[1]

How We Will Learn

By Danny Hillis, November 9,2000

I have always envied Alexander the Great, because he had Aristotle as a personal tutor. In those days, Aristotle knew pretty much everything there was to know. Even better, Aristotle understood the mind of Alexander. He understood which topics interested Alexander, what Alexander knew and did not know, and what kinds of explanations Alexander preferred. Aristotle had been a student of Plato, and he was himself a great teacher. We know from his writings that he was full of examples, explanations, arguments, and stories. Through Aristotle, Alexander had the knowledge of the world at his command.

Of course no one today knows all that is known, in the sense that Aristotle did. Now there is far too much knowledge for that to be possible. The scientific revolution, and the technological revolution that followed it, led to a self-reinforcing explosion of knowledge. The explosion continues. Today not even the most highly trained scientist, the most scholarly historian, or the most competent engineer can hope to have more than a general overview of what is known. Only specialists understand most of the new discoveries in science, and even the specialists have trouble keeping up.

This problem isn't new. In 1945, Vannevar Bush wrote an essay for Atlantic Monthly about out the problem of too much knowledge. He wrote,

> *There is a growing mountain of research. But there is increased evidence that we are being bogged down today as specialization extends. The investigator is*

[1]

> *staggered by the findings and conclusions of thousands of other workers - conclusions that he cannot find time to grasp, much less to remember, as they appear. Yet specialization becomes increasingly necessary for progress, and the effort to bridge between disciplines is correspondingly superficial.*

Bush's imagined solution to this problem was something he called Memex. Memex was envisioned as a system for manipulating and annotating microfilm (computers were just then being invented). The system would contain a vast library of scholarly text that could be indexed by associations and personalized to the user Although Memex was never built, the World Wide Web, which burst onto the scene half a century later, is a rough approximation of it.

As useful as the Web is, it still falls far short of Alexander's tutor or even Bush's Memex. For one thing, the Web knows very little about you (except maybe your credit card number). It has no model of how you learn, or what you do and do not know--or, for that matter, what it does and does not know. The information in the Web is disorganized, inconsistent, and often incorrect. Yet for all its faults, the Web is good enough to give us a hint of what is possible. Its most important attribute is that it is accessible, not only to those who would like to refer to it but also to those who would like to extend it. As any member of the computer generation will explain to you, it is changing the way we learn.

A New Tool for Learning

Let's put aside the World Wide Web for a moment to consider what kind of automated tutor could be created using today's best technology. First, imagine that this tutor program can get to know you over a long period of time. Like a good teacher, it knows what you already understand

and what you are ready to learn. It also knows what types of explanations are most meaningful to you. It knows your learning style: whether you prefer pictures or stories, examples or abstractions. Imagine that this tutor has access to a database containing all the world's knowledge. This database is organized according to concepts and ways of understanding them. It contains specific knowledge about how the concepts relate, who believes them and why, and what they are useful for. I will call this database the Knowledge web, to distinguish it from the database of linked documents that is the World Wide Web.

For example, one topic in the Knowledge web might be Kepler's third law (that the square of a planet's orbital period is proportional to the cube of its distance from the sun). This concept would be connected to examples and demonstrations of the law, experiments showing that it is true, graphical and mathematical descriptions, stories about the history of its discovery, and explanations of the law in terms of other concepts. For instance, there might be a mathematical explanation of the law in terms of angular momentum, using calculus. Such an explanation might be perfect for a calculus-loving student who is familiar with angular momentum. Another student might prefer a picture or an interactive simulation. The database would contain information, presumably learned from experience, about which explanations would work well for which student. It would contain representations of many successful paths to understanding Kepler's law.

Given such a database, it is well within the range of current technology to write a program that acts as a tutor by selecting and presenting the appropriate explanations from the database. The automated tutor would not need to create the explanations themselves-- human teachers would create the explanations in the Knowledge web, and the paths that connect them. The program would merely find the appropriate paths between what a student already knows

and what he or she needs to learn. Along the way, the automatic tutor would quiz the student and respond to questions, much as a human tutor does. In the process, it would improve both its model of the student and the information in the database about the success of the explanations.

For example, imagine yourself in the position of an engineer who is designing a critical component and wants to learn something about fault-tolerant design. This is a fairly specialized topic, and most engineers are not familiar with it; a standard engineering education treats the topic superficially, if at all. Fault-tolerant design is an area normally left to specialists. Unless you happen to have taken a specialized course, you are faced with a few unsatisfactory alternatives. You can call in a specialist as a consultant, but if you don't know much about the field it's difficult to know what kind of specialist you need, or if the time and expense are worth the trouble. You could try reading a textbook on fault-tolerant design, but such a text would probably assume a knowledge you may have forgotten or may never have known. Besides, a textbook is likely to be out of date, so you will also have to find the relevant journals to read about recent developments. If you find them, they will almost certainly be written for specialists and will be difficult for you to read and understand. Given these unsatisfactory choices, you will probably just give up. You will go ahead and design the module without benefit of the proper knowledge, and hope for the best.

Learning with Aristotle

Now let’s assume that you have access to the automatic tutor. Let's call it Aristotle. Aristotle would begin by asking you how much time you're willing to devote to this project and the level of detail you want. Then Aristotle would show you a map of what you need to learn. The tutor program does this by comparing what you know to what needs to be

known to design fault-tolerant modules. It knows what needs to be known because this is a common problem faced by many engineers, and knowledgeable teachers have identified the key concepts many times. Aristotle knows what you know because it has worked with you for a long time. There may be some things you're familiar with that Aristotle doesn't know you know, but you can point these things out to Aristotle when it shows you the learning plan. Aristotle might take your word for what you know, but it is more likely to quiz you about some of the key concepts, just to make sure.

Aristotle plans its lessons by finding chains of explanations that connect the concepts you need to learn to what you already know. It chooses the explanatory paths that match your favorite style of learning, including enough side paths, interesting examples, and related curiosities to match your level of interest. Whenever possible, Aristotle follows the paths laid down by great teachers in the Knowledge web. Aristotle probably also has a model of how you want to be paced: when you have learned enough for one day, when it needs to throw in an interesting side story, etc. Along the way, Aristotle will not only explain things to you but will also ask you questions--both to make you think and to verify for itself that concepts are being learned successfully. When an explanation doesn't work, Aristotle tries another approach, and of course you can always ask questions, request examples, and give Aristotle explicit feedback on how it’s doing. Aristotle then uses all these forms of feedback to adjust the lesson, and in the process it learns more about you.

The process of teaching helps Aristotle learn to be a better teacher. If an explanation doesn't work, and consistently raises a particular type of question, then Aristotle records this information in the Knowledge web, where it can be used in planning the paths of other students. The feedback eventually makes its way back to the knowledge web's

human authors, so that they can use it to improve their explanations.

Like any good tutor, Aristotle allows you to get sidetracked from a lesson plan and follow your interests. If you find a particular example compelling, you may want to know more about it. If some concept you have just learned allow you to appreciate an elegant explanation of some fact you already know, Aristotle may point it out. If you are close to understanding something of critical importance or something that would be of particular interest to you, Aristotle may decide to show it to you even though it is not strictly necessary as part of the lesson. Of course, as Aristotle gets to know you, it will know how much you like this sort of distraction.

Once you have learned the material, and Aristotle has verified that you have learned it, the program will update its database to indicate that you have recently learned it. As you learn more and more, it will continue to connect your recently acquired knowledge to the new concepts you are learning, until you have fully integrated them. Because Aristotle knows which subjects you are and have been interested in, it can consolidate your learning by finding connections that tie these subjects together.

For example, there's a short film of Dr. Richard Feynman explaining a principle of quantum mechanics called Bell's inequality. Most people have little interest in quantum mechanics and no interest at all in understanding Bell's inequality. Most quantum physicists already understand Bell's inequality and would learn little from Feynman's explanation. On the other hand, if you are a student who is just learning quantum mechanics, who has just mastered the necessary prerequisites, Feynman's explanation can be exciting, startling, and enlightening. It not only can explain something new but can also help you make sense of what

you have recently learned. The trick is showing the film clip at just the right time. Aristotle can do that.

I used an engineering example to describe Aristotle because most engineering knowledge is a straightforward, factual type of knowledge. Similar techniques would work for learning other subjects - history, mathematics, or the kind of technical information that might normally be conveyed in a training course or a technical manual. Of course, there are types of useful knowledge that a program like Aristotle would not be suited to: Aristotle would not be of much use in learning how to ride a bicycle or tell a joke. It would not replace hands-on experience, nor would it replace the enthusiasm and wisdom of a great teacher. What Aristotle would do is help you gain mastery of factual knowledge--exactly the kind of knowledge that is overwhelming us.

How the Knowledge Web Changes Education

In his book Diamond Age, the science fiction writer Neil Stephenson describes an automatic tutor called The Primer that grows up with a child. Stephenson's Primer does everything described above and more. It becomes a friend and playmate to the heroine of the novel, and guides not only her intellectual but also her emotional development. Such a Primer is beyond the capabilities of current technology, but even a program as limited as Aristotle would be a step in that direction. Teachers and students both understand that a school is no longer able to "pre-load" its students with the knowledge that they will need for life. Instead, a good school teaches the basics--reading, arithmetic, social skills--and introduces students to subjects that they can learn more about. It gives them an overview of knowledge as a starting point for further learning. A good school education also gives students the skills to acquire knowledge as they need it.

Teachers know that individual attention helps a child learn and they would like like give their students more of it. Even a young child has special interests, special topics that he or she would like to know more about. A good teacher learns to recognize these individual interests and tries to nurture them, but this takes a lot of time. A program like Aristotle will give teachers a tool to help children follow their passions. It will also enable teachers to evaluate a child's progress and to provide individualized instruction in areas in which the child has gaps. A computer program like Aristotle cannot replace most of what goes on in school, but it can complement what goes on there. It can free teachers from the routine job of broadcasting information and give them more time for individual attention to their students.

A system like Aristotle also enables teachers by giving them a way to publish. Any good teacher knows how to teach certain topics especially well, but there are few easy ways for them to share that information effectively with others. A teacher can write a textbook, or develop a curriculum, but each of those efforts is a major undertaking. There is no simple way for a teacher to publish an isolated idea about how to explain something. If a system like Aristotle existed, then with the proper authoring tools a teacher could publish a single explanation--an effort comparable to creating a Web page. In fact, existing Web pages are a good source of initial content for the knowledge web. As Marshall McLuhan said, "The content of the new medium is the old medium." The initial content of the knowledge web will be the old curriculum materials, textbooks, and explanatory pages that are already on the World Wide Web. The existing materials already contain most of the examples, problems, illustrations, and lesson plans that knowledge web will need.

As students gain access to the best explanations from the best teachers on a given subject, their own teachers will be able to take on the role of coaches and mentors. Freed from the burden of presenting the same information over and

over, the teachers will be able to give greater individual attention to their students.

A Better Infrastructure for Publishing

The shared knowledge web will be a collaborative creation in much the same sense as the World Wide Web, but it can include the mechanisms for credit assignment, usage tracking, and annotation that the Web lacks. For instance, the knowledge web will be able to give teachers and authors recognition or even compensation for the use of their materials. Teachers and learner willbe able to add annotation and links to explanation--connecting them to other content, suggesting improvements, or rating their accuracy, usefulness, and appropriateness for children of various ages. For instance, with such a system, it would be possible for the learnier to accept only knowledge that had been certified as correct by an authority such as Encyclopedia Britannica or the National Academy of Sciences.

All of this raises the possibility of a different kind of economic underpinning for the knowledge web, one that is not possible on the document Web of today. The support infrasture for payments would allow different parts of the knowledge web to operate in different ways. For instance, public funding might pay for the creation of curriculum materials for elementary school teachers and students, but specialized technical training could be offered on a fee or subscription basis. Companies could pay for encoding the knowledge necessary to train their employees and customers, consultants would be able to publish explanations as advertising for their services, and enthusiasts would offer their wisdom for free. Students could subscribe not only to particular areas of knowledge but to particular types of annotations, such as commentary or seals of approval. Schools and universities could charge for for interaction with teachers and certification of the

student's knowledge. The system could also become the ultimate hiring tool, since employers could map areas of knowledge that they need in prospective employees.

What Makes the Knowledge Web Different?

One way to think about the knowledge web is to compare it with other publishing systems that support teaching and learning. These include the World Wide Web, Internet news groups, traditional textbooks, and refereed journals. The knowledge web takes important ideas from these systems. These ideas include peer-to-peer publishing, vetting and peer review, linking and annotation, mechanisms for paying authors, and guided learning. Each of these existing media demonstrates the success of one or more of these ideas. They are all incorporated into the knowledge web.

Peer-to-Peer Teaching

One of the reasons why Internet news groups and the World Wide Web have enjoyed such runaway success is that they allow people to communicate with each other directly, without publishers as intermediaries. The great advantage of such a peer-to-peer publishing system is that anyone with something interesting to say has an easy way to say it. The Internet has eliminated the publishing bottleneck and has created a flood of authorship. This basic human desire to share knowledge is what will drive the creation of the knowledge web. The task of recording the world's knowledge is so overwhelming that only peer-to-peer publishing can plausibly accomplish it. Yet the knowledge web is not only a record of knowledge but also a way of imparting it. The knowledge web will do for teaching what the World Wide Web did for publication. Peer-to-peer teaching will allow literally millions of people to help each other learn.

Vetting and Peer Review

One of the downsides of peer-to-peer publishing is quality control. Publishers of textbooks and journals do more than market and distribute; they also edit and select. In the case of peer-reviewed journals, some of the burden of quality control is shifted to the reviewers, but it is still coordinated by the publisher. On the World Wide Web, there is no commonly accepted system of rating and peer review, nor is there a mechanism to support one. The result is chaos. The information you find by searching on the World Wide Web is often irrelevant, badly presented, or just plain wrong. It is difficult to screen out obscene material and propaganda. It is almost impossible to sort the wheat from the chaff.

The knowledge web addresses this problem by supporting an infrastructure for peer review and third-party certification. It supports mechanisms for the labeling, rating, and categorization of material, both by the author and by third parties. The browsing tools will allow information to be filtered, sorted, and labeled according to these annotations. In addition, user feedback tools will be built into the browsing software to help identify material that is particularly good, bad, or controversial.

Linking and Annotation

Anyone who has used the World Wide Web understands the importance of linking. In principle, articles in conventional journals also support a kind of linking, in the form of footnotes and references, but these kinds of links are far less convenient to use than the convenient “click throughs” of hypertext. Even the simple message-thread linking of threaded news groups helps make the information more usable. Students accustomed to hypertext find the linear arrangement of textbooks and articles confining and inconvenient. In this respect, the Web is clearly better. The knowledge web will allow an even more generalized form of linking than the World Wide Web. In the knowledge web,

not only the author but also third parties can create links, comments, and annotations.

Ways to Pay Authors

An advantage of textbooks and journals over the World Wide Web is that they support a mechanism for paying the author. The World Wide Web has demonstrated that many authors are willing to publish information without payment, but it does not give them any convenient option to do otherwise. The knowledge web will provide that option by supporting various payment mechanisms, including subscription, pay per play, fee for certification, and usage-based royalties. It will also support and encourage the production of free content.

Much of the content on the knowledge web will probably be free, but there are a number of other economic models that can coexist with this. One of the most obvious is the paid course, in which a student pays tuition for a cluster of services, including access to teachers, curriculum materials, and interaction with other students, and some form of certification at the end. With the knowledge web, many institutions may choose to offer the curriculum material for free--as a form of advertising--and charge for the other services, especially the certification. The knowledge web would also help solve one of the online course providers' greatest problems, which is marketing. The knowledge web would help direct students to the courses that meet their need.

Another model that may work well is a micropayment system, in which a student pays a fixed subscription fee for access to a wide range of information. Usage statistics would serve as a means to allocate the income among the various authors. This system has the advantage of rewarding authors for usefulness without penalizing students for use. The students' fees would be independent of the amount of use they make of the system. The ASCAP music

royalty system and university payments for student access to Encyclopedia Britannica are examples of how such a system might work.

Guided Learning

Part of the information in a textbook is about the subject matter, but part of it is also about how to learn the subject matter. A good textbook is full of information on the plan of attack, strategies for learning, practice problems, and suggestions for further study. Part of what makes up a good curriculum is not just the material but the plan for moving through it. A textbook often has an accompanying "teacher guide" that contains more of this type of information. For example, the guide may note that if the student makes a certain type of error, the student is probably missing a particular point and needs it to be re-explained. Some of the best computer-aided instructional material also encodes such information. The knowledge web provides an easy mechanism for a teacher to include this type of information.

Table of Affordances

	The Web	News Groups	Text Books	Journals
Peer-to-Peer publishing	Yes	Yes	No	Limited
Supports linking	Yes	Limited	No	Limited
Ability to annotate	No	Yes	No	No
Vetting and certification	No	Limited	Yes	Yes
Supports payment model	No	No	Yes	Yes
Supports guided learning	Limited	No	Yes	No

Achieving Critical Mass

Once the knowledge web achieves critical mass, it is easy to see how it can sustain itself. The real question is how it gets started. Presumably, first users will be adults, and the first adopters will be industry and government. Industry spends $60 billion a year on training for employees. It spends even more on customer support, product liability, poor design, and other costs that could be mitigated by better training. This is the first market for the knowledge web. The second market is probably the military, which is the single largest educator of adults. The third market is continuing education for individuals.

Adults are an easier initial market than children for several reasons. Adults are usually learning because they want to or because they need to. They are already motivated to learn. Adults often need specific knowledge for a specific reason. A recent survey of the Amercan workplace showed that 80 percent of workers feel that additional education is important for them to be successful at their job. Adults are much more likely to treat time as a limited resource. They want to learn efficiently. Also, the mechanisms that pay for adult education are often more rational and less politically charged than the mechanisms that pay for the education of children. For-profit companies are likely to adapt quickly to a more efficient process.

Eventually traditional educational institutions will use this new system of learning. It will be used first by colleges and trade schools; then it will make its way into secondary and elementary education. It is important to emphasize that computers will not replace the teachers; rather, they will give teachers a new tool. Instead of spending most of their time broadcasting information to a group, teachers will have more time to help students integrate their knowledge through discussion and through individualized interaction.

There are three technical components necessary for such a system to work: the tutor (browser), the authoring tool, and the knowledge web itself. This last component is the most difficult to create, but fortunately it does not have to be built all at once. Even a small part of it would be useful. Presumably the knowledge web will get its start in a few narrow areas, probably determined by available funding. For instance, it is easy to imagine a scenario in which a part od the knowledge web is initially funded by a supplier to explain the use of its products. Imagine, for example, that Cisco publishes the knowledge of how to configure and maintain its routers in this form. Another scenario is that the government sponsors the creation of the system for a specific application, such as continuing education for schoolteachers or job training for factory workers. A company might pay for the development of training programs for its workers and customers. A foundation might sponsor an initial effort as a way to have an impact on education.

Summary: An Idea Whose Time Has Come

It seems almost inevitable that such a system will eventually be developed, but why is now the time to undertake such a project? After all, people have been dreaming of such a project for half a century, yet Memex, Xanadu, Plato, Wais, and numerous other schemes failed to achieve critical mass. Why should this one succeed? For one thing, the infrastructure is now ready for it. The knowledge web requires widespread access to network-connected computers capable of handling graphics, audio, and video. This did not exist until recently. But it is not just the technology that is ready for this idea-- people are ready for it. Email, the Web, and video games have all whetted their appetites. The younger generation is more than ready. They expect something better than listening to lessons in class-sized groups.

At the same time that a solution is becoming possible, the problem is reaching a crisis point: the amount ofknowledge is becoming overwhelming, and the need for it is increasing. There is a widespread conviction that something radical needs to be done about education--both the education of children and the continuing education of adults. The world is becoming so complicated that schools are no longer able to teach students what they need to know, but industry is not equipped to deal with the problem either. Something needs to change.

With the knowledge web, humanity's accumulated store of information will become more accessible, more manageable, and more useful. Anyone who wants to learn will be able to find the best and the most meaningful explanations of what they what to know. Anyone with something to teach will have a way to reach those who what to learn. Teachers will move beyond their present role as dispensers of information and become guides, mentors, facilitators, and authors. The knowledge web will make us all smarter. The knowledge web is an idea whose time has come.