Preventing Eviction-Caused Homelessness through ML-Informed Distribution of Rental Assistance

Catalina Vajiac[*] [1], **Arun Frey**[*2], **Joachim Baumann**[*3,4], **Abigail Smith**[*5],
Kasun Amarasinghe[1], **Alice Lai**[1], **Kit Rodolfa**[2], **Rayid Ghani**[1]

[1]Carnegie Mellon University,
[2]Stanford University,
[3]University of Zurich,
[4]Zurich University of Applied Sciences,
[5]NORC at the University of Chicago
cvajiac@cs.cmu.edu, arunfrey@law.stanford.edu, baumann@ifi.uzh.ch, als1@u.northwestern.edu, {kamarasi, alicelai, ghani}@andrew.cmu.edu, krodolfa@law.stanford.edu

Abstract

Rental assistance programs provide individuals with financial assistance to prevent housing instabilities caused by evictions and avert homelessness. Since these programs operate under resource constraints, they must decide who to prioritize. Typically, funding is distributed by a reactive or first-come-first serve allocation process that does not systematically consider risk of future homelessness. We partnered with Allegheny County, PA to explore a proactive allocation approach that prioritizes individuals facing eviction based on their risk of future homelessness. Our ML system that uses state and county administrative data to accurately identify individuals in need of support outperforms simpler prioritization approaches by at least 20% while being fair and equitable across race and gender. Furthermore, our approach would identify 28% of individuals who are overlooked by the current process and end up homeless. Beyond improvements to the rental assistance program in Allegheny County, this study can inform the development of evidence-based decision support tools in similar contexts, including lessons about data needs, model design, evaluation, and field validation.

1 Introduction

Homelessness remains a pervasive and pressing issue across the United States. In January 2022, more than 500k individuals experienced homelessness on a single night (de Sousa et al. 2022). Rising eviction rates and the lack of affordable housing contribute to this problem: as the gap between housing costs and income levels continues to grow, an increasing number of households struggle to pay rent, eventually facing eviction and, in some cases, homelessness.

To curb rates of homelessness, policymakers are increasingly seeking to reduce the rate of entry into homelessness (Culhane, Metraux, and Byrne 2011). One popular prevention strategy is rental assistance programs, which provide temporary financial assistance to individuals facing eviction to keep them stably housed. Experimental evidence suggests that such assistance is effective at reducing subsequent homelessness: using quasi-random variation in funding availability of a rental assistance program in Chicago, it was shown that individuals who called when funding was available were 76% less likely to become homeless in the subsequent 6 months compared to those who called when funding was unavailable (Evans, Sullivan, and Wallskog 2016). A similar program in NYC randomized the provision of financial assistance to households and reduced average days in shelter from 32 nights among the control group to 10 nights among the treatment group in the 2 years following funding (Rolston, Geyer, and Locke 2013).

For rental assistance programs to be a viable homelessness prevention strategy, they need to be effective (i.e. prevent individuals from falling into homelessness), efficient (i.e. target individuals at risk of falling into homelessness) (Burt 2007), and equitable. Currently, however, these programs prioritize individuals based on simple heuristics (e.g. first come, first served) instead of their likelihood of future homelessness. In addition, they place the burden of applying for funding on the individuals facing eviction and overlook those who need help but do not apply. The administrative process can create long delays, so that even eligible tenants often remain on a waitlist for multiple months and can end up evicted before they receive assistance.

In this paper, we describe our collaboration with the Allegheny County Department of Human Services (ACDHS) to improve the allocation of rental assistance by prioritizing individuals with the highest likelihood of falling into homelessness. We combine rich county and state administrative data with court records to predict homelessness among all individuals currently facing eviction, regardless of whether they contacted the county for financial support. Using data from January 2012 through August 2023, we develop a series of machine learning (ML) models that predict a tenant's need for homelessness services in the next 12 months, allowing ACDHS to proactively prioritize rental assistance to the most vulnerable tenants. In particular, we make the fol-

[*]These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

lowing contributions:

- **Enabling a proactive approach:** We consider a significantly higher proportion of residents, namely all tenants in the county facing eviction, instead of only those who call the county for help. Our models identify 28% of people who are overlooked by the current process and end up homeless. By shifting the burden away from those impacted, this proactive approach also reduces administrative effort and is more likely to provide preventative rental assistance to tenants before their eviction.
- **Implementing need-based prioritization:** Our models identify individuals in need of future homelessness support services with at least 20% improvement over simpler baselines, and are 10x better than random selection while also being fair and equitable.
- **Field validation:** We conduct a shadow mode deployment to validate our models on new data, mitigating the risk of leakage in the future. Additionally, we are planning a randomized control trial to compare our proposed solution to the current process and to evaluate the effectiveness of rental assistance in preventing entry into homelessness among targeted individuals.
- **Lessons learned:** We reflect on pitfalls and successes that may inform AI researchers seeking to ethically design predictive decision support tools in other contexts.

Reproducibility. Our code is available at https://github.com/dssg/acdhs_housing_public.

2 Ethical Considerations

Since we are informing the allocation of scarce and critical resources to vulnerable people who could become homeless, we bring into focus the ethical considerations that informed and were embedded in every phase of the scoping, design, and development of our approach.

Equitable outcomes through the allocation of resources. The use of AI in real-world contexts can perpetuate systemic biases such as racial disparities (Chen, Joshi, and Ghassemi 2020). To mitigate this risk, we carefully designed the scope of our work and our formulation and analyzed our model results to guard against bias against certain demographic groups, i.e. race and gender (see Section 5.3). Our field trial will also test the equity in the impact of the downstream decisions made using our system.

Transparency and interpretability. Our goal is to develop a system that supports and informs social workers in making better decisions that lead to improved, more equitable outcomes. To this end, we train models that are more explicit about which features they learn from (see Section 4.3) and analyze which features end up most predictive of future homelessness (see Section 5.4). Our solution is designed to augment the current process, supporting rather than excluding social workers from the decision-making loop and providing them with additional information beyond predictions to help them make better-informed decisions. We also consulted impacted community members in the design and validation process.

Privacy concerns of using sensitive data. The collection of personal information always comes with privacy concerns. In our case, particularly because ACDHS connects administrative data across different facets of residents' lives, we want to be sure that the use of this data offers significant enough improvements to the community, particularly people potentially facing homelessness. In Sections 5 and 6, we extensively evaluate our models with historical data, and a field trial to ensure that this work improves outcomes. We are currently soliciting feedback from groups potentially impacted by the system through a community engagement process to better assess tradeoffs before conducting a validation trial.

3 Current Approaches

We first consider how ACDHS currently prioritizes residents for rental assistance and highlight relevant prior research.

3.1 ACDHS' Current Process

The current process for obtaining rental assistance in Allegheny County is illustrative of how such programs are usually managed throughout the United States. Tenants with an eviction notice can contact the Allegheny County Link helpline to request rental assistance, or they may be referred to the rental assistance program by a mediation program or a housing assistance organization. Applicants are prescreened according to available funding and county-level eligibility requirements, including income and the amount of rent owed. Tenants are then placed on a waitlist and will be considered on a first-come-first-served basis. When the tenant reaches the front of the list, a social worker examines their funding request. If they are deemed eligible and can provide the required documentation, they will receive a payment covering the rent owed.

There are many issues with this reactive process, and the county is seeking our help in improving it. First, it puts heavy logistical strain on tenants facing eviction, who have to know to contact the helpline or other housing stability organizations to apply for assistance, and must be able to prove their eligibility with documentation. This requirement excludes all individuals at risk of homelessness due to eviction who do not proactively seek help. For those who apply, the substantial delay between application and receipt of funding means that many tenants have already been evicted by the time they are considered for assistance, or their rental debt has substantially increased. Since eligibility requirements change over time and are not always easily accessible, the current decision system is opaque to those seeking assistance. Finally, rental assistance is not distributed according to who has the most need, but instead to those who know to call the helpline, get through the waitlist in time, and meet the eligibility requirements before being evicted.

3.2 Related Work

Previous work has aimed to predict future homelessness to identify those most in need and more efficiently allocate homelessness prevention resources (Fowler et al. 2019). In New York City, re-evaluation of an existing homelessness prevention service resulted in a model based on 15 risk factors that would have outperformed human judgment in selecting at-risk individuals (Shinn et al. 2013). In another work focusing on those who were previously homeless, administrative data was used to better predict future homelessness within two years to best match people with programs, though this also disadvantaged other service recipients (Kube, Das, and Fowler 2019). Because of these tradeoffs, full automation of resource allocation is not recommended: human decision-makers should still be in control.

Work by the California Policy Lab predicted the risk of homelessness among all residents in Los Angeles County (Wachter et al. 2019). Several factors were found to be indicative of future homelessness, including repeated and recent interactions with government agencies or county services and prior mental health crises. Their predictive model achieves 29% precision for all residents.

Research Gaps. To our knowledge, no previous work has predicted future homelessness among those who are facing eviction. he California Policy Lab's work likely suffers from label bias due to the low percentage of homeless individuals utilizing county services (see Section 7), while our focus on informing a concrete action, rental assistance, allows our system to fit into the existing support infrastructure.

4 Predictive Approach

We develop ML models to facilitate need-based prioritization for proactive distribution of rental assistance. One limitation we encounter with this type of data is that it is conditioned on past rental assistance allocations and offers no insights into the counterfactual outcomes. As a result, our approach resembles an augmentation of the existing process. This is pertinent not only for the technical implementation but also for the interpretation of our results (e.g., in Section 5.2, we concentrate on the outcomes for vulnerable individuals who are not served under the current practice) and the conceptualization of the overall project, which is centered on social impact. and equity. In Section 6, we outline the transition from this stage to a fully deployed system.

4.1 Problem Formulation

We formulate our task as a binary classification problem: identifying tenants that will interact with homelessness services within 12 months of the date of prediction. At each prediction date, we consider individuals who have an active eviction filing against them within the past four months and are not currently homeless.[1] From this cohort, our models identify the 100 individuals with the highest chance of interacting with homelessness services in the next 12 months. We select 100 individuals since this corresponds to the monthly intervention capacity of ACDHS. Specific inclusion and exclusion criteria are detailed in Figure 1.

4.2 Data and Feature Engineering

ACDHS collects and combines administrative data from various county and state-level programs, including information

[1]More details on the demographic composition of this group are provided in Supplementary Materials B.

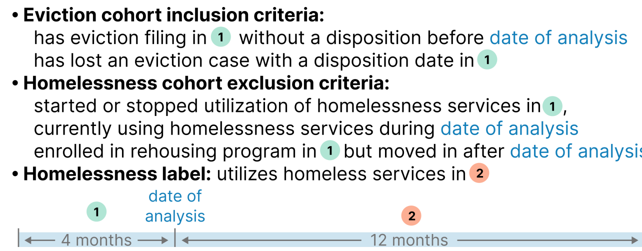

Figure 1: Inclusion/exclusion criteria for homelessness and eviction. Indicators for homelessness in ACDHS data include clients' interactions with homelessness services, such as staying in a shelter. We also include clients enrolled in rehousing programs that have not moved in as of the prediction date, as being homeless is a prerequisite for enrollment.

on individuals who have previous eviction filings, homelessness spells, interactions with mental, behavioral, and physical health institutions, address changes, or who have been enrolled in a variety of other ACDHS and state programs.

For each client at each analysis date, we generate $\sim$7000 features based on the data sources described in Supplementary Materials A. The features can be classified as follows:

Demographics and Event features. For each tenant, we include the total number of interactions with each ACDHS and state program, total number of evictions, and total number of physical and mental health visits. We also include categorical features, such as the particular type of physical or mental health visit and demographic features (race, gender, and age).

Temporal aggregation features. For each data source, we generate several temporal features to capture the dynamic nature of the process and assess how frequently an individual interacts with state and local services within the last 3 & 6 months, and 1, 2, 3, 4, and 5 years. For example, using the eviction data, we generate the following features per individual: number of days since most recent eviction; number of evictions in the specified time period; sum, min, max, average rent owed in eviction cases; min, max, average interarrival times between evictions.

4.3 Model Training and Validation

We use various supervised classification methods to predict entry into homelessness: logistic regression (LR), decision trees (DT), random forests (RF), Adaboost, Light Gradient Boosting Model (LGBM), and XG boost using Scikit-Learn (Pedregosa et al. 2011) and the hyperparameter grid specified in Supplementary Materials C. To most closely reproduce the context in which our models will be deployed, we use temporal validation (Hoptroff 1993), generating different training and evaluation matrices for each analysis date. For example, when evaluating the efficacy of predicting homelessness as of January 2019, we train models on feature label pairs using data up to January 2019, and evaluate them based on how many people become homeless between January 2019 and January 2020 (see Supplementary

Materials D). Temporal validation requires that no information beyond a certain date is used to evaluate an algorithm to avoid data leakage, i.e., predicting the future using data from the future, which can be challenging when combining real-world data from disparate sources that are updated at different frequencies (see Section 8 for a full discussion).

4.4 Baseline Models

We compare the ML models to several simple baselines that either attempt to approximate ACDHS' current system of allocating rental assistance, or that provide simple improvements on the current status quo. Baselines represent simple heuristics that do not require implementing an ML model (i.e., ranking individuals based on a single attribute) and are, therefore, much easier to deploy and understand.

B1. **Previous homelessness.** Prior homelessness is a strong indicator of future homelessness (Glendening and Shinn 2017). With this in mind, ACDHS could prioritize clients by the last date they interacted with any homelessness service. Here, a more recent date would imply a tenant is at higher risk of future homelessness.

B2. **Baserate.** If ACDHS were to randomly select individuals to give rental assistance, the precision of the approach would be equal to the proportion of individuals in our cohort who become homeless in the next year, around 2% during our period of analysis.

B3. **Earliest OFP.** As an approximation of the current process' first-come-first-serve waitlist, we look for tenants who have been waiting for the longest since an *Order for Possession (OFP)* has been granted, which allows the landlord to evict the tenant.

Other baselines were omitted from these results due to poor performance (see Supplementary Materials E for a full list).

5 Key Findings

We trained and validated over 5000 model variants to select the 100 individuals with the highest score, i.e. the highest risk of falling into homelessness within the next 12 months. The model selection metrics are measured at top-k, where $k = 100$.

5.1 Predictive Models Are More Efficient and Effective than Heuristic Baselines

We evaluate the performance of predictive models based on their ability to improve the efficiency (i.e., ensuring assistance goes to those who are most in need) and effectiveness (i.e., maximizing the reach to people who would otherwise become homeless) of the allocation of rental assistance resources while satisfying the equity constraints. Efficiency and effectiveness are measured by precision@100 and recall@100 respectively. For each model, we calculate these metrics across all temporal validation splits and explore a broad hyperparameter space for each model type. Table 1 summarizes the metrics over time. We primarily focus on the average values of both metrics as we aim to select models that will generalize well to the future. Based on guidance from our partners at ACDHS, we exclude validation

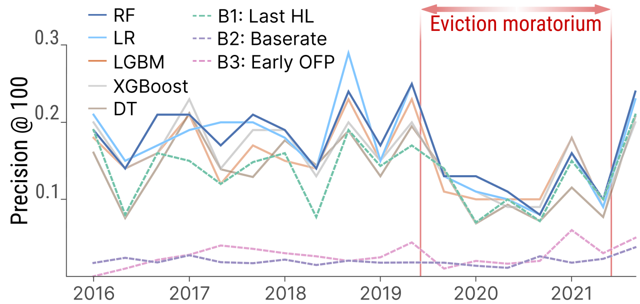

Figure 2: *LR and RF outperform all baselines:* Precision@100 over time shows that out of all baselines, *B1: Previously Homeless* performs best, but LR and RF perform better for all splits outside of the moratorium.

splits that coincide with the COVID-19 eviction moratorium period when measuring average precision[2]. Figure 2 shows precision@100 over time for the best hyperparameter configuration for each model type.

We observe that RF and LR perform best and outperform simple heuristic baselines with respect to precision@100, showing a ∼20% improvement over the best heuristic, *B1*, and performing 10x better than our approximation of the current allocation process, *B3*.

Model type	Precision@100 Avg	Min	Max	Recall@100 Avg	Min	Max
RF	**0.20**	**0.14**	0.25	**0.22**	0.16	**0.34**
LGBM	0.18	0.12	0.23	0.19	0.14	0.33
LR	**0.20**	**0.14**	**0.29**	**0.22**	**0.17**	0.33
DT	0.16	0.08	0.21	0.18	0.09	0.30
XGBoost	0.18	0.13	0.23	0.20	0.16	0.29
B1: Prev. HL	0.15	0.08	0.21	0.17	0.09	0.30
B2: Baserate	0.02	0.01	0.04	0.02	0.02	0.05
B3: Early OFP	0.03	0.00	0.05	0.03	0.00	0.07

Table 1: *RF and LR outperform:* Precision and recall over time for the model types considered show that our models outperform even the best baseline, *B1: Previously Homeless*.

5.2 Predictive Models Identify People Who Are Overlooked by the Current Process

The risk in relying on tenants to proactively apply for rental assistance is that the people most in need will not receive it and will end up homeless. In each cohort, about 75 people become homeless within a year of their eviction filing. Notably, a majority of these people (50 on average) do not apply for help. This *missed group*—those who interacted with the homelessness system within 12 months of the prediction date, did not apply for help, and did not receive rental

[2]During this time, there are very few people facing eviction due to the moratorium on evictions during COVID-19 so our cohort becomes very small. See Supplementary Materials F.

assistance—illustrates the limitations of the current reactive allocation process.

In Figure 3, we examine what percentage of the *missed group* is found by each model, comparing one of our best ML models (RF) to the baselines *B1: Previously Homeless* and *B2: Baserate*. *B2* finds 4% of this group on average, meaning that proactive outreach to a list of 100 people from this random approach would find about 4 people every month who would be overlooked by the current rental assistance practice and would end up homeless. Both *B1* and RF reach a substantially larger group of residents, 23% for *B1* and 28% for RF on average, potentially providing rental assistance to an additional 10–20 people from the *missed group* in each cohort every month.

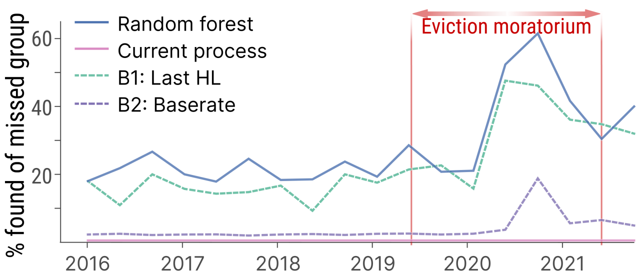

Figure 3: Percentage found of *missed group*, the individuals who become homeless having not contacted Allegheny County Link or received rental assistance.

5.3 Predictive Models Can Promote Fairness and Equity

Efficiency and effectiveness metrics do not reveal whether allocating resources according to these models would be biased against vulnerable groups (Barocas, Hardt, and Narayanan 2023). Since we are informing the allocation of a scarce resource, we use *equality of opportunity* (Hardt, Price, and Srebro 2016) as the fairness principle, which is captured using the true positive rate (TPR) of vulnerable subgroups in the top 100. We consider the following groups:

Race. Black individuals are at higher risk of falling into homelessness in Allegheny County. Out of those that will become homeless in our validation cohorts, on average, 60% were Black, even though Black individuals make up only 14% of the county's population (U.S. Census Bureau 2022). To mitigate these disproportionate impacts, ML models should have a higher TPR compared to white individuals (i.e., serving a higher proportion of Black individuals who are actually at risk of homelessness). We calculate $\frac{P(D=1|Y=1,A=black)}{P(D=1|Y=1,A=white)}$ where D, Y, and A represent the decision, outcome, and attribute of interest (e.g., race) respectively, finding average recall ratios of 1.34 and 1.14 respectively across temporal splits, indicating that our models meet our fairness considerations for race.

Gender. For gender, we consider men and women since data on non-binary and transgender individuals is not well-documented in ACDHS data. We calculate $\frac{P(D=1|Y=1,A=female)}{P(D=1|Y=1,A=male)}$ for our best models and see that RF and LR models are slightly under-serving women, with average TPR ratios of 0.9 and 0.87 respectively. This issue needs to be considered further in the resource allocation process.

5.4 Prior Homelessness and Mental Health Crises Contribute to Homelessness Risk

We find that previous use of homelessness services and interactions with mental and behavioral health services are consistently predictive of future homelessness spells across all validation temporal splits. Most predictive features include:

Previous homelessness service utilization. The number, duration, and recency of past homelessness spells are highly predictive of future homelessness. Past referrals to homelessness services, emergency shelter utilization, and public housing utilization are identified as highly predictive. Interestingly, the most predictive feature was the number of days since the last homelessness spell (baseline *B1*).

We also compared the characteristics of the top 100 individuals versus the rest of the tenants. We found that for the validation split starting on Sept. 1, 2021, the top 100 individuals were *34x more likely to have been in an emergency shelter* in their lifetime, 31x more likely to have been homeless, and likely to have spent 28x as many days in homelessness compared to other tenants.

Mental and behavioral health service interaction. Mental and behavioral health events and related service utilization are highly predictive of future homelessness. In the best performing RF, the number of days since the last mental health or behavioral health crisis event, the duration of mental health service utilization, the number of times one used mental health services, and the number of mental health crisis events were highly predictive features.

Compared to the rest of the cohort, those in the top 100 were *100x more likely to have had a mental health crisis event* in the last three years and 34x more likely to have had one in their lifetime, likely to have had 28x as many days utilizing mental health services, and likely to have had 24x more behavioral health events in their lifetime.

5.5 First-Time Homelessness is Harder to Predict

About half of the people in a cohort who become homeless have not previously been homeless, but we find that our models rely on features about previous homelessness spells, resulting in lower recall for people experiencing first-time homelessness. Our models perform substantially better on people with a history of homelessness: RF's average recall (excluding the eviction moratorium) is 55% for people who have experienced previous homelessness but only 4% for people experiencing first-time homelessness. We are exploring 1) building separate models, and 2) adding additional data sources to better predict future homelessness based on whether or not the person has previously been homeless.

Interestingly, the group who receives rental assistance through the existing practice appears to mostly be people without a history of homelessness: for the validation split starting on Sept. 1, 2021, over 70% of people in our cohort who apply for help and over 80% of people who receive rental assistance have not been homeless before.

5.6 On False Positives: Models Identify Vulnerable Individuals

Out of those facing eviction in Allegheny County, about 2% end up interacting with homelessness services within 12 months. With such heavy class imbalance, it is inevitable that we have false positives in our top-k, i.e. recommend giving rental assistance to those who will not become homeless within 12 months. Who are these false positives? Are they vulnerable to other adverse outcomes, including homelessness beyond 12 months or other crises?

In Figure 4, we investigate the top 100 picked by the RF model as of May 1, 2019 as it provides a sufficient temporal gap from the present to observe long-term outcomes. We see that even though the predictions are "wrong" about who will interact with homeless services within 12 months, they highlight people vulnerable to homelessness and other adverse outcomes and in need of assistance.

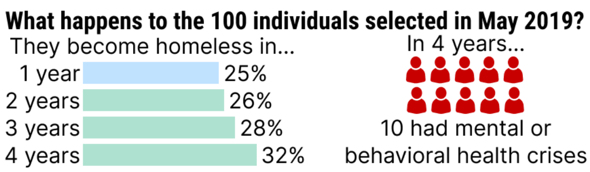

Figure 4: *False positives are still vulnerable:* among these tenants, we see that homeless service utilization and mental or behavioral health crises are common beyond 1 year.

6 Field Validation

Before deployment, we validate our results in two consecutive stages: first, shadow mode deployment (SMD) alongside the current process to validate our solution on live data, and second, a planned randomized control trial (RCT) to compare our proposed solution to the status quo and assess the effectiveness of the treatment (rental assistance).

6.1 Shadow Mode Deployment (SMD)

To validate our model against real-time data with no chance of data leakage (see Supplementary Materials A), we use our model to make predictions in real time, while the decision-making process continues to use the current system. We used the best-performing model that met our equity goals as of September 1, 2022, and trained it on all data available at the time to produce a list of 100 individuals at risk of falling into homelessness within 12 months. We then compared this list to the current rental assistance distribution process.

Between September 2022 and August 2023, 22 of the 100 individuals on the predicted list made use of homelessness services, confirming our model's precision@100 of $\sim$0.20. There is little overlap between individuals targeted by the current approach and those that would have been targeted using our predictive model: among our list of 100 individuals, only 12 received rental assistance.[3] Importantly, our model would have resulted in proactive assistance to 17 people who are *missed* by the current system, i.e. to those who did not reach out for assistance, did not receive rental assistance, and ended up falling into homelessness within a year.

6.2 Randomized Control Trial (RCT)

To accurately assess the *efficiency* of both the current and the ML-based process, we must compare rates of homelessness among individuals who received rental assistance to the counterfactual case where such individuals would not have received assistance. Similarly, to assess the *effectiveness* of rental assistance in reducing homelessness, we must compare rates of homelessness between those that did and did not receive assistance, all else being equal. Such causal identification is not possible through historical observational data.

To compare the impact of the current and proposed allocation process, we are planning a randomized control trial (RCT). In principle, an RCT should:

1. select two sets of eligible candidates, one using the current approach and one using the algorithmic approach,
2. randomly assign $k\%$ individuals to a treatment and $(1-k)\%$ to a control group for each candidate set.

To compare the *efficiency* of our new system to that of the current system, it would then be possible to compare homelessness rates among individuals selected by both approaches who were assigned to the control group (those who did not receive rental assistance), i.e. $P(Y=1|D_{ML}=0) - P(Y=1|D_C=0)$, where $D_{ML}=0$ and $D_C=0$ denotes random assignment to the control group in the algorithmic (D_{ML}) and current list (D_C). Similarly, to assess the *effectiveness* of rental assistance at reducing homelessness, it would be possible to compare rates of homelessness among those in the treatment to rates of homelessness among those assigned to the control group (i.e. $P(Y=1|D_{ML}=1) - P(Y=1|D_{ML}=0)$ and $P(Y=1|D_C=1) - P(Y=1|D_C=0)$).[4]

There are ethical challenges involved with randomizing the allocation of rental assistance. In particular, this RCT would entail not providing assistance to eligible people who applied for funding and who would have received it in the absence of the RCT. Instead, we could rely on a *quasi-random* assignment of individuals to the control group, as some happen to call the helpline on days when the waitlist is full or funding is not available (see Evans, Sullivan, and Wallskog 2016). This would allow us to construct a counterfactual group of eligible individuals who do not receive funding without interfering with the current process. We are currently discussing the RCT design options with ACDHS and plan to finalize them soon.

[3]Of those 12, 4 used homelessness services within 12 months. Assuming rental assistance is effective, our SMD precision@100 of 0.22 represents a lower bound, but this will have to be tested in our RCT (see Section 6.2).

[4]Note that the effectiveness of the treatment may vary between people targeted by the current process and those targeted by the ML-based solution, and even different ranges of risk scores between the two. These potential differences should also be assessed.

7 Limitations

Eligibility Constraints. We decided not to consider eligibility requirements for the following reasons. There are multiple sources of eligibility requirements—those tied to funding sources and the internal ones set by ACDHS—which change over time and are not well-documented. Not only is it impossible to encode historical eligibility requirements given the lack of data, but we also want to alert ACDHS to individuals with the most need, regardless of their eligibility.

Label Bias. Our labels are likely biased due to *outcome measurement errors*, which are common in applying predictive modeling approaches to observational data (Coston et al. 2020; Jacobs and Wallach 2021). Utilization of homelessness services is just a *proxy*, as not all who fall into homelessness use these services. Individuals who are couch surfing, sleeping in their cars, or unsheltered on the street will not be captured by our outcome definition (label).

In Allegheny County, the large majority of homeless individuals use homelessness services, such as emergency shelters,[5] at some point during their homelessness spell, particularly in winter. Studies conducted in warmer climates, however, are likely to systematically miss those who never make use of any homelessness services (e.g. Wachter et al. 2019). By selecting 12 months as our label period, we guarantee that our label captures shelter use during the winter months.

8 Lessons Learned

Throughout this process, we've learned several lessons about how to most effectively use AI methods in a real-world, resource-constrained context, many of which generalize beyond the allocation of rental assistance.

Scoping: the underrated first step. What problem actually *needs* solving? What will *actually* be used in practice? While these questions may seem like the obvious first step, a common question instead asked by many is "what are interesting research problems we can explore with this data?," regardless of their real-world impact. We spent months scoping (Data Science for Social Good 2023) the project with domain experts focusing on societal goals and actions we can inform until ACDHS' specific needs were clear. Only then did we formulate that need as a modeling problem and explore ML-based approaches.

Data leakage: the secret deceptor. What data will be available at the prediction date, when decisions are made? Real-world data is messy: columns are populated on different dates and updated at different intervals. Unwittingly using data that was updated after the prediction date can artificially inflate performance but data leakage can be difficult to detect. Early on, we discovered that the feature denoting if a resident's age was imputed was a strong indicator of future homelessness. Upon further investigation, we found that the source containing age was updated in place, meaning that if a person had interacted with a program that recorded age *after* the prediction date, it would be reflected in our feature. It became a signal for homelessness because they had interacted with social services after the prediction date. Unknowingly, we had run into data leakage.

Field trials: the reality check for ML's social impact. Often, a field trial is necessary to construct counterfactual outcomes and compare the effectiveness and efficiency of the proposed solution to that of the current process. However, randomization can be ethically challenging as it involves not providing help to otherwise eligible individuals in need. In such cases, researchers should aim to construct counterfactual groups in a manner that minimizes interference with the current allocation process if feasible.

Evaluation for real-world impact. Metric selection and validation techniques require careful consideration of the real-world problem. Standard evaluation methods, such as k-fold CV and AUC, may not accurately capture the success of a model for the given task. We started by defining our specific goals, e.g. efficiency, effectiveness, and equity, and then found the appropriate metric for each goal. In resource-constrained policy problems like ours, precision@k represents an appropriate efficiency metric, while AUC or F1 scores are less meaningful.

Communicating to policymakers. Policymakers may not have the technical expertise to decipher ML models or metrics. We need to assist them in making informed decisions, e.g. by providing intuitive explanations of metrics and a palette of model options with their associated policy trade-offs. While this can be time-consuming, the alternative can lead to model misuse or output misinterpretation.

9 Conclusion

We have shown that predictive modeling can improve upon the prioritization of rental assistance to tenants facing eviction in Allegheny County to reduce the rate of entry into homelessness. Our novel approach has four main contributions, providing (a) need-based prioritization of rental assistance (b) in a proactive manner, which is at least 20% more effective than simpler baselines while being equitable. We (c) validated our models by deploying them as a shadow model and are designing an RCT that is being discussed with ACDHS. We also (d) included some lessons learned that other AI researchers can use to ethically design predictive support tools. After the RCT, this process is expected to be used in practice, improving the effectiveness of allocation of rental assistance and equity in outcomes in Allegheny County.

Acknowledgements

This work was started as part of the 2022 Data Science for Social Good (DSSG) Fellowship at Carnegie Mellon University and partly funded by a grant from the Richard King Mellon Foundation. We thank Adolfo De Unanue for his valuable input during DSSG. We also thank our partners at ACDHS, particularly Rachel Rue and Justine Galbraith, for their help throughout this project.

[5] ACDHS' most recent point in time count, where workers manually count all homeless individuals in the county on one day, found that 83% of homeless individuals were staying in shelters that day (Allegheny County Department of Human Services 2023).

References

Allegheny County Department of Human Services. 2023. Point-In Time Count of People Experiencing Homelessness. https://analytics.alleghenycounty.us/wp-content/uploads/2023/05/23-ACDHS-04-PIT-Brief_v7.pdf.

Barocas, S.; Hardt, M.; and Narayanan, A. 2023. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press.

Burt, M. R. 2007. *Homelessness: Prevention Strategies and Effectiveness*. Nova Science Publishers.

Chen, I. Y.; Joshi, S.; and Ghassemi, M. 2020. Treating health disparities with artificial intelligence. *Nature medicine*, 26(1): 16–17.

Coston, A.; Mishler, A.; Kennedy, E. H.; and Chouldechova, A. 2020. Counterfactual Risk Assessments, Evaluation, and Fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 582–593.

Culhane, D. P.; Metraux, S.; and Byrne, T. 2011. A Prevention-Centered Approach to Homelessness Assistance: A Paradigm Shift? *Housing Policy Debate*, 21(2): 295–315.

Data Science for Social Good. 2023. Data Science Project Scoping Guide. https://www.datasciencepublicpolicy.org/our-work/tools-guides/data-science-project-scoping-guide/.

de Sousa, T.; Andrichik, A.; Cuellar, M.; Marson, J.; Prestera, E.; Rush, K.; and Abt Associates. 2022. The 2022 Annual Homelessness Assessment Report (AHAR to Congress) Part 1: Point-In-Time Estimates of Homelessness, December 2022. *The U.S. Department of Housing and Urban Development, Office of Community Planning and Development*, 1–108.

Evans, W. N.; Sullivan, J. X.; and Wallskog, M. 2016. The impact of homelessness prevention programs on homelessness. *Science*, 353(6300): 694–699.

Fowler, P. J.; Hovmand, P. S.; Marcal, K. E.; and Das, S. 2019. Solving Homelessness from a Complex Systems Perspective: Insights for Prevention Responses. *Annual Review of Public Health*, 40(1): 465–486.

Glendening, Z.; and Shinn, M. 2017. Risk Models for Returns to Housing Instability Among Families Experiencing Homelessness. *Cityscape (Washington, D.C.)*, 19(3): 309–330.

Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, 3323–3331.

Hoptroff, R. G. 1993. The Principles and Practice of Time Series Forecasting and Business Modelling Using Neural Nets. *Neural Comput. Appl.*, 1(1): 59–66.

Jacobs, A. Z.; and Wallach, H. 2021. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 375–385.

Kube, A.; Das, S.; and Fowler, P. J. 2019. Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 622–629.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.

Rolston, H.; Geyer, J.; and Locke, G. 2013. Evaluation of the Homebase Community Prevention Program. *NYC Department of Homeless Services*.

Shinn, M.; Greer, A. L.; Bainbridge, J.; Kwon, J.; and Zuiderveen, S. 2013. Efficient Targeting of Homelessness Prevention Services for Families. *American Journal of Public Health*, 103(S2): S324–S330.

U.S. Census Bureau. 2022. 2022 Census. U.S. Department of Commerce.

Wachter, T. V.; Bertrand, M.; Pollack, H.; Rountree, J.; and Blackwell, B. 2019. Predicting and Preventing Homelessness in Los Angeles.