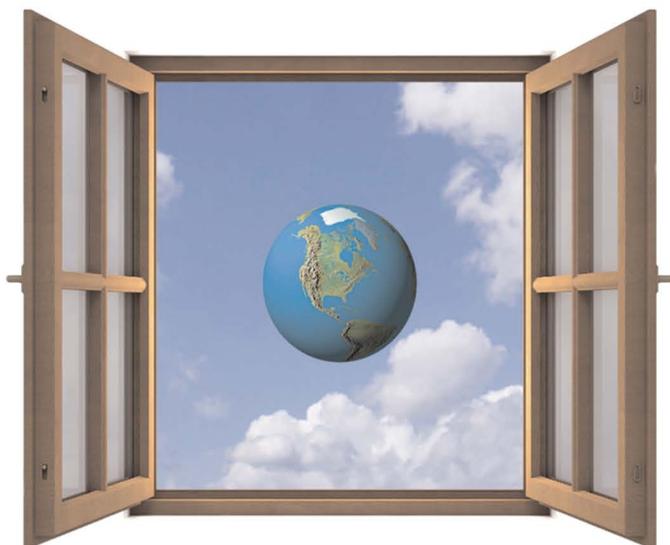# OPENING UP EDUCATION

The Collective Advancement
of Education through Open Technology,
Open Content, and Open Knowledge

edited by

Toru Iiyoshi and M. S. Vijay Kumar

foreword by John Seely Brown

Opening Up Education

Opening Up Education

The Collective Advancement of Education through Open Technology, Open Content, and Open Knowledge

Edited by
Toru Iiyoshi and M. S. Vijay Kumar

Foreword by John Seely Brown

creative commons

Attribution-Noncommercial-No Derivative Works 3.0 United States

You are free:

to Share — to copy, distribute, display, and perform the work

Under the following conditions:

Attribution. You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

Noncommercial. You may not use this work for commercial purposes.

No Derivative Works. You may not alter, transform, or build upon this work.

- For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to this web page.
- Any of the above conditions can be waived if you get permission from the copyright holder.
- Nothing in this license impairs or restricts the author's moral rights.

Disclaimer

To our wives and sons—
Hiromi, Rukmini, Ken, Suhas, and Taku
—whose support and encouragement makes this important work
possible.

Foreword: Creating a Culture of Learning

How could I say no? When approached by Toru and Vijay to participate in the workshop bringing the authors of this influential book together to discuss their preliminary chapter drafts, I was intrigued. After all, the challenges that we face in education today are daunting, and the editors had assembled quite an impressive company who were struggling to find ways through education to respond to those challenges. The world becomes more complex and interconnected at a lightning-fast pace, and almost every serious social issue requires an engaged public that is not only traditionally literate, but adept in a new, systemic literacy. This new literacy requires an understanding of different kinds of feedback systems, exponential processes, the unintended consequences inherent in evolving social systems, etcetera. In addition, the unrelenting velocity of change means that many of our skills have a shorter shelf life, suggesting that much of our learning will need to take place outside of traditional school and university environments. It is also unlikely that sufficient resources will be available to build enough new campuses to meet the growing demand for higher education, at least not the sort of campuses we have traditionally built for colleges and universities. Nor is it likely that current methods of teaching and learning will suffice to prepare students for the lives they will lead in the twenty-first century.

In response, we need to find a way to reconceptualize many twentieth-century education models, and at the same time reinforce learning outside of formal schooling. There may be powerful ways to blur the distinction between formal learning and informal where both turn on the social life of learning. Of course, we need to get back to basics—refining the most

effect modes of teaching, mentoring, and coaching, but we also should put some thought to informal teaching to support the inevitability and necessity of informal learning.

Fortunately, the networked age might be the "silver bullet" that will provide a way to both improve education and to set the stage for a necessary culture of learning. In the digital age, communities self-organize around the Internet, which has created a global "platform" that has vastly expanded access to all sorts of resources including formal and informal educational materials. The Internet has also fostered a new culture of sharing, one in which content is freely contributed and distributed with few restrictions.

The latest evolution of the Internet, the so-called Web 2.0 has blurred the line between producers and consumers of content and has shifted attention from access to information toward access to other people. New kinds of online resources—such as social networking sites, blogs, wikis, and virtual communities—have allowed people with common interests to meet, share ideas, and collaborate in innovative ways. Indeed, the Web 2.0 is creating a new kind of participatory medium that is ideal for supporting multiple modes of learning. Two of those include social learning, based on the premise that our understanding of content is socially constructed through conversations about that content and through grounded interactions around problems or actions. The focus is not so much on *what* we learn but on *how* we learn.

The second, perhaps even more significant, aspect of social learning, involves not only "learning about" the subject matter but also "learning to be" a full participant in the field. This involves acquiring the practices and norms of established practitioners in that field or acculturating into a community of practice. By entering into this community, you are required to assimilate the sensibilities and ways of seeing the world embodied within that community. And this is exactly what happens if you want to join an open source community with their key practices and expected contributions.

A culture of sharing, augmented with a culture of participation now made more feasible than ever with the Internet and Web 2.0 platforms will most likely start with the students themselves, as we see so vividly in both the complex massively multiplayer game worlds and the power

of study groups—whether conducted face-to-face or virtually. But the culture of sharing and participation must also involve content. MIT broke the ice when then President Chuck Vest launched the bold Open-CourseWare (OCW) project. Other universities quickly followed MIT's lead, and both the content and the means of accessing class materials and remixing and repurposing them for different audiences grew. But it is time that we move beyond merely considering content. We must begin to determine ways in which this content can encompass multiple kinds of instructional or learning activities within it. It is, after all, the combination of things we do with content that creates learning platforms.

Technology, of course, is key, and I want to dwell on only two aspects of how technology can now transform our learningscape: immersion and intelligent tutoring systems. Immersion is a concept that has received all too little attention in the learning literature. Consider, for example, how every one of us has learned the immensely complex system that is our own native language. We learn language through immersion and desire. Immersion comes from being surrounded by others talking and interacting with us and is furthered facilitated by our deep desire to interact, be understood and express our needs. We learn language fearlessly and constantly. Nearly everyone with whom we interact is a teacher for us—albeit an informal teacher, encouraging us to say new things, correcting us, extending our vocabulary, and so on. This simple form of immersion is fundamentally social in nature.

In today's high tech, graphically rich world we now have almost limitless opportunities to leverage immersion. We can now build simulation models of cites, historic events, atomic structures, biological and mechanical systems to name just a few. Our challenge becomes how to share the vast simulations and data bases that already exist and share them in a way that others can extend, remix and compose them in order to expand their reach and scope. I still dream of a virtual human system where I can explore any aspect of how our bodies function from organs to cells to membranes. There are promising signs, but as of yet we have no real framework for constructing and sharing modules of such a system. But if we can entertain the semantic web, perhaps we could entertain a vast and recursively interconnected web of simulations. No one group can build it all, but many could contribute, including students themselves.

Although richly visual, immersive three dimensional simulations will help the *born digital* student master complex topics, that will not be enough. We need to augment these systems (as well as the more passive course material) with computer-based intelligent tutors. Intelligent tutoring systems have had a long history stretching back to the 1970s where, in fact, I first became engaged in the issues of education and training. Some of our systems were pretty impressive as long as you overlooked the fact that they required a million dollar computer per student for our most advanced systems. But now we have machines 10,000 times more powerful, and in the last few years we have experienced the commoditizing of computing where vast amounts of computer power can be rendered cheaply through cluster or utility computing or even more recently, cloud computing. This means that our past dreams for building intelligent tutoring systems that could afford open-ended learning under the skillful eye of a tutor/coach/mentor is now becoming realistic. Indeed, the prior work of Carnegie Mellon University (CMU) and now their Open Learning Initiative (OLI), wonderfully described in this book, have reignited interest in this direction by demonstrating the power and utility of such systems.

The third category I want to dwell on is knowledge. For decades we have worked to create better theories of learning and successful models of teaching. We collect small fragments of data and struggle to capture context from which this data was extracted, but it is a slow process. Context is sufficiently nuanced that complete characterizations of it are extremely difficult. As a result, educational experiments are seldom definitive, and best practices are, at best, rendered in snapshots for others to interpret—see the description of the KEEP Toolkit in Toru Iiyoshi's and Cheryl R. Richardson's chapter in section III for the best and most evocative example of this. But we can do more.

If there was one thing I took away from the Carnegie meeting, it was that no one pedagogical or technical approach is the answer to ensuring that students are engaged and prepared. We need to be catholic in our point of view. There are thousands of colleges and universities worldwide, as well as many other institutions of learning, including training centers and technical schools. In addition, there are tens of thousands of institutions that support "informal" learning. And as the section

headings in this book illustrate, we need to think about how technology, content, and knowledge about learning and teaching can be creatively combined to enhance education and ignite students' passions, imaginations and desires to participate in constant learning about (and sense making of) the world around us. And we need to collect shared, distributed, reflective practicums in which experiences are vetted, clustered, commented on, and tried out in new contexts, as several of the authors here posit.

Again, in the spirit of sharing and participation, how might we start to instrument our learningscapes and to start collecting massive amounts of data about what is working and what is not and why. Consider OCW. There may well be millions of students using this material, but we need to ask what we're learning from them. Are there particular paragraphs or problems (in a problem set) that are routinely misinterpreted? How are test questions being misinterpreted? What sequence of materials appears to be working best? Are any systematic error patterns showing up? Can we model these errors? These queries barely scratch the surface of the information we need to collect, but readers who are particularly interested in these questions will be especially informed by Diana Laurillard's chapter. And of course, just as Amazon does in their Web site, experiments can be designed and continuously run and the data collected to see what works best for a class of users. When enough micro experiments are run and the data is continuously collected and analyzed automatically it can lead to some surprising breakthroughs.

But if we do all the above along with all the ideas expressed in this book—will it be enough to meet the educational challenges we currently face? I would suggest we need to extend our thinking around open education to include more of a Learning 2.0 (based on Web 2.0) perspective. There are two primary reasons why this might be true. The first turns on a question first posed to me by John King, associate provost of University of Michigan. He asked how many students I thought that University of Michigan (Ann Arbor) taught each year. I knew that the university had approximately 40,000 students give or take a few thousand, so that was my answer. He responded that 250,000 was closer to the actuality, adding that I had gotten UM's enrollment right. What you forget, he said, is that each year the incoming students bring their social

networks with them. These networks reach back into UM students' communities and schools. Using the social software and social network tools of SMS, IM, Facebook, MySpace, they extend the discussions, debates, bull sessions and study groups that naturally arise on campus to encompass this broader constituency, thus amplifying the effect the university is having across the state.

Now I doubt that there is hard evidence regarding the extent this is currently happening, but nevertheless it does draw attention to the broader learning milieu or learning landscape we need to consider, and it brings attention to extended forms of participation offered by the Internet. These extended forms start to merge tools for doing research with tools for learning—a boundary that needs to be blurred ever more. As a simple example, consider how many students pick up the practice of writing software by joining an open source community of practice such as Linux, Apache, etcetera. There may be small groups on a campus, but generally these communities of practice are highly distributed. Joining one of these communities entails first becoming a legitimate peripheral participant working on a small project improving or extending some piece of code and slowly building up a reputation from the work you do to moving into more central tasks and challenges in the community. Participants are learning new techniques about software practice from watching the work of their peers, defending their own work and participating in community discussions about emergent problems. This peer-based learning process is about *learning-to-be* a practitioner rather than just *learning about* software. Today's students don't want to spend years learning about something before they start to learn-to-be a practitioner in that knowledge domain. Of course, this happens to some extent on today's campuses in the form of practicums and laboratory exercises, but these are usually labor intensive in terms of the instructor's time and effort. They are also labor intensive for the student. But time spent on learning is a funny commodity. If the student is passionately engaged in acquiring the practice, then time seems to disappear. Passion is the key, which brings us to our last point.

Today the Web offers students incredible opportunities to find and join niche communities, niches that make up much of what Chris Anderson in 2004 in *Wired* magazine termed the "Long Tail." Finding

and joining a niche community that ignites their passion sets the stage through productive inquiry and peer-based learning to acquire both the practice of and knowledge about a field. Nearly all the resources discussed in this book, along with the millions of niche amateur (from the Latin word *amateur* meaning lover of) communities could provide a powerful learningscape for life long learning that is grounded in the learning practices acquired on campus. Not only would this provide real leverage for open education, but it would be a major step toward creating a culture of learning for the twenty-first century.

John Seely Brown