

Attention By Proxy? Issues in Audience Awareness for Webcasts to Distributed Groups

Jeremy P. Birnholtz^{1,2}, Clarissa Mak², Saul Greenberg³, Ron Baecker²

¹Department of Communication
Faculty of Computing and
Information Science
Cornell University
www.comm.cornell.edu
jpb277@cornell.edu

²Knowledge Media
Design Institute
University of Toronto
www.kmdi.utoronto.ca
rmb@kmdi.utoronto.ca
clarissa.mak@utoronto.ca

³Department of
Computer Science
University of Calgary
grouplab.cpsc.ucalgary.ca
saul.greenberg@ucalgary.ca

ABSTRACT

Instructor/student interaction in e-learning environments can positively impact both student learning and instructor satisfaction. In online webcast lectures, however, interaction can be difficult because instructors lack basic awareness information about their remote students. Our goal is to better understand the kinds of awareness information that instructors should have if they are to interact frequently and effectively with their students in e-learning environments. We conducted an exploratory study – via interviews and observations – of instructor attention in face-to-face classrooms at a large university. Our results imply that a webcast system should provide instructors with overview and detailed data about their students, but that this detailed information should not be displayed publicly.

Author Keywords

Attention, awareness, e-learning, lectures, webcasting

ACM Classification Keywords

H5.3. Group and Organization Interfaces

INTRODUCTION AND BACKGROUND

One persistent theme in the literature on distance education is interaction. In particular, the amount of interaction in an e-learning setting can positively influence both instructor satisfaction [14] and learning outcomes [1].

There are many types of e-learning settings, however, and these vary in the extent to which they facilitate frequent interaction. For example, interaction is easy in online discussions and forums, but is much harder within lecture-style presentations that are simultaneously delivered to a live audience and distributed group of students.

Our focus is on these live and distributed lectures. In particular, we know that it can be quite problematic for the

presenter to be aware of and naturally interact with remote students who are viewing the lecture as a webcast [9]. When the number of participants is small, this can sometimes be overcome using videoconferencing. As the class grows, however, it becomes more difficult [15]. Problems arise because screen space for visible representations of remote participants is limited, and humans have limited capacity for processing such representations [11]. Thus, the rich information that instructors ordinarily use to find appropriate interaction opportunities, to find good “target” students, and to time these interactions is missing. Clearly it is not possible to replicate all of the real-world cues in a distributed environment. This raises the question of whether there are critical bits of awareness information that would be useful for promoting interaction in these environments.

Given the documented value of interaction in face-to-face lectures, others have attempted to improve audience awareness and interaction in videoconferencing and webcasting [2] environments. For example, the TELEP system [9] displayed video or photographic representations of remote audience members on the wall of the physical presentation room, and these were visible to the instructor and the local audience. Next, Chen [5-7] designed a sophisticated virtual auditorium. Remote participants were again represented on a large screen with video, and interaction was via two-way audio and video, along with some gaze-tracking support for “eye contact” between the presenter and remote audience.

One problem with these approaches is that they only scale to support as many remote participants as can be reasonably displayed on a screen of fixed size, and even this may tax instructors’ perceptual capabilities. Full fidelity awareness also raises issues of privacy [4]: the remote person is *not* in the classroom, and connecting the classroom to that person’s physical space can introduce many problems.

An alternative strategy is to display awareness information *by proxy*. In other words, the display might consist of information that is visualized only in the aggregate (e.g., [7]), or detailed information about a subset of the students, with the intent that this subset represents the whole.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5-10, 2008, Florence, Italy.
Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00.

In some ways, this is likely what instructors already do. In large lecture halls, for example, it is clearly not possible to be aware of or interact with every student. Nonetheless, these lectures are seen as effective and are extremely common [3]. How, then, are instructors able to get what is presumably “enough” awareness information about their students in these environments to encourage interaction? This question is the motivation behind this study, and it is critical to understanding how to provide instructors with visual information about their remote students in ways that promote both awareness and interaction.

METHODS

Interviews with instructors and observations of classrooms were conducted between September, 2005 and August, 2006 at a large North American university.

Interview Data

We conducted 20 semi-structured interviews lasting 20-70 minutes using an iteratively developed protocol. Instructors were recruited via convenience sampling, but were deliberately sought from many departments and at many career stages. They represent 9 disciplines and have taught classes ranging in size from 2 to 1,400 students. Interviews focused on how instructors pay attention to students, and how and when they interact. To help participants think critically and specifically about their behavior, we asked them to focus on specific courses they have taught, and also asked how their behavior varies in different settings.

Observational Data

Observations were conducted in 4 courses, 3 of which were taught by interview participants. A similar sampling strategy was used, resulting in classes of varying size from 4 disciplines. Instructors ranged from advanced graduate students to senior faculty. Four sessions of each class were observed, for a total of 26 hours of observations. The observer sat in the classroom with the students, taking field notes focusing on the amount of interaction taking place, and what the instructors' attention seemed focused on. Notes were recorded on paper and later typed and expanded upon. There was no evidence to suggest that behavior was modified due to the presence of an observer.

Data Analysis

Both interview transcripts and field notes were analyzed using qualitative methods described in [13]. Interview analysis consisted of reading and rereading of transcripts by two of the authors. We first coded responses into categories corresponding to the primary foci of the study: audience awareness and interaction. While the interviews covered a range of topics, we focused our analyses here on issues of instructor-student interaction and attention.

RESULTS

Our main finding is that instructors frequently and seamlessly switch between focusing on individual students or groups of students, and high-level information about the

entire class. Moreover, they seem able to do so while maintaining the appearance of attending to the entire class. We center our discussion around three key points.

Instructors aim to give the impression that they are attending to the entire class. Participants indicated that they attempt to consciously look around the room to give the impression that they are attending to all students. Participant O, for example, said that he is “machine-gunning the audience with [his] eyes,” attempting to be aware of everything taking place in the classroom.” Similarly, Participant A said that he “mov[es] from head to head...the whole time” trying to make eye contact with every student. He also noted that he specifically looks for:

signs in their body language in which they're not understanding the material well or they're not satisfied with the way you are delivering the material [A].

Our observations confirmed what instructors said. It was clear that instructors were paying attention to students all over the room. Even in a large lecture hall, the instructor made comments to and interacted with students in the balcony, the back of the room, and the front.

Instructors actually focus on a subset of students. Despite their intent to appear as if they are attending to the entire class, the instructors indicated that they actually pay closer attention to a subset of the students. Participant T, for example, “tend[s] to make eye contact with and engage a handful, a dozen of the students.” This is not to suggest that instructors aren't looking around. For example, Participant C tries hard to look at all of the students:

you're actually talking to individuals in the audience and try to make sure that you see some sort of glimmer of understanding in their eyes [C].

However, it is the most engaged students that tend to reciprocate with eye contact or affirmative nodding to show understanding and acknowledgement. Participant G said that he notices when students are looking directly at him, and makes eye contact with those students:

I'm, probably looking at, at a set of, you know, 3 or 4 individuals who are actually looking at me, and making eye contact, and so if you have one of these people, in different parts of the room, you get good coverage of the room, but you also get kind of these key people who are actually listening and you can gauge reactions to individual things that you've said and stuff like that [G].

Participant D said that there are certain students she looks at for confirmation, such as when she is explaining a difficult concept or wondering if the students are following along:

If somebody normally engaged is looking really confused, it's often an opportunity to say, 'is something not clear?' because normally you're getting one kind of feedback and now you're getting a different kind of feedback which may stand in for feedback from the whole class [D].

Several other participants also spoke about focusing on students to be sure the lecture was understood by everyone, such as those they know might be having difficulty with the material. Participant K targets the extremes: those who seem to be paying close attention, and those who are not:

[I begin by looking] around the room. I look at the back, I tend to gravitate more to those who are nodding positively, or those who are totally out of it and or whispering to their friend or falling asleep and I'll pay particular attention to that latter group and maybe ambush them by asking a question [K].

All participants indicated that they quickly notice when those students they rely on for visual feedback are not in class, because they could mentally picture the seating plan or recognize faces of students, even if not names.

Instructors spatially divide the room into manageable 'chunks' of students. Participants indicated that they divide the room up spatially – particularly if the room is large – and consciously attend to different regions. Some do this to ensure at least an effort to cover the entire class:

I try to look at sections, so I'll look—try to look at the different balcony sides, side, side, sides, and back, front, sides, but I, I think probably the reality is that it's mostly the front [J].

Others do so because they believe students in certain areas may be more or less likely to be engaged by the material:

I don't have to worry too much about those who deliberately try to come and sit near the front, but I'm concerned about those who are midway or beyond, and also routinely scan the balconies [L].

What is interesting in these descriptions is the implicit transition between scanning the entire room by region, and focusing on particular students. Also interesting is the conscious focus on the extremes – those who are nodding along, and those who are not paying attention at all.

DISCUSSION

Our participants want to be aware of and respond to their students. In large classes, they focus on subsets of students, often looking for reactions from specific students. In other words, these subsets are serving as proxies for the whole class. This returns us to our initial question: can awareness be conveyed by proxy in e-learning environments?

Social and Theoretical Implications

One key difference between face-to-face and distributed classroom environments is highlighted by our first two findings: instructors maintain the image that they are attending to the entire class, but may actually be attending to only a subset. In other words, the details of an instructor's actual attention are masked in face-to-face classes. Gaze direction provides some cues, but it is difficult for students to discern the instructor's exact focus. In this way, subsets or "proxies" in face-to-face environments can be selected dynamically and – from the students' perspective – invisibly, as was evident in our data.

In a mixed environment, on the other hand, where remote student images are displayed in a face-to-face lecture room (e.g., [9]), this display is public. Thus, if only a subset of students are being displayed, others will know who is being focused on by the instructor. In other words, the instructor's attention is no longer private. Moreover, being on public display may also mean additional attention for those students from their peers as well. This threatens both

student privacy and the plausibility of the instructor's impression of attending to the entire class.

Making selected student subsets public also raises issues of how to choose these students and whether or not public display of their images is acceptable. Given that some of our participants indicated that they try to focus on students *not* paying attention, this could be controversial. In the TELEP experience [9], for example, many remote participants elected not to display video images because they did not want to be seen multi-tasking in their offices. Public display of student subgroups is also problematic in the case of perceived disproportionate focus on socially charged categories, such as foreign or minority students.

In this regard, attention in e-learning environments is importantly different from, for example, small group video conferences where focus on a person may be an indicator that it is appropriate to draw additional attention to that individual (e.g., by enlarging their video image), as in [10].

Another social issue to consider is the extent to which an "awareness by proxy" strategy will work in a distributed environment, where context is not shared. Our participants described using a combination of looking around the entire room, and then picking on particular students. In this way, the sampling/proxy strategy works because the sample is picked out with at least some conscious awareness of the whole population. When this shared context is taken away, such a strategy may be less effective.

All of this suggests that a "brute-force" approach using video to mimic the face-to-face environment will likely miss some critical and subtle facets of classroom awareness. We suggest an alternative approach that exploits the affordances of the e-learning environment and provides for multiple displays and levels of privacy.

Design Implications

Awareness displays should allow for overview and detail.

The instructors we studied pay attention to their students with two goals in mind: 1) getting an overall sense of the students' engagement and understanding, and 2) looking more closely at a smaller number of students to gauge their individual reactions to the material. Moreover, instructors described switching regularly and rapidly between these modes. This suggests that awareness systems in e-learning should allow for this switching as well. In particular, it should be possible to gauge the overall engagement and comprehension state of remote students, as well as "zoom in" on particular students for more details.

While our data do not point to a specific implementation strategy, one way to accomplish this goal would be a "focus + context" display [8] that presents visualizations of aggregate data about the class, while also allowing instructors to get more information about specific students. This information could consist of video images, or other cues that could reliably indicate engagement for particular students (e.g., data from sensors). This information could be provided either upon deliberate selection of students by the

instructor, or by using an automated selection algorithm based on a combination of cues (e.g., sensor data, *a priori* interest in a student, random selection). Specific implementations and cues are topics for future work.

Instructor attention must remain private. A second key element in our findings is that actual instructor attention in face-to-face classrooms is private. Yet this is not always true in e-learning environments. This presents a unique challenge for designers. The goal in creating public awareness displays for classrooms, after all, was not to invade remote participants' privacy, but to create a peripheral display [12] that could be seen while delivering a lecture and attending to local students. Overcoming the unintended privacy consequences of this design decision is then a matter of creating a "private peripheral display," and determining what information belongs on it.

It may in fact be desirable to allow the face-to-face students to see some information about their remote peers on a public display. Problems arise when it is possible for the local students to discern the instructor's focus of attention from this display, or see detailed information that remote students want only the instructor to see.

One possible solution could be a system that makes aggregate data available on a public display, but uses private displays for zooming in for more details about specific students. For example, a public display such as that used by Chen [7] or Jancke, et al. [9] could display images or data describing the entire remote class. To get more information, the instructor could use a gaze-based interface that displays details on his or her private screen mounted on the podium or "augmented reality" glasses.

A second facet of this implication is that while students should be given the impression that instructors are attending to the whole class, this should be rooted in some sort of real attention. Providing false information would raise ethical concerns and likely have an impact on social relations.

Limitations

This study has two weaknesses, and thus should be interpreted with caution. First, it is modest in size, is conducted at a single university, and relies heavily on self-reported interview data. While these interviews were partially validated by observations, it is not possible to fully verify the claims about attention at particular times. Second, there is a risk in using face-to-face data to design systems for distributed groups; the two are not the same. One key difference is that remote students may not wish to be attended to in the same way that face-to-face students are; as others have found, they may be multi-tasking and paying only scant attention to the lecture [7]. Additional research is needed to fully appreciate the student perspective.

Still, our study moves research in audience awareness away from intuition and toward empirical practice. We identify issues with a proxy approach to awareness that can be applied to the design and refinement of webcast systems that will yield additional insights.

ACKNOWLEDGMENTS

This work was supported by the NSERC-funded NECTAR research network. We also thank our participants.

REFERENCES

1. Anderson, T. Getting the mix right again: an updated and theoretical rationale for interaction. *International Review of Research in Open and Distance Learning*, 4, 2 (2003).
2. Baecker, R., A principled design for scalable internet visual communications with rich media, interactivity, and structured archives. In *Proc. CASCON* (2003), 16-29.
3. Bligh, D.A. *What's the use of lectures*. Jossey-Bass, San Francisco, CA, 2000.
4. Boyle, M. and Greenberg, S. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM TOCHI*, 12, 2 (2005). 328-370.
5. Chen, M., Achieving effective floor control with a low-bandwidth gesture-sensitive videoconferencing system. In *Proc. ACM Multimedia* (2002), 476-483.
6. Chen, M., Design of a virtual auditorium. In *Proc. ACM Multimedia* (2001), 19-28.
7. Chen, M., Visualizing the pulse of a classroom. In *Proc. ACM Multimedia* (2003), 555-561.
8. Furnas, G., Generalized fisheye views. In *Proc. ACM CHI* (1986), 16-23.
9. Jancke, G., Grudin, J. and Gupta, A., Presenting to Local and Remote Audiences: Design and Use of the TELEP System. In *Proc. ACM CHI* (2000), 384-391.
10. Jenkin, T., McGeachie, J., Fono, D. and Vertegaal, R., eyeView: focus+context views for large group video conferences. In *Proc. ACM CHI* (2005), 1497-1500.
11. Kastner, S. and Ungerleider, L.G. Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, 23, (2000). 315-341.
12. Matthews, T., Rattenbury, T. and Carter, S. Defining, designing and evaluating peripheral displays: An analysis using activity theory. *HCI*, 22, 1 (2007). 221-261.
13. Miles, M.B. and Huberman, A.M. *Qualitative Data Analysis: An Expanded Sourcebook*. Sage Publications, Thousand Oaks, 1994.
14. Shea, P., Pickett, A. and Li, C.S. Increasing access to higher education: A study of the diffusion of online teaching among 913 college faculty. *International Review of Research in Open and Distance Learning*, 6, 2 (2005).
15. Webster, J. and Hackley, P. Teaching effectiveness in technology-mediated distance learning. *Academy of Management Journal*, 40, 6 (1997). 1282-1309.