Steven A. Wolfman
Research Statement

wolf@cs.washington.edu

December 8, 2003

My research goal is twofold: to enhance learning through technology and to use technology to understand learning. Technological interventions in the classroom effect change, exposing important phenomena and practices. Understanding these learning phenomena, in turn, informs the design of new educational technologies. This synergy is the heart of my iterative research method: observation of a learning environment, followed by design of an intervention to improve that environment, and then observation of the intervention in its environment leading to further iterations.

My dissertation work was to study interaction in the university classroom through this research method. My contributions include: leading development of the Presenter system for classroom presentation; designing and developing the Classroom Feedback System and the Structured Interaction Presentation system, two student interaction systems based on Presenter; performing classroom experiments to study the impact of the Classroom Feedback System and Presenter; and analyzing this experimental data to develop important insights into educational practices.

In the rest of this statement: I describe how my early research led to my interest in educational technology. I then describe Presenter and its two successors. Next, I describe two insights into the learning environment exposed by my work on these systems. Finally, I outline my future work.

Early Research

My early research was core AI work on resource-constrained planning. The planning problem is to select an action sequence that leads from an initial state of the world to a desired goal state. Resource-constrained planning introduces limited resources, which are produced and consumed by the actions. My key innovation was to combine the strengths of two approaches to such problems: translation to Boolean satisfiability (SAT) and Linear Programming (LP). SAT-based systems excel for discrete choices among courses of action while LP is effective for optimizing continuous resources given a course. My system LPSAT combined these into an effective engine for resource-constrained planning and similar domains. [IJCAI '99, *Knowledge Engineering Review*]

My first experience as an instructor shifted my research focus by inspiring me to apply lessons learned in the classroom to the SMARTedit system. SMARTedit is a text-editing system that learns macro-like programs by demonstration. Unlike traditional macro systems, which parrot editing actions, SMARTedit generalizes from demonstrated actions to programs. Early studies with SMARTedit showed that users were uncertain how best to exploit its learning facilities. In response, I enhanced SMARTedit with new interaction modes based on a model of effective students: the new modes exposed SMARTedit's "knowledge" to its teacher/user through confidence ratings, "asked questions" through active learning[1], and employed decision-theoretic "metacognition" to adjust SMARTedit's learning style to the user's teaching style. The new interface helped users understand how best to use SMARTedit and increased learning efficiency. [IUI '01]

My second experience as a course instructor — for a large introductory programming class — moved my interests further toward teaching and learning. The giant class size, with 200 students, initially intimidated me, and I searched the pedagogical literature for advice on large classes. Almost all the advice I found was dedicated to *mitigating the problems* of large classes rather than *exploiting their advantages*. Intrigued, I focused my teaching efforts that quarter on discovering and emphasizing such advantages and shared my results with the computer science education community [SIGCSE '02]. My interest in pedagogy led me to join my department's Education and Educational Technology (EdTech) research group just as the development of Presenter began.

Presenter System

The Presenter system is an example of my work on enhancing learning through technological interventions. Presenter is a slide-based presentation system that supports high-quality digital ink in the context of slides. It runs as a distributed application, maintaining synchronized but distinct presentation state across a multicast channel. It renders slides using a modular, layer-based architecture, making it easy to augment slide display, *e.g.*, by superimposing extra information over the instructor's view of a slide. [WACE '03]

Presenter has a broad and growing impact on instruction. With the EdTech group, I documented its use and its success in increasing instruction quality in 25 courses taught by 15 instructors at three universities [SIGCSE '04]. Others adopted Presenter without formally participating in our studies.

Presenter's success is founded on the iterative design process I described above. Our study of predecessor technologies suggested that flexible modification of prepared material would be a key enabler for classroom interaction. So, offering this flexibility became our key design goal. Over two years, I adapted Presenter to address other lessons learned during iterative studies. For example, in response to an instructor's request for private notes, I developed an extension allowing instructors to create different views of slides for instructors and students. This is now a favorite feature for managing brainstorming sessions, scaffolding complex discussions, or otherwise offloading cognitive effort to pre-class prep time.

Seeing instructors benefit from computers in the classroom leads naturally to a vision of classrooms a few years from now in which each student has her own computing device. To explore this vision, I designed and developed two systems based on Presenter that support student interaction.

The Classroom Feedback System supports student interaction without much increase in cognitive load on the busy classroom participants. I achieved this low cognitive load by exploiting my observation that classroom participants use slide context heavily in communication, *e.g.*, making deictic (pointing) references to slides. In the feedback system, students use just two pen taps to position simple feedback from a few fixed categories (*e.g.*, "Example" or "More explanation") directly in the context of the slides. The instructor's display shows a customized aggregation of this

[1] This is the machine learning technique called active learning, not the educational technique by the same name.

feedback. This simple feedback is expressive because the slide context connects naturally with the key ideas under discussion. [CSCL '03]

My development of a second Presenter successor, the Structured Interaction Presentation system, was also guided by classroom observations. I found that instructors who were new to the use of interactive exercises had trouble managing the flow of class from interactive exercise back to lecture. The structured interaction system folds interactive exercises into the class's flow by incorporating them into the class slides. Instructors create the exercises by laying out simple interactive widgets in PowerPoint — just as they normally lay out static elements — and linking these widgets together. During class, the system manages student contributions through a database back-end. This design allows instructors to offload some of the cognitive effort of creating and organizing interactive exercises to pre-class preparation and unify this effort with their design of lecture flow (in the form of the static slides). The design also incorporates interactive exercises into an artifact (the slides) that instructors already share with each other across terms, encouraging reuse of the exercises.

Patterns of Practice

In tandem with the development of educational technologies, my iterative research process also exposes important patterns of practice. My technological interventions effect change, and participants' adaptations to these changes often surface hidden aspects of the learning process. In this section, I describe two patterns I discovered while studying the systems described above.

"Feedback lag" is one pattern I identified. A student suffering feedback lag formulates a question about a topic under discussion but lacks the confidence to ask it because he believes the instructor might answer the question without prompting. The student waits for the discussion to move on to the next topic (often signaled, in slide-based lectures, by a slide transition) to ensure that his question is "necessary." Unfortunately, by this time the student often feels it is too late for his question and leaves it unasked.

I first identified this pattern based on participants' free-response survey items in a pen-and-paper prototype of the Classroom Feedback System. I redesigned the feedback system and the survey instrument to investigate feedback lag directly and found further evidence of the phenomenon. Later, review of combined audio/video/ink logs from classes using Presenter confirmed that many student questions do occur just after slide transitions. Together, these data establish a pattern of lagged feedback that can be tapped with a technological intervention.

One study class developed a surprising interaction pattern to address feedback lag. Students left feedback just after a slide came up, and the instructor used these to guide the amount of discussion on the annotated points as he reached them. This pattern gives students an early, socially acceptable opportunity to provide feedback that might otherwise be "lagged." The pattern is particularly exciting since it would be infeasible with spoken or written feedback. [CHI '03]

More recently, a study of existing audio/video/ink Presenter logs led to a new framework for understanding the use of ink during presentations. With three colleagues — including an undergraduate I supervise — I performed a careful study of how instructors used ink in archived courses. Surprisingly, the meaning of much of the ink was ephemeral: clear in the moment but impossible to understand after its spoken context had passed. By coding instances of this ephemeral inking, we were able to establish the prevalence of ephemeral markings and a correspondence with a framework previously developed for physical gestures. [CHI ’04]

Future research

My future research will continue to explore technologies and practices of learning. Further studies of the student interaction systems are a natural starting point. These studies will be excellent opportunities to train students in participant observation and qualitative research. The gestural framework for inking described above suggests research on alternative ink renderings that express context information, *e.g.*, visually grouping ink strokes drawn cotemporaneously. A larger design challenge would be to “explode” Presenter, constructing ink-based presentations dynamically from any set of normal applications on the instructor’s computer.

Beyond my thesis work, I am intersted in exploring untapped aspects of the learning environment and developing technologies. Management of physical and social classroom factors is one interesting direction. For example, instructors and students often consciously choose goals that are difficult to *monitor* consciously such as an instructor limiting audible pauses in speech or trying to maintain gender equity when calling on students, or a student managing time on specific tasks or shifting notetaking from transcription to interpretation and summary. These goals are easily forgotten during the stress of teaching and learning. Imagine a learning environment instrumented to support these choices. This might someday be accomplished through networks of sensors but could be studied today by soliciting data from students through handheld devices. The privacy concerns raised by this vision are all the more reason to understand such systems *before* we find ourselves surrounded by cheap, networked sensors and RFID tags.

My research interests also extend beyond educational technology. Human-computer interaction, ubiquitous computing, software engineering, and educational theory are natural companion areas, and I believe my iterative research method is a good fit for these disciplines. I will also maintain my connections with AI research. For instance, the translation methods I worked with in LPSAT lead to a diverse and easily accessible set of student research projects such as crafting translations from new domains like temporal planning or improving the simple but powerful satisfiability algorithms. Finally, as with my work on large classes, much of my research inspiration will come from scholarly treatment of my own and others’ teaching. In the end, I will keep an open mind and choose my research directions in collaboration with my most important stakeholders: my students.