

The Social Life of Learning: How can Continuing Education be Reconfigured in the Future?

John Seely Brown

We are facing some very exciting times. In less than a decade we have witnessed the rise and fall of the internet economy, the creation of the knowledge age, and the transformation of the entire “learningscape,” ranging from universities to the corporate world. But as practitioners in education, especially continuing education, we are more than witnesses, we can be key participants in designing environments that enable us to distinguish what is important from what is incidental, the effective from the merely efficient, and most of all, knowledge from information.

TECHNOLOGICAL DRIVERS

Let me begin by establishing a common framework that shows the dynamics that are driving us. The media dwell only on Moore’s Law, which states that computing power doubles every 18 months. Moore’s Law is important and will be in effect for another two decades, but it is the least spectacular and slowest law at work. Every nine months—twice the speed of Moore’s Law—our ability to increase the bandwidth of optic fibers and optical amplifiers doubles, according to a fiber law. Multiply that by the ability to store information, which doubles every year, and the result is more useful information generated, flowing, and accessible.

Copyright 2002, John Seely Brown. John Seely Brown, Director of the Xerox Palo Alto Research Center from 1990–2000, is a member of the National Academy of Education, a fellow of the American Association of Artificial Intelligence, and a trustee of the MacArthur Foundation and of Brown University.

Coupled with all of that is the law that is almost never talked about. It's not Metcalfe's Law, which says the power of the Net grows as the square of the number of people. Bob Metcalfe created that law to describe the number of person-to-person interactions enabled by the Net. But the power of the Internet is also described by an additional law, the law of community, or perhaps the law of content, because it is community that creates most of our content today. It says that when you completely eliminate distance, the number of possible virtual communities that can emerge out of n people is not n^2 , it's 2^n , an astronomically larger number.

Looking at the dynamics of these laws, you realize that the forces at work are not just exponential, they're hyperexponential because these are not independent laws, they interact with each other. And because they rule our technological infrastructure, they enable powerful changes in the way we live, work, and learn. I want to dwell a moment on the law of storage to show how it can change something fundamental. My digital camera contains a one-gigabyte rotating disk, slightly larger than a quarter, that costs about \$300 today. This is made possible by a new technology called microelectrical/mechanical systems. IBM and several other companies have been looking into using nanotechnology to increase the storage capacity of a system of similar physical size up to a thousand gigabytes. Why do I bring this up?

These developments mean that we will be able to hold our entire world of information in our cellphones and personal digital assistants (PDAs): a thousand gigabytes can store everything we've read, every movie or television program we've ever seen.

Once you begin to take that in, you become aware of a major challenge: with so much information literally at our fingertips or in the palm of our hand, how do we use that information to act intelligently and wisely? We need to pay very close attention to the distinction between information and knowledge, because it gets to the heart of what it means to be literate in an age where we are flooded with information. And when one can get easy access to limitless information, the challenge for lifelong learners is to be able to find

*We need to pay
very close attention
to the distinction
between
information and
knowledge, because
it gets to the heart
of what it means
to be literate in
an age where we
are flooded with
information.*

what you want, to judge whether or not to believe in it, and to keep an open mind on evaluating information that could alter your own point of view.

The root of tunnel design is our tendency to focus on individuals and information. But it should lie in the balance between our focus on individuals/information and the larger notions of communities, communication, and context.

THE LEARNING CHALLENGE

Most of us have some hope that the cause of the problem—technology—will facilitate the solution. You might think that the road ahead would be straight, that I'm talking like a technological determinist who is saying that with this kind of technology, we can zoom down the information highway, powered by Moore's Law, disk law, bandwidth law, fiber law, community law. On the cover of Bill Gates's 1995 book, *The Road Ahead*, the road in the background is straight. And in fact, I can see why the publicists might have thought that the road ahead was straight; maybe even Bill did, at that point, because there's a lot missing on his road—not just cars—but people, cities, institutions—in a word, communities. When you factor them in, you find that the road ahead is not straight and that a very complex digital culture is unfolding, creating all kinds of unexpected twists and turns. We should never let our focus get away from the complex co-evolution between society and technology. Technological advances influence society, society influences technological advances—a spiral develops, sometimes going in some pretty crazy ways, but it's usually not a straight line nor a narrow path.

And I'm going to argue that one of the things we have to be extremely careful about, especially those of us who are technoenthusiasts, is tunnel vision. That perspective makes us think that the road ahead is

straight because we are not willing or able to take in the broader context, the larger periphery. Tunnel vision too often leads to tunnel design, which explains why our technologies don't necessarily increase the quality of life as we'd like them to. So as we march forward in designing new learning environments, we have to be careful that we don't engage in tunnel design.

The root of tunnel design is our tendency to focus on individuals and information. But it should lie in the balance between our focus on individu-

als/information and the larger notions of communities, communication, and context. In some sense they all come from the same root, because information taken out of context doesn't have much meaning.

Some time ago researchers at Palo Alto Research Center (PARC) ran a basic experiment that relates to tunnel vision and tunnel design. Working with the School of Architecture at MIT in Cambridge, our research center in California wanted to see if we could do remote critiques: the master in California would critique students working in a studio in Cambridge. We had wonderful technology—optical fiber gave us good video—but the first few attempts didn't work. And then one day we happened to have an open link when someone at MIT was moving the carousel supporting the remote camera system. As he moved from one student area to another, the architect at PARC said: "Stop! I want to look at the walls." The camera was pulled back and allowed to pan over the walls. For a few minutes, the PARC-based architect studied the sketches on the walls, and from those he began to understand a little bit more about the context for that particular project. The ability to bring in more of the peripheral context to create the basis for interpreting the center made all the difference in the world, and suddenly the power of these critiques went up dramatically. So it's an interesting example of how the kind of tunnel vision that zeroed in on the student project without taking in the context dampened the master's ability to offer meaningful reviews to the student.

So what does this say about learning and the design of learning environments? A lot of our thinking about pedagogy in learning has a Cartesian foundation. Three hundred years ago or so, Descartes's philosophical stance was "I think, therefore I am." Within that framework, the more abstract the knowledge is, the better. Knowledge that is decontextualized and disembodied is all the more applicable to different situations. That seems like a very reasonable position to take. But in some ways it falls short—very short. It views knowledge as a kind of substance and leads to a pedagogy encouraging the optimal pouring of knowledge from a storage device into the head of a student. Now

*Knowledge that
is decontextualized
and disembodied
is all the more
applicable to
different situations.
That seems like
a very reasonable
position to take.
But in some ways
it falls short—
very short.*

I say “pedagogy” and “optimal” because this not a trivial problem, even if you bought into the Cartesian duality, because after all, students already know a lot. So the outstanding problem is getting a good impedance match between the pourer of the knowledge—the “pitcher” of knowledge—and the student’s mind. But in my opinion, it is not the right way to frame learning because it confuses information with knowledge. And note that just because we now have optical fiber, a fiber-to-the-head is not much better than a pitcher; in other words, maximizing our ability to pour information is not the same as creating learning.

THE SOCIAL CONSTRUCTION OF UNDERSTANDING

So let me propose another way of looking at learning, an approach that has riveted my attention and influenced a great deal of my thinking over the past decade. Instead of Descartes’s disembodied thought, the separation between mind and body, “I think, therefore I am,” I want to suggest that there is something much more fundamental at stake. This framework, coming more from Dewey and psychoanalytic theory than from Descartes, says “We participate, therefore we are”: we come into being through participation with others, and our own understanding is socially constructed. That means we have to start looking at how conversations can be evoked by content. This is not saying that all knowledge is relative—far from it—and it’s not the standard constructivist point of view according to which you construct everything yourself. What we’re really saying is that you often construct and augment your own understanding of the content socially through conversations. (Note how this shift of emphasis might have special importance for learning strategies for lifelong learners.)

I’m proposing a pedagogical approach that has to do with designing evocative knowledge objects and spaces that foster focused conversations that scaffold a student’s ability to construct his own understanding of what that object is about. This suggests that very often what makes a great simulation is not the absolute fidelity of the simulation, but how it encourages dialogue, how it leads students to play with it collaboratively, and how it functions as a boundary object for constructing their own understandings of it. Knowledge is information that has been internalized and integrated into our frameworks. To facilitate a student’s learning we need to design spaces that encourage and scaffold conversations that do that. This means considering not only the social and informational spaces, but also the physical space and how all three of these spaces combine to promote the kinds of conversations we are talking about.

At this point, let's introduce a cautionary note about knowledge. It has an explicit dimension that you read about in books and that you pour from one head to another, or one pitcher to another, and it has a tacit dimension. These are not different kinds of knowledge—these are different dimensions to what we mean by knowledge and knowing. Cognitive psychologist Jerome Bruner drew our attention to the difference between learning *about* and learning *to be*. Knowledge is like an iceberg, where ten percent is explicit and visible, and 90 percent tacit and invisible. Perhaps a more useful metaphor for knowledge—or the dimensions of knowledge—is that it is like a tree that has a trunk, branches, and leaves that you see. But for the tree to endure, it has to have a root system. And it is the roots that are deeply embedded in the earth that helps to keep that tree robust. Those roots are like the tacit structures, the practices—the professional practices, the social practices, and the institutional practices that underlie what you see—the visible and explicit. And it's all too easy to forget the importance of the things that lie beneath, until you try to pull up a tree and try to replant it. You can't do it very well. So the challenge we face in education is not just looking at how to add new leaves and new branches to the tree, but sometimes going back and thinking about the root structure, the tacit structures, the social practices underlying the explicit.

*Knowledge
is like an iceberg,
where ten percent
is explicit and
visible, and
90 percent
tacit and
invisible.*

DESIGNS FOR LEARNING

At this point, I'd like to describe experiments and examples of different learning environments, beginning with one of our most brilliant failures. Every week at our research center, we bring in an interesting outside speaker. We're all overloaded today, so we took advantage of the bandwidth inside our building to create a high quality web broadcast to the researchers, who then wouldn't have to leave their offices. This seemed like an efficient use of everyone's time, both speakers' and researchers'. Perhaps I was the first to stumble into some of the downsides of this, because after I invited these august speakers to come, they would walk into an auditorium, and there would be no one there, or maybe two or three people. You would feel a bit silly, but you could reassure your speaker (and yourself) that there was a meter on the wall measuring eyeballs, and a

thousand pairs of eyeballs were watching. Needless to say, that didn't work too well. Unlike a movie, lectures are, in theory, an ongoing conversation with the audience; the sense of engagement and involvement between the speaker and a

visible audience makes the event come alive—or not. But the separation between the researchers and the speaker kept the magic from unfolding. For me, this was the most telling example of the difference between efficiency and effectiveness. The setup was highly efficient, but not effective.

*A lecture in
an auditorium
or a classroom
is probably still
the norm, but
even assuming
the presence
of an audience,
how do you
capture it as
an intellectual
asset that can
be re-used in
other ways?*

A lecture in an auditorium or a classroom is probably still the norm, but even assuming the presence of an audience, how do you capture it as an intellectual asset that can be re-used in other ways? Videotaping a lecture and the social interaction is a start, but watching a videotape of a lecture is not an entirely stimulating experience and, as it turns out, not the most effective and efficient use of technology, especially at a place like the PARC. PARC is a high tech environment where everyone takes notes on laptops, and those notes can be time-stamped and indexed into the digitized video stream of a lecture. And then you can create a “soup” of annotations from the community itself to capture the context of what actually went on. For example, any time a question is asked, any time people laugh, any time people clap, everything can be marked automatically, because it's easy to build technology that can segment those events from the standard audio streams. In fact, you can even segment out changes in the speaker's emotional intonations, converting them into additional explicit annotations on this soup.

The audio stream can be sent out for quick transcription, and the text and slides indexed and inserted. Subsequently, when someone wants to verify reactions—laughter, excitement, questions—or review a particular point, it's a simple matter of finding a slide.

This is done with virtually no overhead and using naturally occurring signals coming from the community about this content. By captioning, wrapping up, and structuring to a very limited amount the natural context, we create an intellectual asset, transforming something that was never used into

something that is now useful for just-in-time learning and other purposes as well. For example, although almost no one ever really watches a video of a missed lecture, points often arise in the talk that one wants to explain or argue about. With this system, just pull up the slide, touch the point on the slide and up comes the talk, the questions from the audience, or their reactions to the point. Suddenly this medium jumps to life as does the context around the content point. New conversations are evoked, new learning starts to flow. It's a simple but powerful demonstration of how to turn a never-used artifact—the video of the event—into an evocative object around which various conversations get stimulated and anchored. And a rich ecology of comments can ensue.

Even without PARC's resources, something as seemingly uninspiring as a videotape can be made evocative. A while back, Jim Gibbons, who was dean of engineering at Stanford, was teaching a course in which there were several students from Hewlett-Packard. They were transferred as a group to northern California, but they wanted to continue. Since the course was being broadcast to other companies in Silicon Valley, it was easy to videotape the broadcast and send it out. Hardly a breakthrough idea, you might think. But Gibbons also understood intuitively at that time a little bit about the social construction of knowledge and learning. So he said to the group of five, "Choose someone to be the 'chauffeur' for a particular session. Play the tape, but about every two or three minutes, I want you to stop the tape and have a conversation about what you just saw. If there's any confusion, replay the tape. Replay it until you can figure out for yourselves what was happening." Now this is something that can never happen in a class, which is a time-based medium that can't be interrupted. But the remote group *could* interrupt it at any time, and *had* to interrupt it after two or three minutes to use it as an evocative object for a conversation.

This idea is simple, practical, immediately executable, and incredibly powerful. When these students came back, they did outstandingly well in the final

*. . . We already
know about the
power of study
groups to accelerate
learning following
workshop classes.
But more pertinent
is the ability to
disrupt a time-
based medium
like a lecture so
that students can
converse in a
focused way.*

exams. This has now been replicated in about 10,000 different learning situations, including prisons in California and schools in Alaska. Enough data are coming in to show that students who learn this way come out about a half a

*Learning that
lasts stems from
enculturation into
a practice. . . .*

*So how do you
build environments
that facilitate
our ability to
enculturate into
a practice?*

grade point better on average than those who are physically in a class. When this was being done, it had to be done within the format of study groups. And we already know about the power of study groups to accelerate learning following workshop classes. But more pertinent is the ability to disrupt a time-based medium like a lecture so that students can converse in a focused way. This technique turns a foible into a feature.

What I find interesting is this: If we can create the ability to foster conversations between each other watching this content, how might instant messaging (IM) play out in this same kind of a context? Might it be nice to have an IM conversation during the lecture as the lecture is going on; better yet virtually stop the lecture? There may be a bit of a contradiction here, but if you start to tear these things apart, you begin to see some naturally occurring technologies that might be used to pragmatically foster the social basis of this kind of learning.

Let me switch my focus slightly and talk about a great evocative learning environment: the architecture studio, a physical social space where all work in progress is made public. In the academy, we are accustomed to keeping a lot of our work in progress secret, until one day I suddenly surprise my colleague next door when he reads my article in *Nature*. In a studio, it's impossible to keep anything private because designed objects are out in the open. Moreover, the master's critiques of one student can be overheard by his neighbor. But the richness of what is overheard stems from the fact that the neighbor has had a chance to see the project grow, starting with the root structures that are no longer visible. The other students have monitored the progress that has been made, and have witnessed the unfolding process that leads to this current reified artifact. And that whole context provides an enriched learning experience, one that enables focused, efficient conversations about what is being discovered, what failed, and why. The second intriguing aspect is that what these students are doing is not so much learning about architecture, they're enculturating into a practice of architecture. Learn-

ing that lasts stems from enculturation into a practice. Those are not terms we usually use to think about learning environments, but we should. So how do you build environments that facilitate our ability to enculturate into a practice?

A simple example of a powerful learning environment that never gets talked about is open source and the creation of the Linux operating system. The open source environment has thousands of people engaged in writing code for the Linux operating system. True, the ones who write the core of the kernel are a very small number, but as you branch out there are thousands of people engaged in this. What is unusual about this community is they all write code that's meant to be read by other people. If you can't write code that can be read, you will not become a member of this community. For 30 years those of us in computer science lived in a culture in which we wrote code in the most obscure way, because if I could write code that nobody else could read, I obviously had to be brilliant, assuming the code ran. This game has now reversed, in the open source community; it works by social capital and the amount you contribute.

My purpose in talking about this is not to talk about Linux open source as a new kind of industrial phenomenon, but to do a slight twist and view the open source phenomenon as a new form of cognitive apprenticeship. New members learn by opening up code, reading it, tinkering with it, and perhaps expanding it or making it more robust and efficient. This is tinkering on the periphery and learning by being a legitimate peripheral participant. These folks are being enculturated to the aesthetics and best practices of those who created the kernel of Linux. Here is a naturally occurring phenomenon—a common effort to build an open system—that has become a learning community based on enculturation and cognitive apprenticeship.

LEARNING ON CAMPUS

The challenge of e-learning is to find ways to honor the fact that as much learning happens outside the classroom on campus as happens inside the classroom. The discussions that happen as you leave the classroom or over a beer or the

*The challenge
of e-learning is
to find ways
to honor the
fact that as
much learning
happens outside
the classroom
on campus as
happens inside
the classroom.*

late-night bull sessions help you to internalize information in a way that makes it actionable knowledge. If you think about the kinds of learning that actually happen in university education, they appear to break down into two kinds—

*This is what makes
campus life such a
profound learning
experience; not
the facts you learn,
but the interpretive
and analytic
practices you pick
up or at least
become aware of.
But picking these
practices up is as
much a part of the
social life of campus
as the purely
academic side.*

extensive versus intensive. Extensive learning involves exposure to multiple communities of scholarship/practice experiencing, for example, the different kinds of explanations and sensibilities that occur in philosophy, mathematics, physics, engineering, sociology, and even what constitutes an explanation, a problem, or a solution. It helps to be exposed to these different practices. But the last two years of undergraduate, and surely graduate, education have more to do with the intensive learning, which involves enculturation into a particular practice. More generally, the sense of acquiring the aesthetics of a practice as well as the practice itself is part of intensive learning.

On campus you get to rub shoulders both in and out of class with these multiple sensibilities, not always in pleasant encounters. Dr. Ruth Simmons, the new president of Brown University, says that campus life is meant to be intellectually quarrelsome: "I don't care what you claim as long as you can argue and support those claims; you can hold even the most unpopular points of view as long as you're willing to defend your point of view. Campus is a space—a safe space—to have intellectual quarrels." Think about what that really means in terms of being able to enculturate into practices and to experience different sensibilities. This is what makes campus life such a profound learning experience; not the facts you learn, but the interpretive and analytic practices you pick up or at least become aware of. But picking these practices up is as much a part of the social life of campus as the purely academic side.

STUDENTS WHO GROW UP DIGITAL

Let's consider today's digital student, who simultaneously does e-mail, instant messaging, chat room, cell phoning, ripping music, and maybe even watching

MTV. I bring this up because it's all too common to focus high tech on campus just on academic affairs. But on today's campus, students view this digital infrastructure more like the oxygen they breathe, affecting and enabling every aspect of how they live, work, interact, and learn. So the idea of focusing only on technology's impact just on the academic side of an institution is probably wrongheaded. The issue extends to thinking, for example, about the ways to install administrative systems so that they become part of the student's life as well, instead of setting up firewalls between academic and administrative uses of computing.

A NEW KIND OF DIGITAL DIVIDE

The digital divide that is rarely discussed is the one between today's digital student and yesterday's analog professor. We have an interesting situation today because there's a new vernacular developing, one that most faculty are unaware of: a digital vernacular coming out of digital interactive media. It has to do with the extension of the language of film with the interactivity of multiple media and computer games. When I go on campus today, I feel like I'm back in Italy in Galileo's time, where the professors are speaking Latin and the students are using the vernacular. And that kind of vernacular shift can't be overlooked.

This idea of a vernacular is more serious than you might think. I throw out a warning note, because I used to think video games were junk. I wouldn't even take hypertext seriously: I love narratives, I don't want to jump from place to place. I don't want to invent my own story. I want to read stories constructed by authors who have the will and skill to construct great narratives. So I approached this world with hostility until I came across the works of Stephanie Barish of the USC Annenberg Center and J. C. Herz (*Joystick Nation*). From them I began to understand that there are new genres developing in this medium that we ought to be thinking about, because these are the vernaculars of today's students. And let me comment on USC, where I now spend more and more of my time. Elizabeth Daley, the dean of the film school and also the Executive Director of the Annenberg Center, made an interesting observation to the provost about four years ago: "can you really

*When I go on
campus today,
I feel like
I'm back in Italy
in Galileo's time,
where the professors
are speaking Latin
and the students
are using the
vernacular.*

*Picasso does not
traffic in commas
and parentheses;
an image has
to get to the
very essence.
What Picasso
or Matisse can do
in five lines takes
me a hundred
pages of a book.*

argue that in turning out today's students it is any less important to understand Hitchcock than Hemingway? If you don't understand Hitchcock, you're as disabled as if you don't understand Hemingway." You can imagine the

furor this created. But from that starting point, USC launched an ambitious multimedia literacy project. It has nothing to do with learning the tools; it involves going to all higher level courses, selecting professors who want to experiment with multimedia, and working with them to create classroom settings and projects. Their students no longer write theses; they do a film or a multimedia project, either a time-based or nonlinear interactive system that they present publicly at the end of the class.

Let me describe a specific example: the first assignment in a women's literature course at Berkeley. During opening week, the first assignment is for every class member to pick the figure—a feminist, in this case—of whom they think most highly, choose four photographs that manifest something about what makes this feminist special to them, and write one paragraph under each photograph. Why? Picasso does not traffic in commas and parentheses; an image has to get to the very essence. What Picasso or Matisse can do in five lines takes me a hundred pages of a book. Picking a photograph and limiting the prose to one paragraph is a beautiful interplay between the

concrete and the abstract. You have to know exactly what you're going after. We in the academy have learned how to obfuscate much of what we do, whether we are willing to admit it or not. Or we add layers of parentheses on any kind of an idea. But as I do more and more work in film, I've discovered that when you do a storyboard, there are no commas, there are no parentheses; you have to get right to it in a storyboard. And in fact one of the exercises is for a professor to do a storyboard of the class.

There's another digital divide separating generations: instant messaging, which was mentioned earlier. Most people of a certain age probably don't do IM because we didn't grow up with it. But if the game we're talking about is amplifying your ability to socially construct your understanding, IM may have a significant role in learning environments. I've just come from a biology

course where all the students were together in a classroom having an oral discussion about different forms of evolution. At the same time, they were using IM to check out ideas with each other. Now that's potentially dangerous, because pretty soon all they're doing is passing notes about what a stupid comment so-and-so made. But the culture of this classroom was completely the opposite. The amount of sharing before somebody made a comment was high with this group. We don't yet know very much about how to structure the sociology in the class with IM, but I think that it is going to be a valuable tool to augment, for example, distributed study groups.

LESSONS FOR A DIGITAL WORLD

The Internet started out as just being a network of networks, but it has unquestionably become a new medium or new collection of media, and each medium has its own set of genres. It's the first medium that we have seen in a long time that has technological tools to support multiple forms of intelligence. Almost all the tools in the past have supported our ability to engage in text, engage in abstract mathematical reasoning. But now we're beginning to find tools that start to honor the multiple intelligences of our students.

But this is not just for students. For example, think of IM as a technology that supports and leverages things going on in the periphery. It creates an amazing sense of being connected to your buddy group. Even if you're not in close proximity or physically present, you're aware that they are there for you.

Instant messaging was thought of as a joke, and many corporations tried to keep it from happening, but it is now becoming a breakthrough tool for knowledge management inside the corporation. It is a way for cross-functional groups to stay in contact with each other if they remain in their core competency groups, their communities of practice. Or if they all assemble physically in one cross-functional war room to work, it enables them to stay connected to their communities of practice. This kind of technology lets you live in two spaces

Almost all the tools in the past have supported our ability to engage in text, engage in abstract mathematical reasoning. But now we're beginning to find tools that start to honor the multiple intelligences of our students.

simultaneously; it enables you to work together and stay connected to other people highly significant in some aspect of your life.

Another youth-oriented technology that we tend to pay very little attention to, particularly in the Western world, is massively multiplayer online games (MMO) like Sony's *EverQuest*. It is easy to dismiss these games until you see the numbers: in Korea, a game called *Lineage* has two million players in Seoul, where you can play it and other MMOs in 26,000 game parlors; hundreds of thousands of players are online playing *Lineage* at the same time. These games can also be played at home, and perhaps not surprisingly, MMOs have substantially eroded the watching of TV. You can laugh at these games—I did—but what makes them interesting is not the game itself, but the knowledge ecologies being formed around them. Participants are constantly conversing with each other, discovering tricks, building extensions to the worlds, inventing and implementing new plug-ins, training their “avatars” to have more skills, even selling the avatars they have trained on eBay. In effect, they are engaged in a rich social world that has many of the properties of the open source Linux world we discussed previously. So don't just think about the games themselves—the content—but about the knowledge ecologies developing around these games—the context.

Both the military and the corporate worlds are beginning to see that the ways these massive multiplayer games are evolving suggest a new design approach, if not a fundamentally different notion of design. The design challenge is figuring out how to structure the ecology around the game as an evocative experience. This is the design that is not top down, it's a question of building things that nurture and influence how these games develop. I've been surprised as I spend my life around architects—not computer architects, real architects—that the thinking that surrounds these games is very much like urban design: how do you start to lay out things that influence the way these games are developed? It's a fundamental change of thought from designing top down to influencing bottom up.

There is something very important happening here, and as I conclude, I want to suggest that our goal should be to think carefully about how we can let the virtual augment the physical, and not replace it. As we contemplate a new charter for reinventing the university for the twenty-first century, we are going to have a lot of failures, because none of us know how to navigate this new space flawlessly. We have to listen, experiment, and reflect. We are all faced with the issue of figuring out what can be done on versus off campus, using our resources to extend the reach of the university both in time and in

space. But we also have to use digital means to transform ourselves into a true learning institution. How can we learn from the community around us as well as help nurture it? How do we learn from our alumni as well as help provide them with lifelong education? We must learn to listen and learn from the outside world as much as the outside world will profit from us. The key message that I have learned and that led to our book, *The Social Life of Information*, is that the way forward is paradoxically not to just look ahead, but to look around at how we actually do learn with and from each other.

QUESTION AND ANSWER SESSION

Q: Paul Edelson, State University of New York-Stony Brook: Just an observation. To hold the world in the palm of your hand or to see the world in a grain of sand are really two different things. One is a much larger, inclusive, more beautiful, more complex world. I agree with you that this world is a provocative and powerful tool to understand the other. But this opens up windows that we never even dreamed of before.

A: Absolutely. I am also a strong believer that when we follow a problem to its root we can not only solve the problem but also uncover truths of nature.

Q: Peg Wherry, Weber State University: About six years ago, I was at the American Studies Association conference session on using the Web and hypermedia for teaching and creating assignments. The question came up: if you abandon linearity, what do you have left for an organizing principle? Many of us work on designing learning experiences. Could you address the question about the role of structure, or providing a structure or framework for the students as they work about constructing knowledge for themselves?

We are all faced with the issue of figuring out what can be done on versus off campus, using our resources to extend the reach of the university both in time and in space. But we also have to use digital means to transform ourselves into a true learning institution.

A: The design challenges we face now concern how we structure the ecologies. When you think ecology, as opposed to a mechanism, you think about the way you build a garden and nurture it. Those are deeper design

and structural issues. Let me give you an example. I once visited a high school where they were trying to engage students with in-depth reading of poetry, but their approach was oblique to that task. They asked the students to create a three-minute interpretive film of a poem under study and to do it in two days' time. An interpretive film is a little bit like MTV. It is a vernacular. It is not a narrative or Hollywood-type film, but it is about combining images, sound, and voice overlays into a film that communicates what the poem meant to them. It's an example of a way to construct an in-depth reading of a poem, and then having powerful emotive tools to express your own interpretation of it. Now that's a relatively simple task and one that is open-ended in a very powerful way.

I like to think about how we can take design as a powerful motif for all kinds of learning. In this case I never would have thought about a design exercise for how to read a poem. I just didn't have the imagination. This group did. So I think part of the challenge we all face here is to be entrepreneurs, but not entrepreneurial in the classic venture capital sense, but how to break out of the box and think about new kinds of powerful learning environments.

... Part of the challenge we all face here is to be entrepreneurs, but not entrepreneurial in the classic venture capital sense, but how to break out of the box and think about new kinds of powerful learning environments.

Q: *Patricia Book, Penn State:* From my own experience, I multiprocess all the time: I'm on audio conferences, I'm checking my e-mail, and I'm talking to my staff. But the truth is I'm not present in all those multiprocessing events. I'm not listening to much of what I'm multiprocessing. So could you talk a little bit about the issue of what it really means to be present as you observe the social and human behavior and the implications of that?

A: Unfortunately we don't have very good ethnographies of kids who grow up digital in these highly multiprocessing environments. I'm very sympa-

thetic to what you say because what you're saying is true of myself. I do a very fast context switching, like anybody running complex organizations does. You know what you're not paying attention to. Let me just add one cautionary note. I've been really struck with the tunnel vision we find in the design of information systems and user interfaces. Note how today we all experience information overload. Now the curious thing is that when I'm driving my high performance motorcycle, I am processing a thousand times more information than I ever do in front of my screen, but most of that information is being processed subconsciously, subliminally. Our brain attunes to huge amounts of information and continuously decides what we should attend to or focus on. We need to design systems that enable us to attune to them as opposed to always having to attend to them.

Q: Tom Phelan, Syracuse, New York: I'm on the corporate side, working in emergency planning, response to terrorism, and instances of mass destruction. There's a lot of adult learning going on in those fields. My question goes to a new tolerance that we experience in continuing education perhaps more than other places. Let me give you a quick example. I was doing a lecture at a police academy. My pager went off from my corporation. I was in Syracuse, my people were in Puerto Rico rebuilding an electrical system after a storm. And the page came from a wife having to get in touch with her husband working in Puerto Rico. I took a five-minute break and five minutes later she was talking to her husband because we had set things up to make that happen. So my question concerns tolerance. If our students are adult learners who have real jobs, who are connected to a pager or the Blackberry, are we as teachers of adult learners going to have the tolerance necessary to deal with those kinds of dual connections, interruptions, people sitting in class with a pager or some kind of device, students working at home who might have to leave the online program momentarily to handle a business call? Are we prepared as adult educators to tolerate the dual connection of our students, often while they're sitting right in front of us?

A: Great question. I certainly hope so, but it will not be easy for us.

Q: Unidentified member of the audience: Let me talk about tolerance of a different kind: How do we ensure that these technologies are used to support citizenship, tolerance, and the creation or maintenance of an open society?

*The essence
of being educated
in the twenty-first
century has a lot
to do with our
ability to learn
how to listen
and to hear
multiple points
of view.*

A: I would change the question slightly, because the notion of ensurance reflects, perhaps, the wrong point of view, namely a mechanistic one. The question is how do we increase the chances for things to evolve in a way that supports a liberal society? The essence of being educated in the twenty-first century has a lot to do with our ability to learn how to listen and to hear multiple points of view. Let me give you a simple example of something I find encouraging. It's probably more important than ever that our students spend time abroad, especially in a third world country, so they "marinate" in the practices of a substantially different culture. But students often have prerequisites that prevent them from being able to take off half a year or a year. What an obvious way to use distance learning: they can stay connected to their school for prerequisite courses, and yet experience the other culture. And that ability to absorb multiple cultures is fundamental in addressing the issue of tolerance, because we become aware that there are people with different points of view that are sensible and can be argued out. When you live in their culture, you understand their practices, not their rhetoric. This is when we'll begin to appreciate the complexity of the grain of sand that was mentioned by the first question.

Personally and professionally, my interest is in exploring how we use this technology to amplify our ability to listen and the ability of our students to listen and learn from other cultures.

Q: *Joseph DiGrigorio, Georgia Tech:* You made me think philosophically and more theoretically about what I do on a day-to-day basis, but as you were speaking, I was trying to see the practical applications for me, my job, my university, and how we get to that future you're describing. It occurred to me that the skill sets and the qualifications for the people who are going to be working with the students ten years from now are going to be significantly different based on what you said. And it occurred to me, if we were sitting here ten years from now and you were writing a job description for a faculty member to teach in this new environment, what would that job description look like?

A: Both students and teachers will need to feel comfortable swimming in a sea of information and be fluent at making good judgments about what information to believe in. We have institutional warrants behind much of the printed information we read. In a published book the publisher's institution stands behind the validity of the information in the book, more or less. So publishing is a warranting mechanism. That's nonexistent in most of the Internet. Most of what you'll find there, you wouldn't necessarily want to believe in very much. So now you're faced with a sea of questionable information. How do you start to feel comfortable in that sea? Swimming here becomes a kind of deliberation or triangulation. But how do you do that deliberation? How do you guarantee that you're bringing in enough different points of view to believe that you are not fooling yourself? And as we become more global, we need to be able to bring in and feel comfortable with increasingly diverse points of view.

In the last three or four years, I've moved from focusing on becoming a better learner to spending more time learning how to unlearn. My conceptual eyeglasses limit my ability to understand something that is deeply different. And so in the age of discontinuity we have to be very much aware of how our own lenses create a form of tunnel vision. We must learn new strategies to overcome the tendency to interpret the world with narrowly construed assumptions even if they worked for us in the past. I guess this is a twenty-first century twist to what Dewey always claimed about the critical need for an educated public if a democracy is going to be effective. 🐼

*. . . In the age
of discontinuity
we have to be very
much aware of how
our own lenses
create a form
of tunnel vision.
We must learn
new strategies
to overcome the
tendency to interpret
the world with
narrowly construed
assumptions even
if they worked
for us in the past.*

*This paper is a modified version of a keynote address
delivered at the 87th Annual UCEA Conference on
April 19, 2002, Toronto, Ontario.*