



Using data mining as a strategy for assessing asynchronous discussion forums

Laurie P. Dringus *, Timothy Ellis

*Nova Southeastern University, Graduate School of Computer and Information Sciences,
3100 College Avenue, Ft. Lauderdale, FL 33314, USA*

Received 12 August 2003; accepted 6 May 2004

Abstract

The purpose of this paper is to show how data mining may offer promise as a strategy for discovering and building alternative representations for the data underlying asynchronous discussion forums. Presently, the instructor's view of the output of a threaded forum is limited to reviewing a transcript or print version of the written dialogue produced by participants. With potentially hundreds of contributions to review for an entire online course, the instructor lacks a comprehensive view of the information embedded in the transcript. In this context, the authors attempt to sort out the question, "what is data from an online forum?" among other key questions. The present work seeks to intersect the information (i.e., participation indicators) an instructor may wish to extract from the forum with viewable and useful information that the system could produce from the instructor's query. Temporal participation indicators are used to show how using data and text mining techniques in the query process could improve the instructor's ability to evaluate the progress of a threaded discussion.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Collaborative learning; Computer-mediated communications; Teaching-learning strategies; Distributed learning environments; Distance education and telelearning

1. Introduction

Asynchronous discussion forums are used increasingly in courses in which students and instructors interact in academic and social contexts. In instances where online discussions replace or

* Corresponding author. Tel.: +1-954-262-2073; fax: +1-954-262-3915.

E-mail address: laurie@nsu.nova.edu (L.P. Dringus).

supplement face-to-face class participation, instructors often wish to assess the level and quality of student activity in forums. Research continues to focus on assessing student activity within asynchronous discussion forums, particularly with regard to how forums can be used to support active engagement and discourse in instructional contexts (Garrison, Anderson, & Archer, 2000; Goldman, 2001; Graesser, Gernsbacher, & Goldman, 2003; Jarvela & Hakkinen, 2003; Jeong, 2003). However, many problems contribute to the difficulty of assessing activity in a forum and providing students with meaningful feedback about their progress and performance. The instructor needs to know what information is useful to extract from the transcript to begin a valid evaluation of student performance in the forum and develop meaningful feedback to the student. The instructor also needs to have a comprehensive view of the information that is contained in the written dialogue. In this context, the authors attempt to sort out the question, “what is data from an online forum?” among other key questions. The emphasis on assessment in this article relates to how data and text mining concepts and techniques can be used to reduce the difficulties instructors currently face in using the output from their course related forums as data to assess student progress and performance. The authors discuss how the manual process of assessing threaded discussion forums can be simplified by merging data mining concepts with assessment criteria and select participation indicators.

Threaded discussion forums are widely used, but there is not an accepted and tested method for assessing student participation. There are many issues and challenges in assessing threaded discussion forums. A lack of theoretical grounding of Web-based communication tools for academic use (Jarvela & Hakkinen, 2003; Koschmann, 1994) has been noted, indicating that the tools (from an instructional view) have not delivered the extent of discourse instructors desire to achieve in their online courses. Furthermore, from a system-related view, given the textual nature of most asynchronous discussion forums, assessment is hindered by limitations of the query and reporting toolset within the forum that most often produces rudimentary transcribed texts or frequency count outputs of discourse (Jarvela & Hakkinen, 2003). Jarvela and Hakkinen noted that there is an urgency to develop ways to organize and analyze data in Web environments to show the dynamics of online learning and interaction processes. An essential question is, “What are the possible contextual and pedagogical contributors for high quality conversations?” (p. 93). The authors of the present article contend that this question cannot be answered fully until the forum’s query and reporting toolset advances to provide instructors with alternative representations for the data underlying the forums. With different views of the data that can be extracted from a forum, the instructor then has the opportunity to explore in further detail the possible contextual and pedagogical contributors to online discussions.

For example, an important challenge in assessing threaded discussion forums is analyzing interactivity in terms of its historical progress as a community of learners, or how discussion has progressed over a period of time within a certain subset of participants (e.g., either for a course or some other purposeful group collaboration). To determine the quality of engagement of a group of participants over time, it is not enough to review each contribution simply as a separate and distinct “mini-essay”. Although the content of the posting is important, contributions as a whole (primarily the flow and exchange of ideas) in the forum discussion process cannot be ignored. This context can only be determined by viewing the way the forum progresses historically, or over a determined period. From an instructional view, questions such as: “When did the student make postings?” “Did the student respond to postings of other students?” “How immediate were those

responses?”, and “Did other students respond to this student?” provide a part of the criteria used to assess a participant’s contribution as a member of the forum community. Many of these criteria are not readily obtained from the forum system in an organized form. Rather, the instructor often has to create a manual mechanism for extracting these “data” from the forum output, making coding processes difficult and time consuming.

Researchers have developed models for analyzing the process of learning in asynchronous computer conferencing. (Henri’s (1982) *Analytical Model* and Garrison, Anderson, & Archer, 2001, 2000) *Practical Inquiry Model of Cognitive Presence*, are extensive content analysis models partially based on discourse theory, cognitive theory, and interaction theory. Garrison et al. (2001) and Jeong (2003) provided an invaluable assessment tool for analyzing learning and interaction constructs reflected in text-based computer conferencing transcripts. One key recommendation that Garrison et al. (2001) made was that there is a need to develop tools that effectively manage large numbers of messages in longer running online courses. Schrire’s (2003) follow up study to Garrison, et al.’s work also indicated that the coding process relies heavily on manual interpretation by the human rater, suggesting that more formidable content analysis tools should be embedded within the threaded forum assessment toolset to automate coding processes. Roblyer and Wiencke (2003) provided a rubric for assessing interactive qualities in distance courses. Their rubric serves as “a tool to allow for more meaningful examination of the role of interaction in enhancing achievement and student satisfaction in distance learning courses” (p. 77, abstract). However, presently these and other rubrics are not implemented within the forum toolset to produce a manageable view of the data that may represent the rubric criteria. The instructor is relegated to performing a parallel process of mapping the rubric in some way to whatever form the forum data is presented.

The research related to assessing student interactivity in computer conferencing and asynchronous discussion forums continues to expand to provide new techniques that could be used as standard assessment practice. In many cases, these works (Garrison et al., 2001; Henri, 1982; Jeong, 2003; Roblyer & Wiencke, 2003, and others) point to the need for assessment tools to reduce the cumbersome manual assessment process that burdens the instructor. In the present article, the authors foresee yet another viable alternative, one that combines data and text mining concepts within the discussion forum toolset, to simplify the burden the instructor carries in manually assessing forums.

1.1. Purpose of the paper

The purpose of this paper is to show how data mining may offer promise as a strategy for discovering and building alternative representations for the data underlying asynchronous discussion forums. Data mining is a process for examining databases to discover and display previously unknown interrelationships, clusters, and data patterns with the goal of supporting improved decision-making (Benoit, 2002). Businesses have used data mining to analyze customer demographics and transaction history to better target direct marketing efforts (Tsantis & Castellani, 2001). Although not yet widely used in education, several promising areas for data mining have been suggested and at least partially implemented in academic administrative applications such as a system to analyze transfer student records to identify predictors of success (Luan, 2002) and to automate co-author citation analysis to support scholarly research (He & Hui, 2002). To

date, no previous work has been found in the literature that combines data mining techniques with the assessment of asynchronous discussion forums. The present article extends the research by integrating data and text mining techniques in the forum query and reporting toolset and extending the instructor's ability to view various representations of data from a forum. The value of the present work is to show that simple mining operations can produce useful information for the instructor such as data related to time, pace, and sequence of contribution exchange. The strategy for using data and text mining is to support an instructor's method for providing meaningful feedback to students.

Three goals are intended by the authors of this paper: (1) to discuss the general system-related problems that contribute to the difficulty of assessing asynchronous discussion forums, (2) to identify common participation indicators that instructors may wish to extract from the forum as "data" to use to assess student progress and performance in online discussions, and (3) to describe data and text mining as a strategy for assessing forums, particularly in providing manageable views of the data. The first two goals are intended to provide background for the discussion on the strategy for mining data. The third goal, which is also the primary focus of the paper, is intended to describe and demonstrate the data mining concepts in detail. The data mining strategy is described by detailing the appropriate steps of mining data as applied to an analysis of threaded discussion forum contributions. Demonstrations are given to show how mining provides extended views of information from the forum. For the demonstrations, the authors have selected specific temporal-related participation indicators to show how the data mining process can be used to simplify the manual assessment of the forum, and in essence, to extend how student contributions and progress can be viewed as data in the output of a threaded discussion. The present work seeks to intersect the information (i.e., participation indicators) an instructor may wish to extract from the forum with viewable and useful information that the system could produce from the instructor's query. The present work only provides the concepts for data and text mining applied to a forum. The discussion on text mining is limited to demonstrating one example indicator. In addition, the authors do not provide an extensive interpretation of the example data in the form of student feedback.

1.2. Context

This paper should be viewed in the context of the following issues:

The authors discuss the system-related problems of assessing threaded discussion forums. The application of data mining concepts to the assessment process is demonstrated from the view of one single threaded discussion forum implementation. The authors' institution provides an online environment with a customized version of Allaire's Cold Fusion forum software. This forum software shares similar features common with those discussion forums widely used in course management systems (CMSs) such as WebCT, Blackboard, or AltaVista. The techniques for extracting data from the toolset of these CMSs may differ because of the lack of a standard toolset for threaded discussion forums. For example, the actual code used to execute the data mining operations is not included in the paper since the specific queries necessary to extract the desired information will vary with the software. The data mining *concepts* that are demonstrated using the Cold Fusion software can be applied to most CMS threaded discussion forums. Instead of actual code used for the Cold Fusion software, the authors provide an algorithm for showing the underlying logic of the code.

The authors attempt to provide a sample of participation indicators that have been noted in the literature. Common indicators are drawn to use as examples for how data can be extracted either manually by the instructor, or mined through the forums dataset. The table of indicators (Table 1)

Level 0: Topic initiation	→	Discussion topic 2 is now open for contributions	Tim	02/07/03
		↳ Multimedia Enhancement on Learning	Carol	02/08/03
		↳ Your note of caution seems wise	Dean	02/16/03
		↳ Restraint	Lorraine	02/25/03
Level 1: Position statement	→	↳ Does Multimedia Enhance Learning?	Dean	02/16/03
		↳ Response to Dean	Patricia	02/19/03
Level 2: Response statement	→	↳ More support for multimedia's effectiveness	Cathene	02/23/03
		↳ Research	Lorraine	02/25/03
		↳ Reply to Dean	Tammy	02/26/03
		↳ Enhancement of multimedia	april	02/22/03
		↳ Response to April	Walterzene	02/27/03
Level 3: Response to a response	→	↳ How about this?	april	03/13/03

Top
 Previous
 Next
 Print
 Reply
 Edit
 Move
 Delete

Date: February 16, 2003 06:42 PM

Author: Dean

Subject: Your note of caution seems wise

I think you're right to be concerned that the bells and whistles of multimedia, and the enthusiasm that some learners show for them, may entice courseware authors to defeat their own purpose by inadvertently distracting learners from the tasks at hand (Hede, 2002). Measuring the course objectives obtained would help steer authors to the right mix of show and substance. It certainly underscores the importance of thorough instructional design procedures and accurate measurement of goal achievement (Kemp, Morrison, and Ross, 2001).

Hede, A. (2002). An integrated model of multimedia effects on learning. *Journal of Educational Multimedia and Hypermedia*, 11(2), 177-191.

Kemp, J., Morrison, G., and Ross, S. (2001). *Designing effective instruction*. New York, John Wiley & Sons.

Top
 Previous
 Next
 Print
 Reply
 Edit
 Move
 Delete

Date: February 25, 2003 09:42 PM

Author: Lorraine

Subject: Restraint

Essentially the right combination of appropriate multimedia may be more effective than the use of a specific multimedia type.

When utilizing multimedia products in an attempt to "wow" the audience a presenter may become more involved in the delivery method than in the content. The strength in any multimedia presentation is the ability to effectively apply communication tools for each component of a communication problem, and to do so in a cohesive fashion. Successfully integrating multimedia into a project requires careful consideration of specific products, available resources, and most importantly the users (Eilers, 2002).

Eilers, Perthenia (2002). *Developing effective multimedia for users*. Retrieved February 15, 2003, from <http://www.utah.edu/stc/PERTSSTCPAPER.pdf>

Fig. 1. Threaded discussion board organization – outline view (top) and compiled view (bottom).

is intended to show only the *range* of indicators that may be used. The table is not intended to serve as a rubric for evaluation or as a tool in itself. The authors have selected indicators (Table 2) that are temporal, such as time, pace, and sequence, to provide a simple demonstration of mining concepts. To address the aspects of quality assessment in various quantitative and qualitative forms is beyond the scope of the article.

2. Background

2.1. General system-related problems of assessing forums

Many problems contribute to the difficulty of assessing threaded discussion forums. Two critical system-related problems of assessing forums are presented in this section: the limitations of common thread and message organization in a forum and the problems of using forums output as data for assessing student participation and performance in a forum.

The online environment of the authors' institution provides a customized version of Allaire's Cold Fusion Forum software. Such is similar to the main screen of a typical Web-based asynchronous threaded discussion forum that contains an open screen where topics, commonly referred to as threads, are posted. Typically, within a main thread, a group of messages on a similar topic expands in the form of a response-to-response linear-type list, creating a branching effect. The system provides an automatic tab indent of posted messages that brings visibility to the linear progression of the response list. However, the extent of linear progression also depends on manual placement of messages within a thread, allowing participants to manage the organization of threads, messages, and discussion flow. In this manner, participants select the thread to which they wish to respond and prepare a response to post directly under the previous posting. Fig. 1 contains an example of a typical asynchronous thread organization pattern from an outline view and a compiled view. To show an organized flow pattern from an outline view, levels are indicated to show where topics are initiated (Level 0), where position statements are given (Level 1), where response statements are placed (Level 2), and where there is a response-to-response message placement. The top portion of Fig. 1 shows an outline view of the organization pattern.

Manual placement by participants produces problems in data organization and flow. An exact and efficient organization scheme is not ensured given many messages are not placed in the proper sequence of thread or message response origin. Some threads and message postings are not well organized, creating a complex branching reply tree and difficulty in following the flow of discussion. Similarly, threads often become fragmented by theme and by temporal aspects of time, pace and sequence (Cazden & Beck, 2003), frequency, and duration of discussion as contributions build over time. These types of problems inherent in the arrangement of a typical asynchronous threaded discussion forum restricts thread and message organization, and to some extent, the level of engagement by participants. The bottom portion of Fig. 1 shows a compiled view of the organization. The compiled view only contains a few contributions. With potentially hundreds of contributions to review for an entire online course, the instructor lacks a meaningful view of the information embedded in the transcript.

Using forums output as data presents a challenge for the instructor who wishes to assess the level of contribution and exchange in threaded discussions. The restricted organization of threaded discussions makes assessing data (i.e., written dialogue) difficult to manage. The forum design, in its native, textual, and static form, limits the options the instructor has from which data to support assessment can be queried from the system. At minimum, assessment can begin through reviewing a transcript or print version of a part of or the entire written dialogue exchanged by the class participants. However, the forums transcript is generally static, as messages or other data cannot be easily moved or reorganized to improve the content flow as needed. Scale is a related problem in that the dialogue is mainly textual and is varied in size and scope, i.e., where contributions typically contain a range from a few words to a few sentences and from short to very long essay type postings. In addition, the fragmented nature of asynchronous forums in general, restricts the assessment strategy to primarily a manual level of simple reading and reviewing contributions and discussion flow (Williams & Murphy, 2002). Discontinuity, fragmentation, and loss of context of discussion are often evident in discussion threads, particularly when contributions build over time. When threads become large and deeply nested in a response-to-response linear list form, the instructor has a difficult challenge to assess the content in different ways.

Asynchronous forums, like Allaire's Cold Fusion Forum, may have embedded search tools in the software that assist with the manual critique of discussion flow, but there are limitations to the types of data that can be extracted with these tools. The search tools generally restrict assessment to a query level such as counting the number of threads, messages, and frequency or time stamps of when responses were posted. The coding scheme of messages is generally regulated to simple counting and categorization of contributions that are easily visible within the forums transcript (Jarvela & Hakkinen, 2003).

2.2. Participation indicators as data to assess student progress and performance in forums

A survey of the literature has not revealed a comprehensive participation indicator dataset from which to draw consistent assessment practice. A defined participation indicator dataset could be useful to show common or standardized approaches to analyzing threaded discussion forums. In the present article, the authors use the term "participation indicators" to describe "data" that can be used to represent the various ways threaded discussion forums can be assessed in qualitative and quantitative forms. Garrison et al. (2000) used the term "indicators" to represent key words and phrases that were organized into distinct groups from an analysis of computer conferencing transcripts. In the present article, the authors contend that the number or type of participation indicators are essentially unlimited in what instructors could use to define how activities in a forum are used as data points for assessment.

However, identifying key indicators that are well defined, useful, and extractable in assessing forums is difficult. Precise definitions for interaction and the levels of interaction, for example, are not easily identified, as Moore (1989) indicated that the construct of interaction has been given so many different meanings that the term is often misused. Further, it is difficult to isolate the extent to which certain participation indicators are extractable from textual transcripts and other usable forms given the textual arrangement of current threaded discussion forums. The promise of participation indicators is a vast range of assessment data points from which instructors can

choose. The challenge in current research is to identify and define clearly a standard set of key participation constructs and indicators that can be commonly used to assess activity in forums.

In the present article, the authors have attempted to list some common participation constructs and indicators that have been identified in the literature, specifically focusing on those that may pertain to asynchronous discussion forums. Some “common” indicators identified in the present article, having been drawn in “piecemeal” from previous studies, most likely represent only a subset of the many that actually exist. Nevertheless, the present article draws out the common indicators to use as examples for how data can be extracted either manually by the instructor, or mined through the forums dataset. Table 1 provides a sample list of participation constructs and associated indicators that are commonly used by instructors to assess forums. Many of these are often undefined or are used in different contexts, making it difficult to appropriately categorize and match patterns of constructs and indicators. In Table 1, literature citations are noted when key constructs, phrases, or indicators have stood out in current research. The table is a start for illustrating the need for a comprehensive list of participation constructs and indicators to be addressed in future work.

Many of the indicators shown in Table 1 are manually derived by the instructor and not readily available from the forum system as specific data points. Even in forums where highly organized threads and manual placement of messages follow the proper sequence of discussion flow, content analysis becomes difficult to manage. Primarily, the decision criteria an instructor may use to assess forums vary with the instructor’s experience and training using forums and the level of content analysis output the system provides the instructor. Content analysis of the forum becomes a tedious process. If instructors wish to perform a content analysis on a forum, data must be extracted from the system that represents a wider range of participation indicators than are currently extractable.

3. A strategy for mining data and text in assessing forums

Since threaded discussion forums are built upon a computerized database foundation, an examination of that technology would be helpful in an investigation of how to better derive meaningful information from the forum. A database is essentially a system for managing related data. Data are pieces of meaningful fact that can be observed and recorded. The database provides structure to these pieces of fact to create models of one or more aspects of the real world with the goal of making those aspects more easily interpreted and understood (Elmasri & Navathe, 2000).

In the case of threaded discussion forums, the real world model produced by the database is typically limited to a listing of entries organized by comments and the associated responses. This structure shows the flow of discussion – tracking who responded to whom, but other meaningful constructions could be developed from the data. One of the primary strengths of computerized databases is the capacity to construct multiple views. The twofold challenge is to identify alternative views that represent the flow of discussion and restructure the display of the data to show clearly, what actually occurred in the forum.

In both business and educational applications, data mining has been based on a computerized, statistical analysis of extremely large databases (Benoit, 2002). Since the forum database is a relatively small structure, a fully automated implementation of data mining technology is not

Table 1

Common participation indicators used to assess student progress in forums

Participation construct	Example indicators (data derived manually by instructor)
Level of interaction in forum;	<ul style="list-style-type: none"> • Level of discussion – high, progressive, low level (Jarvela & Hakkinen, 2003)
Learner–learner interaction activity (Moore, 1989)	<ul style="list-style-type: none"> • Message type: messages, contributions, responses, and general postings • Amount of interaction (Wentling & Johnson, 1999)
Degree of “presence” in forum Cognitive and Social presence (Garrison et al., 2000)	<ul style="list-style-type: none"> • Early, middle, late, or last minute contributions • Critical thinking skills – creativity, problem solving, intuition, insight (Garrison et al., 2000) • Practical inquiry – triggering event, exploration, integration, resolution (Garrison et al., 2001) • Encouraging collaboration among peers (Garrison et al., 2000)
Timing and pace: respond to others in a timely fashion	<ul style="list-style-type: none"> • Instructor’s definition of the immediate vs. latent continuum • Interval “wait” time (Cazden & Beck, 2003) • For responses, time between initial posting and response
Staying on topic; Learner-content (Moore, 1989)	<ul style="list-style-type: none"> • Meaningful and relevant keywords, phrases used to stay on topic
Transitions; Turn taking activity Shifts in topical focus	<ul style="list-style-type: none"> • Transitions/continuous discussion • Interval “wait” time (Cazden & Beck, 2003) • Initiating versus responding to topics (Williams & Murphy, 2002) • Within/outside the boundaries of the thread domain; where the changes occurred; subtle and dramatic shifts
Extent of instructor interaction; facilitation	<ul style="list-style-type: none"> • Facilitation rating– highly facilitative, informative, useful, non-facilitative
Instructor–learner activity (Moore, 1989)	<ul style="list-style-type: none"> • Instructor lag time; feedback (Caspi, Gorsky, & Chajut, 2003) • Defining and initiating discussion topics; focusing discussion (Garrison et al., 2000)
Teaching presence (Garrison et al., 2000)	<ul style="list-style-type: none"> • Instructor–learner knowledge building processes (Schrire, 2003)
Mandatory/non-mandatory participation	<ul style="list-style-type: none"> • Lag time between postings • Increased or decreased interaction and mandatory or non-mandatory directives about topics and group size (Caspi et al., 2003)
Lurking (Beaudoin, 2002); Vicarious interaction (Fulford & Zhang, 1993)	<ul style="list-style-type: none"> • Time spent on reading, writing, composing messages • Extent of reviewing vs. participating (Beaudoin, 2002) • Extensive/exclusive use of summarizing postings • Absence/lag time between postings
Shared resources	<ul style="list-style-type: none"> • Citations from scholarly literature • Usable resources such as links to informative websites
Accuracy of message content	<ul style="list-style-type: none"> • Content is accurate, valid, and relevant; writing quality
Accuracy response placement	<ul style="list-style-type: none"> • Responses are placed in the proper sequence
Contribution to “group process” (Gestalt view); group size	<ul style="list-style-type: none"> • Quality of contributions/the number of contributions made; conflict/negotiation: agreements/disagreements made in responses (Jeong, 2003; Gunawardena, Lowe, & Carabajal, 2000)

Table 1 (continued)

Participation construct	Example indicators (data derived manually by instructor)
	<ul style="list-style-type: none"> • Position statements versus new points • Group think versus individual effort • Group size and proportion of interaction (Caspi et al., 2003)

indicated; a manual application of the concepts underlying the data mining process should be adequate. The strategy for mining data can be applied to an analysis of the forum database to provide multiple views of the historical progress of a forum. The multiple views of the historical progress would enable the instructor to assess student contributions in the context of the group process inherent in the forum.

There is not a single model for mining data. Typically, however, nine sequential steps are included in the mining process (Benoit, 2002). Because the database underlying the threaded discussion forums is a relatively small and well-structured dataset, all nine steps are not necessary to extract successfully the information necessary to support improved assessment of contributions. This may not be the case for all types of forums databases. The following section provides details on how the appropriate steps of mining data as applied to an analysis of threaded discussion forum contributions using Allaire's Cold Fusion Forum software.

Step 1: Clearly identify the task

The first step in any data mining procedure is to examine thoroughly the reason(s) prompting the analysis. In the case of the forum database, the general task to be accomplished is to improve the assessment of student performance in threaded discussion forum assignments by placing the student's participation in the context of the forum's group processes. Table 1 presents many of the specific indicators commonly found in forum participation. The goal for this mining operation is

Table 2

Enhanced assessment of student progress in forums: Participation indicators/data extracted from the forum system

Participation construct	Example indicators (data extracted from the forum dataset)
Degree of "presence" in forum	What is the distribution and frequency of participant's contributions through the time frame in which the forum was open?
Lurking	<ul style="list-style-type: none"> • Number of contributions • Categorization of contributions – early, middle, late, last minute
Level of interaction in forum; learner–learner interaction activity	What is the distribution and frequency of a participant's contributions in terms of initiating versus responding to topics?
Transitions; turn taking activity	<ul style="list-style-type: none"> • Number of Level 1 contributions
Extent of instructor interaction; facilitation; instructor–learner activity	<ul style="list-style-type: none"> • Number of Level 2, 3, etc. contributions • Categorization of contributions – early, middle, late, last minute
Timing and pace: respond to others in a timely fashion	For responses, how soon after the initial posting was the response made? <ul style="list-style-type: none"> • Listing of each Level 2, 3, etc. contribution • Number of hours and days after the initial posting lapsed
Shared resources	Text mining for keywords, phrases related to topic area

to discover patterns of data that can provide insight into performance on one or more of those indicators, such as the indicators listed in Table 2.

In Table 2, the authors selected seven indicators to show how data mining techniques can be used to increase an instructor's efficiency in extracting data from the forum system. For the purposes of providing a simple demonstration of mining concepts, the authors selected indicators that are temporal, such as time, pace, and sequence, to represent the historical progress of a community of learners, or how discussions progress over time. These indicators include the distribution and frequency of the participant's contributions, represented by number of contributions and how contributions are categorized in time frames; the distribution and frequency of contributions in terms of what contributions were initiated versus responded to, represented by the number of levels (1, 2, 3, etc.) of contributions; time frames of postings and between postings, represented by number of hours and days lapsed between a response posting and the initial posting; and, text mining for keywords and phrases related to the topic area to assess shared resources.

Step 2: Get to know the data available

“Knowing” the data available for the data mining entails understanding the underlying structure of the database. What categories of people, places, or things – entities – are being tracked? For each entity, what characteristics or attributes – fields – are being recorded? For each field, what type of data – alphanumeric, long text, integer, date, etc. – is recorded? Finally, how do the various entities interact – relate? Although one might try to infer this information by reviewing the output from the database, the only way of knowing the nature of the data available is to examine that structure. A database administrator would generally handle this task.

The next three *steps* – *Acquisition of data*, *Integration and checking*, and *Data cleaning* – are necessary in large-scale data mining of poorly structured data. Because the forums database is small, well structured, and includes measures to protect data-integrity, these steps are not of great concern for this discussion.

Step 6: Develop initial questions

The data mining process itself will generate previously unformulated questions that can be answered through an analysis of the data. However, it is important to start the process with at least some general questions to be explored. The initial questions for the forums data mining, as detailed in Table 2, include:

1. What is the distribution and frequency of a participant's contributions through the time frame in which the forum was open?
2. What is the distribution and frequency of a participant's contributions in terms of initiating versus responding to topics?
3. For responses, how soon after the initial posting was the response made?

Step 7: Mine the data

The actual mining of the data entails writing code to extract the desired information from the dataset. Since the actual code used to extract the information necessary to answer the three questions listed above is rather arcane and specific to the Cold Fusion software implementation used for this study, it is of limited general interest. Of greater interest are the algorithms or logic underlying the code shown in Fig. 2.

The final two steps of the data mining process – *Step 8: Verification* and *Step 9: Interpretation* – focus on determining what the knowledge derived from the mining process means and how it can

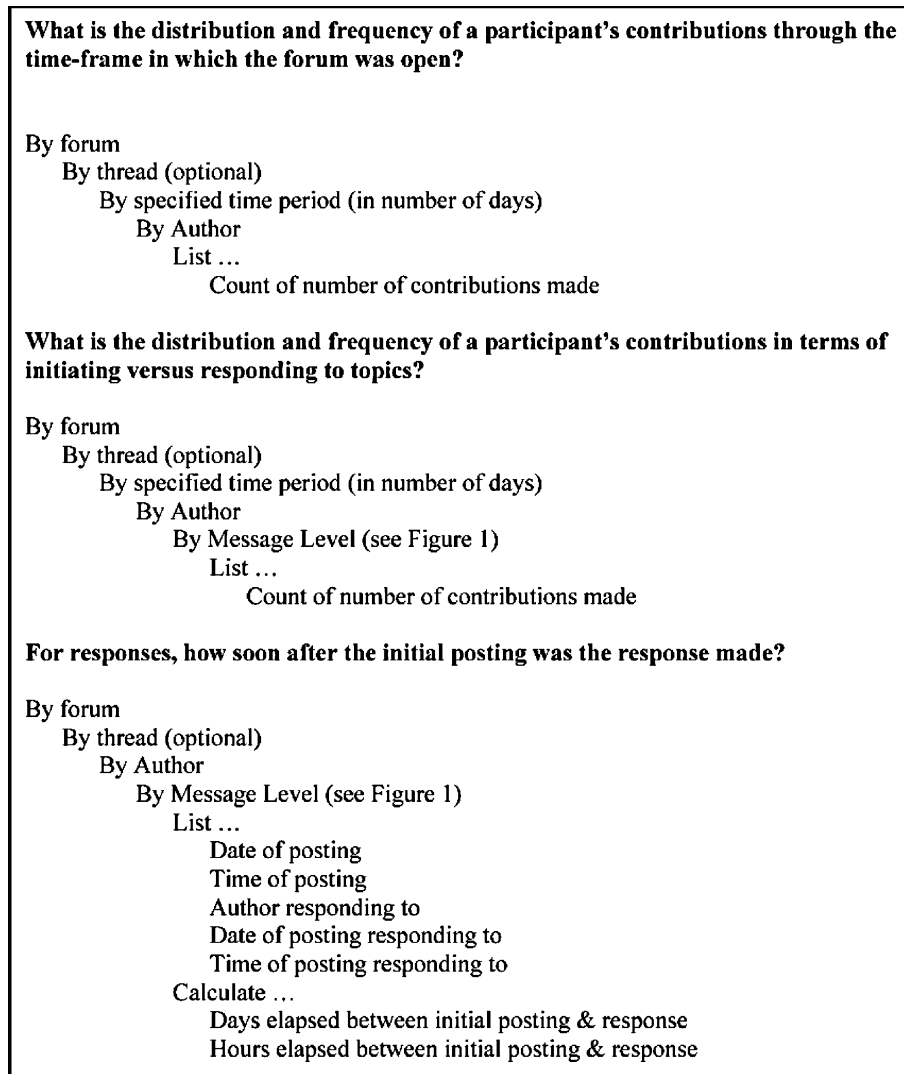


Fig. 2. Data mining algorithms.

be used to promote better decision-making. These steps relate back to the start of the process in which the *Tasks* were specified. Figs. 4–6 present the decision-support information mined from the forums database using the algorithms detailed in Fig. 2. How this information will be combined with the more qualitative measures of a student's contributions to the forum to assess her or his performance remains for the individual instructor to determine.

The authors selected indicators that are temporal, such as time, pace, and sequence, to represent the historical progress of a community of learners or how discussion progresses over a period of time by group or by individuals. To provide the reader with a context for the demonstration of the mining of these indicators, an example is given (Fig. 3) of a specific course objective that is facilitated by requiring students to participate in a debate or a similar group

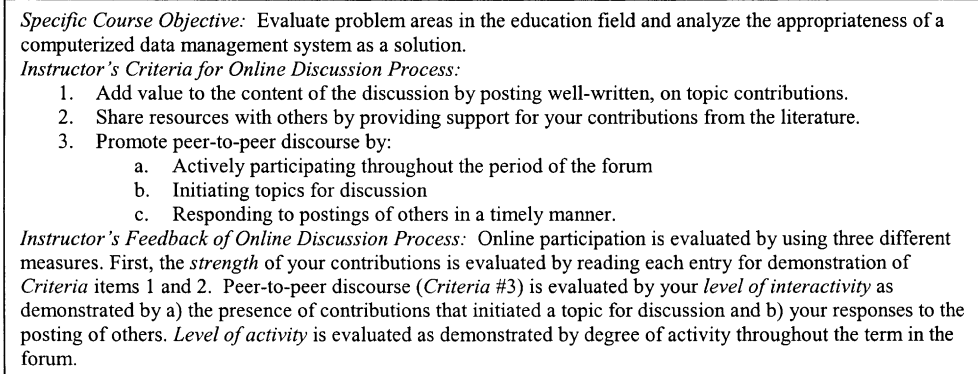


Fig. 3. Example instructions and feedback on student performance in the forum.

interaction activity. The example also includes the instructor's criteria and feedback of the online discussion process.

Mining demonstration of distribution and frequency of contributions (time indicators)

The following demonstration is related to question #1 in Table 2 and *Criteria* item #3 (level of interactivity) in Fig. 3:

1. What is the distribution and frequency of a participant's contributions through the time frame in which the forum was open?
2. When did the student make postings?
3. How immediate were those responses?

For example, to assess the degree of presence in a forum over a period, such as a week, several weeks or an entire term, the instructor may be interested in reviewing the indicators relative to distribution and frequency of contributions in specific time frames. To query the system, the instructor would input possible date ranges for when and how many contributions were produced by group or by individual. The scenario in Fig. 4 shows the output of that query: a time frame range in date format, the author, the number of contributions by the author, and the total number of contributions a group of participants have made to a specific thread. This information gives the instructor an organized view of the distribution and frequency of contributions. The instructor is then able to determine how often a student contributed and if these contributions were early, middle, late or last minute contributions.

Mining demonstration of distribution and frequency of initiation vs. response (level of interaction and frequency indicators)

The following demonstration is related to initial question #2 in Table 2 and *Criteria* item #3 (level of interactivity) in Fig. 3:

4. What is the distribution and frequency of a participant's contributions in terms of initiating versus responding to the topic?
5. Did the student initiate responses?
6. Did the student respond to postings of other students?

To assess participation constructs such as level of interaction, learner–learner activity, transitions in discussion threads, and extent of instructor facilitation, the instructor may wish to query the system to show the distribution and frequency of when contributions were made, what

02-07-2003 to 02-14-2003		
Andrea	2	
Carol	2	
Lorraine	2	
Rhonda	1	
Tina	1	
	TOTAL	9
02-15-2003 to 02-21-2003		
Dean	7	
Patricia	3	
Patricia Y	2	
Rhonda	1	
Tina	3	
	TOTAL	16
02-22-2003 to 02-28-2003		
Aprill	2	
Brian	6	
Caroline	5	
Cathene	5	
Dean	1	
fran	3	
Lorraine	5	
Patricia	2	
Patricia Y	2	
Rhonda	1	
Tammy	6	
Walterzene	5	
Tina	2	
	TOTAL	45

Fig. 4. Distribution and frequency of contributions.

contributions were initiated and what responses were position statements or responses to other responses. In essence, the instructor can query the system for a view of the levels of interaction (Levels 0–3), and how those patterns were of value to the participant's effort to sustaining discussion throughout the thread. The scenario in Fig. 5 shows the output of the query – the week of contribution, the author, the message level (level of interaction) and the number of contributions posted. The instructor would then be able to better gauge the level of the student's engagement and the depth of exchange on the topic.

Mining demonstration of latency between initial posting and response (timing and pace indicators)

The following demonstration is related to initial question #3 in Table 2 and *Criteria* item #3 (level of activity) in Fig. 3:

7. For responses, how soon after the initial posting was the response made?
8. To what degree did the student make steady progress in the forum?

Timing and pace, and responding to others in a timely fashion are common indicators used to assess a participant's progress. The instructor would query the system to specify the level of in-

Week	Author	Message Level	Number Posted	
Week 1	Rhonda	Level: 1	1	
	Carol	Level: 2	2	
	Lorraine	Level: 1	1	
	Lorraine	Level: 2	1	
	Andrea	Level: 1	1	
	Andrea	Level: 2	1	
	Tina	Level: 3	1	
	Week 2	Patricia	Level: 1	1
		Patricia	Level: 2	2
		Dean	Level: 2	4
Dean		Level: 3	2	
Tina		Level: 1	1	
Tina		Level: 3	2	
Rhonda		Level: 3	1	
Patricia Y		Level: 1	1	
Patricia Y	Level: 3	1		

Fig. 5. Distribution and frequency of initiation vs. response contributions.

teraction, the time posted, who the participant replied to, the original date the message was posted, and the time elapsed (latency) between the initial posting and the response. The scenario in Fig. 6 shows the output of an organized view for the instructor to determine the extent of steady progress in the forum, by the individual and by the group.

3.1. Mining the text to assess on topic and shared resources

The majority of the forum data, and the information necessary to answer some of the most interesting questions that can be generated from Tables 1 and 2 – “Was the contribution on topic?” and, “Did the student share resources in the contribution?” for example – are contained in a single, large text field in which the content of the contribution is stored. Since these data are in an unstructured, text format, they are not accessible through data mining techniques. Text mining technology is necessary to access this type of information.

Text mining offers access to “. . . information in unstructured textual form [that] is not readily accessible to be used by computers” (Dorre, Gerstl, & Seiffert, 1999, p. 398). Included in text mining are the dual functions of extracting items of information – features – from textual documents and analyzing the distribution of those features across multiple documents to identify

AUTHOR	Level	Time Posted	Reply to	Original Message Posted	Time Elapsed
Rhonda	Level: 1	02/07/2003 09:00 AM			
	Level: 3	02/17/2003 03:15 PM	Andrea	02/11/2003 12:58 PM	6 days 2 hours
		02/28/2003 11:21 AM	Dean	02/16/2003 04:54 PM	11 days 18 hours
Patricia	Level: 1	02/15/2003 12:49 AM			
	Level: 2	02/15/2003 12:56 AM	Rhonda	02/07/2003 09:00 AM	7 days 15 hours
		02/15/2003 01:19 AM	Andrea	02/10/2003 08:59 PM	4 days 4 hours
		02/27/2003 07:24 AM	Tammy	02/26/2003 08:49 PM	0 days 10 hours
	Level: 4	02/23/2003 02:21 AM	Dean	02/17/2003 09:22 AM	5 days 16 hours

Fig. 6. Latency between initial posting and response.

patterns of value. These functions can be accomplished through either an automated process (clustering) in which a set of rules instruct the computer how to identify blocks of text that might be of interest or a more manual procedure (categorizing) in which the user must first develop a taxonomy of expressions of interest by which the computer would order blocks of text.

Unfortunately, text mining does not yet offer promise as a tool to evaluate threaded discussion forum contributions to determine if the contributions were well written and on topic since those concepts cannot yet be adequately constrained in the current technology. Text mining techniques can, however, be used to examine the large blocks of text for the presence of specific content that could address the quality of the posting.

Mining for shared resources. For example, text mining may be useful to assess the participation construct “shared resources”. As there is not a commonly accepted definition for “sharing of resources” in a discussion forum entry, the authors operationally defined the term as including Web references or citations from the literature. Web references can be easily mined by searching the “Message” field for entries that include an “http://” string, and literature citations can be identified by looking for strings that contain a “(####)” string, in which the # character is a wildcard representing any numeral. It should be emphasized that a text mining operation of this

nature will not guarantee the quality of the resources shared, only the presence of an indicator of sharing of resources.

3.2. Mapping multiple views to student performance and student feedback

It is beyond the scope of this article to provide an extensive interpretation of the data that result from the mining operations. The intent is to show the need for better visibility of the data. In short, what is visible from the mining demonstrations of the selected temporal indicators is that the instructor now has clearer information to work with to determine student performance based on preset criteria pertaining to level of interactivity and degree of activity. A substantial contrast is evident when comparing the view of the transcript output in Fig. 1 to the expanded views provided in the mining demonstrations of the temporal indicators. The expanded views show clearer patterns of activity and peer-to-peer discussion, ranging from the number of contributions to the distribution of contributions and responses over a specified time. The expanded views also provide time ranges, frequency of contributions, and what contributions were position statements or responses to other responses. With this information, the instructor could begin to answer such initial questions as “Did the student respond to postings?” “How immediate were those responses?” The output could be used to determine a student’s collaborative effort in a forum. Instructors could assign course points or percentages to this effort. Meaningful feedback could then be generated for the student.

Establishing criteria for forum participation may help students learn how to engage effectively in an online forum. Therefore, the instructor needs to establish the criteria in advance and decide what data is necessary to extract from the forum system that represents the criteria. However, even with tools to assist instructors with managing forum data, providing meaningful feedback to students will remain to a large degree a highly subjective and inexact method. In other words, an enhanced tool will not replace the need for the instructor to manage and engage in the process of assessing student performance and providing meaningful feedback to students. The tool will serve, at best, to provide and organize information for the instructor.

4. Conclusions

The assessment of threaded discussion forums offers the instructor a number of challenges. Since two common goals for a discussion forum assignment are typically to foster collaborative learning and build a community of learners, it is not enough to evaluate each contribution in isolation; contributions as a whole in the forum discussion process cannot be ignored. In essence, the forum, as an instance of human discourse, must be at least partially assessed from a discourse analysis perspective. Unfortunately, current discussion forum technology imposes a double bind effect on instructors. Although the threaded discussion forum is widely used to support a variety of course requirements, the instructor is faced with the difficulty of interpreting and evaluating the learning and quality of participation reflected in the student contributions. When postings to a single topic can run into the hundreds, span several weeks, and represent many types of contributions, the view of an individual student’s experience in the learning activity is not clear. The student’s success is represented by the content of the individual contributions posted and the

context in which the postings were made. The sheer volume of postings makes it difficult, if not impossible, to analyze content and context for each student. The instructor is faced with the dilemma of having access to a great amount of data that could potentially be of use in evaluating a student's performance but has limited capacity for processing that data into meaningful information.

Data mining and text mining techniques offer the promise of providing the instructor with information upon which an analysis of the process aspects of a forum assignment can be based. Although the results of the mining operations cannot, and should not, replace a careful reading of each forum contribution, the results can show objective information to questions such as: "When did the student make postings?" "Did the student respond to postings of other students?" "How immediate were those responses?" "Did other students respond to this student?" "Did the student share resources with others in his or her posting?" and "Did the posting cite a specific reference?" This information could lead to meaningful feedback about student's performance in a discussion forum. The authors used several time-related participation indicators to show how mining techniques could be useful in capturing an organized view of the historical progress of a community of learners or how discussion progresses over time. Other participation indicators should also be visible from the forums reporting toolset.

The discussion about participation indicators revealed an urgent need for future research to identify key indicators that are well defined, useful, and extractable in assessing forums. Without highly defined participation indicators, instructors lack common ground or a standard set of criteria from which assessment can be consistently made. Participation constructs and the associated indicators are found in the literature, but in piecemeal, lacking comprehensive organization and consistent use by educators. This lack of research raises a few important questions, "What is data from an online forum?" "How can data be extracted from the forum toolset?" "How should data from a forum be analyzed?" Until there is a more comprehensive and standardized participation indicator dataset that can be easily extracted from the forum system, instructors will mostly likely continue to extract manually their own indicators and data points from transcribed texts. The manual effort can certainly decrease consistency in evaluations and decrease efficiency of instructors' work, requiring a substantial amount of time and effort an instructor must expend in assessing forums. Data and text mining techniques can be one part of the solution to reduce the difficulties instructors face in using their course-related forums output as data.

The strategy for discovering and building alternative representations for the data underlying asynchronous discussion forums is important to achieve to extend the instructor's ability to evaluate the progress of a threaded discussion. The authors contend that a mining strategy can be effective, with various participation indicators, and with any formidable assessment model or rubric that an instructor may wish to use insofar as the forum toolset is programmed to support the model or rubric.

However, there are also challenges in the data and text mining technology that currently cannot be mapped to many common discussion forum systems. In text mining, for example, the complexities and inconsistencies inherent in natural language text places advanced text mining implementation outside the scope of current technology (Dorre et al., 1999). Since there are virtually an unlimited number of ways in which the same concept could be expressed, computers cannot yet analyze and evaluate the content of text passages unless severe constraints are placed on the

domain being investigated. Mining within highly specified domains using limited, preset vocabularies has, however, proven to be not only possible, but also productive (Dorre et al., 1999; Li & Yamanishi, 2001).

Current asynchronous forums and embedded search tools restrict how online instructors can apply specific decision criteria to assess progress and participation in online discussions. The deep information discovery level of discourse and participation is not often retrievable from the reporting tools. The modeling opportunities coordinated by diverse visualization and data mining techniques in threaded discussion forums could help provide such meaningful output showing trends, patterns, summaries, and statistics to demonstrate the extent of contribution and progress in the forum in short and long-term durations, and for other participation indicators.

The authors demonstrated in this paper a strategy for using data and text mining approaches to extract temporal information from a threaded discussion forum based on Allaire's Cold Fusion forum software. The general technique should be equally applicable to forums managed by other systems such as WebCT, Blackboard, and AltaVista. Approaching the problem of mining forums for useful information with external tools is a workable solution. It is also essential to expand the query and reporting capabilities into the standard toolset within the forum system. It is recommended that assessment capabilities become an integral element of a course management system with the integration of SCORM and other standards for developing CMSs (Advanced Distributed Learning, 2001; Learning Technology Standards Committee of the IEEE Computer Society, 2001).

Acknowledgements

The authors thank Will Ferri and the GSCIS IT staff for their help in the development work of the forums reports as shown in the figures.

References

- Advanced Distributed Learning (2001). *Shareable content object reference model (SCORM): The SCORM overview*. Retrieved July 24, 2003. Available: <http://www.adlnet.org>.
- Beaudoin, M. F. (2002). Learning or lurking. Tracking the "invisible" online student. *The Internet and Higher Education*, 5(2), 147–155.
- Benoit, G. (2002). Data mining. *Annual Review of Information Science and Technology*, V36, 265–310.
- Caspi, A., Gorsky, P., & Chajut, E. (2003). The influence of group size on non-mandatory asynchronous instructional discussion groups. *The Internet and Higher Education*, 6(3), 227–240.
- Cazden, C. B., & Beck, S. W. (2003). Classroom discourse. In A. C. Graesser, M. A. Gernsbacker, & S. R. Goldman (Eds.), *Handbook of discourse processes* (pp. 165–197). Mahwah, NJ: Lawrence Erlbaum Associates.
- Dorre, J., Gerstl, P., Seiffert, R. (1999). Text mining: Finding nuggets in mountains of textual data. *Proceedings of the fifth ACM SIGKDD international conference on knowledge discovery and data mining, August 15–18, 1999, San Diego, California* (pp. 398–401).
- Elmasri, R., & Navathe, S. B. (2000). *Fundamentals of database systems* (3rd ed.). Reading, MA: Addison-Wesley.
- Fulford, C. P., & Zhang, S. (1993). Perceptions of interaction: The critical predictor in distance education. *The American Journal of Distance Education*, 7(3), 8–21.

- Garrison, D. R., Anderson, T., & Archer, W. (2001). Critical thinking, cognitive presence, and computer conferencing in distance education. *The American Journal of Distance Education*, 15(1), 7–23.
- Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical inquiry in a text-based environment: Computer conferencing in higher education. *The Internet and Higher Education*, 2(2–3), 87–105.
- Goldman, S. R. (2001). Professional development in a digital age. *Interactive Educational Multimedia* [On-line serial], Issue 2. Available: <http://www.ub.es/multimedia/jem/>.
- Graesser, A. C., Gernsbacher, M. A., & Goldman, S. R. (Eds.). (2003). *Introduction to the handbook of discourse processes. Handbook of discourse processes* (pp. 1–23). Mahwah, NJ: Lawrence Erlbaum Associates.
- Gunawardena, C. N., Lowe, C., Carabajal, K. (2000). Evaluating online learning: Models and methods. In *Society for information technology & teacher education international conference: Proceedings of SITE 2000 (11th, San Diego, California, February 8–12, 2000)* (Vols. 1–3, pp. 1677–1684).
- Henri, F. (1982). Computer conferencing and content analysis. In A. R. Kaye (Ed.), *Collaborative learning through computer conferencing* (pp. 117–136). Berlin: Springer-Verlag.
- He, Y., & Hui, S. C. (2002). Mining a Web citation database for author co-citation analysis. *Information Processing & Management*, 38(4), 491–508.
- Jarvela, S., & Hakkinen, P. (2003). The levels of web-based discussions: using perspective-taking theory as an analytical tool. In H. v. Oostendorp (Ed.), *Cognition in a digital world* (pp. 77–95). Mahwah, NJ: Lawrence Erlbaum Associates.
- Jeong, A. C. (2003). The sequential analysis of group interaction and critical thinking in online threaded discussions. *The American Journal of Distance Education*, 17(1), 25–43.
- Koschmann, T. D. (1994). Toward a theory of computer support for collaborative learning. *Journal of Learning Sciences*, 3(3), 219–225.
- Learning Technology Standards Committee of the IEEE Computer Society (2001). *IEEE P1484.11D8, 2001-04-06 Draft standard for learning technology – Learning technology system architecture*. IEEE Standards Department: Piscataway, NJ.
- Li, H., Yamanishi, K. (2001). Mining from open answers in questionnaire data. In *Proceedings of the seventh ACM SIGKDD international conference on knowledge discovery and data mining, August 26–29, 2001, San Francisco, California* (pp.443–449).
- Luan, J. (2002). Data mining and its applications in higher education. *New Directions for Institutional Research* (Spring), 17–36.
- Moore, M. G. (1989). Three types of interaction. *The American Journal of Distance Education*, 3(2), 1–6.
- Roblyer, M. D., & Wiencke, W. R. (2003). Design and use of a rubric to assess and encourage interactive qualities in distance education courses. *The American Journal of Distance Education*, 17(2), 77–98.
- Schrire, S. (2003). *Journal of Instruction Delivery Systems*, 17(1), 6–12.
- Tsantis, L., & Castellani, J. (2001). Enhancing learning environments through solution-based knowledge discovery tools: Forecasting for self-perpetuating systemic reform. *Journal of Special Education Technology*, 16(4), 39–52.
- Wentling, T. L., Johnson, S. D. (1999). The design and development of an evaluation system for online instruction. In *Web '99 World conference on the WWW and internet proceedings (Honolulu, Hawaii, October 24–30, 1999)* (pp. 2–7).
- Williams, C. B., & Murphy, T. (2002). Discussion groups: How initial parameters influence classroom performance. *Educause Quarterly*, 25(4), 21–29.