

**Enhancing Learning Environments Through Solution-based Knowledge Discovery Tools:
Forecasting for Self-perpetuating Systemic Reform**

Linda Tsantis

John Castellani

Johns Hopkins University

Introduction

The purpose of this article is to outline a future direction for special education through emerging uses of data mining and knowledge discovery. We will discuss how data mining can be used within special education based on our advancing technically equipped society, the explosion of information now contained within school databases, and a need to enhance instruction and improve delivery of special education services. It is our intent to explore how data mining is being used in industries such as finance, public health, and telecommunications and connect current innovations using data mining to special education.

In direct response to the rapid growth of these information databases, business and health industry leaders have turned to computer applications that can increase the efficacy of treatment and services. One of the most effective applications comes from the emerging science and industry of knowledge discovery in databases, also known as data mining.

Data mining applies modern statistical and computational technologies to the problem of finding useful patterns hidden within large databases. Data mining has proven to be a powerful tool capable of providing highly targeted information to support decision-making and forecasting for scientists, physicians, sociologists, the military and business. The predictive power of data mining comes from its unique design – it combines techniques from machine learning, pattern recognition, and statistics to automatically extract concepts, and to determine interrelations and patterns of interest from large databases (Edelstein, 1997).

The data mining process is most efficient when there exists a wealth of quality databases, a prior culture of analysis, and standards to safeguard privacy and promote the ethical use of information. To take advantage of the results, a system needs to be in place for transforming new knowledge into successful models for teaching and learning to develop and improve student relationship management.

The most appealing aspect of data mining is what it produces – complex statistics transformed into usable visualization charts that convey large amounts of information to a user in

the most meaningful and transferable form. This type of information transfer helps to reduce time spent with data and increases how data can connect to models of service delivery. The telecommunications and financial industry provide good examples of how data visualization can help organizations interact with data available on individuals they serve.

Once a model is built, the data-mining tool looks for patterns in the data and identifies these patterns through markers. These markers indicate patterns in collected data. In the case of industries dealing with finance, data mining tools can gather information on financial transactions. Those patterns that deviate from typical customer patterns are displayed as markers. The markers indicate that a specific account is either in jeopardy or needs specific attention. The data marker tells the responsible financier about the account and allows the individual to implement company policies about what to do in the case of an overdrawn balance, fraud detection, or other transactions. In the case of telecommunications, data visualization charts can assist interested parties to observe television viewing habits across the course of a week, by specific time of day from Sunday through Saturday. In this scenario, a data visualization chart not only communicates that on Thursday evenings from 9 to 11 is when most viewing occurs but that recently new programming has led to a shift in time from 9 to 11 on Tuesday to that same time bracket on Thursday. It may also tell us that viewing has increased for females within the age group of 20 to 45 rather than previous models indicating 8 to 11 for this same group. These charts can be accessed any day of the week during any season of the year, meaning that as spending or TV viewing habits change, the model will capture the pattern and inform the user of the change, degree of change, and sustainability of the change. The immediacy of these data allows for more precise decision-making within an organization.

Data mining tools monitor these complex relationships, which are then translated to interested parties through the use of data visualization charts, which turn several complex correlations and patterns into 3 dimensional graphs. These multidimensional graphs translate complex statistics through visual representations that can easily be understood by the end user. Advanced machine learning occurs when the model is programmed to look for non-intuitive

patterns that become future models, thereby learning from markers that subsequently are fed back into the initial system. This is typically known as artificial intelligence and has been pursued through computer science in the form of neural networks.

Currently, most educators use variations of knowledge discovery techniques through correlation and probability research. Educators at all levels know the degree of time required to turn existing school data into some form easily translated to principals, directors of special education, and ultimately, classroom teachers. In many cases, the data exist, the tools to translate the data exist, but the limited time available for practicing educators to organize even descriptive statistics remains limited. Time for advanced statistical analysis is even more restricted. Other organizations dealing with large data sets have turned to data mining as a solution since data mining continually refines available data that can be used to determine markers for decision-making and model development.

For a moment, consider the financial industry being inundated with the ever-increasing number of accounts and daily transactions through online banking and the ubiquitous bank machine or ATM. Data mining allows the bank to keep track of all of these transactions that can be queried at any given time; however, only the transactions that need immediate attention are marked and kicked out of the system through a data visualization chart indicating that something needs dire attention. It is timely, immediate, and has saved the financial industry millions of dollars. If we apply this same idea to schools, it is possible to imagine a system that automatically displays a data visualization chart informing a program director (e.g., principal or special education chair) of students receiving special education services that have 3 consecutive weeks of consistent D's, reached their 6th absence, and received their tenth discipline referral for the semester. Even though all of the data on special education students exist, it would be ideal to be able to target a group of students who may need immediate attention rather than waiting until a more serious situation has presented itself. This is not only a cost savings of time, but also a proactive effort to reduce staff burnout, and improved delivery of special education services. The

same scenario could be imagined for students who have received consistently high grades, perfect attendance, and no discipline referrals or any combination of student indicators.

School and Organizational Reliance on Data Mining Technologies

Many schools are ready for integrating data mining technologies into systemic approaches to education. The national networking infrastructure and personal connectivity to digital educational environments has made the collection, transfer, and dissemination of information more accessible than at any other point in history. However, the importance of the new information lies in its being refined to the degree that it can help predict and forecast trends. It is this predictive ability that distinguishes data mining from other more one-dimensional data management and number crunching applications. Data mining includes tools that can provide a meta-layer presentation of data that gives a snapshot of interesting data that might otherwise go unnoticed. The power of this technology to target interesting snapshots or markers empowers its users in making proactive, knowledge-driven decisions based on markers found within the data.

As educators, both authors recognize the degree of high quality research that education professors and school personnel deal with on a yearly basis. The use of data mining does not negate the use of educational research within schools. There are decisions relative to how and when each should be used. The National Center for Data Mining draws a distinction between research and data mining. Formal statistical inference is assumption-driven in the sense that a hypothesis is formed and then tested against the data. Data mining, in contrast, is discovery-driven in the sense that the hypothesis is automatically extracted from the data and therefore is data driven rather than research-based or human driven.

However, one of the most successful uses of data mining is the breakthrough of identifying markers in the Human Genome Project, which left off where researchers could not deal with the data in a timely and consistent manner. As we look toward the future in the field of special education, we may benefit from the advances from other disciplines to find more effective ways to interpret data.

Emerging Issues in Special Education

Students with disabilities in the year 2000 are increasing in number and type. As the special education classroom shifts to include greater access to the regular education curriculum, the roles and responsibilities of the special education classroom teacher are evolving so drastically that the term special education does not capture the instructional and behavioral demands. Teachers are working to prepare students to participate in a highly technical information society where opportunities for students with disabilities have expanded. As more students with disabilities enter the public education system, a crisis of professionalism in special education has become apparent.

Special educators have struggled to determine best practices for teaching and learning. Many of the obstacles for documenting appropriate instructional intervention have been (a) discovering what teachers are doing, (b) researching and evaluating special education models, and (c) transforming this into a discussion of best practices. The variables involved for teaching students with special needs can be extremely complex and difficult to document and assess. In the regular education classroom, there is already a multitude of student variables that influence student learning and determine success in differentiated educational programs. If we include the additional variables that students with disabilities bring to the learning environment, it is easy to see how educators have difficulty providing accommodations, individualizing teaching practices, and assessing student learning.

If we placed an educator from the early 1900s in the classroom of today, the student population would look extremely different. If that same educator were to look at how teaching and learning have changed as a result of the increased number of students with disabilities in the classroom, the instructional repertoire would look almost the same. Why? While there may be many reasons, one of great concern is the ability of the educational community to keep pace with changes in the educational landscape. New legislation, models for effective teaching practice, and highly influential special education research remain outside of the realm of teacher knowledge.

New understandings for teaching and learning with technology, efforts for inclusion, and advances in cognitive science have grown at an unprecedented rate; however, in many cases the day-to-day teaching practice for students with disabilities remains the same. In many cases, antiquated paper-based individualized educational programming (IEP) has evolved into computerized student information systems. Emerging electronic data systems have the potential to transform our ability to provide anticipatory guidance and predictive modeling opportunities for teaching and learning. Educators continue to struggle with two main questions.

1. What is the optimal learning environment for students with disabilities?
2. How can special and regular education systems work together to promote student achievement, teacher development, and organizational effectiveness?

Future Directions for Special Education

For the past 30 years, our society has increased educational programmatic access for persons with disabilities. This access has been promoted through arguments of inclusion, and legally through the integration of regular education practices in IEP planning. Regular educators now team with special educators to develop and implement the IEP. There is more of an emphasis on collaboration and systemic thinking on teaching and learning practice. Communication between educators has become more consistent, especially in the collection, analysis, and dissemination of information that appropriately categorizes learning outcomes for students with disabilities. Questions for increasing student access to successful teaching and learning practice within the regular education environment may naturally prove to accommodate students with disabilities. In some cases, it may even serve to reduce school failure.

More holistic views of professionalism in the field are required for determining best practices, including new ideas and integrated technologies that impact school settings, teacher instructional strategies, and ultimately student learning (Pugach, 1996; Rosenberg, Griffin, Kilgore, & Carpenter, 1997). There is a wealth of information available to the special educator, including behaviorist, cognitivist, and constructivist positions.

Kenneth Zeichner (1991) observes in *Teacher Education and the Social Conditions of Schooling* that:

We should be developing teachers who are able to identify and articulate their purposes, who can choose the appropriate instructional strategies or appropriate means, who know and understand the content to be taught, who understand the social experiences and cognitive orientations of their students, and who can be counted on for giving good reasons for their actions. These justifications should take into account the activity of teaching, the larger communities of educators, and a greater understanding of the social and political context of schooling. (p. 188)

In order to capture the efforts of systems, it is necessary to harness the available data collected for students with disabilities and transform these data into informed intervention and predictive teaching practice.

The Role of the Special Education Teacher

As we begin the new millennium, it is apparent that the demands and pressures on the special educator are extensive. Special education teachers are confronted with managing student behavior, dealing with regular education systems, and managing large quantities of paperwork to comply with federal and state requirements (Rosenberg et al., 1997). The paperwork burden for these teachers and systems is often extreme. If systems are required to collect data on students, then it seems imperative that knowledge discovery techniques be used to inform teaching and improve learning. Data collection should assist our efforts and not become another regulatory or accountability activity. The entire idea of the IEP and subsequent data collection should be responsive to the complex and daily changes in student growth and educational progress. Data

collection and individualized educational programming should reduce the demands on the special educator, not increase frustration and parent dissatisfaction and litigation.

To make this idea a reality, it is necessary to build a comprehensive data collection and discovery system that is responsive to individual intricacies in the teaching and learning process. This system should provide educators with the information necessary to adapt learning and continually respond to the classroom environment. A broader discourse on teaching and learning and an expansion of current pockets of success should include the greatest number of schools, teachers, students, and parents.

The Learning Community

It has been suggested that school reform is successful and sustainable when administrators, counselors, and teachers share a common vision, build a collaborative culture, and weave changes into the organization of the school (Fullan, 1991; Ruddick, 1991). Peter Senge (1991) comments that within learning communities the success or failure of one person influences others' successes or failures. At this point, there is no single model under which a discussion of teaching and learning can be contained. This should be seen as a positive attribute to past efforts in education. Education is a systemic process that includes many individuals. Larry Cuban, in an historical analysis of teaching from 1890 to 1980, documents changes in educational beliefs and practices over the last 100 years (Cuban, 1993). Cuban's focus is on those "organizational mechanisms that attempt to integrate new ideas into the teachers instructional repertoire," and the fact that many school-wide strategies for instruction remained the same. Cuban's final analysis reveals that many instructional decision making procedures in the regular and special education environment are often not solely the responsibility of the classroom teacher. As the number of student variables increases, the intricacies of determining what to teach and the changing landscape of how to learn becomes increasingly complex. A further effort in school and individual decision-making processes requires our most creative, responsive, and thoughtful attention.

Decision Making and Predictive Solutions

Decision making remains one of the more challenging roles of every special education professional. The importance of bringing expertise to the decision making process is obvious, but the best methods to acquire it may be more illusive. In his article, "The Mind's Journey from Novice to Expert," Bruer (1993) makes a convincing case that:

If we understand the mental process that underlie expert performance in school subjects, we can ask and answer other questions that are important for education. How do students acquire these processes? Do certain instructional methods help students acquire these processes more quickly and more easily? Can we help students learn better? (p. 6)

According to Bruer, it is imperative that as educators, we understand what is needed to move students to success ". . . if we know the route, we can help students negotiate their way" (p. 6). Or put another way, if we know the pattern, we can predict the behavior. But how do we discern the patterns and what do we need to know in order to make a judgment that predicts and shapes the future? It is the next step, decision making, that ultimately makes the difference. Looking to other disciplines that use knowledge discovery techniques, and which apply data mining solutions to systemic problem solving, may define the route and answer important educational questions.

Related Disciplines and Applicable Knowledge

Cognitive science and educational research may not provide a direct answer to how we can solve special educational problems. However, data gleaned from the wealth of research can guide special education practice in much the same way that biology guides medical practice. Granted, there is more to medicine than biology, but it is the basic medical science that drives progress as well as contributes to the doctor's decision making. It is the expertise in decision-making that translates into successful treatment and results in the patient's well being. Today, physicians can sharpen their decisions with technology that allows them to reap the benefits of sophisticated databases, to access electronic medical records, examine patient profiles, and guide

patient recovery. Data mining has become so important in many medical arenas for diagnosis and treatment, that proponents of its use suggest that, “it’s time to write a data management clause into the Hippocratic oath” (Stodder, 1996, p. 7).

Today’s medical services are characterized by complex and interwoven activities that can either contribute to or confound patient treatment. The complexity of the system calls for vigilant monitoring and support of patient treatment and recovery and helping physicians with anticipatory guidance. The added value of predictive technology in the medical field is essential to improving decisions that determine the quality of treatment and service to patients.

One method of applied data mining practice is called predictive data mining. Predictive technology leads to anticipatory guidance. While this concept is not new to medicine, it can now have a more precise role through the addition of predictive technology. Predictive solutions learn from data and their predictive capabilities improve as data experiences increase (Biafore, 1999). By identifying patterns in these complex databases, predictive solutions allow doctors to make more meaningful decisions. Predictive solutions support better medical treatment by gleaning information from all available data and by objectively looking across many data variables.

Predictive technology allows analyses of variables and interactions between variables to recognize complicated and subtle patterns, patterns that are too complex to be expressed using traditional techniques that deal with individual variables in isolation. Intelligent knowledge extractions use predictive solutions that include comprehensive behavior profiles of individual entities, which in turn are the inputs that fuel predictive models (Aigner & Aigner, personal communication, July 24, 2001). The models produce the information needed to make better decisions and quickly and expeditiously identify interestingness in a multitude of patient data.

The use of anticipatory guidance is supported by applying information from patient profiles and patterns that have been identified by data-mined applications of patients and their treatment. For example, when a treatment and/or medication is prescribed to a patient, information is available about specific side effects or other consequences that could occur such as nausea, a rash, or some other expected symptom. Having access to this guidance allows the

patient to feel better prepared to deal with the condition and, more importantly, to know that certain reactions are anticipated. Providing this type of guidance also helps patients feel more accepting of the outcome of their treatment and leads to better doctor-patient communication. An added value of anticipatory guidance is in its contribution to empower patients to participate in and take greater responsibility for their recovery.

The success of this guidance is built on the decision-making expertise of the doctor. It is further supported by the doctor's confidence in what he/she judges to be the best match of medication and treatment. For most doctors, their decisions are based on definitive examination and lab results that help support, if not confirm, a diagnosis. In addition, they have a selection of researched treatments matched to patterns of symptoms, database information that provides both patient and drug profiles, and specialized database findings from pharmaceutical companies that provide an immediate alert for potential drug interactions, allergies and/or risk factors.

At the health care systems level, data-mining efforts are in place to qualify symptoms before illness sets in and to help physicians instill best practices. Data-mining applications can identify which treatments work, determine which patients should receive attention earlier, reduce expensive medical treatment, and ascertain the interventions that can help avoid treatment complications. In some systems, data are gleaned from patient records and result in specific patient profiles. In other systems, data miners work with derivatives from reported patient care rather than from individual patient records. Data are also collected to identify patterns in doctors' treatment profiles to check whether they are encouraging patients with certain chronic illness to have certain blood tests or other treatment on a regular basis. Within the health care industry, data mining is becoming one key to preventive care.

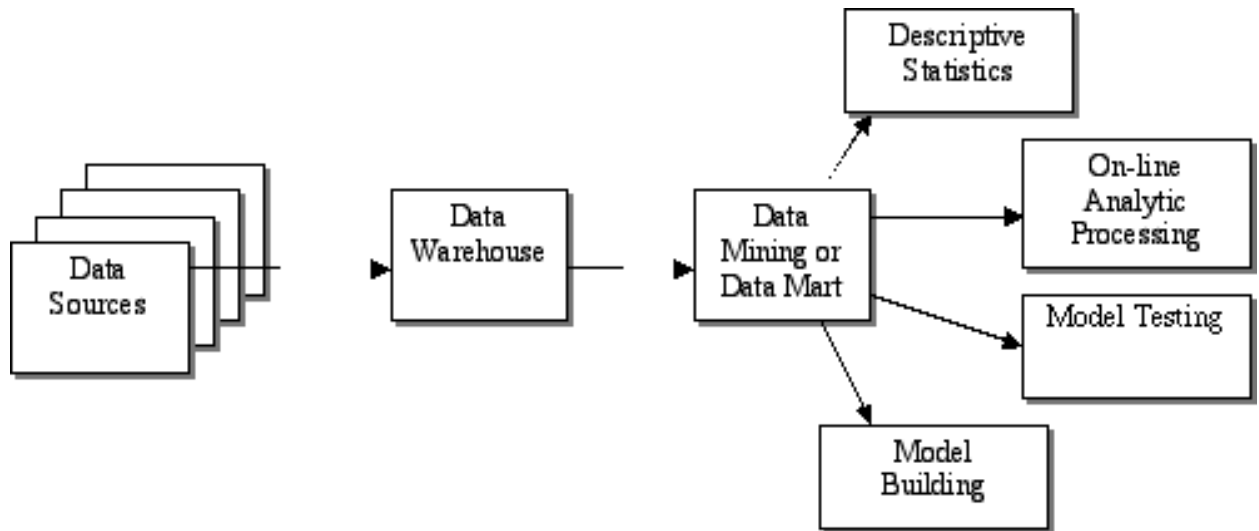
Data Mining Examples from Industry and Public Health

Over the past three decades, computers have been used to capture details of business transactions such as banking and credit card records, retail sales, and telecommunications. The data from these transactional systems leave a trail of the key trends that impact various aspects of each business such as products that sell together, sources of profits, factors that affect

manufacturing quality and other factors that correlate. Data mining has evolved from sifting these relational databases for information to fine-tuning marketing efforts. These databases are mined for the right information, as the computer applies specific statistical and computational applications that sift through data, detect patterns, and distill the results.

Many for-profit and non-profit agencies and companies make extensive use of data mining technologies, such as credit card companies, health maintenance organizations (HMOs), and grocery stores. One highly successful example of data mining is direct mail marketing via unsolicited, and often-expensive gift catalogs promoting special sale items. These mailings offer items that correlate to the interest and buying patterns of a specific customer set. When the customer responds with an order, the purchasing choice, method of payment, and other consumer behaviors are observed, collected as data, and then mined to determine future buying potential. The same scenario occurs when HMO's mine medical records to determine individual and group usage patterns that will then be used to formulate patient benefits and medical coverage. Grocery stores use data mining application to determine product shelf location, to select weekly specials, and to predict customer traffic and product availability and to determine employee schedules. These decisions result from information generated by scanning product bar code data, then collecting and mining the data to identify customer profiles and buying patterns.

Data-intensive businesses and organizations use data mining to extract value from massive amounts of transactional information they process daily. Businesses are expected to invest \$200 billion in data warehouse technology between 1997 and 2001 in order to turn this operational data into information that will ultimately increase their bottom line. Data mining is among the most important applications of a data warehouse since it uses a variety of data analysis tools to discover patterns and relationships hidden in the data. The last step in the data mining process involves the extraction of interpretable information that can be used to make valid predictions through the use of descriptive statistics, on-line analytical processing, and model testing and building (see Figure 1).



While there are many steps in data mining, the first and simplest step is to describe the data by summarizing its statistical attributes. Based on patterns found from the data, a predictive model is built which is tested outside the original sample. We are warned, “a good model should never be confused with reality (you know a road map isn’t a perfect representation of the actual road), but it can be a useful guide to understanding your business” (Edelstein, 1997, p. 14). The final step is to empirically verify the model to raise confidence in its predictive value.

Data Mining

The goal of most data mining processes is to enhance organizational, professional, and business practices. The potential for educators to use available data for analyzing and predicting the impact of programs, classroom practices, and student progress is becoming easier with the advent of new tools, from available statistical software programs useful for research, to Internet-based educational plans, database software applications, electronic IEP software, and using data warehousing techniques. An additional, but yet unexplored new tool available to do transactional predictive modeling for education practice is data mining.

The purpose of engaging in data mining is the same, to look for new patterns in large amounts of data to turn new knowledge into actions. These actions might include improved delivery of services, enhanced information about market trends, improved educational technology software products, or in the case of special education, influencing curriculum

development, monitoring student growth, and enhancing academic progress. The same model could apply to the recruitment, training, and retention of special education teachers.

Assumptions About Data Mining

For the purpose of data mining, data are any facts, numbers, or text that can be processed by a computer. Today, organizations are accumulating vast and growing amounts of data in different formats and different databases. Databases can range in size from a few megabytes to terabytes of data. Data collected within business or market driven organizations includes:

1. Operational or transactional data, the fundamental building blocks such as, sales, cost, inventory, payroll, and accounting,
2. Non-operational data, such as industry sales, forecast data, and macro economic data, and
3. Meta-data (data about the data itself), such as multidimensional and logical database design or data dictionary definitions.

Data mining supports knowledge management (KM). Through KM, it is possible to (a) seek to minimize information overload and (b) focus on how to get the right information into the right hands at the right time and in the right form.

The intent is to develop system's applications that will (a) filter out information that does not support an immediate or long term need or goal, (b) present information in the form appropriate for the context and the audience needing the information, and (c) make information accessible to those who can act upon it and give it value (Cowley-Durst, 1999).

Data Mining in Process

While large-scale information technology has been evolving separate transaction and analytical systems, data mining provides the link between the two. Data mining software applies algorithms to detect and identify apparent patterns and anomalies and then analyzes the relationships and patterns in stored transaction data based on open-ended user queries. Several types of analytical software are available such as statistical, machine learning, and neural networks. Generally, four types of relationships are explored.

1. Classes: Stored data are used to locate data in predetermined groups.
2. Clusters: Data items are grouped according to logical relationships or consumer preferences.
3. Associations: Data are mined to identify associations.
4. Sequential patterns: Data are mined to anticipate behavior patterns and trends.

These relationships can be represented in the form of decision trees, neural networks, rule induction, nearest neighbor, and genetic algorithms.

Today, data mining applications are available on all size systems for mainframe, client/server, and PC platforms. System prices range from several thousand dollars for the smallest applications up to \$1 million a terabyte for the largest. Enterprise-wide applications generally range in size from 10 gigabytes to over 11 terabytes. There are two critical technological drivers.

1. Size of the database: The more data being processed and maintained, the more powerful the system required.
2. Query complexity: The more complex the queries and the greater the number of queries being processed, the more powerful the system required.

Assessing the potential of investing in data mining tools requires that systems investigate the status of relevant databases and examine how much data is available for mining. In some instances, it is necessary to build these databases from scratch or clean existing data sources so they are useful for applying data mining tools.

Magnifier Effect: Data Mining and Knowledge Discovery

Although data mining and knowledge discovery are designed to reveal patterns in data by sifting out nuggets of information from masses of unanalyzed or under-analyzed data, it is the next step that transforms these mined nuggets into gold. Interestingness involves identifying specific interactions that produce a magnifying effect on whatever is involved in the experience or experiment (Piatetsky-Shapiro, 2000). While there are many ways to examine the concept of interestingness, there are some examples that help put its impact in perspective.

In the early 1980's, it was discovered that children recovering from influenza and chickenpox could become severely or fatally ill if they also took aspirin. Taking aspirin interacted with, and magnified some aspect of the original illness and brought about a new and more serious illness, Reyes Syndrome. The interestingness in the data centered on the magnification aspect of aspirin interaction. Less traumatic interactions can occur when grapefruit juice is taken with some medications. Certain combinations can magnify the impact of the medication and are therefore recognized to have interestingness features.

One of the most compelling and dramatic examples of interestingness in data is found in the medical breakthroughs from the Human Genome Project. This critically acclaimed effort supports the cataloging of tens of billions of bits of data into the online gene sequence repository known as GenBank. However, DNA sequences aren't the only types of data being mined. Scientists are using "DNA chips" to detect patterns as thousands of genes are being turned on and off in a living cell. These data are then added to the flood of findings that are becoming available at a "mind-blowing pace" (Regalado, 1999, p. 58). The urgency of this research challenges pharmaceutical companies to refine data-mining tools to expedite pattern recognition and investigate potential magnifier effects. Finding hidden structure and interesting interactions in massive amounts of molecular data can contribute to "gene expression profiles" (Regalado, 1999, p. 59). The added value of data-mining is the ability to identify interactions that possess sufficient interestingness.

There are also some near misses in the medical field that may have been more quickly accepted if stronger data mining tools had been applied, such as the refusal of the medical establishment to acknowledge that peptic ulcers are, for the most part, not caused by stress, smoking, alcohol, and genetic predisposition (Cowley, 2000; Mirsky, 2001). This wrong thinking continued even when Australian researcher, Barry Marshall risked his own health by "swilling beakers of bacteria" and giving himself ulcers and then proceeding to cure himself with antibiotics (Cowley 2000). Building on Marshall's and others research, biologist Paul Ewald (2000) is shaking up the medical community with a new book, *Plague Time*, in which he

documents bold assertions that the primary causes of today's slow-burning plagues are parasites (i.e., viruses, bacteria and other infectious microbes) whose long-term effects we have simply failed to recognize (Cowley, 2000). While some experts doubt Ewald's beliefs, his recent findings that examine the interestingness of interactions, are of great interest to others. Ewald has identified several patterns: (a) the link of childhood strep infection to obsessive-compulsive disorder and Tourette's syndrome, (b) traces of a virus that caused mammary cancer in mice have been recovered from human breast tumors, (c) the link of borna virus, a brain infection in horses, sheep, and cats, to schizophrenia and bipolar disorder in people (this has been supported by researchers in Japan and Germany), and (d) data that a "growing body of evidence supports the theory that Chlamydeous pneumonia, a common respiratory bug, may play a key role in coronary artery disease" (Cowley, 2000, p. 67).


Perhaps one of the most futuristic data-mining applications is a rather clunky replica of the old Ovaltine Decoder Ring with a Dick Tracy/James Bond twist. The ring was christened Monitor 911 and does just that; it gathers vital sign measurements and beams the data to just about any device, whether it's a computer or a Palm Pilot (Levine, 2000). In addition, the ring alerts patients and their doctors about any data that raise a red flag or identify a risk pattern, which in turn, will be followed with emergency protocols. Data from the ring are collected and secured on a Web site where it can be stored and analyzed. This ingenious device uses wireless technology to circumvent hospital and doctor office barriers as well as providing one of the best examples of just in time, just enough, and just for you collection and delivery of data.

Data Mining for Schools

If data mining allows the business and medical sector to accurately identify and predict the most effective way to deliver their respective services, shouldn't educators understand how to use their data to match (a) instruction to student, (b) teacher style to learner needs, and (c) learning environment to student requirements? Shouldn't administrators mine school or district data for the characteristics of the most successful classes, teachers, pedagogy, schedule, or other variables that identify a significant pattern? While data mining should be applied at all levels of

education, there are additional legal and ethical requirements for teachers and support service personnel to address issues of universal access and individual accommodation. Students bring additional variables to the classroom and it is sometimes difficult to sift through patterns of interestingness, as described above, and use these patterns to inform instruction.

Guttentag and Clark (1999) have documented the impact of information overload in a variety of education settings. Their findings sound an alert that schools are drowning in the hundreds of daily transactions including attendance, student and teacher schedules, and individual student assignments. A student carries with him or her historical data, random sets of demographic information, standardized test scores, unique family situations, health records, extra curricular participation, and grades. Each student also interacts with teachers, administrators, academic specialists and career counselors who in turn have their own data histories with details such as salary, educational attainment, certification and professional development in addition to regional, school-based, and student domains (see Figure 2).



Effective instruction often requires consideration of variables within the previously mentioned domains.

1. Student – Type and severity of disability, gender, academic performance, emotional status, behavioral and performance data, IEP goals and objectives, and daily attendance.

2. Teacher – Degree, professional development training, number of years teaching, and certification
3. Educational – Number of faculty and students, available software and curriculum resources, curriculum offerings, class size, student computer ratio, and school day structure
4. Societal – Workforce data, rural or urban setting, and regional literacy rates.

Many of these variables are captured by schools or at the state level for accountability purposes. These data are not typically captured electronically, in the same database, or available for analysis, especially at the classroom level. Currently most schools collect and process data in separate systems, but need to report and analyze it in one place. Despite recent claims by some vendors, at least for the foreseeable future there is the need to run multiple applications. Even applications from the same vendor are rarely fully integrated. In most educational systems, the data is in relational databases rather than the multidimensional databases. There is no application currently available that will aggregate information such as student attendance, course information, curriculum goals, and individualized education plan data.

Schools find themselves under increasing pressure from parents, school boards and politicians to meet educational standards set by state, regional, and national authorities. As students in special education increasingly work with regular education content, school accountability for special education students could be strengthened. With the emphasis on school improvement and performance accountability, schools must understand how teachers and school administrators can best use curricula and educational programs to ensure students receive the quality education they deserve. Even the most basic K-12 data warehouse strategy could alleviate a district's need to painstakingly aggregate data from different systems (e.g., student information systems; financial management decisions) for local, federal and state reporting. It could increase efficiency, enhance accountability, and perhaps save money better spent on instruction.

Current State of Educational Practice for Data Mining

In Missouri, the Department of Elementary and Secondary Education is developing their data warehouse with a focus on data reporting. The greatest strength of their data warehouse is providing users with the ability to retrieve data themselves rather than having to contact someone from the Department of Education. With an established data warehouse, a district can begin answering questions such as:

1. How much does a particular program cost?
2. Do higher costs correlate with effectiveness?
3. Why is one class outperforming other, similar groups of students?
4. As planning and forecasting improves; it furthers development of accountability frameworks.

South Carolina's Beaufort County Public Schools is using their data warehouse to determine the most appropriate student groupings.

Under the old model, students were locked into a track: College prep, general, or vocational education. Nowadays, we use continuous progress – continually challenging the student. Now when a student has mastered first grade reading, he doesn't spend the next two or three months wasting time. He moves on to second grade reading. He may be behind in math, so he will get extra instruction during related arts time – whatever he needs for continuous progress and keeping the child interested in school (Guttentag & Clark, 1999, p. 4).

Dennis P. Doyle (1997), author of *Raising the Standard: An Eight Step Action Guide for Schools*, explains that as a result of the significant data analysis that led to this new approach to student groupings, "district scores are climbing robustly all over (p. 135)." Not only will Beaufort Public Schools be using their warehouse to facilitate continuous progress, they will be able to explain how each and every dollar is being spent within each of their schools (Guttentag & Clark, 1999).

When a district has developed a mature warehouse environment, with a substantial store of longitudinal data, the power of data mining can be exploited. As described earlier, the data mining process is used to find high-value information that has been extracted from large volumes of data through the development of various statistical models. Data mining technologies allows for sifting through of enormous amounts of data for unusual patterns, surprising trends and interesting outcomes. With this information, a district can then move beyond simply answering questions posed by educators and researchers and begin to uncover key questions previously not even considered.

Dr. Phil Streifer, University of Connecticut professor and former superintendent, has been an advocate for data-driven decision-making for over a decade. He has a vision for the power of data warehousing and in particular, data mining. Student achievement in his school districts were drastically improved through a systems integration process with benchmarking that integrated key areas, such as strategic planning, program evaluation, personnel evaluation, and supervision, staff development and training, and leadership and organization.

While data mining is powerful, it is not magic. Data mining assists analysts with finding patterns and relationships in the data, it does not place a value on the patterns to the organization. To be of value the patterns must be verified in the real world. It is important to remember that the predictive relationships found through data mining are just that, relationships, and not necessarily the causes of an action or behavior. Data mining will provide highly refined information but it will not automatically discover solutions without guidance -- and it does not replace the skills of an analyst or researcher. There is another limiting element to data warehousing that can pose a serious problem, that is, the questioning ability of its users. However, Streifer (2001) and his colleagues suggest advances in data mining technology will eventually remove even that barrier. "Down the road, the warehouse will be predictive; data mining will uncover complex relations that we can only speculate upon now, but will be a reality a year or so down the road."

With 16,000-plus students, the Tyler independent School District is the largest public district in east Texas. Since 1997, the district's assessment and accountability group has

measured student performance and evaluated educational programs. To improve the future of education, the quality of its programs and ensure future community support, Tyler must measure student learning, track educational progress and target areas for improvement through several processes.

1. Expanding relationships with current students and parents,
2. Analyzing test scores of 16,000 students in grades K-12,
3. Developing reliable, data-based methods to evaluate program effectiveness,
4. Winning more state and federal grants for increased funding, and
5. Communicating clearly and often to keep community support high.

Using data mining applications, the district analyzes individual student test scores and other data on its educational programs. Sharing results with key district decision makers helps Tyler better identify strengths and weaknesses to improve educational effectiveness. It also increases opportunities for future funding from public and private sources, and strengthens community support. Tyler was successful in several areas.

1. Established a district-wide system to track student performance,
2. Reduced labor costs with a streamlined analytical system,
3. Forged a path for stronger educational programs,
4. Increased the likelihood of winning external funding, and
5. Strengthened the community support and communications,

While these systems differ in location, student composition, leadership, and other factors, they all used data mining to provide the information needed to improve their schools. The following section illustrates how data mining might be used meet current needs and provide anticipatory guidance for the future.

Applicability of Data Mining to Special Education

The issues inherent in the special education arena are not quite the same as business and medicine. Special education teachers are continually faced with individualized teaching and learning considerations and are required to continually integrate prescriptive educational

treatments into the IEP and their curriculum. Teachers are rarely prepared to fathom what type of best practice should be implemented for the children and youth they service. Guidance for instruction may be in the form of traditional IEPs while other teaching considerations flow from a constellation of professionals addressing issues related to test scores, observational data, interviews, reports, and more. In some situations there may be information relative to the specific educational environment or a description of the therapeutic milieu. There is no shortage of data but there is often a lack of requisite information for quality decision making to take place.

Predicting behavior is critical in the development and implementation of successful teaching-learning strategies. In special education, the array of information available will not have been mined to a level that would allow decision makers to discover potential patterns or to develop profiles that could be used to provide a base for anticipatory guidance of both teacher and student. In many cases, much of the information in student files, regular or special education, may not be reviewed for the purpose of determining how to improve the teaching-learning environment, reinforce and continue past successes, identify what strategies are most successful, or analyze how best practices in similar classes could be replicated.

Without adequate access to, or practitioner-friendly translation of student data, special education professionals' decisions are sometimes compromised by information isolation. The diminished quality of these decisions directly influences factors relevant to selecting and implementing instructional strategies, learning environments, electronic technologies, and other critical teaching-learning strategies. The digital presence and mining of these data could change decisions that would dramatically impact the dynamics of teaching and learning in special education settings. Most importantly, knowledge discovery could serve as a catalyst to significantly enhance the lives of the children with special needs.

Special Education Scenario

Steven is in the 11th grade and participates in the regular education curriculum. Steven has cerebral palsy and uses an augmentative communication device (ACD) with a built in data collection program called a language activity monitor(LAM). He uses the ACD to communicate

and interact with other students and teachers. The LAM function records each language event (e.g., letters, words, phrases), as the ACD is used, and assigns a time code, which is available for download on a daily basis. In order to meet his individualized educational goals, the school has invested in universally accessible computer programs to be used in conjunction with classroom textbooks. All digitized assessments, end of chapter questions, and student worksheets link into the school data mining system. Other students within the special education program use computer emulation keyboards to access digital text and supplemental materials off the Internet. Students are actively involved in computerized assessments stemming from IEPs. Students' passwords and IEP are housed on the schools local area network, and as a result, class work, tests, and quizzes are all entered in a database and matched to individual objectives as students progress through the semester.

On Friday, the knowledge management specialist, working with the special education chair, pulls up the data mining application and clicks analyze. The data mining tool accesses students' IEPs and other individualized student data, searches through current and historical classroom performance data, such as speech and language activities, student characteristics, and academic tracking information available from historical records and current classroom activities. The screen indicates that most of the students have met their goals for the week, but that overall, student work across specific course content has fallen slightly. The tool recognizes, through available search domains, that student work across content areas remains strong. However, the data indicate that all quizzes related to word problems show abnormally high failing grades. The knowledge specialist and his classroom teacher look at the data visualization chart. They notice that many of the students had difficulty with prior coursework in word problems and that it always involved the same information-processing skill.

The teacher decides that two afternoon tutorial sessions are necessary. The department chair orders math software to be used as support during the sessions. At the next series of IEP meetings, the chair informs the team of the tutorials and current student progress. Data from the networked version of the math software is automatically entered into the school's online data

mining system. The knowledge manager continues to monitor student performance data and has the data mining application continually analyze which students continue to need support with the afternoon tutorial sessions.

School Staff Development Training Scenario

Mr. Lewis is a special education supervisor in the area office. He looks at the teacher data available within his database each month and runs charts to explain how many vacancies exist. He uses the information to allocate teachers to schools based on teacher-student ratios. Mr. Lewis is also responsible for allocating training dollars within the system. He usually targets those teachers who are working with severe populations and who are working on initial teaching certification. Available data-visualization charts inform Mr. Lewis that schools in certain areas of the county are having difficulty with academic performance on required state assessments. As a result, he has been targeting teachers in these schools for training on a new math initiative, involving recently acquired math software. He has hired two master teachers to develop online training materials that are linked to the state math content standards. He has required that all teachers in these schools participate in 20 hours of ongoing inservice training over the two fall quarters.

After the first quarter, Mr. Lewis accesses available data on teacher certification status, previous coursework completed to date, number of online hours of training teachers have completed, and their performance within the course. The data mining tool compares teacher rates of completion in the online course with historical teacher professional development data. The computer finds a pattern between teacher technology uses, personal skills level for integrating technology in the classroom, and overall student progress within math courses at the targeted schools. A data visualization chart indicates that 20 special education teachers have limited or no background in technology integration although they have all been teaching for 10 years or more. Mr. Lewis uses the information to speak with the staff development coordinators and school principals. The administration updates the school's professional development model and

institutes a separate unit on the integration of instructional math software. Staff development personnel work together with teachers in the classroom to integrate model strategies for teaching.

After teachers complete the required hours of training, another data visualization chart is produced in the professional development office. Mr. Lewis considers the available data and indicates to principals that the teachers have completed the required training. The principals agree to look at student performance data charts at the end of the month to analyze progress and allow for teacher release time to continue staff development training as needed. At the end of the month, school principals look at student performance data, which indicates progress toward state standards. Based on teacher performance data and an updated model, Mr. Lewis is able to anticipate the same issues with teachers involved with training in another school and makes sure his staff integrates the new module into the training. At the end of the year, increasing performance data indicates the training has been successful. He communicates with the superintendent about the new model for teaching and learning. The superintendent decides that this training should be available to all new teachers beginning in the fall.

Considerations for Future Practice

Although the level of specificity for data collection in the above scenarios may be extreme, it is necessary to acknowledge that the dynamics in special education have become increasingly detailed and the management of student information has become unmanageable. As new knowledge management and discovery systems are integrated in schools, the level of specificity will be determined according to need. In the future, data mining will be used to anchor the role of information in providing anticipatory guidance to the field of special education. In order to realize how special educators could adopt new knowledge discovery systems, several tasks need to be undertaken.

1. Bring together information from the data mining and the special education community in order to understand the current capabilities and educational possibilities available through data mining and knowledge discovery.
2. Anticipate the possible future of data mining for special education,

3. Conduct extensive research on schools implementing data mining systems, and
4. Identify requirements for policy development to support data mining priorities and research.

Special education practice will continue to be influenced by innovations that occur from related disciplines. New ideas from other disciplines will need to be analyzed in order to decide the appropriate use of these innovations for future practice. Based on trends and new technology applications in the field of information technology, it is possible that the following will be true.

1. Available data warehouses within education students will include specific domain data collected from school, home, and work environments.
2. Data mining applications will be adapted or invented to specifically manage the special education data warehouse.
3. Data mining will help educators “magnify” instructional practice, program development, and will inform new and existing teaching and learning models.
4. The IEP will become a transparent document and will be transformed into effective models of practice beneficial to all students.
5. The development of new and universally acceptable instructional and assistive technologies will feed into a common data warehouse.
6. The use of Web-based learning will transform education into a seamless environment where data is collected anytime and anywhere on individuals as needed.

Privacy Safeguards

The development of regulations, research, and practice to back knowledge networks needs to be instituted. At this point, there are few, if any, requirements for the ethical use of online data. Currently, there are no recognized standards or regulations that specifically address privacy issues for data warehouses in the educational community. Organizations, such as the Organization for Economic Cooperation and Development (OECD), have responded to privacy issues by developing a regulatory framework for the ethical use of data. In addition, business and public health industries are working to ensure appropriate access to and use of confidential data.

For example, it has been suggested that as the information available within society becomes more pervasive, privacy concerns will become more important and impact the way data can be used and analyzed. It also seems appropriate to exclude data, such as race, gender, and age in the decision-making process (Clarke, 1999). Within the field of special education, information access may be limited as it is now, to special education teachers and members of the IEP team. Student performance data may be open to all educators within a specific designated community or based on the individuals right to make decisions. It will become increasingly important to minimize information overload, getting the right information into the right hands, at the right time, and in the right form. The intent is to develop a systemic process across educational environments, specific to the audience needing the information. Protecting special education students from the potential misuse of data will become an increasing challenge for educators.

We have learned many things about the privacy issues given the expansion of interactivity on the Internet. It is painfully clear that appropriate protections need to be in place to deal with those individuals who are predators or would capitalize on marginalizing certain groups of individuals. For those individuals who might use existing data for intentions not even thought of, safe harbors, as well as federal, state, and local privacy and protection acts need to be instituted for all participants in the educational community. The Departments of Education and Commerce must work together to ensure standards for privacy and protection.

One of the most significant concerns with the use of student learning profile data is that information collected and disclosed should only be permitted in certain specific circumstances with well-defined, clear guidelines for disclosure. Collection and disclosure must be the exception, not a mechanism for routine use. Although privacy safeguards exist for student data, the development of student learning profiles and the value and effect of the information requires the privacy issue to be closely evaluated and monitored. The development of privacy rights reflects the necessity of well-defined ethical standards that embodies the goal to do the right thing.

Important issues such as notice, choice, access, and enforcement must be crafted to protect students' privacy and secure their data. While there is no perfect model, the Patients' Bill of Rights legislation may offer an opportunity to craft the development of a more comprehensive structure to ensure students' privacy rights. In addition to privacy safeguards, it is critical to have security provisions with appropriate firewalls and other technology protection to ensure data safety, access, and use (Tretick, 2001). The work of the Computer Ethics Institute and collaboration with data mining companies specializing in advanced security design, such as ASC Database Marketing Corporation, are other potential models and resources as schools consider data mining tools. Such a collaboration could be piloted to identify how to customize privacy safeguard systems to test the development and use of student learning profiles.

A Call to Action

One way to re-structure the way data is documented in schools is to evoke alternative thinking strategies about existing resources, teacher training, and personnel needed for the future. Business and financial communities have used data management systems, like data mining, to deal with similar issues. The solutions for data management are growing. For schools, *application server providers* are currently being marketed as low cost solutions available to deal with advances in networking. As this market evolves, we need to prepare individuals to institute these systems. New methods of knowledge management have encouraged differential approaches to change. In special education, we must identify new forms of leadership for teachers that use available technology for systemic change. This type of change can only be acquired by having systems that advocate for knowledge management, with the teacher trained as the facilitator of that change and held in regard as an expert on instructional strategies and student variables. In the new order of special education systems thinking, teachers could be prepared as specialists focused on specific areas of expertise or interests. In this regard there needs to be a new title to identify the new role. A new title will do more than identify the role; it will serve to recruit a new type of professional into service. In an era of dwindling numbers of special educators, this is of some significance.

The new *knowledge specialists*, as they may be called, will benefit from a professional development experience based on some of the most successful models from medical and corporate sectors. All special education teachers who become knowledge specialists will have to complete rotations in basic and advanced educationally challenging settings. Like their resident physician counterparts, the knowledge specialists will be immersed in level 1-5 education and treatment environments and will have extensive experience in at least two settings. In addition, like hospital-based training, the knowledge specialist's skills will be sharpened through a required series of grand rounds led by content, behavior, instructional, or other experts who use a case study approach. If face-to-face opportunities do not exist, virtual grand rounds and rotations will be available. The virtual environment will include field experiences with all learner types in a variety of instructional and treatment settings. Knowledge specialists will contribute to a repository of "transportable" lesson plans, teaching-learning strategies, and other related instructional materials. All of this will be available online and downloadable for others to customize as needed.

Knowledge specialists will be prepared as teacher mentors to help colleagues move from a one-size-fits-all instructional model to highly individualized instruction matched to student learning profiles. All of this information, including the learning profiles, will be available as part of each student learning card which will be based on a smart card concept. The learning card will hold important student information, medical needs (e.g., asthma, seizure medication directions), learning strategies (e.g., kinesthetic, concrete, call on last-needs time to rehearse), teaching dynamics (e.g., limited oral lecture linked to visuals, immediate feedback, collaborate projects), and other information that is critical to student success. Since the learning card will be continuously updated, it will allow the new teacher, substitute, or other related professional (e.g., speech language pathologist, occupational/physical therapist) to be kept apprised of each student's status, needs, and progress.

Tomorrow's special educators will be faced with increasingly complex teaching-learning issues that will require collaboration across all education related disciplines. Fortunately, there

are successful models that can guide us in developing solutions to meet these unique needs. We can begin by identifying respected and recognized institutions such as the Brookings Institution, the Aspen Institute, and the Rand Corporation, that have developed think tanks to address issues of national importance. We can use their blueprints to guide the development of an education think tank to provide a repository of information and resources to identify and support special education research, teaching practices, student learning, partnerships, instructional design, assistive technology, and more.

The contributing disciplines could include education/special education, computer science, biomedicine, cognitive science, arts, and management. Delivery channels will represent more traditional and emerging systems: (a) Web-based electronic learning communities, (b) Web-symposia and dedicated Web interest sites, (c) virtual reality, (d) online and traditional publishers/vendors, and (e) hardware and software developers. The audience will include system and school based special educators/technologists, higher education, health care, parents, and policy makers. The special interest areas will represent research, curriculum and instruction, partnerships, metrics (e.g., evaluation/assessment) and will help us develop a systems approach, supported by data-mining and knowledge management tools. This will provide powerful new decision-making strategies to special educators faced with increasingly complex teaching and learning challenges to making high quality decisions.

Conclusion

Based on the dynamic nature of information and data, it seems reasonable to assume that the role of knowledge managers will be central to the success and growth of most institutions and industries. The field of special education will need to reengineer the environment (a) to take advantage of the exponentially growing data to be mined for key decision making information and (b) to develop knowledge managers at the system level to expedite information and knowledge sharing to generate individual education plans for every student.

This paper has presented some ideas and strategies that reflect practitioner, researcher, administrator, and student interests. The goal of this article is to synthesize the success within the

field of special education, systems change, business, medicine, and technology and motivate the current educational system to consider alternative movements and improve the lives of students with disabilities. Both special and regular education professionals must meet the challenge to erase the margins that prevent children with special needs, teacher, and schools from achieving their full potential. Like the glass ceiling that marginalized women and minorities in the corporate structure, the margins within education sometimes keep students with special needs, their teachers, and entire schools on the fringe of effective teaching and learning practice. Erasing the margins will happen only when we actualize new methods for change by preparing all teachers to meet the educational needs of all students and working with systems to facilitate new models for education.

References

- Biafore, S. (1999, November) Predictive solutions bring more power to decision makers. *Health Management Technology*, 10(20),12.
- Bruer, J. (1993). The mind's journey from novice to expert. *American Educator*, 17(2), 6-15.
- Clarke, R. (1999, February). Internet privacy concerns confirm the case for intervention. *Communications of the ACM*, 42(2), 60-67.
- Cowley, G. (2000). The real 'hot zone'. *Newsweek*, 136 (22), 66-67.
- Cowley-Durst, B. (1999). Gathering knowledge for your knowledge management system. *Performance Improvement*, 38(7), 23-27.
- Cuban, L. (1993). *How teachers taught: Constancy and change in american classrooms, 1890-1990* (2nd edition). NY: Teachers College Press.
- Doyle, D. and Pimentel, S. (1997). *Raising the standards: An eight-step action guide for schools and communities*. Thousand Oaks, CA: Corwin Press.

- Edelstein, H. (1997). Data mining: Exploring the hidden trends in your data. *DB2 Online Magazine*. Retrieved from <http://www.db2mag.com>.
- Ewald, P. (2000). *Plague time*. New York, NY: The Free Press.
- Fullan, M. G. (1991). *The new meaning of educational change* (2nd edition). New York, NY: Teachers College Press.
- Guttentag, S., & Clark, S. (1999, September). A brief lesson in data warehousing. *Converge Magazine*, 2(9), 2-10.
- Levine, S. (2000). With this ring, I thee MED. *Business Forward*, 19, 32-34.
- Mirsky, S. (2001, May). A host with infectious ideas. *Scientific American*, 284(5), 32-33.
- Piatetsky-Shapiro, G. (2000, Feb.) Knowledge discovery in databases: 10 years after. *SIGKDD Explorations*, 1(2), 32-34.
- Pugach, M. C. (1992). Uncharted territory: research on the socialization of special education teachers. *Teacher Education and Special Education*, 15(2), 133-147.
- Regalado, A. (1999, September/October). Mining the genome, *Technology Review*, 56-63.
- Rosenberg, M. S., Griffin, C. C., Kilgore, K. L., & Carpernter, S. L. (1997). Beginning teachers in special education: A model for providing individualized support. *Teacher Education and Special Education*, 20(4), 301-321.
- Ruddick, J. (1991). *Innovation and change*. Philadelphia, PA: Open University.
- Senge, P. (1991). *The fifth discipline: The art and practice of the learning organization*. New York, NY: Doubleday.
- Streifer, P. (2001, April). The 'drill down' process. *The School Administrator* [Web edition], American Association of School Administrators. Retrieved from http://www.aasa.org/publications/sa/2001_04/streifer_april%202001.htm.

Stodder, D. (1996). High-tech hipporates. *Database Programming & Design*, 11(9), 7.

Tretick, B. (2001). Can you keep a secret? *Intelligent Enterprise*, 4(1), 68-69.

Zeichner, K, and Liston, D. (1991). *Teacher education and the social conditions of schooling*. New York, NY: Routledge.