Mining Meaningful Patterns from Students' Handwritten Coursework

James Herold
Department of Computer Science
University of California, Riverside
900 University Ave
Riverside, CA 92521
jhero001@ucr.edu

Alex Zundel
Department of Mechanical Engineering
University of California, Riverside
900 University Ave
Riverside, CA 92521

Thomas F. Stahovich
Department of Mechanical Engineering
University of California, Riverside
900 University Ave
Riverside, CA 92521
stahov@engr.ucr.edu

ABSTRACT

A key challenge in educational data mining research is capturing student work in a form suitable for computational analysis. Online learning environments, such as intelligent tutoring systems, have proven to be one effective means for accomplishing this. Here, we investigate a method for capturing students' ordinary handwritten coursework in digital form. We provided students with Livescribe™digital pens which they used to complete all of their homework and exams. These pens work as traditional pens but additionally digitize students' handwriting into time-stamped pen strokes enabling us to analyze not only the final image, but also the sequence in which it was written. By applying data mining techniques to digital copies of students' handwritten work, we seek to gain insights into the cognitive processes employed by students in an ordinary work environment.

We present a novel transformation of the pen stroke data, which represents each student's homework solution as a sequence of discrete actions. We apply differential data mining techniques to these sequences to identify those patterns of actions that are more frequently exhibited by either good- or poor-performing students. We compute numerical features from those patterns which we use to predict performance in the course. The resulting model explains up to 34.4% of the variance in students' final course grade. Furthermore the underlying parameters of the model indicate which patterns best correlate with positive performance. These patterns in turn provide valuable insight into the cognitive processes employed by students, which can be directly used by the instructor to identify and address deficiencies in students' understanding.

1. INTRODUCTION

Educational data mining has typically been applied to data extracted from students' interactions with Intelligent Tutoring Systems (ITS) and Course Management Systems (CMS). This research has been used to improve the way students interact with these interfaces and has led to a better understanding of the ways students learn when using these systems.

In this study, we investigate a method for capturing students' ordinary, handwritten coursework in digital form. In the winter quarter of 2012, undergraduate Mechanical Engineering students were provided Livescribe™digital pens with which they completed all their coursework. These pens record students' handwriting as time-stamped pen strokes, enabling us to analyze not only the final image, but also the sequence in which it was written.

We have developed a novel representation of a student's handwritten assignment which characterizes the sequence of actions the student took to solve each problem. This representation comprises an alphabet of canonical actions that a student may perform when solving a homework assignment. Each action is characterized by its duration, problem number, and semantic content. This representation allows us to apply traditional data mining techniques to our database of students' handwritten homework solutions.

Our analysis focus on two separate groups of students: those who scored in the top third of the class on exams, and those who scored in the bottom third. We applied a differential data mining technique to the sequences of each of these groups and identified behaviors that are more frequently exhibited by one group than the other.

These patterns serve as the basis for a number of features used to train a linear regression model to predict students' performance in the course. This model achieves an R^2 of 0.34. More importantly, the underlying parameters of this model provide valuable insights as to which of the patterns most correlate with performance. Using these most-predictive patterns, we are able to identify high-level, cognitive behaviors exhibited by the students.

2. RELATED WORK

Data-driven educational research has traditionally been limited by the time-consuming process of monitoring students' learning. For example, substantial research has been performed which investigated the correlation between performance and the amount of time and effort spent on homework assignments [2, 5, 8, 16, 18]. Manually watching each student solve each homework assignment would require an intractable amount of time and, additionally, may skew the results of the study. Instead, each of these researchers relied on students or their parents to self-report the amount of time spent on each homework assignment.

Cooper et al. [6] compared the results of each of these studies and found an average correlation of $r = 0.14$ with a range from -0.25 to 0.65. Cooper et al. summarize this inconsistency in findings when they state that, "to date, the role of research in forming homework policies and practices has been minimal. This is because the influences on homework are complex, and no simple, general finding applicable to all students is possible." This underlies the impact that Educational Data Mining can have on the educational research community. By instrumenting students' natural problem-solving processes, we are able to capture a precise measurement of the actions students perform when solving their homework assignments.

More recently, researchers have applied data mining techniques to ITS and CMS data. For example, Romero et al. [17] applied data mining techniques to data collected with the Moodle CMS. This system allows students to both view and submit various assignments, e.g., homework and exams, and records detailed logs of students' interactions. These interaction logs were mined for rare association rules, that is, patterns which appear infrequently in the data. The resulting rules were then manually inspected to identify fringe behaviors exhibited by students.

Similarly Mostow et al. [14] applied data mining techniques to interaction logs taken from Project LISTEN's Reading Tutor, an ITS. This system tutors young students as they learn to read by listening to them read stories aloud and providing feedback. The authors developed a system which automatically identified meaningful features from these logs which were then used to train classifiers to predict students' future behavior with the system.

Sequential pattern mining [1] is a technique used to identify significant patterns in sequences of discrete items, e.g., consumer transaction records [1] or DNA transcripts [4]. These techniques have typically been used to mine patterns from a single database of sequences. In Educational Data Mining, it is often the case that researchers seek to find patterns that best distinguish students who do and do not perform well in the course. Thus there is a need for novel pattern mining techniques aimed at differentiating between two databases of sequences.

More recently, Ye and Keogh [20] developed a novel technique which identifies patterns which best separate two time-series databases. This technique identifies frequently occurring patterns within each database, as traditional pattern mining techniques have, but furthermore, evaluates each pattern by using it to separate sequences from the two databases. If a sequence contains the pattern, that sequence is identified as being part of the same database that the pattern came from. The pattern which provides the greatest information gain is kept as the "shapelet" that best separates the two databases.

Similarly, Kinnebrew and Biswas [13] have developed a novel differential pattern mining technique used to identify patterns that differentiate between the interactions of different groups of students with the Betty's Brain ITS. This technique begins by using SPAM [3] to identify patterns that occur in a significant number of sequences in either database. A t-test for each pattern is then performed to determine if there is a significant difference in the frequency of that pattern in each sequence of each of the two databases. This algorithm can identify patterns that occur significantly frequently in one database and not the other.

The work of Oviatt et al. [15] suggests that natural work environments are critical to student performance. Their examination of computer interfaces for completing geometry problems suggests that, "as the interfaces departed more from familiar work practice..., students would experience greater cognitive load such that performance would deteriorate in speed, attentional focus, meta-cognitive control, correctness of problem solutions, and memory." Thus, our goal is to apply Educational Data Mining techniques to data collected in natural work environments.

To that end, recent research has focused on mining ordinary, handwritten coursework data. For example, Van Arsdale and Stahovich [19] demonstrated that a correlation exists between the temporal and spatial organization of students' handwritten problem solutions and the correctness of the work. The organization of exam solutions was characterized by a set of quantitative features, which were then used to predict performance on those problems. On average these features accounted for 40.0% of the variance in students' performance on exam problems.

Similarly, Herold and Stahovich [12] presented a study in which data mining techniques were applied to students' handwritten coursework to identify how self-explanation affected students' solution processes. In this study, students from a Mechanical Engineering course were split into two groups, one which provided handwritten self-explanation along with their homework assignments and one which did not. Digital copies of the students' handwritten homework were mined for commonly occurring n-grams, revealing that students who generated self-explanation solved problems more like an expert than did those who did not generate self-explanation.

In this work, we build upon these prior efforts by applying differential pattern mining techniques to students' ordinary, handwritten problem-solving processes. In so doing, we aim to identify successful and unsuccessful solution habits and infer the higher-level cognitive processes they indicate.

3. DATA COLLECTION

In the winter quarter of 2012, students enrolled in an undergraduate Mechanical Engineering Statics course were given Livescribe$^{\text{TM}}$ digital pens. Students completed all their coursework with these pens, creating a digital record of their handwritten homework, quiz, and exam solutions. A typical exam problem is shown in Figure 1. Each problem includes a figure describing a system subject to external forces. The

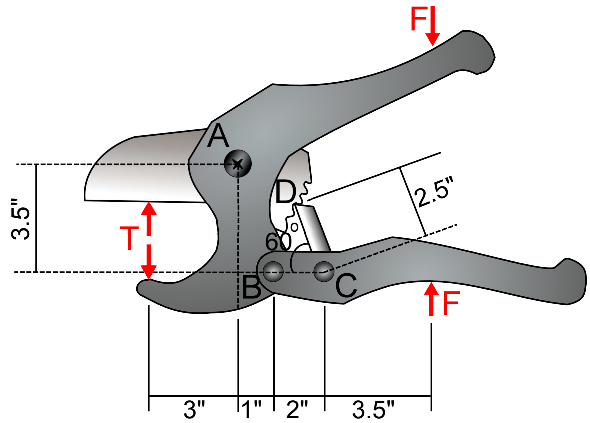

Figure 1: A typical Statics problem. The problem statement reads, "The device shown is used for cutting PVC pipe. If a force, $F = 15lb$, is applied to each handle as shown, determine the cutting force T. Also, determine the magnitude and the direction of the force that the pivot at A applies to the blade."

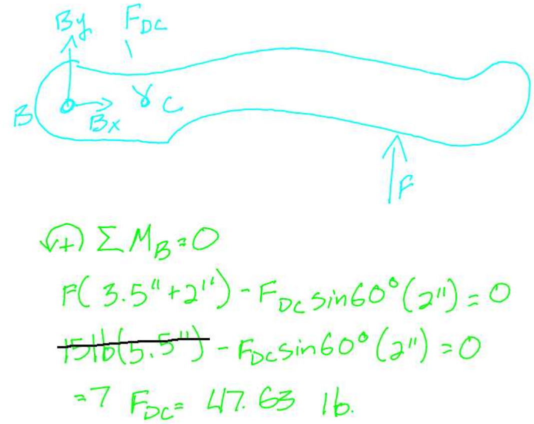

Figure 2: A hypothetical solution to a Statics problem. The color of each pen-stroke identifies the component to which it refers: cyan = FBD, green = equation, and cross-out = black.

student must apply the principles of equilibrium to compute the resulting reaction forces.

Students typically begin solving problems by drawing a free body diagram (FBD) representing the boundary of the system and the forces acting on it. The free body diagram is then used as a guide for constructing force and moment equilibrium equations. Most pen strokes in a solution correspond to either a free body diagram, an equation, or a cross-out. (Because the Livescribe pens use ink, students cannot erase errors and must instead cross them out.) Figure 2 shows a hypothetical solution to a Statics problem.

In this study, we focus on the data for homework assignments three, four, five, six, and eight, as these are the ones focused on equilibrium analysis. Assignments one and two, by contrast, focused on basic vector math, while assignment seven was a review of centroids.

The resulting data set comprises 556 sketches from 132 students. Each sketch corresponds to a single page of work from a student. Each sketch, $K = \{s_1, ..., s_m\}$, comprises a series of pen strokes. Each pen stroke, $s_i = \{p_1, ..., p_n\}$, comprises a series of points. Each point $p_j = \{x, y, t\}$ is a triple where x and y are two-dimensional Cartesian coordinates, and t is the time-stamp of that point. All points within a pen stroke, and all pen strokes within a sketch, are ordered by increasing time-stamp. The time-stamp of the first point in a pen stroke signifies the start time of that pen stroke and the last point is used to signify its end time. A sequence of labels also exists for each sketch, $L = \{l_1, ..., l_m\} | l \in \{FBD, EQN, CRO\}$. Each label, l_i, identifies stroke, s_i, by its semantic content: free body diagram (FBD), equation (EQN), or cross-out (CRO). We manually labeled the pen strokes of each sketch, but it has been shown in recent work that this process may be automated reliably [11].

While these labels account for virtually all the ink written by students, more fine-grained labeling schemes could have been developed by subdividing each label. For example, instead of labeling a pen stroke as being part of an equation, it could be labeled according to the type of equation to which it corresponds, namely, sum of forces in the X or Y direction or the sum of moments. We chose the labeling scheme presented for two major reasons.

First, this labeling scheme is sufficient for investigating hitherto unverifiable intuitions about the ways students solve Statics problems, such as the intuition that students who possess a strong understanding of the material will complete their FBD entirely before beginning their equation work. Similarly, we may corroborate the intuition that students who possess a strong understanding of the material will complete their problems in problem-number order, that is, they complete problem one entirely before completing problem two and so on.

Second, by subdividing each of the labels, we risk increasing the granularity of the resulting action sequences too far, increasing the number of total discrete actions to a point that prevents patterns from being identified.

4. ACTION SEQUENCES

In this section, we describe how each sketch may be transformed into an *action sequence*, comprising discrete actions, that is suitable for differential pattern mining. Each action is an element of a predefined alphabet of canonical actions. Each element in the alphabet represents an uninterrupted period of problem-solving performed by a student as he or she solves a homework assignment. We seek to characterize the duration, semantic content, and homework problem number for each action.

We begin by segmenting the pen strokes of each sketch by semantic type. To do so, we simply identify each index, i, in L such that $l_i \neq l_{i+1}$, and segment the series of pen strokes at each identified index. Each resulting segment contains

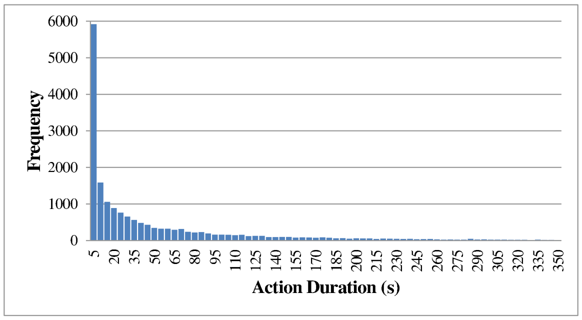

Figure 3: A histogram of the durations of FBD actions across all homework assignments. For example, the first(leftmost) bar indicates that approximately 6,000 FBD actions were between zero and five seconds long.

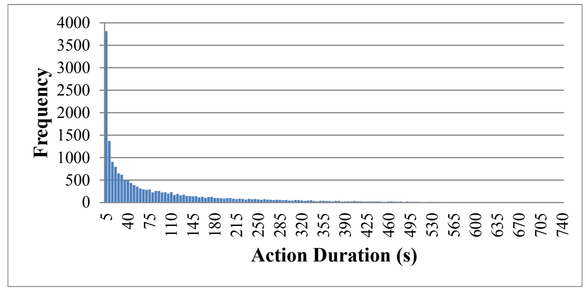

Figure 4: A histogram of the durations of EQN actions across all homework assignments. For example, the first(leftmost) bar indicates that approximately 3,500 EQN actions were between zero and five seconds long.

a sequence of actions corresponding to the same semantic type. The resulting segments do not yet satisfy the above definition of an action, as they do not necessarily contain uninterrupted work. Thus, we further segment the sketch at each index, j, such that the difference between the start time of s_j and the end time of s_{j-1} is greater than a specified threshold. In this paper, we use a threshold of five minutes, which was determined *a priori*; five minutes is a sufficiently large gap to be considered an intended break in the problem-solving process.

Each segment is then labeled with an element from the alphabet of canonical actions. If the segment comprises cross-out pen strokes, then it is given the cross-out label, C, regardless of its length or problem number. The remaining groups are labeled with a triple, $\{P, T, D\}$, where P represents the problem number, T represents the semantic type, and D, represents the duration of the action. $P \in \{1, ..., 8\}$ as there are never more than eight problems on a given homework assignment. $T \in \{F, E\}$ where F represents a FBD action and E represents an equation action. Lastly, $D \in \{S, M, L\}$, where S, M, and L indicate an action of small, medium, or large duration respectively. Take for example, the label <1-E-S>. This indicates a small action on the equations from problem one of an assignment.

The cut-off points for each duration category were determined by studying the distribution of lengths of all the FBD and equation actions. Figures 3 and 4 show a histogram for the duration of FBD and equation actions respectively. We partition each distribution into three segments such that the area under the curve for each segment is equal. The resulting thresholds are 11.26 and 80.1 seconds for FBDs and 29.59 and 147.82 seconds for equations. There are 49 unique labels in the canonical action alphabet, comprising the 48 possible combinations for a given triple and the additional cross-out label.

We seek to assign the action sequences of a student to a *performance group* based on that student's performance. In particular, we group a student's action sequence for an assignment by that student's performance on the most relevant exam, which we defined as the one that occurred most recently after that assignment was due.

Students completed homework assignments three and four prior to the first midterm exam. Students completed homework assignments five and six after the first midterm exam and before the second. Students completed homework assignment eight after the second midterm exam and before the final exam. Each midterm exam only comprised problems similar to those encountered on the homework assignments leading up to it. Thus the first midterm exam required that students solve problems similar to those found on homework assignment three and four and the second midterm exam required students to solve problems similar to those found on homework assignments five and six. The final exam comprised problems similar to all those encountered on all homework assignments.

Using this schedule of exams and homework assignments, we assign each action sequence to a group based on performance. An action sequence is assigned to the top-performing group if the student who performed those actions scored in the top third on the relevant exam. Similarly, an action sequence is assigned to the bottom-performing group if the student scored in the the bottom third of the class. The differential mining technique employed in this paper requires exactly two databases as input, thus the remaining middle-performing students are excluded from our analysis to help accentuate the differences in problem-solving behaviors of top- and bottom-performing students.

Descriptive statistics of the lengths of the action sequences for the two performance group for each assignment are shown in Table 1. It is interesting to note that the average action sequences of the bottom-performing group are always longer than those of the top-performing group, and in two cases this difference is significant ($p < 0.01$).

5. DIFFERENTIAL MINING

To identify patterns that distinguish good performance from poor performance we employ the differential pattern mining technique developed by Kinnebrew and Biswas [13]. This algorithm identifies patterns that are differentially frequent with respect to two databases of sequences, called the left and right databases.

This algorithm uses two metrics to measure the frequency of a pattern, s-frequency and i-frequency. s-frequency is de-

Table 1: Average, median, and standard deviation of the sequences for each grouping of sequences on each assignment. The fourth column contains the p-value of a t-test comparing the bottom-performing and top-performing groups on each assignment.

Group	Average	Median	Std. Dev.	t-test
HW3 Bot.	89.52	88	43.30	0.21
HW3 Top	75.38	54.5	56.12	–
HW4 Bot.	130.25	128	63.97	0.00
HW4 Top	83.28	78	46.67	–
HW5 Bot.	127.88	119.5	70.89	0.01
HW5 Top	87.14	72	54.42	–
HW6 Bot.	144.73	140	52.55	0.171
HW6 Top	126.52	122	66.05	–
HW8 Bot.	82.45	72	73.94	0.17
HW8 Top	62.28	53.5	39.64	–

fined as the number of sequences in a database that contains a specific pattern. i-frequency is defined as the number of times a pattern appears within a single sequence. Take for example, a database of ten sequences in which the first seven sequences contain one instance of a particular pattern and the last three sequences contain two instances of that same pattern. This pattern would then have a s-frequency of 10. This pattern would have an i-frequency of one in the first pattern and an i-frequency of two in the last pattern.

This algorithm begins by finding all patterns that meet a specified s-frequency threshold in the left and right database separately. Each such pattern is called an s-frequent pattern. A modified implementation of the SPAM algorithm [3] is used to identify the initial set of s-frequent patterns constrained by the a maximum gap between subsequent elements within a pattern. We use a maximum gap constraint of two in our study.

The i-frequency of each s-frequent pattern is computed for each sequence in each database. A separate t-test is computed for each s-frequent pattern to determine if the i-frequency values computed using the left database are significantly different from those computed for the right database. If the resulting p-value of the t-test is below a certain threshold, called the p-value threshold, it is considered to be *differentially frequent*. This algorithm identifies four types of differentially frequent patterns: those that are s-frequent in both sets but whose average i-frequency is higher in the left database; those that are s-frequent in both sets but whose average i-frequency is higher in the right database; those that are are only s-frequent in the left database; and those that are only s-frequent in the right database. In this study, we consider only the sequences from the last two cases as they are the most most useful for distinguishing between good- and poor-performing students.

In our implementation, we use the set of sequences from the bottom-performing group as the left database and those from the top-performing group as the right database. We use a s-frequency threshold of 0.6, meaning that a pattern must appear in at least 60% of the sequences in a database in order to be considered s-frequent. We use a p-value threshold of 0.1.

6. PERFORMANCE PREDICTION

The differential pattern mining technique identified 98 patterns in total: 6 that were s-frequent in the top-performing group but not in the bottom-performing group, and 92 that were s-frequent in the bottom-performing group but not in the top-performing group.

Our goal is to use these 98 patterns to construct a model to distinguish between good- and poor-performing students. We represent each student with 98 binary features. Each feature indicates whether a particular differential pattern from a particular assignment is contained within a student's action sequence for that assignment. To avoid computing a model that over-fits the data, we used the Correlation-based Feature Selection (CFS) algorithm with 10-fold cross-validation to identify the subset of the 98 features with the most predictive power. Those features that were selected in more than six of the ten folds by the CFS algorithm were included in the final feature subset. Table 2 shows the 20 features that were ultimately selected in this way.

We then used these 20 features to construct a linear regression model which predicts students' overall performance in the course. While more robust, non-linear classifiers could have been used, e.g., AdaBoost [7] or Support Vector Machines [9], we use a linear regression model because of the ease of interpretation; the coefficients that comprise the model give insight into the predictive power of the features used to train it. We used the linear regression package available in the WEKA machine learning software suite [10] to train the model. Our predictive model achieves an R^2 of 0.343 and includes seven features with non-zero coefficients. Table 3 lists these seven features.

7. DISCUSSION

We manually inspected each of the 98 patterns identified by the differential pattern mining algorithm and categorized the different types of cognitive processes they demonstrate. We identified seven distinct categories. *Difficulty* is the category in which students seem to encounter difficulties with a particular problem, evidenced by either repeated cross-outs or repeated attempts at the same component of the same problem. For example, the pattern <C, 1-E-S, C> describes a scenario in which the student crossed out work, worked on equations for problem one for a short time, and then again crossed out work.

Three categories describe patterns in which actions are repeated: *Repeated Equation*, *Repeated FBD*, and *Repeated Cross-out*. For instance, <2-E-S 2-E-S> is an example of a *Repeated Equation* action. Such sequences may be an indication that a student is taking a break in the middle of a particular activity to think more carefully before continuing with that activity.

Two categories describe patterns suggesting that a student may be revising either a FBD (*FBD Revision*) or an equation (*Equation revision*). These patterns comprise a cross-out followed by either the FBD or equation they are most likely revising. Also, when a student moves from working on an equation back to a FBD, this is likely an indication that the FBD is being revised; students typically attempt to complete their FBD before moving on to equations.

Table 2: Features selected using the CFS algorithm. Each feature corresponds to a pattern identified by the differential pattern mining algorithm. Each line shows the homework number and group (top or bottom) from which the pattern was identified. The final column shows the pattern that was used to compute the feature.

HW No.	Perf. Group	Sequence
3	Top	1-E-M 1-F-S
3	Top	1-F-M 1-E-M
3	Bot	2-F-M 2-E-S
3	Bot	C 5-E-S
3	Bot	5-E-M 5-F-S
3	Bot	5-E-S 5-F-M
3	Bot	C 4-E-L
4	Bot	C 1-E-M
4	Bot	1-E-L C
4	Bot	C 5-E-L
4	Bot	1-E-M 1-E-S
4	Bot	C C
5	Bot	4-E-S 4-E-S
6	Bot	1-F-M 1-E-S 1-E-M
6	Bot	1-E-S 1-E-M 1-F-M
6	Bot	1-F-M 1-F-S
6	Bot	1-F-S 1-F-M
8	Bot	5-F-M 5-F-S
8	Bot	5-F-S 5-E-M
8	Bot	5-F-M 5-E-S

Table 3: Non-zero feature coefficients for the linear regression model trained to predict student performance.

Sequence	HW	Weight	Category
1-F-S 1-F-M	3	48.8	Repeated FBD
C 5-E-L	3	51.0	EQN Revision
C 5-E-S	4	51.2	EQN Revision
1-E-M 1-F-S	4	55.5	FBD Revision
C 1-E-M	6	62.7	EQN Revision
1-F-M 1-E-S 1-E-M	6	63.1	Difficulty
5-F-M 5-F-S	8	73.1	Repeated FBD

Lastly, is the *Normal* category. This is the category for all patterns in which a FBD is followed by an equation of the same problem number. A differential pattern belonging to the *Normal* category is particularly informative when one group exhibits significantly more normal sequences – it is an indication that the other group is solving their homework assignment out-of-order more often.

The non-zero weighted features of the linear regression model (Table 3) help identify the patterns which are most predictive of students' grades, and thus provide insight into the behaviors which best correlate with students' performance. In Table 3, Patterns 1, 4, 5, and 6 are all similar in that they comprise actions pertaining to the first problem on a homework assignment, and suggest that a student may be having difficulty or is frequently revising his or her work. This is an indication that when students encounter difficulty on the first problem, which is typically the easiest problem of the homework assignment, that they may continue to encounter those difficulties throughout the quarter.

Patterns 2, 3, and 7 in Table 3 are all similar in that they pertain to problems that are very similar to problems that appear on either a later midterm, the final exam, or both. (These problems differ only superficially from exam problems. For example, the geometry may be rotated.) These patterns all describe situations in which the student is revising his or her equations or FBDs. The features suggest that students who frequently revise problems which are similar to an exam problem are likely to have difficulty with those problems later on during an exam.

It would be difficult to use the linear regression model to predict performance for students of a future section in Statics. To do so would require that the instruction, assignments, and exams, be identical. This is not likely to be the case, as some of the homework problems are modified each year to prevent copying solutions from the previous offering.

Instead, the patterns and correlations discovered in this paper may be used to guide future offerings of this course. For example, if a student's work contains patterns which indicate difficulty, similar to those found in this study, on the first problem of an assignment or on a problem which is similar to one that will appear in a future exam, the instructor can provide targeted materials for that student to address that difficulty. Furthermore, the results here indicate which problems have a strong bearing on students' performance. For example, students who seemed to have difficulty constructing a FBD on problem five of homework eight typically did not perform well in the course. This indicates to the instructors of future offerings this course, that more time should be spent in class reviewing how the FBD for this problem should be constructed.

8. CONCLUSION

We have presented an application of data mining techniques to educational data extracted from a novel environment. We have given undergraduate Mechanical Engineering students Livescribe™ digital pens with which they completed all their coursework. These pens record students' handwriting as time-stamped pen strokes enabling us to not only analyze the final image, but also the sequence in which it was written.

We developed a novel representation of students' handwritten work on an assignment which characterizes the sequence of actions the student took to solve that problem. This representation comprises an alphabet of 49 canonical actions that a student may make when solving his or her homework assignment. Each action is characterized by its duration, problem number, and semantic content. This representation allows us for the first time, to apply traditional data mining techniques to sequences of students' handwritten problem solutions.

We assigned these sequences into top- and bottom-performing groups according to performance on each sequence's most relevant exam. The most relevant exam for a sequence from a particular homework assignment is the exam which occurs most recently after that homework assignment was due. Sequences from students who performed in the top third of the class on that assignment comprise the top-performing group and sequences from students who performed in the bottom third comprise the bottom-performing group. We applied a differential data mining technique to the sequences from the students in each of these groups and identified patterns that are more frequently exhibited by one group than the other.

These patterns serve as the basis for features used to train a linear regression model to predict students' performance in the course. This model achieves an R^2 of 0.34. Furthermore, the underlying parameters of this model provide valuable insights as to which of the patterns best correlate with performance. From these best-correlating patterns, we have manually identified high-level cognitive behaviors exhibited by the students. These behaviors provide insight as to when students may be experiencing difficulty in the course. These techniques may be applied in future sections of this course to identify when students are having difficulty in class, enabling the instructor to rapidly address those difficulties.

9. REFERENCES

[1] R. Agrawal and R. Srikant. Mining sequential patterns. In *Data Engineering, 1995. Proceedings of the Eleventh International Conference on*, pages 3–14. IEEE, 1995.

[2] J. L. Antonek. Interactive homework: Linking parents to the foreign language learning of their children. In *Dissertation Abstracts International*, volume 57, 1996.

[3] J. Ayres, J. Flannick, J. Gehrke, T. Yiu, et al. Sequential pattern mining using a bitmap representation. In *Conference on Knowledge Discovery in Data: Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, volume 23, pages 429–435, 2002.

[4] S. Bai and S. Bai. The maximal frequent pattern mining of dna sequence. In *Granular Computing, 2009, GRC '09. IEEE International Conference on*, pages 23 –26, 2009.

[5] N. K. Bowen and G. L. Bowen. The mediating role of educational meaning in the relationship between home academic culture and academic performance. In *Family Relations*, volume 47, pages 45–51, 1998.

[6] H. Cooper, J. Robinson, and P. Erika. Does homework improve academic achievement? *Review of Educational Research*, 76, 2006.

[7] C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[8] R. Deslandes, E. Royer, P. Potvin, and D. Leclerc. Patterns of home and school partnership for general and special education students at the secondary level. In *Exceptional Children*, volume 65, pages 496–506, 1999.

[9] Y. Freund and R. Schapire. A desicion-theoretic generalization of on-line learning and an application to boosting. In *Computational learning theory*, pages 23–37. Springer, 1995.

[10] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. In *SIGKDD Explorations*, 2009.

[11] L. Han-lung, T. Stahovich, and J. Herold. Automatic handwritten statics solution classification and its applications in predicting student performance. In *Proceedings of the American Society for Engineering Education Annual Conference*, 2012.

[12] J. Herold and T. Stahovich. Characterizing students' handwritten self-explanations. In *Proceedings of the American Society for Engineering Education Annual Conference*, 2012.

[13] J. Kinnebrew and G. Biswas. Identifying learning behaviors by contextualizing differential sequence mining with action features and performance evolution. In *Proceedings of the Fifth International Conference on Educational Data Mining*, 2012.

[14] J. Mostow, J. Gonzalez-Brenes, and B. H. Tan. Learning classifiers from a relational database of tutor logs. In *Proceedings of the Fourth International Conference on Educational Data Mining*, 2011.

[15] S. Oviatt, A. Arthur, and J. Cohen. Quiet interfaces that help students think. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 191–200, 2006.

[16] K. Pezdek, T. Berry, and P. A. Renno. Children's mathematics achievement: The role of parents' perceptions and their involvement in homework. In *Journal of Educational Psychology*, volume 94, pages 771–777, 2002.

[17] C. Romero, J. Romero, J. Luna, and V. S. Mining rare association rules from e-learning data. In *Proceedings of the Third International Conference on Educational Data Mining*, 2010.

[18] K. Singh, M. Granville, and S. Dika. Mathematics and science achievement: Effects of motivation, interest, and academeic engagement. In *Journal of Educational Research*, volume 95, pages 323–332, 2002.

[19] T. Van Arsdale and T. Stahovich. Does neatness count? what the organization of student work says about understanding. In *Proceedings of the American Society for Engineering Education Annual Conference*, 2012.

[20] L. Ye and E. Keogh. Time series shapelets: a new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 947–956. ACM, 2009.