Optimal and Non-Discriminative Rehabilitation Program Design for Opioid Addiction Among Homeless Youth

Amulya Yadav[1] , **Roopali Singh**[1] , **Nikolas Siapoutis**[1] , **Anamika Barman-Adhikari**[2] , **Yu Liang**[1]

[1]Pennsylvania State University, University Park, PA 16802
[2]University of Denver, Denver, CO 80208
[1]{amulya, rus82, nzs30, luy70}@psu.edu, [2]Anamika.BarmanAdhikari@du.edu

Abstract

This paper presents CORTA, a software agent that designs personalized rehabilitation programs for homeless youth suffering from opioid addiction. Many rehabilitation centers treat opioid addiction in homeless youth by prescribing rehabilitation programs which are tailored to the underlying causes of addiction. To date, rehabilitation centers have relied on ad-hoc assessments and unprincipled heuristics to deliver rehabilitation programs to homeless youth suffering from opioid addiction, which greatly undermines the effectiveness of the delivered programs. CORTA addresses these challenges via three novel contributions. First, CORTA utilizes a first-of-its-kind real-world dataset collected from ∼1400 homeless youth to build causal inference models which predict the likelihood of opioid addiction among these youth. Second, utilizing counterfactual predictions generated by our causal inference models, CORTA solves novel optimization formulations to assign appropriate rehabilitation programs to the correct set of homeless youth in order to minimize the expected number of homeless youth suffering from opioid addiction. Third, we provide a rigorous experimental analysis of CORTA along different dimensions, e.g., importance of causal modeling, importance of optimization, and impact of incorporating fairness considerations, etc. Our simulation results show that CORTA outperforms baselines by ∼110% in minimizing the number of homeless youth suffering from opioid addiction.

1 Introduction

Opioid addiction is a chronic disease which is characterized by a compulsive urge to use opioid drugs (even when they are no longer required medically), and it can cause devastating social, economic and health problems. The United States is in the midst of an opioid overdose epidemic (or "opioid crisis"), with ∼2.1 million cases reported in 2017 alone, leading to 47,600 deaths [CDC, 2018]. In particular, opioid addiction is highly prevalent among homeless youth, who often choose opioid drugs over other recreational drugs (e.g., methamphetamine, cocaine) as opioids are far cheaper and relatively easier to obtain [Fischer *et al.*, 2006]. In fact, previous studies show that the rates of opioid addiction among homeless youth are 7X higher than the general youth population [Hadland *et al.*, 2014]. Thus, any attempt at tackling the "opioid crisis" crucially depends on our success at minimizing rates of opioid addiction among homeless youth.

To tackle this problem, rehabilitation centers are tasked with prescribing and implementing COR-12 rehabilitation programs for addicted homeless youth, which represent the most commonly used method of treatment for opioid addiction [Hazelden, 2013]. In order to be effective, it is critical for COR-12 programs to be tailored to the needs of each patient, based on an evaluation of the underlying causative issue/problem that "caused" the opioid addiction. Therefore, in addition to prescribing long-term medication for treating addiction, these rehabilitation programs crucially include behavioral counseling and evaluation/treatment for the underlying causative issue/problem. For example, if poor financial condition (or unstable mental health) is the underlying reason behind the youth's opioid addiction, then career counseling (or psychiatric care) can be added into their rehabilitation plan. Table 1 shows several potential issues (e.g., poor mental health, low education levels, etc.) that commonly have a causative effect on opioid addiction among homeless youth.

Unfortunately, rehabilitation centers face several challenges in their work. First, it is very difficult for health workers at rehabilitation centers to uncover causal associations between potential issues/problems faced by a homeless youth and opioid addiction. Thus, while rehabilitation centers do deliver customized rehabilitation programs, they determine the causative issue/problem for each homeless youth in an ad-hoc manner, and there is little evidence to support the accuracy of their decision making. Second, most rehabilitation centers operate under manpower constraints (i.e., limited number of health workers) due to which strategic assignment of their limited resources to the correct set of homeless youth is critical, yet most centers do this assignment in an unprincipled manner, which leads to gross inefficiencies (we validate this in our experiments). Finally, in sensitive domains impacting life/death (like the ones that motivate this work), rehabilitation centers need to ensure that their decision making does not discriminate among people with respect to protected or sensitive characteristics such as race, ethnicity, disability, etc.

Thus, these challenges significantly limit the effectiveness of the centers rehabilitation efforts.

In this paper, we address these challenges by proposing CORTA (**C**omprehensive **O**pioid **R**esponse **T**ool Driven by **A**rtificial Intelligence), a novel software agent which optimizes the delivery of opioid rehabilitation services to homeless youth. CORTA collects data about opioid usage behaviors of homeless youth and uses that data to train high-dimensional causal inference models which can predict susceptibility of homeless youth to opioid addiction. CORTA specifically trains causal inference models due to the aforementioned need to uncover underlying reasons behind opioid addiction of homeless youth. Finally, using counterfactual estimates derived from these causal inference models, CORTA solves novel Integer Linear Program (ILP) formulations to determine the optimal assignment of homeless youth to the correct rehabilitation programs. CORTA's ILP formulation finds such optimal assignments by minimizing the expected number of homeless youth suffering from opioid addiction, while respecting fairness and limited capacity constraints faced by rehabilitation centers. Our simulation results show that CORTA improves upon existing state-of-the-art by resulting in $\sim$110% fewer homeless youth suffering from opioid addiction, and is robust to handling fairness considerations.

Related Work To the best of our knowledge, there is no prior work on using Artificial Intelligence (AI) techniques to minimize the likelihood of opioid (or any other drug) addiction among homeless youth. However, there has been a lot of recent interest in developing AI solutions for different problems faced by the homeless youth community. This line of work started with [Yadav *et al.*, 2016; Wilder *et al.*, 2017; Yadav *et al.*, 2015; Yadav *et al.*, 2017], who proposed and deployed AI algorithms for influence maximization to raise awareness about HIV prevention among homeless youth. Next, [Rahmattalabi *et al.*, 2019a] proposed algorithmic interventions on friendship based social networks of homeless youth to minimize substance abuse in this population. Similarly, [Rahmattalabi *et al.*, 2019b] proposed social network based interventions for suicide prevention among homeless youth. However, collecting accurate social network data for homeless youth is very difficult, due to which their approach may not always be applicable in the real-world. Finally, [Chan *et al.*, 2018; Azizi *et al.*, 2018; Kube *et al.*, 2019] proposed methods to assign different homeless services to homeless youth in order to minimize their probability of re-entry into homelessness.

2 Real World Dataset

CORTA was used to collect a first-of-its-kind dataset from $\sim$1400 homeless youth about their educational background, exposure to gang violence, mental health history, STI infection status, etc., along with their opioid misuse history. Using this data, CORTA builds causal inference models to predict likelihood of opioid addiction in homeless youth. These models allow CORTA to generate counterfactual estimates which are used inside novel ILP formulations which provide optimal assignment of rehabilitation programs to homeless youth. We now describe the generated real-world dataset in detail.

Data Collection The dataset was collected by surveying a total of 1426 homeless youth in the United States. Each homeless youth was asked to undertake a computer based survey questionnaire in which they were asked about their opioid use behaviors, i.e., they were asked whether they suffer from opioid addiction or not. In addition, each homeless youth was asked to report his/her sociodemographic characteristics (age, gender, ethnicity, sexual orientation, etc.), psychological characteristics (which evaluate their mental health), street victimization experiences (physical/sexual abuse and exposure to gang violence), and their sexual risk behaviors (e.g., their opinions about unprotected sex, their awareness about different STIs, and best practices). All procedures and data collection methods used in CORTA's data collection were reviewed and approved by an Institutional Review Board. Please refer to [Barman-Adhikari *et al.*, 2019] for more details on the data collection procedure.

Data Pre-Processing The data described above could not be used as-is with Machine Learning (ML) algorithms because of two issues. First, the collected data had lots of missing entries for several features, as homeless youth could choose to not answer survey questions that made them feel uncomfortable. Second, it is necessary to identify exogenous features in the data, as only these exogenous features can potentially be modified via targeted rehabilitation programs. Thus, while there could be other endogenous causative issues for opioid addiction, CORTA only focuses on modifying exogenous causative issues (or features).

After data collection, the raw dataset contained 1426 datapoints (one for each homeless youth). Each datapoint contained 212 features and a single binary label which denotes whether the homeless youth suffers from opioid addiction (1=yes) or not (0=no). Finally, CORTA uses MICE [Azur *et al.*, 2011], a state-of-the-art data imputation algorithm to infer missing feature values in this dataset.

In addition, we identify nine exogenous features in our dataset, each of which could represent a potential issue/problem that causes opioid addiction among homeless youth (and can be treated via rehabilitation programs). All of them were categorical in nature. For ease of exposition, we collapsed these nine categorical features into binary features using simple thresholding rules. Table 1 shows a list of these nine exogenous binary features in our dataset. At a high level, CORTA's goal is to find an optimal many-to-one matching between homeless youth and these exogenous features. We can then create custom rehabilitation programs for each homeless youth based on the matched exogenous feature in the optimal matching solution, i.e., if a homeless youth is matched with "*Mental Health Condition*" in the optimal solution, then psychiatric care is added into the youth's custom rehabilitation plan. CORTA finds this optimal matching solution by solving novel ILP formulations that we describe later in the paper.

After pre-processing, our dataset contained 1426 datapoints, and each datapoint consisted of 212 columns and a binary label. Our final dataset suffered from a class-imbalance problem, as 80% of the 1426 homeless youth did not suffer from opioid addiction. On the other hand, only 20% of our

ID	Causative Issue	Explanation of Corresponding Binary Feature in Dataset
1	Atttitude towards LGBTQ+	Youth is pro-LGBTQ+ (1=yes) or anti-LGBTQ+ (0=no)
2	Educational History	Youth has passed high-school (1=yes) or not (0=no)
3	Current Schooling Status	Youth currently attends school (1=yes) or not (0=no)
4	Foster Care History	Youth is in foster care (1=yes) or not (0=no)
5	Condom Accessibility	Youth can access condoms before sex (1=yes) or not (0=no)
6	Condom Usage Behavior	Youth uses condoms during sex (1=yes) or not (0=no)
7	PrEP Awareness	Youth has awareness about PrEP HIV medication (1=yes) or not (0=no)
8	Exposure to Street Violence	Youth has experienced street violence in the last month (1=yes) or not (0=no)
9	Mental Health Condition	Youth suffers from a mental health disorder (1=yes) or not (0=no)

Table 1: Potential causative issues/problems that can be the underlying cause of opioid addiction in homeless youth. Each of these nine causative issues corresponds to a binary exogenous feature in our homeless youth dataset. Note that while there could be other endogenous causative issues for opioid addiction, we only focus on these nine exogenous features (i.e., features which can potentially be modified via targeted rehabilitation programs) in our dataset.

homeless youth (datapoints) suffered from opioid addiction.

3 Causal Inference Models

Our next contribution is to build ML algorithms which can predict the likelihood of each homeless youth (datapoint) suffering from opioid addiction. These predictions can guide rehabilitation centers in choosing appropriate rehabilitation programs for at-risk homeless youth. Since these rehabilitation programs need to be tailored to the individual needs of each homeless youth (e.g., based on the underlying cause of their opioid addiction), it is necessary to build causal inference models (as opposed to traditional ML models) which can uncover hidden causal relationships between input features and response variables. An additional benefit of causal models is that they allow us to derive counterfactual estimates for each datapoint, which can be used to formulate Integer Linear Program (ILP) formulations for optimal rehabilitation delivery (described in the next section).

We use state-of-the-art models Bayesian Additive Regression Trees (BART) [Chipman *et al.*, 2010] and Honest Causal Forests [Wager and Athey, 2018] that can successfully generate heterogeneous treatment effect estimates (or counterfactual predictions) in addition to ATE (average treatment effects). Previous studies have shown that BART and Honest Causal Forests outperform several baseline algorithms (e.g., propensity score, nearest neighbor matching algorithms, etc.) for causal inference on complex observational data [Hill, 2011], hence we compare only among these two models.

Learning Results We compare the predictive performance of BART and Causal Forest against two standard non-causal ML models: (i) Lasso [Tibshirani, 1996]; and (ii) a C4.5 decision tree [Quinlan, 2014]. In order to train these ML models, we divide our dataset into training/test sets using a random 70:30 split. The hyper-parameters for all our models are optimized using K-fold ($K = 10$) cross-validation. The trained models are then used to get probabilities of opioid addiction for each homeless youth in the test set. Using these predicted probabilities and the ground-truth labels, we plot Receiver Operating Characteristic (ROC) curves to analyze the predictive performance of the various models.

None of our models overfit the data. For example, training/test AUROC (Area under ROC Curve) was 0.7765/0.6981 for BART and 0.7834/0.6729 for Decision Trees. The key takeaway is that all four models significantly outperform random classifiers (AUROC = 0.5). Out of the two causal inference models, BART (AUROC = 0.69) outperforms Causal Forest (AUROC = 0.67), hence we use BART as our model of choice inside CORTA. Also, this figure shows that BART, Lasso and C4.5 achieve similar levels of predictive performance, with AUROC values of 0.6981, 0.6839, and 0.6729 (respectively). While this suggests that standard ML models (e.g., C4.5 decision trees, Lasso, etc.) suffice for this problem, we illustrate the limitations of using standard ML models for optimal rehabilitation delivery in our experiments.

Uncovering Causal Relationships For the larger enterprise proposed in this work to make sense, it is important that different rehabilitation programs tailored according to different "underlying causes of opioid addiction" actually lead to different treatment effects on the homeless youth. For example, if homeless youth A is enrolled in a rehabilitation program R_1 tailored according to cause C_1 (as compared to program R_2 tailored for cause C_2), and all other aspects of the rehabilitation program (features) are left unchanged, then the likelihood of opioid addiction for youth A should be different when they are assigned in programs R_1 and R_2. Further, our model should also allow us to estimate the heterogeneous impact of a rehabilitation program (say R_1) across the homeless youth population (characterized by some features w that we explain in Section 4). Recall that Table 1 shows nine exogenous features present in our dataset, each of which represents a possible "underlying cause of opioid addiction". Accordingly, we assume that we have nine different rehabilitation programs. Next, we try to uncover causal relationships between these nine exogenous features in our dataset and the binary target variable (i.e., whether homeless youth suffers from opioid addiction or not). These uncovered causal relationships can be used by rehabilitation centers to select their key decision variable, i.e., selecting the "underlying cause of opioid addiction" in order to customize the rehabilitation program of each homeless youth.

Since our observational data comes by surveying homeless youth, it is susceptible to measurement errors in treatment assignments. Hence, differentiating between the efficacy of all

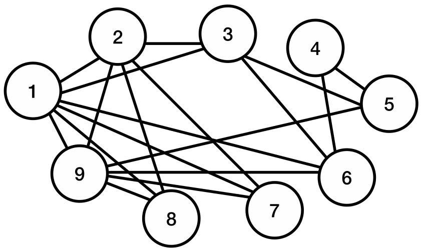

Figure 1: Pairwise Treatment Effects

rehabilitation programs together could amplify the bias (due to measurement endogeneity) and lead to misleading inferences. Thus, to compare the effects of different programs, we propose a novel pair-wise four-step identification strategy: (i) We compare these nine rehabilitation programs by doing pairwise inference, i.e., we look at differential treatment effects between pairs of rehabilitation programs (as opposed to considering heterogeneity of treatments in all nine rehabilitation programs at once). For each of the $\binom{9}{2}$ program pairs, we only keep the datapoints (from our full dataset) corresponding to homeless youth assigned to only one (or both) of these two programs (in the pair) and remove the remaining datapoints. We train a BART model based on this pruned dataset; (ii) Next, we use BART to approximate the counterfactual distribution of opioid usage based on this model for both the programs (i.e., estimate the heterogeneous treatment effect of both programs). At the end of this process, we have two counterfactual distributions, one for each rehabilitation program in the pair. (iii) We then conduct $t-test$ across these two distributions to find statistically different treatment effects for our two programs. First, we compute the difference across the treatment effects of the two programs (by computing a difference distribution that captures the difference between the two counterfactual distributions). Then, we take the mean and 2.5% and 97.5% quantiles of the difference distribution and reject the null (i.e., distributions are the same) if 95% credible region includes 0. (iv) Finally, this pairwise procedure is repeated for all $\binom{9}{2}$ pairs of rehabilitation program types in order to differentiate the effect of each rehabilitation program on homeless youths' opioid usage.

Out of $\binom{9}{2} = 36$ program pairs, we observe statistically significant ($p = 0.05$) population-wide treatment effects in 18 pairs. All these heterogenous treatment pairs are represented as a graph in Figure 1. There are nine nodes in Figure 1, one for each rehabilitation program. The 18 edges represent program pairs which exhibit population-wide treatment effects. For instance, we find that Programs 1 (Attitude towards LGBTQ+) and 4 (Foster Care History) do not have statistically different efficacy and hence, are represented by a missing edge between them. Due to lack of space, we move the actual 95% credible intervals for these 18 program pairs to the appendix[1]. We find that Program 1 (Attitude towards LGBTQ+) has a statistically stronger efficacy on opioid usage over the other six programs that it is connected to in Figure 1.

[1]http://amulyayadav.com/Papers/ijcai2020-appendix.pdf

4 Optimization of Rehabilitation Delivery

Rehabilitation centers have limited manpower to conduct their different rehabilitation programs, e.g., only a limited number of homeless youth can receive career counseling as part of their rehabilitation program, as the rehabilitation center has a limited number of career counselors. Thus, it is crucial for centers to select the right set of homeless youth to rehabilitate, and provide them with the correct rehabilitation program (based on an assessment of the underlying cause of their opioid addiction) for efficient utilization of their limited resources. As a result, CORTA looks at two operational objectives that rehabilitation centers care about. First, CORTA tries to maximize the number of youth suffering from opioid addiction (in the ground-truth test set) that are chosen by the center for rehabilitation, and are assigned to the correct rehabilitation program (as per counterfactual estimation). Second, CORTA also tries to minimize wastage of centers limited resources. To achieve these objectives, CORTA formulates the problem of optimal rehabilitation program delivery among homeless youth as an Integer Linear Program (ILP).

4.1 ILP Formulation

For clarity, we denote decision variables and constants in CORTA's ILP formulation with bold and plain symbols, respectively. Let N denote the total number of homeless youth (i.e., number of datapoints in our test set). Let M denote the number of available rehabilitation programs (we use $M = 9$ based on the nine exogenous features in our dataset). Let $C_j \; \forall j \in \{1, M\}$ denote the maximum capacity of rehabilitation program j, i.e., C_j denotes the maximum number of homeless youth that can be rehabilitated in program j. Let p_i^0 denote the "base" probability of homeless youth i suffering from opioid addiction (i.e., probability of addiction without the effect of any rehabilitation). Let p_{ij} denote the counterfactual probability of homeless youth i suffering from opioid addiction if they are assigned to rehabilitation program j. Note that p_i^0 and p_{ij} represent the output from running CORTA's trained BART model on the test set. Let $\boldsymbol{x_{ij}}$ be a binary decision variable that denotes whether homeless youth i is "rehabilitated" in program j ($\boldsymbol{x_{ij}} = 1$) or not ($\boldsymbol{x_{ij}} = 0$). Let $\boldsymbol{w_i}$ be a binary decision variable that denotes whether homeless youth i is chosen for rehabilitation by the center ($\boldsymbol{w_i} = 1$) or not ($\boldsymbol{w_i} = 0$).

Finally, let y_{ij} be a binary constant which denotes if homeless youth i already "enjoys" the effects of rehabilitation program j ($y_{ij} = 1$) or not ($y_{ij} = 0$), before the center chooses to rehabilitate that youth. Thus, $y_{ij} = 1$ denotes that the rehabilitation center does not need to assign homeless youth i to program j, as youth i already "enjoys" the effects of program j. For example, let rehabilitation program j be tailored to people suffering from mental depression. Thus, the direct effect of assigning program j to homeless youth i is that he/she is no longer depressed, hence the value of feature ID 9 (Table 1) for youth i is set to 0. However, if youth i is not depressed to begin with (i.e., feature ID 9 for youth i is already set to 0 in the test dataset), then he/she is said to already "enjoy" the effect of intervention j. In this case, we set $y_{ij} = 1$, otherwise $y_{ij} = 0$. The ILP formulation is given as follows:

$$
\begin{aligned}
\min_{\boldsymbol{x_{ij}},\boldsymbol{w_i}} \quad & \sum_{i=1}^{N}(1-\boldsymbol{w_i})p_i^0 + \sum_{i=1}^{N}\sum_{j=i}^{M}(1-y_{ij})p_{ij}\boldsymbol{x_{ij}} \\
\text{s.t.} \quad & \sum_{j=1}^{M}\boldsymbol{x_{ij}} \leq 1,\ \forall i \in 1\ldots N \\
& \sum_{j=1}^{M}\boldsymbol{x_{ij}} = \boldsymbol{w_i},\ \forall i \in 1\ldots N \\
& \sum_{i=1}^{N}\boldsymbol{x_{ij}} \leq C_j,\ \forall j = 1,\ldots,M \\
& \boldsymbol{x_{ij}} \leq 1-y_{ij},\ \forall i=1,\ldots,N, j=1,\ldots,M \\
& \boldsymbol{x_{ij}},\boldsymbol{w_i} \in \{0,1\},\ \forall i=1,\ldots,N, j=1,\ldots,M
\end{aligned}
$$

where the objective function minimizes the expected number of homeless youth suffering from opioid addiction. Note that this ILP is equivalent to minimal weighted assignment subject to two different capacity constraints, which is an NP-Hard problem (through a trivial reduction from the knapsack problem). *In practice, our ILP runs quickly – the largest experiments reported in Section 5 look less than one minute to solve*. We solve this ILP using Gurobi 8.1 to find optimal assignments of rehabilitation programs to homeless youth.

5 Experimental Evaluation

We evaluate the effectiveness of CORTA in a variety of settings. All our experiments were run on a 2.2 GHz Intel Core i7 machine having 16 GB of RAM. All experiments are averaged over 30 runs. All our experiments assume equal capacities for each rehabilitation program, i.e., $C_i = C\ \forall i \in \{1, M\}$. We use two metrics of comparison in this section: (i) number of opioid addicts in the test set who have been chosen for rehabilitation, and are assigned to the correct rehabilitation program (as per counterfactual estimation); and (ii) percentage of unused capacity of rehabilitation centers (or wastage). Throughout this section, the output from our trained BART model (i.e., base probabilites (p_i^0) and the counterfactual probabilities (p_{ij})) is assumed to be the ground truth probabilistic model of opioid addiction among homeless youth, and this is used to compute the expected number of opioid addicts in the population. This is because actual ground-truth heterogenous treatment effects of different rehabilitation programs on the likelihood of opioid addiction in homeless youth is not available in our dataset. Finally, all experiments are statistically significant under bootstrap-t ($\alpha = 0.05$).

We present results in three stages. First, we present results to illustrate the importance of causal modeling in CORTA, i.e., how do causal models improve over standard non-causal ML models. Second, we present results to illustrate the importance of ILP optimization in CORTA, i.e., how would our results differ if we used causal models, but did not use explicit ILP formulations to assign rehabilitation programs to homeless youth. Finally, we present results to show the impact of fairness constraints on the ILP optimization.

Baselines We propose several non-trivial heuristic baselines in this section. In order to evaluate the importance of causal modeling in our approach, we use the following heuristic method to assign rehabilitation programs to homeless youth: (i) we train a standard non-causal C4.5 decision tree on our training data, which outputs probabilities of homeless youth suffering from opioid addiction (for each youth); (ii) Let K denote the total capacity available to the rehabilitation center, i.e., $\sum_{j=1}^{M} C_j = K$. We select the homeless youth with the top-K probabilities of suffering from opioid addiction (as per the output of the C4.5 decision tree) for rehabilitation; and (iii) we select the rehabilitation programs for these K homeless youth. Since C4.5 decision trees do not provide counterfactual estimates, there is no way to choose appropriate rehabilitation programs for these K chosen homeless youth. Thus, we evaluate the best (*C4.5-Best*) and worst-case (*C4.5-Worst*) performance on our metrics with the K-chosen homeless youth. In *C4.5-Best* (*C4.5-Worst*), we select the rehabilitation program (for each homeless youth) with minimum (maximum) counterfactual probability as per the output of BART (we use counterfactual probabilities output by BART since we have no way to derive ground-truth counterfactual estimates). Note that choosing rehabilitation programs with minimum (maximum) counterfactual probabilities minimizes (maximizes) the ILP's objective function, hence these heuristics represent the best-case (worst-case) performance achievable with C4.5 decision trees.

In order to evaluate the importance of ILP optimization in our approach, we use two heuristic baselines. In both of our baselines, we assume that we have access to the output from CORTA's BART model in deciding assignments of programs to youth, however, we replace our explicit ILP formulation with heuristic assignments. For our first baseline (*BART-Min*), we select the homeless youth with the top-K probabilities of suffering from opioid addiction (as per the output of the BART model) for rehabilitation, and for each of these youth, we select the rehabilitation program with minimum counterfactual probability. For our second baseline (*BART-Max Difference*), we select the K-homeless youth who hold the greatest potential for "improvement", i.e., youth whose probabilities of suffering from opioid addiction can be decreased by the greatest amount (as a result of rehabilitating them in some program). More formally, we sort homeless youth on the basis of $\left(p_i^0 - \min_j p_{ij}\right)$ and pick the top-K youth from this sorted list, and assign rehabilitation programs to these K youth accordingly.

Importance of Causal Modeling We evaluate the performance of CORTA against *C4.5-Best* and *C4.5-Worst* in Figures 2(a), 2(b). The X-axis in these figures shows increasing capacity per rehabilitation program. The Y-axis in Figure 2(a) shows the number of opioid addicts (i.e., ground-truth label in test set is 1) who have been chosen for rehabilitation by the different assignment strategies. The Y-axis in Figure 2(b) shows the percentage of centers capacity unused by assignment of programs to youth. For example, when the capacity per intervention is 10, CORTA rehabilitated 27 opioid addicts and resulted in a wastage of 10% of the overall ca-

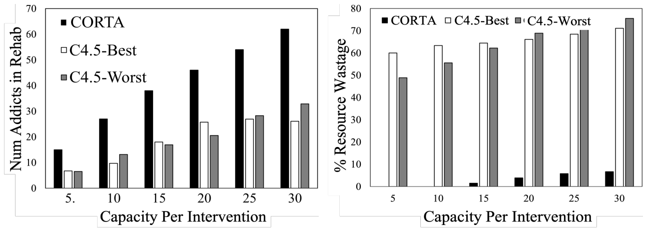

(a) Opioid Addicts in Rehab (b) % Resource Wastage

Figure 2: Comparing BART with C4.5 Decision Trees

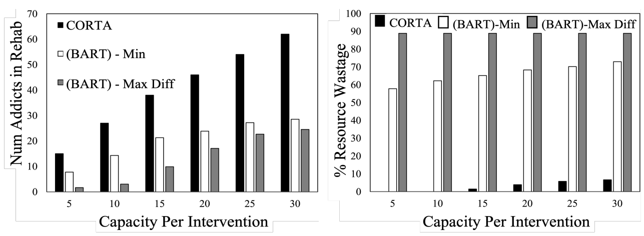

(a) Opioid Addicts in Rehab (b) % Resource Wastage

Figure 3: Importance of ILP Optimization

pacity K. Figure 2(a) shows that CORTA significantly outperforms *C4.5-Best* and *C4.5-Worst* by rehabilitating 110% more opioid addicts (on average). Figure 2(b) shows that CORTA leads to minimal ($\sim$4%) resource wastage as compared to *C4.5-Best* and *C4.5-Worst* ($\sim$66% wastage). Thus, Figures 2(a), 2(b) establish the importance of causal modeling in our approach, as it leads to significantly more ($\sim$2X) opioid addicts being rehabilitated while wasting a small fraction of resources in comparison (4% vs 66%).

Importance of Optimization Next, Figure 3(a) and 3(b) evaluate the performance of CORTA against *BART-Min* and *BART-Max Difference*. The X and Y-axes in Figures 3(a), 3(b) are defined as before. Figure 3(a) shows that CORTA significantly outperforms *BART-Min* and *BART-Max Difference* by rehabilitating $\sim$390% more opioid addicts (on average). Figure 3(b) shows that CORTA leads to minimal ($\sim$3%) resource wastage as compared to *BART-Min* and *BART-Max Difference* ($\sim$75% wastage). Thus, Figures 2(a), 2(b) and Figures 3(a), 3(b) establish the combined importance of causal modeling and ILP optimization in CORTA; both components are essential to achieving desired outcomes.

Impact of Fairness Considerations Each homeless youth in our dataset belongs to one or more protected (minority) categories. For example, homeless youth may belong to LGBTQ+, non-white, and/or female categories, etc. In real-world domains involving low-resource communities such as homeless youth, it is important to ensure that AI-driven algorithmic assignments of rehabilitation programs to homeless youth do not unfairly discriminate against any protected minorities in the population. CORTA's ILP framework allows us to explicitly mitigate against such concerns with the inclusion of proportional fairness constraints [Tsang *et al.*, 2019].

Let L be the total number of different protected categories

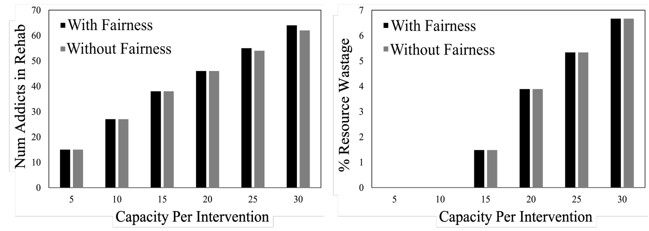

(a) Opioid Addicts in Rehab (b) % Resource Wastage

Figure 4: Incorporating Fairness Constraints

in our population. Let c_{il} be a binary constant which denotes whether homeless youth i belongs to protected category l (we get c_{il} values from our dataset). Let $\alpha_l \in [0,1]$ denote the fraction of homeless youth in our dataset that belong to class $l \in [1, L]$ - atleast an α_l fraction of the total capacity K must be chosen from protected category l in the output solution. Then, we can add explicit fairness constraints in CORTA's ILP framework as follows:

$$\sum_{i=1}^{N} c_{il} \sum_{j=1}^{M} \boldsymbol{x_{ij}} \geq \alpha_l K \; \forall l \in 1, \ldots, L \tag{1}$$

Figure 4(a) and 4(b) show the impact of adding fairness constraints on CORTA's solution quality (under varying capacities per intervention). To generate these figures, we restrict our attention to three protected minorities: gender, sexual orientation and race. The X and Y-axes in Figures 4(a), 4(b) are defined as before. These figures show that incorporating fairness constraints does not impact CORTA's effectiveness. In Figure 4(a), we observe minimal differences in the number of opioid addicts in rehabilitation between solutions with and without fairness constraints for all values of capacity per intervention. Further, Figure 4(b) shows no difference in resource wastage between solutions with and without fairness constraints regardless of capacity per intervention.

Surprisingly, this suggests that the "*price of fairness*" in the optimal rehabilitation delivery problem is negligible, i.e., incorporating fairness constraints in the problem do not lead to solutions with reduced effectiveness (at least until these constraints are satisfiable). Therefore, high-quality solutions do not need to be sacrificed in order to achieve fairness.

6 Conclusion

This paper proposes CORTA, an AI agent to assign appropriate rehabilitation programs to the correct set of homeless youth in order to minimize the expected number of homeless youth suffering from opioid addiction. We develop causal inference models that can predict the likelihood of opioid addiction among homeless youth. These causal inference models generate counterfactual treatment estimates, which are used to formulate novel ILP formulations for finding optimal assignments of rehabilitation programs to homeless youth. Our results show that CORTA outperforms non-trivial baselines by rehabilitating 110% more homeless youth suffering from opioid addiction, while ensuring minimal resource wastage.

References

[Azizi *et al.*, 2018] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 35–51. Springer, 2018.

[Azur *et al.*, 2011] Melissa J Azur, Elizabeth A Stuart, Constantine Frangakis, and Philip J Leaf. Multiple imputation by chained equations: what is it and how does it work? *International journal of methods in psychiatric research*, 20(1):40–49, 2011.

[Barman-Adhikari *et al.*, 2019] Anamika Barman-Adhikari, Hsun-Ta Hsu, Daphne Brydon, Robin Petering, Diana Santa Maria, Sarah Narendorf, Jama Shelton, Kimberly Bender, and Kristin Ferguson. Prevalence and correlates of nonmedical use of prescription drugs (NMUPD) among Young adults experiencing homelessness in seven cities across the United States. *Drug and Alcohol Dependence*, 200:153–160, 2019.

[CDC, 2018] CDC. Centers for Disease Control and Prevention wonder database. http://wonder.cdc.gov/, December 2018.

[Chan *et al.*, 2018] Hau Chan, Eric Rice, Phebe Vayanos, Milind Tambe, and Matthew Morton. From empirical analysis to public policy: Evaluating housing systems for homeless youth. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 69–85. Springer, 2018.

[Chipman *et al.*, 2010] Hugh A Chipman, Edward I George, Robert E McCulloch, et al. Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010.

[Fischer *et al.*, 2006] Benedikt Fischer, Jürgen Rehm, Jayadeep Patra, and Michelle Firestone Cruz. Changes in illicit opioid use across canada. *Cmaj*, 175(11):1385–1385, 2006.

[Hadland *et al.*, 2014] Scott E Hadland, Kora DeBeck, Thomas Kerr, Cindy Feng, Julio S Montaner, and Evan Wood. Prescription opioid injection and risk of hepatitis c in relation to traditional drugs of misuse in a prospective cohort of street youth. *BMJ open*, 4(7):e005419, 2014.

[Hazelden, 2013] Hazelden. Comprehensive Opioid Response with Twelve Steps (COR-12). https://www.hazelden.org/web/public/document/COR-12_Opioid-Response.pdf, October 2013.

[Hill, 2011] Jennifer L Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.

[Kube *et al.*, 2019] Amanda Kube, Sanmay Das, and Patrick J Fowler. Allocating interventions based on predicted outcomes: A case study on homelessness services. 2019.

[Quinlan, 2014] J Ross Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.

[Rahmattalabi *et al.*, 2019a] Aida Rahmattalabi, Anamika-Barman Adhikari, Phebe Vayanos, and Milind Tambe. Social network based substance abuse prevention via network modification (a preliminary study). In *Proceedings of the Strategic Reasoning for Societal Challenges (SRSC) Workshop*. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

[Rahmattalabi *et al.*, 2019b] Aida Rahmattalabi, Phebe Vayanos, Eric Rice, Anthony Fulginiti, Bryan Wilder, Amulya Yadav, and Milind Tambe. Exploring algorithmic fairness in robust graph covering problems. In *Proceedings of the 33th Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

[Tibshirani, 1996] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

[Tsang *et al.*, 2019] Alan Tsang, Bryan Wilder, Eric Rice, Milind Tambe, and Yair Zick. Group-fairness in influence maximization. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, IJCAI'19, pages 5997–6005. AAAI Press, 2019.

[Wager and Athey, 2018] Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

[Wilder *et al.*, 2017] Bryan Wilder, Amulya Yadav, Nicole Immorlica, Eric Rice, and Milind Tambe. Uncharted but not uninfluenced: Influence maximization with an uncertain network. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 1305–1313. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[Yadav *et al.*, 2015] Amulya Yadav, Leandro Soriano Marcolino, Eric Rice, Robin Petering, Hailey Winetrobe, Harmony Rhoades, Milind Tambe, and Heather Carmichael. Preventing hiv spread in homeless populations using psinet. In *Twenty-Seventh IAAI Conference*, 2015.

[Yadav *et al.*, 2016] Amulya Yadav, Hau Chan, Albert Xin Jiang, Haifeng Xu, Eric Rice, and Milind Tambe. Using social networks to aid homeless shelters: Dynamic influence maximization under uncertainty. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 740–748. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[Yadav *et al.*, 2017] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. Influence maximization in the field: The arduous journey from emerging to deployed application. In *Proceedings of the 16th conference on autonomous agents and multiagent systems*, pages 150–158. International Foundation for Autonomous Agents and Multiagent Systems, 2017.