Chapter 7
The Data-Assisted Approach to Building Intelligent Technology-Enhanced Learning Environments

Christopher Brooks, Jim Greer, and Carl Gutwin

7.1 Introduction

The purpose of this chapter is to describe the data-assisted approach to embedding intelligence in technology-enhanced learning environments that leverages the sensemaking process of instructional experts. In this approach, instructional experts are provided with summaries of the activities of learners who interact with technology-enhanced learning tools. These experts, which may include instructors, instructional designers, educational technologists, and others, use this data to gain insight into the activities of their learners. These insights lead experts to form instructional interventions that can be used to enhance the learning experience. The novel aspect of this approach is that it frames the learning environment as a system that is not just made up of learners and software constructs, but also of the educational experts who may be supporting the learning process. This approach demonstrates how the sensemaking process in the field of learning analytics can be used to affect teaching and learning.

Higher education increasingly makes use of courses with large cohorts of learners and smaller instructor-to-learner ratios. Bloom (1984) demonstrated that learners who are taught in one-on-one learning have, on average, summative assessment marks two standard deviations higher than those taught in a traditional classroom setting. This finding is a principal motivator in the academic research area of *intelligent tutoring systems* (ITS), where software systems form models of learners and adapt the learning environment based on learner performance, much the way a human tutor would. Despite several commercial successes (e.g. Carnegie Learning 1998; Suraweera and Mitrovic 2002) and continued

C. Brooks (✉)
School of Information, University of Michigan, Ann Arbor, MI, USA
e-mail: brooksch@umich.edu

J. Greer • C. Gutwin
Department of Computer Science, University of Saskatchewan, SK, Canada

J.A. Larusson and B. White (eds.), *Learning Analytics: From Research to Practice*,
DOI 10.1007/978-1-4614-3305-7_7,

development of the research field[1], ITS have been deployed only in a few isolated cases and have introduced minimal changes in day-to-day teaching and learning in higher education. In part this is an effect of the expense and labour required to build an intelligent tutor: it requires the work of a domain expert to outline how concepts relate to one another, a pedagogical expert to provide a method for determining when a learner has made a mistake and what should be done to mediate the issue, and a content expert to build initial and remedial content to be delivered to learners. Instructors in higher education rarely have such resources available to them, and instead have to rely upon their own understanding of each of the aforementioned areas, understanding which may be limited and which may be difficult to implement as an automated computational tutor.

Instead of ITS, universities and colleges have focused on scalability with learning technologies and have widely deployed *learning management systems*. These solutions include content creation and delivery solutions, synchronous and asynchronous discussion forums, and multimedia streaming, capture, and playback systems, and provide infrastructure support for learning activities. Learning management systems are used to augment traditional teaching and learning experiences as well as to provide distance and online learning. The environments are, in a sense, non-ITS, where personalization of course content is difficult and the instructor is relied upon to make an a priori identification of the kinds of problems learners may have. While these benefit from a reduction in the number of experienced programmers and knowledge engineers needed, helping scalability and thus leading to greater adoption in higher education, these systems suffer from some of the same resource issues that affect ITS. In order to ensure that all learners are supported, a breadth of content needs to be developed that fits the expected needs and goals of those learners. In sufficiently large or diverse courses, this results in a team of instructional design experts being used to build a comprehensive and pedagogically sound course offering. The end result is the same in terms of a priori effort: building instructional interventions requires significant up-front modelling of learners, pedagogical approaches, and domain-specific content.

In most technology-enhanced learning situations, the interactions happen between the learner and the learning environment, and instructional experts are often limited in their ability to see these interactions and make necessary interventions. Instead of the traditional method of relying heavily on a team of experts who have mapped out the space of possible challenges, difficulties, and misconceptions learners might encounter a priori, we argue that the instructional experts involved in delivering a course can use learner interaction data and make use of their contextual knowledge of the content, cohort, and pedagogy to provide a more individualized learning experience. By considering the issues that learners face while the course is being offered, a smaller highly contextual problem space (e.g. how to teach a specific concept that a specific group of learners are having issues with) need be considered instead of a broader

[1] In particular see the *Intelligent Tutoring Systems* (*ITS*) and *Artificial Intelligence in Education* (*AIED*) conference series, as well as the *Journal of User Modeling and User-Adapted Interaction* (*UMUAI*).

general problem space (e.g. the design of a whole course for learners with a variety of backgrounds). Instructors often employ pedagogical interventions as a part of their normal teaching practice in traditional lectures—it is well known that master teachers observe the behaviours of their students and adapt instruction on demand (Buskist 2004)—but as course sizes grow and interaction moves online, both the observation of social cues that inform the instructor and the ability of that instructor to intervene based on those behaviours becomes more difficult. In this chapter we will show how making the hidden behaviours of learners visible to the instructional expert allows him or her to form *insights*, and enables him or her to react to these insights from technology-enhanced learning environments with *instructional interventions*.

The key aspect of the data-assisted approach to supporting instructional interventions in technology-enhanced learning environments is that both the insights and the interventions are the result of a dialogue between the intelligence in the system and intelligence of the instructional expert. Traditionally there are two forms of intelligent learning environments: those that are *adaptive*, and those that are *adaptable*. Adaptive environments are those described as systems that automatically change in response to user activity based on some pre-programmed knowledge base and rule set. This classification places intelligence in the software system itself. Adaptable environments are those that are personalizable based on direct requests from the learner. This represents a view that the intelligence exists in the learner alone. The continuum between these two approaches is referred to as the *locus of control*, and historically has placed artificial intelligence techniques at one end and human–computer interaction techniques at the other. In the data-assisted approach, interaction information is collected about learners and made available to a third kind of actor, the instructional expert. This expert goes through a sensemaking process, which may be while a course is being run or in between courses, and acts on insights they find by providing pedagogical interventions in the learning system.

The next section of this chapter will describe the data-assisted approach in more detail, ending with three motivating scenarios describing how it can be applied. Each of these scenarios purposefully considers different instructional expert roles (instructors, instructional designers, and educational technologists) and the reasons they might interact with the technology-enhanced learning system. Aspects of these scenarios have each been investigated in depth, and Sect. 7.3 provides case studies demonstrating the outcomes of applying the data-assisted approach. This is followed with a brief conclusion of our findings in realizing this approach.

7.2 The Data-Assisted Approach

The data-assisted approach is intended to generate *insight* that leads to and supports *instructional interventions*. It is different from traditional methods of building intelligent educational systems in that it explicitly acknowledges the role of instructional experts. It is through these experts that insight is generated, and from these experts that instructional interventions come.

Before talking about methods of generating insight or interventions, it is useful to talk about the experts themselves. Instructional experts for a particular environment can be many different people with different tasks. These experts can be instructors, instructional designers, content creators, tutors, markers, or peer learners. The experts can have a formal role in the course, or can be informal actors that engage with the environment through happen-stance. There is already a body of work that examines how learners gain insights about their own actions (e.g. open or scrutable learner modelling), as well as how learners can adapt an environment to their own needs (e.g. adaptable systems). But how instructors tune learning environments through a sensemaking process has been largely unaddressed by the technology-enhanced learning environment community.

The data-assisted approach starts with the collection of information about learners as they interact with tools in the learning environment. This information describes the behaviours of learners, and may be collected through a variety of means across a variety of different tools (see Sect. 7.2.1). The data must then be summarized and correlated to identify groups of learners based on some educational attributes, tasks, or goals.

The process of aggregating and correlating data with pedagogical goals requires that the instructional experts interact with computational elements of the system. A number of different computational approaches can be used, and this chapter will demonstrate how two approaches in particular, *information visualization* and *unsupervised machine learning*, can be successful in identifying aggregations of learners. Regardless of the computational technique involved, it is the instructional expert who gives meaning to the groups, and identifies or parameterizes appropriate instructional interventions.

One of the differences between the data-assisted approach and traditional intelligent learning environments is that the interventions are largely based on the expertise of the instructional expert and not on a priori domain, curriculum, or pedagogical knowledge that has been formalized and loaded into the system. The intelligence is in the overall system that includes the learners, the software, and the instructional expert, and not in software alone. This does not forbid traditional intelligent software (e.g. ITS) from interacting inside a data-assisted approach, of course, but changes the focus of intelligence to be system-wide instead of being located in software alone.

As the data-assisted approach relies upon the instructional expert to provide pedagogical knowledge, it is compatible with many learning theories and instructional design approaches. For instance, an instructor who observes a deficiency of learning in one set of learners might decide that a social constructivist approach is appropriate as a particular intervention, and perhaps configure an online learning environment to send learners to a blogging activity. Another instructor might, after seeing the same grouping of learners, modify the curriculum to include another assignment that focuses on experiential learning. Regardless, both instructors can use the data-assisted approach to get insights into student activity and to customize the course for those learners.

7.2.1 *What Is Meant by* Data *in the Data-Assisted Approach?*

As learners interact with educational environments, traces of their activities can be logged. These traces link actors (e.g. learners, instructors, instructional assistants) and artefacts (e.g. videos, tests, web pages) with interaction behaviours (e.g. watching, answering, clicking). Most educational environments already have some form of trace logging, although many don't separate functional traces (those that pertain to the correct operation of the system) from informational traces (those that are useful specifically for data analysis).

The granularity of the data collected is an important factor to consider when creating an educational environment. One approach is to capture low-level interactions such as key presses or mouse movements (sometimes referred to as clickstreams (McCalla 2004; Peckham and McCalla 2012)), while another is to use more coarse-grained events such as whether or not an answer to a question was correct. The choice between a fine or coarse level of granularity affects the types of behaviours that are available for analysis. A fine level of granularity (e.g. individual key presses) may require the end-user to aggregate data in order to form meaningful insights that they believe are useful (e.g. concepts in a course), while a coarse level of granularity (e.g. a grade on an online examination) may not be decomposable in order to relate to the same meaningful insight. A number of pragmatic concerns around the capture and storage of user data also exists—collecting mouse pointer data as a user interacts with an online content management system, for instance, would create a significant amount of data that would need to be transferred back to a central location for storage. In addition, this level of trace data would almost certainly need to be summarized in order to be useful as an attribute in either automated reasoning or visualization techniques.

A number of authors have considered how semantics can be added to learner trace data to make analysis easier, including (Najjar and Wolpers 2006; Brooks et al. 2004). Building tools to support the data-assisted approach requires a consideration of how data is collected and labelled such that it is meaningful, but this activity takes place outside of the adaptation process, and is largely one of traditional knowledge engineering and system design.

7.2.2 *What Does It Mean to* Assist *in the Data-Assisted Approach?*

The summarization of learner traces into meaningful insights can be considered a form of learner modelling, where the underlying data for the modelling process comes from the interactions learners have had with the learning environment, and the model itself is represented by the insights that are generated from this data. An important differentiator between the data-assisted approach and other forms of learner modelling is that in the data-assisted approach the instructional expert is considered the key actor in the modelling processes. Instead of loading instructional intelligence into the software a priori, there is a reliance on the instructional expert to form hypotheses of learning activity, validate the pedagogical relevance of patterns, and form instructional interventions as appropriate. Thus, software to support data-assisted investigation must collect data about learners, allow instructors to parameterize the analysis of this data as needed, summarize and present patterns of behaviours to the instructors, and map instructional interventions to groups of learners as instructed. The intelligence in such software may be quite limited depending on the data being collected and the mechanisms being used to present this to instructors; on one end of the spectrum, a data-assisted system may rely totally on information visualization techniques with minimal filtering to provide a summary to the instructor, while at the other end of the spectrum a data-assisted approach may provide sophisticated artificial intelligence techniques such as unsupervised machine learning mechanisms to aggregate learners into cliques. It is through the application of these techniques that assistance is given to instructional experts to complete the modelling process.

Embedding the instructor in the learner modelling process aims to increase the scalability of adaptive e-learning systems. While current intelligent learning environments scale well to many learners, they do not scale well between domains, and require a significant amount of domain and pedagogy modelling when being applied to new curricula. By helping instructors to form groups of learners based on behaviours, the data-assisted approach seeks to be a generalizable approach to building adaptive learning environments. Human intelligence in the form of the instructor is used to provide the labelling of groups and their relationships to adaptive components, while artificial intelligence and information visualization can be used to provide a statistical and graphical understanding of the relationships between observed behaviours and groups of learners.

The use of human intelligence in the data-assisted approach fits well with the notion of *sensemaking* as a principal goal of the field of learning analytics as described by Siemens (2012). Klein et al. (2006) describe sensemaking as a process by which events can be understood with consideration of perspectives, which they refer to as a *frames*:

> We can express frames in various meaningful forms, including stories, maps, organizational diagrams, or scripts, and can use them in subsequent and parallel processes. Even though frames define what count as data, they themselves actually shape the data (for example, a house fire will be perceived differently by the homeowner, the firefighters, and the arson investigators). Furthermore, frames change as we acquire data. In other words, this is a two way street: Frames shape and define the relevant data, and data mandate that frames change in nontrivial ways. (Page 88 of Klein et al. 2006)

The process of the instructional expert interacting with the data collection and analysis aspects of the learning environment can be thought of as a dialogue; as the instructional expert elicits the formation of groups or classifications of learners from the system, he or she can form new hypotheses as to the state of learning happening in each group, and modify how course attributes such as content, sequencing, activities, or tools are made available. The dialogue is bidirectional—the instructor uses the environment to better understand learners, and shares this understanding of learners with the system through labelling. The system takes these labels, and the associated adaptations provided, and applies them to learners who fall within the groups based on learner behaviours.

The dialogue between system and instructional expert is also an ongoing one. As learners continue to interact, or as new learners begin to interact, with the tools in a learning environment, new data and new behaviours can be presented to the instructional expert. This expert creates, deletes, and modifies instructional interventions as they apply to learning objectives and pedagogical approach. Thus, the modelling process is intended to fit within the day-to-day activities of an instructional expert such as a course instructor, and not be an activity that takes place before learners are present, as much traditional instructional design activity does. This is not to say there is no place for a priori design in electronic learning environments or that, when available, resources to build ITS should not be utilized. Instead, the suggestion here is that the data-assisted approach is a complementary method to both instructional design activities and intelligent tutoring methods, and requires a different kind of resource that may be more readily available in the context of higher education (e.g. the instructional expert).

7.2.3 Motivating Scenarios

The rest of this chapter will consider how the data-assisted approach might be applied in three different real-world contexts. To aid in this, we present here three scenarios that explore the needs of different kinds of instructional experts. While the scenarios presented here are fabricated, each will be paired with details from actual investigations (Sect. 7.3). The emphasis of these case studies is not on completely satisfying the requirements of each scenario (a monumental task), but in providing evidence that demonstrates how the data-assisted approach has been used to solve particular problems.

7.2.3.1 Scenario One: Visualizing Community Interactions

Katheryn is a faculty member who regularly teaches introductory Computer Science for non-majors. This course is typically moderate in size (50–100 students) and made up of learners from a wide variety of disciplines. This year, Katheryn is teaching the course in an online capacity instead of in a traditional lecture format. Learners have access to an online content management system that includes sets of web pages describing content as well as an asynchronous discussion forum. There are 20 learners enrolled, and because of the distance component the majority of evaluation is weighted on the final examination and assignments.

The content for the course is broad in nature, and Katheryn feels that encouraging group discussion will be one key to keeping learners engaged. She is concerned that the distance modality of the course will cause learners to shy away from interpersonal interactions, and scaffolds this aspect of the course by making weekly reading assignments. Each week, she will post questions to the discussion forums and learners will have to reply with their thoughts on the issues. By having learners read one another's posts, she hopes they will form a shared sense of community, resulting in greater engagement and deeper learning.

In a traditional discussion forum system, Katheryn can see only that students have written messages, as well as the content of those messages. All other interactions, such as traces related to the reading of postings, are inaccessible to her. Using the data-assisted approach, tools can be built to visualize these hidden interactions, enabling Katheryn to more deeply understand how learners are collaborating.

An important aspect of the data-assisted approach is that the underlying data is summarized as it is presented to the instructional expert, and the expert is actively involved in the summarization process. The intelligence in the system is thus a *mediated* and *emergent* property; it exists because of the interaction between the domain and pedagogy expert and the data-rich learning environment. In this scenario, Katheryn has some specific thoughts as to what active engagement means; i.e. that learners are reading and considering one another's messages. Katheryn might be able to aid in the summary of the data by modifying the visualization to render traces differently based on her goals. She also might be able to compare her current class visualizations with those taken in previous years, or those of other classes where she knows different pedagogical techniques are being employed.

The data-assisted approach for supporting instructional interventions is made up of two activities: the summarization of usage data to generate insight, and the support for making instructional interventions based on this insight. This scenario has focused on the first of these, and demonstrates that by revealing to instructors the hidden behaviours of learners (such as the reading of a message), an instructor can understand the affects of their pedagogical practice. The data-assisted approach does not require that instructional interventions take place directly in the technology-enhanced learning environment; insights can be leveraged to form traditional instructional interventions, such as curriculum changes or changes to assignment activities (as Katheryn did). By providing mechanisms to compare visualizations over time, instructional experts can compare the effects of their actions and deepen their understanding of the learning cohort as well as their pedagogical practice.

7.2.3.2 Scenario Two: Measuring Educational Impact

Michelle is an instructional designer supporting a second year undergraduate Chemistry course. This course serves as a service course for other colleges (Agriculture, Engineering, Medicine) as well as for the core Chemistry programme. The course is taught by five different instructors with a common curriculum and common set of examinations. The course reaches a total of 600 learners in a given semester. One instructor has agreed to have his lectures recorded using a lecture capture and playback system. These recordings are then made available (asynchronously) to all students enrolled in the course.

Michelle's responsibility with the course is to manage both the short-term and long-term instructional design. The lecture recording system is new this year, and Michelle is interested in better understanding how learners use recorded lectures and the effect they have on performance. In particular, Michelle is interested in understanding whether the students who use the lecture capture system incorporate it regularly into their study habits, or whether they just use it only during the examination period. If there is a difference in performance between these groups she will adjust the course content as appropriate, guiding the learners to more effective study habits.

The data-assisted approach is ideal for Michelle as she is interested in exploring the effect of introducing a new tool as well as tuning the course to take best advantage of that tool. Michelle has deep knowledge of different instructional interventions, but lacks an awareness of the underlying activities of learners on which to base these. In a traditional classroom Michelle could go to lectures and observe attendance, a time-consuming process. Further, the lack of time shifting tools in a traditional setting (such as lecture recording) can cause learners to rely on traditional methods for consuming content (e.g. notes or textbooks) that Michelle cannot easily observe. With the data-assisted approach, learning tools are augmented with tracking features that log the activities of learners and make them available to experts like Michelle. While there may still exist modalities of learning outside these, the set of tools that collect learner traces along with methods of summarizing and acting upon these traces form the basis of the learning environment. Michelle is now able to watch in real-time as learners use the system, and leverage her pedagogical expertise to change the way in which the course is designed. Further, Michelle is able to use these traces of learner activity to expand her understanding of the knowledge domain and how learners interact with it.

Michelle's interest in seeing groups of related students fits well with unsupervised machine learning techniques (*clustering*). For instance, Michelle might be interested in grouping learners by their weekly access patterns in the lecture capture system and comparing these to midterm evaluation marks. Depending on her ability, she might directly interact with the logging subsystem of the lecture capture product, or a user-friendly graphical interface (e.g. as in Brooks et al. 2012) may provide her the ability to choose the kinds of data she is interested in clustering. Regardless of the underlying tool used, as the system returns results for her to consider, she may tweak her clustering parameters and attributes based on her background knowledge of what educational theory is appropriate.

This scenario has illuminated two aspects of the teaching and learning environment that are relevant to the data-assisted approach. First, courses are often large multi-section offerings made up of hundreds of learners where a visual in-person analysis of activity is either difficult or impossible. As learning opportunities increasingly happen outside of the classroom (e.g. through time shifting of lectures using lecture capture), the need to capture interaction between learners and technology grows. The data-assisted approach is built on the premise that these interactions allow instructional experts to gain valuable insights into how learners interact with the educational environment.

Second, the data-assisted approach puts the instructional expert at the centre of the *sensemaking* process. In this scenario, it is Michelle who determines whether the clusters fit with her pedagogical approach and are relevant to her investigation. The educational environment helps aggregate and correlate behaviours, but it is the expert who validates the results and forms interventions. How these findings might be tied to interventions is largely omitted from this scenario, but will be touched upon in the next scenario (Sect. 7.2.3.3).

7.2.3.3 Scenario Three: Adapting Learning Environments to Tasks

The previous scenarios have described how an instructor and an instructional designer might use data-assisted approach tools in their daily activities. Intelligent learning environments typically offer automatic adaptivity in response to learner behaviours. The data-assisted approach is appropriate for configuring and building of these systems by bringing technologists into the sensemaking process. Whereas in a traditional intelligent learning environment the sensemaking is done a priori by technologists, a data-assisted intelligent learning environment encourages the creation of adaptive features, or the configuration of existing adaptive features, in response to new patterns of behaviours observed.

This scenario follows Adam, an educational technologist with a software engineering background, who is supporting a lecture capture system similar in nature to the one Michelle is using. Adam has received a lot of feedback from instructors that the navigation in the application is hard for students, especially since the principal method of navigation through the video is by selecting from a series of index thumbnails that are generated from the captured data projector every 5 min of recording (Fig. 7.1). Some instructors want index thumbnails more often, while others want fewer thumbnails that are more like chapter markers in traditional DVD media.

Adam understands the perspectives the instructors have shared, and he believes that students navigate through the videos according to their needs. He may examine the behaviours of learners using visualization tools for the lecture capture system that have similar capabilities as the ones Katheryn used in Sect. 7.2.3.1. Students rarely use the same indices when navigating, but he could use clustering tools (such as described in Sect. 7.2.3.2) to find three clusters of activity which he might label as *reviewing* for frequent use of many indices, *watching* for those students who seem to only use a few indices that refer to significant breaks in the lecture, and *images* for a cluster of access that appears to focus on images in slide content (e.g. chemical drawings in material sciences courses, paintings shown in art history courses, or diagrams used in paediatric nursing courses). He may then either parameterize the indexing functionality of the lecture capture environment based on these clusters, or change the functionality of the system as appropriate.

In this scenario, Adam has used the data-assisted approach to understand how learners are using the lecture recording tools. He has identified patterns of behaviour using tools similar to those used by Katheryn and Michelle. As shown in this section, the data-assisted approach does not have to end with just the detection of patterns.

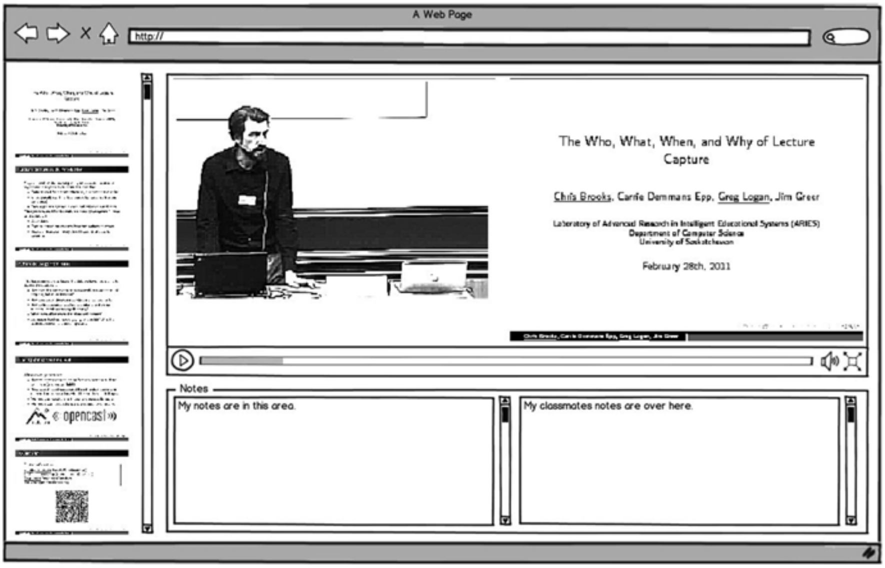

Fig. 7.1 A mock up of a lecture capture environment. The images down the *left hand side* control navigation through the video, and allow a user to seek to a particular *chapter* directly. The video playback component on the *upper right* includes video of both the instructor and the data projector, as well as a traditional scrubber that allows for non-chapter seeking. Additional tools such as note-taking components, discussion forums, or suggested readings are available underneath the video

While Katheryn and Michelle were able to leverage insights to change their pedagogical approach and instructional designs respectively, Adam is making the results of the patterns available to the educational environment by labelling clusters. With labels attached to prototypical behaviours (the student *centroids*), learners can be classified into groups and the appropriate indices can be shown automatically.

7.2.4 Conclusion

Here we have described how the data-assisted approach works, and how it might be used to aid in teaching and learning process. The scenarios presented here are intended to be motivational, and describe how different instructional experts might interact with learning environments that collect end-user behaviour data. To this end, a number of different actors (instructors, instructional designers, and educational technologists) with different use cases (supporting learners, analyzing effectiveness of tools, and customizing the learning environment) using different instructional modalities (online distance learning and traditional face-to-face lecture learning) teaching to different-sized class have been considered.

7.3 Case Studies

7.3.1 Introduction

Over the last 7 years, the laboratory for Advanced Research in Intelligent Learning Environments (ARIES) at the University of Saskatchewan has been applying the data-assisted approach to different technology-enhanced learning environments. The goal here has been to validate that the approach is appropriate for enabling new technology, as well as to realize the teaching and learning gains of these next-generation environments. To this end, thousands of higher education learners from a variety of disciplines have taken part in studies that have used the tools we will describe.

Due to the scope and depth of the studies we have engaged in, a full account of the findings from these studies would be difficult to provide here. Instead, we focus this section on describing three cases that illuminate how the data-assisted approach has engaged experts similar to those presented in the previous scenarios.

7.3.2 Case Study One: Visualizing Community Interactions

One way the data-assisted approach is different from other methods of creating intelligent learning environments is in the way insight is generated. In the data-assisted approach, insight is created through an interaction between an instructional expert and data that represents a particular cohort of interest (a group of learners). A more traditional intelligent learning environment approach would be to fully describe the learning space before the system is deployed, then categorize learners and react accordingly. The data-assisted approach is more reactive, and focuses on in situ exploration of learner activities.

Similarly, instructional interventions can be constructed and delivered through multiple methods. A traditional intelligent learning environment such as an ITS or adaptive hypermedia system changes content delivered to a learner in reaction to their activity. These systems will not often react by changing pedagogical approach, however, unless they have been specifically pre-programmed to do so. In contrast, human instructors often react by considering how they might restructure aspects of a course. By giving these experts control over how instructional interventions are formed, larger pedagogical changes can be made with minimal a priori consideration.

Based on these differences, this section explores two questions that arise when considering the scenario involving Katheryn in Sect. 7.2.3.1:

1. Is it possible to augment a discussion forum system to capture the hidden traces?
2. Can an instructor derive enough meaning from a visualization that they are able to modify or improve upon their pedagogical practice?

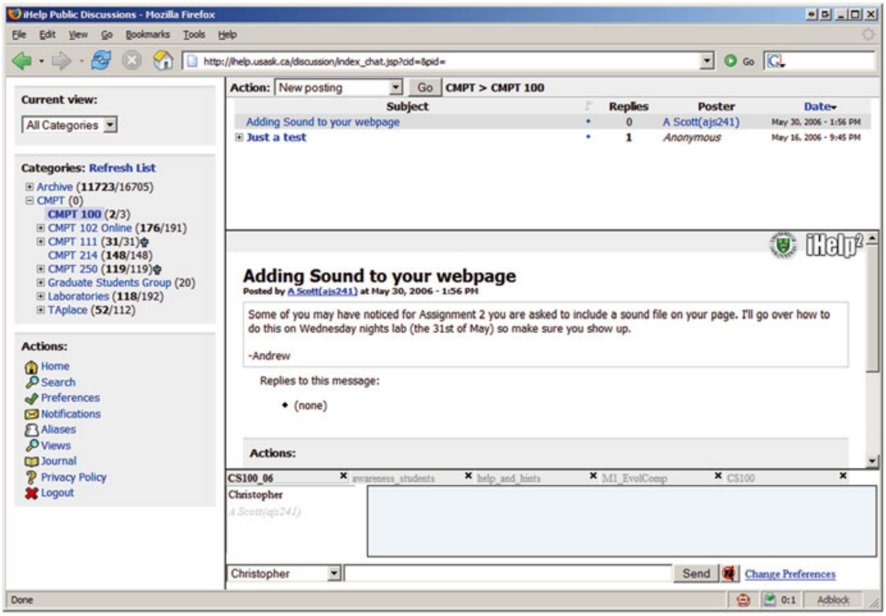

Fig. 7.2 The iHelp Discussion forum system c.2006. The *left side* shows the different forums that a user has access to, in this case a number of course-related forums as well as social and employment forums. Forums are heirarchical, and may include sub-forums related to special topics. Each forum is made up of heirarchical *threads* of discussion as shown in the *top window*. Content for a given *posting* in a thread is shown in the *middle* of the application, and other e-learning tools (such as a synchronous chat system) can be shown at the *bottom* of the application

7.3.2.1 The iHelp Discussions Environment

The iHelp Discussions learning environment (Fig. 7.2) is a Web 2.0 asynchronous discussion forum intended for use within higher education. It was developed out of the iHelp research project (Greer et al. 2001) with the goal of increasing usability and scalability of a technology-enhanced learning environment supporting peer help. The system was deployed within the Department of Computer Science from 2004 to 2010, and was used by thousands of students in dozens of courses annually. When it was retired from general use by the Department in 2010, it contained over 75,000 messages in over 2,200 forums created by over 3,300 users.

This environment is different from other web discussion forums in that it has been augmented to record learner interactions at a fine-grained level. Where other systems typically would show all of the messages within a thread at one time, iHelp Discussions has nested lists of messages which allow the system to record when a user requests a message and how long the user stays with the message open. Further, the forums are hierarchical in nature, and access to forums is logged using similar mechanisms. Over the 6 years it was deployed, more than 3,000,000 read requests were issued by users, with the top message being read over 1,700 times.

The use of iHelp Discussions as a forum ranged broadly by instructor and course. It was commonly used in large cohort undergraduate courses, and there were several department-wide “off topic” forums that students used to discuss issues of technology, philosophy, and politics (amongst others). Access to the forums was restricted through an institutional username and passwords, and public access to messages was not available. In some forums learners were permitted to post anonymously, though the back-end system is still able to disambiguate usernames if needed.

Instructors were free to structure sub-forums however they wanted, and it was common for topic-based, course structure-based, or a flat single forum environment to be used.

One unique aspect of the iHelp Discussions deployment situation is how access is granted to instructional experts. It was not uncommon for instructional experts such as instructors and tutorial/lab assistants to be routinely given access to all of the undergraduate forums. This, along with the detailed usage tracking, makes it possible to study the behaviour of instructional experts as well as learners.

7.3.2.2 Visualization of Learner Activities

Sociograms are a common method of visualizing asynchronous discussions, and have been used by a number of researchers to visualize email correspondence in particular (e.g. Weskamp 2003). These visualizations are graph-based structures where nodes represent individuals in the community and arcs between nodes represent the creation of replies to a message. Nodes are typically rendered using different sizes to represent the status of an individual in the community or the amount of discussion the individual has contributed (e.g. Weskamp 2003). Nodes can be arranged in a number of ways: a strict hierarchy which outlines the abilities or status of groups of nodes is common, as is a physics-based “force graph” which moves nodes closer to one another depending on the characteristics they share.

A representation for the iHelp Discussion forums was formulated using sociograms where nodes indicate particular persons involved in a course and edges between those nodes indicate a discussion replied to relationship. To clearly indicate the difference between learners and instructional experts (e.g. Instructors, Tutorial Assistants, and Markers), each node is colour coded to be either a learner (light grey), or an expert (red). As the iHelp Discussion forums are available to many instructional experts (faculty and assistants assigned to other courses) regardless of the content of the forum, there are typically many red circles in the sociogram.

In large courses (e.g. those with more than 100 students) this formulation quickly became unwieldy. To address this, individuals are further broken up into membership in one of three categories:

- Participants: Those individuals who have written messages, either on their own or as replies to other messages.
- Lurkers: Those individuals who have read postings but have not written any.
- Non-users: Those individuals who have never read nor written a posting.

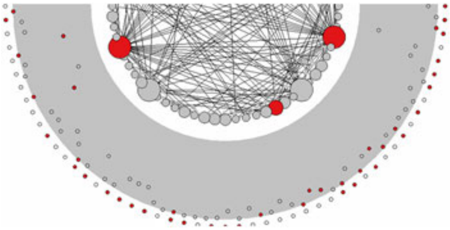

Fig. 7.3 Portion of a sociogram from an introductory computer science course that was taught in a blended fashion and had approximately 200 students and contained 254 discussion postings. Darker (*red*) nodes indicate facilitators, while lighter (*grey*) coloured nodes indicate learners. The *inner circle* is made up of participants, four of which are very important to the community (as shown by having a larger node size). A casual observation of this network indicates that, while some learners write a fair bit (many interconnected nodes in the middle), there are lots of learners who haven't ever read anything (the outer ring of non-users), and many lurkers who read very little (as they tend to be closer to the outside of the middle ring instead of the inside of that ring). Note that the ring of non-users includes a disproportionately high number of facilitators as our current deployment gives access to this forum to most staff and faculty in the department (Colour figure online)

Each category of users is put into their own sociogram that aligns nodes along the exterior of a circle. The different sociograms are then layered on top of one another such that the non-users are farthest from the centre of the sociogram, the participants are closest to the centre of the sociogram, and the lurkers are in between (Fig. 7.3). This corresponds well with the perceived participation rate of individuals (*participants* are more central than *lurkers who* in turn are more central than *non-users*), as well as with the sizes of the different categories of individuals (generally there are more non-*users* than there are *lurkers*, and more *lurkers* than there are *participants*).

Lurking is modelled as a continuous variable—one individual can lurk more than another by reading more forum postings. To support this in the visualization, lurking ratios are calculated and node distances are varied from the outer edge of the lurker region to the inner edge, where lurkers who are closer to the inner edge of the sociogram have read more content. This further reinforces the idea that users who are central in the overall visualization are participating more in the course than learners that are close to the edge of the screen.

Initial feedback from instructors indicated that the act of posting a message does not mean that a user has contributed in a meaningful way. To represent a measure of *importance* in a course, the size of an individual node is varied by the number of persons who have read a user's postings. The calculation for an individual's *importance* is given in Eq. 7.1.

$$\text{Importance} = \frac{\text{Number of people who read my postings}}{(\text{Number of participants} + \text{Number of lurkers}) \times \text{Number of postings}} \tag{7.1}$$

A number of observations about the sociograms can be made based on informal interactions with instructors who reviewed the prototype. In particular, it was observed that:

- A highly connected graph indicates that learners are communicating with one another, while a graph where many nodes are connected only to an instructional expert indicate little peer collaboration.
- The degree of edges coming out of expert (red) nodes indicates how much direct control an instructor has over conversations. Instructors who wait to answer questions have very few arrow heads pointing at their node, while instructors who provoke discussion have many arrow heads point at their node.
- Lurking rates are highly variable, and the majority of lurkers read fewer than 30 % of the postings. This includes non-learners (e.g. tutorial assistants or other instructors), a result that surprised many of the instructors who saw the visualizations.

The next four sections provide specific examples of how this visualization has been used by instructors, and demonstrate the effect of making learner traces more readily available.

7.3.2.3 Example One: Online Small Cohort Course

Many instructors change their pedagogical approach when teaching in online environments, in order to match the needs of learners. In these situations, learners can more easily lose a sense of shared community or shared purpose as the activities and actions of their classmates may be hidden from them. The instructor feels this change too, and many of the consequential awareness indicators that they would normally get from an in-person teaching environment (e.g. attendance and interaction in the lecture, or students coming to tutorials and office hours for assistance) are missing. By revealing indicators of activity in discussion forum systems, the data-assisted approach can help instructors to understand what the virtual classroom interaction looks like.

The Department of Computer Science offered an introductory course for non-majors on the topic of basic Computer Science history, principles, and techniques. In 2006 this course was taught simultaneously to a large cohort of over 100 learners in a blended manner, and to a small online group of 20 learners who had no face-to-face instruction. While taught by different instructors, the online instructor was also the creator of content, assignments, and examinations for both sections.

The online instructor was a firm believer in using the iHelp Discussion forums to build a sense of community amongst the students. Especially in the online course, she saw the discussion forums as the main method of engaging with learners. To this end, one of the weekly requirements was to write an online forum posting about a prescribed course topic and this requirement was mandatory for the online learners only. She would post the initial message indicating what the weekly topic was, and learners were expected to respond with details of content they found on the Internet.

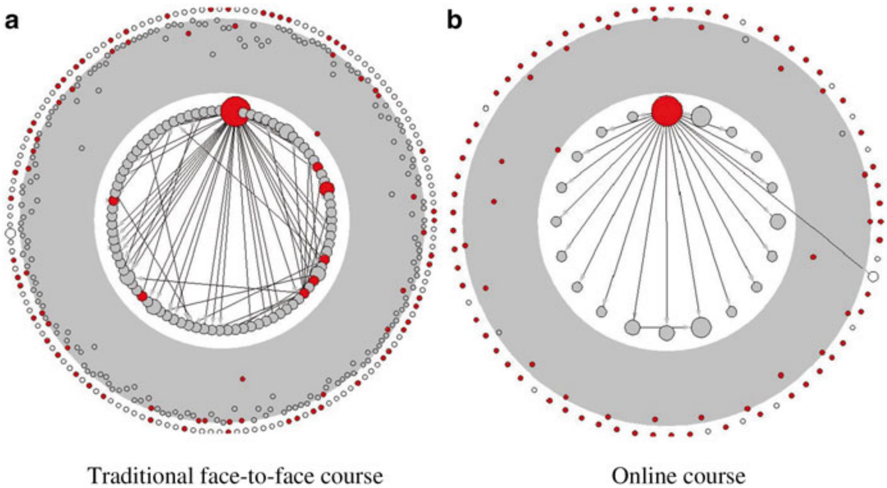

Fig. 7.4 Visualizations of an introductory Computer Science course for non-majors. (**a**) Shows interactions amongst learners who attend traditional lectures, while (**b**) shows interactions amongst learners in a completely online environment. In the traditional course there are many lurkers who have read only a few messages, while the online course required participation of learners and thus has very few lurkers (only other instructional experts). In both cases the instructor is the principal actor in the social network (*large dark* (*red*) node), and learners rarely reply to anyone but the instructor (Colour figure online)

Her reasoning was that by sharing the results of this activity publicly, learners would form a sense of community with one another. A number of weeks into the course, the instructor was shown visualizations for both the in-class and online discussion forums (Fig. 7.4).

The visualization was explained to her, and she showed particular interest to the way size of nodes was generated. She was bothered by the fact that she could easily identify herself as the large dark (red) node who was connected to most students. She wanted students to communicate and form a community with one another, and not just herself. That the pattern existed in both the online and the traditional cohort didn't surprise her, but she felt that traditional learners had other mechanisms by which they form a shared sense of community and thought that making the assignment mandatory for the online learners would be enough to encourage broader use of the discussion forums. In a subsequent offering of the course, the instructor changed the weekly assignments to be more problem-solving based, and set the evaluation criteria such that learners would interact more. The result was a discussion graph that was more fully connected and where learner nodes varied in size (thus *importance*) relative to the instructor (Fig. 7.5).

The use of the data-assisted approach provided insight that was otherwise lost to the instructor. By visualizing the hidden traces learners leave, the instructor became able to understand the effect of her pedagogy, and make interventions that were reflected through subsequent visualizations. Further, the instructor was

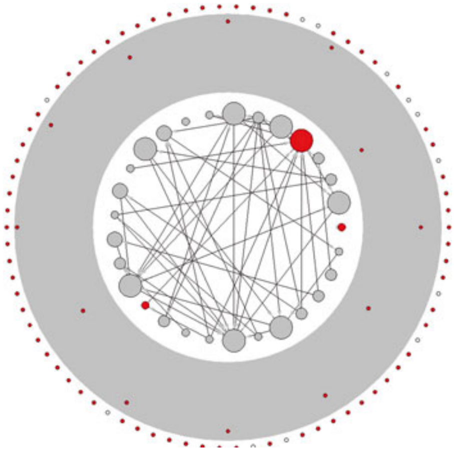

Fig. 7.5 A visualization of a subsequent offering of an online introductory Computer Science course for non-majors. This course was offered after the instructor made pedagogy changes aimed at having learners interact more with one another. In this example, there are several learners who are *important* to the community as shown by their large node size and high level of connectedness

able to compare the familiar traditional face-to-face cohort to a new mode of teaching, and create the kind of learning environment she thought was important for online learners.

7.3.2.4 Example Two: Face-to-Face Large Cohort Course

In traditional face-to-face courses, instructors have the ability to make direct observations of learners. But these observations are minimal in bandwidth, especially when dealing with large cohorts, and tend to only give instructors a broad understanding of the issues learners face. Especially in introductory courses where many of the learners are new to both the discipline and higher education, there is often a hesitancy to speak up in the classroom. The hidden traces of online activities can be leveraged in these circumstances if the traditional course is augmented with technology and taught in a blended mode. Blending a course generally changes the lectures minimally, but offers other avenues for help and exploration (such as discussion forums).

A visualization (Fig. 7.6) of a large cohort of learners involved in introductory Computer Science courses for majors was shown to one of the instructors who teaches in a blended fashion. Unlike the study described in Sect. 7.3.2.3, all learners taking this course (regardless of section) were shown one discussion forum. The instructor shown the visualization was a lecturer as well as the overall administrative organizer for the course, and he would often leave the discussion forums open throughout the day so he could answer student questions. He immediately identified the nodes that represented himself and one of his tutorial assistants (large dark (red) nodes), and theorized on the identities of several of the learner nodes using the size and in-degree. He seemed comfortable with the visualization as an interpretation of the community formed in his course, and felt that the results he had achieved were those he set out to achieve.

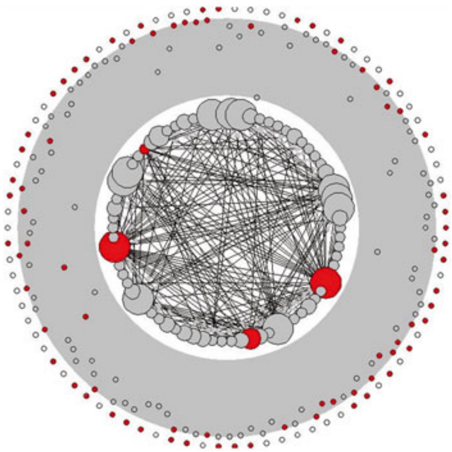

Fig. 7.6 A visualization of an introductory course in Computer Science for majors. This visualization is made up of several sections of learners, each with their own instructor. The high degree of connectedness along with large student nodes (*light grey* in colour) supported the instructor's notion that course discussion would be sustainable if he reduced his level of involvement

A follow-up interview with the instructor provided surprising results; despite being pleased with the interpretation that resulted from the visualization of his course, he still modified his pedagogical approach. He indicated that for the first portion of his course he would typically answer questions as soon as they appeared in the discussion forums. He knew from the volume of questions being asked that learners used the system heavily, and he wanted to address those questions publicly as fast as possible in order to minimize their waiting. However, after seeing several large learner nodes in the visualization (e.g. several learners were read regularly by their peers), he felt that the community had reached a level where peer help would be sustainable without his intervention. In short, he felt he could reduce the speed at which he replied to learners without negatively affecting their learning experience, as their peers were active both in writing responses to questions and in having those responses read by classmates. He was shown a follow-up image of the social network for his course after some time, and remained satisfied that interactions were still happening at a good pace.

In this case, the data-assisted approach was used to confirm a belief an instructor held which would have been otherwise untestable. By aggregating the low-level read events captured by the forum system into a simple metric (*importance*), a small feature of the visualization became enough to convince the instructor of the self-sustaining nature of his course discussions.

7.3.2.5 Example Three: Comparison of Visualizations based on Granularity

The heirarchical nature of the iHelp Discussion forum system allows instructors to create niche topics for their courses. One such example of this comes from an introductory course in Computer Science for majors, where most of the discussion

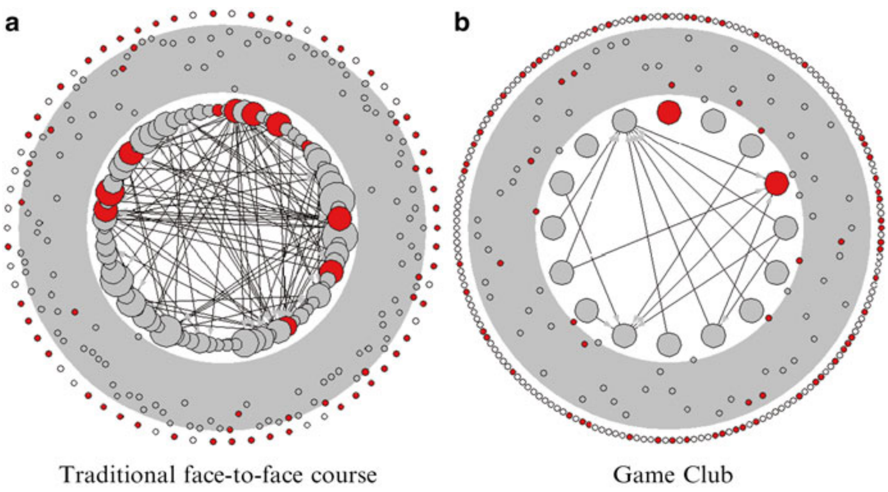

Fig. 7.7 Visualization of a general discussion forum (**a**) and a more specialized discussion forum which is unrelated to course content (**b**). The user cohort is the same, though qualitative remarks about participation in and importance of the discussions can be made by comparing visual aspects of the two renderings

happened in a general discussion forum (726 of 901 messages), but some happened in topic-specific forums. This course taught general principles of Computer Science, but based on the interests of the instructor and the students a sub-forum called the *Game Club* was created to talk about video game-related issues. These two discussion forums can be visualized individually (Fig. 7.7) and, while no discussion with the instructors about this forum took place, some high level qualitative comments can be made.

First, it is easy to see that there is much more activity in the general discussion forum than the game forum by looking at the in-degree of participants (higher in general forum) and the number of persons in the non-user sociogram (higher in game forum). Second, the rate of lurking is higher in the more specialized game forum, though the cause of this is unknown (it might be that there are just fewer postings, so it is easier for the keen learners to read them all, or there is an initial surge of activity where the community has promise but then dies out). Finally, it is worth noting that a number of instructional experts sit very close to the participation ring in the game forum, but do not write new messages or reply to messages to become participants. The precise role of the instructional experts here is not clear and it is difficult to draw specific conclusions; they may be instructors who are keeping an eye on developments, or tutorial assistants who are peer students and very much interested in the content of the discussions.

Insight coming out of use of the data-assisted approach is very much about making visible the invisible—instructional experts cannot regularly see interactions learners make and thus do not involve them when making pedagogical decisions.

In this example, an end-user looking at discussion forum postings without the visualization would assume that only the two instructional experts who have written messages in the game forum are interested in the topic. The visualization, however, makes it clear that this is not the case, in that there are eight other instructional experts who have read almost every message posted in that forum. By making visible the hidden traces users leave as they interact with the learning environment, instructional experts are better positioned to make both short-term and long-term pedagogical decisions. Given the keenness of learners to view this game-related forum, it might not be unreasonable for similar content to find its way into the standard curriculum for the course in future offerings.

7.3.3 Case Study Two: Measuring Educational Impact

As shown in the previous section, the data-assisted approach can be used to generate insight which instructors can use to modify their teaching activity, and each of the examples given describes how broad changes in pedagogy that could be made based on understanding the hidden behaviours of learners. But the data-assisted approach is not limited to just instructors or broad changes; it is possible for instructional designers, for instance, to leverage similar techniques to make changes that affect only a portion of the learner population. This section[2] describes how learner behaviour data can be statistically described and related directly to educational outcomes. It further provides methods to model learners and associate them with pedagogically sound groups where more individualized interactions can take place. In doing this, this section addresses three questions that come from the motivational scenario in Sect. 7.2.3.2:

- Can learners be clustered based on their viewing habits into pedagogically relevant groups?
- If so, do these groups differ in their formal assessment measure?

7.3.3.1 The Recollect Environment

A key consideration of the data-assisted approach is whether learner interactions within the learning environment can be correlated with pedagogical goals and measures of learning outcomes. These interactions can be either explicitly made by learners (e.g. through the filling out of a survey) or implicitly made as a by-product of the learning activity itself (e.g. navigating through content). Sometimes referred to as *clickstream* data or *traces*, these interactions are difficult to understand on their own in part because of the large amount of data collected (potentially millions of data points) and the low level meaning that the data represents (e.g. the clicking of

[2] Portions of this section appear in Brooks et al. (2011a).

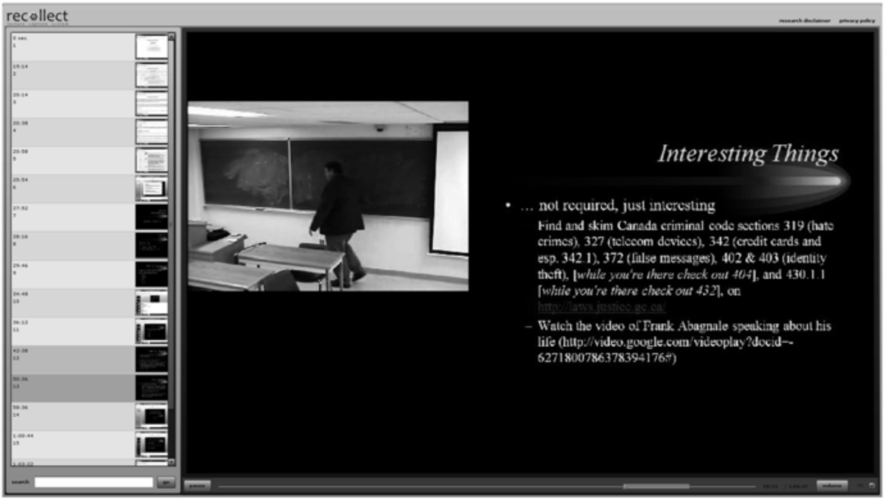

Fig. 7.8 A screenshot of the Recollect lecture capture system. The system shows thumbnails of upcoming slides on the *left* of the window to allow for navigation. The *right side* of the player shows video of the classroom as well as video captured from the data projector

a single link or keystroke on the keyboard). The data must first be summarized, and then linked to learner goals in order to be made actionable.

One learning environment that collects this low-level learner behaviour data is the Recollect lecture capture solution (Fig. 7.8), developed at the University of Saskatchewan in part by the authors. This system records in-person classroom lectures and stores them for playback by students for both the initial viewing of and reviewing of content. Recollect records a number of user behaviours, including the time spent streaming a lecture (discretized into 30 s intervals), clicks on any buttons in the user interface (e.g. volume change), searching through lecture slide content, seeking within the video using the video scrubber, and navigating within the video using section thumbnails. Each of these behaviours is linked to the time in which they were observed, the student who initiated the behaviour, and the particular video that was being watched.

7.3.3.2 Formal Assessment

The Recollect system was deployed for a number of sections of a second year Chemistry course in both the 2010 and 2011 academic years. Students were allowed to use the system how they saw fit, and every lecture from a single section taught by one professor was shared with students in all sections of the course. Instructors did not change grading criteria based on the presence of the recorded lectures, and mid-term and final examinations were common across all sections of the course.

Lecture capture is only one resource learners had available to them, and consistent patterns are difficult to see from the raw data. The viewing behaviours for Chemistry 2010 learners were cleaned[3] and summarized into viewing habits broken down by calendar week. In this model, learners were deemed to have either watched or not watched lecture content during the 12 weeks of the course. Out of 636 learners registered in the course, 133 were included in the study (participation rate of 36.6 %) by virtue of their use of the Recollect system. The weekly viewing rates of learners were used as attributes with k-means clustering ($k=5$) to create a model of learner behaviours. Five clusters were chosen based on preconceived hypotheses of how learners might use the system (see Brooks et al. 2011a for more details). The results of the clustering activity provide a number of insights into learner activities. For instance, some learners only watch lectures during the week of the midterm, while others watch fairly regularly. Regardless of viewing patterns, the last 2 weeks of the course (corresponding to the time between the end of classes and the final examination) tended to have a high amount of disagreement between participants and the centroids. The disagreement for these weeks, ranging from 19 to 40 %, suggests that activity throughout the term isn't indicative of behaviours between the end of term and the final exam, and thus only data during the teaching portion of the term was used for further analysis.

Using the data provided from the initial 2010 students, a high level model for five idealized clusters was developed. In this model, first cluster has learners who habitually watch lectures throughout the term (*high activity learners*), the second cluster is made up of learners who observed the lecture the week before the midterm examination (*just-in-time learners*), the third and fourth cluster appear to correspond to (roughly), the first and second half of the course (*disillusioned learners* and *deferred learners*) respectively, and the last cluster is made up of learners who did not watch many lectures, though they must have watched at least 5 min of video in a week to be included in the study (*minimal activity learners*).

With this high level model defined, a learner from any cohort can be placed into a particular group based on similarity. For instance, a learner who watches video every week except for the first and sixth weeks will be placed into the *high activity* cluster. Despite this learner not fitting perfectly with this cluster, his or her activity patterns are most closely related to it. Thus the centroids are not the only interesting aspects of the clusters, the amount of error is as well.

Instructional goals are often represented by midterm and final examinations as a proxy for learning. Correlating patterns of behaviours with differences in grades provides some evidence of learning from activity. Learners use lecture capture as one tool to aid in learning, but many other tools and methods contribute to learning (e.g. online quizzes, in-class lectures, textbooks, study groups) and make identifying the effect of any single tool difficult. A pairwise tukey test for the midterm

[3] A threshold of at least 5 min of viewing was arbitrarily chosen to remove behaviours that were deemed to be tool experimentation over tool use for learning. As the time period for this course was in the second semester of the academic year, the 1 week of data over midterm break was excluded from analysis.

Table 7.1 Midterm, final examination, and overall grade averages and standard deviations broken down by cluster in percentages for the Chemistry 2011 course

Cluster label	n	$\text{Midterm}_{\bar{\chi}}$	Midterm_{σ}	$\text{Final}_{\bar{\chi}}$	Final_{σ}	$\text{Overall}_{\bar{\chi}}$	Overall_{σ}
High activity	14	77.32	15.71	75.43	22.19	80.14	14.49
Disillusioned	18	64.71	15.72	58.98	17.81	66.82	13.85
Just-in-time	86	68.14	15.92	60.49	23.68	70.69	14.39
Minimal activity	191	64.30	15.36	58.98	22.89	68.41	14.71
Deferred	24	63.33	12.20	59.83	20.21	69.04	12.43

examination, final examination, and overall marks for the initial cohort demonstrated that there is an effect on marks for one cluster of learners in particular, the *high activity learners*, and that the effect's significance ranges between the levels of $p=0.021$ and $p=0.240$. Table 7.1 shows the difference in grades between the different learner clusters. Not shown in this table are the average incoming GPA values of each cluster, for which there was no statistically significant difference.

7.3.3.3 Learner Goals

Formal assessment is not the only indicator of learning, and a similar approach using machine learning allows one to form a relationship between the behaviours of learners and subjective questionnaire data. This can be useful in many ways—for instance, if a relationship is discovered between low use of lecture capture and negative opinions of the technology, the instructional expert may be able to change the learning environment in the future to accommodate learners who would prefer alternative tools based on their activity alone.

Armed with knowledge of the domain and pedagogy, domain experts like Michelle can query learning environments to gain a deeper understanding of how learners are acting within groups. Such queries are likely to be driven by hypotheses based on curiosity, preconceptions based on training, and instincts based on years of practice. To emulate this investigation, seven questions about the usefulness of the system and perceived workload were examined with respect to how well they fit clusters[4] ($k=2$) based on two behaviours: The number of minutes the learner watched and the number of unique videos the learner watched. The goal in doing this was to see if activity could be linked to statistically significant differences in learner opinions.

Learners ($n=636$) in the Chemistry 2011 cohort were surveyed as to the relevance and usefulness of the Recollect system in this class. The questions asked covered a mixture of technical, pedagogical, and policy issues, and a number of these questions were designed to elicit beliefs learners had about their learning (response rate of $n=229$, 30 %). The full survey instrument can be found in Brooks

[4] The choice of the number of clusters (i.e. the value of k) to make affects outcomes greatly. This was an initial investigation to determine if unsupervised machine learning approaches can be used for clustering of subjective responses to data. Given the results shown here it is reasonable to continue exploration with an aim to find ideal values for k.

Table 7.2 Student behaviour clusters based on the number of minutes watched and unique videos watched with $k = 2$

	Clusters				
	Keen ($n = 12$)		Less keen ($n = 115$)		
Attribute	$\bar{x}$	σ	$\bar{x}$	σ	ANOVA p
Minutes watched	1143.25	414.37	187.30	236.51	≤ 0.001
Unique videos	28.25	7.10	6.93	6.43	≤ 0.001
q4	1	1.04	1.16	1.01	0.61
q5	1.5	1.31	1.70	1.08	0.56
q6	0.59	0.80	0.79	0.77	0.04
q9	0.42	0.51	1.19	1.03	0.01
q21	0.25	0.87	0.44	0.97	0.51
q22	0.17	0.39	0.99	1.30	0.03
q24	0.50	0.80	1.25	1.28	0.05
Midterm	72.08	18.49	71.33	15.14	0.87
Final	72.22	21.15	70.41	20.00	0.77
Overall	77	15.32	75.75	13.15	0.76

Cluster labels, *keen* and *less keen* were added by the author as descriptive elements only, and are arbitrary. The only strong correlation with questionnaire data existed for questions 6, 9, 22, and 24

(2012). The seven questions looked at were all based on a five-point scale, with the first two being an increasing scale of marks, and the last five being a likert-based scale where zero represented the strong affirmative (e.g. "Very High" or "Very Important") and four representing the strong negative (e.g. "Very Low" or "Not Important"). The question text was:

- Q4: I am working in this class to try and get a mark in range:
- Q5: Reflecting on my performance in the class so far, I think my mark will actually be in the range:
- Q6: My workload this term including all of the courses I am in as well as other commitments is:
- Q9: How important do you feel that watching the recording of the lecture was for your success in this class?
- Q21: If you used the lecture capture system, how important was it for reviewing content you hadn't seen (e.g. missed classes).
- Q22: If you used the lecture capture system, how important was it for reviewing content you saw but didn't understand or couldn't remember?
- Q24: If you used the lecture capture system, how important was it for studying for examinations?

The analysis of these questions with the clusters formed is shown in Table 7.2. Labels on clusters were chosen by the authors, and learners were segmented into a smaller cluster ($n = 12$) of users who watched a large number of videos (on average, 28) for a mean time of 19 h and 3 min. A larger number of learners ($n = 115$) watched fewer videos (on average, 6) for a mean time of 3 h and 56 min. Only questions 6, 9, 22, and 24 showed statistically significant results ($p \leq 0.05$), though the means between clusters for question 6 were of little meaningful difference.

7.3.4 Case Study Three: Adapting Learning Environments to Tasks

The scenario presented in Sect. 7.2.3.3 followed an educational technologist, Adam, who used the data-assisted approach to understand how learners are using the lecture recording tools. More than just gaining *insight*, Adam was interested in building *instructional interventions* of an automated manner. Using clustering techniques like those described in the previous chapter, it is possible for an instructional expert like Adam to identify interesting groups of learners, and create an intervention. Thus far, however, only broad pedagogical interventions executed by instructors or instructional designers have been described. One of the interesting aspects of traditional intelligent learning environments (such as ITS) is that they respond automatically to learner actions. In this section[5] we consider whether this ability is lost in a data-assisted approach, where instructional experts are expected to be involved in the sensemaking process. In particular, we look at two questions:

- Do groups of learners really agree on where indices should be placed, or are their preferences for navigational aids more varied? If the former is true, then the clustering methods described previously may well yield a more personalized and efficient navigation structure.
- Is it appropriate to use supervised machine learning to build indices, and how might such an approach compare to algorithms that already exist?

7.3.4.1 Navigation in Recollect

The Recollect lecture capture environment (Fig. 7.9) has multiple methods a learner can employ to navigate through content. For instance, thumbnails across the left hand side of the environment allow for quick "chaptering" of the content with image preview, while the scrubber along the bottom allows for precise navigation throughout the video based on time. It is not unreasonable to think that navigational style might differ depending on learning goal; a learner watching a lecture for the first time might not use either of these navigational aids; a learner who is searching for a particular topic in the lecture might use the thumbnails provided; a third kind of learner might use the scrubber to quickly replay video about a critical concept they missed while watching.

Thumbnails in the Recollect environment were originally generated using a naive algorithm based on Time—every 5 min of video a still image would be copied from the video and metadata for the video would be updated linking the image and its position in the video. More sophisticated methods have been proposed for the same purpose; for instance, the Opencast Matterhorn system (Brooks et al. 2011b) uses a frame differencing algorithm with thresholds for RGB colour values, while

[5] Portions of this section appear in Brooks and Amundson (2009).

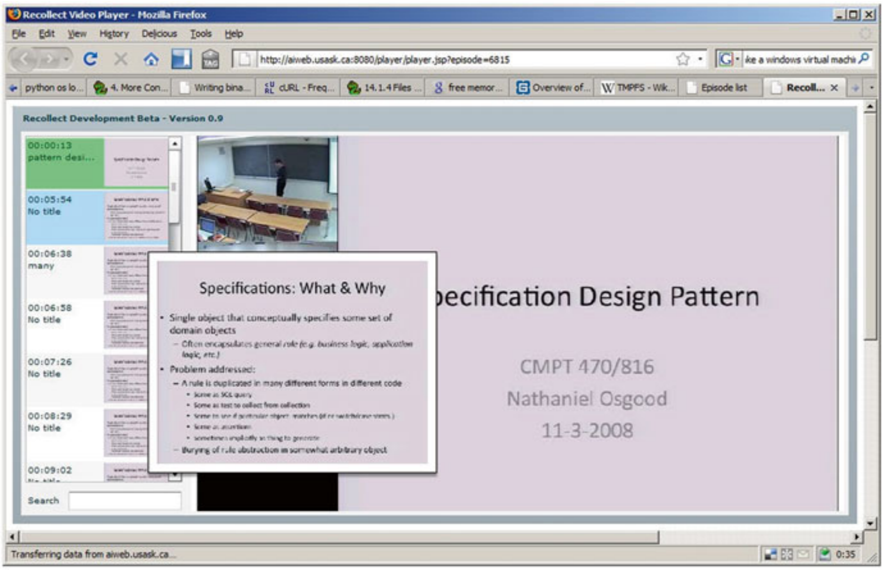

Fig. 7.9 The Recollect lecture capture system, showing navigational thumbnails on the *left hand side*. As users mouse over a given thumbnail, a small image opens up and shows what the data projector feed recorded at the corresponding time in the video. All interactions involving the thumbnails such as mousing over, clicking, or scrolling through the list, are recorded

Dickson's algorithm (Dickson et al. 2006) is a multi-pass image processing function that examines both pixel and block characteristics of video to determine stable events. Both of these algorithms were designed to work with lecture video captured by similar hardware as used by the Recollect system, making the potential for comparative study possible.

7.3.4.2 Comparing Users Actions to Traditional Algorithms

A laboratory study was undertaken to determine whether indices of video could be created for lecture video based on learner opinions of relevance. Indices are used by Recollect as thumbnails for navigating through a lecture, as shown in Fig. 7.9, and correspond roughly with DVD chaptering. If successful, such a method may be appropriate for generating indices in an ad hoc manner—an important result when applying the data-assisted approach.

Six human subjects who were unfamiliar with lecture capture systems were asked to go through four different lectures and identify where significant events occurred. The tool provided to study subjects allowed for navigating through a video linearly both forwards and backwards in one, five, 10, and 30 frame increments (each frame was equivalent to 1 s of video).

Observations and survey results from the participants identified that they used two distinct mechanisms for identifying significant events: visual structure (e.g. slide advancement in PowerPoint, or extending the canvas in the Sympodium) and semantics (topics being taught in the slides). Out of the six participants, five used primarily visual structure to identify events, while the sixth used the semantics of the lecture material. The level of agreement between participants excluding this sixth rater ranged from $\kappa = 0.18$ (slight, according to Landis and Koch 1977) to $\kappa = 0.87$ (almost perfect, according to Landis and Koch 1977). The videos that held traditional PowerPoint slides all had high inter-rater reliability ($\kappa \geq 0.66$).

This suggests that the previous methods of unsupervised machine learning we have discussed may be appropriate for generating groups of learners who navigate through lecture video similarly. This is useful for an educational technologist like Adam, who is seeking to modify the learning environment in order to improve indexing. Subsequent investigation as to the performance of indexing algorithms like Time, Opencast, and Dicksons's indicated they only poorly matched human raters. Instead, we look to leverage that rating data with supervised machine learning in order to personalize the indexing method. In a real-world environment, this rating data might come from logging data or social bookmarking behaviour of learners.

7.3.4.3 Adapting Navigation Based on Supervised Machine Learning

Most thumbnailing methods approach the issue of forming indices in lecture video as an image recognition problem. The goal of these methods is to measure the difference between two or more frames of video, and use this with some threshold value to determine when a significant change has occurred. The problem with this approach is in the selection of and weighting of attributes that make up the difference function; a data-assisted approach argues that the attributes should not be chosen a priori, but should be customized based on the learner (or cohort of learners) who are using the system.

Supervised machine learning methods take a set of instances, a set of attributes, and a set of classifications and build a model that can be used to predict new classifications for further instances with similar attributes. In the case described here, the set of instances are the video frames shown to subjects, the set of attributes are the image characteristics for these frames which are determined automatically (see Brooks and Amundson 2009; Brooks 2012) and the set of classifications is whether a given image is an index or is not. The output of a supervised method is a set of rules that can be applied to new images to determine if, based on this *training data*, those images are or are not indices.

Having determined that end-users are in a reasonable level of agreement when coming up with video indices, we ran a second study (Brooks et al. 2013) with six new participants in order to collect detailed indexing information. With this data, we formed 6 tenfold cross-validated J48 decisions trees. The trees were formed on modified versions of the training set, adjusting the threshold for the minimum agreement among raters before an instance was considered an index. The thresholds were

Table 7.3 Group κ between raters and algorithms

Comparison algorithms			Our trained algorithms					
Time	*Opencast*	*Dickson*	T_1	T_2	T_3	T_4	T_5	T_6
0.391	0.370	0.448	0.574	0.565	0.565	0.537	0.530	0.487

Upper and *lower* are the max and min values any algorithm could provide for κ. The κ between the six expert raters without an algorithm was 0.577

set such that T_1 was given a positive classification on all instances where at least one rater indicated there should be an index, T_2 is a tree trained on data where the threshold was two raters, and so on, until T_6 required perfect agreement between raters that an instance was an index. The goal was to see what effect different aggregation methods would have on the resulting level of agreement. As shown in 3, the trained algorithms all outperform the traditional methods of indexing, with the most stringent training method (T_6) providing the worst results ($\kappa = 0.487$) and the most lax training method (T_1) providing the best results ($\kappa = 0.574$). Further, each trained algorithm outperformed the static algorithms used for comparison (Table 7.3).

7.3.5 Conclusions

Previously (Sect. 7.2.3), we considered the needs of three different kinds of instructional experts; Katheryn the Instructor, Michelle the Instructional Designer, and Adam the Educational Technologist. Each of these experts is interested in gaining insight into the interactions learners have with technology, and leveraging this insight to create instructional interventions.

In this section, we have described three real-world educational systems that employ the data-assisted approach. In the first of these, human–computer interaction techniques of information visualization were used to aggregate traces of learner activities and make them available to instructional experts. Through augmenting an asynchronous discussion forum, instructors have been able to modify their pedagogical practice and gain insight into how learners interact in niche communities. This demonstrates how applying the data-assisted approach can lead to insights in instructional experts that they can use to modify their teaching practice.

The second system we looked at demonstrated that the traces learners leave behind when using a lecture capture and playback system can be data-mined and related to educational outcomes and goals. Clusters formed the basis for an abstract model, where each cluster represented different learning strategies. One group in particular, the *high activity learners*, correlated well with an increased achievement compared to other groups. With this knowledge, the instructor could apply clustering to future students and build instructional interventions aimed at particular groups. For instance, if it is the instructor's belief that the correlation relationship between regular lecture video watching and higher marks is a causal relationship, he or she might send out an alert to all learners who are not watching videos to encourage them to watch more consistently.

Continuing with a look at the Recollect environment, the third case study used a mixture of qualitative and quantitative laboratory studies to create a method for adapting the presentation of navigational indices in the user interface of a lecture capture system. By combining learner opinions of significance with supervised machine learning techniques, we demonstrated that substantially higher levels of accuracy of navigational indices can be achieved. Such a result demonstrates that there is value in using the data collected from learning environments to change the environment itself through instructional experts.

7.4 Discussion and Conclusions

Students learn better in more individualized tutoring situations (Bloom 1984), a result that has spawned two decades of intensive research in intelligent learning environments. These environments, such as ITS and adaptive hypermedia systems, deliver content to learners and form models of them based on the a priori definition of pedagogical approaches, learner traits, and content semantics. This allows for personalization of the learning environment which can be realized by changing the content, navigation, or structure of the learning environment for a particular group of learners.

This method of personalizing learning environments is expensive. It requires up-front cost in design and development and, as such, these methods are usually used within a single discipline or course. To scale across different domains, institutions of higher education use simplified learning content management systems. These systems offer instructors a thin technological wrapper around their existing content and learning activities, and provide only minimal support for personalization.

The data-assisted approach presented here supports learning in technology-enhanced learning environments by generating *insight* for instructional experts and enabling this insight to be used for *instructional interventions*. Instead of replacing instructional experts, the data-assisted approach enables them to see the hidden traces learners leave behind as they interact with the learning environments, to understand these traces in light of educational goals, and to apply this insight to form instructional interventions.

7.4.1 Discussion of Findings

There are many different kinds of instructional experts who might use data-assisted approaches: for example, instructors, instructional designers, tutorial assistants, and educational technologists. Each of these groups has different needs. For instance, an instructor might need to be able to understand what problems are faced by a particular cohort of students they are instructing, while an instructional designer or an educational technologist might want to generalize trends across cohorts and obtain insight about particular approaches or tools.

The data-assisted approach is broad enough to address these different cases. Section 7.3.2 described a situation where a variety of different instructors were shown visualizations of student interactions. These instructors taught different courses with different modalities (e.g. online versus blended in Sects. 7.3.2.3 and 7.3.2.4) and different scopes (e.g. communities of interest in Sect. 7.3.2.5). In each of these cases instructors were able to form insights into how learners in their course were interacting, and were able to use this insight to change their teaching practice.

Instructional designers and education researchers are also actors that can engage with data-assisted approaches. Section 7.3.3 describes the use of unsupervised machine learning methods to discover clusters of learners based on their lecture video viewing habits. These clusters correlate well with both pedagogical expectations and educational outcomes. By making visible the hidden viewing habits of learners using a lecture capture system, statements about the efficacy of lecture capture as a study aid can be made. For instance, the evidence that learners who watch lectures regularly have higher outgoing grades suggests that lecture capture may have an impact on learning, an important consideration when designing support for large cohort courses.

The data-assisted approach creates insight with instructional experts through dialogue. In Sect. 7.3.2 this dialogue takes the form of information visualization, and instructors could see different discussion forums in their courses at different times. In Sect. 7.3.3 this dialogue was more interactive, and allows experts to parameterize clustering and select attributes of interest. Regardless, it is the method of explicitly including the instructional expert in the sensemaking process that makes the data-assisted approach suitable for building intelligent educational environments in higher education.

Once insight has been formed, instructional experts need a way to improve learning through instructional interventions. Here again the different roles of experts change the way that instructional interventions are made. For instance, in Sect. 7.3.2 instructors largely developed interventions outside of the technology-based learning environment. One instructor, for instance, changed her assignment requirements which caused learners to interact differently—an interaction pattern she saw as more pedagogically sound. Another instructor used the insight generated from the data-assisted approach to reduce his level of interaction in the class based on a perception that the current discussion environment was already sustainable. In both of these scenarios, the instructor made these interventions based on visualizations resulting from the application of the data-assisted approach.

Instructional interventions are not always broad pedagogical changes, and software systems such as adaptive hypermedia systems often focus on small customizations to the learning environment to improve learner experience for individuals or groups of learners. Section 7.3.4 demonstrated that the data-assisted approach can be used to provide these forms of adaptations as well. Working from data representing ideal indices in lecture video, supervised machine learning was to be used to take prototypes of ideal indexing and apply these to different video content. Many different prototypes for various situations can be formed—those learners who want an overview might get one set of indices, while those who want visual navigation might get another.

Table 7.4 Outline of data, data-processing techniques, and the insight and instructional interventions they might lead to when using the data-assisted approach

Data	Data-processing technique	Insight	Intervention
Student reading data of asynchronous discussion messages	Sociogram-based information visualization	Discover community of practice and level of social engagement	Scaffold discussion (instructor), contribute to communities of interest (instructor)
Student viewing of lecture videos	Clustering of students with *k*-means and statistical treatment of assessment	Identify students with suboptimal study habits	Prompt change in study habits (instructor), recommend lecture video usage (system)
Student navigation in lecture videos	Clustering of students with *k*-means	Discover popular portions of lectures (instructional expert), identify segments of lecture video (system)	Review heavily studied concepts with extra material (instructor). Provide better indexing of video and adaptive navigation (system)

These prototypes do not have to be formed by an expert; instead, they can come directly from the learners themselves, and the expert (in this case an educational technologist) employs insight in designing a technology-enhanced learning environment.

7.4.1.1 Connecting Data to Insights and Insights to Interventions

While in this work we demonstrate several different insights and interventions, it is less clear exactly which data and data-processing techniques lead to these insights and interventions. This question is particularly salient in light of the software engineering task of building personalized learning environments—to form repeatable design patterns (“recipes” for successful software development) that can be used by software developers to build data-assisted software, these developers need to understand what data and data-processing techniques will provide instructional insight.

This issue is multidisciplinary in nature, and requires the consideration of education researchers, human–computer interaction designers, and information retrieval experts. Further, each set of data, data-processing technique, insight, and intervention can be considered at different levels of granularity, making the issue potentially more complex. Table 7.4 provides initial formulations of what such a taxonomy might look like, using the investigations provided in this chapter. The spirit of the data-assisted approach is that the intelligence of the system is the result of a dialogue between software and the instructional expert. In keeping with this, Table 7.4 should be seen as some general guidelines towards design patterns, and not an exhaustive list of which data and techniques lead to specific insights and interventions.

7.4.2 Conclusion

Unlike most other methods of building intelligent learning environments, the data-assisted approach does not seek to replace instructional experts but to actively engage with them. It does this in two ways; by generating *insight* from data through a dialogue with the expert, and supporting experts as they act on this insight to form *instructional interventions*. This allows institutions of higher education to leverage the intellectual support resources they already have (e.g. instructors, instructional designers, and educational technologies) to provide more personalized learning experiences in technology-enhanced environments.

References

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4–16.

Brooks, C. (2012). *A data-assisted approach to supporting instructional interventions in technology enhanced learning environments*. PhD thesis, University of Saskatchewan, Saskatoon.

Brooks, C., & Amundson, K. (2009). Detecting significant events in lecture video using supervised machine learning. In V. Dimitrova, R. Mizoguchi, B. du Boulay, & A. Graesser (Eds.), *2009 Conference on Artificial Intelligence in Education (AIED09)* (pp. 483–490), Brighton, UK. Amsterdam: IOS Press.

Brooks, C., Epp, C. D., Logan, G., & Greer, J. (2011). The who, what, when, and why of lecture capture. In *Proceedings of the First International Conference on Learning Analytics and Knowledge—LAK'11* (pp. 86–92). New York, NY: ACM Press.

Brooks, C., Greer, J., & Gutwin, C. (2012). Using an instructional expert to mediate the locus of control in adaptive E-learning systems. In *Second International Conference on Learning Analytics and Knowledge 2012 (LAK'12)* (pp. 84–87), Vancouver, BC.

Brooks, C., Johnston, G. S., Thompson, C., & Greer, J. (2013). Detecting and categorizing indices in lecture video using supervised machine learning. In *26th Canadian Conference on Artificial Intelligence*, Regina.

Brooks, C., McKenzie, A., Meyer, D., Moormann, M., Rihtar, M., Rolf, R., et al. (2011). OpenCast Matterhorn 1.1. In *Proceedings of the 19th ACM International Conference on Multimedia—MM'11* (pp. 703–706). New York, NY: ACM Press.

Brooks, C., Winter, M., Greer, J., & McCalla, G. (2004). The massive user modelling system (MUMS). In *Intelligent tutoring systems 2004 (ITS04)* (pp. 73–112). Maceio-Alagoas, Brazil: Springer.

Buskist, W. (2004). Ways of the master teacher. Association for Psychological Science: Observer. Available online at http://www.psychologicalscience.org/index.php/uncategorized/ways-of-the-master-teacher.html.

Carnegie Learning. (1998). *The cognitive tutor: Applying cognitive science to education. Technical report*. Pittsburgh, PA: Carnegie Learning, Inc.

Dickson, P., Adrion, W., & Hanson, A. (2006). Automatic capture of significant points in a computer based presentation. In *Proceedings of the eighth IEEE international symposium on multimedia* (pp. 921–926). San Diego, CA: IEEE Computer Society.

Greer, J., McCalla, G., & Vassileva, J. (2001). *Lessons learned in deploying a multi-agent learning support system: The I-Help experience*.

Klein, G., Moon, B., Hoffman, R. R., & Associates, K. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems, 21*(5), 88–92.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33*(1), 159–174.

McCalla, G. (2004). The ecological approach to the design of e-learning environments: Purpose-based capture and use of information about learners. *Journal of Interactive Media in Education, 2004*(May), 1–23.

Najjar, J., & Wolpers, M. (2006). Attention metadata: Collection and management. In *Workshop on logging traces of web activity: The mechanics of data collection*, Edinburgh, UK.

Peckham, T., & McCalla, G. (2012). Mining student behavior patterns in reading comprehension tasks. In *Proceedings of the Fifth International Conference on Educational Data Mining* (pp. 87–94), Chania, Greece.

Siemens, G. (2012). Sensemaking: Beyond analytics as a technical activity.

Suraweera, P., & Mitrovic, A. (2002). KERMIT: A constraint-based tutor for database modeling. In *Intelligent Tutoring Systems: Sixth International Conference, ITS 2002* (pp. 2012–216), 2–7 June, 2002. Biarritz, France: Springer.

Weskamp, M. (2003). *Social circles*. Retrieved March 2006, from http://www.marumushi.com/apps/socialcircle