Investigating institutional influence on graduate program admissions by modeling physics Graduate Record Examination cutoff scores

Nils J. Mikkelsen¹, Nicholas T. Young^{2,3}, and Marcos D. Caballero^{1,2,3,4,*}

¹Center for Computing in Science Education & Department of Physics,
University of Oslo, N-0316 Oslo, Norway

²Department of Physics and Astronomy, Michigan State University, East Lansing, Michigan 48824, USA

³Department of Computational Mathematics, Science, and Engineering, Michigan State University,
East Lansing, Michigan 48824, USA

⁴CREATE for STEM Institute, Michigan State University, East Lansing, Michigan 48824, USA

(Received 29 September 2020; accepted 24 December 2020; published 19 February 2021)

Despite limiting access to applicants from underrepresented racial and ethnic groups, the practice of using hard or soft Graduate Record Examination (GRE) cutoff scores in physics graduate program admissions is still a popular method for reducing the pool of applicants. The present study considers whether the undergraduate institutions of applicants have any influence on the admissions process by modeling a physics GRE cutoff score with application data from admissions offices of five universities. Two distinct approaches based on inferential and predictive modeling are conducted. While there is some disagreement regarding the relative importance between features, the two approaches largely agree that including institutional information significantly aids the analysis. Both models identify cases where the institutional effects are comparable to factors of known importance such as gender and undergraduate GPA. As the results are stable across many cutoff scores, we advocate against the practice of employing physics GRE cutoff scores in admissions.

DOI: [10.1103/PhysRevPhysEducRes.17.010109](https://doi.org/10.1103/PhysRevPhysEducRes.17.010109)

I. INTRODUCTION

While recent studies have called into question the overreliance on Graduate Record Examination (GRE) scores in physics graduate admissions [1,2], filtering applicants based on a strict or effective minimum score is still a popular practice today [3]. Given the role of the GRE in admissions, understanding the factors influencing GRE scores may provide insight into how, when compared to other science, technology, engineering, and mathematics (STEM) disciplines, the physics graduate admissions process has failed to improve gender, racial, and ethnic diversity by systematically excluding these applicants [4,5].

A number of studies have investigated correlations between GRE scores and demographics [1,2], but little attention has been given to the institutional backgrounds of applicants. An applicant's undergraduate background could play a significant role in their graduate application [6]. Institutions offering a Ph.D. program themselves would likely place more emphasis on both preparing and

motivating undergraduate students for further studies. Larger physics departments with more resources are able to offer students more advanced course work and hands-on experimental work as well as provide a larger variety of staff expertise. Larger undergraduate programs can facilitate network building, both between students and faculty members, and collaboration via projects and study groups. Although attributes such as motivation and opportunity cannot be appropriately measured, their effects on the GRE can be linked to metrics such as the size and type of institutions as was done in Halley *et al.* [7]. In order to estimate these institutional effects, we have analyzed the physics GRE subject test (P-GRE) scores of graduate program applications from four public universities and one private university.

The applications include a variety of information, but the present study will focus on numerical and categorical data, all of which constitutes a mixture of data structures. A number of recent studies working with similar data have approached the problem using machine learning methods [8,9]. Many machine learning methods lend themselves to problems with mixed data, albeit they do not share the interpretability of more conventional modeling methods. The present study will employ both approaches, contrasting and comparing the results.

The aim of this study is to continue the discussion on the practice of employing formal or informal P-GRE score cutoffs in graduate admissions using a combination of

*Corresponding author.
caballero@pa.msu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

modeling and machine learning methods. The idea is to analyze the P-GRE scores of Ph.D. program applicants with respect to applicants' undergraduate grade point average (U-GPA), demographics and institutional background. Our guiding research questions (RQs) are as follows.

- (1) To what extent does the applicant's undergraduate institution influence whether they are able to attain a minimum P-GRE score expected by an admissions committee?
- (2) To what degree do the institutional effects compare to known effects such as U-GPA, gender and race?
- (3) How do the results depend on the specific cutoff chosen by the admissions office?
- (4) How well do the conventional and machine learning approaches agree on RQs 1, 2, and 3?

II. BACKGROUND

Following the calls for increasing diversity in STEM disciplines, there has been a steady growth of women's and ethnic or racial minorities' representation over the past couple of decades [10]. Despite the progress however, physics has seen particularly poor development in comparison. Since the late 1990s, the percentage of bachelor's and Ph.D. degrees awarded to women in physics has stagnated at about 20%, mirroring similar numbers of engineering and computer science [4]. The numbers are even more concerning for racial minorities who during the three-year period 2014–2016 earned 11% of bachelor degrees and only 7% of Ph.D. degrees [10]. The discrepancy in female, racial, and ethnic representation likely stems from a variety of factors involving admission and retention issues, many of which are rooted in cultural and structural problems including sexual harassment and systemic racism [11–13].

In her extensive review of the general practices of graduate program admissions, *Inside Graduate Admissions* [14], Posselt notes that most admissions (in the natural sciences as well as in the humanities and social sciences) measured students' merit primarily on the basis of their undergraduate GPA (U-GPA) and GRE scores alone. Indeed, Young and Caballero were able to predict the admittance of prospective physics Ph.D. students with 75% accuracy using machine learning methods based only on their U-GPA and P-GRE score [8]. The GRE test makers, Educational Testing Service (ETS), recommends against the use of GRE scores as the sole basis for admissions decisions, particularly emphasizing the practice of filtering applicants based on a minimum cutoff score [15]. Despite this, Potvin *et al.* found that 32% of physics graduate program admissions state they filter applicants with a minimum P-GRE score [3]. Furthermore, of the programs that say they do not filter applicants, several reported using a “*rough cutoff*” or wanting a “*preferable score*,” suggesting that more than 32% of programs filter applicants in practice.

As highlighted by Miller and Stassun in Ref. [1], on average, women score 80 points lower than men on the GRE in the physical sciences, while Black test takers score 200 points lower than white test takers. The authors further note that the practice of filtering prospective students with a minimum score, which is in violation with ETS's own guidelines, thus “*adversely affects women and minority applicants*”. In addition to limiting access for minority applicants during the application process, the GRE also acts as a barrier to apply. In a survey of prospective students from underrepresented racial and ethnic groups interested in pursuing a Ph.D. in physics who ultimately chose not to apply, Cochran *et al.* notes that the GRE was the “*most common theme*” expressed by students as a barrier to apply [16].

In spite of its established popularity in admissions, the GRE's ability to identify promising students has recently been called into question. One study found that while requiring a minimum P-GRE score limits access to physics graduate program applicants from minority groups, GRE scores were incapable of predicting Ph.D. completion [2]. In a 2015 survey of prize-winning postdoctoral fellows in astronomy [17], Levesque *et al.* found that the P-GRE scores of fellows did not adhere to any minimum percentile score, suggesting that the GRE is also a poor estimator for future research excellence. The authors further point out that a minimum percentile score of 60% would have eliminated 44% of participants, including 60% of female fellows. The inability of the GRE to identify promising students has also been noticed by other groups such as the National Science Foundation, which recently decided to drop the GRE from the application to their Graduate Research Fellowship Program (see FAQ No. 52 [18]).

Prior work has typically focused on admissions committees' overreliance on the GRE and the consequences of using cutoff scores in graduate admissions [1,2,19,20]. Missing from the conversation is an understanding of what institutional factors, which come into play during applicants' undergraduate study (or even earlier), may influence GRE scores. In a 1991 study, Halley *et al.* investigated how the topics covered by P-GRE compared with physics major curriculum by analyzing the P-GRE scores of students from different institution types [7]. The authors noted that the portion of correct answers was higher for students from “top” institutions, and highest for students from top institutions with graduate programs. However, this study is both nearly 30 years old and worked with an imbalanced sample (701 test takers in total, 21 of which attended a top undergraduate institution). Since then, the GRE has evolved and the number of physics degrees awarded annually has almost doubled [21]. Nowadays, the GRE does not penalize incorrect answers, i.e., guessing, which has likely changed the way students approach the test. To our knowledge, there has not yet been a modern study analyzing how institutional factors may affect GRE scores.

A. Overview of the admissions process

For the unfamiliar reader, the typical admissions process for graduate physics programs in the U.S. starts with students preparing an application package. This package usually contains various requested documentation such as demographic information on the student, written statements outlining their research interests and why they are interested in attending graduate school, their undergraduate transcript(s), and test scores on the Graduate Record Exams including the physics subject test. For some graduate programs, GRE scores are optional, but for the schools in this study, they were required as part of the application. In addition, students request several letters of recommendation to be submitted confidentially to the graduate program. In the U.S., most students apply to Ph.D. programs directly after earning their bachelor's degree. This is in contrast to European systems where a student typically has to have earned a master's degree to apply. Admissions committees review applications packages based on the criteria set by their program and offer admission to students meeting those criteria. In some cases, students are denied admission if they fail to earn physics GRE scores above some threshold. What predicts earning those threshold scores is the principal research question of this paper.

III. METHODS

The target for this investigation is to explain whether a student scores *above* or *below* a P-GRE cutoff score selected by an admissions committee. This is encoded using a binary response variable named ABOVE with the interpretation that an applicant with a score *above or equal* to the cutoff has ABOVE = 1, and an applicant with a score *below* the cutoff has ABOVE = 0. That is, given a test score of x and a cutoff score of C , we define

$$\text{ABOVE} = \begin{cases} 1, & x \geq C, \\ 0, & x < C. \end{cases} \quad (1)$$

One should recall that the possible scores on GRE subject tests range from 200 to 990 in 10 point intervals. We have focused on P-GRE cutoff scores ranging from 620 to 800 points, corresponding to the 32nd and 67th national percentiles [22]. Typical P-GRE cutoff scores lie in the region of 700 [2].

The data used in this study consist of 2017–2018 physics graduate program admissions records from 4 public universities in the Big Ten Academic Alliance and one private Midwestern university, all of which are members of the Association of American Universities (AAU). The records contain unidentified profiles of program applicants with information regarding their GRE performance, undergraduate GPA, ethnicity and race, gender, etc. To justify combining the records into a single pool of applicants,

two comparisons are presented: Table I(a) provides a brief comparison of the admissions records and Table I(b) compares the National Research Council's R-rank and S-rank ratings [23]. The former demonstrates the similarity between the applicant pools, while the latter shows the degree to which the physics graduate programs share characteristics with top tier programs (R rank) and possess characteristics that are highly regarded by faculty (S rank).

In addition to personal information, the records also include which institution the applicants attended during their bachelor's degrees. Complementary data describing the bachelor institutions have been added from three sources: the 2015 Carnegie Classification of Institutions of Higher Education [24], Barron's selectivity index [25], and 2017–2018 surveys of American universities by the American Institute of Physics (AIP) [26,27]. The additional data describe several aspects of the institutions such as institution-wide admissions selectivity and the size of physics programs. The main idea is to study the statistical effects from applicants' institutional backgrounds using this complementary data.

A. On the data

The admissions records contain 5738 applications in total, but only 5314 (ca. 93%) of them include the students' P-GRE scores. Applications without P-GRE scores are ignored to avoid influencing the P-GRE distribution. Of the remaining applications, 2575 are domestic (ca. 48%). This study will focus entirely on domestic students for two main reasons. First, the P-GRE distribution for international students is much more saturated with perfect scores. The saturation problem is visualized in Fig. 1: The percentage of international students scoring above the selected cutoff scores both starts off much higher and falls off much slower than the percentage of domestic students. Second, because there is not a systematic collection of graduation records for non-U.S. schools, it is difficult to reliably collect the necessary information from every international student.

Because the applicants are not identified, several applications may come from the same student. While these applications are unique in the sense that each application addresses a different school, they count as duplicated applications in this analysis by virtue of being from the same student. Duplicate applications could have an effect on the results, most notably in the logistic regression model, which relies on independent observations (see Supplemental Material [28]). By comparing applications according to demographics and academic performance, a number of possible duplicate applications have been identified. In case all candidates are duplicates, roughly 18.8% of the data should be ignored. Because the applications are anonymous, i.e., it is impossible to determine whether two applications belong to the same student, we will conduct the analysis both with and without the possible duplicates.

TABLE I. A comparison between the four public Big Ten universities and the private Non-Big Ten Midwestern university based on the admissions records [Table (a)] and the National Research Council's R-rank and S-rank ratings [Table (b)]. (a) A comparison of the admissions records. The Non-Big Ten records share several similarities with the records from the other universities, especially in terms of women's representation, share of domestic applicants, and academic performance. Additionally, the discrepancy seen in the distribution of racial and ethnic applicants can be explained by the applicants who did not report ethnicity or race. Further details regarding the contents of the admissions records data are provided in Sec. III A. (b) A comparison of the R-rank and S-rank from the National Research Council [23]. For the sake of anonymity, only ranges of ranks are presented. The Non-Big Ten school certainly does not stand out as a contrast to the Big Ten schools. The similarities between the profiles of the universities justify analyzing the records as a single pool of applicants.

Quantity	Non-Big Ten School	Big Ten School 1	Big Ten School 2	Big Ten School 3	Big Ten School 4
Percentage of women	16.4	14.7	15.4	18.4	18.3
Percentage of not reported gender	0.0	0.8	0.0	1.3	0.0
Percentage of White	41.0	30.0	37.2	36.8	40.1
Percentage of Asian	37.7	9.8	42.4	5.5	44.0
Percentage of Black	1.5	0.9	1.0	1.2	1.0
Percentage of Latinx	4.1	0.6	0.9	2.7	7.9
Percentage of multiracial	2.8	5.8	6.4	2.0	2.0
Percentage of not reported ethnicity or race	12.7	52.9	12.0	51.7	4.9
Percentage of domestic applicants	53.3	42.9	45.6	49.4	54.0
Percentage of not reported nationality	0.0	0.5	0.0	0.0	0.0
Median undergraduate GPA	3.75	3.74	3.70	3.67	3.62
Median physics GRE score	880	880	880	800	820
Median analytical writing GRE score	4.0	4.0	4.0	3.5	4.0
Median verbal reasoning GRE score	160	157	158	156	157
Median quantitative reasoning GRE score	167	167	167	164	166

Rank	Non-Big Ten School	Big Ten School 1	Big Ten School 2	Big Ten School 3	Big Ten School 4
R-rank 5th percentile	1–10	1–10	1–10	11–20	31–40
R-rank 95th percentile	11–20	1–10	11–20	31–40	61–70
S-rank 5th percentile	1–10	1–10	11–20	21–30	31–40
S-rank 95th percentile	31–40	41–50	51–60	61–70	101–110

FIG. 1. A comparison of the P-GRE distribution between national data [22] and data used in this study. The analysis is primarily concerned with domestic applicants (yellow curve).

1. The raw features

In addition to the P-GRE score, thirteen features, or variables, have been selected for analysis. A summary of the features and their sources is given in Table II.

The features from the admissions records include the applicants' P-GRE score, U-GPA, gender, and race. Note that the gender feature is encoded as a binary variable; while we acknowledge that gender is not binary, more detailed descriptions were not collected by the admissions offices [29]. Similarly, different practices regarding the collection of data on racial and ethnic backgrounds has limited the scope of the race feature. See Posselt *et al.* for more details regarding the collection of data on racial and ethnic backgrounds by admissions offices [20]. The features from the admissions records constitute the applicant-specific component of the models, while the remaining features comprise the institutional component.

Of the Carnegie features, the two most prominent are the (2015) Carnegie (basic) classification of institutions and the (2015) undergraduate population profile classification. The basic classification is an overall categorization of the

TABLE II. A summary of the features used in this study.

Feature	Type	Source
Physics GRE score	Continuous	Admissions
Undergraduate GPA	Continuous	Admissions
Gender	Binary	Admissions
Race	Categorical	Admissions
Carnegie classification	Categorical	Carnegie
Undergrad population profile	Categorical	Carnegie
Funding category	Categorical	Carnegie
ACT selectivity category	Categorical	Carnegie
Minority serving institution	Binary	Carnegie
Barron's selectivity index	Categorical	Barron's
No. bachelor graduates (2017)	Continuous	AIP Survey
No. bachelor graduates (2018)	Continuous	AIP Survey
No. Ph.D. graduates (2017)	Continuous	AIP Survey
No. Ph.D. graduates (2018)	Continuous	AIP Survey

academic degrees offered and awarded by the institutions, e.g., *Doctoral university with high research activity* and *Master's college with large programs*. The undergraduate population profile classification characterizes the typical undergraduate population according to three metrics: portion of full-time undergraduates, academic achievements of first-year and first-time students, portion of entering transfer students. In addition, the Carnegie features also include the institutions' funding category and ACT selectivity category, and whether the institutions are minority serving institutions (MSI). The ACT category measures the entry selectivity of admissions offices by grouping all institutions according to the ACT scores of first-year bachelor students, and MSI indicates whether an institution satisfies the requirements for a minority serving institution [30].

Lastly, Barron's provides the Profile of American Colleges [25], which is an index for institution-wide admissions selectivity, and the AIP surveys provide the numbers of bachelor's and Ph.D. students graduating in physics.

The data will be analyzed using two different data analysis methods based on logistic regression modeling and predictive machine learning analysis (described in Sec. III B). As they stand, the raw features are not well suited for logistic regression due to computational issues as well as modeling-related difficulties. The remaining part of this section describes our data preprocessing and modeling choices. See Sec. V C for a discussion of potential issues. Because the predictive analysis requires less preprocessing than logistic regression, we provide a summary of all the models used in this study in Sec. III C to avoid confusion.

2. Underrepresented racial and ethnic minorities

The small representation seen of applicants from racial and ethnic minorities (Black, Latinx, Multi, and Native) is of computational concern because logistic regression fails poorly with low-frequency categories [31]. Because initial

FIG. 2. Estimated P-GRE distributions by racial and ethnic groups (number of applicants indicated in parenthesis). Note that the combined distribution normalizes the differences between the combined groups.

tests including every racial group produced results with limited statistical power (e.g., infinite p -value confidence intervals), we combined racial and ethnic minorities in an underrepresented minority (URM) category despite Teranishi's warning [32]. This also combines their P-GRE distributions (see Fig. 2), leading to loss of information. This issue is further discussed in Sec. V C.

3. The Carnegie classification and undergraduate population profile

While the Carnegie classification and undergraduate population profile support 34 and 16 unique categories each, the limited pool of applications leaves many categories empty or with only a handful of applicants. Most of the categories are difficult to combine into meaningful groups. Thus, to avoid computational issues the features are replaced by the binary labels: *Doctoral university w/ highest research activity* and *Most selective undergraduate population*.

4. Funding category and ACT selectivity category

Similar to the Carnegie features, both *Funding category* and *ACT selectivity category* have categories with too few applicants. To avoid complications, the features are reduced to binary labels *Public Funding* and *Most ACT-selective*, which, respectively, indicate whether the institution is publicly funded and if the institution is in the most selective ACT category.

5. Barron's selectivity index

Barron's selectivity index is an admissions selectivity measure that categorizes institutions according to school competitiveness. In decreasing order of competitiveness,

the categories include *most competitive*, *highly competitive*, *very competitive*, *competitive*, *less competitive* and *non-competitive*. Additional “plus” categories such as *highly competitive plus* have been collapsed into their corresponding ordinary levels. In this study, admissions selectivity is used as a metric for an institution’s resources and staff experience. Because admissions selectivity is expected to have an effect only for the most selective schools, the selectivity categories less competitive than *most* and *highly competitive* are combined to a *not as competitive* category.

6. Number of bachelor’s and Ph.D. graduates 2017–2018

In this study, the AIP features (see Table II) provide a measure of the size of undergraduate physics departments. As larger departments typically have more financial resources available and may offer students more opportunities for advanced coursework or research, the P-GRE scores of applicants from larger programs is expected to be higher [7]. However, because of the variety of institutions and physics programs, a systemic effect is expected to only emerge for very large physics programs. Instead of analyzing the raw number of graduates, a physics program is classified as *large* if the number of graduates is above the 75th national percentile [21].

While the typical size of physics departments is unlikely to change on a yearly basis for most institutions, the exact number of graduates is much more sensitive to variation. Moreover, the applicants spent several years at the undergraduate institutions, thus it is unreasonable to estimate the general size of the physics departments using data from a single year. Because the statistical models cannot include data on both years simultaneously (i.e., as individual features) due to correlation issues, the 2017 and 2018 data must be combined (bachelor’s and Ph.D. features separated). For most institutions, the difference in the number of bachelor’s and Ph.D. graduates between 2017 and 2018 is not significant enough to have any effect on the analysis. However, because the difference is large for some institutions, naively selecting, say, the average could overestimate or underestimate the size of some departments. In addition, there are some institutions with data missing for either 2017 or 2018. To avoid inaccurate single-point estimates of department sizes, the maximum and minimum cases are considered separately. In the *maximum graduates models*, the maximum number of bachelor’s and Ph.D. students between the 2017 and 2018 data are included, and vice versa in the *minimum graduates models*. For institutions with missing data, any available data are used for both models.

B. Methods for data analysis

The following section provides a brief overview of the methods used in this study. Additional details are provided as Supplemental Material [28]. Because logistic regression

is likely familiar to a greater audience, more time is spent on the machine learning methods.

1. Logistic regression modeling

Logistic regression analysis is a technique for modeling a binary response $y \in \{0, 1\}$ with respect to explanatory variables $x_1, \dots, x_k$, which may consist of a mixture of continuous and discrete data. While binary data are naturally handled by logistic regression, categorical (discrete) data with $M > 2$ number of categories must be encoded using $M - 1$ binary variables according to the *one-hot encoding* scheme (see Supplemental Material [28] for details). The response is modeled according to the odds equation,

$$\text{odds}(p) = \exp(\beta_1 x_1 + \dots + \beta_k x_k + \epsilon), \quad (2)$$

where p is the probability of the outcome $y = 1$, β_i is the regression coefficient of x_i and ϵ is an error term. The regression coefficients are determined numerically using an iterative scheme based on maximum likelihood estimation. In our study this is handled by the `glm` function in R [33].

A major benefit of logistic regression modeling is the interpretability of its regression coefficients. When x_i increases by 1 unit, the odds change by a factor of $\exp(\beta_i)$ called the odds ratio:

$$\text{OR}(p; x_i) = \frac{\text{odds}(p; x_i + 1)}{\text{odds}(p; x_i)} = e^{\beta_i}. \quad (3)$$

The interpretation of the odds ratio depends on whether x_i is continuous or categorical. For continuous features, the change is associated with a unit increase in x_i . For binary features, the change is associated with a switch in x_i from category 0 to category 1. Because multileveled categorical features are encoded with binary features, each binary represents a change from the reference category to the category associated with the binary. Odds ratios below 1 are inverted so that $1/\text{OR}(x_i)$ is the odds ratio associated with a unit decrease in x_i or a switch in x_i from category 1 to category 0. In order to avoid interpretation issues relating to very large or very small continuous features, it is customary to standardize continuous features by centering the mean about 0 and normalizing the variance to 1. For standardized features, the odds ratio is associated with an increase in the original feature by 1 standard deviation.

Alongside the regression coefficients, the `glm` function provides the corresponding p values. To avoid multiple comparisons problems, the p values are adjusted according to the Bonferroni correction. For a logistic regression model with N features, the Bonferroni-adjusted p value is $\tilde{p} = pN$. We follow common practice and include three levels of significance: $\alpha = 0.05$, $\alpha = 0.01$, and $\alpha = 0.001$.

Because logistic regression is unable to handle missing values, we follow Nissen *et al.*’s recommendation of

imputing the missing data instead of discarding it [34]. Our approach employs the multiple imputation by chained equations (MICE) algorithm, which is handled by the MICE package in R [35]. MICE is an iterative algorithm that applies linear and logistic regression techniques in order to impute the data while conserving the relationship between the features as well as possible. The algorithm constructs N individual datasets to be modeled separately, the results of which are *pooled* (combined) according to Rubin's rules [36]. In this study, 5 imputation sets were created using 20 iterations (leaving other `mice` parameters to their defaults). Because the raw features are processed, the transformation must occur either before, after, or during the imputation. To our knowledge, there are no recommended strategies for the kinds of transformations used in this study. We therefore follow the general recommendation of von Hippel of “*impute, then transform*” [37]. As recommended by Moons *et al.* [38], the P-GRE scores are included in the imputation before preparing ABOVE.

2. Machine learning analysis

Whereas logistic regression favors interpretability (via the odds ratios), machine learning analysis (MLA) focuses on making accurate and reliable predictions. Given inputs $x_1, \dots, x_k$ and an output y , the goal of MLA is to identify a map f such that

$$y = f(x_1, \dots, x_k) + \epsilon, \quad (4)$$

where ϵ is a prediction error. When y is categorical (e.g., binary), f is called a *classifier* because it classifies a set of inputs into discrete outputs. As classifiers are seldom perfect, a major component of MLA consists of finding the optimal f , i.e., minimizing ϵ . To measure “how well” a classifier is able to classify inputs we use *performance metrics*. Different metrics highlight different types of behavior, meaning a classifier can score well according to one metric, but poorly according to another metric. This study employs two metrics: *prediction accuracy score* and *AUC-ROC score*.

The prediction accuracy score of a classifier is the portion of correctly classified cases. In terms of our data, a correctly classified case is any application for which the classifier successfully predicts whether the applicant scores above or below the cutoff score. It is typically referred to as simply the accuracy and it is often reported as a percentage. Accuracy is a number between 0% and 100%, where 100% signifies a perfect classifier. While easy to interpret, accuracy is very sensitive to unbalanced output classes (see the “Domestic applicants” curve in Fig. 1 for the class imbalance faced in this study) because it does not distinguish between the output classes. For instance, if 80% of applicants score above the cutoff, then a naive classifier predicting *above* regardless of the inputs will have an accuracy score of 80%. For this reason, accuracy should always be considered

relative to class imbalance. Furthermore, because the class imbalance changes as the cutoff increases (Fig. 1), the interpretation of the nominal accuracy score changes. Hence, the accuracy scores of two classifiers using different cutoffs should not be compared nominally.

The AUC-ROC score is a more complex metric than accuracy. Here, ROC refers to a receiver operating characteristic curve and AUC means taking the area under the ROC curve. For more details regarding ROC curves, consult the Supplemental Material [28]. The AUC-ROC score, or simply the AUC, is a measure of a classifier's ability to distinguish between output classes. AUC is a number between 0 and 1, where 1 signifies a perfect classifier, but a score of 0.5 is equivalent to complete guesswork. There is no universal scheme for judging AUC scores, but Hosmer *et al.* provides a rough guide: $0.7 \leq \text{AUC} < 0.8$ is acceptable, $0.8 \leq \text{AUC} < 0.9$ is excellent, and $0.9 \leq \text{AUC}$ is outstanding [31]. In contrast with the accuracy score, AUC is more robust towards imbalanced output classes [39], and thus AUC scores can be more reliably compared across different cutoff scores.

MLA typically consists of two phases: training and testing. Here, *training* refers to the construction of a classifier, and *testing* refers to its evaluation based on performance metrics. A typical problem in MLA known as *overfitting* arises when a classifier is trained to recognize “too many details” of a dataset. Thus, instead of replicating the general trend of the dataset, the classifier replicates the random errors. To avoid this, it is standard practice to use different datasets for the training and testing phases by splitting the (complete) dataset at random. Because random splits can have unforeseen consequences, it is common to conduct several training-testing procedures and average the performance metrics, using the standard errors of the averages as indicators for the confidence intervals. This study employs the K -fold cross-validation algorithm with $K = 10$ to prepare the random splits [40].

It is important to note that to find a perfect classifier is typically considered impossible, even if $\epsilon = 0$ for all known data. Thus, there is no correct algorithm for constructing f , and in fact, there are many unique algorithms to choose from. This study employs the conditional inference forest (CIF) algorithm, which is variant of the earlier random forest algorithm [41,42]. A random forest is comprised of an ensemble of decision trees, each of which is an independent classifier. A decision tree is an algorithmic approach to decision making (predictions) that asks a series of yes-no questions based on the input data (e.g., male? and $\text{GPA} > 3.0$?). The questions are determined during the training phase and are chosen to optimize performance. Each tree is given a random sample of the training set and a random selection of the input features. Predictions of the forest are then based on a majority vote among the predictions of the trees. A CIF is similar to a random forest in principle, but differs in its construction.

This study employs the CIF algorithm via the `party` package in R [43–45]. The forests were built using 200 trees and 3 features per tree (following the recommended $\sqrt{p}$ [46]), all other parameters kept at their defaults. One of the selling points of the CIF is that it provides a natural way of measuring the importance of each feature in the model. The process of preparing the importance measures for each feature is also handled by `party`. The idea is to remove a feature from the forest and measure the resulting change in a performance measure, interpreting a larger change as the feature being more important. As described in Janitza *et al.*, measuring AUC loss is preferred due to its robustness with imbalanced data [47]. The importance measure is a tool for comparing the relative importance of features and should not be interpreted further [48].

Because the importance measures focus on the impact of removing each feature separately, a backwards recursive feature elimination (RFE) procedure is conducted to study the effect of removing several features. (see e.g., Ref. [40]) To restrict the scope, the procedure is only executed for P-GRE cutoffs in intervals of 30 point RFE is an iterative process that involves training a forest, estimating its performance, and removing the least important feature from the set of active features. Starting with all features, the process is repeated until one feature remains. The order of removal is determined by the importance measures of the forest model. The importance measures are computed using the complete model, i.e., *not* during the procedure, to avoid overfitting [49]. Because the importance measures vary depending on the cutoff, one would ideally prepare a removal order separately for each cutoff and conduct a unique RFE for each cutoff. However, because the importance measures are similar for different cutoffs, an average removal order is used for all cutoffs.

C. A summary of the models

Most of the data preprocessing described in Sec. III A is done for logistic regression. This includes combining racial and ethnic minorities in an underrepresented minority category; reducing the Carnegie features *Carnegie Classification*, *Undergraduate Population Profile*, *Funding Category*, and *ACT selectivity category* to binary labels; combining the Barron’s selectivity categories less competitive than *most* and *highly competitive* to a *not as competitive* category; and categorizing physics programs (both undergraduate and graduate) as *large* if the number of graduates is above the 75th national percentile. Because the computational difficulties of logistic regression related to multicollinearity and low-frequency categories are circumvented by the decision-tree construction of the CIF algorithm, none of these preprocessing procedures are required for the data to be compatible with the CIF models. With the exception of combining the 2017 and 2018 graduates data into minimum and maximum cases, the data are only preprocessed for the logistic regression models.

To avoid data preprocessing for the CIF models is actually in line with the philosophy of the predictive modeling approach. In contrast with how logistic regression emphasizes interpretability, machine learning is only interested in the relationship between the features and the response. Preprocessing the data dilutes the available information, and thus may negatively affect the predictive analysis (e.g., as in Fig. 2).

Overall, $19 \times 2 \times 2$ logistic regression and CIF models are studied: there are 19 unique P-GRE cutoffs under consideration, and for each cutoff, a model is constructed with and without the potential duplicate applications (Sec. III A), and using the maximum and minimum number of graduates between 2017 and 2018 (Sec. III A 6).

IV. RESULTS

Because of the large volume of similar results, we will primarily present the results (odds ratios of logistic regression and feature importance measures of the CIF) for the models including the potential duplicate applications. We discuss deviations from these results where relevant.

A. Key findings

The statistical effects from the institutional features become more involved, both in the logistic regression models and the CIF models, as the P-GRE cutoff increases. In particular, applicants from well-funded institutions with large physics programs and high research activity are more likely to score above the cutoff. The logistic regression models and the CIF models identify several examples where the effect of an institutional feature is comparable to U-GPA or gender. While U-GPA and gender are integral components of every model (as expected), the race and ethnicity of applicants did not contribute as much to our models as anticipated based on the differences in scores between racial and ethnic groups found by Miller *et al.* [2]. Overall, the logistic regression approach and the machine learning approach typically agree on whether a feature has any relevance in the model. Having said that, the odds ratios typically identify a larger set of important features that, in addition, changes as the P-GRE cutoff increases.

When it comes to the maximum and minimum graduates models, the contributions from the AIP features are devalued in the minimum graduates models in favor of Barron’s selectivity index and high research activity. Of the maximum models, the logistic regression models favor attending an institution with a large bachelor program over a large Ph.D. program, while the opposite is true in the CIF models. Finally, the analysis as a whole is similar for the models with and the models without possible duplicate applications. Specifically, by removing the possible duplicates, some features become less significant in the logistic regression models and the performance of some CIF models are slightly reduced.

FIG. 3. Summary of odds ratio significance levels of 19 independent models: Each P-GRE cutoff score (horizontal coordinate) denotes an independent maximum graduates logistic regression model that includes possible duplicates (see Sec. III C). In each model, the fields of the statistically significant features are marked with a symbol indicating the Bonferroni-corrected significance level: circle indicates $\alpha = 0.05$, triangle indicates $\alpha = 0.01$, and diamond indicates $\alpha = 0.001$. Blank fields indicate statistically insignificant results. The dotted separation lines categorize the features as applicant-related, institution-wide metrics and concerning Physics departments.

B. Logistic regression

1. Significance analysis

Consider first the maximum number of graduates models. Figure 3 shows a diagram indicating how the set of significant features changes between the logistic regression models as the P-GRE cutoff score is increased from 620 to 800. The features are typically significant for every or almost every cutoff, for no or few cutoffs, or for higher cutoffs. In the following we provide an overview of the applicant-related and institutional features that are statistically significant.

Of the applicant-related features, the odds ratios of U-GPA and gender are always statistically significant. However, contrary to expectations, odds ratios between applicants from different racial groups were only statistically significant for some cases. In particular, when compared to applicants identifying as White, the odds ratios for applicants identifying as Asian are significant for higher cutoffs, while the odds ratios for applicants identifying as Black, Latinx, Multi, or Native are only significant for a few cutoffs in the region of ≈ 710 . This is further discussed in Sec. V C.

When it comes to the institutional features, those statistically significant to a majority of the P-GRE cutoffs include attending a most competitive institution, an institution practicing some of the highest amounts of research activity, and an institution with a large physics bachelor's program. Interestingly, attending a highly competitive institution is only significant for cutoffs between 640 and 690, while attending one of the most competitive institutions is significant for all cutoffs up to 760. Additionally, attending private universities or universities with large Ph.D. programs becomes significant when the cutoff increases beyond ≈ 740 . In contrast, attending an MSI, a most ACT-selective institution or to graduate in a most selective undergraduate population profile is never significant, regardless of the cutoff.

In order to provide a rough overview of the difference between the maximum and minimum number of graduates models, Table III shows the fraction of P-GRE cutoffs for which each feature is significant. Note that the table also separates models with and without the possible duplicate applications. By removing the possible duplicates, the general significance of the features decreases. The change does not seem to originate in any particular feature as, with

TABLE III. Fraction of P-GRE cutoffs for which the odds ratio of each feature is statistically significant in the logistic regression models. There are 19 logistic regression models in each category (see Sec. III C). The first column (the maximum graduates models with possible duplicates) corresponds to the significance diagram (Fig. 3).

Variable	Maximum		Minimum	
	With	Without	With	Without
(Intercept)	0.63	0.58	0.58	0.58
Undergraduate GPA	1.00	1.00	1.00	1.00
Gender	1.00	1.00	1.00	1.00
Combined B, L, M, and N	0.26	0.00	0.21	0.00
Asian	0.53	0.37	0.47	0.42
Highly competitive	0.26	0.26	0.63	0.26
Most competitive	0.79	0.89	1.00	1.00
Highest research activity	0.89	0.63	1.00	0.63
Most selective UG p. profile	0.00	0.00	0.00	0.00
Most ACT selective	0.00	0.00	0.00	0.00
Privately funded	0.42	0.21	0.32	0.21
Minority serving institution	0.00	0.00	0.00	0.00
Large bachelor's program	1.00	0.74	0.37	0.21
Large Ph.D. program	0.32	0.05	0.32	0.11

the exception of U-GPA, gender and most competitive, the fraction of significant cutoffs is reduced for all features. Compared to the maximum graduates models, the typical significance of attending a large bachelor program is considerably lower in the minimum graduates models. Notably, the difference corresponds with an improvement in the fraction of significant cutoffs for attending a competitive school or an institution with high research activity, thus suggesting the variables may suffer from a confounding issue (see Sec. V B for a discussion).

Considerable changes in the set of significant features are only observed for large changes in the cutoff score. We therefore only discuss the odds ratios corresponding to cutoffs 650, 710, and 770, representing the lower, middle, and higher regions, respectively.

2. Odds ratios

The odds ratios for the maximum and minimum number of graduates models are shown in Tables IV(a) and IV(b), respectively. First and foremost, improving one's undergraduate GPA by 1 standard deviation, roughly equivalent of improving a B to a B+, improves the odds of scoring above the cutoff by at minimum a factor of 2.5 (increases to ≈ 2.8 for higher cutoffs). This substantial increase in odds reflects the importance of U-GPA in admissions expressed by both admission committees and prospective students [3,50]. Additionally, the odds of scoring above the cutoff is $1/0.17 \approx 5.9$ times greater for male applicants than for female applicants. The odds ratios of U-GPA and gender are consistent for all P-GRE cutoffs in both the maximum and minimum number of graduates models.

While the benefit of attending a competitive institution diminishes as the P-GRE cutoff increases from 650 to 710 and 770, attending one of the most competitive institutions is always preferable to a highly competitive institution. For cutoffs 650 and 710, the odds increase from attending a most competitive school is similar to the applicant increasing their U-GPA from a B to a B+. The model also finds institutional funding and high levels of research activity to be important factors. For high P-GRE cutoffs (e.g., 770), the odds of scoring above the cutoff is about 2 times as large for applicants who attended a private university compared to applicants who attended a public university. Similarly, for applicants attending a university that practices some of the highest levels of research activity, the odds ratio is roughly 1.6–2.0 depending on the cutoff.

Applicants from institutions with large physics programs typically also score higher. In the maximum graduates number of models, having attended a university with one of the largest undergraduate physics programs improves the odds of scoring above the P-GRE cutoff by a factor of about 1.7–2.0 (typically closer to 2.0). When the cutoff is high, a similar effect is seen for students attending a university offering a large graduate program (an odds ratio of about 1.6). In the minimum number of graduates models, the odds ratios are only statistically significant for the highest cutoffs (≥ 760). They are also typically smaller than the corresponding odds ratios in the maximum graduates models. The only statistically significant example in Table IV(b) is the odds ratios for attending an institution with one of the largest Ph.D. programs.

The remaining variables, i.e., most ACT selective, most selective undergraduate population profile, and MSI, contribute little to none.

C. Conditional inference forest

The general performance of the CIF models is shown in Fig. 4. Alongside the accuracy score is the class imbalance, which provides the baseline from which the accuracy score is interpreted. Because the imbalance is considerably high for lower cutoffs, the accuracy score is more representative of the CIF's ability to identify applicants scoring above the cutoff when the cutoff is higher (as the imbalance decreases as the cutoff score increases, the accuracy becomes increasingly more representative). However, because the imbalance level is outside the standard errors of the K -fold estimate, it is reasonable to conclude that the CIF is not simply predicting the majority class. Additionally, the AUC score is mostly *outstanding* (>0.9) and, more importantly, very stable with respect to changes in the P-GRE cutoff. The stability of the AUC coupled with the high score suggests that the results of the model may be reasonably interpreted, that is, that the feature importances provide a reasonable picture of the relationship between the features and the output for all P-GRE cutoffs. Because of the similarity in performance between the maximum and

minimum graduates CIF models, we present only the maximum graduates models going forward.

Figure 5 graphs the change in the importance measure of the features as the cutoff increases. The plot shows evidence of distinct groups of features with similar importances. The first group consists only of undergraduate GPA, whose importance measure is about 2 times higher than any other feature. The next group consists of gender and the number of Ph.D. graduates, which stand out when compared to the remaining group of the least important features (see Sec. III B 2 for how to interpret the importance measure). With the exception of some minor variation, the importance measure of U-GPA is fairly stable across all models. Notably

however, while the importance measure decreases for gender as the cutoff increases, it simultaneously increases for the number of Ph.D. graduates. Hence, for cutoffs greater than ≈ 750 , the model finds a greater statistical difference between applicants scoring above and below the cutoff when given the number of Ph.D. graduates compared to an applicant's gender. Proportional to their own importance measures, several features in the remaining group undergo large changes in importance measures. However, because these variations are small when compared to U-GPA, gender and the number of Ph.D. graduates, they should not be overemphasized.

The results of the feature elimination procedure are shown in Fig. 6. The diagram is arranged such that the

TABLE IV. Odds ratios for P-GRE cutoff scores 650, 710, and 770 of the logistic regression models with possible duplicates (see Sec. III C). The maximum and minimum graduates models are separated in tables (a) and (b), respectively. Statistically significant odds ratios are marked with asterisks [see below (b)]. Note that $\tilde{p} = 14p$ refers to the Bonferroni-corrected p values.

(a) Bachelor maximum's and Ph.D. graduates models						
Variable	Cutoff 650		Cutoff 710		Cutoff 770	
	OR	95% CI	OR	95% CI	OR	95% CI
(Intercept)	1.72	[1.09, 2.72] **	0.81	[0.53, 1.25]	0.29	[0.18, 0.47] ***
Undergraduate GPA	2.53	[2.12, 3.03] ***	2.63	[2.22, 3.11] ***	2.87	[2.40, 3.42] ***
Gender	0.16	[0.10, 0.24] ***	0.18	[0.12, 0.27] ***	0.17	[0.11, 0.25] ***
Asian	1.37	[0.73, 2.54]	1.61	[0.95, 2.71]	2.20	[1.34, 3.59] ***
Combined B, L, M, and N	0.68	[0.41, 1.11]	0.65	[0.42, 1.00]	0.76	[0.49, 1.19]
Highly competitive	1.98	[1.10, 3.58] *	1.55	[0.95, 2.52]	1.22	[0.76, 1.93]
Most competitive	3.40	[1.76, 6.56] ***	2.79	[1.53, 5.07] ***	1.62	[0.97, 2.73]
Doc. inst. with highest research activity	1.68	[1.00, 2.81] *	1.94	[1.23, 3.04] ***	1.63	[1.06, 2.51] *
Most selective UG population profile	1.72	[0.68, 4.37]	1.85	[0.73, 4.68]	1.75	[0.71, 4.32]
Most ACT selective	0.66	[0.26, 1.69]	0.49	[0.19, 1.27]	0.66	[0.25, 1.75]
Privately funded	0.96	[0.57, 1.62]	1.42	[0.91, 2.21]	2.07	[1.38, 3.11] ***
Minority serving institution	1.17	[0.62, 2.21]	1.05	[0.60, 1.82]	1.13	[0.64, 1.99]
Large bachelor's program	1.96	[1.23, 3.12] ***	1.74	[1.15, 2.65] **	1.94	[1.28, 2.92] ***
Large Ph.D. program	1.28	[0.72, 2.25]	1.31	[0.83, 2.08]	1.62	[1.07, 2.45] **

(b) Minimum bachelor's and Ph.D. graduates models						
Variable	Cutoff 650		Cutoff 710		Cutoff 770	
	OR	95% CI	OR	95% CI	OR	95% CI
(Intercept)	2.03	[1.31, 3.15] ***	0.94	[0.61, 1.45]	0.34	[0.22, 0.53] ***
Undergraduate GPA	2.50	[2.08, 2.99] ***	2.59	[2.19, 3.08] ***	2.82	[2.35, 3.39] ***
Gender	0.17	[0.11, 0.25] ***	0.19	[0.13, 0.28] ***	0.17	[0.12, 0.26] ***
Asian	1.41	[0.72, 2.75]	1.63	[0.96, 2.77]	2.19	[1.35, 3.56] ***
Combined B, L, M, and N	0.66	[0.41, 1.07]	0.65	[0.42, 0.99]	0.77	[0.50, 1.17]
Highly competitive	2.21	[1.25, 3.93] ***	1.68	[1.00, 2.80] *	1.35	[0.83, 2.18]
Most competitive	3.69	[2.07, 6.55] ***	3.09	[1.89, 5.05] ***	1.87	[1.17, 3.00] **
Doc. inst. with highest research activity	1.72	[1.02, 2.92] *	2.10	[1.33, 3.33] ***	1.83	[1.20, 2.79] ***
Most selective UG population profile	1.56	[0.60, 4.05]	1.72	[0.62, 4.76]	1.59	[0.62, 4.07]
Most ACT selective	0.75	[0.28, 2.03]	0.56	[0.19, 1.67]	0.74	[0.28, 1.99]
Privately funded	0.88	[0.53, 1.45]	1.27	[0.81, 1.98]	1.88	[1.27, 2.79] ***
Minority serving institution	1.11	[0.60, 2.06]	1.00	[0.58, 1.74]	1.11	[0.64, 1.93]
Large bachelor's program	1.44	[0.91, 2.29]	1.28	[0.84, 1.94]	1.47	[0.99, 2.20]
Large Ph.D. program	1.43	[0.83, 2.45]	1.34	[0.86, 2.08]	1.54	[1.05, 2.26] *

*** $\tilde{p} \leq 0.001$.

** $\tilde{p} \leq 0.01$.

* $\tilde{p} \leq 0.05$.

FIG. 4. Overall performance of the conditional inference forest. The standard errors of the K -fold ($K = 10$) estimates are indicated by the error bars. While the ABOVE class imbalance is very high for lower cutoffs, the accuracy standard errors are always above the imbalance level. The AUC score is mostly above 0.9, which Hosmer *et al.* categorizes as “outstanding” [31].

FIG. 5. Conditional inference forest feature importance measures. The standard errors of the K -fold ($K = 10$) estimates are indicated by the error bars. Features included in the CIF are listed in the legend in decreasing order of average (across all cutoffs) decrease in AUC upon removal.

FIG. 6. Conditional inference forest feature elimination procedure. The standard errors of the K -fold ($K = 10$) estimates are indicated by the error bars. The features are eliminated from left to right, where the named feature is currently the least important feature, and thus the next to be dropped from the model.

FIG. 7. Performance comparison between conditional inference forests with all features, with two institutional features, and without institutional features (only U-GPA, gender, and race). The standard errors of the K -fold ($K = 10$) estimates are indicated by the error bars. The significant improvement in performance by the simple addition of two institutional features suggests that the contribution from the institutional features is captured by a few features.

features are removed left to right, starting from a complete model and ending with a model that only includes U-GPA (i.e., the named feature at a given horizontal coordinate currently has the lowest importance measure). The accuracy and AUC scores largely agree on the effect of removing a feature. Specifically, removing MSI through Barron's selectivity index has no detrimental effect to the CIF's accuracy and AUC, and despite the high importance measure of the Carnegie classification, the model does not perform worse once it is removed either. Using only the three features with the highest importance measures (U-GPA, gender and number of Ph.D. graduates), the CIF is able to score ≈ 0.9 on the AUC metric and roughly between 80% and 90% on the accuracy metric.

Figure 3 shows that the set of statistically significant features in the logistic regression models changes as the P-GRE cutoff score increases (e.g., whether the institution is privately funded is only significant for cutoffs ≥ 740). A similar change is not present in the importance measures of the CIF models (Fig. 5), which, in contrast with the odds ratios, preserve the feature groups described above. In particular, the features U-GPA, gender, and number of Ph.D. graduates, are the three most important features for every cutoff score. Because the importance measures of the remaining features are consistently lower by a considerable margin for all cutoff scores, the set of important features in the CIF models is very robust towards changes in the cutoff score.

As a final check for whether the added performance can be attributed to including the institutional features, the performance of the full CIF is compared to a CIF excluding all institutional features, and a CIF including the number of Ph.D. graduates and the Carnegie classification. The results of the comparison is summarized in Fig. 7: The addition of only two institutional features makes a considerable improvement for both metrics, regardless of the cutoff. Hence, the added performance is reasonably attributed to the inclusion of institutional features.

V. DISCUSSION

A. Research questions

This study investigated four research questions (RQs) that we address in order.

- (1) To what extent does the applicant's undergraduate institution influence whether they are able to attain a minimum P-GRE score expected by an admissions committee?
- (2) To what degree do the institutional effects compare to known effects such as U-GPA, gender, and race?
- (3) How do the results depend on the specific cutoff chosen by the admissions office?
- (4) How well do the conventional and machine learning approaches agree on RQs 1, 2, and 3?

Regarding RQ 1, the institutional background helps explain whether a student scores above a given P-GRE

cutoff. Consider a cutoff score of 710, which is just above the most common cutoff score of 700. In the logistic regression models (see Table IV), applicants from competitive institutions with large physics programs, practicing high levels of research are statistically more likely to score above the cutoff than other applicants. Similarly, the size of physics programs (number of graduates) and the institution-wide Carnegie classification are integral components of the predictive capacity of the CIF models (see Sec. 6). Hence, the models suggest that to employ a cutoff score of 710 not only limits access to racial and ethnic minorities [2], but also to applicants from smaller universities with less resources that are less competitive and practice lower (not necessarily among the lowest) levels of research. Similar observations are found for every other cutoff in the CIF models. In the case of the logistic regression models, the set of statistically significant institutional features varies depending on the cutoff, but the overall interpretation is similar: To include institutional data in the analysis certainly helps explain whether a student scores above the cutoff, regardless of the chosen cutoff.

Now, is it necessary to include a complete description of an applicant's undergraduate background? Figure 6 suggests that this is probably not the case as a large portion of the institutional data do not contribute to the models. Moreover, because the performance of the CIF does not decrease as the Carnegie classification is removed, there is also reason to suspect that the institutional features may share information. The independence of the features is discussed more in detail in Sec. V B.

The modeling and machine learning approaches disagree somewhat with respect to RQ 2. In the logistic regression model, the odds ratios for U-GPA is comparable to admission competitiveness (roughly 2–3), while the odds ratios for gender is just shy of 6.0. In contrast, U-GPA is by far the most important feature in all CIF models. Meanwhile, the feature importance measure of gender is similar to the number of Ph.D. graduates, particularly for higher cutoffs (≥ 750). Because neither approach placed as much emphasis on race and ethnicity, it is unreasonable to judge the overall effect of institutional data by comparing it to the effects of race and ethnicity in the models. Despite disagreeing on some of the finer details, both approaches find examples where the effects from institutional data, e.g., admission competitiveness and the size of physics departments, are comparable to U-GPA and gender. The most clear-cut example is shown in Fig. 7, which demonstrates that to replace a CIF model without institutional features with a similar CIF model that includes the Carnegie classification and number of Ph.D. graduates provides a blanket improvement in the accuracy and AUC scores for every P-GRE cutoff.

Finally, we address RQs 3 and 4 together. First and foremost, both approaches have identified statistically significant differences in P-GRE scores of applicants with

different institutional backgrounds. Having said that, the specifics regarding the statistical difference and the extent to which it is explained by different institutional backgrounds depends on the model and cutoff in question. For instance, the significance level of odds ratios vary to such an extent that some features are only relevant for a select few cutoffs (e.g., private or public institution for higher cutoffs). The importance measures of the CIF models are much more stable across cutoffs, but lacks the interpretability of the odds ratios. Nevertheless, while the set of useful features changes with the cutoff, institutional features always contribute to the analysis. Here, logistic regression disagree with the CIF on the set of useful features and their importance to the model, but both recognize useful institutional features for every cutoff score.

B. Limitations

Central to this study is the question of whether the institutional background of an applicant can be reliably measured, or estimated, with the available data. Here, “institutional background” is used in an extended sense that includes the applicant’s experiences in relation to attending a particular institution. Our data certainly do not allow for quantifying the effects of such experiences as studying in an encouraging environment or at an institution with a large array of opportunities. However, data such as the Carnegie features and the number of graduating bachelor’s and Ph.D. students likely capture some aggregate effect of studying at different types of institutions. In addition, these features were found to be important in our models, suggesting that there is a statistical difference between the applicants that is dependent on the institutions.

Because the universities considered in this study are typically highly regarded, the data likely suffer from a selection bias effect, favoring prospective students with higher grades and GRE scores. In a 2018 survey of prospective students from racial and ethnic minorities, Cochran *et al.* identified concerns regarding GRE scores and undergraduate GPA as commonly expressed barriers to apply to physics graduate programs [16]. Indeed, this is reflected in the P-GRE distribution of the applicants in our dataset: Figure 1 shows that the applicants consistently score as high or higher than the national averages, thus implying our data set consists of a biased selection of all prospective students (the data set comprises an upper limit of $\approx 18\%$ of all P-GRE test takers in 2017–18 [22]). Because of this selection bias, the distributions of the other features in our dataset are likely also biased. Most prominently, the selection bias will disproportionately affect women, and racial and ethnic minorities [4,10]. The problem of selection bias and its consequences for physics education research as a whole was recently discussed in Kanim and Cid [51]. Our findings should thus be considered in light of our biased sample and their discussion.

A related, but different issue is that applicants are more likely to have attended large programs by virtue of there being more prospective students from larger programs than smaller programs. This can be seen in our data from the median number of bachelor graduates. Whereas the national median was 8 in both 2017 and 2018 [26,27], the median in our data is 27 (2017) and 30 (2018), i.e., more than 3 times as many. Consequently, our data consist of a larger fraction of applicants from larger programs than usual, and thus the distributions of all the features in our data set are likely primarily determined by applicants from larger programs. This also contributes to the selection bias discussed above.

Another methodological problem is the question of whether the different institutional variables attempt to describe the same effect, implying a possible problem of correlation, or even multicollinearity, between the features. The number of bachelor’s and Ph.D. graduates are particularly sensitive to this issue as they both represent a measure of the size of physics departments. Indeed, the features share a positive correlation of roughly 0.7. Both approaches present evidence in favor of there being some degree of relationship between the features. For instance, when comparing the minimum and maximum graduates logistic regression models, Table III shows that the difference in the fraction of P-GRE cutoffs for which the size of bachelor’s and Ph.D. programs are significant is similar to the same difference for attending a competitive school or an institution with high research activity. As it is not uncommon for institutions with larger programs to be more competitive or practice higher levels of research, we suspect that some statistical relationship between these features is likely. A more direct example is seen in Fig. 6, where the removal of the Carnegie classification during the feature elimination procedure does not deteriorate the performance by any measurable amount. This indifference suggests that the information contained in the Carnegie classification, which is known to be considerable due to Carnegie’s high importance measure (see Fig. 5), is also contained within the remaining set of features (U-GPA, gender and number of Ph.D. graduates). As a final example, the performance comparison (Fig. 7) shows that most of the overall effect of the institutional influence can be described by a limited selection of institutional features.

C. Data processing and modeling choices

A major difficulty for the logistic regression approach is the need for data processing, especially in the context of losing information by unfortunate modeling choices. The most prominent example in this study is the combination of racial and ethnic groups into a single underrepresented minority group. As suggested by Fig. 2, the lack of race features being important in the logistic regression model may actually be a case of Simpson’s paradox (information loss due to combining data [52]). That is, because the

combined P-GRE distribution of URM applicants resembles the P-GRE distribution of white applicants (see Fig. 2), and because the race feature was one-hot encoded using “White” as reference level, the difference between the distributions is not large enough to be statistically significant. In comparison, the distribution is much more skewed for Asian applicants, and thus the difference becomes statistically significant for higher cutoffs. Other examples include the Carnegie classification and undergraduate population profile, which were essentially reduced from multilevel categorizations to simple binaries. Estimating the amount of meaningful information lost for these features is particularly complicated because of the high number of low-frequency categories.

Compared to the logistic regression approach, the CIF avoids the data processing issues described above. When processing categorical features for inferential modeling, the features must remain interpretable. However, because the CIF does not require the combination of categorical levels to be meaningful, a tree node can find the optimal grouping of categories without regard to interpretation. Indeed, the construction of the CIF algorithm allows it to naturally handle unprocessed data without suffering the same issues as logistic regression (and other machine learning methods that require preprocessing the data). As a result, the CIF is able to identify statistical properties much more easily than logistic regression. An example of this effect is seen in Barron’s selectivity index: Whereas the odds ratios decrease and become less significant as the P-GRE cutoff increases (Table IV), the feature importance is relatively stable with respect to changes in the cutoff (Fig. 5).

Furthermore, compared to the odds ratios of logistic regression, the importance measures of the CIF are more effective and provide a clearer picture. The framework of logistic regression assumes that every feature is a distinct component of the response [Eq. (2)]. In contrast, a tree in the CIF will only include a feature if its found to be important enough (see Sec. III B 2). Hence, if a particular feature is always less important than the other features in every tree (recall each tree is built on a subset of the features), then its importance measure will be 0. A similar mechanism is not present in the logistic regression framework, which will always try to interpret every feature as an integral component of the model. Accordingly, the importance measures more accurately reflect the degree to which the features are associated with the response. Indeed, note that the set of features essential to the model is always larger in the logistic regression models, and in addition, changes as the P-GRE cutoff increases. For example, the odds ratio for attending a privately funded institution is only statistically significant for cutoff scores ≥ 740 (Fig. 3). By relaxing the necessary assumptions of the logistic regression framework, we get a more effective tool for identifying the relationship between the features and the response, albeit one that is harder to interpret.

The effects of unfortunate modeling choices in logistic regression models depends, in the end, on the data. In our case, the combining of racial and ethnic minorities in an underrepresented minority category has likely influenced how racial and ethnic information is treated (in the computational sense) in the model. Similarly, the significance of other processed features may also have been diminished. That being said, we have conducted two very different analyses (inferential vs predictive modeling) and found similar results. It is therefore unlikely that the choices unique to each approach have affected the overall results of the analysis.

D. Future work

The present study has looked into how the undergraduate institutions of applicants may influence the physics graduate admissions process by studying its statistical relationship with P-GRE cutoff scores. Lacking from this analysis is an understanding of whether institutional influence may exert its primary effect at a different stage in the admissions process. For example, it is known that a number of bachelor students that are interested in further studies eventually decide not to apply [16]. While some cases arise due to personal or financial concerns, some students may not have received the preparation or encouragement necessary for motivating further studies. If such motivation plays a significant role for students unsure of whether to pursue a career in physics, then one would expect that prospective students from institutions with Ph.D. programs would be more likely to apply to graduate programs. Additionally, it is worth considering whether these prospective students are more likely to apply to any graduate program in general, or simply the program at their undergraduate institution.

VI. CONCLUSION

The present work has studied the effects of institutional influence on graduate program admissions by modeling a hard physics GRE cutoff score with application data from five Midwestern universities. For completeness, all possible cutoff scores between 620 and 800 (32nd and 67th percentile) have been analyzed, although most admissions employ a cutoff of 700. The analysis has been conducted using both inferential and predictive modeling based on logistic regression and the conditional inference forest algorithm respectively. Both approaches identify the known effects of undergraduate GPA and gender, but do not emphasize a statistical difference between applicants from different racial and ethnic minorities as expected from earlier work [2]. However, this apparent contradiction with past work can likely be understood as a combination of a Simpson’s paradox and selection bias among the applicants. Both approaches identified cases where the impact of institutional features were comparable to the known effects of undergraduate GPA and gender. Overall, the two

approaches agree on the analysis as a whole, but disagree on the result of increasing the P-GRE cutoff. In terms of the odds ratios, increasing the cutoff places more significance on institutional features associated with competitive schools, private funding, large physics programs and high research activity. On the other hand, the added performance when including institutional features can be attributed to a small number of features.

In conclusion, when analyzing graduate program applications we recommend including information regarding the applicants' bachelor institutions. Moreover, due to the innate flexibility and precision of the conditional inference forest algorithm, combined with the large variety of data structures seen in application data, we also recommend the forest algorithm as well as the predictive analysis approach

in general. Based on these findings and its known problems of limiting underrepresented racial and ethnic minorities, we advocate against the practice of using GRE cutoff scores in admissions.

ACKNOWLEDGMENTS

This project was supported by the Michigan State University College of Natural Sciences, the Lappan-Phillips Foundation, and the Norwegian Agency for Quality Assurance in Education (NOKUT), which supports the Center for Computing in Science Education. This project has also received support from the INTPART project of the Research Council of Norway (Grant No. 288125) and the Thon Foundation.

- [1] C. Miller and K. Stassun, A test that fails, *Nature (London)* **510**, 303 (2014).
- [2] C. W. Miller, B. M. Zwickl, J. R. Posselt, R. T. Silvestrini, and T. Hodapp, Typical physics Ph.D. admissions criteria limit access to underrepresented groups but fail to predict doctoral completion, *Sci. Adv.* **5**, eaat7550 (2019).
- [3] G. Potvin, D. Chari, and T. Hodapp, Investigating approaches to diversity in a national survey of physics doctoral degree programs: The graduate admissions landscape, *Phys. Rev. Phys. Educ. Res.* **13**, 020142 (2017).
- [4] A. M. Porter and R. Ivie, Women in Physics and Astronomy, 2019, Tech. Rep. (American Institute of Physics, New York, 2019).
- [5] L. Merner and J. Tyler, African American, Hispanic, and Native American Women among Bachelors in Physical Sciences & Engineering, Tech. Rep. (AIP Publishing, Melville, NY, 2017).
- [6] N. T. Young and M. D. Caballero, The Physics GRE does not help applicants "stand out", [arXiv:2008.10712](https://arxiv.org/abs/2008.10712).
- [7] J. W. Halley, A. Adjoudani, P. Heller, and J. S. Terwilliger, The Graduate Record Examination as an indicator of learning of the curriculum taught to physics majors in US institutions, *Am. J. Phys.* **59**, 403 (1991).
- [8] N. T. Young and M. D. Caballero, Using machine learning to understand physics graduate school admissions, [arXiv:1907.01570](https://arxiv.org/abs/1907.01570).
- [9] C. Zabriskie, J. Yang, S. DeVore, and J. Stewart, Using machine learning to predict physics course outcomes, *Phys. Rev. Phys. Educ. Res.* **15**, 020120 (2019).
- [10] R. Ivie, *Beyond Representation Data to Improve the Situation of Women and Minorities in Physics and Astronomy* (AIP Publishing, Melville, NY, 2018).
- [11] L. M. Aycock, Z. Hazari, E. Brewster, K. B. Clancy, T. Hodapp, and R. M. Goertzen, Sexual harassment reported by undergraduate female physicists, *Phys. Rev. Phys. Educ. Res.* **15**, 010121 (2019).
- [12] K. Rosa and F. M. Mensah, Educational pathways of Black women physicists: Stories of experiencing and overcoming obstacles in life, *Phys. Rev. Phys. Educ. Res.* **12**, 020113 (2016).
- [13] S. Hyater-Adams, C. Fracchiolla, N. Finkelstein, and K. Hinko, Critical look at physics identity: An operationalized framework for examining race and physics identity, *Phys. Rev. Phys. Educ. Res.* **14**, 010132 (2018).
- [14] J. R. Posselt, *Inside Graduate Admissions* (Harvard University Press, Cambridge, MA, 2016).
- [15] Educational Testing Service, *Guide to the Use of Scores* (Educational Testing Service, Princeton, NJ, 2019).
- [16] G. L. Cochran, T. Hodapp, and E. E. A. Brown, Identifying barriers to ethnic/racial minority students' participation in graduate physics, in *Proceedings of the 2018 Physics Education Research Conference, Cincinnati, OH* (AIP, New York, 2018), pp. 92–95, <https://doi.org/10.1119/perc.2017.pr.018>.
- [17] E. M. Levesque, R. Bezanson, and G. R. Tremblay, Physics GRE scores of prize postdoctoral fellows in astronomy, [arXiv:1512.03709](https://arxiv.org/abs/1512.03709).
- [18] National Science Foundation, *Frequently Asked Questions (FAQs) for NSF 20-587, Applicants for FY2021 Graduate Research Fellowship Program (GRFP)* (National Science Foundation, Washington, DC, 2020).
- [19] G. Attiyeh and R. Attiyeh, Testing for bias in graduate school admissions, *J. Health Hum. Resour.* **32**, 524 (1997).
- [20] J. R. Posselt, T. E. Hernandez, G. L. Cochran, and C. W. Miller, Metrics first, diversity later? Making the short list and getting admitted to physics PhD programs, *J. Women Minorities Sci. Engin.* **25**, 283 (2019).
- [21] P. J. Mulvey and S. Nicholson, Physics Bachelor's Degrees, Tech. Rep. (American Institute of Physics, New York, 2015).
- [22] Educational Testing Service, GRE® Subject Test Interpretative Data (2019).

- [23] National Research Council, *A Data-Based Assessment of Research-Doctorate Programs in the United States (with CD)* (The National Academies Press, Washington, DC, 2011).
- [24] Center for Postsecondary Research, *The Carnegie Classification of Institutions of Higher Education* (Indiana University Bloomington, Bloomington, IN, 2016).
- [25] Barron's Educational Series, inc. College Division, *Barron's Profiles of American Colleges* (B.E.S. Publishing, New York, 2014).
- [26] S. Nicholson and P. J. Mulvey, *Roster of Physics Departments with Enrollment and Degree Data, 2017*, Tech. Rep. (American Institute of Physics, New York, 2018).
- [27] S. Nicholson and P. J. Mulvey, *Roster of Physics Departments with Enrollment and Degree Data, 2018*, Tech. Rep. (American Institute of Physics, New York, 2019).
- [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevPhysEducRes.17.010109> for a brief overview of binary response modelling using logistic regression and predictive machine learning analysis with conditional inference forests.
- [29] A. L. Traxler, X. C. Cid, J. Blue, and R. Barthelemy, Enriching gender in physics education research: A binary past and a complex future, *Phys. Rev. Phys. Educ. Res.* **12**, 020114 (2016).
- [30] U.S. Department of Education, *Lists of Postsecondary Minority Institutions* (U.S. Department of Education, Washington, DC, 2020), <https://www2.ed.gov/about/offices/list/ocr/edlite-minorityinst.html>.
- [31] D. W. Hosmer and S. Lemeshow, *Applied Logistic Regression*, 2nd ed. (John Wiley & Sons, Inc., New York, 2000).
- [32] R. Teranishi, Race, ethnicity, and higher education policy: The use of critical quantitative research, *New Dir. Institutional Res.* **2007**, 37 (2007).
- [33] J. Friedman, T. Hastie, and R. Tibshirani, Regularization paths for generalized linear models via coordinate descent, *J. Stat. Softw.* **33**, No. 1 (2010).
- [34] J. Nissen, R. Donatello, and B. Van Dusen, Missing data and bias in physics education research: A case for using multiple imputation, *Phys. Rev. Phys. Educ. Res.* **15**, 020106 (2019).
- [35] S. van Buuren and C. Groothuis-Oudshoorn, mice: Multi-variate imputation by chained equations in R, *J. Stat. Softw.* **45**, No. 3, 1 (2011).
- [36] D. B. Rubin, *Multiple Imputation for Nonresponse in Surveys* (John Wiley & Sons, Inc., New York, 1987).
- [37] P. T. von Hippel, How to impute interactions, squares, and other transformed variables, *Sociol. Methodol.* **39**, 265 (2009).
- [38] K. G. M. Moons, R. A. R. T. Donders, T. Stijnen, and F. E. Harrell, Using the outcome for imputation of missing predictor values was preferred, *J. Chronic Dis.* **59**, 1092 (2006).
- [39] C. X. Ling, J. Huang, and H. Zhang, in *Advances in Artificial Intelligence*, Lecture Notes in Computer Science, edited by Y. Xiang and B. Chaib-draa (Springer, Berlin, Heidelberg, 2003), pp. 329–341.
- [40] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. (Springer, New York, 2017).
- [41] L. Breiman, Random forests, *Mach. Learn.* **45**, 5 (2001).
- [42] T. Hothorn, K. Hornik, and A. Zeileis, Unbiased recursive partitioning: A conditional inference framework, *J. Computational Graphical Stat.* **15**, 651 (2006).
- [43] T. Hothorn, P. Bühlmann, S. Dudoit, A. Molinaro, and M. Laan, Survival ensembles, *Biostatistics (Oxford, England)* **7**, 355 (2006).
- [44] C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn, Bias in random forest variable importance measures: Illustrations, sources and a solution, *BMC Bioinf.* **8**, 25 (2007).
- [45] C. Strobl, A.-L. Boulesteix, T. Kneib, T. Augustin, and A. Zeileis, Conditional variable importance for random forests, *BMC Bioinf.* **9**, 307 (2008).
- [46] V. Svetnik, A. Liaw, C. Tong, J. C. Culberson, R. P. Sheridan, and B. P. Feuston, Random forest: A classification and regression tool for compound classification and QSAR modeling, *J. Chem. Inf. Comput. Sci.* **43**, 1947 (2003).
- [47] S. Janitzka, C. Strobl, and A.-L. Boulesteix, An AUC-based permutation variable importance measure for random forests, *BMC Bioinf.* **14**, 119 (2013).
- [48] L. Auret and C. Aldrich, Interpretation of nonlinear relationships between process variables by use of random forests, *Miner. Eng.* **35**, 27 (2012).
- [49] V. Svetnik, A. Liaw, C. Tong, and T. Wang, in *Multiple Classifier Systems*, Lecture Notes in Computer Science, edited by F. Roli, J. Kittler, and T. Windeatt (Springer, Berlin, Heidelberg, 2004), pp. 334–343.
- [50] D. Chari and G. Potvin, Understanding the importance of graduate admissions criteria according to prospective graduate students, *Phys. Rev. Phys. Educ. Res.* **15**, 023101 (2019).
- [51] S. Kanim and X. Cid, Demographics of physics education research, *Phys. Rev. Phys. Educ. Res.* **16**, 020106 (2020).
- [52] E. H. Simpson, The interpretation of interaction in contingency tables, *J. R. Stat. Soc. Ser. B* **13**, 238 (1951).