the tenth decimal digit of the initial  $x$  produces nearly the identical iteration, but after one tour around the cycle this one eventually settles down to a different root.




Figure 2. Near a point of period six.

Throughout the activity, student interaction with the computer consists mainly of clicking push buttons and watching the results. There is a bit of typing when a new function is to be defined. It is significant that the interaction is completely orchestrated.

Students are not required to give the computer any commands, and they do not have to use a specialized syntax to communicate with the program.

**3. PEDAGOGICAL RATIONALE.** A good part of the underlying rationale for my use of these activities comes from a constructivist learning model. That does not imply that students must essentially rediscover everything for themselves. But it does mean that true learning occurs only after active engagement with ideas. In particular, it is not sufficient to tell your students what you want them to understand.

Teachers should sympathize with this perspective, because most of us have had the experience many times of completely losing the thread of meaning while listening to a lecture or colloquium. At first we can understand the ideas as they are presented. But with each new definition we are less able to keep the important concepts clear. Before long, we are totally at sea. The carefully prepared explanations of the speaker cease to have any meaning, because they describe interconnections between mathematical constructs that we have not internalized.

Now consider the student's experience of Newton's method. After some sort of presentation in class about the main idea, and possibly after reading a similar explanation in the book, the student spends some concentrated time on homework exercises. To what is his attention drawn by this experience? To the recursion formula, to the process of entering numbers in a calculator, perhaps to convergence of the iterates. I believe it is this experience that gives Newton's method meaning for most of the students, and the experience defines what that meaning will be. Students may remember the formula, they may remember an arithmetic algorithm and computing successive estimates of the root. But they are unlikely to retain that compelling picture of sliding down the tangent line.

The Newton's method activity is intended to remedy this problem. It illustrates a number of the principles I try to incorporate into all of my computer activities. Among the most important, in my view, is creating an illusion of tactile interaction. I do not want my students to be distracted by cryptic commands into which they must embed parameters. They should not need to translate questions into a highly abstract computer syntax. Instead, the ideal is for questions to lead directly to action, just as it does for physical manipulatives. Students should be able to conduct experiments to answer questions. With physical manipulatives, the experiments are performed by immediate physical actions: pick them up; arrange them in a line; compare their sizes, shapes, and colors. The act of carrying out each experiment is completely unconscious. The student does not think about how to reach out and pick something up. She just *does* it, with the attention completely focussed on seeing what happens. Just so for my students. I hope that they will quickly understand how to click buttons, manipulate slider bars, change numerical values, and so on, so that they can naturally and thoughtlessly conduct experiments. That is what I mean by *virtual empirical investigation* in the title of the paper.

I have an underlying conviction here that cannot be stated too strongly: it is not sufficient to *tell* the students what you want them to know. No matter how clearly you state an idea, and no matter how attentively students listen, that transaction does not generally result in deep conceptual learning. Certainly there are counterexamples. The very best students in our classes probably do absorb key ideas and concepts from lectures. Indeed, to my recollection, I learned quite a lot in this way as an undergraduate mathematics major, a recollection very likely shared by nearly all college mathematics faculty. But we are not representative models for the majority of our students. Our collective experiences in the classroom should convince us of the ineffectiveness of merely telling students what we want them to know.

**Hare's irate calculus student.** This idea is dramatically illustrated by an incident reported by Hare [2]. In her research, pairs of students worked with interactive computer activities similar to the ones described in this article. In one activity, students were presented with two graphs and asked to determine which is the derivative of the other. Clicking a point of either graph with the mouse revealed the  $x$ - and  $y$ -coordinates of the point, as well as the slope of the curve. Working with this activity, one student came to a sudden realization that the *slope* on one graph was numerically equal to the *y-coordinate* for the other. However, the reaction to this discovery was surprising. The student became quite incensed, even angry. If that was all there was to it, he wanted to know, why had no one ever told him so? It never occurred to him that in fact he probably had been told, and more than once. For this student, true understanding came with interactive experience. Moreover, that understanding was utterly compelling. We see this student asking himself how he could not have known so simple an idea. Ironically, his conclusion is that he would have known it earlier if only someone had bothered to tell him.

**Concept formation.** The discussion so far has concerned the acquisition of meaning. The constructivist model says that meaning obtains from interconnections among many concepts, illuminated by a wealth of examples, viewpoints, and representations. A student learns by active engagement with all of these, and it is the experience of this interaction that creates meaning. That is what is meant by *constructing knowledge*. The motivation for the computer activities, then, is to provide experiences that will contribute to the construction of knowledge in a way that complements the experiences of listening to a lecture, reading a book, or working on homework problems. To the extent that the computer activities enrich the student's experience of a subject, I believe they can support construction of deeper understanding.

**Theory justification.** While conceptual understanding is vitally important, it is not all that we are after. We want our students to know more than what is true—we want them also to understand *why* it is true. That is not quite correct. On some level, empirical investigation *does* convey insight about *why* something happens. But there is a somewhat different epistemological issue in play here. Different academic disciplines are characterized (in part) by their distinct approaches to substantiating knowledge. Factual knowledge in physics is established by different means and with different standards of evidence than in mathematics, for example. For the core courses in the traditional undergraduate mathematics curriculum, we wish our students to come to some understanding of the standards of evidence in mathematics. We want them to understand what we call proof.

For this goal, the reliance on empirical investigation for concept development introduces a potential hazard. If students find the experiential evidence sufficiently compelling, they may see no need for proof. Worse, they may consider the idea of proof to be an unwelcome distraction, imposed apparently as a matter of pedantic excess. This already common perception might be reinforced by an instructional approach that places too great an emphasis on learning by experiment.

By its nature, mathematics is an abstract subject. Developing intuitions about what is true is only the first half of the mathematical method. Equally important is the articulation of an appropriate abstract formulation and the deductive proof within that formulation of the principles that we have discovered intuitively. I hope to convey the importance of both aspects to my students throughout the undergraduate curriculum.

I am concerned that by working so hard to create an illusion of tangibility, I may be distorting the understanding that my students ultimately obtain. Accordingly, in most

of my activities, I try to include examples that highlight the limitations of the empirical approach. Thus, these experiences have two intended goals. First, I hope that by engaging in the computer activities students will acquire a strong conceptual understanding of mathematical ideas. Second, I try to arrange experiences that show why a theoretical formulation is needed. In this way, the computer activities are intended to support both concept formation and theory justification.

In the next section, I will present several more examples of computer activities I use with my classes. For each activity there are exercises for both concept formation and theory justification. Unfortunately, no verbal description of these activities can really convey what it is like to experience them. Both the activities and the software under which they operate are freely available over the Internet. Details will be presented in section 6.

**4. ADDITIONAL EXAMPLES.** I usually have students work with these computer activities in class, for a time that can range from half an hour to the full class period. I provide students with a set of instructions that detail how each activity operates. The instructions also include an outline of experiments to perform and questions to answer. The following discussion describes both aspects of each activity—how the activity operates, and what the students do with it.

**Limits.** This example, one of several computer activities exploring various facets of the limit concept, is intended to convey one particular conception of what  $\lim_{x \rightarrow a} f(x)$  means. The *big picture* I hope students will obtain can be stated in experiential terms:

Examine the values of  $f(x)$  when  $x$  is near, but not equal to,  $a$ . Is there one obvious value suggested for  $f(a)$ ? If so, that value is what we call the *limit*. If not, no limit exists.

In this activity the screen layout (Figure 3) features a graph window and text boxes for tabular output of function values. Students define a function by text input, the point  $a$  at which the limit is to be investigated, and a step size. The main interaction is to click a button and produce a graph of several equally spaced points to the left and

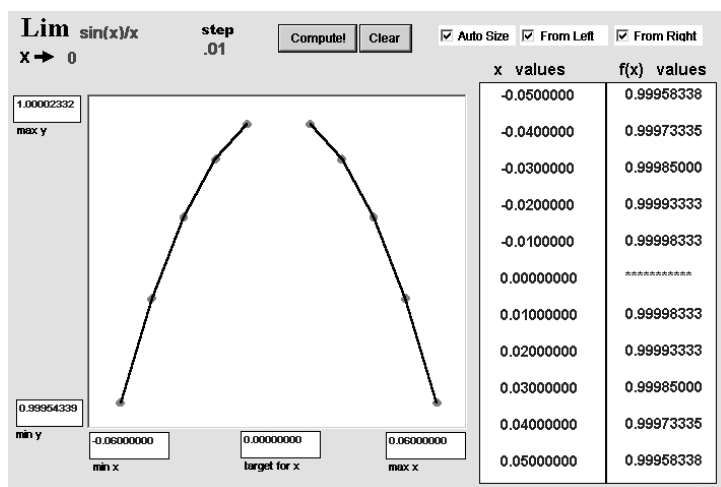


Figure 3. Screen layout for limit exploration.

right of  $a$ , as well as the printed numerical coordinates of the points. The points are joined with straight lines. The instructions explain that this is purely for the sake of visualization and is not an accurate graph of the function between the plotted points. By changing the step size, students can examine the behavior of the function on finer and finer scales. They can also specify that  $a$  should be approached just from the right or just from the left. The graph window automatically sizes itself to fit the ranges of values of  $x$  and  $f(x)$  that occur.

Since most functions defined by simple algebraic expressions are differentiable, on a sufficiently fine scale the points will appear to fall on a straight line. In this situation, there is an obvious interpolated value for  $f(a)$ , and students easily identify it visually and numerically. The activity outline calls for students to estimate values of various limits, which are carefully selected to linearize on different scales. Students find that the step size required to estimate the limit reliably depends on the function. The activity also includes several examples for which limits do not exist. The standard example  $f(x) = \sin(1/x)$  provides one interesting case. Of course, students begin with no preconception of the graph of this function. They simply observe that on any scale, the computer activity fails to reveal any obvious interpolation point. This leads to a discussion of graphing by *connecting the dots*, and its limitations. It also motivates the standard discussion of the behavior of the graph.

Examples with jump discontinuities are also provided, so that students can observe distinct limits from the left and right. Here, I feel it is better to avoid functions defined by cases, or even examples like  $|x|/x$ , because I believe students recognize these as different from “normal” functions and tend to dismiss them as unrepresentative examples. Instead, I use functions like  $f(x) = x/\sqrt{1 - \cos(10x)}$ . Near  $x = 0$ , the quadratic approximation of  $1 - \cos x$  is  $x^2/2$ , so the function  $f$  behaves like a multiple of  $x/|x|$ . But students find it much more credible as the sort of function one might plausibly encounter.

For theory justification, I ask students to examine the behavior of

$$f(x) = |x|(1.01 - .01^{10|x|})/x$$

as  $x$  approaches 0. This function has been carefully contrived so that the apparent limiting behavior changes at different scales. A selection of graphs with the step size ranging from .1 to .000001 appears in Figure 4. At first glance, the graph suggests that

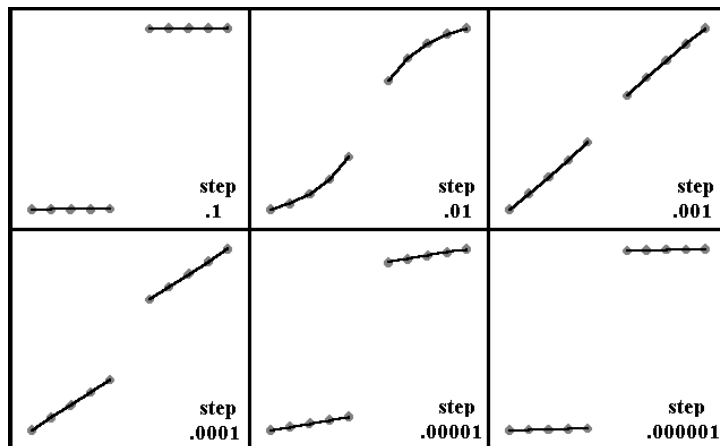


Figure 4. Graphs for  $f(x)$ .

this function has a jump discontinuity at  $x = 0$ . Magnification by a factor of 10 or 100, however, suggests that the limit may exist, and equal 0. Still greater magnification reveals what appears to be a discontinuity after all. This example illustrates dramatically that empirical investigation can be misleading. After my students have looked at this example, I ask how fine a scale must be used to be sure of getting the right answer. They readily observe that no matter how fine the scale is, we cannot really be sure that we are observing the limiting behavior.

I use this final observation as a motivation for a more theoretical treatment of limits. As the students themselves have observed, repeated magnification of a graph can not conclusively show the value of a limit. They begin to appreciate that defining a limit must correspond, in principle, to magnifying the graph by an arbitrary amount. On first exposure, most students find it difficult to see how the formal definition of a limit accomplishes this aim. But at least they can see why some sort of theoretical formulation is needed.

**Derivatives.** For derivatives, the big picture that I try to convey is the following:

Under extreme magnification at a point, a curve will either appear to straighten into a line, or not. If it does, the slope of that line is the derivative at the point; if it does not, there is no derivative at the point.

In the activity students zoom-in on a point of the curve until the curve seems to be a straight line. Then they measure the slope of that line to determine the value of the derivative. The screen layout for this activity is seen in Figure 5. Note that the graph shows both a curve and a straight line. The curve is the graph of a function  $f$ , centered at a point  $(a, f(a))$ . The line is depicted for reference purposes. It is fixed to the point  $(a, f(a))$ , but is free to rotate about that point.

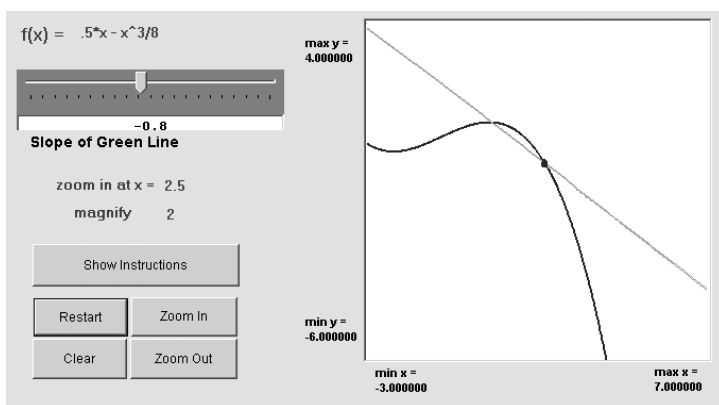


Figure 5. Derivative activity.

Students find the derivative of a function at a point as follows. First, they enter an expression for  $f(x)$ , the value of  $a$ , and a magnification factor. By clicking a button, they repeatedly zoom-in on the graph of  $f$  near  $a$  until the graph appears to become a straight line. Then they rotate the reference line until it coincides with the graph of the function. A slider-bar is used for this purpose. As the slider is dragged with the mouse, students see the reference line rotate. When the reference line and the graph are aligned, the slope can be read off from the settings on the slider bar.

The activity outline directs students to compute the derivative at several different points for a particular function  $f$ . During this activity, their attention is repeatedly drawn to the intended big picture concept, because the process of finding the derivative at each point adheres exactly to the desired image of the derivative. There are also some related activities, such as organizing the derivative values in a table and graphing the results, that the students perform by hand with paper and pencil. These provide background for the concept of a derivative *function*, and for the idea of visually estimating the graph of  $f'(x)$  from the graph of  $f(x)$ .

The activity outline also includes examples where the derivative fails to exist. For these, the student observes that the function fails to become a straight line under repeated magnification. In one example, students examine the function  $f(x) = .8x - \sqrt{e^x - x - 1}$  (Figure 6). Then they compare its behavior with that of  $g(x) = .8x - \sqrt{e^x - x} - .99999$ . In the first case, under repeated magnifications, the graph takes on the appearance of a fixed angle. For  $g(x)$ , the first several magnifications reveal the same behavior as for  $f(x)$ , but after several more magnifications, the curve does eventually straighten out.

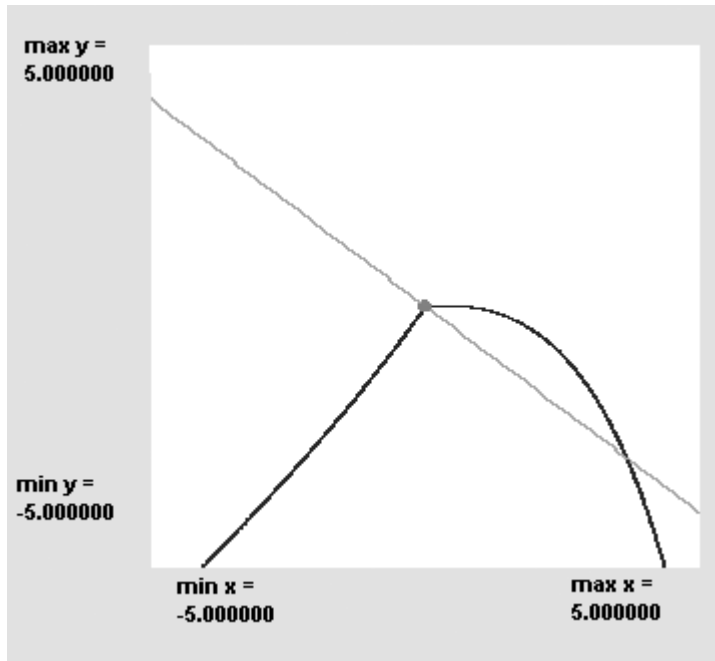


Figure 6. Graph with a cusp.

The example for theory justification in the case of the derivative activity involves the function  $f(x) = \sin x + \sin(50000x)/50000$ . Here, a very low-amplitude, high-frequency oscillation is superimposed on a sine curve. Zooming-in seems to reveal a straight line in the primary oscillation before the effect of the secondary oscillation appears. Several magnifications are shown in Figure 7. The students typically zoom-in until the graph is almost straight, then adjust the reference line to determine the slope. But then they are directed to zoom-in even further, and they find that the curve reveals a new level of variation, only to straighten out a second time. And this time, the slope is something quite different from what was first observed. In later discussion they readily






Figure 7. Six views of  $f(x)$ .

repeat what they observed for limits: no matter how far one zooms in, there can be no certainty that the empirical examination has given the correct answer. As before, this is used as a motivation for introducing a more theoretical formulation of the concept of derivative.

**Riemann integration.** The final sample activity constructs Riemann sums to approximate definite integrals. In this example, there is no simple concise statement of the concept that I seek to reinforce for students. That is because the big picture conception of Riemann integration is very complex.

Consider the typical introduction of the subject. Students are told that the goal is to find the area under a curve; that this area will be approximated using a collection of rectangles; that the bases of the rectangles are defined by partitioning an interval on the  $x$ -axis; and that the height of each rectangle is found by evaluating the function at an  $x$  along the base. They are told about upper sums, lower sums, left-endpoint sums, right-endpoint sums, and midpoint sums. They are told that the true area must lie between any upper and lower sums. They are told something about the limiting process. This is a tremendous amount to take in, and it is frequently all presented in a single lecture or a single section of a textbook.

Naturally, once one has mastered this complex of ideas, the entire development appears intuitive and straightforward. But on a first exposure, the number of ideas to coordinate and organize is likely a bit overwhelming. The students I have observed do not, in general, emerge with anything approaching a coherent understanding of all of the various facets of Riemann integration just indicated.

The objective of the computer activity is to engage students directly in building Riemann sums. Working with an interaction that is nearly tactile enables students to think about the big picture of what they are doing and why they are doing it. In the course of a class period, students work through a progression of three computer screens. The first introduces the idea of approximating the area under a part of a curve using a single rectangle. In the second, students create a series of rectangles. The third automatically produces left- and right-endpoint sums with equal subdivisions for a user-specified number of terms, in order to investigate the limiting process for defining the integral. In the interest of brevity, I focus only on the second, a portion of whose screen layout appears in Figure 8.

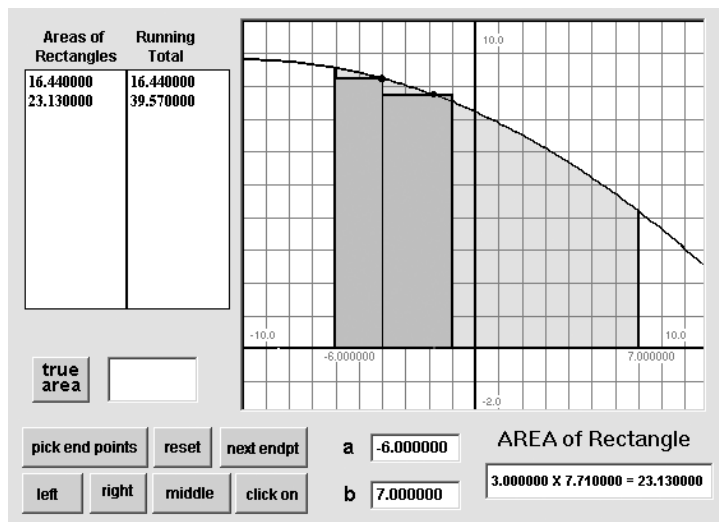


Figure 8. Area activity.

The activity proceeds as follows. First, the student selects a function. A predefined set of functions is included in the activity, and there is a push button for choosing among them. The student clicks this button repeatedly until the graph of the desired selection appears in the graph window. Next, the student selects the endpoints for the integral by clicking with the mouse on the  $x$ -axis in the graph window. Rectangles are defined one at a time. The right side of each rectangle is defined by clicking on the  $x$ -axis. The height of the rectangle is the value of  $f(x)$ , where there are four options for  $x$ : the left or right endpoint of the interval, its midpoint, or a point of the interval selected by mouse click. The student selects one of these options by clicking a push button. At this point, there is an animated display in the graph window. A colored line is shown growing vertically from the  $x$ -axis to the curve, and from that point horizontally to create the top of the rectangle. This rectangle then fills up with a new color. The displays are intended to dramatize the process of determining first the top of the rectangle and then the area within the rectangle. When the rectangle is complete, its area is shown in a textual display and added to a running total. The figure shows the result after completing two rectangles. When the final rectangle is complete, the student can obtain a value for the true area by clicking a button.

The activity outline directs students to construct a variety of rectangular approximations to the area under the graph using this setup. In completing these tasks, the students are repeatedly required to think about the underlying ideas of rectangles, how the width and height are determined, and how the resulting area relates to the area under the curve. Note that the process of constructing the rectangles is virtually a physical manipulation. Students do not have to carry out the mechanics of computation, and they see the results of their actions essentially instantly.

The Riemann sum computer activity does not include any theory justification exercises. For this topic, the big picture concept is quite complicated, and I am content if the students acquire a clear understanding of it.

**5. ARE COMPUTER ACTIVITIES EFFECTIVE?** The reader may wish to know what evidence I am able to provide in support of the use of these computer activities.

Indeed, what evidence is there that students are any more likely to achieve significant conceptual understanding using computer activities than they would without them?

At this point, my evidence is anecdotal and judgmental. I have observed repeatedly that class time spent working with computer activities results in a higher level of engagement and interest than time spent in traditional lecture or class discussion. In a recent linear algebra class, for example, quite a few of the students were reluctant to leave at the end of the period because they were so involved with the computer activity. I have never experienced that sort of engagement in a traditional classroom activity. So, one point that can be made in favor of these activities is that students appear to value them and consider them interesting and worthwhile.

But are the activities educationally worthwhile? In addition to finding the activities interesting and captivating, are the students actually *learning* anything? This is difficult to determine. In fact, it may be nearly impossible to quantify the precise role of any experience or set of experiences in a student's construction of knowledge. Even the students themselves are unlikely to recognize how a computer activity contributed to their ultimate understanding of a concept. Once they have acquired conceptual understanding, they may easily discount either its conceptual depth or the difficulty involved in achieving understanding. (This point is dramatically illustrated by the anecdote presented earlier of Hare's irate calculus student, who inferred from the simplicity of his insight that it had never been pointed out during previous instruction.) And it would obviously be silly to claim that conceptual understanding can be achieved *only* through the use of computer activities. These observations suggest that some sort of empirical demonstration of the impact of virtual empirical investigation would be very difficult to arrange.

Another source of evidence resides in the judgment of teachers like myself who have embraced this sort of classroom activity. I can sincerely report that these computer activities *feel* right to me. Other teachers who have seen them report similar reactions. That is highly subjective, but it is evidence and should not be dismissed.

**6. CREATING COMPUTER ACTIVITIES.** The primary focus of this article is to present a pedagogical rationale for a certain type of instructional computer activity and to describe several sample activities. I hope that other teachers will want to experiment with these activities (instructions for obtaining them appear in what follows). I also hope that some teachers will want to develop their own instructional activities. Accordingly, this section discusses the design and implementation methodology that I use.

All of my computer activities were created using a software product called Mathwright. That is not the only approach available. One alternative is to program in Java, for example, to create applets that are accessed via the Internet (see [5]). Similar kinds of interactions also appear in commercial educational software, such as ODE Architect [1]. On the other hand, learning to use Mathwright is easy (and quick). Recently, Maple and Mathematica have begun to support webpage capabilities, promising to offer another alternative for creating applets. At present, however, I am not aware of any option other than Mathwright that offers the same combination of ease of use and expressive power.

Mathwright includes separate programs for creating and operating interactive computer activities. Activities are created with the *author* program. The composition style is highly intuitive. Screen components (graph windows, text windows, etc.) are selected from menus, then dragged to the desired location and sized using the mouse. Attributes are defined by filling out simple interactive forms. For example, in creating a graph window, the user can select menu items to change the graph color, coordinate

system, appearance of axes and gridlines, and so on. A simple script can be attached to each screen component, specifying actions to carry out when the component is clicked with the mouse. The *reader* program allows users to open and operate activities developed using the author program. For a more detailed description of the Mathwright paradigm see [2]. Note that Mathwright operates only under a Windows operating system and is not available for Macintosh or Unix environments.

Mathwright was a featured software package for several years in the MAA's Interactive Mathematics Text Project (IMTP) [3]. Quite a number of Mathwright activities were developed by participants in the IMTP, as well as other Mathwright users, and are currently available over the Internet from the Mathwright Library website (<http://www.mathwright.com>). An earlier version of the library, supported by an NSF grant, provided free access to the reader software and to the entire collection of activities. This collection was reviewed in [4]. At the current library, necessary software and a variety of sample activities are distributed at no charge. Access to the complete holdings of the library requires payment of a nominal membership fee.

Originally, Mathwright activities were saved as documents on the user's computer and opened with the reader software, in much the same way that word processing or spreadsheet software can be used to create and open the corresponding types of documents. There is now an additional option for hosting Mathwright activities on webpages. For this approach, the user must install a free software extension for the webbrowser, currently available only for use with Internet Explorer. Reference [6] provides a detailed description of the use of Mathwright to create webpages.

I maintain two webpages where visitors can obtain free access to the activities that I have developed, as well as to the software necessary to use the activities. At <http://www.dankalman.net/mwweb> are several Mathwright activities available as either webpages or for download and use offline. This website also includes instructions for downloading and installing necessary software. Following the terminology of [6], the activities at this webpage are referred to as *microworlds*. The *Magnify* microworld includes both the activities for limits and for derivatives, and the *Area* microworld includes the Riemann integration activities described earlier. Three linear algebra microworlds available at my webpage have also been included in the MAA's Mathematical Sciences Digital Library (at <http://www.mathdl.org/mathDL/>).

The Newton's method activity discussed at the start of the paper is currently not available in a microworld format. Instead, it operates under an earlier version of the Mathwright reader. This software, together with the Newton's method activity (among others), is available at <http://www.dankalman.net/mathwright>. The webpage also includes the necessary download and installation instructions.

**7. CONCLUSION.** There are many different instructional uses of computers. They can provide structured drill problems with immediate feedback to students, as well as the closely allied activity of interactive testing. They support dramatic graphics and visualization for classroom demonstrations. Computers can take over the drudgery of computation or symbolic manipulation so that students can proceed quickly from questions to answers. And computers offer a variety of enhancements in the management of textual information, including hypertext linking and electronic searching. But all of these uses either duplicate or extend only slightly capabilities that teachers already have without computers.

By contrast, the interactive computer activities of the sort described here are something completely different. Under their spell, students have the illusion of exploring mathematical ideas by sight and touch. Computer environments have the ability to simulate reality, with the teacher dictating exactly what the fabric of that reality will be.

What could be better suited to mathematics instruction? This is a new opportunity for students and teachers. To me it is the most exciting arena for instructional computing.

**DEDICATION.** This paper is dedicated to the memory of James White, creator of Mathwright and publisher of the Mathwright Library. My ideas about mathematical pedagogy and instructional computing bear the unmistakable stamp of his influence.

#### REFERENCES

---

1. R. L. Borrelli and C. S. Coleman, Modeling and visualization in the introductory ODE course, in *Revolutions in Differential Equations*, M. J. Kallaher, ed., Mathematical Association of America, Washington, D. C., 1999, pp. 1–12.
2. A. Hare, Software review: Mathwright author and player, *College Math. J.* **28** (1997) 140–144.
3. Interactive Mathematics Text Project, *FOCUS* **13** (1993) 28.
4. D. Kalman, Software review: New Mathwright Library, *College Math. J.* **30** (1999) 398–405.
5. F. Wattenberg, B. Stewart, and S. Alejandro, Lite applets, *Journal of Online Mathematics and its Applications* **2** (2002); available at <http://www.joma.org>.
6. J. E. White, Introducing Mathwright microworlds, *Journal of Online Mathematics and its Applications* **2** (2002); available at <http://www.joma.org>.

**DAN KALMAN** received his Ph.D. from the University of Wisconsin in 1980. Before joining the mathematics faculty at American University in 1993, he worked for eight years in the aerospace industry in Southern California. His involvement with interactive computer activities began in 1994 when he participated in an MAA IMTP workshop. Kalman is a past associate executive director of the MAA, author of a book published by the MAA, and frequent contributor to MAA journals. He delights in puns and word play of all kinds, and is an avid fan of Douglas Adams, J. R. R. Tolkien, and Gilbert and Sullivan.

*Department of Mathematics and Statistics, American University, 4400 Massachusetts Avenue NW, Washington, DC 20016-8050*  
*kalman@american.edu*