

To appear in *Designing Virtual Communities in the Service of Learning* edited by Sasha Barab, Rob Kling, and James Gray. Cambridge University Press, forthcoming 2003. An earlier version of this article was published in the Proceedings of ICLS 2000, Ann Arbor, MI, June 2000 (36% acceptance rate).

Co-Evolution of Technological Design and Pedagogy in an Online Learning Community

Amy Bruckman

College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280

Email: asb@cc.gatech.edu

The Design Process

The design of any piece of technology intended for human use—whether for entertainment, work, or education— is ideally iterative and user-centered. Designers can not anticipate all the needs of users, but most begin with a prototype and revise it based on user feedback. This is even more true of online learning communities, where designers must understand the needs not just of individual users, but of groups of users and their complex inter-relationships, as facilitated by the technology. Designers begin with theory, prototype, test, and then revise. However, it is not just the technology that can be revised, but also the underlying theory. Technological design and pedagogy have the potential to co-evolve in this new medium.

In this chapter, I will describe in detail one example of this co-evolution: a new perspective on motivation in constructionist learning environments which evolved through

quantitative and qualitative observations of an online learning community¹ called MOOSE Crossing. These observations led to a significant design change to the environment (the addition of a system of “merit badges”), and this in turn led to further reflections on pedagogy.

Background: Constructionism in a Community Context

In *Mindstorms*, Seymour Papert has a vision of a "technological samba school." At samba schools in Brazil, a community of people of all ages gather together to prepare a presentation for Carnival. “Members of the school range in age from children to grandparents and in ability from novice to professional. But they dance together and as they dance everyone is learning and teaching as well as dancing. Even the stars are there to learn their difficult parts” (Papert, 1980). People go to samba schools not just to work on their presentations, but also to socialize and be with one another. Learning is spontaneous, self-motivated, and richly connected to popular culture. Papert imagines a kind of technological samba school where people of all ages gather together to work on creative projects using computers.

Papert calls his approach to learning “constructionism,” and views it as an extension of Jean Piaget’s “constructivism:

We understand “constructionism” as including, but going beyond, what Piaget would call “constructivism.” The word with the v expresses the theory that knowledge is built by the learner, not supplied by the teacher. The word with the n expresses the further idea that this happens especially felicitously when the learner is engaged in the construction of something external or at least shareable... a sand castle, a machine, a computer program, a book. This

¹ I use the word “community” here in the loosest possible sense: a group of people interacting in some fashion over an extended period of time. I concur with Kling et al. (this volume) that the value of “online community” is often over-hyped. One solution to this problem is, as Kling proposes, to be vigilant to only apply the word “community” to groups that meet appropriate criteria. A second solution is to choose to use the word “community” in a looser, non-value-laden way, separating the issue of value from the issue of word choice. I prefer this second solution.

leads us to a model of using a cycle of internalization of what is outside, then externalization of what is inside. (Papert, 1991)

Much research on constructionist learning is highly individualistic in nature, focusing on providing the individual with “tools to think with” and the freedom to pursue self-selected goals (Kafai & Resnick, 1996). Papert’s samba school metaphor goes beyond that individualistic focus, highlighting the social nature of learning (Newman, Griffin, & Cole, 1989) and the role that a community context can play in scaffolding learning (Rogoff, 1994).

MOOSE Crossing is a text-based virtual reality environment (or “MUD”) whose design was inspired by the idea of a technological samba school. In this online world, children create magical places and creatures that have behaviors. In the process, they practice creative writing and learn object-oriented programming. Past research on MOOSE Crossing has focused on the power of the Internet to create a supportive, community context for constructionist learning (Bruckman, 1997, 1998). In many experiments in constructionist learning done to date, exciting learning results are observed during trials; however, those gains are not sustained after the researchers go home. The researchers provide important support for both learners and teachers. Online communities have the potential to provide ongoing support for innovative pedagogy under realistic settings. Past research on MOOSE Crossing has found that the community provides:

- Role models,
- Situated, ubiquitous project models,
- Emotional support to overcome technophobia,
- Technical support, and
- An appreciative audience for completed work. (Bruckman, 1998)

“Situated, ubiquitous project models” require special explanation. A good example of such models is HTML coding on the World Wide Web. On the Web, you can often

view the HTML source for a document. When you want to know how to do something yourself, you can remember something you saw that uses that technique, and look at the HTML source code. As you are going about your daily business using the Web, you are developing an ever richer vocabulary of models you can refer back to. Most software products come with detailed instruction manuals and sometimes libraries of examples; however, it is a truism that users tend not to take the time to read them. You can not learn from an example you have not seen. On the Web, a user's knowledge of available examples increases continually in the course of normal activity.

Similar to the Web, every object on MOOSE Crossing is an example to learn from. A learner in a virtual place can look at every object present—objects, other characters, the room itself—and use those as models to learn from. Since the set of examples to learn from is the set of all things everyone in the community has ever made, those examples reflect the interests of the community. That set of examples is not static; it grows and evolves as the community grows. A single, centralized author of a set of examples could never reflect the community's interests and needs as well as the decentralized process of using everyone's creations as examples to learn from.

Note in particular that examples on MOOSE Crossing are created by specific people. A child using another's work as a model typically has the opportunity to meet the more experienced programmer and ask questions. As a result, role models, project models, technical support, and emotional support for learning are all linked and become mutually reinforcing sources of support. Internet technology has unique affordances to help make a constructionist approach to learning viable in realistic settings.

Children and adults began using MOOSE Crossing in October 1995. It continues to be an active site at the time of this writing, with participants connecting from both home, school, and after-school programs. An after-school program was held at the MIT Media Lab from that date until February 1997 for the purpose of helping us to better understand the environment in use and the learning that takes place there. Children participating in the program were observed closely, and interviewed regularly. In addition to those local users, the environment was made freely available on the Internet to anyone interested. Many new participants simply heard about the project online or from friends and participate from home. A few teachers have brought their classes online. Teachers hear about MOOSE Crossing by word of mouth, and through presentations about the environment at conferences like the National Educational Computing Conference (NECC). Note that participation from home or after-school programs is almost always voluntary, while from schools may or may not be. In some classes, MOOSE Crossing is offered as just one of many options during unstructured computer lab time. In other classes, students are assigned to participate. Where participation is mandatory, in some cases teachers assess student work, and in other cases they chose not to. To further complicate the picture, many students who participate from school also have access from home, and a significant number of students whose in-school participation is required chose to participate voluntarily from home as well. However, taking all these factors into account, required participation in MOOSE Crossing is unusual and forms only a small fraction of total activity.

Every command typed on MOOSE Crossing is recorded, with written informed consent from parents and assent from kids. The log files comprise 3.1 GB data as of the

time of this writing (December 2001), and provide an opportunity for quantitative analysis to complement the ongoing process of qualitative observation and interviewing.

Like many educational online communities, participation in MOOSE Crossing is primarily voluntary. A few kids participate from classes, but most do so from home or from after school programs. Even those who participate in school are often given total or near total freedom in to what to do while online. Voluntary participation has advantages but also raises significant challenges for learning.

Challenges to Voluntary Participation and Learning

The 1989 movie “Field of Dreams” popularized the phrase “If you build it, they will come.” In the early 1990s, this was an apt description for developing online communities. For example, in 1992, a few email messages was all that was necessary to attract over a thousand members and create a thriving community on MediaMOO, a MUD designed to be a professional community for media researchers (Bruckman & Resnick, 1995). By the mid-1990s, the situation had changed: an exponentially increasing number of sites competed for users’ time and attention. Participation in most sites is voluntary, and a wide variety of sites compete for users’ limited time.

Voluntary participation is the norm for entertainment forms, but is a little more unusual in the field of education. After all, school work is typically required. Progressive pedagogy, however, suggests that students should learn better when they do so in the context of self-selected, self-motivated activities. At the radical extreme, Ivan Illich argues that school “makes alien the very ability of the individual to act or create. In school the ability to act is no longer an individual matter but is turned over to the experts who grade, rank, and prescribe. Activity itself no longer belongs to the individual but to the

institution and its experts” (Illich, 1970). Taking this radical rethinking of school and power relations seriously, generations of educational researchers have experimented with new ways to foster learning on a purely voluntary, self-motivated basis. However, moving from an extreme authoritarian to extreme anti-authoritarian view doesn’t eliminate problems—it simply introduces different ones. Given the freedom to direct their own learning, what if some students chose to learn little? What if they learn a great deal in some areas but neglect others? For example, Friskolen 70 is an elementary through middle school in Denmark with absolutely no requirements for students. In this open environment, children tend to excel at topics related to history, literature, and culture, but often fail to learn much mathematics. Students typically have difficulty moving to a traditional school for high school because of gaps in their mathematical background (Falbel, 1989).

As early as the 1930s, John Dewey was aware of this tension and struggled to find a resolution:

What is the place and meaning of subject-matter and of organization *within* experience? How does subject-matter function? Is there anything inherent in experience which tends towards progressive organization of its contents? What results follow when the materials of experience are not progressively organized? A philosophy which proceeds on the basis of rejection, of sheer opposition, will neglect these questions. It will tend to suppose that because the old education was based on ready-made organization, therefore it suffices to reject the principle of organization *in toto*, instead of striving to discover what it means and how it is to be attained on the basis of experience. We might go through all the points of difference between the new and the old education and reach similar conclusions. When external control is rejected, the problem becomes that of finding the factors of control that are inherent within experience. When external authority is rejected, it does not follow that all authority should be rejected, but rather that there is need to search for a more effective source of authority. Because the older education imposed the knowledge, methods, and the rules of conduct of the mature person upon the young, it does not follow, except on the basis of the extreme *Either-Or* philosophy, that the knowledge and skill of the mature person has no

directive value for the experience of the immature. On the contrary, basing education upon personal experience may mean multiplied and more intimate contacts between the mature and the immature than ever existed in the traditional school, and consequently more, rather than less, guidance by others. The problem, then is: how these contacts can be established without violating the principle of learning through personal experience. The solution of this problem requires a well thought-out philosophy of the social factors that operate in the constitution of individual experience. (Dewey, 1938)

Many online learning communities are organized around the principle of voluntary participation in either whole or part. Consequently, they face the challenge of finding a middle ground between the “extreme either-or philosophies” Dewey describes. Designers of online learning communities face the challenge first of eliciting any form of participation, and second of creating a high likelihood that participation leads to learning gains.

Portfolio Scoring

How effective is voluntary participation on MOOSE Crossing in fostering learning? Based on our first four years of observation and analysis of MOOSE Crossing from 1995 to 1999, the system designers and administrators developed an impression that that achievement in the environment is quite uneven. In fall 1999, we set out to explore variations in achievement systematically.

To begin our analysis, we decided to use portfolio scoring (Baron & Wolf, 1996; Chi, 1997) to examine achievement levels systematically. We randomly selected 50 children from the 803 total MOOSE Crossing users (as of November 1999). All chosen kids were under the age of 18 during their time of participation, and all had logged into the system at least once. Of the 50 kids selected, we had 23 girls and 27 boys. Further information about the children and their level of involvement can be found in Table 1.

Degree of participation is a key factor to track, and in most research is typically measured by time on task. On MOOSE Crossing, as in many online environments, time on task is not equivalent to total connection time. Many users may be multitasking, or just leave themselves logged in and inactive for long periods of time. Therefore, in order to get a more accurate measure of time on task, we chose to count the number of commands issued by each child.

We performed a portfolio-based assessment of each participant's scripting ability. While children do significant creative writing as well as programming in this environment, the writing data is harder to interpret since children can learn to improve their writing through a variety of activities and sources of support. For this reason, we are focusing this analysis on their programming ability. Each child's portfolio contains all the scripts the child has written on MOOSE Crossing. Two independent judges reviewed and scored the portfolios. In cases where they did not agree, a third judge scored the student's work, and then compared the three scores and assigned a final score. The children were scored on their highest level of achievement using the following scale and criteria:

- 0: Wrote no scripts
- 1: Demonstrated understanding of basic input/output
- 2: Used variables and properties
- 3: Performed list manipulation and flow control
- 4: Demonstrated mastery of all aspects of the system

	Minimum	Maximum	Median	Mean (<i>std. dev</i>)
Age	7	17	12	12 (2.3)
Period of Participation	7 minutes	4 years, 1 month	3 months, 25 days	9 months, 12 days (1 year, 1 month)
Commands Typed	6	51,850	788	6,638 (11,958)
Scripts Written	0	234	2	19.5 (43.3)
Portfolio Score	0	4	1	1.36 (1.47)

Table 1

Results

The children we examined exhibit uneven levels of both participation and programming achievement in MOOSE Crossing (Table 1). It is clear that while some of the users have attained a high level of programming mastery, a larger subset have not written any MOOSE scripts at all (Figure 1).




Figure 1

This supports our informal observation that a small subset of MOOSE users are deeply involved and learn significant programming skills, but a large portion of the community learns only the most basic programming concepts. While over a quarter of the students (categories 3 and 4) are achieving a level of programming knowledge comparable to that taught in a high school class, over 40% do not appear to be learning anything. Despite a reasonable sample size (50 of 405, 12% of children participating), the standard deviations on all mean values are high. The difference between the mean and median scores shows that a small group of kids are greatly offsetting the average. The mean time-on-task of 6638 commands typed and the median of 788 shows that a few enthusiastic users are spending enormous amounts of time on MOOSE Crossing, while the majority are much less active. The median number of scripts written is only 2, as opposed to the

mean of 19.5. Again, a very active part of the community is skewing the average score, whereas the median presents a more realistic measure of participation.

In order to address this uneven level of achievement, we have considered several influences on children's programming performance in MOOSE Crossing. We recognize that users with previous exposure to some form of programming are likely to do better. In addition, we must consider the obvious relation between the users' time-on-task and their achievement. Finally, we pose some interesting questions about the relevance of gender and home versus school use on the children's participation and programming scores.

Prior programming experience is an influential factor on achievement. For the purposes of this study, we considered use of any programming or markup language (such as HTML) as prior programming experience. Of the 50 kids, 16 had some form of previous programming experience (4 girls and 12 boys). As expected, the kids with previous exposure achieved significantly higher scores than those without ($p < 0.05$ (Mann-Whitney Test)). There was no significant difference between the time-on-task between the two groups ($p > 0.05$ (Non-Pooled T-Test)). The children with previous experience had a mean score of 2.4 and a median of 3; those without achieved only a mean score of 0.9 and a median of 0 (Figure 2). It is interesting to note that the group with prior experience consisted mostly of boys. Only 17% of the girls had previous experience, compared with 44% of the boys.




Figure 2

We are not surprised that our data supports the established concept that time-on-task is directly related to achievement. However, we are encouraged by the logarithmic trend in the children’s programming scores based on the number of commands typed. The users who made it past the initial threshold usually went on to create more complex projects. This illustrates that the amount of participation required to gain understanding of more complex programming concepts in MOOSE Crossing is low. It does not appear that gender affects this trend, as shown by the nearly identical curves for both boys and girls (Figure 3). It should be noted, however, that children generally self-select their time on task. We can not say for sure whether the children who chose not to continue would learn more if they were encouraged or required to persist. From this data, we can not separate the element of self selection from time on task.




Figure 3

Noting the heavy gender bias with previous programming experience, we examined whether programming achievement on MOOSE Crossing is directly related to gender (Table 2). While the boys had a higher mean programming score (1.63 boys, 1.04 girls) and a higher median score (1 boys, 0 girls), the curves for both boys and girls have the same approximate shape as the composite (Figure 4) The slightly higher performance of the boys may be explained by their prior programming experience and slightly higher time on task, but these differences are not significant. Our data indicates that gender does not affect the kids' level of achievement or involvement ($p > 0.05$ (Mann-Whitney Test for programming scores and non-pooled T-Test for involvement)).

	β		MAXIMUM		MEDIAN		MEAN (STD. DEV.)	
	Boys	Girls	Boys	Girls	Boys	Girls	Boys	Girls
Age	7	7	1 7	1 6	1 2	1 1	1 2 (2.23)	1 2 (2.36)
Period of participation	7 min	1 3 min	4 yrs, 1 mo	3 yrs, 5 mo	4 mo, 28 days	3 mo, 1 day	1 0 mo, 23 days (1 year, 1 mo)	7 mo, 24 days (11 mo, 21 days)
Commands typed	6	4 1	4 0,436	5 1,850	1 ,632	2 41	7 ,267 (10, 974)	5 ,900 (13,233)
Scripts written	0	0	2 34	1 24	8	0	2 6.6 (52.8)	1 1.3 (27.2)
Portfolio score	0	0	4	4	1	0	1 .63 (1.50)	1 .04 (1.40)

Table 2




Figure 4

MOOSE Crossing is used in-school (42% of participants), in after-school programs (8%), and as a free-time activity (50%). We found no significant differences in time-on-task or programming achievement (Figure 5) among these groups ($p > 0.05$ (Non-Pooled T-Test for time-on-task and Mann-Whitney Test for achievement)). (Note that some school users also have access at home. These students were classified as school users. This complicates the interpretation of this data.)

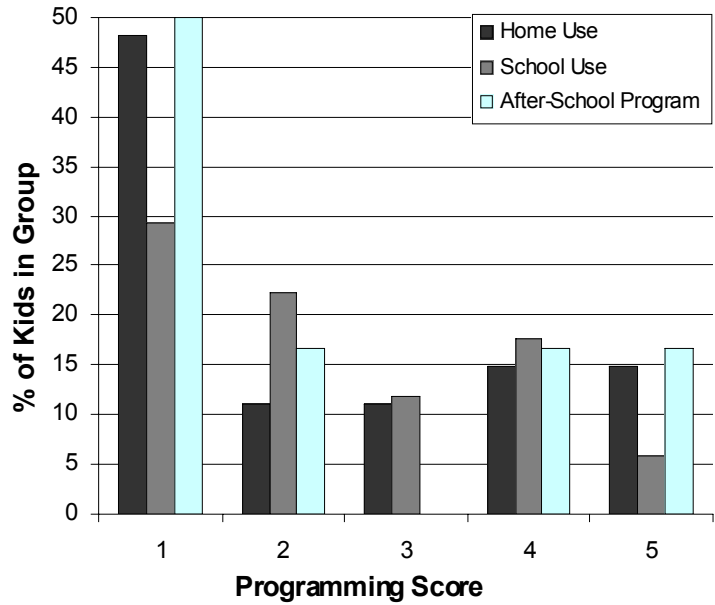


Figure 5

Combating Unevenness: a Change to Both Theory and Design

We do not believe this unevenness to be a problem unique to MOOSE Crossing, but rather a fundamental issue that arises in the design of many learning environments, especially self-motivated ones. In a self-motivated learning environment, students are given the freedom to choose whether to get involved in the hopes that they will embrace the opportunity with genuine enthusiasm. However, inevitably some will decline what is offered.

This unevenness of achievement is undesirable, particularly for an in-school activity. A free-time activity can not be expected to appeal to all children, and unevenness in both interest and achievement is acceptable. However, for an in-school activity, we have a greater expectation that a large percentage of children should benefit at least to some degree, or the activity should be removed from the curriculum.

This led us to a question: how can we begin to remedy this problem without spoiling the open-ended, self-motivated nature of the learning environment? Children using MOOSE Crossing in school are generally given a great deal of freedom in how they spend their time online. We do not want to require specific programming tasks be accomplished, but to encourage the students to chose to learn more.

In winter 2000, this led to the design of a system of “merit badges.” Children are able to earn badges for a variety of programming and writing, tasks. Special badges are awarded for community service. To earn a badge, a child first finds a sponsor. A sponsor is someone who already has earned that badge (or one of a group of community leaders designated to give badges to the first applicants). With help from the sponsor, the applicant prepares a project portfolio demonstrating understanding of the designated concept (for example using a conditional statement, using a property reference, writing something funny, writing something beautiful, etc. A list of currently offered badges appears in Table 3). When the student and sponsor feel the portfolio is ready, it is

submitted for anonymous review. If the reviewer approves the application, the student is awarded the badge. The sponsor is also rewarded for this achievement: sponsoring one other person changes your badge in that area to silver, sponsoring five others changes it to gold. When you have a gold badge, you are eligible to be an anonymous reviewer for that badge

<p>Programming Badges: Rooms Properties Conditionals List1 Input Operators</p> <p>Writing Badges General Descriptive Funny Poetry Spooky</p> <p>Table 3: Merit Badges Offered (as of July 2001)</p>
--

Thus, the new merit-badge system combines elements of portfolio assessment and peer tutoring (Johnson & Johnson, 1987). Our design goal is to raise the level of achievement without spoiling the self-motivated nature of the learning environment. This is a delicate balance. Seven kids and one adult member have earned badges as of July 2001. Sarah² (girl, age 13) is the most enthusiastic of the initial badge earners. She has been participating in MOOSE Crossing regularly since she was nine years old. Since the introduction of the badge system, she has earned six badges and at the time of this writing is working on two more. (A total of eleven badges are currently offered.) She has also sponsored three other members for badges, including her father, Jim. Jim is technology coordinator for a suburban school district, and takes a special interest in educational technology. Neither Sarah nor Jim are “typical” of either the population at large or the MOOSE population; they both have unusual interest in and knowledge of educational technology. However, Sarah’s involvement with the badge system puts her in a good position to offer insights into both the strengths and weaknesses of the system.

In a telephone interview in July 2001, Sarah captured the complexity of the change that the badge system represents:

Amy: Do you think the merit badge system has changed the mood of the place at all?

Sarah: Slightly. Kind of. Well, I don't know why, but it's just like... it seems to a little bit.

Amy: In what way?

Sarah: It's like, more... I mean, I like the merit badge system, but it's more of a... well, I don't know. There's nothing that says you *have* to do merit badges,

²All real names and online pseudonyms of subjects have been changed.

but everyone's like 'merit badge! merit badge!' Like, the talk of the town. It's more, like, learning? That sounds kinda weird. But um, it's like, the mood is like more directed towards learning new things.

Amy: As opposed to what was it before?

Sarah: Slackers? <giggles>

Amy: <laughs>

Sarah: No, it's... before it was too, but it's just more, like, you have to, like... teaching to the test. We're learning what we need to know on the requirements.

Jim said in an interview that he feels that the badge system has had a positive impact on Sarah. He and Sarah participate in MOOSE together often. Before the introduction of the badge system, they had just finished working on a project together. After finishing the project, he noticed that Sarah lost interest to some extent. However, the badge system revived her enthusiasm. Like his daughter, he has mixed feelings about the badge system, but is on the whole positive:

Jim: I see MOOSE Crossing as really having two purposes. One is, it's a playground—an after-school club, a place where you can go and develop an online community. And it's free form. You have your own ideas about what you want to do and you do them. So having this thing that you have to do that has requirements seems very antithetical to that. [...]

But I think that the badge system is very motivating, and if one of your goals is to get students to become more proficient in the programming, then this is certainly a way to do that. Because you set not a very high bar toward attaining these things and some of them I think are quite challenging—you know, doing the ones for lists, and the if-thens, and all that. If you haven't had exposure to that, well that's a good motivation to learn it. So, I think it goes back and forth. The fact that MOOSE Crossing is so open-ended means that you could chose to do the badges or you could chose not to. And the fact that they're there is fine.

We hope that the merit badge system will help teachers to structure the use of MOOSE Crossing in their classrooms. In informal conversations, teachers using MOOSE Crossing have expressed great enthusiasm for this addition to the learning environment. One teacher commented that she would like to use the badge system as a form of assessment for her students: they will be assigned to earn a certain number of badges during the course of the school year. Unfortunately, using the badge system as an assessment detracts from the self-motivated nature of the learning environment. It is possible that this change will negatively impact the atmosphere of the online community.

We are attempting to achieve a delicate balance between a radical constructionist perspective (which eschews assessment and insists that the project should be its own reward) and a more traditional perspective (which sees assessment as an essential component of any learning environment that can meet the demands of functioning in real classrooms). At the time of this writing, the badge system is undergoing iterative design.

We are conducting interviews, observing interactions online, and trying to refine the model to support learning both through earning a badge and being a sponsor. In future work, we plan to perform a detailed study of changes in both students' perceptions of the learning environment and in their level of achievement. Our ultimate goal is not just to create a successful learning environment, but more importantly to test our evolving ideas about motivation and arrive at a more generally applicable advance to the underlying theory of learning.

Conclusion: Co-Evolution of Theory and Technological Design

The MOOSE Crossing project began in 1992 as an exploration of the application of constructionist learning theory to the design of an online learning community. Online communities have the potential to provide social support for learning and teaching, making this approach to learning more scalable and sustainable in real-world settings. Initially, the project began with a commitment to self-motivated learning. In the first few years of using the software with kids from our first working version in 1995 through 1999, we informally observed that achievement in the environment was uneven. Portfolio scoring techniques confirmed this observation.

Taking the time to confirm informal observations is of great importance to progress in the theory and design of online learning communities. Researchers are prone to make mistakes in informal observations. Perception is guided by expectations, and all researchers are vulnerable to seeing what they hope to see. Administrators of online communities are particularly vulnerable to being overly influenced by community regulars. Regulars form the heart of any "third place" (neither work nor home) community (Oldenburg, 1989). Most of an administrator's daily interactions are with

these regulars; however, they are often not representative of the population at large. For example, in informal observations of MOOSE Crossing users, we developed an intuition that girls' achievement exceeded that of boys. In reality, gender does not predict programming achievement (Bruckman, Jensen, & DeBonte, 2002). The highest achieving regular members are disproportionately female, but this does not extend to the broader population. While systematic analysis refuted our informal ideas about gender and programming achievement, this analysis confirmed our informal observations about unevenness of student achievement.

These findings have guided our ongoing work on improving both the system's technological design and our understanding of underlying learning phenomena. These findings support the conjecture that fully self-motivated learning environments can be successful for free-time use, but are problematic for in-school use. To address this issue, we are experimenting with a form of compromise between traditional and radical constructionist approaches to student motivation. Over the next year, we will continue to improve the design of the merit badge system through continual formative evaluation and iterative design. After this lengthy process is completed, we will be in a position to once again undertake more systematic data collection to test our hypothesis that our hybrid approach to motivation can be an effective compromise. Improvements in both technological design and pedagogy for online learning communities can best be achieved through a combination of quantitative and qualitative methods over an extended timeframe.

References

- Baron, J. B., & Wolf, D. P. (Eds.). (1996). *Performance-Based Student Assessment: Challenges and Possibilities*. Chicago: University of Chicago Press.
- Bruckman, A. (1997). *MOOSE Crossing: Construction, Community, and Learning in a Networked Virtual World for Kids*. Unpublished PhD, MIT.
- Bruckman, A. (1998). Community Support for Constructionist Learning. *Computer Supported Cooperative Work*, 7, 47-86.
- Bruckman, A., Jensen, C., & DeBonte, A. (2002). *Gender and Programming Achievement in a CSCL Environment*. Paper presented at the CSCL, Boulder, CO.
- Bruckman, A., & Resnick, M. (1995). The MediaMOO Project: Constructionism and Professional Community. *Convergence*, 1(1), 94-109.
- Chi, M. T. (1997). Quantifying Qualitative Analyses of Verbal Data: A Practical Guide. *Journal of the Learning Sciences*, 6(3), 271-315.
- Dewey, J. (1938). *Experience and Education*. New York: Macmillan Publishing Company.
- Falbel, A. (1989). *Friskolen 70: An Ethnographically Informed Inquiry Into the Social Context of Learning*. Unpublished doctoral dissertation, Massachusetts Institute of Technology.
- Illich, I. (1970). *Deschooling Society*. New York: Harper & Row.
- Johnson, D. W., & Johnson, R. T. (1987). *Learning Together & Alone*. Englewood Cliffs, NJ: Prentice-Hall.
- Kafai, Y., & Resnick, M. (Eds.). (1996). *Constructionism in Practice*. Mahwah, NJ: Lawrence Erlbaum Associates.

- Newman, D., Griffin, P., & Cole, M. (1989). *The Construction Zone: Working for Cognitive Change in School*. Cambridge, England: Cambridge University Press.
- Oldenburg, R. (1989). *The Great Good Place*. New York: Paragon House.
- Papert, S. (1980). *Mindstorms: Children, Computers, and Powerful Ideas*. New York: Basic Books.
- Papert, S. (1991). Situating Constructionism. In I. Harel & S. Papert (Eds.), *Constructionism* (pp. 518). Norwood, NJ: Ablex Publishing.
- Rogoff, B. (1994). Developing Understanding of the Idea of Communities of Learners. *Mind, Culture, and Activity*, 1(4), 209-229.

Acknowledgements

Parts of this paper were originally presented at ICLS 2000, with additional co-authors Elizabeth Edwards (who contributed to portfolio scoring and creation of charts and graphs), Jason Elliott (scoring), and Carlos Jensen (scoring and statistical analysis). The author would like to thank current and past contributors to the MOOSE Crossing project, especially Mitchel Resnick, Alisa Bandlow, Austina DeBonte, Elizabeth Edwards, Tysen Perszyk, Will Scott, Steve Shapiro, Adam Skwersky, Trevor Stricker, Steve Tamm, and Adam Tegen. Special thanks to the kids of MOOSE Crossing and their parents. We'd also like to thank the teachers who have tried it in their classrooms, especially BJ Conn. Research in the Electronic Learning Communities group at Georgia Tech is supported by grants from IBM, Intel, Microsoft, Neometron, Ricoh, and The National Science Foundation. For more information, see <http://www.cc.gatech.edu/elc/>