Prescriptive Methods for Adaptive Learning

by

Galit Lukin

B.S., Technion – Israel Institute of Technology (2016)

Submitted to the Sloan School of Management

in partial fulfillment of the requirements for the degree of

Master of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2020

Author . . . .
Sloan School of Management
May 1, 2020

Certified by . . . .
Dimitris Bertsimas
Boeing Leaders for Global Operations Professor of Management
Thesis Supervisor

Accepted by . . . .
Georgia Perakis
William F. Pounds Professor of Management Science

Prescriptive Methods for Adaptive Learning

by

Galit Lukin

Submitted to the Sloan School of Management
on May 1, 2020, in partial fulfillment of the
requirements for the degree of
Master of Science in Operations Research

Abstract

It is undeniable that recent world events and globalization have transformed online learning into one of the main channels for education. Online learning has become a necessity, not a luxury. Universities, schools, and pre-schools have transformed into the online learning space holding classes of hundreds of students concurrently. However, online learning has yet to reach its full potential. Although educators understand the benefits and effectiveness of online learning platforms, the lack of engagement and evaluation are clear. None the less, these challenges can be solved through machine learning.

In this thesis, we present novel, interpretable prescriptive methods to the online learning setting. We apply these techniques to adaptive learning and test them in real online course settings. We show that using an interpretable, optimal tree-based approach improves both the engagement and the learning rates of the learners. We present PLOpt, a full-stack web app that leverages machine learning models and learner, content knowledge to create assignments that best suit each individual learner. We describe the models, how they were tested, and their evaluation. We demonstrate that by using PLOpt, learners achieved higher engagement and proficiency levels. In addition, we show how PLOpt created assignments that matched the correct difficulty level of the learners so that the learner could remain engaged with challenging questions, yet not frustrated by questions too difficult to answer. Altogether, this work demonstrates that applying interpretable machine learning to online learning builds personalized learning platforms and solves the challenges raised in today's online learning world.

Thesis Supervisor: Dimitris Bertsimas
Title: Boeing Leaders for Global Operations Professor of Management

Acknowledgments

I am grateful to all the people in my life who have helped make this experience possible.

I would like to recognize the invaluable guidance that my advisor, Professor Dimitris Bertsimas, has provided during my Master's degree. From our first meeting, he supported me in my vision of improving education and changing the future of learning. I am fortunate to have combined my passion with his in applying machine learning to improve education. We have worked on various projects in the fields of both education and health and it has been a very enlightening and unique experience to work in both areas with him. I hold with high regards his enthusiasm and advice, both professional and personal. Thank you for your support, understanding, and positively at every meeting we've had.

Deploying PLOpt was a team effort. I thank edX for providing the platform and for the support in integration. Thanks to Omer Goldberg for assisting in building the user interface and infrastructure of PLOpt. I also thank Jonathan Amar and Omar Skali Lami for assisting me in the course forum throughout the experiment.

I was also very fortunate to work on other projects throughout my Master's degree. I pay my special regards to Colin Pawlowski and Yuchen Wang who I truly enjoyed working with on our projects with Quest Diagnostics. I thank Quest Diagnostics, especially Dov Shiffman, for their collaboration and support throughout my two years at MIT. I also thank Professor Susana Mondschein for her collaboration in building personalized mammograms and look forward to continuing this work with her and Professor Bertsimas.

Lastly, I would like to pay my special regards to my family in Israel and the United States for all of the support throughout my career, and especially during my degree. To my parents, without whom I would not be here and this work could not have been done, thank you for your constant support, love and advice. To my siblings, thank you for always having my back. Finally, I wish to show my gratitude to my friends at the ORC for the great and lasting friendships I've built during the last two years.

Contents

1 Introduction **9**

1.1 Contributions . . . 10

1.2 Summary . . . 11

2 PLOpt **13**

2.1 Introduction . . . 13

2.2 Data and Measures . . . 14

2.2.1 Offline and Online Data . . . 14

2.2.2 Measures . . . 15

2.3 Optimal Prescriptive Trees for Personalized Learning . . . 16

2.3.1 Outcomes . . . 18

2.3.2 Prescriptions . . . 19

2.4 Implementation . . . 20

2.4.1 PLOpt Design . . . 20

2.4.2 Personalized Learning OPT Models . . . 22

2.5 Conclusion . . . 23

3 The Analytics Edge 15.071x Experiment **25**

3.1 Introduction . . . 25

3.2 Semester 1: OPT . . . 26

3.2.1 Experiment Setup . . . 26

3.2.2 Models . . . 27

3.2.3 Results . . . 30

3.3 Appendices 39
3.4 Semester 2: Clustering 42
3.4.1 Experiment Setup 42
3.4.2 Models 42
3.4.3 Results 44
3.5 Conclusion 49

4 Summary **51**

Chapter 1

Introduction

Online learning has become the main form of education in the last months and has transformed worlwide views on education. However, online learning has been around for over a decade. MOOCs (Massive Open Online Courses) initially appeared in 2008 and were an educational buzzword by 2012 [3], "the year of the MOOCs" [9], offering a higher education to all at scale. By 2014, the hype had diminished, due to its decrease in press coverage, yet the number of data and analytics related articles increased, suggesting the rising importance of analytics for enhancing the MOOC learning experience [6]. Although the number of new learners who signed up for at least one MOOC decreased from 23 million in 2017 to 20 million in 2018, the constant improvement in the courses has increased revenue for MOOC providers [13]. Thus, extensive research has been done on the factors leading to successful courses and the reasons for low retention rates.

The widely researched topic of low completion rates can be attributed to learners' initial goals and motivation, suggesting that future MOOCs understand and adapt to the diverse student needs [16]. A significant amount of students who intend to complete a course, drop out due to difficulty with the subject matter or unchallenging activities [5]. This has lead to research and experimentation in personalized MOOCs, introducing adaptive forms of instruction and AI; technologies and systems which will soon be taking on a prominent role in MOOCs [2].

1.1 Contributions

In this thesis, we built the Personalized Learning Optimizer, or PLOpt, an engine that integrates with MOOCs to provide dynamic and personalized content. We present this system in Chapter 2. Specifically, we provide the data available in this study and the measures we use to evaluate our work. We describe the prescriptive method used to build the personalized learning models and discusses the implementation and integration of PLOpt. PLOpt uses prescriptive, optimal models to create personalized content which will maximize a learner's expertise. PLOpt is particularly built to work with assignments comprised of a series of questions, each with a certain difficulty level. PLOpt supports grading, feedback and personalized paths so that users can easily interact with the engine.

PLOpt models were trained with data pertaining to past course instances and the engine was integrated into the 2018 and 2019 fall semesters of The Analytics Edge 15.071X, a successful edX MOOC which was first run in 2013, attracts thousands of students every year, and was one of the first machine learning courses available in edX. As a proof of concept, PLOpt was integrated into three individual assignments as an A/B experiment. In Chapter 3, we present the settings, models, and results of both experiments where the second experiment leverages the knowledge gained from the prior results to improve upon the models. From the experiments' analyses, we conclude the following:

1. Learners are more likely to complete an assignment and thus more engaged when learning through personalized content.

2. Personalized assignments increase the learners' proficiency levels.

3. In the initial learning units of a MOOC, many learners are unchallenged; a problem solved through personalized learning.

4. Personalized learning models create assignments in which learners' grades are more reflective of their expertise level.

1.2 Summary

In this thesis, we present the Personalized Learning Optimizer, or PLOpt, an online learning engine that integrates with any educational content and online learning platform to create adaptive content that best suits each individual learner. The engine leverages interpretable machine learning approaches to prescribe the learner with the most appropriate content that suits the learner's understanding level and goals. PLOpt requires only the educational content and user related features to build dynamic personalized assignments that maximize knowledge gain and as a secondary, non-intentional result, also improves learner engagement.

We introduce two versions of PLOpt. The first version builds Optimal Prescriptive Trees [1] to create models that are used to prescribe a learner with a question given their location within the curriculum, their personal features, and past performance. These models are trained using static data, since no previous dynamic, personalized assignments were created. The second version of PLOpt leverages the data created while testing PLOpt V1. We used the personalized assignments created and the experiment's results to improve upon the models. PLOpt V2 uses both OPT and Clustering methods to improve upon the first version's results.

Finally, we present the results of the experiments implemented in real-world settings which included over one thousand participants. We demonstrate the importance and the potential of applying interpretable machine learning approaches to online learning. Doing so increases engagement, improves learning rates, and transforms education systems into personalized ones at scale. Such a revolution opens up education to all, in any circumstance.

Chapter 2

PLOpt

2.1 Introduction

Research and implementations of adaptive learning have become increasingly popular in the last decade with the increase of online learning tools and platforms. As online learning becomes an integral part of education, its challenges become more apparent, leading to machine learning solutions such as adaptive learning.

[14] presents a promising, first framework for an adaptive MOOC, yet does not provide empirical evidence that the framework proved to be beneficial or effective for students. Other personalized learning and differentiated instruction methods in MOOCs are proposed such as [10] and [4], yet not implemented. [11] implemented an adaptive learning method which included a behavioral method. [12] implemented an adaptivity-based framework based on logistic, methodological and technological models to improve drop-out rates. However, both works do not report the models' effect on user performance or engagement.

In this Chapter, we present PLOpt which utilizes data, modeling, and engineering to create a system which creates online, dynamic, personalized assignments for learners based on the educational curriculum, learners' features, and learners' past performance. In Section 2.2 we describe the data used to build the models tailored to the educational content of a specific MOOC. We also present the metrics we use to evaluate our models. In Section 2.3, we explain the Optimal Prescriptive Trees framework used to build the models. Finally, Section 2.4 presents the engineering behind PLOpt and the learning experience it provides.

2.2 Data and Measures

We first describe our data sources for this study and the measures we use to evaluate the success of PLOpt.

2.2.1 Offline and Online Data

Offline data is used by the engine as static data and is comprised of three data sets; data from past courses, from sections prior to the sections of the experiment, and the educational content used by the engine. The fourth data set is the online data gathered by PLOpt.

1. Data from previous instances of the course was used to build the models in PLOpt. It includes information on 3,233 learners' demographics, course performance, and exam results. An exercise's performance is described by the attempts taken to solve it and whether or not it was solved correctly.

2. The second data set includes data from the course in which PLOpt is implemented. It is comprised of the learners' demographics and their course performance prior to the use of PLOpt. This data set also includes the exam results which were used for evaluation.

3. The content, a set of questions that evaluate understanding, constitutes the third data set. For this work, the only metadata required for each question is the question's difficulty level. Data on questions that have been solved in the past can be used to define this parameter. Content experts must define the difficulty level for new questions. The content used in this work also includes the maximum amount of attempts per question, dependency between questions, and different assessments such as multiple choice and free text; all supported

by PLOpt. For the purpose of this work, new content was created to amplify the bank of questions. The difficulty level of existing questions was defined by data from previous course instances. Since the course is an introduction to analytics and classic machine learning models, we created the new content and set a difficulty level for each new question.

4. The online data consists of the learners' performance while solving a personalized learning assignment. This data is logged and updated in real time as it directly affects the models' features.

2.2.2 Measures

We measure the learners' models' success using the following metrics calculated per homework assignment $h \in H$ and per learner $u \in U$. We denote Q_{hu} to be the set of questions offered to learner u in assignment h, C_{qu} to be a binary variable indicating if learner u answered question q correctly, and A_{qu} to be the amount of attempts learner u used to answer question q. Lastly, L is the the set integer values expressing the possible difficulty levels available to each question and Q_{hlu} is the set of all questions $q \in Q_{hu}$ with level $l \in L$.

- The *Score* metric is considered our main metric in evaluating the models' success since it describes the learners' understanding level, or efficacy of PLOpt to maximize a learners' knowledge. The metric considers a learner's performance in each level, assigning linearly more weight as the question's difficulty level increases. It is calculated per assignment and per user as follows:

$$Score_{hu} = \sum_{l \in L} \left(l \cdot \frac{\sum_{q \in Q_{hlu}} C_{qu}}{\sum_{q \in Q_{hlu}} A_{qu}} \right).$$

- The *Grade* metric is the percentage of questions answered correctly. Unlike scores, grades disregard amount of attempts or a question's difficulty. This metric is used to calculate the learners' grades in the course.

$$Grade_{hu} = 100 \cdot \frac{\sum_{q \in Q_{hu}} C_{qu}}{|Q_{hu}|}.$$

- We use Score and Grade to evaluate learners' comprehension level using their performance on the exam. The exam is a set of static questions E, offered to all learners. We use the *Exam Grade* metric and the *Exam Score* metric to measure the learners' level of proficiency at the end of the course.

$$ExamScore = \sum_{l \in L}(l \cdot \sum_{q \in E} \frac{C_{qu}}{A_{qu}}). \qquad ExamGrade = \sum_{q \in E} w_q \cdot C_{qu}.$$

- An assignment's *Completion Rate* metric is used to compare the learners' engagement. The metric is calculated per assignment h, yet is challenging to calculate for a personalized assignment. If a learner does not complete a personalized assignment, it is unclear how to estimate the amount of questions the learner had left. Therefore, we treat this metric as the following binary variable and measure each learner's overall completion rate.

$$CompletionRate_{hu} = \mathbb{I}_{\{\text{User } u \text{ completed assignment } h\}} \forall u \in U, h \in H.$$

$$CompletionRate_{u} = \mathbb{I}_{\{\text{User } u \text{ completed all 3 assignments}\}} \forall u \in U.$$

- We measure the *Difficulty* of an assignment offered to a learner. Since learners are offered different sets of questions, the assignment h offered to a learner u varies. Thus, so does its level of difficulty which is defined below.

$$Difficulty_{hu} = \sum_{l \in L} l \cdot \frac{|Q_{hlu}|}{|Q_{hu}|} \qquad \text{st } \cup_{l \in L} Q_{hlu} = Q_{hu} \text{ and } \cap_{l \in L} Q_{hlu} = \emptyset.$$

2.3 Optimal Prescriptive Trees for Personalized Learning

The models used by PLOpt prescribe a learner with a question given past data, maximizing the learner's mastery level. We use a tree based algorithm, Optimal Prescriptive Tree (OPT) [1], to prescribe a learner with a question in an assignment. OPTs are a natural fit for our purpose since they are generally used for personalized decision making. They use observational data $(x_i, y_i, z_i)_{i=1}^n$ which includes features $x_i \in \mathbb{R}^d$, the outcomes $y_i \in \mathbb{R}$, and the real prescription $z_i \in [m] = 1, ..., m$ given n observations.

OPTs' objective balances optimality and accuracy, simultaneously predicting unknown counterfactual outcomes and prescribing the best treatment or in our case, question. The best next question is determined by minimizing $E[y(\tau(x))]$ where $\tau : \mathbb{R}^d \to [m]$ is the next best question policy which provides the next question out of m options given features x. When minimizing $E[y(\tau(x))]$, we choose $\tau(x)$ such that the next question is that which minimizes the outcome. In other words, to choose the next best question, we minimize the expectation of $y(\tau(x))$ taken over the distribution of outcomes for a given next question policy $\tau(x)$. That is, we minimize

$$\sum_{i=1}^{n} (y_i \mathbb{I}[\tau(x_i) = z_i + \sum_{t \neq z_i} \hat{y}_i(t) \mathbb{I}[\tau(x_i) = t]) \tag{2.1}$$

The objective function (2.1) is considered the prescription error where $\forall i,\ \hat{y}_i(t)$ is unknown and is the counterfactual outcome that would have been observed if sample i had been assigned treatment t. OPTs simultaneously attempt to accurately estimate the counterfactual outcomes by solving a second objective function considered to be the prediction error. The prediction error minimizes the squared prediction error for the observed data. Namely, it minimizes

$$\sum_{i=1}^{n} (y_i - \hat{y}_i(z_i))^2 \tag{2.2}$$

Since OPTs balance optimality and accuracy, they seek the next best question policy $\tau(x)$ by minimizing a convex combination of the objectives (2.1) and (2.2), using a hyper-parameter μ to control the trade-off between the prescriptive and the prediction error. Specifically, OPTs minimize

$$\mu \left[\sum_{i=1}^{n} \left(y_i \mathbb{I}[\tau(x_i) = z_i] + \sum_{t \neq z_i} \hat{y}_i(t) \mathbb{I}[\tau(x_i) = t] \right) \right] + (1 - \mu) \left[\sum_{i=1}^{n} (y_i - \hat{y}_i(z_i))^2 \right] \tag{2.3}$$

In practice, the OPT algorithm uses coordinate descent to train a decision tree according to a loss function of the form

$$\underset{T}{min} \text{ error(T,D)} + \alpha \cdot \text{ complexity(T)}$$

where T is the decision tree being optimized, D is the training data, and error(T,D) is the function measuring how well the tree T fits the training data D. In other words, error(T,D) is replaced by (2.3). Complexity(T) is a function penalizing the complexity of the tree and α is the complexity parameter that controls the trade-off between the fit and the size of T. The algorithm repeatedly attempts to find changes in T that improve the objective value received from (2.3) by changing splits in the tree; deleting, adding, and altering splits in the tree.

The decision tree divides D into neighborhoods containing similar samples; samples that have similar features and fall under the same leaf in T. OPTs exploit this idea to estimate the counterfactual $\hat{y}_i(t)$ by using the outcomes y_j for all samples that fall into the same leaf as i whose treatment in the data is t. In this work, we use an immediate method for estimation such that if $X_{lf(i)}$ is the leaf of the prescription tree into which the sample x_i falls into, then

$$\hat{y}_i(t) = \frac{1}{|j : x_j \in X_{lf(j)}, z_j = t|} \sum_{j:x_j \in X_{lf(j)}, z_j = t} y_j. \tag{2.4}$$

We adopt OPTs for personalized learning where each observation i is comprised of a question q_i and a learner u_i. Observation i's prescription z_i is question q's difficulty level, its outcome y_i is learner u_i's performance on question q, and x are features belonging to question q and learner u_i.

2.3.1 Outcomes

In our models, we maximize the mastery level of the learner using the same intuition used for the score metric. We define learner u's outcome on question q to be

$$y_{qu} = l \cdot \frac{C_{qu}}{A_{qu}} \ st \ q \in Q_{hlu}, \ h \in H, \ l \in L, \ u \in U. \tag{2.5}$$

When prescribing the next best question, we consider three possible questions: the first unsolved question that will allow for review, practice, or challenge. Since OPTs minimize outcomes, we substitute (2.5) into (2.3) as follows. We solve this problem to find the optimal prescriptive tree T given an assignment $h \in H$ and difficulty level l st 1$< l < |L|$.

$$\begin{aligned}
\min_{\tau(.)} \quad \mu\Bigg[\sum_{l'=l-1}^{l+1} \sum_{u \in U, q \in Q_{hlu}} \Bigg(&- l' \cdot \frac{C_{qu}}{A_{qu}} \mathbb{I}[\tau(x_i) = z_i] + \\
&\sum_{t \neq z_i} \frac{\sum_{j: x_j \in X_{lf(j)}, z_j = t} -l' \cdot \frac{C_{qu}}{A_{qu}}}{|j : x_j \in X_{lf(j)}, z_j = t|} \mathbb{I}[\tau(x_i) = t]\Bigg)\Bigg] + \\
(1-\mu)\Bigg[\sum_{l'=l-1}^{l+1} \sum_{u \in U, q \in Q_{hlu}} &\Bigg(- l' \cdot \frac{C_{qu}}{A_{qu}} - \frac{\sum_{j: x_j \in X_{lf(j)}, z_j = t} - l' \cdot \frac{C_{qu}}{A_{qu}}}{|j : x_j \in X_{lf(j)}, z_j = t|}\Bigg)^2\Bigg] \qquad (2.6)
\end{aligned}$$

For example, if the learner will easily succeed a question of a higher difficulty level, we expect the model to prescribe an "Increase". However, if the model predicts that the learner will struggle with questions in a higher and current level, we expect the model to prescribe a "Decrease". If a learner currently completed a question belonging to level l st 1$< l < |L|$, then the outcomes of the first question not yet solved in levels $l + 1$, l, $l - 1$ will be predicted and the question with the highest outcome value will be prescribed. The treatment for the lowest and highest difficulty levels is explained in Section 2.4.2.

2.3.2 Prescriptions

When a learner completes a question, PLOpt provides the learner with the next question that is most suited for the learner's current understanding of the material. If the learner is knowledgeable, the model should recommend that the user move on to the next assignment or material. Alternatively, the model should recommend a struggling learner an easier question to prevent frustration, allowing for the review of simpler concepts. Lastly, for learners who are in neither one of the mentioned groups, we expect the model to prescribe the learner with more practice in the current difficulty level. Therefore, we implement OPTs such that there are three possible prescriptions, namely "Decrease", "Increase", and "Continue". In other words, if the last solved question is of level l, the model prescribes the learner with a difficulty level l' which is equal to $l-1$, l, or $l+1$. The first unsolved question of each level l' is used to calculate the treatment's outcome and is served if it results as the highest outcome value. Note that two continuous questions served to a learner cannot differ by more than one difficulty level.

2.4 Implementation

2.4.1 PLOpt Design

The models described above were implemented as a third party full-stack web service, named Personalized Learning Optimizer (PLOpt), which was integrated into the edX platform. Offline, the stand alone engine receives the second and third data sets described in Section 2.2.1 from edX using anonymous user IDs. During the learners' online use of PLOpt, the engine communicates with edX at the start of its use, receiving an anonymous user ID and at the end, returning a grade to edX.

PLOpt provides the user with questions, automatic grading, answers, and explanations. The service allows for multiple attempts per question and displays the answer and explanation after the correct answer was submitted or the amount of attempts were exhausted. Between each question served, PLOpt updates the learner's features, runs the corresponding model described in Section 2.3, and serves the learner the prescribed question.

For any assignment, the learner begins at the lowest difficulty level and escalates a level if the model has decided to do so or if all of the questions with the current difficulty level have been answered. The learner completes the assignment when the learner's current difficulty level is the highest and the model has decided to increase it or when the learner has answered all of the questions with the highest difficulty level. Figure 2-1 displays two examples of how learners can progress through the material.

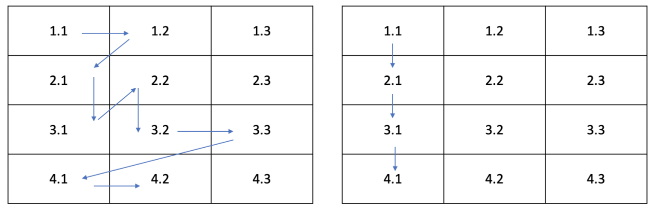

Figure 2-1: Two examples of possible learner progressions through the curriculum. Each row represents a level containing three questions each. An example of a strong learner is shown on the right. The left flow belongs to a more challenged learner who at first remains in level one, advances to levels two and three, returns to level two, and so on.

The content creator can also choose to have mandatory questions which will be served in the correct chronological order independent of the user's performance.

As the learner progresses through the curriculum, PLOpt logs which questions were offered to the learners, the learners' answers, and their attempts. The logged data is used for the online features described in Section 2.2.1, for calculating the learner's grade upon completing the assignment, and for the experiment's evaluation presented in Section 3.2.3.

Figure 2-2 describes the system flow. Learners are directed from edX to PLOpt and back through anonymous user IDs. Through these IDs, PLOpt maintains the student and online data. Pre-calculated offline features and educational content is included into the engine prior to its use. Upon entering PLOpt, learners are served an initial question which is either the first question in the assignment or the last question that the learner encountered in a previous session. As the learner progresses through the material, PLOpt either returns a mandatory response such as feedback on a correct or incorrect response, an additional attempt, or a mandatory question. If no mandatory response is required, the system runs the trained model described in Section 2.3 to receive the next question which the engine then serves to the learner. If the model prescribes that the learner complete the assignment, the learner is notified and redirected back to edX.

2.4.2 Personalized Learning OPT Models

Since some levels are harder to overcome than others, it is clear from the prescriptions and the outcomes that the models differ from level to level. Features can differ as well since questions' appearances are dependant on their level. Similarly, assignments require different features since they measure comprehension in various concepts. Therefore, we build $|L| \cdot |H|$ models, one model for every level and every assignment.

We note that the models for the first and last difficulty levels differ slightly from the rest. The first model does not allow for a "Decrease" prescription, resulting in a model with only two possible prescriptions, "Increase" and "Continue". Models of the highest difficulty level have three prescriptions, yet the "Increase" prescription essentially means that the assignment is complete. This model differs in its outcome variable; as opposed to being the outcome of the next question in a higher level, in models of the highest difficulty level, we define the outcome for learner u to be $(|L|+1) \cdot \sum_{q \in E} \frac{C_{qu}}{A_{qu}}$ where E is the set of questions in the exam.

Furthermore, since personalized learning was unprecedented in the course, the data included only one order of questions that all learners received, that is a static, non-adaptive path. In other words, all users received the same prescription after each question, resulting in challenges when predicting counterfactual outcomes. To resolve the unanimity, we used K-Means Clustering [7] to cluster the learners into three groups which can be described as strong, average, and weak learners. For each cluster, we then randomly assigned a prescription to each learner, question pair such that the prescriptions assigned were at random and not biased towards any type of learner. Since, the clustering was done to assign different prescriptions, we now consider only the chosen prescription as the "real" outcome in the model. For future models, this step can be skipped since learners will have experienced different paths in the data gathered from the experiment presented in Section 3.2.

Finally, all models are trained and tested on data in which each observation is a learner and question pair, where the question is that which the learner last answered. The features include features on the learner and their performance in the course.

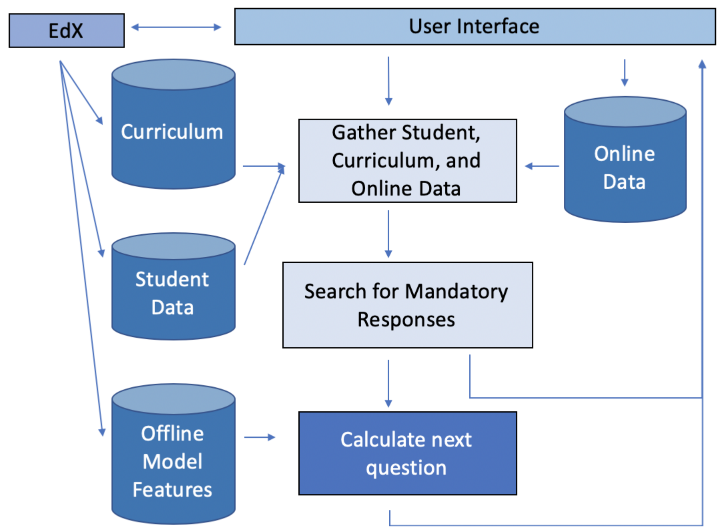

Figure 2-2: The system flow.

Features belonging to the questions can be used as well, yet present challenges when new content is created. The following section describes how these models were implemented in a real-world, online experiment.

Learners can use the PLOpt across multiple sessions, returning each time to their corresponding place in the curriculum. The engine returns a grade only when the assignment is complete. However, since that data is logged, grades for incomplete assignment can be calculated offline.

2.5 Conclusion

In this Chapter, we present PLOpt which utilizes data, modeling, and engineering to create a system which creates online, dynamic, personalized assignments for learners based on the educational curriculum, learners' features, and learners' past performance. The data used to train the models must be specific to the curriculum that PLOpt is being applied to. The Optimal Prescriptive Trees provide interpretable models that can be adopted by any curriculum and support any features. Finally, the generic system is built so that it can integrate and communicate with any educational online learning platform such as edX.

Chapter 3

The Analytics Edge 15.071x Experiment

3.1 Introduction

To test PLOpt, we ran two A/B experiments on one of edX's largest machine learning MOOCs, The Analytics Edge 15.071x. This MOOC runs twice year and over 5,000 learners enroll in the course each semester. Despite the course's popularity, its largest challenge is the low completion rate, which is the issue all MOOCs tackle. This is due to the wide variety of learners that enroll in the course. By applying PLOpt to three of the course's assignments, we solve this issue by providing the learners with adaptive assignments tailored to their knowledge level. In return we see larger knowledge gains and higher grades as expected. An additional, non-intentional result is the increase in engagement. We show that by using, PLOpt learners' completion rates increase, and we solve the greatest challenge online learning faces today.

In Section 3.2, we present the setup, models, and results of the first experiment that tests the effect of PLOpt V1 on the learners' experience. Similarly, Section 3.4 provides the details of PLOpt V2 and its impact on learners.

3.2 Semester 1: OPT

3.2.1 Experiment Setup

The Analytics Edge 15.071x experiment spanned two weeks throughout the Linear Regression unit of the online edX MOOC. The Linear Regression unit is the first non-introductory unit and was chosen due to the high drop out rate in online courses [8]. The unit's assignment is divided into three assignments ($H = \{1, 2, 3\}$), each working on an independent data set and composed of different types and amounts of questions.

We chose to set four proficiency levels for each assignment ($L = \{1, 2, 3, 4\}$). Since the assignments are increasing in difficulty, the corresponding proficiency levels are not equivalent. Each proficiency level is comprised of a constant amount of questions and since the amount of questions in each assignment is increasing, we chose the amount of questions per level to be five, six, and seven for each of the assignments accordingly.

For questions which existed in past courses, the difficulty of the question was set using its correct vs attempt rate. For a question q, as the value $\sum_{u \in U} \frac{C_{qu}}{A_{qu}}$ increased, so did its proficiency level $l \in L$. For proficiency levels that had less than the stated amount of questions, we created content according to the appropriate difficulty level. Each proficiency level had mandatory questions, depending on the goal of the question and its dependence with future questions. Questions within assignments can be dependent, thus any dependence between two questions was solved by making the first question mandatory.

So as to not rely on default values for features, only learners who had solved the introductory unit were considered in the experiment analysis. Learners were divided into two experiment groups at random by edX.

1. The *Baseline* group received a static experience. Learners in this group, denoted U_b, were all given the same set of static questions.

2. The *Personalized* group received a personalized experience. Learners in this

group, denoted U_p, were directed to PLOpt and received questions dynamically as they progressed through the curriculum. The questions served to the learners in this group depended on the static, offline features and their dynamically changing online features.

Learners in both experiment groups received a grade for each of the three assignments. The grade was calculated using the metric, $Grade_{hu}\ \forall h \in H, u \in U_b \cup U_p$. In both groups, there were learners that did not complete assignments. Grades for the Baseline group could be easily calculated since the set of questions was static. However, grades for incomplete assignments in the Personalized group were more difficult to calculate since the remaining amount of points was unknown. We chose to calculate such grades depending on the drop out level of the learner. If the learner stopped answering questions at level i, then $\sum_{l=i}^{4} l$ points were subtracted from their grade. Although learners in different experiment groups experienced different assignments, all learners were given the same exam, testing their knowledge on the material. Due to the course's length, the exam was provided three months after having completed the three assignments.

3.2.2 Models

There were twelve models built all together; one model for every level in every assignment. Table 3.1 describes the features used to build each model which are comprised of demographic, offline and online performance features. Offline features such as location, initial evaluations, and ungraded questions were omitted due to their sparsity in the data. The twelve models trained offline were integrated into PLOpt and are described in Section 2.4.2.

In training the models, we observed that they were very aggressive in progressing learners, suggesting almost always that the learner increase the difficulty level. This is explained by the high success level of the assignments overall due to their early stage in the course. Since most learners correctly answered questions, the score metric was biased towards higher difficulty levels. Therefore, to create more conservative models,

Type	Features	Description
Demographic	Gender	M/F/Null
	Enrollment	Audit/Verified
	Level of Education	JHS/HS/Bachelors/ Masters/Doctorate/Other/Null
	Year of Birth	$\leq$ 1970/1970-1985/1985-1990 /1990-1995/>1995
Offline Performance	Correctness vs Attempts	$\sum_{q \in Q_{hlu}} \frac{C_{qu}}{A_{qu}} \quad \forall u \in U_p, l \in L,$ $\forall h \in \{\text{Introductory Assignments}\}$
Online Performance for level l and assignment h	Current Level Score	$l \cdot \sum_{q \in Q_{hlu}} \frac{C_{qu}}{A_{qu}} \quad \forall u \in U_p$
	Previous Level Score	$l' \cdot \sum_{q \in Q_{hlu}} \frac{C_{qu}}{A_{qu}} \quad \forall u \in U_p, l' < l$
	Previous Assignment Score	$\sum_{l=1}^{4} l \cdot \sum_{q \in Q_{h'lu}} \frac{C_{qu}}{A_{qu}} \quad \forall u \in U_p,$ $\forall h' < h$

Table 3.1: Features used in the model. Demographic features are per learner. Performance features are per learner and question. Offline features are calculated for every graded question prior to the Linear Regression unit. Online features are calculated for every question being served by the model.

instead of minimizing (2.6) presented in Section 2.3.1, we first learn a hyper-parameter r_l. Then, by substituting l with $r_l \cdot l$ in (2.6), we solve the following problem.

$$\begin{aligned}
\min_{\tau(.)} \quad & \mu \Bigg[\sum_{l'=l-1}^{l+1} \sum_{u \in U, q \in Q_{hlu}} \Bigg(- l' \cdot r_{l'} \cdot \frac{C_{qu}}{A_{qu}} \mathbb{I}[\tau(x_i) = z_i] + \\
& \sum_{t \neq z_i} \frac{\sum_{j : x_j \in X_{lf(j)}, z_j = t} - l' \cdot r_{l'} \cdot \frac{C_{qu}}{A_{qu}}}{|j : x_j \in X_{lf(j)}, z_j = t|} \mathbb{I}[\tau(x_i) = t] \Bigg) \Bigg] + \\
& (1 - \mu) \Bigg[\sum_{l'=l-1}^{l+1} \sum_{u \in U, q \in Q_{hlu}} \Bigg(- l' \cdot r_{l'} \cdot \frac{C_{qu}}{A_{qu}} - \frac{\sum_{j : x_j \in X_{lf(j)}, z_j = t} - l' \cdot r_{l'} \cdot \frac{C_{qu}}{A_{qu}}}{|j : x_j \in X_{lf(j)}, z_j = t|} \Bigg)^2 \Bigg] \\
\text{s.t.} \quad & (l-1) \cdot r_{l-1} < l \cdot r_l < (l+1) \cdot r_{l+1} \quad \forall l \in L.
\end{aligned} \tag{3.1}$$

Figures a and b display two of the twelve models used in the system. Each node (excluding the leaves) in the tree consists of a feature posed as a statement.

If the statement is true, the learner, question pair continues to the node on the right and left otherwise. Starting from the root and ending at a leaf, the learner, question pair progresses through the decision tree and is finally prescribed a level and the corresponding question. Each node contains the amount of samples, or learner, question pairs, which belonged to that node in training. The models differ in amount of samples used to train since not all learners answer every question and each level varies in amount of questions.

Each node also contains a feature which can be a demographic feature or a performance feature. Features of the type *scoreX* where $x \in L$ are online features describing the current score of the learner at level x. Features that contain a different string and a value $x \in L$ can be offline or online features which correspondingly indicate the learner's "correct vs attempts" on level x in an introductory assignment or on a previous Linear Regression assignment. The offline features are constant since the learner completed the introductory assignment prior to the Linear regression assignment. Lastly, each node contains a table with a value for each possible prescription. The value is the prediction of the outcome for an observation in that node. Since OPTs minimize outcomes and our model maximizes mastery levels, the models minimizes the negative value of the prescriptions' predicted outcomes. Therefore, each leaf chooses the prescription with the minimum value.

For example, Figure b displays the model trained for the first, easiest assignment and for the second difficulty level. 9696 learner, question pairs were used to train the model. Since this is the first assignment, it does not rely on results from previous Linear Regression assignments, but does use features from previous assignments in the introductory unit such as the assignment represented by "ad". If the learner's correct vs attempts ratio on questions of level four in assignment "ad" is at least 0.7167, we continue to the right side of the tree. If the learner's correct vs attempts ratio on questions of level three in assignment "ad" is less than 0.5583, the learner is given a question of level one, decreasing the difficulty. Note that both of these features are offline features and that such a learner will continuously be decreased a level. This implies that the success in the third level of assignment "ad" is critical for the comprehension of the current level. Eventually, the learner will exhaust the five questions in level one, in which case he will be provided the next question in level two. Once all questions in levels one and two are answered, the learner will move on to level three, after having had sufficient practice. Learners who succeeded in levels three and four in "ad", continue towards the right of the tree. If the learner's score in level one of the current assignment is less than 0.4167, the learner is directed back to this level to improve his mastery level. Otherwise, the learner will either continue in the current level or increase the difficulty level. If the correct vs attempts ratio for level two in assignment "ad" is less than 0.883 and the same metric for level three in the introductory assignment "de" is less than 0.875, then the learner receives the next question in level two. Otherwise, the learner will receive the next unanswered question of level three. All twelve models can be found in Appendix 3.3.

3.2.3 Results

We present the results of the experiment described in Section 3.2 using the metrics described in Section 2.2.2. The results of the experiment show that compared to learners solving static, conventional assignments, learners solving personalized assignments using PLOpt

1. Achieved a higher completion rate which was constant across all assignments (97%),

2. Attained a higher expertise level,

3. Were challenged when the static assignments were too elementary and

4. Received grades which were correlated higher with their level of comprehension.

For assignment $h \in \{1, 2, 3\}$ or the exam, when comparing performance and assignment difficulty, the learners considered are all of those who

1. Completed assignment h or the exam and

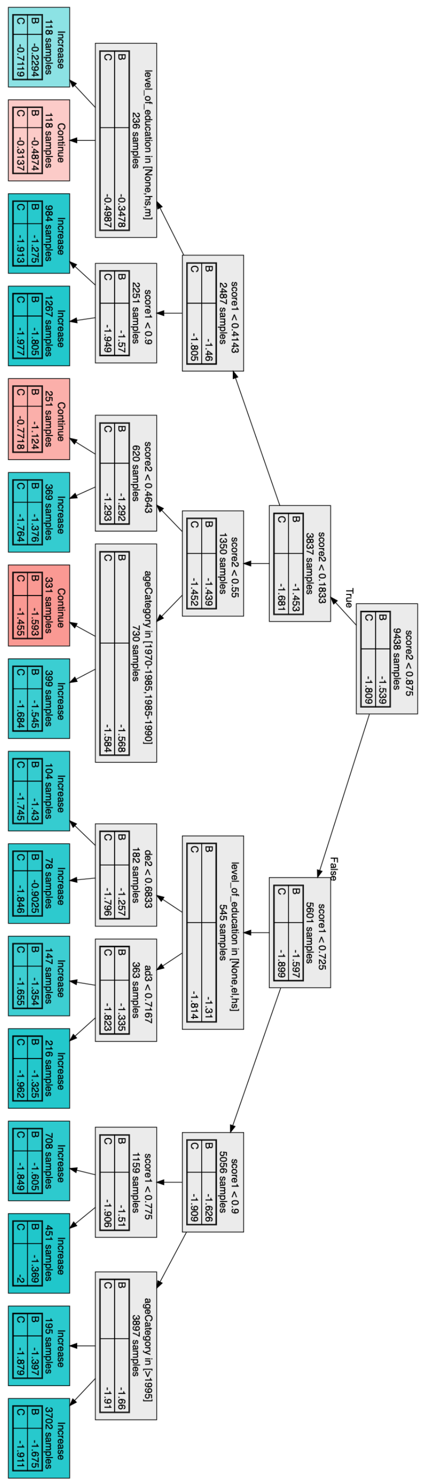

(a) The top model was used for the first level of the second assignment.

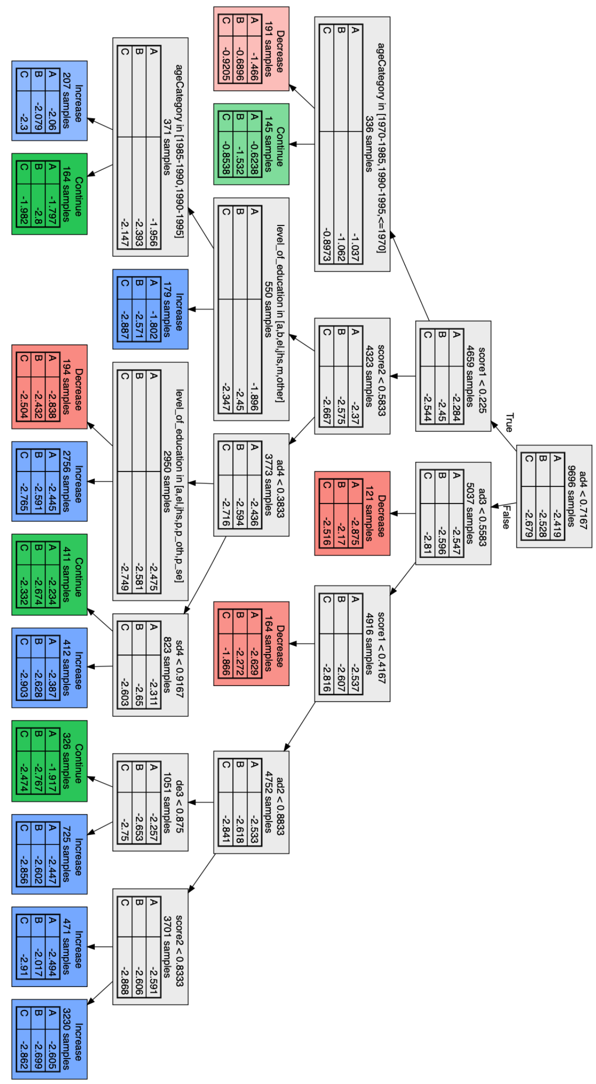

(b) The bottom model was used for the second level of the first assignment.

2. Engaged in the introductory assignment, which had been available to learners the week before the experiment.

The first condition is set so that the metrics are not biased towards the Personalized group which showed a higher completion rate. The latter limits to a population of learners who received assignments which did not only rely on default features regarding their past performance. When comparing completion rates, we include all learners that comply with only the second condition. The resulting group sizes for each assignment are displayed in Table 3.3. Similarly, the amount of learners considered for the analysis of the exam in displayed in Table 3.7. The size of the population considered for the exam is significantly smaller due to the expected drop out rate during the large time period between the assignments and exam. Table 3.2 displays the amount of learners that adhere the second condition.

Engagement

The learners' engagement is measured using the completion rate metric. Engagement was not intended to be maximized, yet is crucial due to today's high drop out rate in MOOCs [8]. The completion rate of each assignment and of each experiment group is the proportion of learners who completed the assignment out of the learners who began the assignment. In Table 3.2, this metric is presented and we show that the Personalized group's completion rate is significantly higher in the second assignment with a 0.1 p-value and in the third assignment with a 0.01 p-value. This implies that as the assignments increased in difficulty, personalized assignments succeeded in retaining more learners. Furthermore, the engagement level in the Personalized group is almost constant as the learners progress in the curriculum. The completion rate of the Baseline group, as expected in a MOOC, declines as the assignments progress and increase in difficulty. Overall, when observing the learners who began the first assignment, with a p-value of 0.05, significantly more learners in the Personalized group completed all three assignments (87.5%).

Proficiency Level

We show that learners using PLOpt achieved higher proficiency levels, which are calculated using the score metric. This metric takes into account correctness, amount

Assignment	Count		Completion Rate		Z-Score	P-value
	Baseline	Personalized	Baseline	Personalized		
1	101	88	0.9604	0.9773	0.6601	0.5093
2	98	83	0.9082	0.9759	1.8639	0.0629
3	92	80	0.8804	0.975	5.185	<0.01
All	101	88	0.7624	0.875	1.9882	0.0466

Table 3.2: Two-tailed z-test of the assignment completion rates for each experiment group showing that learning with PLOpt resulted in higher engagement levels.

of attempts, and the level of difficulty for each question solved. Since there were four levels of difficulty in each assignment, the values of scores range from zero to ten.

A Welch's t-test was performed to compare the scores between the Baseline and Personalized groups (presented in Table 3.3). In the first two assignments, learners in the Personalized group received statistically significant higher scores (p-value=0.01). In the third assignment, the average score of the Personalized group learners was higher, yet not statistically significant. However, in Section 3.2.3, we explain that the third assignment was significantly more difficult in the Personalized group. The difference in difficulty and similarity in proficiency level indicates that at least one of the four models used to create the third assignment challenged the learners excessively, causing a large increase in attempts or a large decrease in correct answers.

Figure 3-2 displays the distribution of scores for each assignment. Linear Regression is the first lesson of the course and we can see how most learners achieve a high level of proficiency. None the less, it is noted that the Baseline group has a larger tail of learners with low proficiency levels in Linear Regression. These results are unsurprising since the OPT models used by PLOpt maximize a metric similar to the score.

Assignment	Count		Average Score		T-Statistic	P-value
	Baseline	Personalized	Baseline	Personalized		
1	97	86	8.1333	9.1353	4.2393	<0.01
2	89	81	8.2222	8.9393	3.6004	<0.01
3	91	78	8.4134	8.655	1.0283	0.3055

Table 3.3: Welch's t-test for the Scores of each experiment group showing that learners in the Personalized group achieved high proficiency levels.

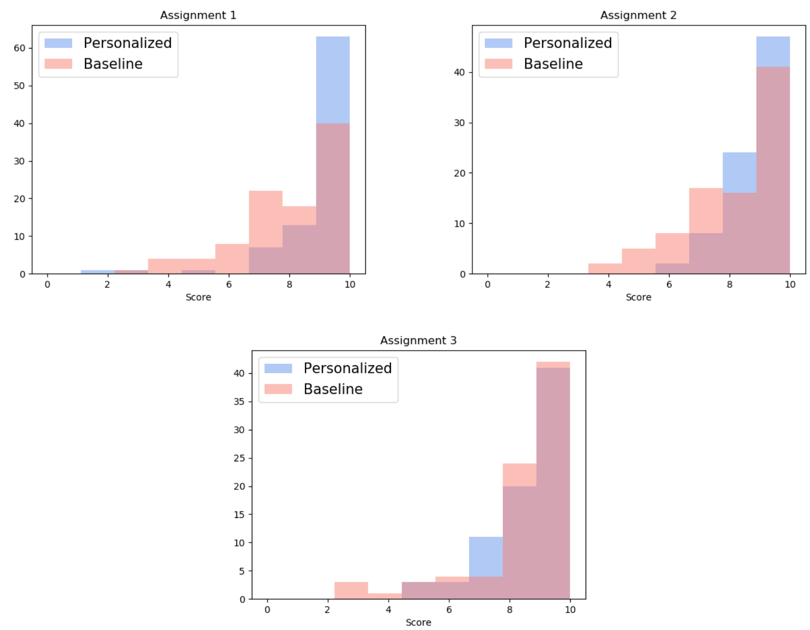

Figure 3-2: A histogram of scores for each assignment showing that the level of understanding varies as expected in a MOOC. None the less, the PLOpt succeeds in minimizing the variance, bringing learners to their maximum comprehension level.

Assignment Grades

In this section, we present the Grade metric and discuss the correlations between grades and scores. A learner's grade on an assignment is the method in which the learner's success is measured in the course. This metric, ranging from zero to one-hundred, takes into account only the correctness of each question. The grade is the metric learners are aware of, thus we assume all learners attempt to maximize their grades in the course.

Table 3.4 presents the results of a Welch's t-test performed to compare the grades between the Baseline and Personalized groups. In the first assignment, there is no significant difference in the learners' grades. In the remaining two assignments, the Baseline learners' grades were significantly higher with a p-value of 0.01 and 0.05 accordingly. That said, we explain in Section 3.2.3 that the third assignment provided to learners in the Personalized group was significantly more difficult. In Figure 3-3, we observe that while in the Baseline group, most learners received very high grades, Personalized learners were more challenged due to the more difficult first and last assignments.

Assignment	Average Grade		T-Statistic	P-value
	Baseline	Personalized	Baseline	Personalized
1	93.52	94.72	0.8203	0.4132
2	94.98	90.13	-3.654	<0.01
3	93.24	89.46	-2.2407	0.0264

Table 3.4: Welch's t-test for the grades of each experiment group. The last two assignments resulted in a significant difference, favoring the Baseline group.

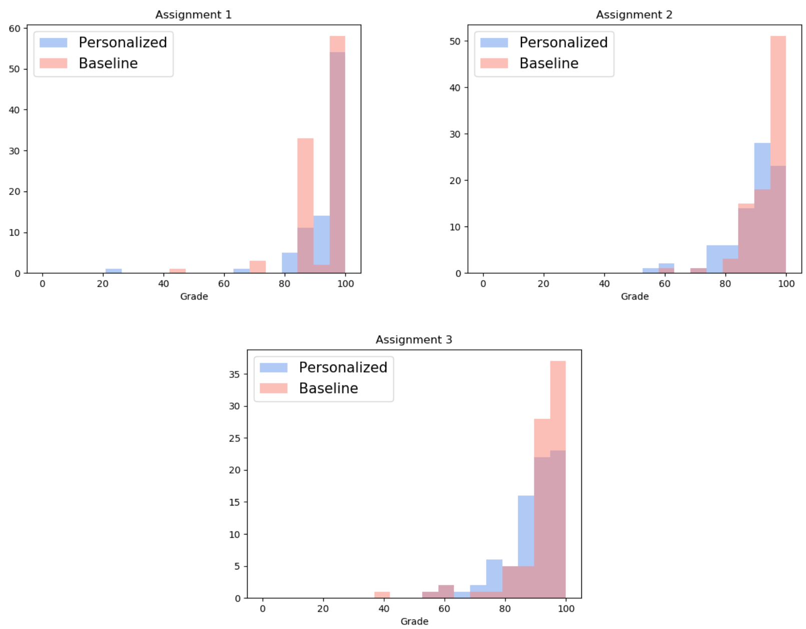

Figure 3-3: A histogram of grades for each assignment shows that the assignments were considered easy by most learners, yet PLOpt challenged learners.

Furthermore, we explore the correlation between the grade and score metrics. A grade should reflect a learner's understanding level; as the proficiency level increases, so does the learner's grade. We calculate the correlation between the two for both experiment groups. In Table 3.5, we present the results of a two-tailed z-test. In all three assignments, the correlation between scores and grades is higher in the Personalized group. The difference is significant for the easiest ($p=0.01$) and hardest ($p=0.1$) assignments. Figure 3-4 displays the relationship between the two metrics where the Baseline group learners achieve a high grade with an average or low score. This is explained by the attempts and weight components which is ignored in the grade metric. Specifically, the discrepancy between the Personalized group learners' low grades and high scores in the second assignment is seen in Figure 3-4 and is explained by the weights. In other words, in the second assignment, Personalized group learners' erred in low difficulty levels, yet performed better in higher ones.

Assignment	Pearson Correlation		Z-Score	P-value
	Baseline	Personalized		
1	0.6183	0.8215	2.916	<0.01
2	0.8501	0.8187	-0.663	0.5073
3	0.8723	0.9253	1.743	0.0814

Table 3.5: Two-tailed z-test of the correlations between scores and grades in each experiment group and each assignment, showing that the grades in the Personalized group were more correlated with the learners' proficiency levels.

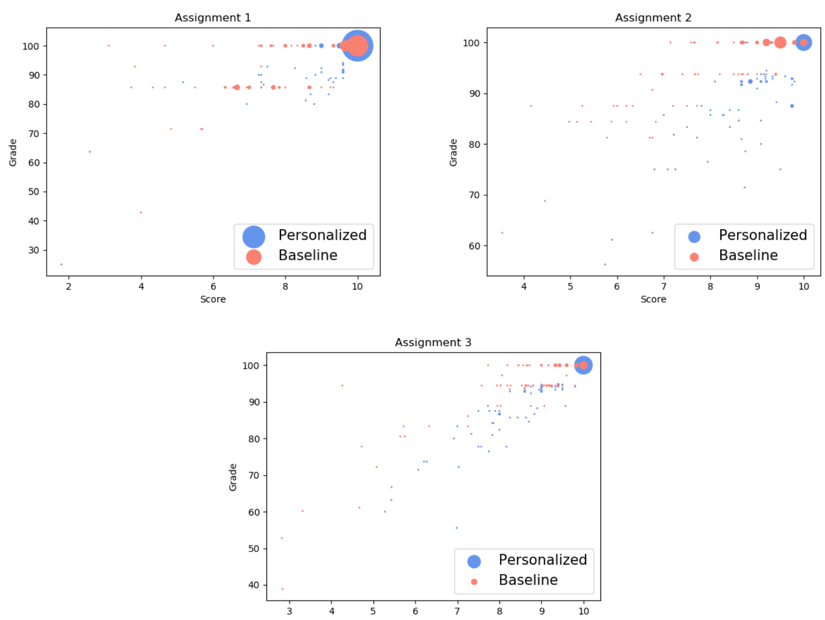

Figure 3-4: The relationship between the learners' scores and grades displaying that PLOpt's dynamic assignments led to grades that were more correlated with the learners' mastery levels.

Assignment Difficulty

The assignments created by PLOpt had a higher difficulty level which is calculated using the difficulty metric as defined in Section 2.2.2. The metric presents the difficulty level of the assignment using the proportion of questions answered in each level.

The results of two-sided t-tests comparing the Personalized assignment difficulty levels to that of the static assignments are presented in Table 3.6. In the first and last assignments, there is a significant difference in the assignment difficulty level with a p-value of 0.01, showing that the assignments offered to the Personalized group were significantly more difficult. In the second assignment, there is no significant difference in the difficulty level. We consider a higher difficulty level a good result due to the learners' high success rate (grades) in the Baseline group.

Assignment	Difficulty Baseline	Average Difficulty Personalized	T-Statistic	P-value
1	2.1667	2.6499	19.0885	<0.01
2	2.5625	2.5494	-0.5878	0.5583
3	2.5	2.6532	11.6702	<0.01

Table 3.6: Two-tailed t-test comparing the difficulty levels presented to the Personalized group to that of the static assignments. We note that PLOpt created challenging assignments which were otherwise easy for learners.

Exam

In this section, we compare the grades and scores of the Linear Regression exam. An identical exam was provided to all learners at the end of the course. Since Linear Regression was the first topic taught, the exam took place three months after the assignments. Table 3.7 shows that there was no significant difference in the distribution of grades or scores between the two experiment groups. We believe this is due to the large amount of time between the assignments and the exam, resulting in the effect known as Ebbinghaus' forgetting curve in learning [15] which describes the exponential decay of knowledge as time passes. Although both groups reached different proficiency levels, the knowledge decay effect led to similar knowledge levels three months after the material was learned.

Metric	Count		Mean		T-Statistic	P-value
	Baseline	Personalized	Baseline	Personalized		
Score	52	44	83.97	83.12	-0.1852	0.8535
Grade	52	44	87.18	89.84	-0.4289	0.669

Table 3.7: Welch's two-tailed t-test for the for the Linear Regression exam scores and grades presenting no significant difference between the groups.

Analysis Per Assignment

Results slightly differ between assignments since each personalized assignment created was a result of four unique models. The differences in results per assignment are explained by their difference in difficulty and the different models used for each level in each assignment, some being more aggressive than others.

The data shows that the first assignment is elementary for most students, resulting in high grades for both experiment groups with no significant difference between them. Nevertheless, the Personalized group achieves a higher proficiency level. This discrepancy is explained by the significantly higher difficulty assignment level served to the learners. Learners in the Personalized group answered significantly more questions of levels three and four and less questions of levels one and two. Lastly, we show that the correlation between grades and scores in the Personalized group is significantly higher than that of the Baseline group, indicating that the assignments provided resulted in grades that better indicated the learners' understanding level.

In the second assignment, we observe that the Baseline group achieves significantly higher grades, and significantly lower scores. Unlike the first assignment, there is no significant difference in the assignment difficulty levels. However, learners using PLOpt erred more in easy questions and slightly less in challenging question. Thus their mastery levels were higher, yet more points were deducted from their grades. Finally, learners in the Personalized group were significantly more engaged in the second assignment.

In the third assignment, we see no significant difference in the scores and higher grades in the Baseline group. However, we show that the average difficulty of the assignments served by PLOpt were significantly higher. This result indicates that

Personalized group learners could have been too challenged. None the less, their scores and grades are significantly more correlated than those of the Baseline group with a very high correlation of 0.925, demonstrating that the assignment provided correctly measured a learners understanding of the material. Moreover, learners were not overly challenged by the engine since the completion rate remained constant as it was in the previous two assignments and significantly higher than that of the Baseline group.

3.3 Appendices

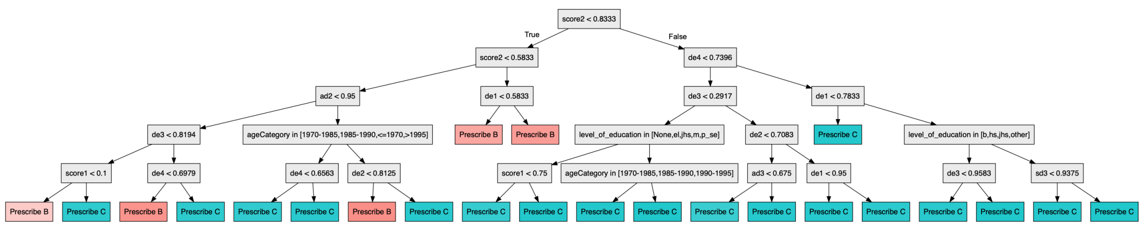

Figure 3-5: Level 1 Assignment 1 OPT

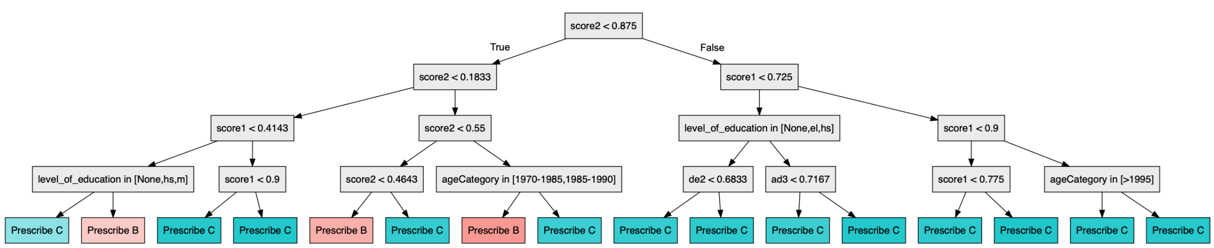

Figure 3-6: Level 1 Assignment 2 OPT

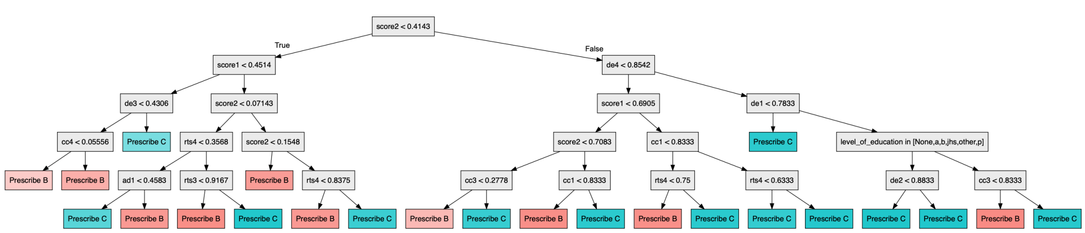

Figure 3-7: Level 1 Assignment 3 OPT

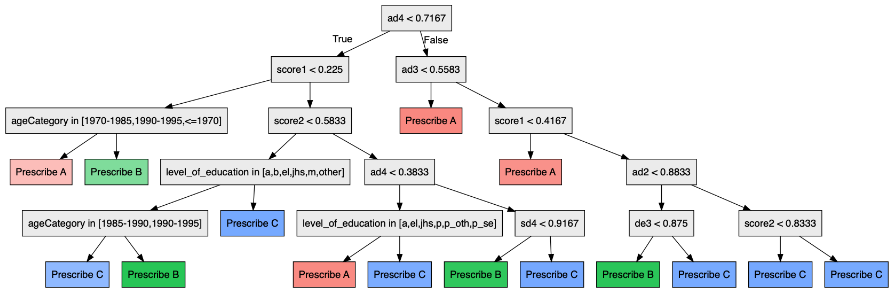

Figure 3-8: Level 2 Assignment 1 OPT

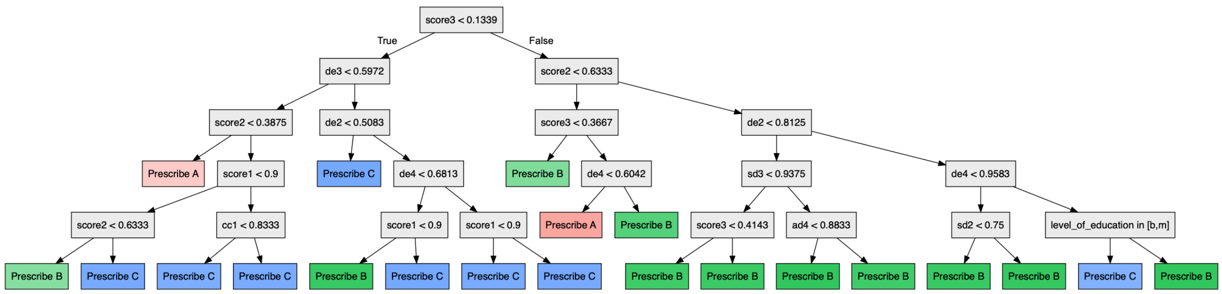

Figure 3-9: Level 2 Assignment 2 OPT

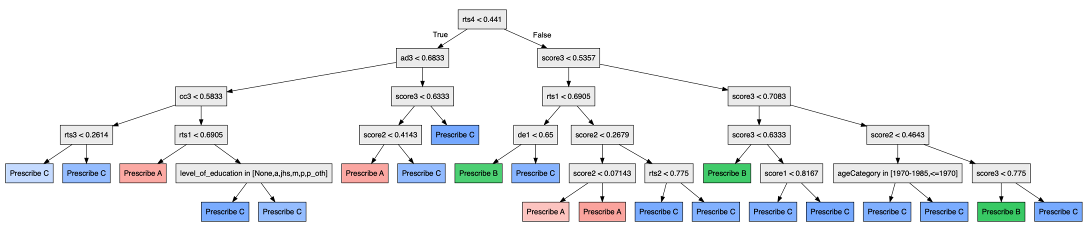

Figure 3-10: Level 2 Assignment 3 OPT

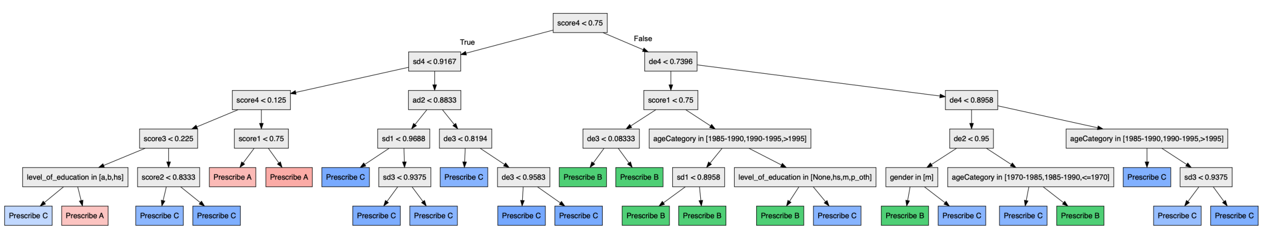

Figure 3-11: Level 3 Assignment 1 OPT

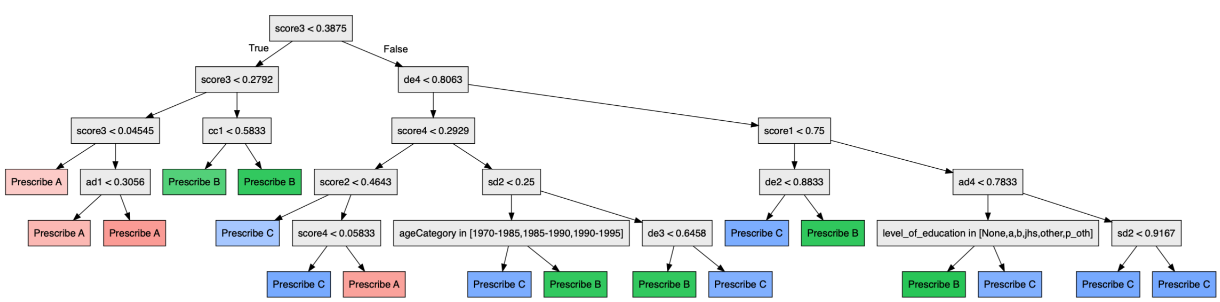

Figure 3-12: Level 3 Assignment 2 OPT

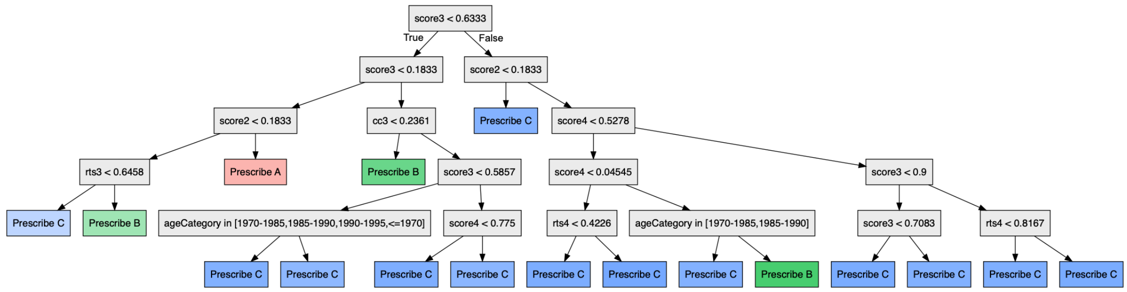

Figure 3-13: Level 3 Assignment 3 OPT

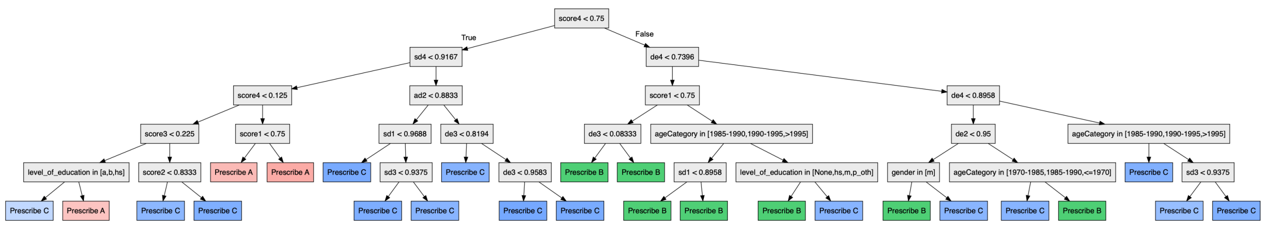

Figure 3-14: Level 4 Assignment 1 OPT

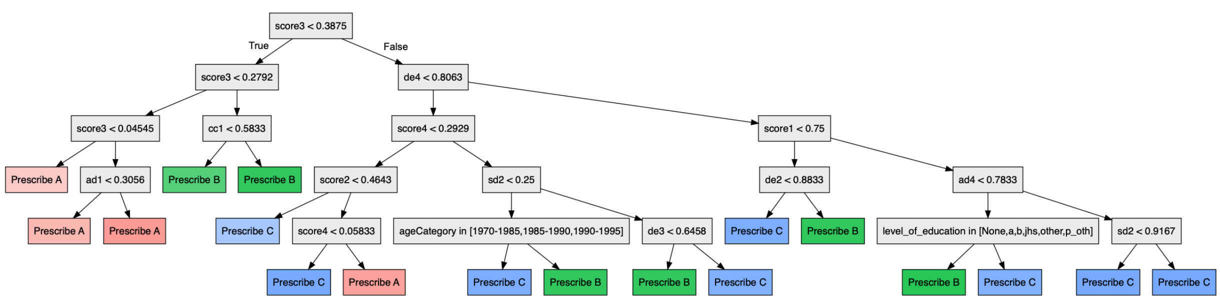

Figure 3-15: Level 4 Assignment 2 OPT

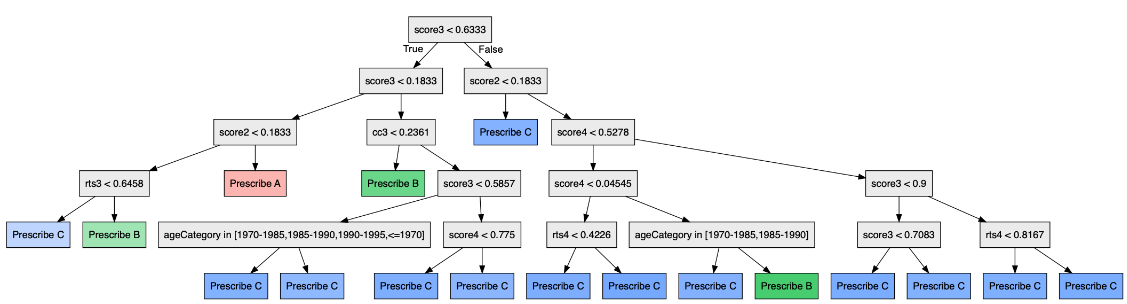

Figure 3-16: Level 4 Assignment 3 OPT

3.4 Semester 2: Clustering

3.4.1 Experiment Setup

We held an additional experiment in the subsequent semester, attempting to improve upon our results. The largest challenge faced in the prior experiment was the models' dependency on the learners' offline features. In practice, offline features were very sparse or non existent because learners do not fill out their personal information and did not solve introductory sections of the course. In this experiment, we improve upon our models, considering only online features, allowing the models to be generic and applicable to additional settings.

The experiment settings remained the same in terms of setup, content and metrics. One difference was that performance data was now available for all the questions in the curriculum and difficulty levels were adjusted accordingly.

3.4.2 Models

To build upon the twelve OPT models, we granularize our modelling so that it is per question as opposed to per level. We utilize the results from the first experiment and evaluate the prescriptions made by the Optimal Prescriptive Trees. For each learner, question pair, we cluster the pair together with similar data points, observe the prescriptions made by our initial models, and evaluate the different prescriptions based on the observations' immediate response (or performance) to the decision. For example, if similar observations were prescribed an increase in difficulty and were successful in the question prescribed, then a similar prescription is given to the learner, question pair. However, if the observations struggled with the increase in difficulty, our new model will prescribe a question of similar difficulty to that of the current question.

We use KNN to to decide which decision needs to be taken. Once an observation's, K-nearest neighbors are found, the decision and outcome of the K neighbors is used to decide the correct decision for new observation.

Traditional KNN chooses the decision z that was most popular within the K neighbors. Instead, we observe the decision taken took for each neighbor and evaluate it. The evaluations are aggregated over the K nearest neighbors and the decision with the best evaluation is chosen. Therefore, this method prescribes a treatment by observing which decision has been taken in the past and its success.

In this work, learners are appointed K neighbors from the previous semester (or current semester if the model is retrained during the experiment) based on their performance features. Therefore, we evaluate decisions taken by OPT (and by this current algorithm). The modeling is done per question q such that the evaluation $e(z_{i,q}(\cdot), p_{i,q'})$ takes into consideration the decision taken for learner i and their performance p in the following question q'. Learners should not be over or under challenged so the evaluation function takes into account the learner's performance relative to the performance distribution of the question.

We first solve KNN for question q where x is data comprised of performance data on q, $p_{i,q}$ $\forall i, t$ and on the learner's current score.

$$\begin{aligned} \underset{k_1 \ldots k_N}{\text{minimize}} \quad & \sum_{i=1}^{N} k_i ||x_i - x|| \\ \text{subject to} \quad & \sum_{i=1}^{N} k_i \geq K \\ & k_i \in \{0, 1\}, \; i = 1, \ldots, N. \end{aligned}$$

We then search for the optimal decision where D are the amount of possible decisions, e_d is the evaluation function for decision d, and y_d indicates which decision is best fit for the cluster chosen above.

$$\begin{aligned} \underset{y_1 \cdots y_D}{\text{minimize}} \quad & \sum_{d=1}^{D} y_d e(z_{i,q}, p_{i,q'}) \\ \text{subject to} \quad & \sum_{d=1}^{D} y_d = 1 \\ & y_d \in \{0, 1\}, \; d = 1, \ldots, D. \end{aligned}$$

We denote the performance features used in the model as follows for question q, learner i

1. c - Correctness
2. a - Attempts
3. s - Time to completion

The evaluation function is then:

$$e = \frac{1}{D} \sum_{i \in \nVdash \{z_{q,i}=m\}} (p_{q,i}), \quad p_{q,i} = \\ ||(\alpha \frac{c_{q,i}}{a_{q,i}} + (1-\alpha)\frac{1}{s_{q,i}}) - (\frac{1}{N}\sum_{j=1}^{N} \alpha \frac{c_{q,j}}{a_{q,j}} + (1-\alpha)\frac{1}{s_{q,j}})|| \quad \forall i$$

Figure 3-17 demonstrates two clustering models for two seperate questions. On the horizontal axis is the time feature. On the vertical axis, attempts and correctness were combing by dividing attempts by correctness. The values of the axes are interpretable since the features were normalized in order to perform KNN. Since attempts and correctness are discreet values, the points are not much scattered across horizontally. The blue points are the learners and the red shaded points are the centroids. The lighter the color of the centroid, the larger the increase in difficulty assigned to that group. We not that the right plot belongs to a level one question and thus, there is no pink centroid, or option to decrease a level. We also note that learners with less attempts increase in difficulty while others do not.

3.4.3 Results

We present the results of the second experiment using the same metrics. The results of the experiment show that compared to learners solving static, conventional assignments, learners solving personalized assignments using PLOpt

1. Achieved similar completion rates even though their assignments were significantly longer and

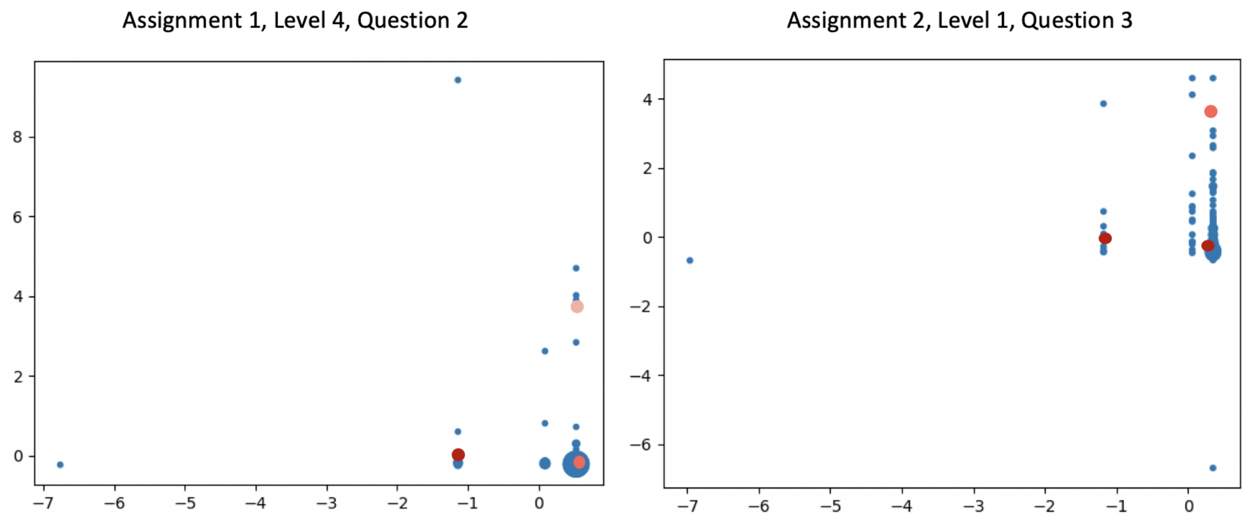

Figure 3-17: Clustering models for two different questions. The blues points are the population. The other points are the centroids of each cluster. Red indicates a decision to decrease the level, orange clusters remain in the same level, and pink clusters increase in difficulty.

2. Attained a higher expertise level,

3. Were challenged when the static assignments were too elementary .

For assignment $h \in \{1, 2, 3\}$ or the exam, when comparing performance and assignment difficulty, the learners considered are all of those who completed assignment h or the exam. By limiting relaxing constraints to our population, we note that more learners were considered in the second experiment that in the first for which learners had to complete the prior assignments.

Engagement

Although in this experiment, there is no significant difference in the completion rates of both groups (Table 3.8), we note in Table 3.9 that in assignment one, the Personalized group received and responded more than twice as many questions as the Baseline group. A similar difference is seen in assignment three, while in the second assignment, we see an inverse ratio. As we continue to present the results, the differences appearing in assignment two will be more apparent.

Assignment	Count		Completion Rate		Z-Score	P-value
	Baseline	Personalized	Baseline	Personalized		
1	231	270	0.9974	0.948	1.512	0.131
2	224	251	0.955	0.924	1.512	0.131
3	217	225	0.963	0.933	1.417	0.157

Table 3.8: Two-tailed z-test of the assignment completion rates for each experiment group showing that there is no significant difference in the completion rates of each group.

Measure	Assignment	Baseline	Personalized	Z-Score	P-value
Responses/Person	1	6.08	13.36	-48.10	<0.01
	2	14.62	11.64	11.95	<0.01
	3	17.14	20.56	-13.09	<0.01
Questions/Person	1	6.08	11.93	-63.56	<0.01
	2	14.62	10.96	12.60	<0.01
	3	17.14	20.01	-13.83	<0.01

Table 3.9: Difference in the amounts of questions and responses between the two groups.

Proficiency Level

We present that once again that the proficiency levels of those learning with a personalized assignment are at least as high as those that aren't. These results are presented in Table 3.10 and are visualized in Figure 3-18.

Assignment	Count		Average Score		T-Statistic	P-value
	Baseline	Personalized	Baseline	Personalized		
1	225	256	7.378	7.8	2.295	0.022
2	224	232	8.046	8.5	1.294	0.197
3	209	210	8.22	8.47	1.608	0.109

Table 3.10: Welch's t-test for the Scores of each experiment group showing that learners learners in the Personalized group achieved at least as high proficiency levels than those learning through a static assignment.

Assignment Grades

In comparing the grades of each assignment, we note that there is no significant difference in assignments two and three. This aligns with the similarities in their proficiency levels. However, learners who learned in the traditional form scored higher

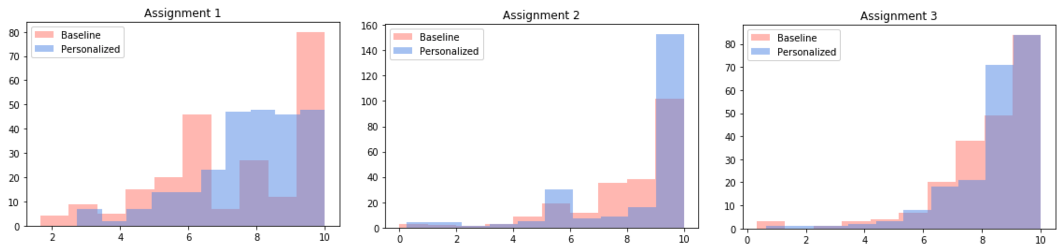

Figure 3-18: Histogram of scores for each assignment.

grades and achieved a lower proficiency level. In the following results, we show that due to the assignments low difficulty, learners in the personalized group were further challenged. The results are presented in Table 3.11 and are visualized in Figure 3-19.

Assignment	Average Grade		T-Statistic	P-value
	Baseline	Personalized		
1	88.21	83.63	3.77	<0.01
2	88.26	87.2	0.806	0.42
3	84.39	86.11	-1.245	0.214

Table 3.11: Welch's t-test for the grades of each experiment group.

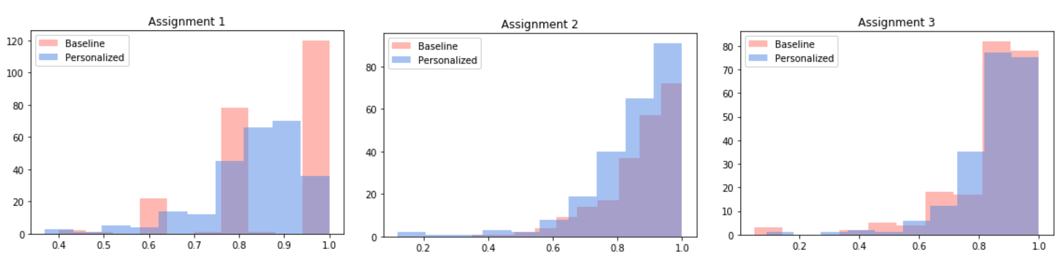

Figure 3-19: Histogram of grades for each assignment.

Assignment Difficulty

Consistently with the first experiment, assignments one and three served to the personalized group were significantly harder. This resulted in higher proficiency levels for the group and grades that were at least as high. However, assignment two yielded different results. In the first experiment, there was no significant difference in the difficulty of this assignment. In this experiment, the assignments created were significantly easier. It seems this assignment is more challenging and when adapted to the learner, results in an easier assessment. Although the assignments were easier, the proficiency levels were similar. This indicates that learners were challenged st the correct level. Additionally, the grades were not affected.

Assignment	Difficulty	Average Difficulty	T-Statistic	P-value
	Baseline	Personalized		
1	2.1667	2.52	46.09	$<$0.01
2	2.5625	1.779	-69.01	$<$0.01
3	2.5	2.511	2.1768	0.0306

Table 3.12: Two-tailed t-test comparing the difficulty levels presented to the Personalized group to that of the static assignments. We note that PLOpt created challenging assignments which were otherwise easy for learners.

Exam

Similarly to the first experiment, the exam took place three months post the assignment deadlines due to curriculum constraints. None the less, the percentage of users in each group that completed the exam was similar (26%). Additionally, as in the first experiment, there was no significant difference between the grades or the scores of the groups. Once again, we attribute this to the Ebbinghaus' forgetting curve in learning.

Metric	Count		Mean		T-Statistic	P-value
	Baseline	Personalized	Baseline	Personalized		
Score	84	104	43.6	45.12	-1.18	0.238
Grade	84	104	89.46	91.6	-1.027	0.306

Table 3.13: Welch's two-tailed t-test for the for the Linear Regression exam scores and grades presenting no significant difference between the groups.

Analysis Per Assignment

As in the first experiment, results varied per assignment due to the assignments different difficulty levels. None the less, as in the first semester, proficiency levels were of the personalized group were at least as high as those of the baseline group.

Engagement was significantly different, although the personalized group answered a significant amount of more questions.

The first assignment, in consistent with the first experiment, resulted easy for learners. We can see in the results that the personalized group had harder assignments, higher proficiency levels, and higher grades. There was no difference in the proficiency levels and grades in the second assignment. However, this assignment had no difference in difficulty in the first experiment and now showed that the original, static assignment was harder than the ones created for the personalized group. On the other hand, the last assignment was more difficult for the personalized group and yielded similar proficiency levels and grades between the groups.

3.5 Conclusion

Since PLOpt is a flexible platform that can integrate into any online learning system, we tested PLOpt in a real-world setting though an online A/B experiment implemented in the edX platform though the famous MOOC, The Analytics Edge 15.071x. We compared the engagement, grades, and proficiency levels of learners who had received personalized assignments through PLOpt to that of learners who had received the static, conventional assignments. From the experiments' analyses, we concluded the following:

1. Learners are more likely to complete an assignment and thus more engaged when learning through personalized content.

2. Personalized assignments increase the learners' proficiency levels.

3. In the initial learning units of a MOOC, many learners are unchallenged; a problem solved through personalized learning. In assignments for which learners are too challenged, PLOpt creates easier assignments, improving proficiency levels.

4. Personalized learning models create assignments in which learners' grades are more reflective of their expertise level.

Chapter 4

Summary

In this thesis, we present a personalized learning engine, PLOpt, which given educational content, past performance and demographic data, can dynamically create assignments that are tailored to a learner's knowledge level. The engine leverages interpretable machine learning approaches to prescribe the learner with the most appropriate content that suits the learner's understanding level and goals. The content's only requirement is that each question be flagged with its difficulty level, yet additional features such as mandatory questions, weighted questions, and multiple attempts are supported as well. Optimal Prescriptive Trees were trained to create the personalized assignments, prescribing learners with questions that maximize their expertise.

PLOpt was integrated into the edX platform as an online A/B experiment in The Analytics Edge 15.071x MOOC across three independent assignments. We compare various metrics such as engagement, grades, and proficiency levels of learners who used PLOpt to that of learners who had received the static, conventional assignments. From our analyses, we derive the following:

1. Learners are more likely to complete an assignment and thus more engaged when learning through personalized content.

2. Personalized assignments increase the learners' proficiency levels.

3. In the initial learning units of a MOOC, many learners are unchallenged; a

problem solved through personalized learning. In assignments for which learners are too challenged, PLOpt creates easier assignments, improving proficiency levels.

4. Personalized learning models create assignments in which learners' grades are more reflective of their expertise level.

We conclude that PLOpt, an engine utilizing prescriptive models for personalized learning, not only allows learners to increase their mastery in the material beyond what was originally provided, yet also increases their engagement in the course. PLOpt can be integrated into MOOCs, advancing the personalized vision, increasing completion rates, and adapting courses to the student's knowledge level so that each learner can maximize their expertise. By integrating such intelligent systems into online learning, education can be accessible to anyone in any location at any time.

Bibliography

[1] Dimitris Bertsimas, Jack Dunn, and Nishanth Mundru. Optimal prescriptive trees. *INFORMS Journal on Optimization*, pages 1,2,164–183, 2019.

[2] Curtis J Bonk, Meina Zhu, Minkyoung Kim, Shuya Xu, Najia Sabir, and Annisa R Sari. Pushing toward a more personalized mooc: Exploring instructor selected activities, resources, and technologies for mooc design and implementation. *International Review of Research in Open and Distributed Learning*, 19(4), 2018.

[3] John Daniel. Making sense of moocs: Musings in a maze of myth, paradox and possibility. *Journal of interactive Media in education*, 2012(3), 2012.

[4] Aracele Garcia de Oliveira Fassbinder, Marcelo Fassbinder, and Ellen Francine Barbosa. From flipped classroom theory to the personalized design of learning experiences in moocs. In *2015 IEEE Frontiers in Education Conference (FIE)*, pages 1–8. IEEE, 2015.

[5] Christian Gütl, Rocael Hernández Rizzardini, Vanessa Chang, and Miguel Morales. Attrition in mooc: Lessons learned from drop-out students. In *International workshop on learning technology for education in cloud*, pages 37–48. Springer, 2014.

[6] Vitomir Kovanovic, Srecko Joksimovic, Dragan Gasevic, George Siemens, and Marek Hatala. What public media reveals about moocs: A systematic analysis of news reports: What public media reveals about moocs. *British Journal of Educational Technology*, 46:510–527, 04 2015.

[7] James MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA, 1967.

[8] Daniel FO Onah, Jane Sinclair, and Russell Boyatt. Dropout rates of massive open online courses: behavioural patterns. *EDULEARN14 proceedings*, 1:5825–5834, 2014.

[9] Laura Pappano. The year of the mooc. *The New York Times*, 2(12):2012, 2012.

[10] Gilbert Paquette, Olga Mariño, Delia Rogozan, and Michel Léonard. Competency-based personalization for massive online learning. *Smart Learning Environments*, 2(1):4, 2015.

[11] Zachary A Pardos, Steven Tang, Daniel Davis, and Christopher Vu Le. Enabling real-time adaptivity in moocs with a personalized next-step recommendation framework. In *Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale*, pages 23–32. ACM, 2017.

[12] María Luisa Sein-Echaluce, Ángel Fidalgo-Blanco, Francisco J García-Peñalvo, and Miguel Ángel Conde. imooc platform: Adaptive moocs. In *International Conference on Learning and Collaboration Technologies*, pages 380–390. Springer, 2016.

[13] D Shah. Year of mooc-based degrees: A review of mooc stats and trends in 2018. class central. *Class Central's MOOC Report*, 2019.

[14] Nishikant Sonwalkar. The first adaptive mooc: A case study on pedagogy framework and scalable cloud architecture—part i. In *MOOCs Forum*, volume 1, pages 22–29. Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA, 2013.

[15] Brian Subirana, Aikaterini Bagiati, and Sanjay Sarma. On the forgetting of college academice: at" ebbinghaus speed"? Technical report, Center for Brains, Minds and Machines (CBMM), 2017.

[16] Saijing Zheng, Mary Beth Rosson, Patrick C Shih, and John M Carroll. Understanding student motivation, behaviors and perceptions in moocs. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, pages 1882–1895. ACM, 2015.