Optimizing Financial Aid Allocation to Improve Access and Affordability to Higher Education

Vinhthuy Phan
The University of Memphis
vphan@memphis.edu

Laura Wright
lmstockbridge@gmail.com

Bridgette Decent
The University of Memphis
bdecent@memphis.edu

The allocation of merit-based awards and need-based aid is important to both universities and students who wish to attend the universities. Current approaches tend to consider only institution-centric objectives (e.g. enrollment, revenue) and neglect student-centric objectives in their formulations of the problem. There is lack of consideration to the need to improve access and affordability to higher education. Previously, we contributed a metaheuristic and machine learning approach for optimizing strategies that allocate merit-based awards and need-based aid. The approach can be used to optimize both institution-centric (e.g. enrollment and revenue) and student-centric objectives (affordability and accessibility to higher education). We now employed an improved version of this approach to explore comprehensively a recent admission dataset from our university. We showed that current applicants depended very much on financial sources other than federal and institution aid to attend the university. This potentially created a financial burden for many of these applicants. We identified seven budget-friendly strategies that promise to increase access to higher education significantly by more than 100%, while still keeping it affordable for students and limiting a budget increase to less than 7%. Additionally, we identified a total of 111 strategies, including those that benefit from more aggressive changes in the budget to obtain higher increases in enrollment, revenue, and/or higher affordability and accessibility for students. This method may be used by other institutions in ways that best fit their institutional objectives and students' profiles.

Keywords: institutional data analytic, financial aid allocation, financial aid optimization, scholarship distribution, enrollment management

1. Introduction

Public universities in the United States (US) create recruitment policies amidst decreased state funding and increased costs of attendance. Decreased funding was predicted as a result of state deficits (Hovey, 1999). Aside from budget deficits, political factors and governmental structures play an important role in decreased support and funding for public higher education (McLendon et al., 2009; Tandberg, 2010). This is now evident in data showing falling state appropriations as a percent of total revenue for public 4-year institutions from academic year 2008-09 (23.8%) to academic year 2018-19 (16.5%), according to the National Center for Education Statistics (2021). Decreased state funding necessitates a greater reliance on tuition and fee revenue from enrolled students.

From academic year 2011-12 through 2016-17 tuition and fees increased by 9% at public 4-year institutions, making college too costly for an increasing proportion of students (especially those from low- and lower-middle-income families) (Assalone et al., 2018; Delaney, 2014; Perna and Li, 2006). Students are being priced out of college or choosing to take out loans and/or work while pursuing a course of study (Perna et al., 2007). This rising cost has resulted in a national college debt crisis in which 57% of debtors with federal student loan debt owe up to $20,000 (Assalone et al., 2018).

The combination of decreased funding and increased costs of attendance have implications for enrollment management activities. An increased reliance on tuition and fee revenue means that setting policies surrounding student recruitment efforts and accurately forecasting their outcomes is increasingly important to universities' income and budget projections (Antons and Maltz, 2006; Trusheim and Rylee, 2011). However, these recruitment efforts are hindered by the increased prices charged to students. Higher prices reduce enrollment rates particularly for low-income students in public institutions (Leslie and Brinkman, 1987). The use of financial aid as an enrollment management tool to counteract these increased prices is well-established (Hossler, 2000). Further, the positive impact of both federal and state need-based financial aid programs on students from low-income families has also been explored (Assalone et al., 2018; Davis et al., 2013; Delaney, 2014). With increasing costs of attendance leading to crises of affordability and access, universities should explore the role of need-based aid in their enrollment management strategies.

The call to optimize enrollment management activities with the consideration of access and affordability is one that appears difficult to achieve. However, the student and financial data needed for such predictions and optimizations are stored in universities' existing data systems (Shacklock, 2016). Despite having access to this data, universities are not quick to leverage it themselves perhaps due to untrained personnel or poor data infrastructure (Shacklock, 2016; Siraj and Abdoulha, 2009). Instead, the data and these analyses are outsourced to external consulting firms who organize and process the data to be used in proprietary models (Hossler, 2009). With this type of industry environment there is limited published work to demonstrate how machine learning prediction and optimization can be used in tandem to inform university recruitment policies.

Existing work on financial aid optimization has been limited in two regards: (1) the optimization only takes enrollment as the sole objective and (2) the resulting strategy is for merit-based aid only (Aulck et al., 2020; Sarafraz et al., 2015). The limiting of the optimization to enrollment is short-sighted in that universities are targeting multiple objectives that may not move with enrollment. Additionally, when only considering merit-based awards, student-centric outcomes like access and affordability are neglected despite their increasing importance in the dialogue surrounding higher education.

Previously, Phan et al. (Phan et al., 2022) introduced an approach to address some of these shortcomings. This approach consists of an interactive feedback loop between two core components: a gradient-boosting classifier and *auto-start* local search optimizer. First, given a current strategy for award and aid allocation (together with many other features), the classifier is used to predict enrollment, revenue, the affordability and accessibility of higher education, and other outcomes. These predicted outcomes are then used by an optimizer to revise the current strategy and suggest a possibly better one. The interaction between the classifier and optimizer forms a feedback loop that is the basis of the optimization. This loop aims to identify optimal strategies for allocating awards and aid with specified expected outcomes. The novelty of this work includes (1) the allocation both merit-based awards and need-based aid in the optimization process; and (2) the inclusion of the affordability (i.e. how much a student needs to borrow, on average, to attend the university) and accessibility (i.e. the likelihood that a student attends the university, given certain amounts of merit-based awards and need-based aid) of higher education in the optimization of allocation strategies.

In this work, we made some improvements upon the previously introduced method (Phan et al., 2022) to increase its applicability to a wider range of admission data. We employed this method to explore comprehensively a recent admission dataset from our university. We showed that current applicants depended very much on financial sources other than federal and institution aid to attend the university. This potentially created a financial burden for many of these applicants. We identified seven viable strategies that promise to increase accessibility greatly while requiring less than 7% increase in the financial aid budget, while keeping all most objectives mostly intact. Several other aggressive strategies, which require additional spending, but yield higher outcomes, were also identified and discussed. Lastly, we provided a comprehensive analysis of how these strategies impact the redistribution of merit-based awards and need-based aid.

2. Related Work

The following describes research on enrollment prediction, scholarship allocation, and financial aid optimization. There is significant work on predicting enrollment, but far less with regard to optimization and allocation of scholarships. Most authors looked at only one of these issues and most work has been primarily concerned with merit-based aid alone.

Predicting Enrollment. Universities' increased reliance on tuition and fee revenue has increased pressure on their limited resources to accurately predict and increase enrollment (i.e. yield). Recruitment counselors need to allocate their limited time to contact students "on-the-fence" about their enrollment decisions (DesJardins, 2002). This business need led to the literature on enrollment prediction at the individual applicant level using data mining and machine learning techniques.

There is a host of research that relies on logistic regression techniques to create these individual enrollment probabilities. Basu et al. compared various machine learning techniques and found that logistic regression performed better than most (Basu et al., 2019). Sugrue used logistic regression to predict overall enrollment yield (Sugrue, 2014). DesJardins used a similar approach to reveal the enrollment probabilities of applicants before any financial aid offers had been made. The resulting model, which relied heavily on demographic features, had a correct classification rate of approximately 67% (DesJardins, 2002). In 2006, Goenner and Pauls used a logistic regression to predict the enrollment of students who had inquired about the university. The resulting model was 89% accurate at out-of-sample predictions (Goenner and Pauls, 2006).

In another vein of literature, other prediction techniques are explored. Two related works explore the effectiveness of neural network models in predicting individual enrollment probabilities for applicants. A 1998 paper by Walczak revealed that a backpropogation neural network model could result in a 56% decrease in recruitment counselor caseloads due to its accuracy

(Walczak, 1998). In a follow-up study, Walczak and Sincich explicitly compare the performance of neural network models to that of logistic regression and find that neural networks produce better results (Walczak and Sincich, 1999). Chang similarly explores the efficacy of various prediction techniques, specifying models for logistic regression, neural networks, and classification and regression tree (CR&T). The work showed that both neural networks and CR&T outperformed logistic regression when judged based on prediction accuracy (Chang, 2006).

Improving access and diversity in college admission. Recent work by Rebecca Zwick utilized constraint optimization to influence the decision making process in college admission to improve access, and increase diversity and representation of low-income students (Zwick et al., 2019; Zwick, 2020; Zwick et al., 2021). This approach places quotas on different criteria (e.g. a certain percentage of student population from an ethnic group, or from a certain gender) and employs mathematical optimization techniques such as linear programming and integer programming to achieve desirable academic standards and diversity goals. These efforts, however, do not focus on improving strategies to provide financial support to students.

Scholarship disbursement and allocation. Not all scholarship policies employed by universities are rule-based. In these cases, the business process of selecting the applicants to receive an award may be ineffective and inefficient. Literature has shown that data mining techniques can be effective in creating rule-based scholarship allocation policies to help reduce business process inefficiencies. In 2019, Rohman et al. illustrated how an ID3 decision tree algorithm could generate rules to select the scholarship applicants most likely to be awarded a scholarship. This general rule allowed for the efficient identification of applicants so that offers could be made (Rohman et al., 2019). Alhassan and Lawal similarly used a tree-based data mining classification technique to determine a generic rule for scholarship disbursal. They found the technique to be effective and efficient (Alhassan and Lawal, 2015).

Optimization of financial aid policies. Mathematical programming models are effective tools for generating and evaluating financial aid strategies. Spaulding and Olswang use discriminant analysis to test the efficacy of various aid strategies (Spaulding and Olswang, 2005). Linear programming was used by Sugrue et al. in 2006 as an aid decision tool where the goal was to maximize net revenue with budget, average Scholastic Assessment Test (SAT) scores, recruitment pools, and enrollment targets as constraints (Sugrue et al., 2006). In later work, Sugrue again employed a linear programming approach to optimize the quality of the incoming class with estimates of yield rates being derived from previous years' yield rates (Sugrue, 2010).

More recent research has incorporated enrollment prediction models with optimization techniques to recommend financial aid strategies. In 2015 Sarafraz et al. used a neural network model to predict enrollment and then employed a genetic algorithm to find a scholarship strategy that maximized total enrollment (Sarafraz et al., 2015). In 2019 Sugrue used data from the University of Miami to develop a merit-based aid allocation model that predicted enrollment via logistic regression and maximized the quality of the incoming class via a linear programming model (Sugrue, 2019). In 2020 Aulck et al. tested a group of machine learning tools to predict the enrollment decisions of admitted, domestic nonresident first-time students and then used those results in a genetic algorithm to recommend an optimal disbursement strategy for a domestic non-resident merit scholarship that would maximize enrollment (Aulck et al., 2020).

None of these previous works have considered merit-based award and need-based aid strategies simultaneously as our paper proposes to do. This joint consideration demands the consideration of what an appropriate optimization objective should be, given that the recruitment goals of merit-based awards and need-based aid differ in some regards. Though the joint consideration of merit and need strategies in one optimization problem does complicate the process, a multi-objective approach needs to be researched since these aid strategies compete for the same limited budget resources.

Local-search optimization. Local search is a heuristic method for solving computationally hard problems. A local search algorithm starts with an initial candidate solution and iteratively moves to a neighbor solution in hope of finding better and better candidate solutions. The algorithm stops when it cannot find a neighbor solution that is better than the current candidate solution. For local search to work, a *neighbor relation* must be defined so that from an arbitrary candidate solution, a neighbor solution can be generated.

Stochastic hill climbing (Juels and Wattenberg, 1995) is a fast local search method because it greedily moves from one candidate solution to a better one. It is similar to a popular method, stochastic gradient descent (Johnson and Jacobson, 2002), but it is faster because it does not need to estimate the gradient of the objective function. The main disadvantage of stochastic hill climbing is that it is often stuck in locally optimal solutions.

Two popular methods that can find globally optimal solutions are simulated annealing (Kirkpatrick et al., 1983) and genetic algorithms (Whitley, 1994). In a number of applications, genetic algorithms produced slightly better solutions than simulated annealing (Sexton et al., 1999). Nevertheless, simulated annealing and its variant, simulated quenching (Ingber, 1993), seem to suit our technical approach better than genetic algorithm because it is not obvious to us how we can meaningfully apply the genetic algorithm's crossover operator to allocation strategies.

3. METHODS

3.1. OVERVIEW OF THE PROBLEM AND CHALLENGES

Merit-based awards and need-based aid. At the University of Memphis, in Tennessee, USA, a certain amount of financial award or aid is offered to each applicant based on his or her profile on merit (academic performance) and need. Merit-based award eligibility is calculated based on a combination of standardized test scores (i.e. American College Testing (ACT) and Scholastic Assessment Test (SAT)) and high school GPA. Need-based aid eligibility is determined by expected family contribution from the Free Application for Federal Student Aid (FAFSA) and then adjusted in an ad hoc manner based on items such as cost of attendance and how much federal aid and other merit-based awards are promised to students. A student may receive both merit awards and financial aid. The allocation strategy of merit awards and aid strives to be fair in that two students with the same residency and academic performance profile will receive the same merit offer and two students with the same residency and need profile will receive the same aid offer. These offers are made in a guaranteed fashion in that each admitted applicant will be offered a merit-based award and/or need-based aid so long as requirements are met.

Allocation strategy of awards and aid. In the context of this work, an allocation strategy consists of a merit-based allocation strategy and a need-based allocation strategy. Each allocation strategy consists of (1) a list of buckets into which students are placed, and (2) how much financial award or aid students in each bucket will receive.

An allocation strategy has an important impact on the budget of the university because it affects enrollment in complex ways. Enrollment affects revenue. Revenue affects the number and amounts of awards and aid the university can give to its students. And awards and aid directly affect the applicants' decision to enroll at the university. Increasing awards and aid

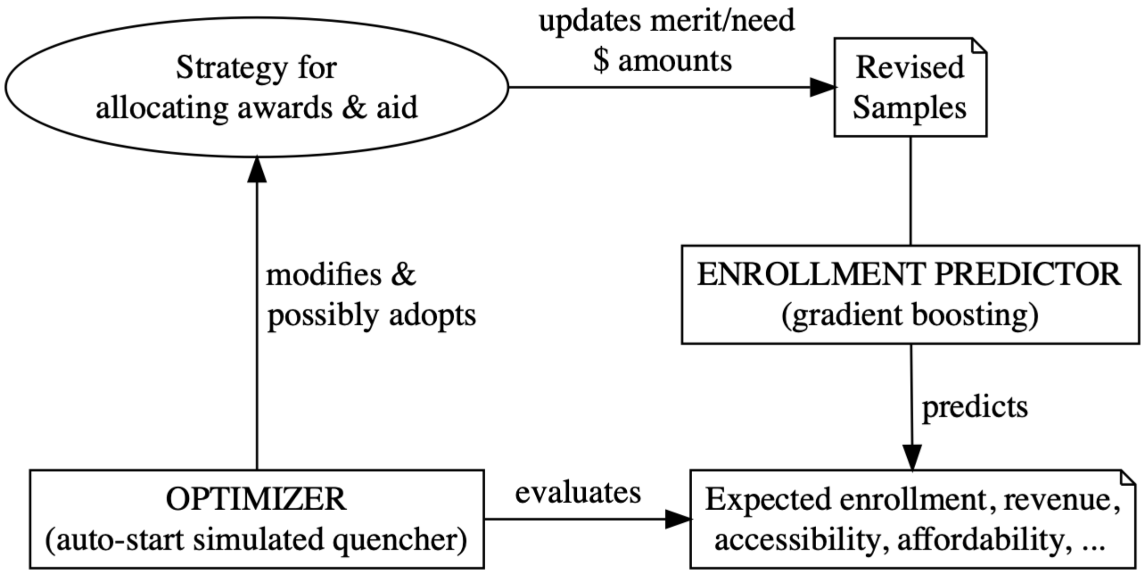

Figure 1: Optimization of allocation strategies.

can increase enrollment, but may reduce revenue, which in turn limits the university's ability to increase awards and aid.

Our mission. Our team was tasked with revising the university's current way of offering merit-based awards and need-based aid. Although the current system is fair, as it should be, the Financial Aid Office has thought that due to multiple reasons the current strategy might not be optimal. We were asked to generate a revised strategy that uniformly impacts the domestic student body in a way that improves multiple, possibly conflicting objectives such as increasing enrollment, increasing revenue, increasing student performance profiles, and making higher education more accessible and affordable.

Technical approach. An overview of our approach is depicted in Figure 1. The approach utilizes a feedback loop that continually revises a feature set from which a classifier learns to predict various outcomes. In contrast to the traditional setting where features stay fixed, the features in our model are continually revised in this feedback loop. This is possible due to the fact that the university can change award and aid offer amounts, which is part of the feature set used to predict enrollment, revenue and other outcomes. The feedback loop is an iterative interaction between the classifier, which predicts expected outcomes, and an optimizer, which evaluates the outcomes and suggests a revised strategy with possibly better outcomes. In each iteration of this interaction, depicted in Figure 1, the current allocation strategy dictates offer amounts for awards and aid. This effectively changes the data, which leads new predicted outcomes (e.g. enrollment, revenue, student affordability and accessibility, etc.). This allows the *auto-start* local-search optimizer to revise and possibly improve the current strategy. This process repeats until the optimizer reaches a near-optimal strategy that does not improve the overall objective.

3.2. Data Collection and Cleaning

Data for this research is all first-time freshmen (FTF) domestic admitted applicants to the university in the Fall 2020. The data on students is compiled from the university's admission application, the FAFSA, and the student information system. Only the students who filed a FAFSA (making them eligible for need-based aid) are included resulting 7,564 observations which is approximately 67% of the total number of domestic FTF students for that term. Data from the admission application includes admission test scores, high school GPA, location/residency, and intended major. Data from the FAFSA includes expected family contribution and family income. Data from the student information system includes whether or not the student had previously enrolled at the institution as a dual enrollment student, aid offer amounts for various categories of aid, and an enrollment indicator. Of the 7,564 students 2,340 chose to enroll (approximately 31%). No demographic variables are used in the study so as to avoid bias in recommending aid strategies to the university.

Financial data points are expected full time tuition, merit-based award offers, need-based aid offers, actual and estimated Pell grant offers, and all other aid offers before loans. Estimated Pell grant offers are included as the university will make offers prior to actual Pell amounts being known. These are estimated based on cost of attendance and expected family contribution using the Pell schedules released each year with the assumption that students will enroll full time. Each of the aid offer features enters the predictive model independently and in a total financial aid variable. These offer amounts are not manipulated when generating the predictive model, but are updated during the optimization process.

3.3. Feature Engineering

Prior to building prediction models and optimizations, feature engineering is done using Pandas (McKinney, 2010). Each 2-value categorical variable with two values (e.g. a gender variable with values Female and Male) is converted to a new binary indicator variable. Each k-value categorical variable ($k > 2$) is converted to k binary variables via a widely-used method known as *one-hot encoding*. An applicant's profile consists of distances to the university and its key market competitors. For each of such distances, two new binary features, *daily-commutable* and *weekly-commutable*, are created. For example, if an applicant is within 45 miles from a university, then the corresponding daily-commutable variable is set to 1. Similarly, if an applicant is within 300 miles from a university, the corresponding weekly-commutable is set to 1. An indicator of familiarity with the university and social networking at the university is calculated based on the ratio of enrollment to admission of students from each applicant's high school from the last five years. The higher this ratio is for an applicant, the more familiar the applicant is with the university and more likely the applicant can connect with peers at the university. We convert the major declared on the application into five major group binary indicators: STEM, Fine Arts, Health, Business, and Humanities. The feature engineering results in 94 features from our data sources on which to build our predictions.

Two important features that affect applicants' decisions to enroll are the promised amounts of merit-based awards and need-based aid. These features are determined by a specific strategy for allocating awards and aid. As indicated in Figure 1, these two features are updated during an optimization process that traverses through the solution space to find an optimal one. As the features are updated, the classification model relearns to update its estimate of enrollment probabilities.

3.4. Modeling enrollment

The data is stored in a matrix D where each row represents an applicant's profile. The first column of D is the binary enrollment variable, y, where y_i being equal to 1 means applicant i enrolled at the university and y_i being equal to 0 means the applicant did not enroll at the university. All other columns of D are features that a classifier uses to learn y. The goal in modeling enrollment is to assess the viability of using popular classification methods to predict enrollment, and to identify and adopt the best method to assist in the optimization of strategies for allocating merit-based awards and need-based aid.

Classifiers. We investigate the performance of several popular classifiers that can predict enrollment probabilities. These include classification methods based on diverse approaches such as support vector machine, logistic regression, and k-nearest neighbors. We also consider the ensemble methods of random forest, AdaBoost (Schapire, 2013), gradient boosting (Friedman, 2001), and a more regularized version of gradient boosting known as extreme gradient boosting (Chen and Guestrin, 2016). Many of these methods do not strictly predict probabilities of a target variable. Rather, they can provide quantities that can loosely be interpreted as probabilities. For example, in case of ensemble methods, which output quantities based on binary decisions of the base learners, we interpret these quantities as probabilities. We used available implementations of these methods in the scikit-learn (Pedregosa et al., 2011) and xgboost (Chen and Guestrin, 2016) libraries.

Data scaling. Classification methods such as support vector machine, logistic regression, and k-nearest neighbors operate on distances between data points. Since distances between different features are not of the same scales, the features need to be standardized first. We employ a popular method of data standardization. For each data point, we subtract from it the mean of the training samples, and then divide that by the standard deviation of the training samples.

Performance metrics. We employ multiple metrics to measure the performance of the classifiers from multiple perspectives. Since the data is imbalanced, a single metric, e.g. accuracy, does not meaningfully reflect different aspects of the performance of a classifier. We consider three compound metrics: F-score, balanced accuracy, and AUC (area under the ROC curve), which show different aspects of performance. F-score is a useful metric when we are interested in the ability to predict class-1 samples (applicants who ultimately enroll at the university). It combines two individual measures and is defined as $\frac{2 \cdot precision \cdot sensitivity}{(precision + sensitivity)}$, where precision is the probability that a positive prediction is correctly predicted, and sensitivity is the true positive rate or the probability that class-1 samples are correctly predicted. Balanced accuracy is used when we are equally interested in the ability to predict both class-1 and class-0 samples. It gives equal weights to true positive rates (sensitivity) and true negative rates (specificity). AUC is useful when we are interested in the trade-offs between sensitivity and specificity at various thresholds as it sums up the area under the curve defined by sensitivity (true positive rate) and 1-specificity (false positive rate).

Comparison to baseline. To gauge the performance of each classification method, we compare its performance to that of two baseline methods. The first baseline is a popular method, ZeroR (Devasena et al., 2011), which always predicts the most frequent label in the training set. The second baseline, Stratified, makes predictions based on the distribution of the labels in the training set.

Cross-validation. To determine the performance of each classifier under a metric, we train the classifier using 90% of the data and test it using the other 10% of the data. We experiment with two popular cross-validation methods of partitioning the data randomly into training and testing sets: (1) repeated random subsampling (up to 50 random splits) and (2) k-fold with k=5 and 10. Although the data partitions are randomly generated, in each iteration, we compare the classifiers using the same random partition.

3.5. Optimization of allocation strategies

Algorithm 1 AutoStart_Simulated_Quencher(s_0 , S_T , δ)

```
1: s ← s_0
2: Run Stochastic Hill Climb to find Δ_max
3: T ← Δ_max / ln 0.5
4: while progress is still being made do
5:    T ← T · δ
6:    for i from 1 to S_T do
7:       t ← random neighbor of s
8:       Δ = objective_value( t ) − objective_value( s )
9:       if t > 0 or with probability e^(Δ/T) then
10:         s ← t
11:         if objective_value( s ) > objective_value( best_strategy ) then
12:            best_strategy ← s
13: return best_strategy
```

Procedure objective_value(S)

1. Update merit & need amounts in applicant data, based on amounts dictated by strategy S.
2. Predict outcomes (e.g. enrollment, ...) using a classifier such as GradientBoosting.
3. $value \leftarrow$ the objective value of S, computed based on newly predicted outcomes.
4. **return** $value$

Local search optimization. Simulated annealing (Kirkpatrick et al., 1983) is inspired from a physical annealing process, in which an initially hot temperature lets the local search explore the solution space more freely, allowing the adoption of neighbor solutions that are not as good as the candidate solution. As the search goes on, the temperature cools down and the search is more aggressive in finding better solutions. When the temperature is low enough, the search essentially becomes a stochastic hill climb. Starting at a high initial temperature, the annealing is known to be able to escape locally optimal solutions and reach a globally optimal solution eventually.

We adopt a modified version of simulated quenching (Ingber, 1993), which is a variant of simulated annealing. It is faster than simulated annealing because it cools temperatures faster. It is shown that in practice simulated quenching was as good as simulated annealing (Ingber, 1993).

While Figure 1 provides a high-level description of each iteration of the optimization, Algorithm 1 provides a more detailed description of how the optimization utilizes simulated quenching to find optimal strategies for allocating merit-based awards and need-based aid. The algorithm has three main inputs: (1) initial allocation strategy s_0, which is an estimate of the university's current formula for distributing merit-based awards and need-based financial aid. Currently, merit-based awards have a few levels with award amounts varying from several hundred to several thousand dollars. Need-based aid amounts are more fluid and also range from several hundred to several thousand dollars; (2) S_T, which is a number that indicates how long the current temperature, T, should be kept constant. Note that T affects the probability of making a *backward move*, i.e. adopting a slightly inferior solution; and (3) δ, which governs how fast the annealing temperature T is reduced after it is kept constant in S_T steps. After S_T steps, T is reduced by a fraction of δ. For our dataset, we experimented and varied δ between 0.90 and 0.95 and S_T between 10 to several hundreds.

The simulated quenching process aims to improve upon the initial strategy, s_0, by moving from one neighboring strategy to the next. A better neighboring strategy replaces the current one. Additionally, a slightly worse neighboring strategy can also replace the current one with a certain probability (lines 9-10, Algorithm 1), which is determined by the current temperature T and how much worse the neighbor is. At the beginning of the process, T is high and the search is more exploratory in making backward moves. Near the end of the process, with T being reduced, the search is more conservative and only adopts strategies that are better than current solutions. The essence of simulated quenching is that this initially exploratory and increasingly focused search will ultimately reach optimal solutions.

Auto-start simulated quenching To explore the search space liberally at the start of the search, the initial temperature needs to be sufficiently high. This initial temperature is, however, problem dependent. As such, users generally have to experiment with different values to find an appropriate one. We employ a stochastic hill climber to derive an initial temperature to start the simulated quenching process without requiring users to initiate the search by specifying an initial temperature (lines 2-3 in Algorithm 1). We set the initial temperature T to $\frac{\Delta_{\max}}{\ln(0.5)}$, where $\Delta_{\max}$ is the difference between the worst strategy and the best strategy (local optimum) that the stochastic hill climber finds. At this initial temperature, the largest backward move is accepted with a probability of 0.5. As T decreases, this probability decreases and a backward move is less likely allowed. When T is small enough, the search effectively turns into a stochastic hill climber.

Solution representation. The local search algorithm navigates through the search space of allocation strategies to find an optimal solution (strategy). An allocation strategy consists of a merit-based award allocation strategy and a need-based aid allocation strategy. Each consists of (1) a list of buckets into which students are placed and (2) how much money students in each bucket will receive. Buckets and amounts must be increasing. This means a higher achievement results in higher merit-based amounts, and higher need results in higher need-based amounts. As an example, consider this allocation strategy:

- Achievement buckets: [0, 0.25, 0.50, 0.75, 1]
- Merit amounts: [$0, $2000, $6000, $10000]
- Need buckets: [$0, $500, $5000, $10000, $20000, $40000]
- Need amounts: [$0, $1000, $1500, $2500, $3000]

If an applicant's achievement index (or need index) falls between bucket[j] and bucket[$j+1$], then the applicant receives amounts[j]. For example, if applicant i has a merit index of 0.4 and a need index of $15000, then the applicant will receive a merit-based amount of $M_i =$ $2000 and a need-based amount of $N_i = \$2500$. Collectively, the offered amounts of merit-based awards and need-based aid make up two features M and N, respectively, in the data.

Construction of neighboring strategies. Local search finds good strategies by navigating through neighboring strategies (line 7, Algorithm 1). The main steps of generating a neighboring strategy t for a given strategy s are as follows:

1: First, select with probability 0.5 either the *merit amounts* list or the *need amounts* list of s. Call this list $A = [a_1, \cdots, a_k]$.

2: Second, with probability 0.5, either add a small amount to a random amount a_r $(1 \leq r \leq k)$ in A, or remove a small amount from a random amount a_r in A.

These steps of adding or removing small amounts from A are repeated until we find a neighboring strategy that satisfies the following constraints: (1) A remains in an increasing order ; (2) the resulting amount, a_r, must be sufficiently different from its adjacent entries (a_{r-1} and a_{r+1}); and (3) the minimum amount in A cannot be too small (in case of a remove) or the maximum amount in A cannot be too big (in case of an add).

Evaluation and adoption of new strategies. Strategies are compared and evaluated based on their objective values. The objective value of a strategy is computed using expected outcomes that are predicted by a classifier. The classifier learns from the existing data, which in our case is admission data from Fall 2020, and predicts outcomes of a strategy using existing applicant data with revised merit and need amounts, which are dictated by the strategy. The calculation of a strategy's objective value is described in procedure $objecitve\_value$ in Algorithm 1.

Calculation of expected outcomes. An allocation of funds to awards and aid affects enrollment, which affects multiple outcomes that the university is interested in. We utilized a classification method, e.g. gradient boosting, to predict enrollment probabilities of applicants. This enables us to compute the expected outcomes that we are interested in. Let $y = (y_1, \cdots, y_n)$ be the binary target variable enrollment, and let $p = (p_1, \cdots, p_n)$, where p_i is the probability that applicant i enrolls at the university.

The expected outcomes we are interested in include:

- Expected enrollment $= \sum_i p_i$. Denote the expected enrollment as E.
- Expected net revenue $= \sum_i p_i \cdot R_i$, where R_i is net revenue obtained from applicant i. $R_i = t_i - (m_i + n_i)$, where t_i is the tuition that the applicant pays, and m_i and n_i are the merit-based and need-based amounts offered to this applicant.
- Expected unmet need $= \frac{1}{E} \cdot \sum_i p_i \cdot (COA_i - EFC_i - Pell_i - M_i - N_i)$, where E is the expected enrollment and is calculated as above; COA, EFC, and Pell are cost of attendance, expected family contribution, federal Pell grant; and M and N are the university's merit-based award and need-based aid. A positive unmet need amount is what the applicant is expected to borrow to pay for attending the university. This definition is widely used as a measure affordability.
- Expected accessibility $= \sum_i p_i \cdot I_i$, where $I_i = 1$ if applicant i is offered some financial aid (i.e. $M_i + N_i > 0$) and has unmet need that exceeds a certain threshold; otherwise, $I_i = 0$. Note that M_i and N_i are the amounts of merit-based award and need-based aid that applicant i is promised. This threshold is set by the university and is believed to represent a level of need that is not surmountable through existing state and federal programs.

- Expected return on investment $= \frac{\sum_i p_i \cdot R_i}{\sum_i p_i \cdot (M_i + N_i)}$. It is the expected revenue divided by the expected total promised amounts of awards and aid.
- Expected achievement $= \frac{1}{E} \cdot \sum_i p_i \cdot A_i$, where E is the expected enrollment; A_i is the achievement index of applicant i and is calculated from a combination of the applicant's standardized test scores (e.g. SAT or ACT) and high school GPA. Applicant i is offered a merit-based award in an amount of M_i based on A_i.

Multi-objective optimization. Algorithm 1 aims to find a strategy s^* that maximizes the weighted sum of expected enrollment (enr), net revenue (rev), unmet need (un), accessibility (acc), return of investment (roi), and achievement (ach) as follows:

$$f(s, D, C) = \alpha_1 \cdot E[\text{enr}] + \alpha_2 \cdot E[\text{rev}] - \alpha_3 \cdot E[\text{un}] + \alpha_4 \cdot E[\text{acc}] + \alpha_5 \cdot E[\text{roi}] + \alpha_6 \cdot E[\text{ach}] \quad (1)$$

where α_i's are the weights of the expected outcomes. Implicitly, the allocation strategy s is applied to the data D, from which the classifier C learns to predict the enrollment probability p, which is used to compute the expected values on the right hand side of the equation.

Although we do not expect the algorithm to find a strategy that is optimal in each individual objective, an overall optimal value of the function should benefit both the university and the applicants. While higher expected values of enrollment, revenue, and achievement benefit the university, less unmet need and high accessibility benefit applicants. On the one hand, it seems that reducing unmet need for applicants may reduce revenue. On the other hand, making attendance more affordable may actually increase both the expected enrollment, which in turn may increase the expected revenue. In other words, increasing merit-based and need-based amounts may reduce the revenue from each enrolled applicant, but may increase the number of enrolled applicants and, ultimately, the overall revenue.

Users can experiment with α's to give different weights to different objectives to obtain realistically acceptable trade-offs and improvements among the individual objectives.

Optimization Constraints. In constructing this research, certain constraints were determined by the university's administration. These constraints reflect the administration's perspective on the competitive environment and internal business processes that will support any proposed strategy. These constraints depend on the award structure under consideration. For merit-based awards, there could be no more than six awards with award amounts restricted to a specific range. The need-based aid strategy specification was limited to four awards with their own maximum and minimum constraints on aid amounts. It is worth noting that the maximum value for a merit-based award was four times higher than the maximum need-based aid amount. Need-based aid eligibility also had a minimum need index cutoff ($COA_i - EFC_i - M_i - Pell_i$). For both allocation strategies there needed to be a minimum difference of \$100 between award buckets.

3.6. Selecting Classifiers for Optimization

The optimization process (Figure 1) requires a classifier to predict enrollment and various expected outcomes. We examined various popular classification approaches that can make probabilistic predictions of enrollment events. We found gradient boosting (Schapire, 2013) was the highest performing classifier and had very high performance in predicting enrollment. This

Table 1: Performance of predicting enrollment.

	Balanced accuracy	**F-score**	**AUC**
Gradient Boosting	**0.91**	**0.88**	**0.96**
XGB	**0.91**	**0.88**	**0.96**
AdaBoost	0.88	0.84	0.95
Random Forest	0.77	0.69	0.92
Linear SVC	0.72	0.61	0.83
KNN	0.61	0.44	0.69
Stratified baseline	0.50	0.30	0.50
ZeroR baseline	0.50	-	0.50

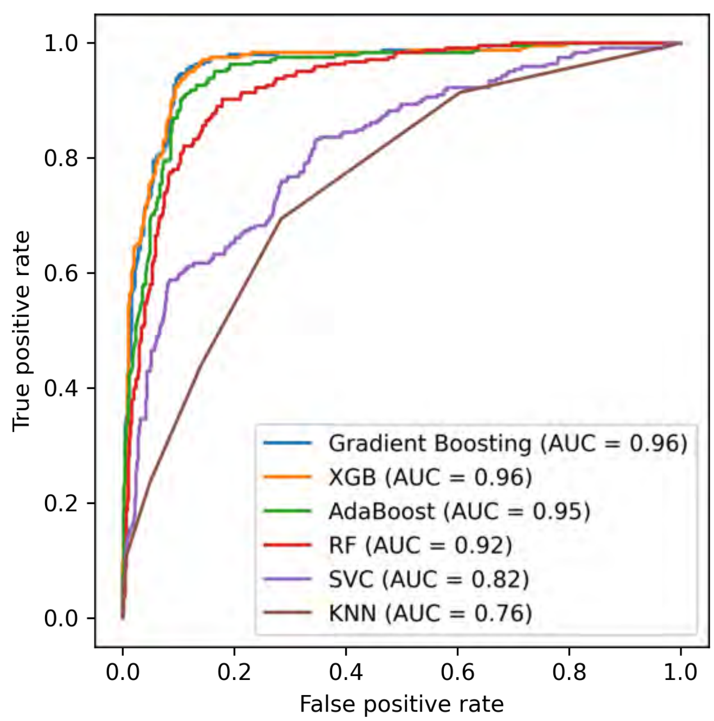

Figure 2: ROC curve of enrollment prediction.

result was obtained by comparing optimized versions of six different popular classification approaches, which could make probabilistic predictions. Figure 2 shows the ROC curve of the classifiers in one random partition of the data into 90% training and 10% testing sets. The figure shows an excellent trade-off between true positive rates and false positive rates for the two top performers (gradient boosting and extreme gradient boosting).

Table 1 shows the performance of the classifiers averaged across 10 folds of cross validation. In validating the classifiers, we experimented with k-fold cross validation and repeated subsampling at various parameters. We ultimately decided that a 10-fold cross validation was a slightly better choice for our study than the others. All classifiers performed significantly better than the two baselines. Note that given the imbalance of the data, ZeroR did not predict any positive label, resulting in undefined precision and F-score. Gradient Boosting (Schapire, 2013) and Extreme Gradient Boosting (Chen and Guestrin, 2016) had the same highest performance across all three metrics (balanced accuracy, F-score, and AUC). gradient boosting, AdaBoost, and random forest were optimized after considering various values of maximum depth and minimum leaf size of the decision-tree base learners.

4. RESULTS

4.1. EXPERIMENTAL SETUP

Starting from an initial set of applicants in the fall semester of 2020, and based on the university's allocation strategy at the time, we employed gradient boosting to predict enrollment events and expected outcomes of applicants for the initial strategy. In each run, the optimizer revises the initial strategy, effectively alters the samples, invokes the classifier to make new predictions of expected outcomes, and sets the optimization process in motion, as depicted in Figure 1 and described in Algorithm 1. We experimented with all combinations of institution-centric and student-centric objectives. We report the results that are gotten from optimizing three objectives: affordability, enrollment, and revenue, with weights of 0.6, 1, and 2. These weights were obtained after an extensive experimentation. They were weighted on features that are not normalized; as such, their values do not necessarily mean that certain features are considered more heavily than the others.

In optimizing for affordability, enrollment, and revenue, we obtained a total of 235 strategies for allocating merit-based awards and need-based aid. As described in Methods, Section 3.5, each strategy consists of 4 things: (i) **achievement buckets** into which a student's achievement profile is binned, (ii) **merit amounts**, which specify how much is given to the student based on merit, (iii) **need buckets**, into which a student's need is binned, and (iv) **need amounts**, which specify how much is given to the student based on need.

To be mindful of the university's concerns, we do not report specific values of these four items for the strategies we obtained. Instead, we report changes between expected outcomes derived from these strategies and the actual outcomes based on the current strategy.

4.2. EVIDENCE OF FINANCIAL BURDEN OF COLLEGE ATTENDANCE

Figure 3 shows the most important features (out of a total of 94 features), as quantified by gradient boosting, which is our chosen model. Although we reported actual numerical amounts of importance that are calculated by the model, these amounts should not be taken in absolute numerical values. In our experience, fluctuations in model parameters and data will result in

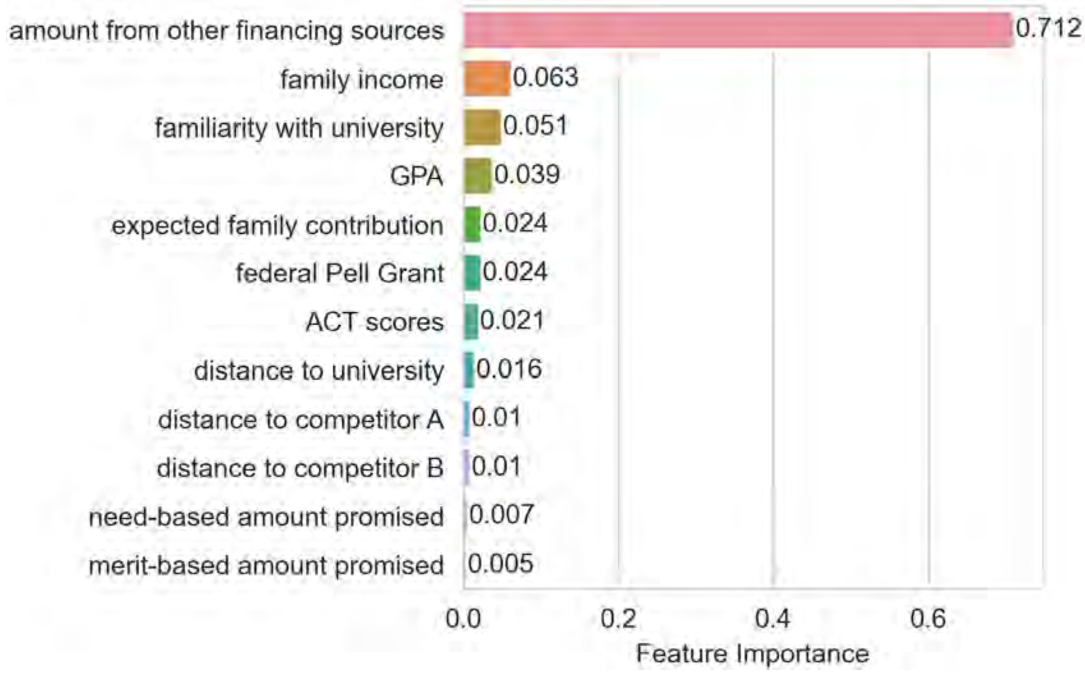

Figure 3: Top most important features derived from a gradient boosting model.

small changes in the actual values of feature importance. With that said, the model's estimate of feature importance captures to a large extent the actual importance of these features.

As seen in Figure 3, the modeling shows that the most important factor in terms of contributing to applicants' decision to enroll at the university is **"other financing sources"**, typically loans, which applicants need to acquire to supplement federal and institutional awards and aid, to be able to attend college. This feature, with an importance amount of 0.712, is far more impactful on enrollment decision than any other feature.

Financial factors, including family income, expected family contribution (EFC), Federal Pell Grant, need-based aid promised, and merit-based award promised, are all impactful on enrollment decision. Interestingly, although institutional aid in the form of need-based aid and merit-based awards are on the top 12-most important features, their impacts are currently quite minor compared to the other financial-related features. This analysis provides inspiration and motivation for paying more attention to awards and aid, as well as for strategic restructuring of their allocation and distribution.

Aside from financial-related features, we observed that academic performance and "social" features also have some importance in enrollment decision. Academic features include GPA and ACT scores, respectively. "Social" features include how familiar an applicant is to the university and how close in distance an applicant is to the university and to some of its competitors. As explained in the section Feature Engineering (Section 3.5), the familiarity indicator is calculated based on how many students from the same applicant's high school had applied to and eventually enrolled at the university within the last five years.

Table 2: Seven budget-friendly, affordable and accessible strategies that are identified in different runs of simulated quenching. For each strategy (each row), reported numbers are changes in expected outcomes from actual outcomes. Outcomes are accessibility (Acc), unmet need, enrollment (Enr), revenue (Rev), return of investment (ROI), and achievement (Ach). Changes in unmet need are reported in dollars ($). Changes in other outcomes are reported in percentages (%). Strategies are ordered in increasing budget changes.

	Student-Centric		Institution-Centric				Budget Change
Strategy	Acc	Unmet Need	Enr	Rev	ROI	Ach	
S1	110%	$479	3.0%	4.2%	0.2%	-0.2%	-0.4%
S2	109%	$336	2.6%	3.2%	0.1%	-0.2%	1.6%
S3	112%	$481	3.6%	3.8%	0.0%	-0.3%	4.4%
S4	105%	$10	1.3%	0.3%	-0.2%	-0.1%	5.8%
S5	105%	$3	1.3%	0.1%	-0.2%	-0.1%	6.5%
S6	108%	$228	2.6%	2.0%	-0.2%	-0.2%	6.6%
S7	110%	$353	3.0%	2.3%	-0.2%	-0.3%	7.0%

4.3. Budget-friendly and Pro-Student Strategies for Increasing Accessibility

Although we obtained 235 strategies, not all of them are, however, desirable, including strategies that resulted in overall good scores and individual expected outcomes. Due to a number of reasons, risk-averse administrators may be hesitant to adopt strategies that promise great expected outcomes at the expense of an expected large increase in budget, and instead prefer solutions that make moderate changes, but can move the needle in some significant way.

From 235 strategies, we identified seven strategies (shown in Table 2) that were budget friendly, affordable for students, and increased accessibility significantly compared to the current allocation strategy of the university. These seven strategies were chosen based on these criteria: (1) the expected total budget was within $\pm$ 7% of the actual budget, and (2) the expected amount of unmet need was within $500 of the actual budget. These strategies have the following characteristics:

- Expected accessibility is between 105% and 110% higher than the actual accessibility. These numbers are shown in the first column of Table 2.
- Expected enrollment and revenue are a few percentages higher than actual enrollment and revenue, respectively.
- Expected unmet need's increased in an amount of less than $500 from the actual unmet need; expected ROIs are between 0.1% and -0.2% of the actual ROI; expected student achievement's are between 0.1%-0.3% less than actual student achievement. Although these figures are worse than the actual outcomes, they are negligible.

This finding suggests that it is possible to design strategies for allocating merit-based awards and need-based aid in a way that increases accessibility to higher education significantly (more

Table 3: 111 strategies (within ±30% of the current budget) identified by different runs of simulated quenching are divided into three different levels of accessibility increase. Changes in expected outcomes are averaged in each group.

	Student-Centric		Institution-Centric				Budget Change
Group	Acc	Unmet Need	Enr	Rev	ROI	Ach	
A	111%	$478.0	3.5%	4.6%	0.3%	-0.3%	0.3%
B	141%	$2066.0	14.0%	15.8%	0.2%	-1.3%	12.6%
C	179%	$3322.0	27.4%	28.4%	0.0%	-1.9%	27.5%

than 100%), while still maintaining a conservative budget and enjoying a small increase in enrollment and revenue.

4.4. Broader Strategies for Increasing Accessibility

Beyond the seven budget-friendly and conservative strategies reported in Table 2, we explored other strategies that may require budget changes within ± 30% of the actual budget. Out of the 235 strategies, we identified a total of 111 strategies that satisfy this constraint. We grouped these 111 strategies into three groups, based on their similar levels of expected accessibility. Strategies in these 3 groups have averaged expected change in accessibility of 111%, 141%, and 179%, respectively. These strategies are reported in Table 3.

Two important correlations can be observed. First, higher accessibility seems to correlate with higher enrollment, higher revenue, and larger budget change. This seems to suggest that in order to increase accessibility, more money may need to be spent on scholarships and financial aid. While this might be true in general, a closer examination of the strategies (Figure 4) shows that strategies in group A are quite diverse in their demands of the budget. In particular, there are a number of strategies that yield approximately 110% increase in accessibility while remaining budget neutral. In fact, we discussed some of these strategies in the previous section (Table 2). Furthermore, it is actually possible to have strategies (indicated as group A* in Figure 4) that can reduce the budget significantly (between 20% to 30%), while still maintaining a high increased accessibility of around 110%.

Second, we observed that higher accessibility seemed to correlate with worse unmet need. For example, while strategies in Group C have higher accessibility than those in Groups A and B, they also had worse unmet needs. This means if they were subjected to strategies in Groups A, B, or C, respectively, students will have to pay an average of $478, $2066, or $3322 more than they do under the university's current strategy. This makes it less affordable for students to attend the university. Although the notion that an increase in accessibility to higher education makes it less affordable for students seems paradoxical, it is not. Giving access has a cost that would not be paid if that access was taken away. A student who is given access may need to pay a cost for that access. And if these costs are added up, the overall result is a decrease in affordability. Fortunately, as reported in the previous section, we were able to identify a few strategies (Table 2) that were affordable, near budget neutral, and yet still increased accessibility significantly compared to the university's current strategy.

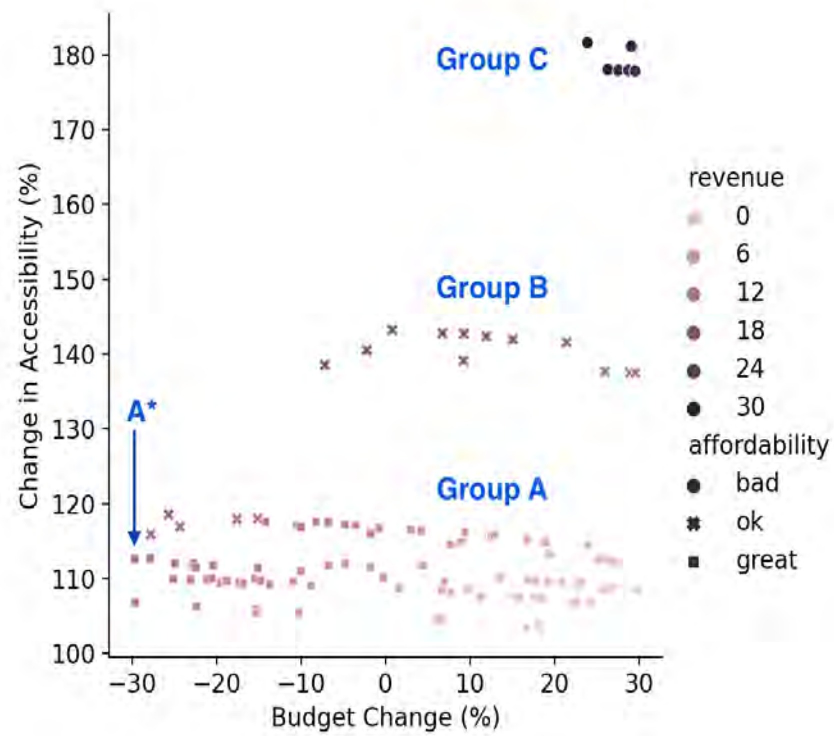

Figure 4: How budget change (x-axis) impacts accessibility (y-axis). Expected outcomes for strategies in Groups A, B, and C are summarized in Table 3. Some strategies require a budget increase up to 30%, yielding increases of +30% in revenue and +180% in accessibility.

Table 4: Allocation of merit-based awards and need-based aid for the seven budget-friendly strategies (Table 2) and strategies in groups A, B, and C (Table 3).

	Accessibility	Merit-based	Need-based	Total Budget
S1	110%	-18.4%	180.9%	-0.4%
S2	109%	-15.5%	173.3%	1.6%
S3	112%	-17.8%	226.7%	4.4%
S4	105%	-2.9%	93.7%	5.8%
S5	105%	-2.2%	93.4%	6.5%
S6	108%	-9.6%	168.7%	6.6%
S7	110%	-8.3%	160.1%	7.0%
Group A	111%	-21.6%	220.1%	0.3%
Group B	141%	-36.2%	502.1%	12.6%
Group C	179%	-46.3%	767.7%	27.5%

4.5. Impacts on the Allocation of Merit-based Awards and Need-based Aid

We report a clear trend in the allocation of merit-based awards and need-based aid in almost all of the strategies we identified. At all levels of changes in the budget, ranging from -0.4% to 27.5%, we saw a clear trend in the reduction in allocation to merit-based awards and the increase in allocation to need-based aid. This is shown in Table 4. Specifically, the seven budget-friendly, affordable strategies called for a reduction between -2.9% to 18.4% in merit-based awards and an increase between 93.7% to 226.% in need-based aid. More aggressive strategies in Groups A, B, C are even more aggressive in reallocating funds from merit-based awards to need-based aid.

4.6. Low Impact of ROI and Achievement

In our experimentation, we observed that the impact on expected ROI and merit profiles is relatively small. This can be seen in Tables 2 and 3. ROI deviated 3% from the baseline in either direction at most, with the majority of deviations being a decrease of 1% or less. Similarly, expected achievement decreases from the baseline by 2% at most, which is equivalent to less than a one point drop in average ACT scores.

4.7. The Importance of Considering Multiple Optimization Objectives

Previous literature focused on optimizing for enrollment alone (Aulck et al., 2020; Sarafraz et al., 2015). In our experimentation, we found that optimizing for enrollment alone can be bad for both the university and students in some aspects. As an example, the first strategy in Table 5 only optimizes for enrollment. Enrollment and net revenue increase 32% and 29% over actual figures, respectively. These results appear attractive if considered in isolation. However, this return comes at the expense of a 48% increase in the total financial aid budget as well as a $3320 increase in the average amount of unmet need carried by matriculates, indicating attending the university has been made less affordable on average.

Table 5: Impacts of having multiple optimization objectives.

Opt. Objectives	Institution-Centric Enr	Rev	ROI	Ach	Student-Centric Acc	Unmet need	Budget Change
Enr	32.0%	29.0%	-1.0%	-2.0%	189%	\$3320.0	48%
Enr+Unmet need	17.0%	-26.0%	-3.0%	-1.0%	132%	-\$403.0	191%
Enr+Unmet need+Rev	8.3%	6.9%	-0.2%	-0.5%	124%	\$825.0	16%

Since optimizing for enrollment alone increases unmet need, we next examine the impact of optimizing for both simultaneously. The second strategy in Table 5 adds "unmet need" as an additional objective for optimization. This strategy sees improvements in terms of enrollment (17% increase) and unmet need (decrease \$403). However, these improvements are costly to the university. To achieve this, total financial aid funding is anticipated to rise by 191%, causing a reduction in net tuition revenue of 26%. This trade-off between affordability and net revenue necessitates a more nuanced optimization specification.

If the university wishes to strike a balance between institution- and student-centric outcomes, enrollment, unmet need, and revenue might be simultaneously optimized. The third strategy in Table 5 optimizes all three objectives: enrollment, unmet need, and revenue. From the university's perspective, this strategy results in a modest increase in enrollment (8.3%) and revenue (6.9%). It, however, requires significantly less investment in total funding than the other two strategies, and could result in cost savings. From the student perspective there is still a marked gain in accessibility of 124% in all of the strategies. Unmet need, as a measure of affordability, on average still experiences an increase of \$825. This is a great improvement over optimizing for enrollment alone.

4.8. Comparing simulated quenching and stochastic hill climbing

To contrast the runtime and performance of the two methods, we set both to have 100 iterations. For simulated quenching, temperature is decreased in each iteration, and in each iteration, the temperature is kept constant for another m steps. Thus, in addition to an overhead for estimating the initial temperature, the runtime of simulated quenching should be approximately m times slower than hill climbing. Table 6 compares the runtime and change in expected outcome between the two methods, when the optimization objective is enrollment. In both methods, we kept the number of iterations the same at $N = 100$. Simulated quenching has an additional parameter M that dictates how many steps the current temperature is kept constant. We compare the running time and expected change in enrollment for five different values of M.

In terms of runtime, simulated quenching is slower than hill climbing by a constant that is directly proportional to M. In terms of expected change in enrollment, simulated quenching has a higher increase (from 31.9% to 32.3%) than stochastic hill climbing (31.4%). Further, this increase grows with higher values of m (from 1 to 11). Keeping the same temperature longer for simulated quenching (i.e. with large values of M) increases objective scores, but does not seem to produce solutions with more meaningful impacts on expected outcomes.

Although we reported the comparison for optimizing a single objective of enrollment, we, however, observed similar results when we compared the two methods with different combina-

Table 6: Performance and runtime from optimizing enrollment of stochastic hill climbing (SHC) and simulated quenching (SQ).

Optimization Method	**Runtime**	**Change in Expected Enrollment**
SHC	8.6 sec	31.4%
SQ, $M = 1$	16.9 sec	31.9%
SQ, $M = 2$	24.4 sec	32.1%
SQ, $M = 3$	34.1 sec	32.2%
SQ, $M = 4$	43.8 sec	32.0%
SQ, $M = 5$	50.0 sec	32.3%

tions of multiple objectives.

5. RESULTS SUMMARY AND DISCUSSION

Our technical contribution includes a method for optimizing the allocation of merit-based awards and need-based aid. Unique characteristics of this computational approach include:

- The optimization is based on simulated quenching, whose general applicability accommodates flexible formulations of multiple optimization objectives. As a result, the method should be applicable to similar circumstances, helping different institutions optimize their strategies for allocating/distributing financial aid. The applicability of the method is further enhanced by an "auto-start" strategy that allows the simulated quencher to automatically determine one of its key parameters, which is the initial temperature to start the optimization process. Nevertheless, users still need to configure, and to some degree, fine-tune two other parameters: δ, the rate of lowering annealing temperatures, and S_T, the number of steps to keep each temperature constant. With smaller values of δ and larger values of S_T, the optimization takes longer to complete, but can potentially find better solutions with better objective values. With different datasets and practical constraints, users will need to experiment to determine the right points of diminishing return to obtain appropriate values for these parameters.

  Technical details can be reviewed in Algorithm 1. The software package is publicly available at `https://github.com/vtphan/fa_optimization`.

- The method utilizes machine learning in its optimization in an interesting way. Throughout the optimization, newly generated solutions (which are strategies for allocating financial aid) are essentially new samples that do not currently exist. This approach is applicable to other problem domains, in which there is a need to optimize certain features that are dependent on an external strategy. In our problem, these features are "merit-based amounts" and "need-based amounts", which are amounts of money given to students based on their merit and need profiles. These amounts are dictated by an allocation strategy. As a strategy is revised, these amounts change. As the two features change, new samples are effectively created. And as such, prediction of the target variable and expected outcomes can be made. The characteristic of this type of problems makes it possible to optimize the strategy that influences certain features. The process that uses machine learning for optimization is depicted in Figure 1.

Although we adopted gradient boosting as a classification method of choice for the optimization, any classification method that makes probabilistic predictions will work. Another method might work better than gradient boosting for a different dataset with different contraints.

In solving a challenging problem of finding good strategies to allocate both merit-based awards and need-based aid, we found several important results.

- Our modeling suggests that financial factors play the most important role in influencing applicants' decision to attend the university. The amount of money that applicants need to come up with (e.g. personal loans) in addition to federal and institutional awards and aid is, by far, the most important factor impacting their decision. At the same time, the university's promised amounts of awards and aid have much less impact.

- We identified seven allocation strategies that promised to improve accessibility to higher education by more than 100%, compared to the current status quo. These strategies also promise that other outcomes won't take a hit. As a matter of fact, the strategies suggest that enrollment and revenue might increase by a few percentages. To accomplish these outcomes, the strategies require an investment up to 7% in the financial-aid budget, which is quite reasonable for a gain of 100% in accessibility.

  Within the context of this work, the notion of accessibility is strictly based on financial means, which we think is an important factor to promote equitable access to higher education. Nevertheless, equitable access encompasses a multitude of elements beyond financial means (Gidley et al., 2010).

- We additionally identified a total of 111 strategies, which might produce high increases in enrollment, revenue, affordability and accessibility for students, etc. While it is possible that large amounts of budget increase might be required, the diversity of the strategies, particular those in Group A as shown in Figure 4, suggests that there exist cost-effective strategies (e.g. those in Group A*).

- We observed that these strategies called for a redistribution of merit-based awards to need-based aid. In some cases, strategies predict need-based aid to be increased by more than 100% to achieve the expected increase in accessibility and affordability.

  It is possible that under the current scheme, need-based aid is badly under-funded by the university. And as such, to increase access to students, it makes sense that need-based aid must be increased significantly. And if there is limit on how much the total budget can be increased, these funds need to be redistributed from merit-based awards to need-based aid.

- We selected a few strategies to demonstrate the necessity of optimizing multiple objectives. In particular, neglecting to optimize student-centric objectives (e.g. affordability or accessibility) might result in strategies that are great for a university (e.g. high enrollment and revenue), but create burdens for students.

- We observed that two of those outcomes, achievement and ROI, did not see meaningful changes for the strategies selected. In the case of achievement, we could not improve it significantly over the actual outcomes. Perhaps this was due to the fact that as strategies drive the need-based aid budget upward we are successfully recruiting students from lower-income backgrounds who have traditionally scored lower in standardized tests (Barrow and Rouse, 2006). Or there could be a limit to our ability to recruit high achieving students due to the university-constrained maximum merit-based award amount. For ROI, we could identify strategies that yielded small amounts of increase. However, those strategies were not as compelling as those we chose to include.

The highest threat to the validity of this approach is that the decision to attend college is complex. In reality, this decision could be influenced by factors that our model did not take into account, e.g. how friends and family influence applicants' decision to attend the university. As such, there is a limit to how accurate our model can predict enrollment, and consequently other expected outcomes.

The fact that our approach yields systematic strategies for uniformly and consistently distributing financial aid is an important advantage that can be taken for granted. These strategies clearly outline the rules for distributing awards and aid; based on a student's achievement and need profile, the student will receive a specific amount. Such a consistent delivery of financial aid can be hard to design without a way (like ours) to predict expected enrollment, revenue, etc.

In a modern world, while a university is finetuning its strategy to attract students, a competing university is also watching and revising its own strategy. Thus, optimizing financial aid distribution is not a one-and-done thing. To be competitive in a world, where social and economic changes are moving a fast pace and where competitors are watching, there is a need to be agile and change quickly. Therefore, there is a need for a systematic approach to optimizing financial aid distribution, instead of relying on heuristics and rules of thumb, whose outcomes are hard to predict.

Although the problem we aim to address is complex, the generality and applicability of our solution suggest it can be a practical solution for universities. This does not mean that universities can take strategies produced by this approach and put them to work without any adjustment. They, however, can serve as a basis for implementing strategies that provide better access to higher education, and for making higher education more affordable. The promise is that this might be achievable without hurting the institutions.

References

Alhassan, J. and Lawal, S. 2015. Using data mining technique for scholarship disbursement. *International Journal of Information and Communication Engineering 9,* 7, 1741–1744.

Antons, C. M. and Maltz, E. N. 2006. Expanding the role of institutional research at small private universities: A case study in enrollment management using data mining. *New Directions for Institutional Research 2006,* 131, 69–81.

Assalone, A., DeShawn, P., and McElroy, B. 2018. Unexpected hurdles: Unpacking the price tag of college affordability. Policy brief, Southern Education Foundation, Atlanta, GA.

Aulck, L., Nambi, D., and West, J. 2020. Increasing enrollment by optimizing scholarship allocations using machine learning and genetic algorithms. In *Proceedings of the 13th International Conference on Educational Data Mining*, A. N. Rafferty, J. Whitehill, V. Cavalli-Sforza, and C. Romero, Eds. International Educational Data Mining Society, Durham, United Kingdom, 29–38.

BARROW, L. AND ROUSE, C. E. 2006. U.s. elementary and secondary schools: Equalizing opportunity or replicating the status quo? In *The Future of Children*, C. Paxson, E. Donahue, C. T. Orleans, and J. A. Grisso, Eds. Vol. 16. Princeton University Press, Princeton, NJ, USA, 99–123.

BASU, K., BASU, T., BUCKMIRE, R., AND LAL, N. 2019. Predictive models of student college commitment decisions using machine learning. *Data 4,* 2, 65.

CHANG, L. 2006. Applying data mining to predict college admissions yield: A case study. *New Directions for Institutional Research 131*, 53–68.

CHEN, T. AND GUESTRIN, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '16. Association for Computing Machinery, New York, NY, USA, 785–794.

DAVIS, D. J., GREEN-DERRY, L. C., AND JONES, B. 2013. The impact of federal financial aid policy upon higher education access. *Journal of Educational Administration and History 45,* 1, 49–57.

DELANEY, J. A. 2014. The role of state policy in promoting college affordability. *The Annals of the American Academy of Political and Social Science 655,* 1, 56–78.

DESJARDINS, S. L. 2002. An analytic strategy to assist institutional recruitment and marketing efforts. *Research in Higher education 43,* 5, 531–553.

DEVASENA, C. L., SUMATHI, T., GOMATHI, V., AND HEMALATHA, M. 2011. Effectiveness evaluation of rule based classifiers for the classification of iris data set. *Bonfring International Journal of Man Machine Interface 1*, 05–09.

FRIEDMAN, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of Statistics 29,* 5, 1189–1232.

GIDLEY, J. M., HAMPSON, G. P., WHEELER, L., AND BEREDED-SAMUEL, E. 2010. From access to success: An integrated approach to quality higher education informed by social inclusion theory and practice. *Higher Education Policy 23,* 1, 123–147.

GOENNER, C. F. AND PAULS, K. 2006. A predictive model of inquiry to enrollment. *Research in Higher Education 47,* 8, 935–956.

HOSSLER, D. 2000. The role of financial aid in enrollment management. In *New Directions for Student Services*, M. Saddlemire, Ed. Vol. 2000. Wiley Online Library, Hoboken, NJ, USA, 77–90.

HOSSLER, D. 2009. Enrollment management & the enrollment industry. *College and University 85,* 2, 2–9.

HOVEY, H. A. 1999. State spending for higher education in the next decade: The battle to sustain current support. Technical report, National Center for Public Policy and Higher Education, San Jose, CA.

INGBER, L. 1993. Simulated annealing: Practice versus theory. *Mathematical and Computer Modelling 18,* 11, 29–57.

JOHNSON, A. W. AND JACOBSON, S. H. 2002. On the convergence of generalized hill climbing algorithms. *Discrete Applied Mathematics 119,* 1-2, 37–57.

JUELS, A. AND WATTENBERG, M. 1995. Stochastic hillclimbing as a baseline method for evaluating genetic algorithms. *Advances in Neural Information Processing Systems 8*, 430–436.

KIRKPATRICK, S., GELATT JR, C. D., AND VECCHI, M. P. 1983. Optimization by simulated annealing. *Science 220,* 4598, 671–680.

LESLIE, L. L. AND BRINKMAN, P. T. 1987. Student price response in higher education: The student demand studies. *The Journal of Higher Education 58,* 2, 181–204.

McKinney, W. 2010. Data structures for statistical computing in python. In *Proceedings of the 9th Python in Science Conference*, S. van der Walt and J. Millman, Eds. Vol. 445. SciPy, Austin, TX, 51–56.

McLendon, M. K., Hearn, J. C., and Mokher, C. G. 2009. Partisans, professionals, and power: The role of political factors in state higher education funding. *The Journal of Higher Education 80,* 6, 686–713.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. 2011. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research 12,* Oct, 2825–2830.

Perna, L. W., Asha Cooper, M., and Li, C. 2007. Improving educational opportunities for students who work. *Readings on Equal Education 22*, 109–160.

Perna, L. W. and Li, C. 2006. College affordability: Implications for college opportunity. *NASFAA Journal of Student Financial Aid 36,* 1, 7–24.

Phan, V., Wright, L., and Decent, B. 2022. Addressing competing objectives in allocating funds to scholarships and need-based financial aid. In *Proceedings of the 15th International Conference on Educational Data Mining*, A. Mitrovic and N. Bosch, Eds. International Educational Data Mining Society, Durham, United Kingdom, 110–121.

Rohman, A., Arif, I., Harianti, I., Hidayat, A. S., Teniwut, W. A., Baharun, H., Dja'wa, A., Rahmi, U., Saragih, J. R., Burhanudin, J., et al. 2019. Id3 algorithm approach for giving scholarships. *Journal of Physics: Conference Series 1175,* 1, 012116.

Sarafraz, Z., Sarafraz, H., Sayeh, M., and Nicklow, J. 2015. Student yield maximization using genetic algorithm on a predictive enrollment neural network model. *Procedia Computer Science 61*, 341–348.

Schapire, R. E. 2013. Explaining adaboost. In *Empirical Inference*, B. Schölkopf, Z. Luo, and V. Vovk, Eds. Springer, New York City, NY, USA, 37–52.

Sexton, R. S., Dorsey, R. E., and Johnson, J. D. 1999. Optimization of neural networks: A comparative analysis of the genetic algorithm and simulated annealing. *European Journal of Operational Research 114,* 3, 589–601.

Shacklock, X. 2016. *From bricks to clicks: The potential of data and analytics in higher education*. Higher Education Commission, London, UK.

Siraj, F. and Abdoulha, M. A. 2009. Uncovering hidden information within university's student enrollment data using data mining. In *2009 Third Asia International Conference on Modelling & Simulation*, D. Al-Dabass, R. Triweko, S. Susanto, and A. Abraham, Eds. IEEE Computer Society, 1730 Massachusetts Ave., NW Washington, DC, 413–418.

Spaulding, R. and Olswang, S. 2005. Maximizing enrollment yield through financial aid packaging policies. *Journal of Student Financial Aid 35,* 1, 27–38.

Sugrue, P. 2014. A model for predicting enrolment yields. *International Journal of Operational Research 19,* 1, 60–67.

Sugrue, P. K. 2010. An optimization model for the allocation of university based merit aid. *Journal of Student Financial Aid 40,* 2, 21–30.

Sugrue, P. K. 2019. A comprehensive merit aid allocation model. *International Journal of Operational Research 36*, 467–476.

Sugrue, P. K., Mehrotra, A., and Orehovec, P. M. 2006. Financial aid management: an optimisation approach. *International Journal of Operational Research 1,* 3, 267–282.

TANDBERG, D. A. 2010. Politics, interest groups and state funding of public higher education. *Research in Higher Education 51,* 5, 416–450.

TRUSHEIM, D. AND RYLEE, C. 2011. Predictive modeling: Linking enrollment and budgeting. *Planning for Higher Education 40,* 1, 12–21.

WALCZAK, S. 1998. Neural network models for a resource allocation problem. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) 28,* 2, 276–284.

WALCZAK, S. AND SINCICH, T. 1999. A comparative analysis of regression and neural networks for university admissions. *Information Sciences 119,* 1-2, 1–20.

WHITLEY, D. 1994. A genetic algorithm tutorial. *Statistics and Computing 4,* 2, 65–85.

ZWICK, R. 2020. Using mathematical models to improve access to postsecondary education. In *Higher Education Admission Practices: An International Perspective*, M. E. Oliveri and C. Wendler, Eds. Cambridge University Press, New York, NY, Chapter 18, 333–346.

ZWICK, R., BLATTER, A., YE, L., AND ISHAM, S. 2021. Using an index of admission obstacles with constrained optimization to increase the diversity of college classes. *Educational Assessment 26,* 1, 20–34.

ZWICK, R., YE, L., AND ISHAM, S. 2019. Using constrained optimization to increase the representation of students from low-income neighborhoods. *Applied Measurement in Education 32,* 4, 281–297.