

Core Course Analysis for Undergraduate Students in Mathematics

Ritvik Kharkar
University of California, Los Angeles
ritvikmath@ucla.edu

Jessica Tran
University of California, Los Angeles
jessicatran@ucla.edu

Charles Z. Marshak
University of California, Los Angeles
cmarshak@math.ucla.edu

April 9, 2016

In this work, we develop statistical tools to understand core courses at the university level. Traditionally, professors and administrators label courses as "core" when the courses contain foundational material. Such courses are often required to complete a major, and, in some cases, allocated additional educational resources. We identify two key attributes which we expect core courses to have. Namely, we expect core courses to be highly correlated with and highly impactful on a student's overall mathematics GPA. We use two statistical procedures to measure the strength of these attributes across courses. The first of these procedures fashions a metric out of standard correlation measures. The second utilizes sparse regression. We apply these methods on student data coming from the University of California, Los Angeles (UCLA) department of mathematics to compare core and non-core coursework.

1. INTRODUCTION

Educators, professors, and mathematicians are continually improving mathematics education [8, 28]. There is a consensus that mathematics education divides into two categories: design of math curriculum and the development of teaching strategies. In particular, members from the Mathematical Association of America and the American Mathematical Society [18] are regularly analyzing the core requirements of undergraduate mathematics and constantly improving the effectiveness of mathematics teachers [7, 24, 29]. In this paper we focus on the design of the curriculum using tools to understand core courses using grade statistics.

The undergraduate math curriculum is the set of courses required for an undergraduate to earn his or her degree. Departmental faculty and counselors carefully select courses which provide a well-rounded mathematics education. Traditionally, core courses within the mathematics curriculum are those that contain the most foundational material and that inform subsequent coursework. These courses are often ones that the department encourages students to take early to ensure they get the most from their major.

We will analyze the core curriculum using techniques and tools in Educational Data Mining (EDM). EDM is an interdisciplinary research area that aims to improve education utilizing data analysis [15, 19, 20, 31]. Although researchers have been collecting and analyzing educational

data for a long time, it has only recently been established as its own field, due to the accessibility of statistical tools and the emergence of EDM journals and conferences [1,21]. Tools developed in EDM research have been more quickly adopted due to the challenge of serving a growing undergraduate population [10]. Recent advancements in education made possible through data analysis tools include course recommendation systems [4,9,30], student clustering [16], grade predictions [3,17], and personalized education through online medium [23,27]. Adapting such data mining tools for assessing student performance is a difficult task but will help administrators even better serve their student body.

The focus of this paper will be to provide statistical criteria that validates the department's core course selection. Specifically, we compared the grade distribution from a single math course to the distribution of GPAs. We applied this methodology to the grades of UCLA mathematics students from 2000-2015. The viewpoint we develop here can be easily adapted to any other department for a similar analysis.

To compare course statistics and GPA statistics we employ two fundamental statistical tools: the correlation coefficient and linear regression. The correlation coefficient is a way to measure the covariance of two distributions [14]. Regression is a tool that linearly models how particular measured inputs influence a response [14]. It informs some of our most important models for business, biology, machine learning, and environmental science [5,13,26]. Ultimately, we will use these tools to compare the grades earned by students in a single class with these students' overall GPAs.

The data-centric viewpoint we develop offers a statistical methodology to validate the core courses within a university's curriculum. Using this methodology, we can identify those courses that especially impact a student's study. Any discrepancy between the core courses defined by the department and those found from this statistical viewpoint can prompt further discussion of an effective way to design the curriculum. Ultimately, by analyzing the university curriculum, we hope to aid departments in better meeting their goals and preparing exiting students better for industry and academic jobs.

We now discuss how this paper is organized. In section 2. of our paper, we describe our data-centric approach to core course analysis providing justification for the use of correlation coefficients and regression. In section 2., we precisely describe the statistical setup and provide some sample output. In section 3., we apply this methodology of determining core courses to pure mathematics majors.

2. CHARACTERIZING CORE COURSES

Core courses are those that are the most foundational for a well-rounded mathematics education and are often allocated additional educational resources. In this study, we consider the core courses selected by the UCLA mathematics department and identify statistical attributes that we expect such courses to have. We develop statistical tools to measure the strength of these attributes across courses. Using these tools, we then compare the core and non-core courses to validate the department's core course selection. At UCLA, the core courses are Linear Algebra and Real Analysis. All mathematics students must take these courses and are recommended to take them early in their course sequence.

The first attribute we consider is the correlation between the course grade of a particular student and his or her overall mathematics GPA. By investigating this correlation, we analyze the relationship between rankings determined by grade compared to rankings determined by overall

mathematics GPA. We expect that the course grades of core courses defined by the department would be highly correlated with one's overall mathematics GPA.

The second attribute we consider is the impact of a particular course grade on overall mathematics GPA. The impact of a course grade on a student's overall mathematics GPA signifies how sensitive overall mathematics GPA is to the course grade. We examine impact by performing lasso regression in order to select the most important courses. When we input courses from the mathematics curriculum as variables into lasso, the algorithm selects a subset of courses which are most relevant to overall mathematics GPA. We conclude that the courses assigned the highest coefficients by lasso regression are those having the greatest bearing on overall mathematics GPA. We now describe the correlation metric used to measure the first attribute.

2.1. MIXED CORRELATION METRIC

To measure the correlation between course grades and a student's overall mathematics GPA, we combine two correlation coefficients. We then analyze these correlation coefficients to measure how well the ranking determined by a single course grade compares to the ranking of overall mathematics GPA. We first consider Spearman's ρ [22] which we will denote by ρ_S . Given a sample of n students, we can rank students both by overall mathematics GPA as well as by grade in a particular course. Let $r_1^{\text{GPA}}, \dots, r_n^{\text{GPA}}$ be the ranking of students by overall mathematics GPA and let r_1^C, \dots, r_n^C be the ranking of students by grade in course C . The measure ρ_S is defined as

$$\rho_S = 1 - \frac{1}{Z} \sum_{i=1}^n (r_i^{\text{GPA}} - r_i^C)^2$$

where $Z = \frac{n(n^2-1)}{6}$ normalizes the sum. The coefficient ρ_S ranges from -1 to 1 in which positive values indicate an agreement between the two rankings. In the context of our study, we use ρ_S to measure how much a rank determined by a student's grade in a course agrees with his or her overall mathematics GPA.

We will also employ the coefficient of determination, R_{sq} , a normalized measure of covariance between grade in a particular course and overall mathematics GPA. Let g_1^C, \dots, g_n^C be the grade in course C of each student and let $g_1^{\text{GPA}}, \dots, g_n^{\text{GPA}}$ be the overall mathematics GPA of each student. The sample means of these quantities are defined as

$$\begin{aligned} \bar{g}^C &= \frac{1}{n} \sum_{i=1}^n g_i^C \\ \bar{g}^{\text{GPA}} &= \frac{1}{n} \sum_{i=1}^n g_i^{\text{GPA}}. \end{aligned}$$

R_{sq} is defined as follows,

$$R_{sq} = \frac{\sum_{i=1}^n (g_i^C - \bar{g}^C)(g_i^{\text{GPA}} - \bar{g}^{\text{GPA}})}{\sqrt{\sum_{i=1}^n (g_i^C - \bar{g}^C)^2} \sqrt{\sum_{i=1}^n (g_i^{\text{GPA}} - \bar{g}^{\text{GPA}})^2}}.$$

Using these two correlation coefficients, we define a new holistic metric as $\rho_H = \sqrt{(\rho_S)^2 + (R_{sq})^2}$. This new metric, ρ_H takes into account both correlation metrics and is the ℓ_2 distance to the origin.

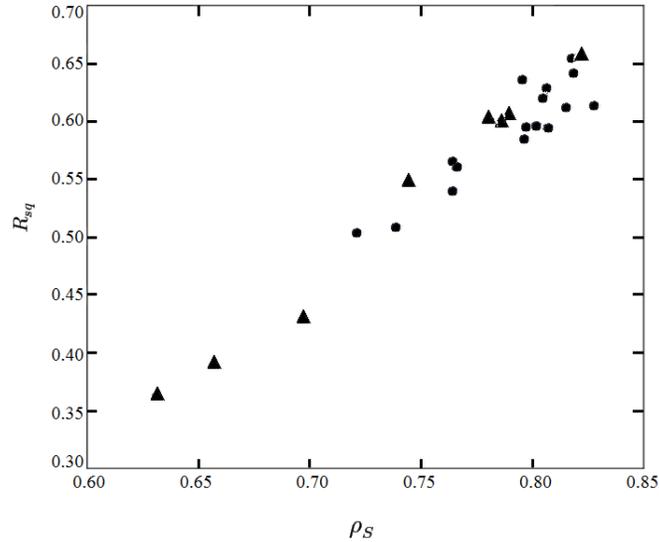


Figure 1: Scatter plot of ρ_S values against R_{sq} values for the Math/Econ, Pure Math, Applied Math courses. The triangular points are courses which are required for all majors and the round points are courses which are not required for all majors.

We give an example of how these metrics are used to identify core courses. We will only consider students who have taken at least five mathematics courses and at least one upper division mathematics course at UCLA so we can identify those courses that might impact students who have made a significant time commitment to their mathematical development. We provide an example that includes the three largest mathematics majors at UCLA: Mathematics/Economics (Math/Econ), Pure Mathematics (Pure Math), and Applied Mathematics (Applied Math) and will only consider students admitted as freshmen. We omit transfer students since these students usually have different training than students that are admitted right from high school. We report the findings in Tables 1 and 2. After generating the ρ_S and R_{sq} values for each student group in each of the courses, we plot the two metrics to understand how well the rankings matched. The scatter plot in Figure 1 plots ρ_S against R_{sq} using the values from Tables 1 and 2.

Table 1: Mixed Correlation Coefficients

Course	ρ_S	R_{sq}	ρ_H
Linear Algebra	0.821	0.655	1.05
Complex Analysis	0.819	0.657	1.049
Ord. Diff. Eq.	0.818	0.654	1.047
Num. Analysis (Part II)	0.827	0.614	1.03
Probability (Part I)	0.806	0.62	1.02
Probability (Part II)	0.814	0.607	1.015
Game Theory	0.805	0.619	1.015
Num. Analysis (Part I)	0.794	0.62	1.007
Nonlinear/Linear Diff. Eq.	0.805	0.603	1.006
Optimization	0.801	0.601	1.001
Combinatorics	0.797	0.601	0.998
Diff. Eq.	0.787	0.61	0.996

Table 2: Mixed Correlation Coefficients

Course	ρ_S	R_{sq}	ρ_H
Diff. Geometry (Part I)	0.802	0.582	0.991
Math Modeling	0.797	0.584	0.988
Real Analysis (Part I)	0.781	0.6	0.984
Multivar. Calc. (Part II)	0.785	0.589	0.981
Algebra for App.	0.765	0.552	0.943
Real Analysis (Part II)	0.764	0.546	0.939
Linear Algebra	0.744	0.548	0.924
Discrete Math	0.738	0.506	0.895
Abs. Algebra (Part I)	0.721	0.512	0.885
Multivar. Calc. (Part I)	0.696	0.423	0.814
Single Var. Calc. (Part II)	0.657	0.383	0.76
Single Var. Calc. (Part I)	0.629	0.34	0.715

We can calculate ρ_S and R_{sq} for each course. We can then calculate the ℓ_2 distance from the origin to (ρ_S, R_{sq}) to get a measure of how accurately grades in each course correlate with overall mathematics GPA values. Thus, the larger this distance, the stronger the correlation between course grade and overall mathematics GPA. Figure 2 shows some of the courses closest to the origin. Since our sample population is from a diverse set of majors, we indicate overlapping required courses with a (*). In the discussion of our results, we use course labellings as assigned by the math department of UCLA, see Appendix A.

From this bar plot, we see that the only required upper division courses that appear from this metric are Linear Algebra, Complex Analysis and Real Analysis. This validates that Linear Algebra is a foundational course and also agrees with the department's decision to recommend this course as a first upper division course. Complex Analysis, although not considered a core course by the department, follows immediately after Linear Algebra. On the other hand, Real Analysis, determined a core course by the department, ranks further down in our list. We will perform regression analysis to find which courses are most impactful.

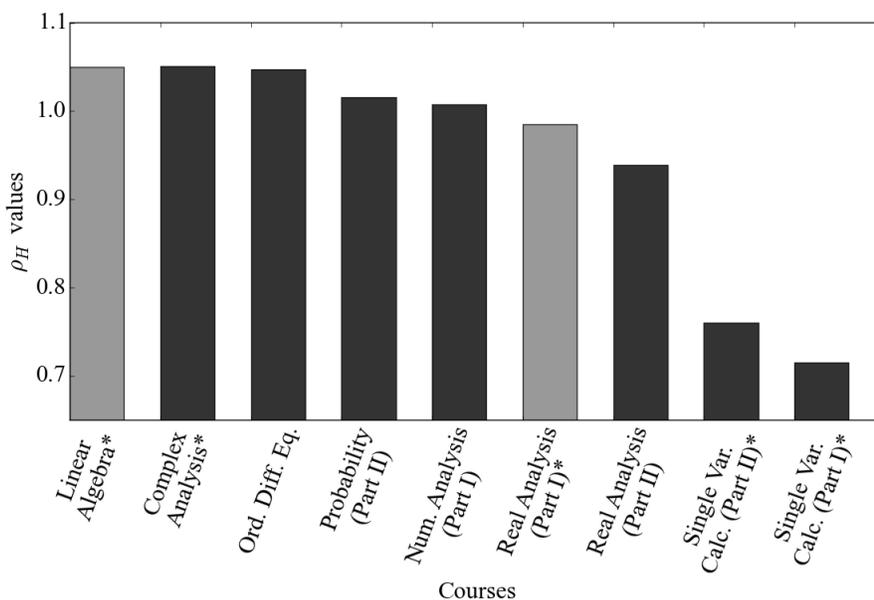


Figure 2: Bar plot of ρ_H values of a sample of courses from the correlation metric. The top courses were calculated based on the distance from the origin. Bars colored grey are core courses as defined by the mathematics department. The asterisk (*) indicates courses that are required.

2.2. LASSO REGRESSION METRIC

In order to further validate the department’s selection of core courses, we will now turn our attention to analyzing the impact of a particular course’s grade on overall mathematics GPA. We model the overall mathematics GPA as a linear response to the different course grades a student earns. We chose to use a linear model due to the versatility of such models, the fact that GPA is a linear function of all course grades, and the apparent linear relationship between core course grades and overall mathematics GPA as seen in Figure 3. We proceed by selecting the most impactful courses using sparse regression. This regression serves a dual purpose as it also assigns zero coefficient values to courses which have minimal bearing on overall mathematics GPA. This effectively allows us to disqualify certain courses from being considered core. A key application of lasso regression is grade prediction. That is, given the impactful courses selected by lasso regression, we can input the grades in these courses into a different regression model aimed towards predicting overall mathematics GPA. This regression model, better suited to calculate the exact values of the coefficients for each course, will allow us to estimate a student’s overall mathematics GPA based on his or her grade in a few courses.

2.2.1. Course Selection with Lasso

We proceed by selecting the most impactful courses with lasso regression [11]. Lasso regression uses an ℓ_1 penalty to ensure the coefficients for the linear model are sparse. A non-zero regression coefficient indicates that the corresponding course is impactful while a zero-coefficient indicates the corresponding course is not. By removing the courses that have little bearing on

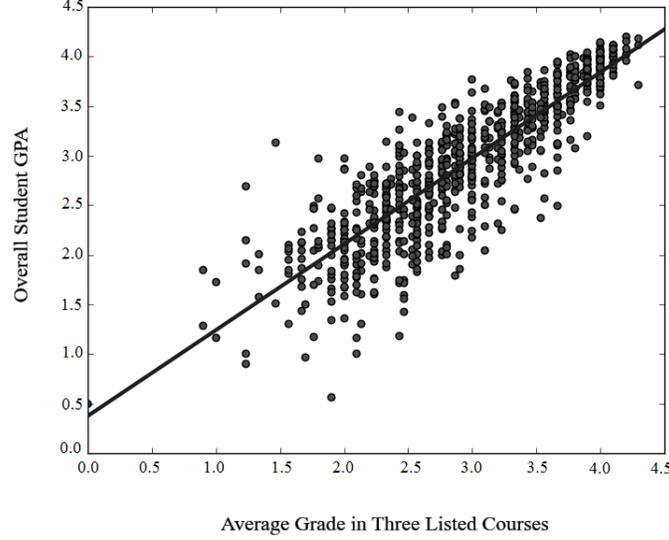


Figure 3: Scatter plot of average grade in Linear Algebra, Real Analysis, and Probability Theory (Part I) against overall student GPA for all students who took all three courses. The ridge regression prediction function is drawn through the points.

a student’s overall mathematics GPA, we are able to extract the most impactful courses and validate core course selection by the mathematics department.

Let C_1, \dots, C_m be the m available courses in the math department. We can model a student i ’s overall GPA g_i^{GPA} as

$$g_i^{\text{GPA}} = \beta_0 + \beta_1 g_i^{C_1} + \beta_2 g_i^{C_2} + \dots + \beta_m g_i^{C_m} + \epsilon_i$$

where

$$g_i^{C_k} = \begin{cases} \text{student } i\text{'s grade in class } k & \text{if student } i \text{ took class } k \\ 0 & \text{otherwise} \end{cases}$$

When considering all n students, we can write this model in the matrix form

$$\mathbf{g}^{\text{GPA}} = \mathbf{G}\beta + \epsilon$$

where the matrix $\mathbf{G} \in \mathbb{R}^{n \times m}$ contains all of our grade data for all m math courses at UCLA for all n students.

The coefficient vector β is determined as

$$\arg \min_{\beta} \|\mathbf{g}^{\text{GPA}} - \mathbf{G}\beta\|_2^2 + \alpha \|\beta\|_1 \quad (1)$$

for $\alpha > 0$. For larger α , we obtain sparser solutions. As α approaches 0, we obtain solutions that resemble those obtained by ordinary least squares. While there is no closed form solution for Eq. 1, the minimization can be solved using standard techniques in convex optimization [2]. Throughout this paper, we select $\alpha = 0.1$ when using a lasso model because, experimentally, it has given that approximately 80% of the coefficients are zero. We note that the β coefficients produced by lasso regression have a smaller variance than those produced by an ordinary least

Table 3: Lasso regression coefficient values for Pure Mathematics, Mathematics/Economics, and Applied Mathematics freshman admits.

Course	Lasso Coefficient
Complex Analysis	0.05734
Diff. Geometry (Part I)	0.0502
Num. Analysis (Part II)	0.04429
Diff. Eq. (Lower Div.)	0.04375
Probability (Part I)	0.03322
Math Modeling	0
Multivar. Calc. (Part I)	0
Math History	0

squares regression [25]. Unfortunately, this comes at the price of the lasso regression β 's being more biased estimators.

The results of the lasso regression for freshman mathematics students from the majors Pure Mathematics, Applied Mathematics, and Mathematics/Economics are shown in Table 3.

2.2.2. Grade Prediction with Ridge

Given the courses selected by lasso regression, we will now utilize ridge regression to assign accurate numerical weights to this subset of courses and thereby gauge precisely how impactful each course in our subset is for overall mathematics GPA. In addition, we will also have a meaningful prediction function for overall mathematics GPA based on grades in a few impactful courses.

We use ridge regression because it will help us account for the multicollinearity between the grades in particular courses, our explanatory variables. We expect multicollinearity to exist due to the fact that students who do well in one course tend to do well in others, and vice versa. This multicollinearity, if unchecked, can skew our regression coefficients, and produce false values [6].

It is important to note that while ridge regression helps to correct for multicollinearity issues by imposing an ℓ_2 penalty, it causes the resulting regression coefficients to be biased [12], systematically shifted towards lower values. But, upon running our tests with both ridge regression and ordinary least square regression, we find that ridge regression preserves the relative ordering of the coefficients, produces a higher accuracy level as measured by the coefficient of determination R_{sq} , and deviates from the ordinary least squares β values only slightly.

In our grade prediction process, we will consider three courses as input into our ridge model. The first two courses will always be Linear Algebra (Part I) and Real Analysis (Part I), because they are considered core courses by the mathematics department. We will cycle through each of the top three courses and last course in Table 3 as our third course in the ridge model, thereby running the ridge model four separate times. We run the model four times because the lasso model selected the top three courses, indicating that overall mathematics GPA is especially

sensitive to grades in each selected course. We run the model a fourth time with a course that lasso indicated as non-core to check whether our accuracy level diminishes as expected. We expect our accuracy level to diminish as we pick courses with progressively lower lasso coefficients.

We use a linear ridge model so that our model for student i is

$$g^{\text{GPA}} = \beta_0 + \beta_1 g_i^{115A} + \beta_2 g_i^{131A} + \beta_3 g_i^{C_3} + \epsilon_i.$$

where C_3 takes on the label of each of the top three courses in Table 3 as well as the last course in Table 3.

We can compactly express the implied system of equations for n students as

$$\mathbf{g}^{\text{GPA}} = \mathbf{G}\beta + \epsilon$$

where $\mathbf{G} \in \mathbb{R}^{n \times 4}$ is a matrix encoding each student's grade in each of the three courses we are using.

The β coefficient vector is given by

$$\arg \min_{\beta} \|\mathbf{g}^{\text{GPA}} - \mathbf{G}\beta\|_2^2 + \alpha \|\beta\|_2 \quad (2)$$

for $\alpha > 0$. We find that $\alpha = 2.5$ gives the highest accuracy in our predicted overall mathematics GPA values, \hat{g}^{GPA} . For each separate run of the ridge model, we consider all students in the population who have taken Linear Algebra, Real Analysis, and the rotating third course. We then take a random 90% sample of this group as our training set and use it to generate a coefficient vector β using Eq. 2. The remaining 10% of the data is used to test this function. We repeat this procedure 100 times, taking a different 90% random sample of the students each time, and then take the average of the 100 resulting β vectors to generate an averaged prediction model for overall mathematics GPA. We also average the accuracy levels, measured by R_{sq} , over 100 runs of a particular combination of three courses.

We find that by using Complex Analysis, the course with the highest lasso coefficient in Table 3, as our third course, we get the highest R_{sq} value. This accuracy level drops by 5% when we use Differential Geometry (Part I), having the second highest lasso coefficient. Furthermore, when we use History of Mathematics as our third course, listed as having a zero lasso coefficient, our accuracy drops by 17% relative to the model with Complex Analysis.

Our ridge model thus serves to validate the results of our lasso regression. The fact that courses assigned higher lasso coefficients also lead to higher grade prediction accuracy levels provides evidence that our lasso coefficients are accurately ranked. Furthermore, the dramatic decline in accuracy attributed to using courses with zero lasso coefficient in the ridge model verifies that these courses have little bearing on a student's overall mathematics GPA.

3. CORE COURSE ANALYSIS FOR UCLA MATH STUDENTS

Now that we have established methods to validate the core courses, we apply these methods on several different math majors at UCLA. The UCLA mathematics department houses seven different majors. While each major has a different set of course requirements, the department considers Real Analysis and Linear Algebra as the core courses for each major and these two courses are required for *all* majors. We expect grades in these two core courses to be highly

correlated with and highly impactful of a student's overall mathematics GPA, more so than non-core courses. We compare the two core and the other non-core courses within a major using the statistical tools outlined in the previous section. We highlight those departmental non-core courses that have high β_i and ρ_H values relative to departmental core courses and interpret this apparent discrepancy.

3.1. CORE COURSE ANALYSIS FOR PURE MATHEMATICS FRESHMEN

We now analyze core and non-core courses for those students majoring in pure mathematics that were admitted to UCLA as freshmen. We found that our course analysis was the most interesting for these students because many non-core courses were both highly correlated to and highly impactful of these students' overall mathematics GPA, more so than the two departmental core courses. We carefully discuss the findings for this major below. We also provide some possible explanations as to why many non-core courses exhibited higher correlation and impact on overall mathematics GPA than core courses did.

First, we study ρ_H of courses when considering pure mathematics students. We display the ρ_H -values in Figure 4. In this bar plot, we show the nine courses with the greatest ρ_H . We see that the two core courses Linear Algebra (Part I) and Real Analysis (Part I) are in this list, but are not the top ranked. Surprisingly, the top ranked courses are Complex Analysis, Topology, and Linear/Nonlinear Differential Equations. While every pure math student is required to take Complex Analysis, they are not required to take it early in their major as it is not required for later coursework. Topology and Linear/Nonlinear Differential Equations are not required at all. While core courses were highly correlated with a student's overall mathematics GPA, we were surprised that Complex Analysis scored highest. Unlike most other courses in the pure mathematics major, Complex Analysis has no upper-division prerequisite courses. Indeed, until recently, it was also required for Electrical Engineering students who generally do not take the two core mathematics courses. Moreover, no other upper division mathematics course at UCLA requires Complex Analysis. On the surface, we expect that Complex Analysis could be taken in a student's last quarter without impacting their performance in other courses. As such, we also expect that such a course would have lower correlation to overall mathematics GPA.

Next, we study the lasso coefficients β_i and display the results in Table 4. Here, the course with the highest β_i value is again Complex Analysis, followed by Probability Theory (Part I), Honors Abstract Algebra (Part I), and Differential Equations. Linear Algebra and Real Analysis, the core courses, both have zero as their coefficients, suggesting that non-core courses have greater impact on overall mathematics GPA. More formally, those courses with non-zero coefficients are the most important inputs for a linear model of overall mathematics GPA. We are surprised to find that Complex Analysis ranks as the top course under both the ρ_H and β_i metrics because rank in these metrics is mutually exclusive. That is, a course can have a high ρ_H value, being highly correlated with overall mathematics GPA, without having a high β_i value, indicating that it is highly impactful on overall mathematics GPA. Then, the fact that Complex Analysis ranks as the top course under both metrics says that it is more correlated and more impactful on overall mathematics GPA than any other mathematics course for our pure mathematics freshman admit group.

To understand why Complex Analysis was the course with highest ρ_H and β_i values, we looked for other explanatory statistics. We found that regular faculty more frequently teach Complex Analysis than temporary faculty do. We believe this is one possible factor in explaining

Table 4: Summary statistics for Pure Mathematics freshman admits. Complex Analysis is the top ranked course under both correlation and regression metrics.

Course	ρ_H Rank	ρ_H	β_i Rank	β_i
Complex Analysis	1	1.039	1	0.3774
Topology	2	1.038	–	0
Linear Algebra	5	0.951	–	0
Real Analysis (Part I)	6	0.9507	–	0
Probability (Part I)	10	0.8224	2	0.00824
Honors Abs. Algebra (Part I)	11	0.811	3	0.01919
Diff. Eqs. (Lower Div.)	14	0.8012	4	0.001573

high ρ_H and β_i values for Complex Analysis. Permanent faculty have more teaching experience and are better able to assess students in their course. As such, Complex Analysis grades more strongly correlate with the overall mathematics GPA. As to whether students are getting more adept training from regular faculty for the rest of their math major is beyond the scope of this work. Moreover, in this data set, we cannot determine whether a regular or temporary faculty assigned a particular student's grade and so do not comment on this further.

Having selected some courses from the lasso model, we can now generate a prediction function to estimate overall mathematics GPA for pure mathematics freshman admit students. We choose to run ridge regression using Linear Algebra (Part I), Complex Analysis, and Probability Theory (Part I). We choose the latter two courses because they are assigned non-zero lasso coefficients as seen in Table 4. We choose Linear Algebra (Part I) because it is a departmental core course and has a relatively high correlation rank in Table 4. From our original student population of pure mathematics students who have taken at least one upper division course, we further sample the students who have taken all three courses. The results of our ridge regression are as follows

$$\hat{g}^{\text{GPA}} = 0.74 + 0.29g^{132} + 0.27g^{115A} + 0.18g^{170A}.$$

We obtain an R_{sq} of 87%. Although we use grades in the above three courses to estimate overall mathematics GPA, in general students take far more courses. In this sample alone, the average number of courses taken is 16 courses with standard deviation of 4.11. Taking three standard deviations below the mean of 16, we still arrive at a number higher than 3 courses, indicating that it is unlikely to find students whose overall mathematics GPA's are determined only by the three courses used in the model. In Figure 5 we can see a scatter plot of the results.

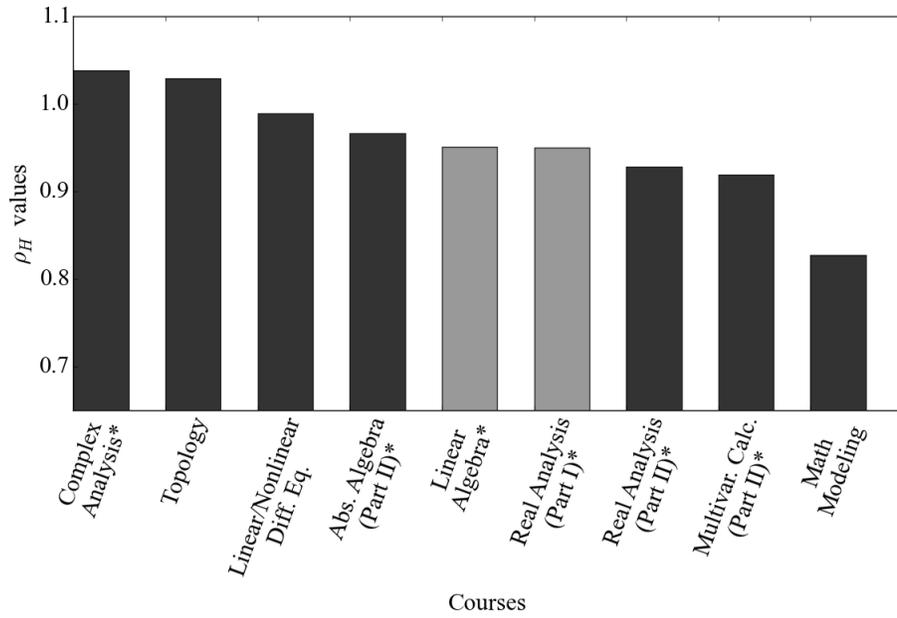


Figure 4: Bar plot of courses of pure mathematics freshman student admits. Bars colored grey are core courses as defined by the mathematics department. Courses marked with an asterisk (*) are required courses.

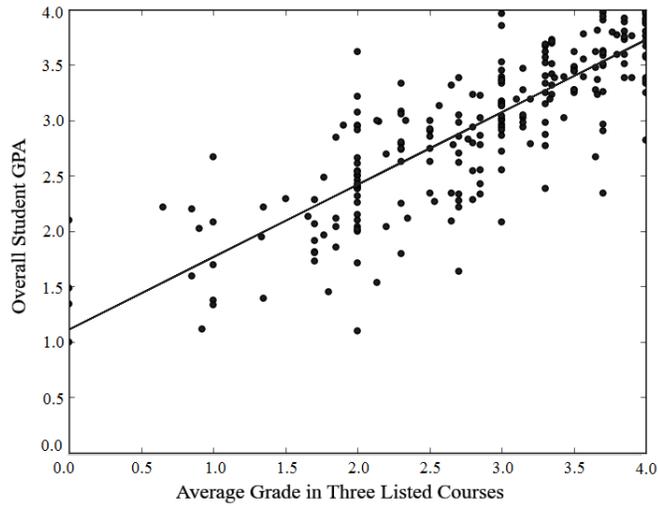


Figure 5: Scatter plot of grade values in Complex Analysis, Linear Algebra, and Probability Theory (Part I) plotted against overall GPA with ridge prediction function.

Table 5: Summary statistics for Math/Econ freshman admits.

Course	ρ_H	Rank	ρ_H	β_i	Rank	β_i
Abs. Algebra (Part I)	1	1.08271	–	0		
Single Var. Calc. (Part I)	2	1.02221	–	0		
Multivar. Calc. (Part II)	3	1.01405	3	0.68589		
Num. Analysis (Part II)	4	1.00849	–	0		
Diff. Eq. (Low. Div.)	5	0.986724	7	0.010956		
Linear Algebra (Part I)	6	0.94111	–	0		
Real Analysis (Part I)	13	0.9101	1	0.08471		
Optimization	12	0.916562	2	0.081447		

3.2. SURVEY OF OTHER MATHEMATICS MAJORS

While we have focused primarily on the Pure Mathematics freshman admit group for our analyses, we have also used the outlined methods on two other large groups of mathematics majors. Namely, we have obtained results for ρ_H for the Math/Econ freshman admit group and the Applied Mathematics freshman admit group. We display the results for these groups in Table 5 and Table 6.

As we can observe from Table 5, within the Math/Econ freshman group, there are two surprising results. Firstly, the course Abstract Algebra (Part I), which is not required for Math/Econ students, has a high ρ_H value. Additionally, Real Analysis (Part I), which is a core course, has a high β_i value. The result of the course grade in Real Analysis (Part I) having high impact is surprising because of the few number of theoretical courses students of the Math/Econ major are required to take. Additionally, the result of a course grade in Abstract Algebra having high correlation with one's overall mathematics GPA is surprising because it is both theoretical and not required within the major. Although theoretical content may prepare students for a well-rounded foundation for their future coursework, most courses within the Math/Econ curriculum tend to be applied.

As we can observe from Table 5 and Table 6 for both the Math/Econ and Applied Mathematics freshman groups, Complex Analysis does not have highly ranked ρ_H and β_i values in either of these student groups. This contrasts with the results found in Section 3.1. Another result to highlight is that in both the Applied Math freshman group and the Math/Econ freshman group, Linear Algebra is not selected as an impactful course by lasso regression. Furthermore, in the Applied Mathematics freshman group, Real Analysis (Part I) is also not selected as an impactful course through lasso regression. Although both courses are considered core by the mathematics department, other courses, such as Probability Theory (Part II) and Mathematical Modeling are more correlated with and impactful on overall mathematics GPA.

Table 6: Summary statistics for Applied Mathematics freshman admits.

Course	ρ_H Rank	ρ_H	β_i Rank	β_i
Probability (Part II)	1	1.0759	5	0.02965
Linear/Nonlinear Diff. Eq.	2	1.0714	–	0
Combinatorics	3	1.0616	–	0
Math. Modeling	4	1.055	1	0.08497
Optimization	5	1.0506	–	0
Num. Analysis (Part II)	11	0.9966	2	0.06327
Complex Analysis	9	1.01803	3	0.04523
Real. Analysis (Part I)	13	0.98749	–	0
Linear Algebra (Part I)	15	0.95237	–	0

4. SUMMARY AND INTERPRETATION OF RESULTS

In this paper, we developed methods to analyze the core courses selected by the mathematics department. We identified statistical attributes we expect the core courses to have. The first attribute was a correlated metric, denoted by ρ_H , which measured the correlation between a course grade and overall mathematics GPA. The second attribute was a metric of impact, denoted by β_i , which measured the impact a course grade has on overall mathematics GPA. We tested our methods on various mathematics majors at UCLA to validate the selection of core courses as selected by the mathematics department.

The most surprising results came from the pure mathematics freshman group. We summarize our results for pure mathematics students in Table 4. In this table, one can see that Complex Analysis appears in both results for the mixed correlation metric and lasso regression. This suggests that the grade in Complex Analysis is both highly correlated to and impactful on one’s overall mathematics GPA. Although this result is unique in that there is a single course that is highly correlated to and highly impactful on overall mathematics GPA within this student group, it may also suggest a weakness in our methods, as other student groups did not have a single course that ranked highest by both measures.

We can interpret the results from the pure mathematics freshmen by looking at instructors throughout the 15 year period. We see that Complex Analysis is more frequently taught by regular faculty than many other upper division courses at UCLA. In particular, about 52% of the professors who taught Complex Analysis were regular faculty. We can compare this to 46% of regular faculty who taught Real Analysis, and 39% of regular faculty who taught Linear Algebra. This may help explain the high correlation attributed to Complex Analysis. For instance, regular faculty may tend to give grades more indicative of a student’s overall performance. Additionally, regular faculty may teach a course more consistently than temporary faculty so that the grade in the course is a more accurate reflection of the student’s performance. If this is the case, the math department may benefit from assigning more regular faculty to courses which are seen as foundational.

Table 7: Sample of data provided by department of mathematics.

ID	Term	Major Code	Classification	Admit Class	Subject	Course	Grade
0x3915	09W	756	USO	UFR	MATH	0003A	3.3
0xE443	09W	671	USO	UFR	MATH	0003A	3
0XC044	09W	00X	USR	UFR	COMPT	10A	3.7

4.1. FUTURE WORK

Our core course analysis could be expanded upon with the inclusion of more data and by refining the students we consider. We purposefully selected versatile models to illustrate this core course analysis and how it could be used to better understand the mathematics curriculum. Going forward, we will utilize more directed and specific statistical models, derive the results from these models and compare them to the results presented in this paper.

We would also like to combine the lasso and ridge models to incorporate the sparsity of the ℓ_1 norm with the metrics coming from the ℓ_2 norm. In addition, we can expand our regression models to include explanatory variables such as average course load, whether or not the student switched majors, and how long a student took to declare a major. The goal of including such variables is to introduce more sources of variation in overall GPA into our model than just grades. We can analyze course rankings under the β_i values from this more comprehensive regression model and compare the associated rankings to those outlined in this paper. We can then investigate any discrepancies between the two rankings and provide the findings to the department which can act accordingly.

A STUDENT DATA

In this appendix, we briefly discuss the data that we investigated in this paper. As mentioned earlier, each row represents a course a student took, his grade in that course, and other pertinent information. The so-called Student Classification in Table 7 indicates their academic standing: freshman (UFR), sophomore (USO), junior (UJR), and senior (USR). The Admit Class is the academic standing when a student *first* entered UCLA. A first-year non-transfer student may enter as as USO or a UFR based on the number of college level credits earned while in high school. The grade values range from 0.0-4.3, where 4.3 indicates an A+ grade. Under the umbrella of the mathematics department there are two possible subject headings: MATH and COMPT, the latter representative of the computer based courses in the program for computation. The Course codes can be viewed at the [UCLA Department of Mathematics](#). The Major Codes and corresponding major names can be viewed at the [UCLA Registrar](#).

B ACKNOWLEDGEMENTS

We would like to thank Dimitri Shlyaktenko and Charles Z. Marshak for all the guidance they have provided us during this project and the writing of this paper. We thank Dimitri Shlyakhtenko and Andrea Bertozzi for their help with administrative issues related to use of this dataset. We also want to thank Mihai Cucuringu for initially helping us parse through the data and M. Puck Rombach for additional oversight and helpful discussions. Finally, we would like to thank our REU group members Dillon Montag, Milica Hadži-Tanović, and Juan Carlos Apitz. Ritvik

Kharkar was supported by NSF grant DMS-1045536 and Jessica Tran was supported by the UC LEADS program.

REFERENCES

- (1) Ryan Baker and Kalina Yacef. The state of educational data mining in 2009: A review and future visions. *JEDM-Journal of Educational Data Mining*, 1(1):3–17, 2009.
- (2) Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- (3) Brent Bridgeman, Laura McCamley-Jenkins, and Nancy Ervin. Predictions of freshman grade-point average from the revised and recentered SAT I: Reasoning test. *ETS Research Report Series*, 2000(1):i–16, 2000.
- (4) Ko-Kang Chu, Maiga Chang, and Yen-Teh Hsia. Designing a course recommendation system on web based on the students’ course selection records. In *World conference on educational multimedia, hypermedia and telecommunications*, volume 2003, pages 14–21, 2003.
- (5) Augustinos I. Dimitras, Stelios H. Zanakis, and Constantin Zopounidis. A survey of business failures with an emphasis on prediction methods and industrial applications. *European Journal of Operational Research*, 90(3):487–513, 1996.
- (6) Robert R. Glauber Donald E. Farrar. Multicollinearity in regression analysis: The problem revisited. *The Review of Economics and Statistics*, 49(1):92–107, 1967.
- (7) Thomas G. Edwards. Current reform efforts in mathematics education. ERIC/CSMEE digest. 1994.
- (8) Mark W Ellis and Robert Q Berry III. The paradigm shift in mathematics education: Explanations and implications of reforming conceptions of teaching and learning. *Mathematics Educator*, 15(1):7–17, 2005.
- (9) Rosta Farzan and Peter Brusilovsky. Social navigation support in a course recommendation system. In *Adaptive hypermedia and adaptive web-based systems*, pages 91–100. Springer, 2006.
- (10) National Science Foundation. Science engineering indicators 2012.
- (11) Chris Hans. Bayesian lasso regression. *Biometrika*, 96(4):835–845, 2009.
- (12) David Harville. Extension of the gauss-markov theorem to include the estimation of random effects. *The Annals of Statistics*, 4(2):384–395, 1976.
- (13) Sarah B. Henderson, Bernardo Beckerman, Michael Jerrett, and Michael Brauer. Application of land use regression to estimate long-term concentrations of traffic-related nitrogen oxides and fine particulate matter. *Environmental science & technology*, 41(7):2422–2428, 2007.
- (14) Dennis E. Hinkle, William Wiersma, and Stephen G. Jurs. Applied statistics for the behavioral sciences. 2003.
- (15) Xing Liu. Fitting stereotype logistic regression models for ordinal response variables in educational research. *Journal of Modern Applied Statistical Methods*, 13(2):31, 2014.
- (16) Robert McGill, John W. Tukey, and Wayne A. Larsen. Variations of box plots. *The American Statistician*, 32(1):12–16, 1978.
- (17) Yannick Meier, Jie Xu, Onur Atan, and Mihaela van der Schaar. Predicting grades. 2015.
- (18) David Lindsay Roberts. E.H. Moore’s early twentieth-century program for reform in mathematics education. *The American Mathematical Monthly*, 108(8):689–696, 2001.

- (19) Cristobal Romero and Sebastian Ventura. Educational data mining: A survey from 1995 to 2005. *Expert systems with applications*, 33(1):135–146, 2007.
- (20) Cristóbal Romero and Sebastián Ventura. Educational data mining: a review of the state of the art. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 40(6):601–618, 2010.
- (21) George Siemens and Ryan Baker. Learning analytics and educational data mining: towards communication and collaboration. In *Proceedings of the 2nd international conference on learning analytics and knowledge*, pages 252–254. ACM, 2012.
- (22) C. Spearman. The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1):72–101, 1904.
- (23) Cem Tekin, Jonas Braun, and Mihaela van der Schaar. etutor: Online learning for personalized education. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 5545–5549. IEEE, 2015.
- (24) Mark Hoover Thames and Deborah Loewenberg Ball. Making progress in us mathematics education: Lessons learned—past, present, and future. In *Vital directions for mathematics education research*, pages 15–44. Springer, 2013.
- (25) Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- (26) Jack V. Tu. Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes. *Journal of clinical epidemiology*, 49(11):1225–1231, 1996.
- (27) Daniel S Weld, Eytan Adar, Lydia Chilton, Raphael Hoffmann, Eric Horvitz, Mitchell Koch, James Landay, Christopher H. Lin, and Mausam Mausam. Personalized online education—a crowdsourcing challenge. In *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1–31, 2012.
- (28) H Wu. The mathematics education reform: Why you should be concerned and what you can do. *The American Mathematical Monthly*, 104(10):946–954, 1997.
- (29) H Wu. *On the education of mathematics majors*, volume 36. MSRI Publications, 1999.
- (30) Jie Xu, Tianwei Xiang, and Mihaela van der Schaar. Personalized course sequence recommendations. *arXiv preprint arXiv:1512.09176*, 2015.
- (31) Yonghong Jade Xu. An exploration of using data mining in educational research. *Journal of Modern Applied Statistical Methods*, 4(1):23, 2005.