Analyzing logfile data to produce navigation profiles of studying as self–regulated learning

Tina Leard & Allyson F. Hadwin
Concordia University, Montreal, QC, Canada
Centre for the Study of Learning and Performance

Paper presented at the Canadian Society for the Study of Education Annual Conference,
May 2001, Quebec City, QC.

DRAFT

Do not cite without permission of the authors.

Support for this research was provided by the Centre for the Study of Learning and Performance and Concordia University. Authors may be contacted at allysonh@education.concordia.ca and tina_leard@education.concordia.ca.

Rationale

Lifelong learners, whether inside or outside the classroom, self–regulate their own learning. Self–regulating learners strategically interact with tasks engaging cognitive, metacognitive, and motivational commitment and expertise (Schunk & Zimmerman, 1994). In addition to engaging a range of study tactics such as notetaking and rehearsing, self–regulating learning (SRL) involves orchestrating desired outcomes by metacognitively reflecting and adapting approaches to a given task. It is precisely this strategic application of skills, tactics, and reflections that form the foundation of strategic learning.

Understanding how the complex learning processes such as SRL unfold poses challenges for researchers. A review conducted with 5 graduate students examined 109 journal articles located using the search term "self–regulated learning." Fifty articles reported empirical studies only two of which explicitly examined aspects of the instructional design and its relationship to the development of SRL. The other 50 articles examined individual and motivational factors associated with SRL such as: (a) self–efficacy (Zimmerman, Bandura, Martinez–Pons, 1992), (b) other motivational factors and achievement outcomes (Pintrich & DeGroot, 1990), or (c) the impact of SRL upon achievement in various instructional contexts (Yang 1993; Young, 1996,). A predominant characteristic of all these studies was the examination of SRL as a product rather than a process. Self–regulation was typically represented as an outcome or something one might achieve in a developmental sequence. Despite the fact that most models of SRL (cf., McCombs & Marzano, 1990; Borkowski & Thorpe, 1994; Zimmerman & Schunk, 1997; Garcia & Pintrich, 1994) emphasize its iterative and recursive nature, empirical studies focus on examining how one achieves or reaches an outcome of being a self–regulated learner or achieving a state of self–regulatory skill. As a result, many of these studies measure SRL at one instance in time with questionnaires, self–reports, or interviews (e.g., Pintrich, Roeser, & De Groot, 1994; Yang, 1993). Examining SRL as an ongoing process that develops in sophistication across time requires that data be collected over time and across task contexts (Winne & Perry, 1999; Hadwin, 2000). Few studies have attempted to examine longitudinal traces of SRL.

Future examination of learning and SRL can be enhanced by collecting precise traces of student engagement with online materials (Barab , Bowdish, & Lawless, 1997; Rouet & Passerault, 1999; Winne , Gupta, & Nesbit, 1994). Traces are artifacts of every tactic and strategy used as well as detailed sub–components of those activities. Tracing is a non–intrusive way to collect information for research about learners as they learn. Other methods, such as students' self reports and interviews, require thinking about actions and thought processes in retrospect and are subject to memory decay. Similarly, think aloud protocols may interfere with the process of engaging with material. However, computers offer some possibilities for unobtrusively collecting traces of student activity through time stamped logfiles. Collecting precise traces of student activity in computer–supported learning environments portrays the dynamic, situated nature of learning, as well as individual differences in engagement (Marchionini, 1990; Winne, et al., 1994). Examining electronically generated patterns of tactics as they are strategically applied in time and content, has potential to inform our understandings of how students self–regulate learning. Analyses of this type of data may help guide students in their studying activities and advance adaptive learning environments toward that endeavor.

How can we collect these precise traces of studying activities?

CoNoteS (Winne, Hadwin & Field, 1997), and CoNoteS2 (Winne, Hadwin, McNamara, Chu, & Field, 1998) are prototype multi–component hypermedia tools developed using an authoring system called STUDY for building adaptive learning environments. CoNoteS2 is a sophisticated electronic notebook that incorporates electronic text and guides students in their notetaking and studying activities. It provides several studying tools including notetaking, glossary making, indexing, highlighting, and organizing tools. The interface can be adapted for any course or learning material. CoNotes is designed to support students to engage in 4 phases of studying as SRL: (a) analyze task requirements and resources, (b) set goals strategically, (c) monitor the implementation and utility of study tactics used to approach goals, and (d) metacognitively adapt studying methods (for a full description see Hadwin & Winne, in press). CoNoteS is programmed to collect data (*logfiles*) about human–computer interaction.

CoNoteS collects data about studying events by documenting: (a) all key strokes and menu selection events, (b) the timing of the event to the nearest second, (c) the type of event, (d) the textbook chapter and section from which events were initiated, and (e) the products of each study event such as the content of a note or glossary definition. Figure 1 illustrates a segment of one logfile. The first column indicates the time. In line 1, the student double clicked on a hyperlink to a section title in the organizer (index). This event initiated two computer initiated events, (Line 2) opening the electronic section of text in a new window and focusing (Line 3) that window by bringing it to the foreground. Two seconds after viewing the objectives, this student closed the window (Line 5). This was a short section of text containing three objectives. Each objective was presented in less than one line of text. Next the student opened a new section of text (chapter 3) (Line 6) from the organizer window. This section was titled “sex differences in the shadows.” The student created a new glossary by selecting the text “artifacts” in the electronic text and choosing “create new glossary” from a pop–up window (Line 9). The student titled this new glossary “artifacts” (Line 10). Once something had been selected from the pop–up window, the system focused the source page (the section titled “sex differences in the shadows”) until the new glossary opened and became active (Lines 12 and 13). When the student closed the glossary (Line 14), the system logged the contents in the glossary field (Line 17–18). This glossary note may have been made anytime after the glossary was opened but the contents were only logged at the time the glossary was closed.

Method

Participants

For this analytical exploration, we drew upon data collected in a prior study (Hadwin, Jamieson–Noel, McTavish, McNamara, & Winne, 2000). Participants were fifty undergraduate students enrolled in a 3rd–year course on instructional psychology. For this paper, we have selected data from two students to provide an illustrative case study of analysis choices. Students participated in this study as part of a course assignment where they were asked to apply knowledge about cognition, metacognition, and motivation to examine and reflect upon their own learning in this computer supported learning environment. Data were collected from students who consented to participate in the research component of this assignment. Fourteen students were deleted from the data set owing to missing data, and two students did not sign consent forms for the study.

Complexity of the hypermedia system used in this study

Using criteria presented in Leard & Hadwin (2001), we classified this hypermedia environment as a complex system because it included a nonlinear structure, as well as opportunities for high interactivity in the context of an ill–defined task. Following, we describe these features of the CoNoteS studying system.

We classified CoNoteS2 as a complex hypermedia structure. CoNotes2 affords opportunities for non–linear presentation of context and enables learners to use its tools flexibly. We wanted to control the order of presentation of these chapters because the second chapter provided some scaffolds designed to assist students in tuning their studying tactics. Using a hierarchical indexing tool, we controlled the presentation of chapters, so students initially progressed through chapters and sections in a prescribed order. However, once a chapter had been read, students could move freely between chapters and sections having multiple windows open simultaneously and accessing other chapters through hyperlinks in the index itself, other chapters, notes, glossaries and indexes.

The CoNoteS environment is designed to promote high interactivity. Students were provided with at least three studying tools that afforded opportunities for them to add to the hypermedia content. They could create notes and glossaries and link those entries directly to indexes, each other, and the source chapter from which the note was created. Students could use these tools wherever and however they chose. The only restriction was that students had to select a string of text and open a pop–up window in order to access those tools. Most students did this from the electronic textbook, but they could have also highlighted text in any of their notes or glossary pages. In studying, students created their own complex structure for organizing and accessing notes, which also added to the complexity of the hypermedia structure.

We constructed an ill–defined task for this study. This type of task contains abstractly defined goals and no–predefined response. Students were asked to study or learn the material in each chapter. They also knew they would be tested on this material after the completion of studying. Students were provided with an introductory set of objectives for each chapter. Objectives stated that students would be required to: (a) remember simple facts and definitions, (b) compare and contrast concepts and terms, and (c) abstract theories and hypothesize about concepts presented. Students had to read each section of the text (including the objectives) prior to having access to the following chapter. They had approximately one hour to study but did not have to complete any specific tasks during that study time.

Chapters

Participants studied three chapters of a text on sex differences. Each chapter was presented in five sections. The first section outlined three levels of objectives for studying: (1) remember facts, restate findings, or define or describe concepts, (2) create examples, explain relations, or compare/contrast concepts, and (3) analyze, summarize, or evaluate theories. Each chapter was approximately 1500 words in length. Chapters were balanced to present comparative numbers of terms, concepts, and comparisons.

Post–test

A post–test for each of the three chapters was created to measure students' achievement of each chapter objective. There were three sections to each test that corresponded to level 1, 2, and 3 objectives. The first section (8 marks) asked students to define four terms or concepts. The second section (4 marks) asked students to compare and contrast information, or ideas from the chapter. The third section (6 marks) asked to students to make an inference about the information in the chapter and explain their rationale of the inference. For example, students were asked to agree or disagree with a particular set of propositions from the text. Students were then asked to justify their position. This required students to understand and question the ideas and propositions of the text.

The post–test also included an efficacy measure. The first question asked students to indicate the amount of the chapter that they understood. The Likert 5–point scale answers ranged from almost none (0–10%) to almost all (90–100%). The second question asked students to indicate how well they would perform on the post–test. The 5–point scale answers ranged from poor (0–10%) to very well (90–100%). A score of 1–5 was given for each efficacy indicator.

Post–test Scoring

The first set of four questions was scored out of two marks for a total of eight. One mark was awarded for a partial definition and two marks were awarded for a full definition. The second set of questions were scored out of four marks: two were awarded for stating each of the two relevant positions, one mark was awarded for a statement of comparison, and one mark was awarded for making an inference about the comparison. The relational questions were worth six marks. Two marks were awarded for explaining the relevant propositions, one mark was given for a statement of agreement, disagreement, or prediction, and three marks were given for justifying their position. Answers were scored independently by two other researchers. Inter–rater reliability of the researchers' marking was $r=.92$, $p<.000$.

Procedure

A preparatory session introduced participants to CoNoteS2 and taught them to use its features to study. This was a guided session where two researchers modeled how to use the features of CoNoteS2 to study. In Session 1 (CHAPTER 1) (approx. 1 hr), students studied the first section of their text. After studying, they answered the post–test questions. First, students completed the efficacy questions. They were then given approximately 20–30 minutes to finish the test. Directions were: "Answer each question to the best of your ability. Answers do not need be long but should be complete and accurate. If you don't know an answer, don't guess wildly. But, if you believe you have some information to answer a question, partial credit will be given." A week later, the same procedure was followed for Session 2 (CHAPTER 2), and a week after that for Session 3(CHAPTER 3).

Results

Challenges to Analyzing Logfile Data in Complex Hypermedia Systems

Following, we introduce five challenges associated with coding, interpreting, and translating logfile data for various analysis programs and purposes. This not only provides context for understanding limitations in our analyses, but also guides the reader toward some issues that must be considered before collecting this type of data. While some of these challenges are specific to complex systems that include high interactivity, many are worth considering whenever logfile data are collected with the intent to examine human–computer interactions.

Since study strategies consist of complex combinations of study tactics applied purposefully to a given task, we were particularly interested in examining occurrences of, and patterns in, students use of studying tools provided in CoNoteS2. We began with raw data files formatted as text files much like the one shown in Figure 1. Separate logfiles were saved for each student per session; there were 150 logfiles altogether.

Challenge 1: Formatting logfiles for a variety of analysis tools

The first challenge we confronted was that the logfile format did not interface easily with the various analysis programs we had chosen. For NUD*IST, logfiles had to be in simple text format and each event had to be followed by a hard return. Conducting any search and replace functions had to be done in either Microsoft Word or Excel. In either case, this required importing each file separately and repeating the search and replace. Once the logfiles were condensed into one document they could not be fed back into NUD*IST as separate cases. Each case had to be saved separately, which was not an efficient use of time for a large number of logfiles. A further complication occurred when particular lines in the logfile were formatted differently than others. Referring to Figure 1, you can see that most events are represented by one line (or unit) of text. Lines 1 to 14 are each contained as separate units as indicated by a paragraph mark at the beginning and end of a line. Line 15 introduces an exception to this rule. Each note and glossary entry was recorded as multiple lines of text. Line 15 introduced the event. That is, a new definition was created. This was followed by a paragraph mark (¶), and a line of dashes followed by another ¶. Each line in the students note or glossary was also separated by a ¶. While this format makes it very easy to visually identify the contents of notes and glossaries, it poses problems for all of the analytical software tools we used. In all cases lines of text that are broken by ¶ are treated as separate units. This means that it is virtually impossible to automatically code everything between the dashed lines as one note. You can search for a line of dashes and then "hand code" the note, but this cannot be done using a search and replace function.

A second related problem was that there were some inconsistencies in the format of logfile output. For example, the line following demonstrates that some informational subunits were separated by semicolons, while others were separated by apostrophes that could indicate the beginning or ending of a string of text.

> 5:35:12 PM; Definition for glossary 'Chp3:Sex Differences in the Shadows:–(artifacts)' has changed. New definition:

We also detected some typographical inconsistencies in our logfiles. In the following logfile segment, notice that a space exists after “Chp3:” in Line 11, but not in Line 12.

> **LINE11** 5:34:54 PM; Focusing section 'Chp3: Sex Differences in the Shadows'¶
> **LINE12** 5:34:55 PM; Opening glossary 'Chp3:Sex Differences in the Shadows:–(artifacts)'¶

While these are all very minor errors to correct in the logging program, they each added significant numbers of hours to the analysis process. Further, these types of inconsistencies are not easily identified during electronic analysis. One must diligently compare electronic counts and coding to hand coding in order to detect these errors or inconsistencies. The only way to avoid such challenges is to decide upon analytical tools that will be used a priori, and then pilot test some sample logfiles with each of those tools. While this may seem evident, we introduce this issue in hope of encouraging researchers to articulate some standard conventions for logging that will afford opportunities for the exchange of data across platforms, and result in efficient means for analyzing large quantities of logfile data.

Challenge 2: Distinguishing between computer generated events and human initiated events.

A second challenge was distinguishing between computer generated and human initiated events. For example, revisiting figure 1, you will notice that the student initiated the creation of the glossary in Line 9, but it was not until Line 13 (four seconds later) that the glossary tool was the active screen. The student initiated the creation of a glossary in Line 9, but the computer initiated a string of responses in Lines 10–13. Further, if the student opened a glossary, but did not write anything in the glossary window, the text recorded between the dashed lines would be blank.

While this precise record of activities is important for reconstructing the student’s experiences using CoNoteS2 tools, it poses problems for the researcher interested in counting user–initiated events, or examining the timing and sequencing of those events. In the example provided here, what counts as a glossary entry? Do we count the number of times a student created a glossary

(Lines beginning with the text “Creating glossary”) or do we count the number of glossary fields that contain text (Lines 16 and 17). If we count the number of times a student filled in information between the dashed lines, what do we do when the student re–opened the same glossary and added some more information to it? Does that count as a new glossary, if not, how do we distinguish it from new glossary entries? Answers to these questions also influence analyses that focus on the precise timing or sequence of events. Examining duration of time between opening and closing this glossary note differs between 21 and 17 seconds. This depends upon whether the duration begins from the moment of intent (creating glossary in Line 9), or from the time the note was focused and available for the user to begin recording information (Line 13).

Challenge 3: Analyzing human–computer interactions as a complete system

CoNoteS2 logs every event to the nearest second. Qualitatively and quantitatively, this affords opportunities to examine patterns of interaction, keyboard entries, mouse movements, and window activation. A simple solution to the problems outlined in challenge 2 would be to parse down the logfiles into key user–initiated events such as “creating glossary”, “creating note”, etc. We argue that this is not a satisfactory solution, particularly when the focus of analysis is on sequences of events. One of the benefits of collecting such precise traces of human–computer interaction, is that the interplay between human cognition and learning context can be examined. Parsing down logfiles to key events, is much like conducting a content analysis on qualitative data. Events are grouped according to category and removed from the context that gives them meaning.

Examining actions as they arise in context is important because it provides invaluable information about how the learning environment may have influenced the user's actions. For example, if the student highlighted a string of text twice and then selected that same text to create a glossary note in session 1, it may indicate that the user was having trouble using the pop–up menu. The difficulty was that the student kept releasing the mouse on "highlight," the item preceding glossary in the pop–up menu. If we coded each action and then looked at those actions as a separate unit, we would only observe 2 repeating highlights.

Perhaps, more importantly, time cannot be accurately analyzed or interpreted if system events are excluded from logfile data. That is, the time between user action 1, and user action 2, may be accounted for almost completely by system events or delays. Without that information, one might mistakenly assume it took the user that amount of time to transition between actions. Sequential analysis is based on calculating probabilities of one action following another. In our system, some user and computer–initiated events are tightly linked with one another. For example, closing a glossary (Line 14) is always followed by the recording of whatever was in the glossary field at the time it was closed (Line 15 onward). Statistically and contextually, this sequence of events forms a repetitive pattern; however this pattern is not particularly meaningful because it does not describe a string of strategically planned user–initiated tactics. Rather, it describes the system reaction to a user–initiated event. Removing the system–initiated event, does not lead toward a more accurate interpretation. Removing system–initiated events listed in Figure 1, produces the logfile displayed in Figure 2. Patterns emerging from this analysis represent only user–initiated actions but ignore any interaction with the system itself. This means that if the system took a long time to focus the glossary window and thereby altered the students intended action, we no longer have record of that. Parsing down logfiles has the most serious implications for time–based sequential analysis. It is inaccurate in the following instance to suggest that it took 21 seconds to create this glossary entry because at least four of those seconds were taken up by slow computer response to a user's request.

The complexity of our logfiles represents the complexity of human–computer interaction using CoNoteS2. Reducing the complexity of the logfiles does not change the complexity of that interaction. Rather, it masks it.

Challenge 4: Overlapping events

The fourth challenge was closely tied to the previous one. CoNoteS, similar to a number of high complexity hypermedia structures, allows for multiple windows to be open simultaneously. Time spent focusing on a given window documents time on task. That is, when a student transitions between multiple windows the transition is logged as "Focusing." Determining the amount of time spent on any given event requires starting the timer from the time the window was activated (Focusing …) until it was either closed or another window was activated. This may mean adding multiple durations together to get total time on task. A difficulty arises because focusing is both a user–initiated event and a system–initiated event. Returning to figure 1, focusing occurs automatically when a new window is opened. However, in Figure 3 focusing occurs because the student is moving between a series of open windows. The bolded lines represent student–initiated events. Here the student moves back and forth between an open glossary and an open section of text. The time delay between focusing the glossary and returning to the section suggests that the student spent some time (1min 18 sec) writing a glossary entry.

Of course, the opening of multiple windows at the same time is further complicated because these windows could be scattered across the computer screen and visible simultaneously. The user could be consulting the contents in one window, but this consultation may not be logged because the user did not focus the window. A logging system that captures events and replays logs of these events would alleviate this problem.

Challenge 5: Deciding which line represents the event.

The final challenge was confronted when we tried to calculate the frequency of a set of user–initiated events. Calculating frequencies meant that we needed to decide which line in a log represented an activity (such as making a glossary). If we focused on the line indicating a note or glossary was opening, "Opening glossary," we were unable to distinguish between new glossary entries and additions to previous glossary entries. Focusing on the content of a note or glossary (the segment between the dashed lines in Figure 1) also posed problems because this information was logged after the window was closed. Since some users did not close all windows before exiting CoNoteS, valuable trace data concerning changing notes, glossaries, and glossary examples were lost. The best we could do was count the number of times that students initiated an activity such as notetaking or glossary making. This was indicated by a phrase such as "creating new glossary".

Furthermore, counts of stocked indexes in session 2 represent non–stocked indexes that were created. The use of stocked indexes could not be coded because the data recorded in the logfile did not distinguish between random clicking of the stocked indexes and actually use of the stocked indexes. By clicking on a stocked index, the string 'new index created' was recorded; it is unclear whether the user was just accessing this index or using it. Many students clicked on the stocked indexes several times in a row to possibly explore the index function. Although some students appeared to have linked the stocked indexes, the trace data are inconsistent. This situation is different from the recording of stocked glossaries and stocked notes because the creation of stocked glossaries and stocked notes was validated through the learner's entry of text.

A Case Study

Collecting precise traces of online activity is critical for investigating how students develop effective forms of SRL. However, the field lacks sophisticated methods for analyzing detailed traces of studying such as those recorded in CoNoteS, or web generated logfiles. A review of literature on logfile analysis (Leard & Hadwin, 2001) uncovered four ways to examine logfile data: (1) frequency counts of tools used, media type accessed, or information nodes followed, (2) total or average time, (3) patterns and sequences of events (4) content analysis.

Building on this review of literature we have applied these analyses to conduct a case study comparing two students' studying activities across three chapters of text. Specifically, we tried to develop understandings of how two students self–regulated across three study sessions. Comparative analyses across study sessions unfold in four sections. First, we compare frequencies of activities that unfold across chapters. Second, we compare duration of time spent engaging in a range of activities. Third, we examine patterns in sequence and extend this analytical technique beyond its application in previous empirical studies by presenting transition matrices of individual user's actions and by coupling patterns with durations through graphical representations; in previous studies, transition matrices were built on counts of a group's actions (i.e., Beasley & Waugh, 1997; Marchionini, 1989), and graphical representations were limited to illustrating patterns of activity and did not include durations (i.e., Horney & Anderson–Inman, 1994). Fourth, we examine contents of notes and glossary entries, coupling content analysis with durations to understand quality of engagement. For each analysis, we interpret or draw some conclusions about this student’s studying. In the final section we compare findings across the three approaches to analysis, illustrating how analytical triangulation is necessary for accurately interpreting this type of data. We follow with some recommendations for future studies.

Selection Strategy

For this study we selected 2 participants for in depth case study analysis. We used 2 criteria for selecting these participants. First they demonstrated large gains in their test scores from Session 1 to Session 3. As illustrated in Figure 4, both Alex and Sam scored less than one standard deviation below the mean on Figure 1 and more than one standard deviation above the mean in Test 3.

Second, Alex and Sam were active studiers; they used the study tools available to them as indicated by the frequency of notes, highlights, and glossaries recorded in their logfiles (see Figure 5). This was an important feature for this study because we were interested in exploring methods of logfile data analysis. The more active students were, the more complex logfile data was produced for analysis. Our goal in this analysis was to better understand how these students self–regulated to improve studying. Comparing against the mean, Figure 5 illustrates that Alex and Sam were not the most active studiers in this study; however with the exception of Chapter 1, these students engaged in an average number of highlighting, indexing, glossary making, and notetaking activities.

Examining the frequency of user–initiated events

Frequency counts provide information about distribution of actions across a spectrum of hyperlinked possibilities. Figure 3 illustrates the frequency with which Alex and Sam used various study tools available in CoNotes. Data illustrated in Figure 6 include raw counts of student engagement. Since we were primarily interested in changes between Chapter 1 and Chapter 3, we focus primarily on those chapters in this analysis. Our goal was to better understand how students self–regulated their engagement in study tactics to produce such improved performance on the test in Chapter 3. Overall, this analysis illustrates a trend for both Alex and Sam in decreasing the use of highlighting and increasing the use of both glossary making and notetaking from Chapter 1 to Chapter 3. For both Alex and Sam indexing was not a well used study tactic and this is consistent with other students profiles. This makes sense because indexing is not an activity that we have a lot of experience with in day to day studying. We use indexes in texts to locate things, but rarely do we have the opportunity or practice in creating our own indexes. Overall, Alex was a more active studier and created more notes and glossaries than Sam, with the exception of Chapter 3 when both students created a comparable number of notes.

Missing from the examination of frequencies is information for comparing the distribution of studying activities relative to other activities in that chapter. This is best examined, by calculating the percentage of total activities allotted to each of highlighting, indexing, glossary making, and notetaking. Pie charts comparing proportions of activity engagement for Alex and Sam are illustrated in Figure 7.

Unlike raw frequencies, examining the proportion of activities across chapters presents very different self–regulating profiles for Alex and Sam. In chapter 1, Alex's activities were fairly evenly distributed between glossary making (26 percent), and notetaking (31 percent), with highlighting being emphasized a little more than the others (43 percent). After completing Chapter 1, Alex adapted studying to eliminate highlighting altogether and instead emphasized notetaking (67 percent) followed by glossary making activities (33 percent). Alex then adapted this approach for chapter 3 where glossary making became a predominant activity (63 percent) and notetaking became a secondary tactic (37 percent). Presumably, the fine tuning of the balance between glossary making and notetaking tactics contributed to Alex's marked improvement on the chapter 3 test.

In contrast, Sam chose not to abandon highlighting after chapter 1. Rather, Sam increased the proportion of highlighting to 60 percent of the total activities. Sam also increased notetaking (25 percent) and proportionally decreased both indexing and glossary making activites to 5 percent and 10 percent respectively. This shift in strategies resulted in an improvement in chapter 2 test performance, but similar to Alex, it was not until Chapter 3 where Sam emphasized notetaking (54 percent) and glossary making (46 percent) that test score improved markedly.

Alex and Sam differed in their studying activities. Sam steadily increased notetaking and glossary making activities, whereas Alex seems to have experimented a bit more with different frequencies of both glossaries and notes. By chapter three, Alex had increased glossary making activities, and reduced notetaking to a more moderate level. This type of experimentation is indicative of self–regulation where students adapt tactics and strategies based upon feedback generated through metacognitive monitoring or externally (Winne, 1997). For these two students, glossary making and notetaking activities dominated their most successful study session. A limitation of this analysis is that it does not provide data with which to explore the similarities and differences in notetaking and glossary making activities. It is possible that both students focused on adapting studying to the chapter objectives, but Alex used the glossary making tool to do this, whereas Sam used the notetaking tool. Alternatively, it is possible that the increased proportion of notes and glossaries resulted because these students focused their study activities on a summary section of the chapter, rather than over emphasizing the body of the chapter. These are all strategies that undergraduate students use when trying to read and process large amounts of text, but we have no means with which to examine those aspects of studying when we limit ourselves to frequency and relative proportions of studying activities.

Duration of Studying Time and Activities

Examining the duration of time spent studying and devoted to each section of the text, provides some information about self–regulating because it illustrates how students alter their focus to particular section. Rather than distributing their time evenly throughout, these students emphasized various sections of the chapter and this changed across chapters.

Both Alex and Sam distributed their time across a spectrum of chapters (see Figure 8). Sam devoted more time each chapter, spending the longest amount of time studying chapter 3. Alex used time more efficiently across the chapters. In chapter 3, Alex engaged in the shortest total study time, compared to the other three chapters. However, the proportion of total time spent studying each section of the chapter was very similar for both Alex and Sam throughout. This is illustrated by comparing pie charts in Figure 8. In chapter 3, the last and most successful study session, both Alex and Sam emphasized the later three sections of text.

Analyzing the proportion of time allotted to any given chapter also has limitations. First, since the contents of each chapter differed, comparisons across chapters must be interpreted with caution. One would expect some variance in time spent on each section of a text depending upon the topic of that section. Even after controlling for section length and idea units, there are a number of confounding variables, such as prior knowledge, concept difficulty. Each of these contribute to the amount of time students spend on individual chapters. Second, the proportion of time spent on any given chapter cannot be directly compared across participants, particularly when there are large differences in the total time spent studying. For example, Sam spent 40 percent of her time on Chapter 1, Section 3, whereas Alex spent 23 percent of her time on the same chapter section. However, both students spent the same amount of time (approximately 12 minutes) on Chapter 1, Section 3. Sam spent considerably less time studying the whole chapter

(29 minutes, 52 seconds). She spent a considerable proportion of that time studying section 3, but it was an average amount of time in minutes and seconds.

Finally, examining duration of time and proportion of time spent studying, does not contribute to our understandings about how students engaged with the non–linear features of a hypermedia environment. It is possible that students did not take advantage of the fact that they could move back and forth between sections of the text. Nor does this assist us in examining how various studying tactics were distributed throughout the study period. It is possible that the timing of notetaking and glossary making activities contributed to improved performance in chapter 3. Duration of time on task does not provide the data with which to explore these issues.

Sequential analysis

One of the most valuable types of data recorded in logfiles, is detailed information about the sequence and context of all activities and actions. For example, our computer generated logfiles record the time of every event, the context or chapter where that action occurred and all events preceding and following. Changes and patterns in sequences of activity provide information about how students use the software. Sequencing of activities in logfiles reveals information about when students opened and closed sections of text, how they navigated between text and notes, and how they moved between various tools and windows that were available to them. Analyzing sequences of events is critical for understanding how students self–regulate their learning activities, but it also provides important information for usability testing. For example, repeated mouse clicks on window may indicate that the student is struggling to activate a new window, that one of the hotlinks is not working, or that the interface is not well designed.

Examining sequences is a challenging task, particularly when authentic logfiles are used. Our logfiles are messy. They contain logging errors and computer glitches, and vary in line structure in a few instances (indexing and highlighting in particular). These are all points mentioned earlier in this paper. We draw your attention to those challenges again because they had the most profound influence on sequential analysis with which we have been struggling for months, formatting and reformatting logfiles only to come up empty handed.

Sequential analysis depends upon identifying each possible event. In our logfiles there were system–initiated events, user–initiated events, and events that could be either user or system initiated (Focusing Section). Our first challenge was to decide whether it was meaningful to include all events in our sequential analyses. After pilot testing with one logfile, we decided that this was not a useful approach. System–initiated events occurred in sequenced patterns, were very frequent, and therefore dominated our analyses obscuring patterns in user–initiated events. Since user–initiated events were central to our research questions we decided to parse down our logfiles to better represent sequences and patterns in student activity. We also acknowledge the limitations in this type of approach. Human–computer interaction is a cybernetic system wherein the computer system, the user, and the context co–evolve. Each event influences and contextualizes another. Our solution, though necessary to illustrate the value of sequential analyses, is not satisfactory because it decontextualizes user–initiated events and potentially misrepresents their relationships with each other and with the system.

We used the Nud*Ist software to deconstruct and reconstruct our logfiles to better reflect user–initiated events. Nud*Ist is an excellent tool for this type of content analysis. It provides tools for searching strings of text and assigning codes. These searches can be operated on using a range of Boolean operators and Nud*Ist will even provide matrixes of the number of times one event follows another. We selected user–initiated events by including the onsets of events only. For example, we selected *Creating new glossary* rather than *Opening Glossary*. *Opening glossary*, and *Focusing glossary* were system responses to the student initiated event *Creating new glossary*. The time and the sequential placement of *Creating new glossary*, was a better indicator of student intent because it occurred when the student decided to make a new glossary. Since, *Focusing section* was both a user–initiated event and a system initiated event, we used Boolean searches to select only instances of *Focusing section* that were not immediately preceded by *Opening section*. This reduced our list of events to instances of *Focusing section* that were user–initiated. This event occurred when students mouse clicked between windows, changing the active window. Following is a list of user–initiated events included in our analyses: (a) *Focusing organizer window*, (b) *Focusing section* (user–initiated only), (c) *Indexing phrase*, (d) *Highlighting phrase*, (e) *Creating new note*, and (f) *Creating new glossary*. Missing from this list is *Opening new section*. Since *Opening section* was always followed by *Focusing section* and focusing section reflected times when that section of text was actually active, it better represented students' time on task.

<u>Observed Frequencies for Two–Event Sequences</u>.

The first step in examining sequences of activities is to collect raw frequencies of the number of times two event sequences occur. Tables 1 and 2 illustrate these frequencies for Alex in Chapter 1 and then again in Chapter 3. Although these tables illustrate the frequency of each two–event sequence, they are very difficult to interpret. The raw frequency only has meaning in relation to the number of event pairings that occurred. In Chapter 1, Alex engaged 40 two–event, user–initiated sequences. The highest frequency occurred when *Focusing section* was followed by *Creating new note*, or *Highlighting.* Nine, or eight instances of this sequence may be more meaningful when fewer event sequences occur and less meaningful when there are a large number of event sequences. However, event sequence frequencies can easily be transformed into simple probabilities such as those used earlier when we examined the frequency of single events and the duration of time spent on any given chapter.

<u>Simple and Transitional Probabilities</u>.

Simple probability is the probability of a target event occurring relative to the total number of events that occurred (Bakeman & Gottman, 1997). It is simply, the number of specific two event sequences divided by the total number of two–event sequences. Referring to Table 1, the simple probability of *Highlighting* being followed by *Highlighting* is 8 divided by the total number of event sequences (40), or 0.20. Tables 3 and 4 display the simple probability of each event sequence occurring in Chapter 1 and in Chapter 3. Each of these probabilities is relatively small. The largest is the probability of focusing glossary following focusing section. We could pursue these analyses further calculating transitional probabilities, but these analyses are not particularly revealing for the purposes of examining changes in studying across sections or across participants.

Transitional probabilities are useful descriptively, but they also have some limitations. They are not very useful for comparing students to each other or across chapters. This is because the probability score is meaningful in terms of the total amount of activity in a given session. If a student was particularly active in one session, then a probability score of .33 may not be very meaningful, but if the student was not very active, the same probability score of .33 may be very meaningful (Bakeman & Gottman, 1997). This makes these scores very difficult to interpret beyond the single case, single session descriptions above. Sequential analysis (analysis based upon transitional probabilities) offers some promise for aggregating across larger samples of logfiles and examining patterns of interaction. We recommend this as a future direction, for more large scales studies. For more explanation we recommend Bakeman and Gottman, 1997.

Timing and Sequencing of Events

Logfile data provide information about the amount of time students spend studying various sections of text, but it also includes information about the timing and sequencing of that studying in relation to various studying activities. Specific to this study, we wanted to examine how Alex and Sam changed the timing and sequence of studying activities across the three chapters. Figures 9 and 10 illustrate these relationships.

For both Alex and Sam, this analysis provides some interesting information about their self–regulation. In session 1, Alex opened each chapter sequentially and left all the chapters open until the end of the studying session. The navy blue bar at the bottom of each figure illustrates the beginning and end of the studying episode. In one instance, Alex moved forward and backward between chapter sections. Her first activity was in Section 3 (green) where she created three consecutive glossary notes. Alex then returned to section 1.2 (red) to create two glossaries and six notes. This was followed by a glossary in section 1.4 (purple) and a series of glossaries and notes in section 1.5 (pale blue). This is an interesting approach to studying. CoNoteS2 affords opportunities for students to move forward and backward between chapters, with one exception: each consecutive chapter can only be opened for the first time after the previous chapter has been opened. When each has been opened once, this feature of linearity disappears and students can re–open chapter sections in any order. Perhaps this wasn't clear to Alex, or perhaps she elected to take advantage of CoNoteS2's potential for viewing multiple chapter section on the same screen.

Sam approached chapter 1 a little differently, although she also left all the chapter sections open until the end of the session (navy blue). Sam opened the chapters sequentially and made her notes and glossaries in that order also.

The study session in Chapter 2 was structured differently. Students had more control over which chapter sections they opened and in what order they did so. Both Alex and Sam took advantage of this learner control. Alex closed sections when she was finished with them, and returned to them at various points in the study session. In most cases, she took notes and glossaries in the first reading, and presumably re–read or reviewed the second time the section was opened. However, on first reading, Alex followed the conventions prescribed in the previous study session by opening each chapter section in order.

Sam on the other hand, did much less reviewing than Alex in Chapter 2. Sam only reviewed section 1.3 (green). However, Sam strayed from the linear convention prescribed in her previous study session. On one occasion, Sam skipped section 1.4 (purple) and read section 1.5 (pale blue), returning to 1.4 later. Without more data this is difficult to interpret. It is possible that Sam clicked on the wrong section in the menu and did not realize until she had completed section 1.5.

The most interesting observation comes in Chapter 3. Alex adapted her studying pattern based upon her Chapter 2 experience. Constrained by CoNoteS2's linear functionality again, Alex, open chapters sequentially, took a series of glossaries and notes and then closed each chapter section before proceeding to the next. In contrast, Sam returned to her old practice of leaving all the sections open for the entire study session. Once the section was opened, it was left open. Sam did draw on her experiences in Chapter 2, however. Toward the end of her study session, she took advantage of the fact that she could review previous chapters and add notes and glossaries. In this example, Sam, took another note for section 1.2 (red) late in her studying session.

While this approach to examining logfile data reveals some interesting features about self–regulation, it is also limited in three ways. First, this figure does not illustrate the event *Focusing*. That is, it clearly indicates when sections were opened and when they were closed, but it does not illustrate when they were active. Second, it is difficult to represent a range of events in this graph without creating a visual maze. You will notice that we selected two events for this graph including notetaking and glossary making. We made this decision because these two events represent conceptually meaningful studying activities, and because we could not include more events without creating an illegible mess. Third, these graphs were painful to produce. To produce these graphs, we reformatted logfiles for excel, recalculated times so that we could determine chronological durations, and hand drew each notetaking or glossary making event on the excel–produced duration graphs.

<u>Quality of engagement</u>

To this point, we have examined the timing, frequency, and temporal context of notetaking and glossary making activities but we have not examined the contents of the notes themselves. CoNoteS2 logs the presence of notes and glossaries by recording: (a) the time when a student opened and closed that entry, (b) a reference to its source, and (c) the contents of the notes or glossaries themselves. The first two items are recorded regardless of the length, accuracy, or quality of the note itself. Therefore, our analyses thus far have not been sensitive to differences in quality.

Examining quality of note and glossary entries illustrates differences in Alex and Sam's approaches. Refinement in the depth of these notes across sessions may be indicative of SRL. This type of refinement was not evident in the quality of Alex's or Sam's notes. Alex consistently copied notes from the text. For the most part her notes were brief (1 or 2 lines of text), and with only one exception they were taken word for word from the original text. Sam, on the other hand, engaged a very different notetaking strategy. Over 90 percent of Sam's notes were written in her own words and on occasion included her own reflections and comments. This strategy continued across studying sessions, and was augmented in chapter 3 with the inclusion of examples for her glossary entries. Theory would predict that translating notes into ones own words facilitates recall of the material because it requires a deeper level of processing.

Comparing these findings regarding quality of notes to the examination of notetaking and glossary making frequency provides some context for interpreting these findings. Alex decreased notetaking activities from chapters 1 to 3 and increased glossary making. We might hypothesize that the types of notes Alex produced would not contribute to test performance, particularly on the second and third questions that required synthesis and reflection of material presented in the text. Increasing glossary making activities, would presumable contribute to increased performance on the first set of questions requiring the restatement of facts and definitions. Sam, on the other hand, increased notetaking and glossary making activities and with that increased attempts to synthesize and reflect on material presented in the text. One would assume this type of activity would contribute to increased performance on test questions that required deeper processing. Examining the relationship between note and glossary quality and test performance on specific questions warrants further investigation.

Examining the quality of notes, is essential for understanding how students self–regulate across chapters. Counting the frequency of studying events, may not provide an accurate representation of “engagement” with hypertext material. However, quality of activities is very difficult to interpret on its own. In this analysis, it became important to juxtapose findings regarding frequency with analysis of depth of processing (or quality).

Discussion

How did students SRL?

Alex and Sam self–regulated their learning in very different ways. One of the advantages of examining navigation profiles such as those provided above is that they provide a breadth and depth of analysis, with which to examine SRL in action. Alex and Sam were similar in that they decreased the amount of highlighting they engaged and increased both notetaking and glossary making activities. A difference however, was that Sam’s notes and glossaries were higher in quality. Sam synthesized ideas presented in the text and added some of her own reflections and comments. In comparison, Alex copied ideas from the text. An interesting follow–up analysis would be to examine and compare each of their tests to see where strengths and weaknesses in recall appeared.

Alex emphasized glossaries over notes in session 3 but both were predominant activities in that study session. Sam emphasized notetaking but also focused primarily on these two study activities in session 3. One of the ways, Sam’s studying changed is that she spent more time studying overall, whereas Alex seems to have decreased studying time while increasing the efficiency of her use of time. Both students emphasized Section 4 of Chapter 3. It may be worth examining the contents of that chapter section.

Examining the timing and sequencing of events illustrated unique ways that the CoNoteS2 environment constrained studying as well as differences in how students adapted to that environment. In session 2 when the chapter sections did not have to be accessed in a linear fashion, both Alex and Sam took advantage of that flexibility and accessed chapters in a non–linear order (on at least one occasion). When the environment had a linear component in Chapters 1 and 3, students maintained a linear approach, even though they could access chapter sections in any order once they had been opened the first time. There was some carry over from Chapter 2 to Chapter 3 for both Sam and Alex but it was different in nature. Alex began to close the chapter section after completing it rather than leaving all sections open throughout the study session. Sam on the other hand, resumed her previous strategy of leaving all chapter sections opened, but took advantage of the fact that she could go back and take notes in previous sections. While these activities are difficult to interpret on their own, they do warrant follow–up studies. It would be particularly revealing to augment this logfile data with student interviews so that we could better understand the interplay between student goals and activity.

It is important to note, that no one analysis could have produced these kinds of understandings of student’s SRL. By producing navigational profiles that examine a range of features of engagement, we were able to abstract rich understandings of both Alex and Sam’s approaches to studying in this hypertext context.

Future Directions and Implications for Logfile Analysis

Examining authentic logfile data is a messy task. We began these analyses with visions of much more complex analyses than those presented here. The challenges we confronted with logfile inconsistencies and format limited analytical directions that could be pursued. The fact remains that we could not efficiently apply analyses to large numbers of logfiles without encountering some kind of logfile error or inconsistencies. Logfiles had to be reformatted for each and every analysis procedure. As a result we conclude with some recommendations for logfile formatting as well as some future directions for large–scale analysis.

First, we recommend developing some logging conventions or standard formats. Our analyses revealed some important conventions such as the following:

(a) Consider all information between the first ¶ and the second ¶ to be one unit of text.

(b) Separate information provided within each unit of text with a symbol, tab mark, or semicolon.

(c) Assign different identifiers to system initiated versus user–initiated events

(d) Record time chronologically starting at 0:00:00 for each session

(e) Avoid recording multiple lines of log for the same event

(f) Use symbols to identify chapters, sections, units, or pages

We also recommend four directions for future research and development.

1. We need to develop some logging conventions (formats).

2. We need to develop some analytical tools for transitioning between logfiles and other programs. We used word, NUD*IST, excel, and SPSS, and re–formatted logfiles for over six months to transition back and forth.

3. We need to identify key indicators of events and transitions. Currently in connotes these are completely different event lines. We suspect this is often the case when logfiles capture the complexity of interaction.

4. Despite the temptation, we do not believe the answer is to parse down logfiles into a simpler form. Rather, we need to develop more sophisticated means for examining complexity of events and activities.

We intend to pursue two further directions in terms of analytical techniques. Bakeman and Gottman (1994) have applied statistical sequential analysis to observational data. They provide some techniques for aggregating across large samples of observational data. Many of the analyses that they recommend are suited to sequential strings of unique, non–overlapping events, however they do provide some recommendations for overlapping events as well as time series analyses. While these types of analysis have primarily been applied to less complex, less precise, and less thorough forms of observational trace data, they do warrant exploration.

Second, suggestions for establishing patterns of activity have been made on the basis of graph theoretical methods (Guzdial, Berger, et al., 1995; Nigemann, 2000; Winne, Gupta et al., 1994) and multidimensional scaling analysis (Guzdial, Berger, et al., 1995; Berger & Jones, 1995). Electronic logfile analysis tools has played a significant role in realizing the suggested analytical techniques for establishing patterns of activity by taking advantage of complex algorithms to search through large logfiles. Graph theoretical methods entail performing analyses on data from transition matrices. Winne, Gupta, et al. (1994) developed three graph theoretic measures for analyzing event sequences through their software logfile analysis tool LogMill (Nesbit & Winne, 1994): 1) density, a measure of regularity and structure in the sequence; 2) S*, a measure of similarity between two sequences; and 3) structrual equivalence, a measure that gauges whether these events are used interchangeabley or equivalently. Similar to performing graph theoretical methods, multi–dimensional techniques also begin with a transition matrix; however, computations on the matrix data differ. Berger and Jones' (1996) software logfile analysis tool Event Recorder supports these computations, which include building a matrix of C correlations to input to a multidimensional scaling algorithm for cluster analysis.

A more recent method of establishing patterns of activity is being made possible through electronic logfile analysis tools (e.g., Jones & Jones, 1997; Okada & Asahi, 1998). These tools support pattern–based sequential analysis and enable searching for specific patterns defined by the researcher. Queries for patterns that are longer than two events can be made, which gives this technique an advantage over analyses based on transition matrices. An example of this type of tool is Jones and Jones' (1997) MacSQEAL. MacSQEAL allows the researcher to transform logfile data by converting events into more abstract representations, recoding events for selective distinctions, and segmenting the data into sequences. MacSQEAL queries generate a list of items specified to display matching lines and the frequency count for the queried pattern. Similar to frequencies of two–event transitions derived from transition matrices, these frequencies have potential to reveal dominant patterns for one particular participant and groups of participants by applying multi–dimensional scaling or graph theoretical methods.

Although multi–dimensional scaling and graph theoretical methods have worked well with test case data, they have yet to be applied to large quantities of logfile data. Further, the application of these techniques has not yet extended beyond two–event transitions. To pursue this analysis, we need to clean up our logfiles and flag all inconsistencies. This is currently underway and these analyses are forthcoming.

References

Bakeman, R., & Gottman, J. M. (1997). Observing interaction: An introduction to sequential analysis. New York, NY: Cambridge.

Barab, S. A., Bowdish, B. E., & Lawless, K. A. (1997). Hypermedia navigation: Profiles of hypermedia users. *ETR&D, 45*, 23–41.

Beasley, R. E., & Waugh, M. L. (1997). Predominant initial and review patterns of navigation in a fully constrained hypermedia hierarchy: An empirical study. *Journal of Educational Multimedia and Hypermedia, 6* (2), 155–172.

Berger, C., & Jones, T. (1995). Analyzing sequence files of instructional events using multiple representations. Paper presented at the Annual Meeting of the American Educational Research Association, San Francisco, CA.

Borkowski, J. G., & Thrope, P. K. (1994). Self–regulation and motivation: A lifespan perspective on under achievement. In D. H. Schunk, & B. J. Zimmerman (Eds.), *Self–regulation of learning and performance: Issues and educational implications* (pp. 45–73). Hillsdale, NJ: Lawrence Erlbaum.

Guzdial, M., Berger, C., Jones, T., Horney, M., Anderson–Inmann, L., Winne, P., & Nesbit, J. (unpublished manuscript). Analyzing student use of educational software with event recordings. Unpublished manuscript based upon a symposium presented at the American Educational Research Association annual meeting, April, 1995.

Hadwin, A. F. (2000). *Building a case for self–regulating as a socially constructed phenomenon.* Unpublished doctoral dissertation. Simon Fraser University, Burnaby, BC, Canada.

Hadwin, A. F., & Leard, T. (2001, April). Navigation profiles: Self–regulating learning examined through 5 analytical representations of logfile data. In A. F. Hadwin (organizer). *Logfile navigation profiles and analysis: Methods for tracking and examining hypermedia navigation.* Symposium to be presented the Annual Meeting of the American Educational Research Association, Seattle, WA.

Hadwin, A. F., Jamieson–Noel, D. L., McTavish, R., McNamara, J. K., & Winne, P. (March, 2000). *Designing Courses to Support Self–Regulated Learning: Teacher as Researcher*. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA.

Horney, M. A., & Anderson–Inman, L. (1994). The ElectroText project: hypertext reading patterns of middle school students. *Journal of Educational Multimedia and Hypermedia, 3*, 71–91.

Jones, T., & Jones, M. (1997). MacSQEAL: A tool for exploration of hypermedia log file sequences. In T. Müldner & T. C. Reeves, *Proceedings of Ed–Media 1997* (pp. 709–716). Charlottesville, VA: AACE.

Leard, T., & Hadwin, A. F. (2001, April). Logfile analysis: A review of the literature. In A. F. Hadwin (organizer). *Logfile navigation profiles and analysis: Methods for tracking and examining hypermedia navigation*. Symposium presented at the Annual Meeting of the American Educational Research Association, Seattle, WA.

Marchionini, G. (1990). Evaluating hypermedia–based learning. In D. H. Jonassen & H. Mandl (Eds.), *Designing hypermedia for learning* (pp. 355–372). Berlin Heidelberg: Springer–Verlag.

McCombs, B. L., & Marzano, R. J. (1990). Putting the self in self–regulated learning: The self as agent in integrating will and skill. *Educational Psychologist, 25,* 51–69.

Niegemann, H. M. (2000). Analyzing processes of self–regulated hypermedia–supported learning: On the development of a log–file analysis procedure. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA.

Nesbit, J. C., & Winne, P. H. (1994). *LogMill* [computer program]. Burnaby, BC: Simon Fraser University.

Okada, H., & Asahi, T. (1999). GUITESTER: A log–based usability testing tool for graphical user interfaces. IEICE Transactions Vol. E82, 6, 1030–1041.

Pintrich, P. R., & DeGroot, E. V. (1990). Motivational and self–regulated learning components of classroom academic performance. *Journal of Educational Psychology, 82,* 33–40.

Pintrich R. P., Roeser W. R., & De Groot M. A. (1994). Classroom and individual differences in early adolescentsí motivation and self-regulated learning. *Journal of Early Adolescence* (14), 2, 139-161.

Reed, W. M., & Oughton, J. M. (1997). Computer experience and interval–based hypermedia navigation. *Journal of Research on Computing in Education, 30,* 38–52.

Rouet, J. –F., & Passerault, J. –M. (1999). Analyzing learner–hypermedia interaction: An overview of online methods. *Instructional Science, 27,* 201–219.

Schunk, D. H. (1994). Self–regulation of self–efficacy and attributions in academic settings. In D. H. Schunk, & B. J. Zimmerman (Eds.), *Self–regulation of learning and performance: Issues and educational implications* (pp. 75–100). Hillsdale, NJ: Lawrence Erlbaum.

Schunk, D. H., & Zimmerman, B. J. (1997). Social origins of self–regulatory competence. *Educational Psychologist, 32,* 195–208.

Winne, P. H. (1997). Experimenting to bootstrap self–regulated learning. *Journal of Educational Psychology, 89,* 397–410.

Winne, P. H., & Hadwin, A. F. (1998). Studying as self–regulated engagement in learning. In D. Hacker, J. Dunlosky, & A. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 277–304). Hillsdale, NJ: Lawrence Erlbaum.

Winne, P. H., & Perry, N. E. (1999). Measuring self–regulated learning. In M. Boekaerts, P. Pintrich, & M. Zeidner (Eds.), *Handbook of self–regulation* (pp. 531–566). Orlando, FL: Academic Press.

Winne, P. H., Gupta, L., & Nesbit, J. C. (1994). Exploring individual differences in studying strategies using graph theoretic statistics. *The Alberta Journal of Educational Research, XL*, 177–193.

Winne, P. H., Gupta, L., & Nesbit, J. C. (1994). Exploring individual differences in studying strategies using graph theoretic statistics. *The Alberta Journal of Educational Research, XL*, 177–193.

Winne, P. H., Hadwin, A. F., & Field, D. (1997). *CoNoteS: An electronic notebook with support for self–regulation and learning tactics* [computer program]. Simon Fraser University, Burnaby, BC.

Winne, P. H., Hadwin, A. F., McNamara, J. K., Chu, S., & Field, D. (1998). *CoNoteS2: An electronic notebook with support for self–regulation and learning tactics* [computer program]. Simon Fraser University, Burnaby, BC.

Winne, P. H., Hadwin, A. F., McNamara, J., & Chu, S. (1998, April). *An exploratory study of self–regulating learning when students study using CoNoteS2*. Paper presented at the Annual meeting of the American Educational Research Association, San Diego, CA.

Winne, P. H., Hadwin, A. F., McNamara, J., & Chu, S. (1998, April). *An exploratory study of self–regulating learning when students study using CoNoteS2*. Paper presented at the Annual meeting of the American Educational Research Association, San Diego, CA.

Yang, Y. C. (1993). The effects of self–regulatory skills and type of instructional control in learning from computer based instruction. International journal of instructional media, 20, 225 –40.

Young, J. D. (1996). The effect of self–regulated learning strategies on performance in learner controlled computer–based instruction. ETR&D, 44 (2), 17–29.

Zimmerman,–B., Bandura,A., Martinez–Pons,M. (1992). Self–motivation for academic attainment: The role of self–efficacy beliefs and personal goal setting. *American–Educational–Research–Journal, 56,* 663–676.

Figures

Figure 1. Logfile output.

```
LINE1 5:34:02 PM; Double clicked section title 'Objectives' in organizer¶
LINE2 5:34:02 PM; Opening section 'Chp3: Objectives'¶
LINE3 5:34:02 PM; Focusing section 'Chp3: Objectives'¶
LINE4 5:34:08 PM; Focusing organizer window¶
LINE5 5:34:08 PM; Closing section 'Chp3: Objectives'¶
LINE6 5:34:10 PM; Double clicked section title 'Sex Differences in the
      Shadows' in organizer¶
LINE7 5:34:10 PM; Opening section 'Chp3: Sex Differences in the Shadows'¶
LINE8 5:34:10 PM; Focusing section 'Chp3: Sex Differences in the Shadows'¶
LINE9 5:34:51 PM; Creating new glossary for phrase 'artifacts' in section
      Chp3: Sex Differences in the Shadows' via popup menu¶
LINE10 5:34:54 PM; New glossary entitled 'artifacts'¶
LINE11 5:34:54 PM; Focusing section 'Chp3: Sex Differences in the Shadows'¶
LINE12 5:34:55 PM; Opening glossary 'Chp3:Sex Differences in the Shadows:-
      (artifacts)'¶
LINE13 5:34:55 PM; Focusing glossary 'Chp3:Sex Differences in the Shadows:-
      (artifacts)'¶
LINE14 5:35:12 PM; Closing glossary 'Chp3:Sex Differences in the Shadows:-
      (artifacts)'¶
LINE15 5:35:12 PM; Definition for glossary 'Chp3:Sex Differences in the
      Shadows:-(artifacts)' has changed. New definition: ¶
LINE16-------------------------------------¶
LINE17 human creations, ¶
LINE18i.e. legal documents, cultural myths¶
LINE19-------------------------------------¶
```

Figure 2. Parsed down logfile

```
5:34:02 PM; Double clicked section title 'Objectives' in
organizer
5:34:08 PM; Closing section 'Chp3: Objectives'
5:34:10 PM; Double clicked section title 'Sex Differences in
the Shadows' in organizer
5:34:51 PM; Creating new glossary for phrase 'artifacts' in
section 'Chp3: Sex Differences in the Shadows' via popup menu
5:34:54 PM; New glossary entitled 'artifacts'
5:35:12 PM; Closing glossary 'Chp3:Sex Differences in the Shadows:-
(artifacts) '
```

Figure 3. System versus student initiated events

7:46:44 PM; Focusing section 'Chp2: Self Regulation'
7:47:35 PM; Focusing organizer window
7:47:36 PM; Double clicked glossary title 'Conceptual tempo' in organizer
7:47:36 PM; Focusing section 'Chp2: Self Regulation'
7:47:37 PM; Opening glossary 'Chp2:Self Regulation:-(Conceptual tempo)'
7:47:38 PM; Focusing glossary 'Chp2:Self Regulation:-(Conceptual tempo)'
7:47:54 PM; Focusing section 'Chp2: Self Regulation'
7:47:57 PM; Focusing glossary 'Chp2:Self Regulation:-(Conceptual tempo)'
7:49:45 PM; Focusing section 'Chp2: Self Regulation'
7:49:47 PM; Focusing glossary 'Chp2:Self Regulation:-(Conceptual tempo)'

Figure 4. Means and standard deviations for test scores

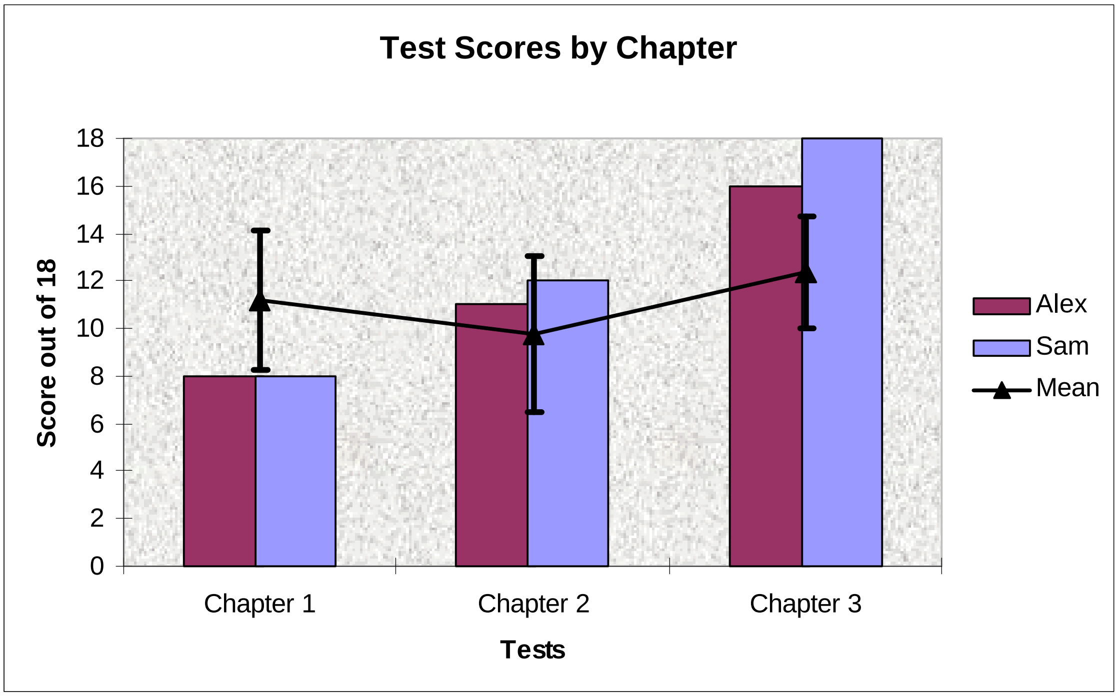

Figure 5. Frequency of total activities compared to group mean and standard deviation

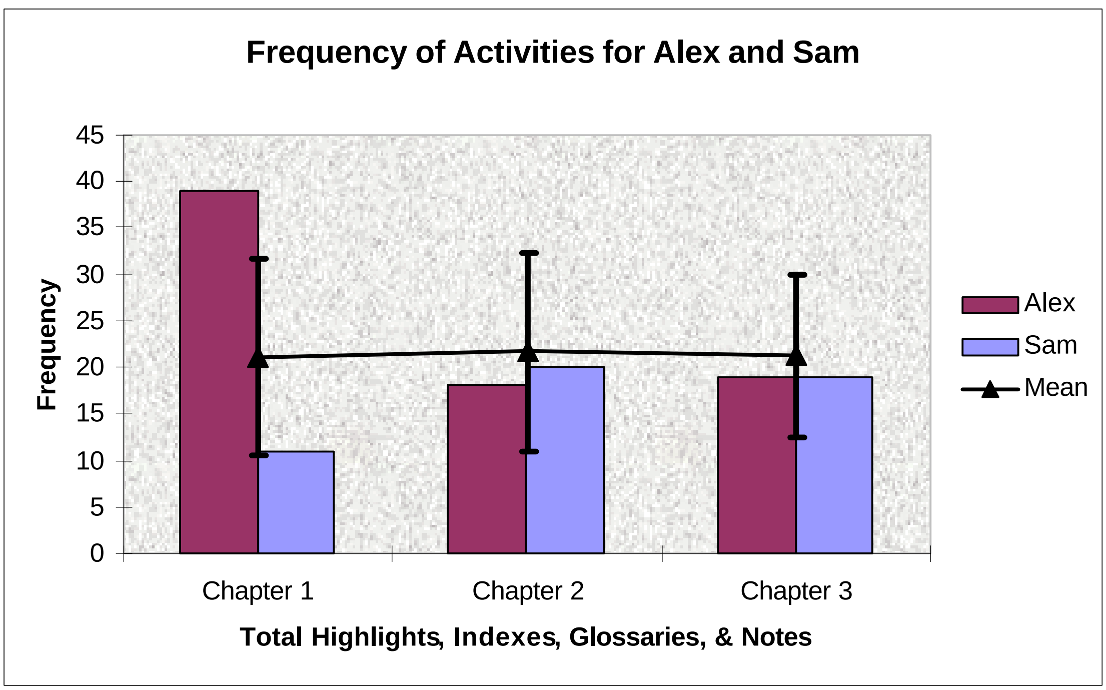

Figure 6. Comparison of frequencies of highlighting, indexing, glossary making, and highlighting between Alex, Sam, and the class mean.

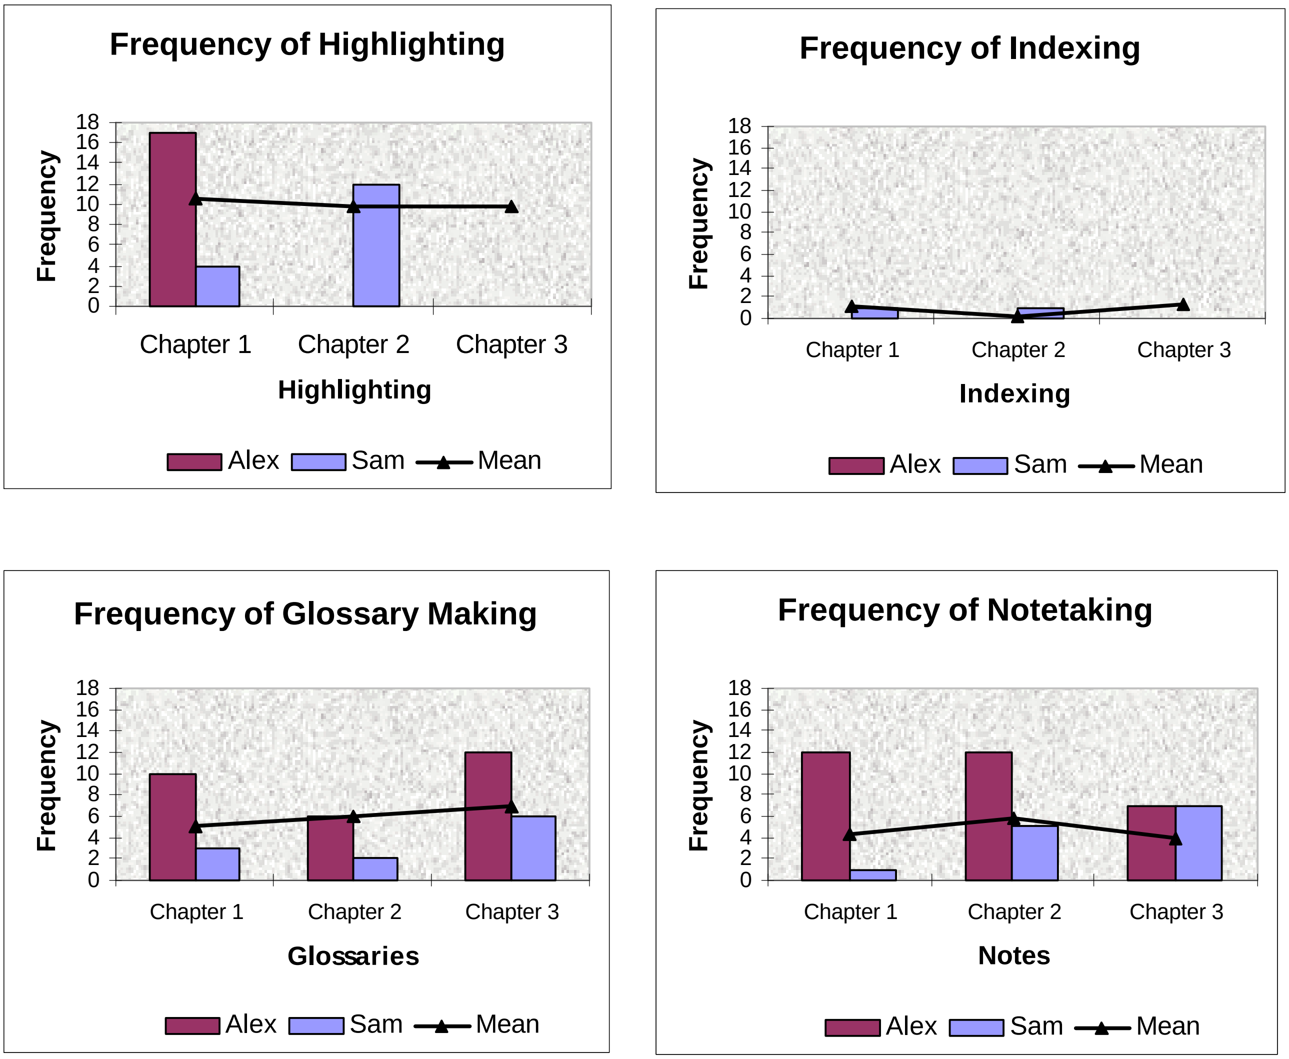

Figure 7. Comparison of proportion of overall activities.

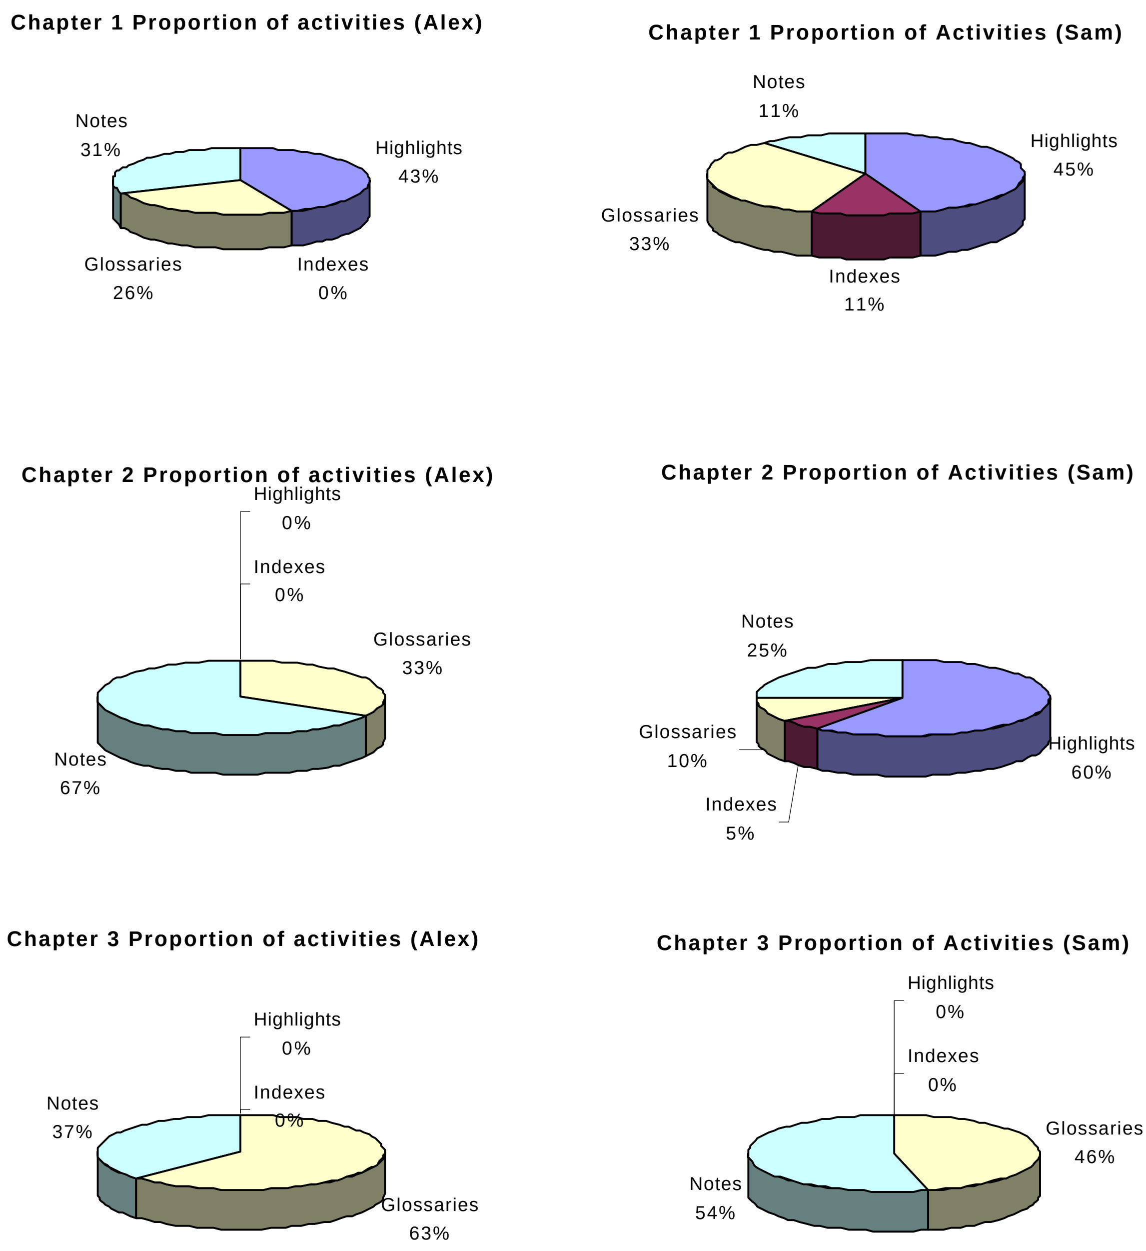

Figure 8. Proportion of total time spent engaging activities related to each chapter section

Chapter 1: Time spent (Sam)

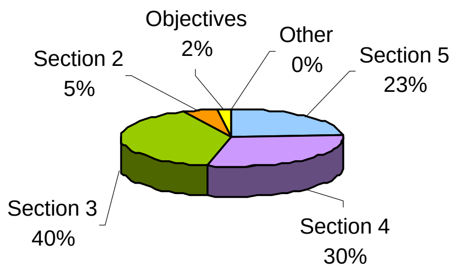

Chapter 1: Time spent (Alex)

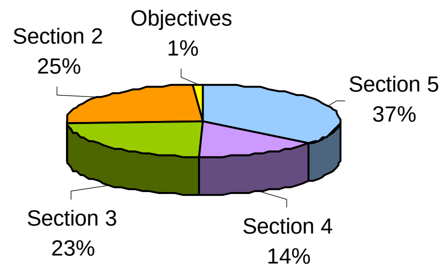

Chapter 2: Time spent (Sam)

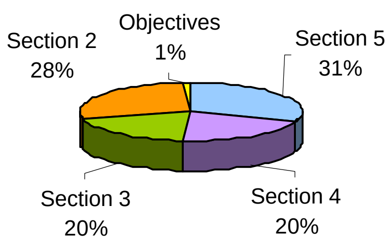

Chapter 2: Time spent (Alex)

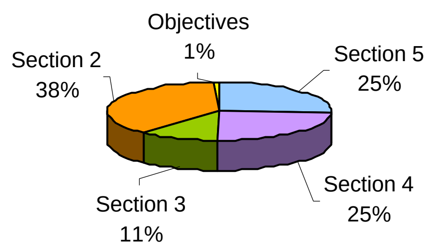

Chapter 3: Time spent (Sam)

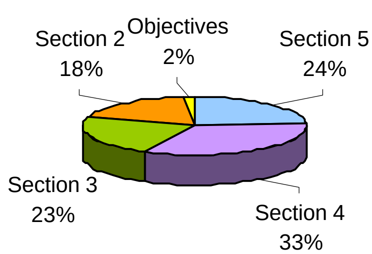

Chapter 3: Time spent (Alex)

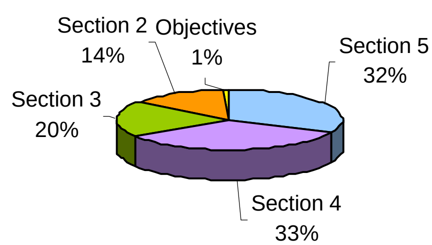

Total Time: 54 min 52 sec

Total Time: 47 min 7 sec

Table 1 & 2. Observed Frequencies for Two Event Sequences in Chapters 1 & 3 (Alex)

Alex Transition Counts (Chapter 1)

Actions	Highlight	Glossary	Note	Index	FOW	F Section	F Glossary	F Note	TOTAL
Highlight	8	4	1	-	4	-	-	-	17
Glossary	-	-	-	-	-	-	-	-	0
Note	-	-	-	-	-	-	-	-	0
Index	-	-	-	-	-	-	-	-	0
Focus Organizer	-	-	-	-	-	3	-	-	3
Focus Section	1	2	9	-	5	3	-	-	20
Focus Glossary	-	-	-	-	-	-	-	-	0
Focus Note	-	-	-	-	-	-	-	-	0
TOTAL	9	6	10	0	9	6	0	0	40

Alex Transition Counts (Chapter 3)

Actions	Highlight	Glossary	Note	Index	FOW	F Section	F Glossary	F Note	TOTAL
Highlight	-	-	-	-	-	-	-	-	0
Glossary	-	-	-	-	-	-	-	-	0
Note	-	-	-	-	-	-	-	-	0
Index	-	-	-	-	-	-	-	-	0
Focus Organizer	-	-	-	-	-	1	-	-	1
Focus Section	-	5	1	-	6	-	13	9	34
Focus Glossary	-	-	-	-	-	2	-	-	2
Focus Note	-	-	-	-	-	1	-	5	6
TOTAL	0	5	1	0	6	4	13	14	43

Tables 3 & 4 Transitional Probabilities

Alex Simple Probabilities (Chapter 1)

Actions	H	G	N	I	FOW	FS	FG	FN
Highlight	0.2	0.1	0.0		0.1			
Glossary								
Note								
Index								
Focus Organizer						0.1		
Focus Section	0.0	0.1	0.2		0.1	0.1		
Focus Glossary								
Focus Note								

Alex Simple Probabilities (Chapter 3)

Actions	H	G	N	I	FOW	FS	FG	FN
Highlight								
Glossary								
Note								
Index								
Focus Organizer						0.02		
Focus Section		0.12	0.00		0.14		**0.30**	0.21
Focus Glossary						0.05		
Focus Note						0.02		

Figure 9. Timing and sequencing of events and chapter sections (Sam)

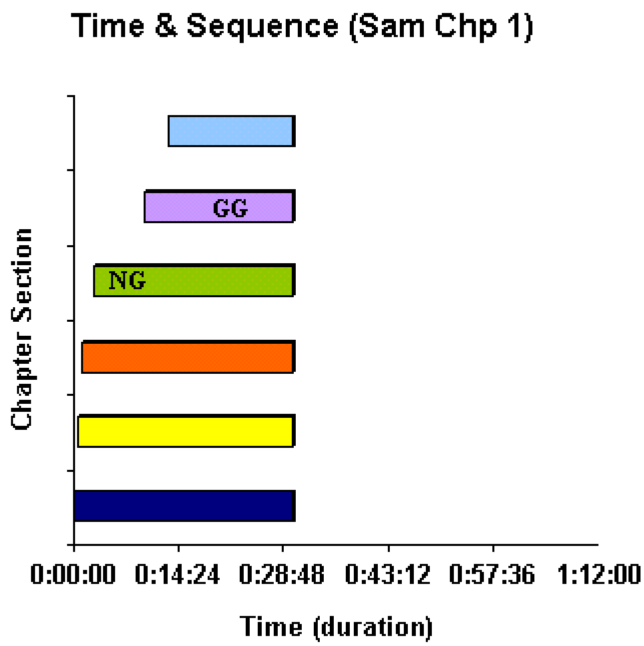

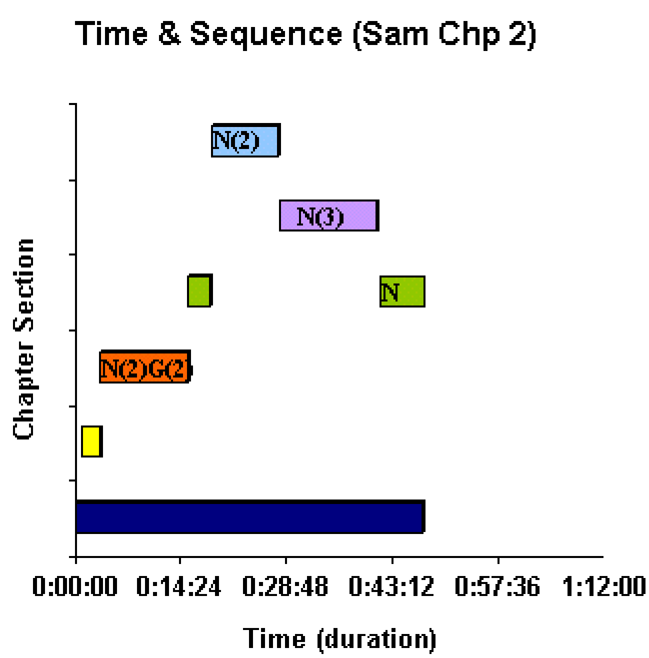

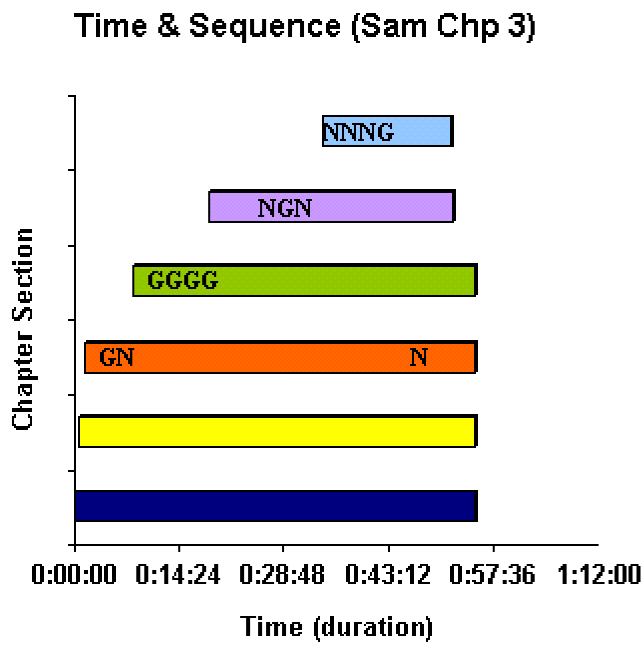

Figure 10. Timing and sequencing of events and chapter sections (Alex)

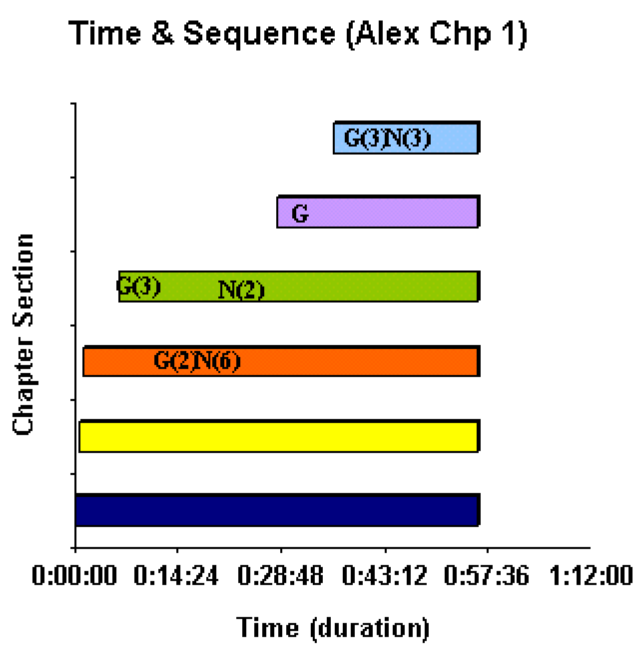

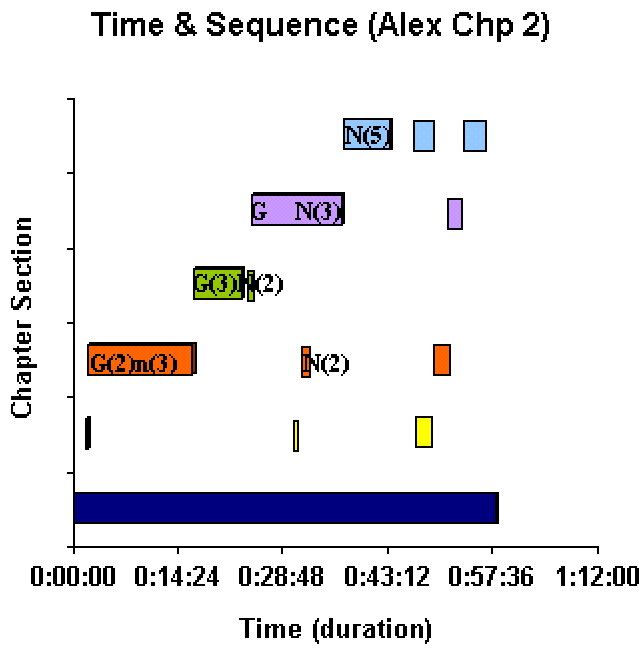
Time & Sequence (Alex Chp 2)
Chapter Section
N(5)
G N(3)
G(3)
(2)
G(2)m(3)
N(2)
0:00:00
0:14:24
0:28:48
0:43:12
0:57:36
1:12:00
Time (duration)

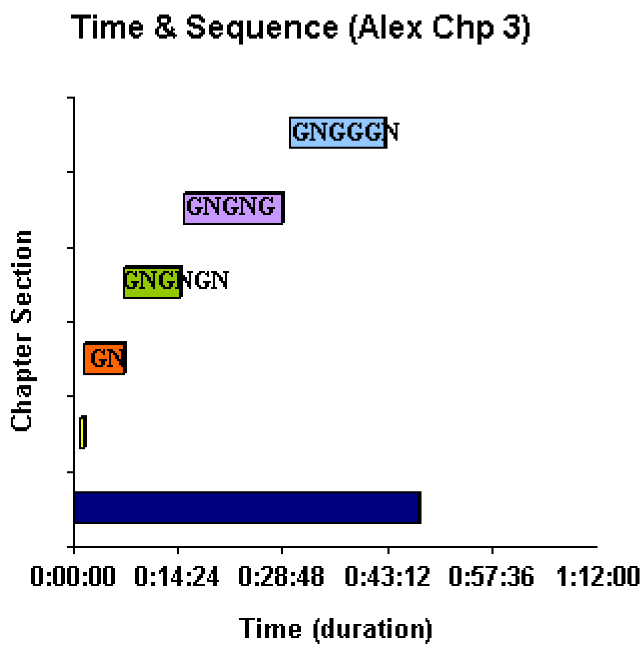
Time & Sequence (Alex Chp 3)
Chapter Section
GNGGGN
GNGNG
GNGNGN
GN
0:00:00
0:14:24
0:28:48
0:43:12
0:57:36
1:12:00
Time (duration)