

# Into the Starting Gate: On Computing and the Curriculum

ROBERT MCCLINTOCK

*Teachers College, Columbia University*

Who knows it half, speaks much, and is always wrong; who knows it wholly, inclines to act, and speaks seldom or late.

—Goethe, *Wilhelm Meister's Apprenticeship*

## THE STARTING GATE

Numerous efforts in higher education and the schools have aimed to make computing an effective tool serving the entire curriculum by helping to make the diverse fruits of academic culture available to students.<sup>1</sup> Despite such efforts, however, computers have yet to prove very useful *substantively* in education. Using computers in teaching a subject is not the same as basing instruction in the subject on computers. Only with the latter use could substantive excellence in the computer-based application be claimed, and this is simply not yet possible. Hence, the goal of excellence in subject matter applications of computers is still distant.<sup>2</sup> Those designing such software undertakings intend that the computer should become an effective tool through which students can study the content of the matter at hand, but more often than not, what happens is that the computer becomes the object of study, not a tool for the study of some subject in depth.

To some degree this conversion of the would-be tool into the object of study occurs because the prerequisites for such courses have not included requiring that students have adequate skill in using computers. But to a larger degree, the computers become the objects of study, not the tools of it, because there is little substantive material available to be studied through the

*This article has been prepared with support from The Center for Intelligent Tools in Education and indirectly from IBM. For this assistance I am most grateful; responsibility for the resulting ideas and opinions, of course, lies solely with me. I thank generous colleagues for their useful comments, suggestions, and encouragement, in particular Brad McCormick, Chris Pino, John Black, Bob Taylor, Amy Heebner, Terri Bush, Janet Asteroff, Frank Moretti, Jinx Roosevelt, and Maxine Greene. A somewhat different version of this article appeared in the Ssigcue Bulletin, published by ACM.*

computers. Thus, the current situation is the one so widely bemoaned: Good educational software is not at hand. Let us try to grasp why pedagogically excellent software is currently so scarce. I believe that the situation is a structural deficiency that can be corrected with serious effort.

In what follows, I am going to argue something that is at once limited and fundamental, and I am going to do it rather single-mindedly, with the risk of seeming to reduce fascinating complexities to a one-thing-needful. To minimize that risk, I would like first to situate the argument. The deficiency of educational software has to do, so to speak, with the importance and difficulty of getting set at the starting point. Computers become the objects of study, not its tools, because nothing yet available is quite ready to serve the educational functions anticipated for it. The structural deficiency is that computer-based education is not ready to begin. This observation has to do neither with what will happen beyond the starting point, once computer-based education has begun, nor with the preparatory sequence that has led up close to the starting point, enabling us to wonder why it has not yet begun. A more elaborate metaphor may help to illuminate this situation.

For matters dealing with the curriculum, metaphors concerning the race track are most appropriate, for the word derived from the Latin term for such a track. The starting point about which I speak is like the starting gate at the race track and the difficulty with educational software is like the difficulty of getting a balky horse into the gate. To expand the metaphor a bit, computing is like a young thoroughbred that has been growing into its physical potential through the work of the hardware and software developers and put through rigorous training by computer science. Everyone says that the young thoroughbred will give a real challenge to the long-dominant steed, the print-based culture, in a match race, and the mature horse stands ready to take the shy challenger on. The problem is to get the challenger to the starting gate, for computing is still frisky and high spirited and we do not yet really know how to bring it into the gate, kicking, bucking, prancing.

This image, getting into the starting gate, situates the argument I want to make. The argument does not pertain to what has built computing up into a potential challenger to the print-based culture, nor does it describe how computing will run the race once it has started. Instead, it concerns one limited but important matter: what needs to be done to get computing into the starting gate so the race can be run. At this juncture, educational software is deficient for bringing computing to the starting point of a contest with print for supremacy in the curriculum. The reason for this deficiency will become clear; correcting it will prove feasible.

## THE STUDY TREE

In order to grasp why educational software is structurally deficient, one needs to analyze as a knowledge structure the intellectual materials used to

support a course. Let us do this loosely in the manner of information theory, abstracting from familiar phenomena various general categories, roughly prototyping them to grasp certain basic points, leaving to some ensuing occasion the refinement of the types and their careful application in describing the information in the curriculum. To begin, I stipulate for the ensuing discussion a controlling definition: The purpose of any course is to evoke in students comprehension of a knowledge structure comprising a quantity of information and skill in the use of it. A course is a knowledge structure comprising an amount of information. To be sure, the character and quality of that information are essential to determining the worth and meaning of the course, its distinctive features. But for our purposes here the distinctive features of courses are not significant; we should be far more abstract, general, descriptive: A course comprises an *amount* of information.

Students acquire a portion of that information and skill in its use by studying the course and teachers seek to impart command of the information by instructing students in the course. When I spoke above of the computer's becoming the object of instruction rather than its tool, I might have put it slightly differently in light of our controlling definition: The computer has so far proved to be a preeminently effective means for studying the information of a course only insofar as the knowledge structure that the course comprises is predominantly about the computer and its uses. Even in these cases, a great deal of the study *about the computer* is still done from books, journals, and manuals—through tools other than the computer. Why is this the case?

To help find an answer, let us develop certain abstract information forms pertaining to the knowledge structure of an academic course. To begin, a course consists of a body of *required materials*, a textbook or set of required readings, lessons, drills, and the like. In those courses where computers naturally have become effective tools of study, *computing*, in one or another form, constitutes the required materials, the primary information form of the course, and hence certain familiar courses, ones well delivered through computers, have become well established, namely those that impart information, elementary and advanced, on how to use computers: computer literacy and computer science. With respect to the general problem of integrating computing into the curriculum, however, these are anomalous. In most courses, in contrast, the fundamental information form—the required materials—is not the computer, but a body of information that has nothing to do with the computer any more than it has substantively to do with printing, photography, or speech, the other basic media through which it might be delivered.

A course, understood as a knowledge structure, consists first of a set of required materials. On the literal level, a finite amount of information will encode these required materials, roughly 2.5 megabytes plus or minus quite a bit according to the character and rigor of the course. This 2.5 megabytes plus/minus of information will be divided into a number of units, sections, and chapters in the texts, as well as assignments and exercises in the se-

quence of classes. This articulation of the required materials really gives rise to two associated knowledge trees, the teaching tree and the study tree.<sup>3</sup> The teaching tree includes the complex of materials that the teacher needs to draw from in order to guide and oversee the students' work with the required materials effectively. The study tree contains that set of recommended and ancillary materials that the students, aggregating their individual activities together, draw on in their efforts to fulfill the requirements of the course.

Both the teaching tree and the study tree can be conceptualized as broadly branching trees descending several levels from the required materials of the course as their root. These knowledge trees will branch down a varying number of levels according to the degree of specialization embodied in the course and the academic quality of the teaching and study associated with it. The teaching tree utilized by an overburdened classroom teacher, who never really has time to do more than read the students' texts with adult care, may branch down only one level and to very few nodes on that. All the same, the quantity of information included within that *shallow, narrow* branching is substantial.

Note how quantitative information theory gives robust precision to these traditional scholastic pejoratives, which predate the theory by millennia. The person whose teaching tree is "shallow," reaching down only to one or two subsidiary levels of reference, will come across as exactly that, just as the person whose teaching tree defines a "narrow" pattern of branching will impress listeners precisely as someone egregiously narrow in view. Likewise, the venerable terms of academic acclaim, *deep* and *broad*, take on rigorous meaning with reference to the teaching tree: The deep teacher can branch down through many levels of organized reference in responding to a query while the broad master can branch out along many possible paths of reflection from any given start.

With respect to the problem of creating good courseware, of getting to the starting point, the study tree is even more important than the teaching tree, for the study tree defines the scope and structure of the information that needs to be made available through computers if we are to realize their potential as educationally effective tools. Too little attention has been paid to the elementary features of the study tree as an abstract information form. The stark feature of it is *size*. Beyond the first few years of schooling, the study tree in any course encompasses a large quantity of information. A study tree would be rooted in the required materials, from which there would be a first level of branching with a separate branch for each student in a class or course, with those branches leading to the information that each student mobilizes in absorbing the required materials, and each of those branches would then further branch one or more times as each student was called on to go beyond the required materials to deal with recommended topics and special assignments. If the contents of this tree are to be computer-based,

that is, accessible through the computer, the computer must have a lot of information stored within it.

Let us roughly estimate the vast quantities of information enclosed within such study trees.<sup>4</sup> For instance, most reading these remarks will have once been undergraduates at quality colleges and they will recall struggling with the knowledge structure of a tough liberal arts course: a book a week, that is, 1-2 megabytes, or a total of about 20 megabytes of required reading, plus a serious term paper, which would require careful study of another 5 megabytes, two or three further books, along with more cursory perusal of a good deal more data, all different from one student to the next. Let us say there were twenty students in the course. The information capacity of the study tree would be the following:

20 meg	Required materials worked with by each of the twenty students
100 meg	Primary materials drawn on in writing papers, 5 meg per student
500 meg	Background materials consulted in preparing papers and class assignments, 25 meg per student
620 meg	Total

In reality, the particular course you recall may well have defined a still much larger knowledge structure, for you could never have gotten away with basing your term paper on a mere two or three books selected from only ten or fifteen possibilities.

Be that as it may, the point here is not to find the exact quantity of information included in either the teaching tree or the study tree of the typical course, whether high school, college, or graduate school. In all cases the structure encompasses a very large amount of information, especially relative to the storage capacities of available academic computing systems. One might imagine a not terribly demanding high school course generating a study tree comprising the following amounts of information:

2 meg	Required materials worked with by each of thirty students
45 meg	Primary materials drawn on in writing papers, 1.5 meg per student
90 meg	Background materials consulted in preparing papers and class assignments, 3 meg per student
137 meg	Total

This is considerably less information than that in the course we imagined earlier, but all the same, 137 megabytes is not a trivial amount.

Naturally, the point of recognizing the large information content of courses is not to argue that the bigger the better, that the size of courseware indicates its quality. Bigger is not necessarily better and quite possibly in

courseware less is more, but, I would argue, only beyond a certain threshold. Should computer-based courseware have an information content below a certain minimum level, that courseware will be limited, thin, too easily exhausted, and even if pedagogically artful it will soon prove to be of minimal use because students will learn too little through it to make it worth their effort. What is the threshold? We cannot be entirely sure, but we give ourselves a good order of magnitude by looking at the amount of information contained in current non-computer-based courseware. This defines a big study tree.

If one uses actual texts on the market to estimate the size of the study tree, the above approximations appear on the low side. My daughter, a senior in high school, is taking an elective English course on the short story. The text for the course, *Story to Anti-Story* edited by Mary Rohrberger,<sup>5</sup> includes stories by some fifty-five authors, giving brief biographical notes on each and citing their major works, say on the average five books per author. It does not cite any critical literature, although the teacher's manual suggests various references that teachers could use to set their students off on the study tree of criticism, should they want to do so. By rough count the text contains 2.8 megabytes of information and points directly to some 275 megabytes of further information (counting 1 megabyte per book cited) and indirectly to a great deal more in the critical literature indicated in the manual, not a small amount for a high school elective.

Study trees described by college texts encompass far larger bodies of information. H. W. Janson's *History of Art*,<sup>6</sup> a text widely used in introductory art history surveys for undergraduates, contains 2.4 megabytes of information in the form of text and it points to well over 600 megabytes of information through its bibliography. In addition, it includes 912 halftone pictures and 143 color plates. The computer equivalent of these, using an IBM PC with a standard color graphics card, would require (depending on the resolution used) between 64 and 256 K of information for each halftone, not accounting for compression techniques in storing the images, or between 57 and 228 megabytes, and for all of the color plates another 36 megabytes. The text thus encodes a large amount of information. I contend that such a large amount of information defines an important order of magnitude; not that being so large makes a text a good text, but that good texts will have, among the many other qualities that make them good, a scale of this size. Because the information content requisite for a course is on this order of magnitude, computers cannot be the primary tools for studying a course unless information in such quantities is accessible through them. Whether then they prove to be good tools for studying the course depends on other factors.

One might argue that a text laden with a such a quantity of information as that in Janson's results uniquely from its being a work in art history. In that subject, illustrations are essential and as quantities of information, a picture is worth, not merely a thousand words, but tens of thousand, if the reso-

lution is reasonably high and data compression is not used. Other sorts of texts are not such smaller as knowledge structures. Were I to teach a survey of modern European history through a general text, I would choose *A Modern History of Europe* by Eugen Weber, for the book has a noteworthy intelligence and comprehensiveness.<sup>7</sup> It is also relatively small as college texts of the type go: about 3.8 megabytes of text, somewhat over three hundred illustrations, forty maps, and about fifteen hundred books in its bibliographies, defining a study tree that totals some 2 gigabytes of data. The various volumes of the Norton anthologies of literature, used in diverse college surveys, each contain 8 megabytes, plus or minus one or so, of text and provide bibliographic pointers to well over a thousand books, at least a gigabyte of information each—and Norton sells many thousands of each volume annually, there being two for English literature, two for American, two for world literature, and one for poetry and another for women's studies, not to mention a somewhat smaller *Reader* and a still smaller *Sampler* for less-focused, less-ambitious surveys.

Textbooks thus define copious study trees, but reaching these proportions is not unique to textbooks, whose *raison d'être* is to give thorough introductions to whole fields, usually the fruit of labor by an extensive staff. Not a few works of serious scholarship, ones that win wide international audiences, can be described as broad and deep, not only in the loose speech of critics but in the sense set forth here, for they set up a study tree for their serious readers that is broad, in that it branches to many topics, and deep, in that it branches down through those topics into the most detailed literature available.

Look, for instance, at Fernand Braudel's master work, his three-volume *Civilization and Capitalism: 15th–18th Century*.<sup>8</sup> The text contains just over 5 megabytes of verbal information and 400 illustrations and 116 maps and graphs, which would require some 100 megabytes of data to represent digitally with some compression. In addition the work includes over 5,700 notes, pointers to a vast field of reference. How should the quantity of information to which these point be counted? These notes are not mere citations of this or that quoted phrase. Generally they are well-targeted references to the monographs, sometimes resolved down to a single page, sometimes general references to entire works. Occasionally they are delightfully imprecise—"Reference mislaid." A sample of 48 notes showed 35 different works referenced, and these ranged from articles to an eight-volume history of Amsterdam.<sup>9</sup> It would probably be fair to say that on the average each reference pointed to a piece comprising .75 megabyte of data. On the basis of these ratios, the notes point to a little over 3 gigabytes of information.

Not every newly minted Ph.D. can expect to create a field of interpretation that copious, but Braudel's achievement is not off the scale of scholarly attainment. Serious intellectual efforts by individuals frequently result in knowledge structures of this scale and there are many works that match

Braudel's in quantity of information, not a few significantly exceeding it. What sets his apart from many is the qualitative range and originality of interpretation, not merely the size of the knowledge structure to which his work in sum points.

We have here the stark and simple reason for the chronic mediocrity of educational software: Educational software has been deficient in the quantity of its information content by a factor of about one thousand, doing in kilobytes tasks that in other media have been done in megabytes. This is not to suggest that good courseware can be crafted solely from information content, any more than a good meal can be cooked from calories alone. It is to suggest that meals habitually prepared with only a small fraction of the calories that people normally consume will not be nourishing. So too courseware: If it is systematically deficient in information content, it will not nourish the minds of those who study with it, no matter how engaging the studying proves to be.

A wide gulf separates *computer-assisted instruction* and *computer-based education*. The former has been tried and generally stigmatized as leading only to variants on drill and practice. The latter has not genuinely been tried for the computer base does not yet have within it the information that the study tree of education must comprise. No computer-based courseware yet presents even 2 megabytes of integrated textual information on a course subject, laid out for effective study by students; yet as we have just seen, 2 megabytes is the typical verbal content for a not-too-demanding survey text and the use of illustrations can quickly make that data content expand by a factor of one hundred or more. Even if the computer-based courseware had the minimum 2 megabytes of information for the root of required readings, what would then happen with the branching of the study tree? As things stand, the branching would be either extremely shallow and narrow or it would cease to be computer-based. Nothing approximating the information content of a good course now exists for study through a computer. Our smallest estimated study tree above was about 140 megabytes. Perhaps the best marketer of educational software, Sunburst Communications, lists seventy programs for the Apple II in its 1985-1986 catalog, eighty-six disks in all. Assuming each disk carries 170 K of information, the entire list would total a bit over 14 megabytes of code, about one-tenth of the information we have calculated is in the study tree of an ordinary high school course.<sup>10</sup>

Note how much work it would entail to convert the information content of a good course into a form that a computer could manipulate. Let us estimate the task for the smallest study tree, 140 megabytes. Once an optical character reader has been prepared to recognize a particular font, a top-of-the-line model will input about 4 kilobytes of information per minute, the equivalent to a secretary who types error free at some eight hundred words per minute. That optical character scanner would have to be kept going at



full capacity for eight hours per day for almost seventy-three days in order, simply, to convert the 140 megabytes of information into machine-readable form. Time for preparing the scanner to read different fonts would add a few days more and then the information would have to be structured and reorganized so that it could be well accessed through computers and a fairly powerful computing system would be needed to make the resultant knowledge structure available for regular use by students.

This example shows that a significant amount of work will be required simply to prepare the information requisite for one small course; devising the computer-based tools of study that will enable students to master the information better than they could with print-based tools is left entirely out of the calculation. If we go back now to the situating of the argument, we can see the point that needs to be made fairly precisely. Creating good computer-based tools of study will be the work of the race round the curriculum to see whether computer-based or print-based education will prevail in the history all of us are now joined in making. But that race will not start until we manage to get computing into the starting gate, which entails our getting the information content requisite for the aggregate study tree of education into a form accessible to electronic technologies.

Some will here jump to the defense of existing educational software, objecting that these calculations are based on quantities of textual information and that they harken unto that great bugaboo, the electronic book, reducing the computer to an expensive page-turner. What is here defined as the starting point, they would like to leave behind, saying goodbye to all that. At its best, this view rests on a faith in the potential information compaction that might be possible in presenting the fullness of our culture through computers. We do not know how far we can go with that great principle of modernism, less is more. We can be certain at the outset that there is a great difference between compacting data and discarding data. Should less is more prove to be the design principle of computer-based courseware that proves better than the print-based, it will operate after the information in the full study tree has been converted.

To be sure, good educational software will have to be more than masses of text that one scrolls onto the screen; good educational software will make creative use of the interactive potentialities of the medium and it will integrate graphics with text and sound and video and who-knows-what. But all this does not resuscitate educational software sized to floppy disks; pedagogically these are flaccid disks, and the integration of graphics and sound and video and who-knows-what with the maximum use of interactivity will simply add further to the information content requisite for quality courseware. Quality courseware must, at the starting point, be measured in hundreds of megabytes, and it must equal or exceed, relative to printed alternatives, the amount of information that can be managed effectively through it, or else the

courseware will continue to be, as it has inveterately been, intellectually and culturally deficient. Once sufficient, relative to the given culture, we can discover in the fullness of time where the cultural creativity that it can nurture and sustain will lead.

For such reasons, I contend, we cannot shirk the point, shying away from the starting gate: We face vast data inputting tasks, and the capacities of adequate educational computing systems must be significantly increased. Or so these reflections seem to force us to conclude.

## A SEARCHING TEXT

Some will here cavil against the direction of development toward which these reflections so strongly point. In essence, they will contend that the matched race anticipated here if computing can get into the starting gate has in fact already been run with the victory decisively going to print-based curricula. Print, they will claim, has proved, in comparative taste-tests with the screen, to be preferable as the means for delivering the bulk of information to be used in any course. Why bother, they will ask, to input all that data and to increase so very significantly the capacities of adequate educational computing systems so that students can work effectively with it all? Let the study tree remain print-based. After all, research so far shows that people read faster with less fatigue from hard copy than from the screen. Why load it all into the computer when it is there in print and more readable at that?

Certainly the objection carries sufficient weight to merit some reflection. Is reading as done with print a suprahistorical skill? Is the character of what people do when they read a constant relative to historical change? Recall Augustine's difficulty understanding why Bishop Ambrose would read silently. "But while reading, his eyes glanced over the pages, and his heart searched out the sense, but his voice and tongue were silent."<sup>11</sup> Normally in Augustine's time, good readers read aloud, undoubtedly for a complex of reasons. For one, a technical reason, text was usually written with little punctuation and words demarcated from one another: Reading aloud would facilitate the grasping of sense and meaning.<sup>12</sup> Second, a socioeconomic reason, reading scarce texts aloud permitted their multiplication, not in the process of production, but in the process of consumption. We multiply numerous printed copies so that these can be read privately; they produced single copies to be "read" through the multiplier of groups. Third, a conceptual reason, reading aloud, particularly in highly image-laden spaces like cathedrals, cloisters, and ornate studies, would facilitate the mnemonic techniques of the *ars memoria*, then essential to efficient reading when books were scarce.<sup>13</sup> The skills good readers use are not constants in history.

How we interpret certain research findings depends on our remembering that the very skills requisite for reading change and evolve through cultural history. Consider, for example, the work of John D. Gould, research done

with great care.<sup>14</sup> Gould confirms the findings of many others that people read faster from hard copy than from cathode-ray-tube (CRT) screens, but to use these findings to conclude that hard copy is a better way to present information to be read than are computers would be to misuse the research subtly but significantly. He and his colleagues seek to understand why the rate of reading from CRT displays is slower than that of reading from paper. Gould concludes that the phenomenon is well documented, the explanation of it moot; explaining the phenomenon is important because then it should be possible to design computer displays from which text can be read as fast as or faster than it is from paper. It would certainly be nice were such better displays designed, but with respect to the long-term integration of the computer into the curriculum, the matter is irrelevant. To see this clearly, let us do a thought experiment in hypothetical history and then look closely at what Gould has and has not tested.

Our thought experiment is this: An important criterion controlling the presentation of text prior to printing concerned techniques for making the text memorable after a careful reading. Manuscripts were scarce; readers had to commit them to memory on the assumption that the text would not be later at hand should one or another point need to be checked in the course of disputation or studious reflection. In the effort to make text memorable, manuscript illumination was not merely decorative; it was highly functional as an aid to the art of memory. Imagine an investigator, circa 1486, when innovators were beginning to produce printed text in a distinctively *printed* form, margins justified right and left, the book equipped with a title page, table of contents, running heads, chapter breaks, page numbers, and an index—in short, the primary features of the printed book as we know it. One could well imagine that a controlled test of the *memorability* of the text presented in the old, illuminated style and the new printed style would produce results significantly in favor of the old style. A meaningful, comparative test simply could not have been done, evaluating the distinctive memorableness of the illuminated text versus the distinctive memorableness of the printed text—the ease with which it can be stored and retrieved, the way it can be indexed and cited from afar. Printed books and illuminated manuscripts properly solve the problem of memorability in radically different ways and they are thus different systems that cannot be compared directly.<sup>15</sup>

With this thought experiment in mind, let us look closely at what Gould and his colleagues have so carefully tested. Most of their tests require subjects to proofread comparable text under carefully controlled conditions on paper and on computer displays. They find such proofreading to be equally accurate when done through both media, but the proofreading from paper is significantly faster than that from the CRT displays. Note first that these tests are emphatically not tests of the speed or accuracy of proofing text by traditional methods versus computer-assisted methods. Such tests would be very different and might have very different results. Gould's tests were tests

of reading speed under the conditions of proofreading text in the traditional manner. Stated most precisely, Gould has found that when told to read text on a computer screen presented as if it were on paper, people will read it more slowly than they would were it really presented on paper. This finding is perhaps not so surprising.

Let us look at the question in a broader context. Why test reading speed? Everyone knows that speed-reading is good and important. Or should we be substituting "has been" for "is"? Rapid reading has been most important in scanning through printed material, sifting rapidly through things to get to what one really wants to know and think about, at which point a high reading velocity may become less functional than a slower, critical, thoughtful reading. System changes occur. Although one may read the screen more slowly, spelling checkers and other proofing aids may actually make computer-assisted proofreading both faster and more accurate than that done unaided on paper. So too powerful search-and-display algorithms can alter the balance between skimming and studying text that a good reader may choose to strike when he can control the text with a computer in place of the page-flicking thumb. Thus reading from CRT displays may be to reading from paper as the memorability of illuminated manuscripts is to the memorability of printed books—they may be different systems, not directly comparable.<sup>16</sup>

One further point about computers and reading should be made. For the present, the question does not concern reading for entertainment under casual circumstances. The market for trade books that one can consume while working on a tan will very likely long endure. Nevertheless, knock-about ease should not be the paradigmatic norm determining the character of serious reading. Those who read seriously usually do it in a relatively settled place, at a desk or reading carrel, or in a favorite chair. If there are distinct gains to be achieved by converting the computer work station into a reading station, the somewhat more restrictive ergonomics of it is not likely to stand as a powerful obstacle to the change.<sup>17</sup> In the end, computers may prove to be, however useful as tools for writing, somehow unsuitable as tools for serious reading. We cannot uncover the fact of that unsuitability, should it emerge, until the computers have been given a serious trial, and we cannot do that until we have available in them study trees worth the serious student's serious effort.

In this way, we return to the starting point, the conclusion claimed above: We face huge data-inputting tasks and we need to increase significantly the capacities of adequate educational computing systems. I shall venture a few remarks about the data-inputting task and then conclude by describing the sort of system one might configure as "the education machine."

With books, the difference between texts in the public domain and those under copyright has been minor. The reason has been a simple one of economics: With books, the bulk of the cost of production has been the costs of

material and labor, not the cost of royalties to holders of the copyright. With books, the end cost of public domain material would be at most 10 to 15 percent lower than the end cost of material under copyright. With computers, this differential can change significantly: The costs of materials and labor needed to manufacture the intellectual content of a course are potentially much much lower than with books. As a result, as a proportion of the cost of the end product, the costs of royalties to the holders of copyright, should there be any, will be much more significant than with books. The effect will be to make the difference between material in the public domain and material under copyright far more salient.

This difference can become startlingly significant. Assume in a developed CD-ROM<sup>18</sup> market that we plan to sell ten thousand copies of a disk that has on it text equivalent to two hundred books. Our rough pricing calculation, ignoring distribution costs, but assuming the materials on the disk are in the public domain, will be astonishingly low:

\$42,000	Cost of data inputting <sup>19</sup>
\$4,000	Cost of mastering the CD-ROM
\$50,000	Cost of 10,000 copies at \$5.00 each
\$96,000	Total production costs, \$9.60 per disk

If the materials on a disk are under copyright, the production costs would still be \$9.60 each, but to that cost we would have to add the cost of royalties. The materials to be put on the disk we stipulated to be the equivalent of two hundred books. Assuming each author expects something less than he or she would get on a hardbound, printed version—say, on the average, a dollar each—the allotment for royalties per disk would be \$200 and the base cost of the disk for the publisher would have leaped well over twenty times. The simple conclusion: Educational applications using CD-ROM and the like will be pioneered in areas in which intellectually sound systems can be developed from materials in the public domain.

With this observation, we can specify broadly from whence the data inputting that needs to occur is likely to emerge. Where ample material in the public domain, sufficiently high in intellectual quality, exists, the opportunity to develop intellectually substantial educational courseware also exists. Here, strangely, the most conservative parts of the curriculum, those working with the “great tradition,” are surprisingly at an advantage. Literature and history offer vast amounts of intellectually significant sources securely in the public domain. They are not, of course, the only quarters from which such quality sources can come: A multitude of studies done with federal support are likewise in the public domain. Further, many, perhaps most, intellectually serious scholars do not expect or receive significant royalty income from their writing, but write instead to exert intellectual influence and to advance their academic positions. If computer-based courseware utilizing public domain materials were to begin to catch on, these authors might be very willing to

permit royalty-free inclusion of their work in substantively excellent courseware simply so that they would not be excluded from the emerging computer-based study trees.<sup>20</sup>

To create adequate computer-based study trees, we need to convert the information in public domain texts into machine-readable form—no small undertaking, but it is no larger and is in fact smaller than the data-inputting task faced in producing the equivalent set of materials in printed form, which goes on all the time. This point is important and should be reinforced. The three volumes of title listings for *Books in Print, 1984-85* have 5,270 pages of listings, each page averaging about 125 titles, or 660,000 different books in print. Assume that on the average, each comprises 1 megabyte of data and that the data in the books in print in any one year was inputted into print form over ten years. This would mean that the current data-inputting capacity of the U.S. publishing industry is about 66 gigabytes per year. This mode of estimating, further, misses a vast quantity of material prepared annually for printing through serial publications and innumerable fugitive catalogs and reports, so the actual data inputting for print is greater by a factor of two or more. Inputting text for print is more complicated than converting it into machine-readable form and the latter task is much more susceptible to effective automation. One could, without creating a huge organization, create a clearinghouse that would input 6 gigabytes of public domain text annually to machine-readable form and maintain and distribute the collection to courseware developers at a reasonable cost.

In the publishing industry, the data-inputting task is highly distributed with numerous authors, secretaries, editors, and printers sharing the task. A somewhat more centralized procedure may make sense, however, in converting the intellectual content of the print-based curriculum into a computer-based form, especially the portions of it in the public domain. The data-inputting task simply needs to be shouldered and it can be efficiently concentrated in one or a few centers from which educational courseware developers can get the requisite materials for their projects at reasonable fees. Such a clearinghouse would encourage the development of a market for quality courseware by greatly lowering its potential cost and it would release creative energies from the data-inputting task so that they can be dedicated to the really interesting one of finding out how best to work with substantively demanding intellectual materials in a multimedia computing environment.

With an annual budget of roughly \$1 million, such a clearinghouse could annually convert approximately 6 gigabytes of textual information in the public domain into machine-readable form, maintaining the growing collection for distribution at nominal fees to courseware developers. It is hard to estimate precisely how likely it would be that such an effort would become self-financing. In the example above, direct data-production costs were estimated at a rate of \$84 dollars per megabyte of material. If the clearinghouse

converted to machine-readable form 6 gigabytes annually, and if it could do that and maintain the complete collection for distribution at an annual budget of \$1 million, then it would cover its costs were it to distribute each megabyte in its collection twice at a rate of \$84 per megabyte during the useful life of the collection. It would not seem to be unreasonable to anticipate that all the material would circulate into courseware production twice at such a fee. In fact it would seem likely to do so much more frequently unless optical-storage technology completely failed as a computer-related technology, in which case much larger sums than those needed for a clearinghouse would be lost in financing the capital development of the technology.

If the information content of computer-based courseware is to be brought up to a level equaling or exceeding that achieved in print-based courseware, the capacities of the computers used in the study of these materials must be greatly expanded. The textual information base of printed works used in education is a necessary minimum that must be matched. But good computer-based materials will go far beyond that minimum, mixing media in intellectually illuminating ways that are simply infeasible given the physical, logistical constraints of print. In beginning to consider what computing systems would work well pedagogically were an adequate information base available, we begin to look out from the starting gate to contemplate the race to be run when the gate clangs open. To do this, we should explore the pedagogical potentials of high-quality computer-controlled multimedia information environments, paying at first relatively little attention to the constraints of cost.

## CAPITAL AND EDUCATION

In the field of education, researchers feel a certain pressure to be directly relevant, which in these matters translates into an imperative to work with the computers presently in the schools. Let us resist this pressure, for existing school computing environments are both obsolete and deceptive. Obsolete perhaps is self-evident, but some readers might wonder why I call them "deceptive." The present environments deceive because they have been pieced together from available products that people found in the marketplace and tried *ad hoc* to use for educational ends. Through some applications, for instance, the use of word processing to facilitate student writing, these fortuitous adaptations have been immensely fruitful. That is part of their deceptiveness, for these chanced-upon tools are partially functional yet do not necessarily look like the tools that would be created for the purposes being served had those tools been *designed* explicitly, expansively, for that purpose. I have driven a nail with my shoe on one or two occasions, but do not conclude that the shoe indicates what a well-designed hammer should be like. A computing environment designed for educational effectiveness has not yet

been created.<sup>21</sup> After it has been created and its functional effectiveness has begun to be optimized, then its potential cost-effectiveness can be evaluated. Let us stick with this point for a bit, for it has latent in it several important implications concerning the potential relation of capital investment with education.

Serious application of capital investment to education must be tried, tried without hesitant ambivalence. In educational history, the application of capital to the pedagogical process has long lagged. The habit of thinking that capital-intensive tools in education have no real place in the pedagogical process has decisively colored the way computing has tentatively entered education. It accounts for the inertial dominance of the Apple II. It accounts for the chronic presumption that educational technologies are to be found among technologies designed for other purposes that happily can be applied to educational activity. "The Apple II forever" is a slogan that appeals to those who would so apply the small personal computer to education without ever entertaining the possibility that computing in education leads inexorably to education's becoming a capital-intensive undertaking, like so much else in modern life. The Apple II may well endure in classrooms like the overhead projector, signifying primarily that the intensive application of capital investment to the improvement of education did not take place.<sup>22</sup>

Cost-effectiveness calculations are often applied prematurely to education. In other domains, when capital is applied with consequential intent to improving human performance, a thoroughgoing adaptation of the tools to the task takes place within the limits of the design capacities available. Industrial investment does not systematically select the immediately most cost-effective procedure; instead, industrial investment backs rational design, forgoing immediately cost-effective procedures and allocating significant resources to creating new procedures thoughtfully. Industrial investment, over and over again, backs reasoned reflection on experience with substantial resources and proves current cost-effectiveness calculations to be systematically short-sighted.

Compare investment in education with that in numerous other domains. We *spend* vast sums on education, but *invest* very little in it. Thus no distinct industry has developed in association with it. One can see how marked this deficiency is in comparison with other domains of life by studying something like the *Value Line Investment Survey*, which provides information about some sixteen hundred corporations that together account for a substantial part of the gross national product (GNP). The survey groups those corporations into some ninety-two industry categories, only one of which, "Toys and School Supplies," nominally involves education, and the corporations included in it primarily produce toys and games. Gross sales for 1986 by companies listed in this category are estimated at \$4.8 billion. Compare to this the estimated sales for companies in various other of the industrial categories: \$11.1 billion



for toiletries and cosmetics, \$6.1 billion for agricultural equipment, \$11.1 billion for shoes, \$17.5 for office equipment and supplies, \$25.0 for medical supplies, and 23.1 billion for precision instruments.<sup>23</sup> A shift in this balance may be imminent.

Education has been a labor-intensive activity, involving primarily the organized use of time by teachers and students. As historical options, people can say that current levels of educational attainment are excessive, sufficient, inadequate. If they are held to be excessive, it makes sense to cut back and decrease by one or another strategem the total proportion of GNP spent on education. If current levels of educational attainment are sufficient, then people can smile and leave things unchanged, holding educational efforts overall on a maintenance regimen. If, however, current levels of cultural attainment are inadequate, here, there, and everywhere, then some improvement in the effectiveness of the system should be sought. A quick survey of the growing complexities in the midst of which we live should convince most that only with an extraordinary complacency can we hold current educational attainments excessive or sufficient. By measures of both equity and excellence, the attainments achieved through education need systematically to be increased. How?

Investment of more labor-intensive education in the training of more and better teachers may marginally improve the overall cultural attainment of the system. So too might the allocation of a larger proportion of the time available to teachers and students to labor at the process, lengthening the school day and the school year. Improvements to be wrought by such strategies will be limited and expensive all the same, for the law of diminishing returns holds here; they are like agricultural improvements achieved by increasing the intensity of cultivation and planting marginal fields. As a labor-intensive activity, education has already been developed to an approximation of its potentiality. If we want improvements in educational attainment to be more than marginally significant, we need to find ways to make the tools used in educational labor systematically more productive. That is the current imperative of capital investment in education.

Existing capital inputs in education do not appear to be promising vehicles for improving education attainment through capital investment, however. Classrooms and textbooks, the main tools of education across all levels, were essentially inventions of the sixteenth century, which have since then been incrementally improved, especially with the advent of functional design in architecture and large-scale production in publishing. Functionally, however, the classroom and the textbook have long been mature technologies: Resource allocation to them turns on narrow calculations of current cost-effectiveness. It is time again to try investing in the tools of education, to apply capital to the thoughtful creation of new tools for the task. One can see this no-nonsense development of capital-intensive tools reflected in the way capi-

tal is being applied to office automation. Look in your mind's eye at the typical office of 1886 and compare it with one of today: The functions are not so different, but the tools and procedures have radically changed. Look then at the typical classroom of 1886 and compare it with one today: Not only are the functions still largely the same, but so too are the tools and procedures—only the dress and decor, and in some places the underlying demographics, may seem different. *Investing* in education has not happened yet, but we assume that it will, that it should, and that we, or others like us, will be the agents of its happening soon.

In view of such considerations, we should not aim to design a *tentative* education machine, one that will fit comfortably into the given structures, cost little, and leave existing patterns essentially unchanged but incrementally improved. Instead, we aim to introduce into the process expensive tools that will not simply improve education incrementally but will radically restructure its character and limits. The eventual cost-effectiveness of these tools will be systemic, not incremental, and it may be of two types: (1) Existing functions may be achieved with less total expenditure of effort; or (2) educational possibilities not formerly feasible may become attainable. What I am calling the education machine, thus, goes all out and tries to put together the sum of the available technologies in a pedagogically useful way. According to established patterns of expenditure in education, it will be expensive, as were steel mills relative to iron foundries. The basic point, however, is that established patterns of expenditure in education are ones that preclude the serious application of capital to the task. This will not change unless educational inventors devise tools that promise to make the serious application of capital to the task at once meaningful, productive, and profitable.<sup>24</sup>

These remarks on the relation of capital investment and education bring us back to the problem mentioned but not discussed above: Computer-based access to data in amounts commensurate, at the least, with those mobilized in print-based courses is essential but not sufficient—it is the starting gate. All the same, the larger task is not merely to put an extensive quantity of information into the computer so that the knowledge structure that can be studied through it is equal to or greater than the one that can be studied through print. The larger task is also to enliven, to invigorate, to energize that study by making it fully interactive and conveying it not only with text but with sound, images, and all the available media of communication.

Technology clearly drives the design process, not the technology of one or another innovation that catches one's fancy, but the technology of the prevailing system, the technology of print-based education. If we cannot exceed what is possible with our existing print-based tools of education, if we cannot design electronic tools, combinations of hardware and software, that make possible—by a significant factor, relative to present possibilities—a more effective, extensive, deeper transmission of information and ideas to students in search of wisdom, learning, and skill, then the application of capital in-

vestment to the task of pedagogical design will be a waste of resources. It is fashionable to insist that one should first specify pedagogical goals, and then let those goals determine the particular features of technologies, but that design sequence will shackle the imagination. Significant changes often occur when people discover an intriguing technical possibility and then slowly, with wondering excitement, begin to realize how the possibility might be put to human use.

So let us let our technical imaginations run free, not into the realm of the unfeasible, but into that of the full-fledged configuration of what is technically feasible for pedagogical purposes. With respect to hardware in an educational setting, too much attention is paid to selecting this feature relative to that, as if it will be the single feature that will prove significant. Thus one encounters proponents of interactive videodisc, of computer simulations, of networking. Such choices might make sense were the various components extremely expensive and their functions highly overlapping. In actuality, electronic hardware has become cheap, particularly relative to the challenge of learning how to use it, and the diverse components generally serve distinctive functions, all of which have a place in the electronic educator's bag of tools.

Forthwith, then, we contemplate the education machine, currently emerging as an ungainly prototype, but something that will be both feasible and powerful. The education machine will be designed on the assumption that all information of value in the culture will be available in binary code. Further, it will be designed on the conviction that education empowers expression and that an education machine is above all else a powerful tool of expression that the student learns to master for his or her human ends.<sup>25</sup> I will close with a few remarks on each of these assumptions, for it is important, not simply to enunciate them, but to do so with full cognizance of their import.

Different events take place on different time scales. One event that we are now in the midst of began sometime in the 1940s and will probably end sometime in the 2020s or so, perhaps not until the 2040s. This long event that we are unfolding involves the conversion of all forms of storage and retrieval of information in our culture to one base coding system, the binary base. When we speak of multimedia systems, what we are really referring to is this unification of media by grounding the implementation of each on a common, shared form of encoding. The benefit is twofold: All the media become manageable through one system and each taken on its own becomes more tractable and effective in use.

When we speak of multimedia systems, we refer to something extremely complicated and potentially very powerful. Storage and retrieval techniques are cumbersome but well developed with respect to text. Libraries house collections of many millions of books and the separate volumes are well cataloged and their contents generally well indexed. The storage and retrieval of images and sounds is substantial, but not so well developed as that of texts.

The different forms of storage and retrieval, however, are still quite separate from one another, as one can see by observing the use of media collections in libraries large and small. Storing and retrieving all examples of all forms of information by using one single, complex system of coding has become technically feasible and certainly no more difficult to implement than the scramjet aerospace plane, for which vast sums are being contracted.<sup>26</sup> We have begun the slow process of implementing this integration of media, but it will not be without great intellectual effort — “the height charms us, the steps to it do not: with the summit in our eye, we love to walk along the plain.”<sup>27</sup>

Until recently, educational technology seemed to place few conceptual demands on its practitioners. It did not appear to be one of the domains of high intellect. But that too changes. What sorts of problems are these of building a unified multimedia system on a common base of binary code? The problems are not primarily hardware problems, or a bit more precisely, our capacity to configure together the requisite hardware is far further developed than is our capacity for dealing with the other dimensions of the task. These other dimensions involve the indexing needed to make feasible the unprecedented storage and retrieval potentialities that the hardware, in principle, can sustain. Were all significant information in our culture stored in an appropriate form of binary code, in principle, any ordinary person should be able to gain nearly instantaneous access to the information pertaining to any matter, consulting it in the form of text, data for calculations, still or moving pictures, sound, graphics, or what have you — whatever is most appropriate to the matter at hand. Implementation of this possibility, however, requires a tremendous extension of our capacity to organize information, to give it useful addresses, to manipulate it purposefully: Extension of this capacity is the intellectual challenge now presented by educational technology.

As we begin to implement the technical possibilities now opening up, it becomes increasingly apparent that educational technology is a pursuit that needs to be redefined as a knowledge-based undertaking.<sup>28</sup> Objective-based instructional technology functionally takes on the role of the teacher, a rather paternalistic teacher intent on shaping the future behavior of its charges. The objective is a way of behaving to be imparted to others and the technology is a means for imbuing others with the proclivity and skill to so behave. Knowledge-based educational technology functionally takes on the role of the curriculum — perhaps even more precisely, that of the library. Its aim is to organize and present knowledge in ways that suit the cognitive structure of people who are thinking while permitting the appropriate inclusion of all that may have cultural significance to them. Knowledge-based education is not designed to serve explicit objectives, no more than is a library, which is designed to be useful to people who come to it with a wide range of unique, divergent purposes.<sup>29</sup>

As the design of educational technology proceeds increasingly from the recognition that all information of significance in the culture is available through binary code, information science and cognitive science will increasingly become the foundations for further innovation in the field. Consequently, in seeking to develop a multimedia system in the full sense of the word, a number of questions need to be given historically actual answers:

Can the doctrine of fair use be extended to the repertory of audio and video productions so that a person can quote from these media as effectively as from the textual media in the process of expressing his or her ideas?

Can we find ways to index images and sounds so that they can become integrated into powerful systems of random-access storage and retrieval in the same way that printed text long has been?

Can we provide retrieval systems that embody sufficient intelligence to ensure that the free play of curiosity and interest in the user is not daunted by the repeated frustration of his or her effort?

Can we overcome the barrier to mutual understanding that has always inhered in the multiplicity of languages without giving up the stimulus to the richness and diversity of possible meaning that that multiplicity imparts?

Can we maintain the incentive to create new ideas and expressions, and the conditions requisite for attaining them, when the complete, continuous recirculation of past creations can be enjoyed with little effort by all?

Answers to such questions will enable the potentials of knowledge-based educational design to be fulfilled, and with that a powerful education machine, one based on the potentials of a unified coding base for the culture, will be created. Such an education machine must be a multimedia system in the strong sense, one commensurate with the depth and complexity of our present culture, one capable of extending those givens far beyond their existing limits. It must be textual in the full sense, allowing one to work, not merely with some text, but any and all text, as a basis for expressing oneself more fully and significantly than one can expect to do through print. It must be an audio medium in the full sense, allowing one to work, not merely with some recorded voice and sound, but with any and all such voice and sound, allowing one to express oneself to the ear more fully than one could through sound recording alone. It must be a pictorial medium, again in the full sense, making available not simply some pictures, but the full range—any and all pictures, be they still or moving, silent or sound, graphically abstract or visually exact. Above all, this complex system must be knowledge-based, a complex generative grammar indicative of all that can be “as we may think.”<sup>30</sup>

If the education machine will be a knowledge-based system, not an objective-based program, what then will be its purpose and use? The proponents

of traditional objective-based methods of instruction will surely ask this question. The answer is fundamental: expression—self-expression, cultural expression, human expression. A library, a curriculum, the education machine, does not imprint on its users the explicit objectives held dear by those who design it. A library serves the formation and expression of all manner of different purposes brought to it by its users. So too will the education machine. Knowledge-based systems are generative, not determinative; they impart tools, not finished structures, tools that people can use to form conviction, to empower action, to sustain reflection, to nurture hope. The purpose of culture is to empower human expression and a fully computer-based education will do that with effect, enabling people to use the tools of expression to pursue their aims in life.

## Notes

1 See James H. Morris et al., "Andrew: A Distributed Personal Computing Environment," *Communications of the ACM* 29, no. 3 (March 1986): 184–201; Edward Balkovich, Steven Lerman, and Richard P. Parmelee, "Computing in Higher Education: The Athena Experience," *Communications of the ACM* 28, no. 11 (November 1985): 1214–24; and Nicole Yankelovich, Norman Meyrowitz, and Andries van Dam, "Reading and Writing the Electronic Books," *Computer* 18, no. 10 (October 1985): 15–30, for descriptions of the sort of initiatives these efforts are leading to on campuses where the integration of computing into the curriculum is well advanced. Even in these cases, the projects are working toward integration of computers into the curriculum and the actual level of subject matter use is limited.

2 With support from IBM to create the Center for Intelligent Tools in Education (CITE), a group of us at Teachers College have been working to integrate computing into graduate education, particularly in five areas: language instruction, social studies, special education, education of the economically disadvantaged, and educational administration. As part of this effort, we have been assessing the available software. The first of these reports (Howard Budin, Robert Taylor, and Diane Kendall, "Computers and Social Studies: Trends and Directions") was presented at the 7th Annual National Educational Computing Conference, San Jose, California, June 4–6, 1986.

3 I use the computer science concept of "tree" somewhat loosely here as the structure being described is not strictly a branching structure, but rather includes, along with much branching, numerous cross-connections characteristic of networks and graphs. The term *tree* seems most appropriate at this stage of the inquiry, at any rate, because it has a long history of usage, going back through the encyclopedists of the Enlightenment, as a concept for describing the organization of knowledge. Those who think of the traditional tree of knowledge as described by D'Alembert, for instance, must bear with my adoption of the directional inversion wrought by computer scientists, who for some reason or other have perceived their trees as branching downward, and I here follow the usage they have established. A priority for further inquiry along the lines charted in this essay should be the creation of an alternative model that gives a more accurate analogy to the full repertory of ways in which connections between phenomena are made within the full domain of knowledge. For now, however, the concept of a knowledge tree, more or less as is used by computer scientists, will have to suffice.

4 In these estimates I am assuming that a byte is a byte is a byte. This is not an entirely fair assumption. Generally, if software is designed well, a quantity of executable program code will usually occupy a student longer than reading an equal amount of American Standard Code for Information Interchange (ASCII) text code. All the same, treating all code as if it were text code gives an adequate approximation of the information content of software. The following estimates are in the number of bytes that would be required to code the information pointed to by various study trees.

For each character is counted as a byte. Representing the amount of information in this way should not obscure the fact that currently most of the information in print-based courses is obviously not coded in binary form. My basic assumption is simply that for a course to be genuinely *computer-based* the full information content within its study tree should be accessible through computers. When we have gotten the information into such a form, then computing is in the starting gate.

5 Mary Rohrberger, ed., *Story to Anti-Story* (Boston: Houghton Mifflin Company, 1979).

6 H. W. Janson, *History of Art* 2nd ed. (Englewood Cliffs, N.J.: Prentice-Hall, 1977).

7 Eugen Weber, *A Modern History of Europe* (New York: W. W. Norton, 1971).

8 Fernand Braudel, *Civilization and Capitalism: 15th-18th Century*, trans. Sian Reynolds (New York: Harper & Row, 1981, 1982, 1984).

9 *Ibid.*, vol. 1, p. 602.

10 This calculation is of course extremely rough as it treats executable code as equivalent to ASCII text code. Owing to the multiple branching and, even more, the recursive possibilities of the executable code, a given amount of it can significantly engage a student for far longer than an equivalent amount of ASCII code, unless, of course, the ASCII code in turn encoded a thought most wondrous and stimulating. On the other hand, by lumping the executable code on the disks in with the information-bearing code, the total amount of information in the set is somewhat exaggerated.

11 *The Confessions*, Book 6, chap. 3, trans. J. G. Pilkington, in *Basic Writings of Saint Augustine*, ed. Whitney J. Oates (New York: Random House, 1948), p. 75.

12 Here we need to recognize that writing words without a clear demarcation between them made such writing more veracious as a representation of speech than modern conventions where words are written with clear demarcations between them even through they are not so spoken.

13 The first few chapters of Frances A. Yates's great study *The Art of Memory* (Chicago: University of Chicago Press, 1966) provide the essential background to this point.

14 See John D. Gould, "Reading Is Slower from CRT Displays than from Paper: Some Experiments That Fail to Explain Why" (ms. from the author, IBM Research Center, Yorktown Heights, NY 10598).

15 The best book on the introduction of printing into Western culture is Elizabeth L. Eisenstein, *The Printing Press as an Agent of Change*, 2 vols. (Cambridge: Cambridge University Press, 1979). M. T. Clanchy examines in *From Memory to Written Record: England, 1066-1307* (Cambridge: Harvard University Press, 1979) how trust in writing as an authoritative document developed prior to the advent of printing. In this context, the functionality of manuscript illumination would be diminished; nevertheless Clanchy explains much about its functionality (especially pp. 226-30), and shows that its use persisted even in documents with a strictly mundane business purpose.

16 A good, brief survey of the efforts to create a computer-based system of reading is "Hypermedia" by Jeffrey S. Young (*Macworld*, March 1986, pp. 116-21). The articles cited in note 1 above are essential. To these one should add Stephen A. Weyer and Alan H. Borning, "A Prototype Electronic Encyclopedia," in *ACM Transactions on Office Information Systems* 3, no. 1 (January 1985); 63-88; and Doug Lenat, Mayank Prakash, and Mary Shepherd, "CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks," *The AI Magazine*, Winter 1986, pp. 65-85.

17 Of course, in judging whether ergonomics are restrictive, one must also be careful what one is comparing. As we have seen, it is becoming possible to get an incredible wealth of material into relatively small computers. The ergonomics of a computer are restrictive compared with a single book, but the ergonomics of searching through two hundred books for a particular point will be very restrictive compared with doing so with a CD-ROM equipped computer. Borland International seems bent on making an entire set of desktop reference tools memory resident with its Turbo Lightening as a query engine and a CD-ROM for storage (see Bill Machrone and Paul Somerson, "Lightening Strikes" and "A Spark of Lightening," in *PC Magazine* 4, no. 25 [December 10, 1985]: 112-23). In a joint-study project with Don Nix of IBM Research, we are equipping a single workstation with videotape, videodisc, audio tape, and a CD-ROM drive all interfacing with an IBM AT with its normal magnetic storage. As a result, the student will be able to manage material through all these media

for the expression of his or her ideas in a way that would be thoroughly impossible logistically were these media not unified under computer control.

18 To explain "lp records" or "videocassettes" in any setting now would be hopelessly pedantic, as they have become everyday objects. CD-ROM disks, and their progeny, CD-I disks, will probably become equally commonplace. Some readers may now need the jargon unpacked a bit, however. CD-ROM stands for compact disk, read only memory, on which information is stored in a very dense way to be read from the disk by a laser beam. One disk, which is just under five inches in diameter, will hold 550 megabytes of information. The data transfer rate from the disk is relatively high; the seek time spent finding a location to be read is relatively slow in comparison with fixed drives but relatively fast in comparison with floppy disks. Owing to the fact that the drives have for practical purposes already succeeded as consumer electronics appliances in the form of compact audio players, the price at which they are coming into the computer peripheral market is very low relative to their capacities, currently between \$700 and \$1,200, with the low end likely to go down to \$350 or so. The disks are durable and cheap to produce in quantity, and not terribly expensive even in limited quantities. For a comprehensive introduction see *CD-ROM: The New Papyrus—The Current and Future State of the Art*, ed. Steve Lambert and Suzanne Ropiequet (Richmond, Wash.: Microsoft Press, 1986). Two articles in the April 1986 *IEEE Spectrum* give an excellent overview: Peter Pin-Shan Chen, "The Compact Disk ROM: How It Works," pp. 44-49, and Tim Oren and Gary A. Kildall, "The Compact Disk ROM: Applications Software," pp. 49-54.

19 This estimate was arrived at as follows and it covers only direct costs, fairly narrowly construed. The cost of a good optical character reader is \$35,000, which for the estimate will be spread over two years. For a labor cost of \$40,000, annual salary for two operators, one should be able to operate the scanner at least 12 hours daily for 240 working days at an output of 4 K of data per minute. This calculates out as an annual output of .691 gigabytes at direct costs of \$57,500. Assuming the disk carries half a gigabyte, the conversion of the data would adjust to \$42,000. Were one to acquire and run the system in order to input data for just one disk, this estimate would be way too low. The costs favor a large, systematic information conversion project.

20 As there is much material clearly in the public domain, some material apparently in the public domain may not really be there. Rights to forms of publication not explicitly stated in contracts remain with the individual authors of pieces, with the result that journals that hold serial publication rights cannot necessarily grant electronic publication rights.

21 Marc S. Tucker comments well on the extramural origins of the push for computers in the schools in his essay "Computers in the Schools: What Revolution?," *Journal of Communication*, Autumn 1985.

22 My references here and elsewhere to specific types of hardware should not be taken to indicate that capital investment in education is primarily a hardware problem. As hardware becomes cheaper and more powerful, the problem of software becomes ever larger and more costly, costly in dollars and cents and costly in the intellectual demands it places on developers and teachers.

23 These sales estimates are for the companies that the *Value Line Investment Survey* groups under each of these headings; they are not an indication of gross expenditures in any of these areas. The *Survey* is particularly suggestive relative to the point being made here, namely that however much we spend on education, we do not invest significantly in it. The *Survey* is meant to be comprehensively informative about investment opportunities and it describes ostensibly no such opportunities explicitly in the domain of education. A large proportion of the GNP is spent on education, compared with toiletries and cosmetics, shoes, precision instruments, or agricultural equipment, yet the sales of companies that use capital to design products for these industries with sufficient effectiveness to be of potential interest to investors are much larger in all these other domains. Estimated sales for companies the *Survey* lists under publishing is some \$15.0 billion, and a significant part of that (say one quarter) would involve the sale of books and materials for use in education, which qualifies the point slightly. Significant allocation of investment capital to expanded production of textbooks, however, is unlikely to change education.



24 In "Some Reasons for the Poor Uses of Technology in Education," John Henry Martin indicates well how the failure of schools to invest in well-designed tools has stunted the development of educational software (*Educational Leadership*, March 1986, pp. 32-34).

25 In thinking about "the education machine," I have benefitted greatly from working with Don Nix of IBM Watson Research Laboratory on a joint study venture using the experimental authoring language, Handy, that Nix has developed to allow children to control a multimedia computing system.

26 Contrast our national reluctance to suspend current cost-effectiveness calculations, in order to employ large-scale resources for the rational design of educational tools, with our willingness to follow such stratagems in other domains. See, for instance, "Spaceplane Work Set to Start Soon" by John Noble Wilford (*New York Times*, Sunday, April 6, 1986, Section 1, p. 1): "In an important step toward developing an aerospace plane, a potential successor to the space shuttle, the Government plans to award contracts this month for the first full-scale test engines and structural components. Air Force officials said the contracts, worth from \$300 million to \$400 million, could be announced next week. The aerospace plane, equipped with scramjet engines that burn their fuel in an airstream that moves at supersonic speeds, is expected to be capable of taking off from a runway and quickly accelerating to speeds 12 or even 25 times the speed of sound." Here is the contemporary, mundane, everyday reality of suspended cost-effectiveness in favor of large-scale investment in possibilities. If in space, why not in education?

27 Johann von Goethe, *Wilhelm Meister's Apprenticeship*, trans. Thomas Carlyle (New York: Collier Books, 1962), p. 447.

28 My colleague, John Black, is developing the distinction between objective-based and knowledge-based instructional design, a distinction he presented in a paper summarized to our Department Colloquium, "Knowledge-Based Instructional Design," March 20, 1986, soon to be published in this series.

29 A wonderful vignette about the befuddlement that arises in the objective-based mentality when confronted with the knowledge-based design of a large library is that by Robert Musil in chap. 100 of *The Man without Qualities*, "General Stumm Invades the State Library and Gathers Some Experience with Regard to Librarians, Library Attendants, and Intellectual Order" (trans. Eithne Wilkins and Ernst Kaiser, vol. 2 [London: Pan Books, 1979], pp. 191-98.)

30 There is a deep, Propertian humanism in Vannevar Bush's phrase and his prescient article "As We May Think," *The Atlantic Monthly*, July 1945, pp. 101-08.