

Beyond Eviction Prediction: Leveraging Local Spatiotemporal Public Records to Inform Action

Tasfia Mashiat
George Mason University
Fairfax, VA, USA
tmashiat@gmu.edu

Patrick J. Fowler
Washington University in St. Louis
St. Louis, MO, USA
pjfowler@wustl.edu

Alex DiChristofano
Washington University in St. Louis
St. Louis, MO, USA
a.dichristofano@wustl.edu

Sanmay Das
George Mason University
Fairfax, VA, USA
sanmay@gmu.edu

ABSTRACT

There has been considerable recent interest in scoring properties on the basis of eviction risk. The success of methods for eviction prediction is typically evaluated using different measures of predictive accuracy. However, the underlying goal of such prediction is to direct appropriate assistance to households that may be at greater risk so they remain stably housed. Thus, we must ask the question of how useful such predictions are in targeting outreach efforts – informing action. In this paper, we investigate this question using a novel dataset that matches information on properties, evictions, and owners. We perform an eviction prediction task to produce risk scores and then use these risk scores to plan targeted outreach policies. We show that the risk scores are, in fact, useful, enabling a theoretical team of caseworkers to reach more eviction-prone properties in the same amount of time, compared to outreach policies that are either neighborhood-based or focus on buildings with a recent history of evictions. We also discuss the importance of neighborhood and ownership features in both risk prediction and targeted outreach.

CCS CONCEPTS

• **Computing methodologies** → **Model development and analysis**.

KEYWORDS

Eviction Prevention, AI for Social Good, Social Applications

ACM Reference Format:

Tasfia Mashiat, Alex DiChristofano, Patrick J. Fowler, and Sanmay Das. 2024. Beyond Eviction Prediction: Leveraging Local Spatiotemporal Public Records to Inform Action. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3630106.3658978>



This work is licensed under a Creative Commons Attribution International 4.0 License.

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658978>

1 INTRODUCTION

One of the standard methodologies for using AI in the social service provision, education, and criminal justice domains is for algorithms to generate risk scores used to direct prioritization metrics or as additional information in bureaucratic decision-making procedures. Examples include the investigation of possible child maltreatment or neglect [6], homelessness resource provision [2, 26], sentencing decisions [35], and targeting interventions to students at risk of dropping out of school [33].

A pressing unresolved question is how these possible applications of AI should be evaluated. First, it is important to compare with the bureaucratic counterfactual [22, 34] as well as to understand the impact of how information is presented to and used by the human-in-the-loop [27, 31]. Even beyond this, though, there is a question of whether improvements in the oft-studied accuracy of the predictive model [39, 45] are truly useful in decision-making [43]. For example, recent work in the context of education suggests that even if a model accurately sorts students in terms of their dropout risk, it may provide no benefit in targeting interventions over a model that only uses information about students' environments [33]. However, the contextual and institutional specifics may be important. Our goal in this paper is to study the problem of targeting interventions based on risk scoring in the context of tenant eviction.

We aim to assess and inform data-driven targeting of assistance to renters at greater risk for eviction. The goal is to keep these renters stably housed. Currently, most communities target tenant outreach at the neighborhood level based on historical eviction rates (where available) or geographic characteristics (e.g., zip code poverty levels). Few communities use dynamic spatiotemporal information or predictive modeling for identifying at-risk properties. Prior research on eviction prediction has used less precise data (e.g., zip code versus property-level prediction, lack of property ownership profiles) and has been measured only in terms of predictive accuracy rather than effectiveness in designing outreach policies [39].

In this work, we leverage a unique community-based research collaboration to build predictive models of eviction and then use these to design and analyze potential outreach policies. Our results also help us better understand the importance of different types of features in prediction and outreach design to keep tenants housed.

Contributions: Our main contribution is an end-to-end methodology that integrates information across diverse societal agencies, utilizes the information for eviction risk prediction, and delegates decision-making on canvassing risk-prone properties. Our community partners have aggregated public and proprietary historical information on evictions (e.g., parties, dates of filing and judgment, etc.) and building information (e.g., units, owner, etc.) across all properties in St. Louis, Missouri, USA. Matching owner addresses from tax assessments allows us to create unique owner profiles (e.g., business status, owner location, number of units, etc.) unavailable in most communities. The geocoded data easily allows the integration of relevant neighborhood characteristics at the block, block group, and tract levels available through the Census and other sources, such as racial segregation, poverty, unemployment, etc.

Using machine learning methods, we leverage the assembled spatiotemporal records to predict court eviction filings at the property level. Modeling employs three machine learning classifiers (i.e., Random Forest, XGBoost, and Feedforward Neural Networks) and iteratively incorporates 1) historical eviction records, 2) neighborhood features, and 3) owner characteristics to assess improvements in predictive performance with additional features. We assess accuracy over time by comparing the change in predictive performance at 3, 12, and 24 months following training.

Finally, we evaluate the effectiveness of risk scores produced by the models in targeting outreach efforts—identifying households at high risk of eviction and routing a theoretical set of caseworkers to them for assistance. Simulations compare risk-score-based outreach policies with commonly used alternatives: outreach based on the recent history of evictions in 1) individual households and 2) neighborhoods.

Preview of Results: Our findings inform the feasibility and potential utility of data-driven outreach for promoting housing stability. Our models accurately predict tenant eviction up to 24 months before court filings. Predictive accuracy incrementally improves when adding features on the owners and neighborhood characteristics of the properties. For example, using only historical eviction records, XGBoost yields an area under the ROC curve of 0.76. Incorporating the neighborhood features with the eviction records improves this by 16% (to 0.84). The AUC further increases by 11% (to 0.89) when we integrate owner features into the model. We show this trend remains relatively stable at 3, 12, and 24 months after training when using all information, whereas performance degrades rapidly when using historical eviction records alone. Models generalize well even when the COVID shock is part of the intervening time frame.

More importantly, we are able to analyze the use case – targeting outreach using our risk scores. When simulating a targeted outreach policy that focuses on higher-risk households, we show that caseworkers would reach 8.5% more eviction-prone properties in the same amount of time as compared with neighborhood-based outreach, and 28% more than an outreach policy focused on buildings with a history of recent evictions. The addition of neighborhood features (demographics, socioeconomic status, and characteristics of the housing stock) adds substantial marginal value to the risk scores when used in our routing policy, while the further addition of owner attributes yields limited additional value. Moreover, the improvement in outreach policy performance shows a much stronger

correlation with improvement in the area under the ROC curve, as opposed to improvement in the area under the precision-recall curve, despite literature claiming benefits for the latter in low base-rate tasks such as ours [9]. Given the context, our results support the feasibility of using risk modeling and data-driven targeting to inform eviction prevention.

1.1 Social Significance and Prior Research

The lack of affordable housing threatens the safety and security of low-income households across the US. Before COVID-19, an estimated 36.5 million US households experienced housing burden, defined as paying more than 30% of income on rent [23]. The Census Household Pulse Survey – a nationally representative biweekly assessment of the economic and health impacts of COVID-19 – shows that 30% of all renters felt little or no confidence in paying the next month’s rent three years into the pandemic, while racial disparities in housing insecurity persist [24, 41].

In addition to its obvious economic implications, eviction has significant negative effects on health and well-being, including higher rates of stress, depression, anxiety, and substance abuse with life-course implications not just for adults but also for the health and well-being of children [11, 20, 40]. Several studies using the Pulse Survey show that households at risk for housing displacement report elevated rates of anxiety and depression [1, 5, 24]. Concerns regarding a tidal wave of evictions at the onset of the COVID-19 pandemic triggered unprecedented federal responses, including federal, state, and municipal moratoria on evictions and allocation of \$46 billion of housing assistance for low-income households. The policies implemented in this response, including moratoria on eviction filings and executions and delivery of rental assistance, improved outcomes [4, 14, 29].

Today, however, eviction filing rates continue to rise beyond pre-COVID levels in communities across the country – putting 2.3 million renters at risk in the near future [19, 30]. Governmental and non-governmental agencies continue to actively search for ways to protect renters from the devastating impact of eviction on health and well-being [3, 11–13].

Community efforts to prevent evictions typically focus on providing tenants with time-limited assistance, such as one-time financial aid to pay rental and utility arrears, landlord mediation to resolve disputes, legal representation in court proceedings, or referral to an array of local services provided by governmental and non-governmental agencies [7, 21, 44]. The lack of affordable housing in many communities challenges the efficient targeting of resources, as demand for stable housing outpaces the availability of limited resources [16, 37]. Governmental and non-governmental agencies struggle to coordinate in identifying and responding to housing risk that warrants intervention with available assistance. Machine learning techniques applied to eviction prediction offer promise for improving the efficiency of prevention efforts; however, technical and ethical challenges impede efforts. Data on rental market dynamics generally lack specificity to inform targeting responses. For example, few municipalities track which residential housing includes rental units, and court eviction records – where available – capture a fraction of places where evictions may occur,

impeding community-wide surveillance. The most proximal information on households collected through the Census, such as race, socioeconomic status (SES), building features, etc., is captured at the block level, or a statistical area of hundreds of housing units, and averaged over five years. Proprietary records on rental units (e.g., average rent listings) collected by Zillow and other sources are reliable at the zip code level – limiting the application of real-time responses to market shifts. Even municipal records collected through tax assessments, code enforcement, and other governmental interactions fail to identify the properties held by a single owner registered under various names and limited liability corporations (LLCs), which obscures shared risk across properties and potential landlord interventions.

Although limited in application, initial efforts to use algorithmic approaches for risk detection show promise. One study in Dallas, Texas, predicted eviction rates within census tracts, representing statistical areas including 4,000 residents on average [39]. A neural network built with historical court filing records and an array of tract-level sociodemographic characteristics produced reasonable predictive accuracy that generalized to unseen census tracts. Yet, the tract-level results provide limited practical utility for targeting households for prevention. Another study conducted before COVID-19 in New York City captured building-level risk of landlord harassment reported by tenants to the Mayor’s Office [45]. A gradient-boosting classifier trained on 92 building-level and neighborhood characteristics outperformed other machine-learning classifiers and expert-driven canvassing in predicting the likelihood of harassment. The results demonstrate the utility of data-driven approaches despite not directly targeting risk for eviction.

Several barriers impede the broader implementation of risk detection and data-driven eviction prevention. Assembling public and proprietary information for model building requires extensive integration of disparate data systems designed for different purposes. For instance, court records of evictions document procedural processes for adjudication without the intention of linking to other municipal data, such as tax assessment and code enforcement. Beyond the nontrivial effort of data collection, much remains uncertain in building accurate models. The likelihood of eviction depends on multiple interacting contexts that prove challenging to capture [39, 45]. Moreover, little is known regarding whether and how rental property owner characteristics (e.g., number of properties, corporate status, location) matter for evictions; sociological research on predatory behaviors of slumlords would suggest so [11]. Likewise, the accuracy of eviction risk predictions over time remains unclear. Pragmatically, it would be easier to deploy interventions based on models that require infrequent updating; however, rapid rental market dynamics could degrade prediction accuracy quickly, requiring rapid adjustments. The fallout from COVID and slow economic recovery warrants careful attention in building accurate models. The present study addresses these questions through a novel community-based research partnership.

2 STUDY CONTEXT, DESIGN, AND DATA

2.1 Study Context

The study leverages an ongoing community-based research partnership aimed at preventing tenant evictions in St. Louis, MO. In

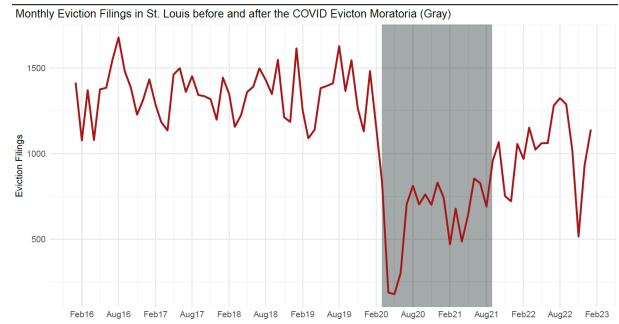


Figure 1: Monthly eviction filings in St. Louis from January 2016 - January 2023.

combination with the neighboring (and non-overlapping) St. Louis County, the midsize, Midwestern city includes 1.3 million residents – 61% of whom identify as White and 30% as Black. Racial and economic segregation remains challenging in St. Louis. Approximately 30% of the population lives in a neighborhood that is more than 90% White or Black, and Black households are nearly three times more likely to fall below the federal poverty level. Among the approximately 210,000 renting households in St. Louis, 46.5% and 47% are rent-burdened in the City and County respectively, meaning that they spend 30% or more of their income on rent [42]. Community partners include academics, housing and social service providers, governmental and non-governmental administrators, local philanthropic agents, and a civic technology firm with extensive expertise in linking and spatially mapping property-level information collected across municipal departments.

Figure 1 illustrates eviction filing trends from January 2016 through January 2023. Before COVID-19, filings averaged eight evictions per year per 100 renters. Eviction moratoria between March 2020 and August 2021 prohibited the execution of evictions and slowed filing rates that gradually returned to pre-COVID levels by the end of 2023.

2.2 Study Design

We present a two-stage empirical study. First, we evaluate property-level eviction prediction performance for different predictive models; specifically, we build binary classifiers of whether an eviction occurs at a residential property. Properties are defined as the smallest separable taxable unit of real estate. In a broader setting, properties take the form of individual homes, entire apartment buildings, or single units of a condominium or townhome. In this work, we restrict the study to properties with two or more units, given differences in the frequencies of evictions; filings occur at less than one-quarter the rate within single-unit properties, according to court records.

Figure 2 illustrates the three feature categories of focus – the historical eviction patterns of the property, neighborhood attributes such as demographics and income, and owner information. We analyze whether incrementally adding these features to eviction records affects prediction. We also evaluate how model performance changes with both increased time between training and testing periods and pandemic-induced shifts in the local housing market, which




Figure 2: Thematic representation of the feature sets in different models. Each property is associated with an eviction record and is situated within a neighborhood, specifically a Census block and block group. Properties are owned by owners who may own multiple properties in different neighborhoods. The number of features incorporated into the binary classifier expands with the radius of the diagram.

include eviction moratoria and allocation of emergency rental assistance.

In the second stage, we utilize risk scores generated from the best-performing model to design and test outreach routing policies that target vulnerable tenants before an eviction filing occurs. Simulations compare our risk-score-based targeted outreach policy with commonly employed outreach practices. We evaluate the policies based on the number of evictions discovered in a fixed period.

2.3 Data

Tolemi, our civic technology partner, assembled a map of all properties in the region and joined these properties with information from multiple sources using standard name and definition protocols. Data came from governmental and nongovernmental open data portals that periodically pulled eviction records (daily), tax assessments (annually), and other information when made available. Data were spatially linked through GIS and merged with Census features. Extensive verification identified owners by linking names and mailing addresses in tax assessments available from 2019 through 2023. Each owner received an owner identification number assigned across all regional properties. Ownership changes were also captured. We combine these data¹ to predict eviction at a property. Table 1 describes the attributes.

Eviction Filings. Court-involved eviction filings represent the primary outcome. By overlaying eviction records across the entire geographic area, we leverage address-level data to capture evictions within properties that can be linked with building, neighborhood, and owner characteristics.

There were 45,197 mappable eviction filings across 15,349 city and county properties from January 2019 through February 2023². Each eviction record includes a filing date, case ID, property ID, and information on the defendant, plaintiff, and plaintiff’s attorney. As detailed below, we counted the monthly evictions for each property

at distinct periods and generated the target labels of whether an eviction occurred in subsequent periods.

Neighborhood Characteristics. We link each property with five-year block group estimates collected through the 2019 American Community Survey (ACS)³. Block groups – the smallest geographic units released by the Census for the ACS – typically represent 600 to 3,000 residents and capture neighborhood features, such as median income, gross rent, households receiving food stamps or cash assistance, and health insurance. We use more proximal block-level attributes for race/ethnicity, age, and occupancy status available through the 2020 Decennial Census⁴. Rates were computed to capture neighborhood characteristics at the lowest level of geography, such as the percentages of each race within a block (Decennial Census) and high school graduates in a block group (ACS). Table 1 lists all neighborhood attributes included in modeling.

Owner Characteristics. We selected properties listed as residential in tax assessments with more than two units within the property. There are 26,770 properties that meet the criteria with linked owner information. Tax assessments record the initiation and termination dates for property ownership, allowing linkages between a property and the owner in a given time frame. Each owner-property pair includes records on whether the owner is a business, the business type (e.g., LLC versus estate), if the owner lives out-of-state, and whether the owners reside at the property. Moreover, we identified whether owners worked with high eviction filing attorneys. Court records enabled counting filings by each attorney, which we ranked from highest to lowest for a given period. Two time-varying indicators captured whether owners filed at least one eviction with (1) a top 25 highest-filing attorney or (2) a top 26 to 50 highest-filing attorney in a given period.

3 PREDICTING RISK SCORES FOR EVICTION

3.1 Experimental Setup

Models. The eviction models built in this study generate property-level scores valued between 0 and 1 representing the risk of an eviction occurring within the next three months. Eighteen classification models were built, each leveraging one of three different learning algorithms, incorporating one of three nested subsets of features, and trained on data from either before or after the start of the COVID-19 pandemic. The three algorithms used were random forests (RF), gradient-boosted decision trees via XGBoost (XGB), and a Feedforward Neural Network (FNN). Implementation details for these models are described in Section A.1 of the Appendix. The sets of features are visualized in Figure 2 and described in detail in Table 1 (Eviction Filings; Eviction Filings and Neighborhood Characteristics; Eviction Filings, Neighborhood Characteristics, and Owner Characteristics).

Train-Test Splits and Model Timeline. Figure 3 illustrates the model timeline with training and test set splits. We drew data to train the models from two ten-month periods: **Pre-COVID**, the time from January 2019 to October 2019, and **Post-COVID**, the

¹Part of our data is proprietary and subject to a specific data sharing agreement that does not allow for public release.

²An additional 4,351 filings (8%) could not be mapped from court records.

³<https://www.census.gov/programs-surveys/acs/data.html>

⁴<https://www.census.gov/programs-surveys/decennial-census/decade/2020/2020-census-main.html>

Table 1: Attribute Description

Category	Attribute Details	Data Type
Eviction Filings	# of evictions during the period by month and quarter	Integer
Neighborhood Characteristics	Decennial Census Block Data: % of population under 18; % of housing units occupied; % of population by racial and ethnic groups American Community Survey Census Block Group Data: median household income; median gross rent; gross rent as a % of household income (GRAPI); % of housing units that are renter occupied; % of renter-occupied housing units with more than one occupant; % of adults below the federal poverty level; % of households with a mortgage; % of households receiving SNAP and/or public assistance; % of households with health insurance; % of female-headed households with children under 18; % of adults with a high school degree; % of adults who are a veteran	Float
Owner Characteristics	# of units in the property	Integer
	# of properties in St. Louis held by the owner	Integer
	is the owner a business	Boolean
	is the property owner-occupied	Boolean
	has the owner used one of the top 25 highest-filing attorneys during the period	Boolean
	has the owner used one of the top 50 highest-filing attorneys during the period	Boolean
	does the owner live in St. Louis, in-state, or out of state	Categorical

time between January 2021 to October 2021. During training, models predict whether an eviction occurred in the last three months of the training period (August-October) by utilizing features from the first seven months of the period (January-July). Feature sets (i.e., historical eviction filings, neighborhood characteristics, and owner characteristics) are included iteratively with five-fold cross-validation.

We test our models using three-month intervals following training. The **Pre-COVID** models test at three months from November 2019 to January 2020 (T_1), twelve months from November 2020 to January 2021 (T_2), and 24 months from November 2021 to January 2022 (T_3). **Post-COVID** models test three months (T_3) following training and twelve months from November 2022 to January 2023 (T_4).

Evaluation Metrics. We first evaluate model performance by plotting the Receiver Operating Characteristic (ROC) curves. The ROC Area Under the Curve (AUC) provides a base-rate agnostic measure for differentiating between positive (Eviction) and negative cases; AUC values reaching 1 indicate perfect differentiation between the two cases in the sense that all the positive examples are ranked higher than the negative examples.

We also assess the AUC of the PR curve, which reports the proportion of positive cases (evictions) correctly identified for different probability thresholds [8]. Precision measures how many of the items predicted to be positive cases are truly positive cases (true evictions), and recall is the fraction of positive cases identified. As the base rates of evictions are low (< 0.1), the area under the PR curve could provide a useful measure for informing the value of interventions when compared against the base rate [9, 38]. We later investigate further the value of both metrics for the actual outreach goal. We use DeLong’s algorithm to test the statistical significance of differences in ROC and PR AUCs between models [10].

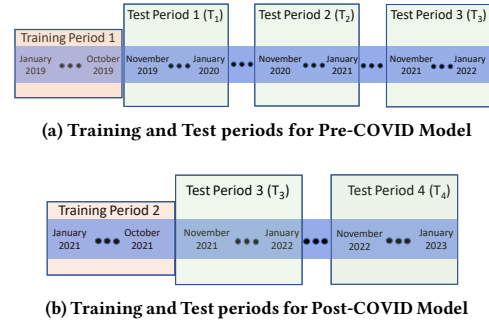


Figure 3: Training and testing timelines. The three black dots represent the interim months between the beginning and end of a period. There are two training periods and four testing periods. We do not train a model on data from 2020 due to the outsized influence of COVID-19 and the corresponding eviction moratoria on eviction patterns. We denote the model trained using Training Period 1 as the Pre-COVID Model and the model trained using Training Period 2 as the Post-COVID Model. We select three-month test periods composed of November, December, and January to compare the ability of the two models to generalize into the future. Data from all testing periods are withheld from both models during training.

3.2 Prediction Results and Analysis

Table 2 compares the predictive performance of the risk models. We analyze performance on different feature sets, model structures (decision trees vs. FNN), timeframes, and evaluation metrics.

Neighborhood and Owner Features Improve Performance. We start by comparing the models with different feature sets in

Table 2: Cross-model performance for predicting tenant evictions. *Training Period* indicates whether modeling occurs before or after the pandemic onset. *Testing Period* and *Base Rate* show the eviction rate during predictions. *Feature Sets* report AUC from Random Forest (RF), XGBoost (XGB), and Feedforward Neural Network (FNN) classifiers when including additional information across testing periods. The top rows summarize AUC for the ROC curves, while precision-recall curves are presented at the bottom. We note that AUC improves significantly ($p < 0.001$) with each additional feature set for most testing periods.

Train Period	Test Period	Base Rate	Feature Sets								
			Eviction			Eviction and Neighborhood			Eviction, Neighborhood, and Owner		
			RF	XGB	FNN	RF	XGB	FNN	RF	XGB	FNN
AUC for ROC Curve											
Pre-COVID Jan 2019 - Oct 2019	T_1	0.032	0.76	0.76	0.76	0.82	0.84	0.82	0.88	0.89	0.88
	T_2	0.019	0.71	0.71	0.71	0.80	0.86	0.78	0.90	0.92	0.89
	T_3	0.021	0.73	0.73	0.73	0.83	0.84	0.82	0.89	0.90	0.89
Post-COVID Jan 2021 - Oct 2021	T_3	0.021	0.75	0.73	0.74	0.83	0.84	0.81	0.89	0.90	0.90
	T_4	0.019	0.75	0.75	0.75	0.83	0.84	0.80	0.89	0.89	0.89
AUC for Precision Recall Curve											
Pre-COVID Jan 2019 - Oct 2019	T_1	0.032	0.37	0.38	0.38	0.42	0.43	0.42	0.45	0.47	0.45
	T_2	0.019	0.37	0.36	0.32	0.39	0.37	0.33	0.40	0.44	0.42
	T_3	0.021	0.39	0.39	0.33	0.36	0.39	0.36	0.42	0.43	0.41
Post-COVID Jan 2021 - Oct 2021	T_3	0.021	0.39	0.38	0.32	0.36	0.36	0.34	0.40	0.40	0.40
	T_4	0.019	0.39	0.38	0.33	0.37	0.37	0.36	0.40	0.42	0.41

* RF = Random Forest Classifier, XGB = XGBoost Classifier, FNN = Feedforward Neural Network




Figure 4: ROC Curves for T_1 and T_3 . (a) The AUC value for T_1 prediction using the Pre-COVID (XGB) model that includes Eviction, Owner, and Neighborhood attributes has the highest ROC AUC (0.89), and outperforms the baseline and model trained on eviction+neighborhood with high levels of statistical significance (p -values of $3.34e^{-72}$ and $2.44e^{-31}$ respectively); (b) Inclusion of Owner, and Neighborhood attributes in the Post-COVID (XGB) model also increases the ROC AUC and the curve is significantly different than the baseline and the model trained on including neighborhood attributes with p -values $1.27e^{-74}$ and $8.00e^{-19}$.

terms of their performance in predicting the immediate future (using features from the last seven months to predict if there will be

an eviction in the next three). For Pre-COVID (XGB) model, when using only prior eviction records without any information about the owner or the property, the ROC AUC for T_1 is 0.76. Adding the neighborhood attributes yields a 8-point increase in the AUC score, while including the owner attributes as features achieves an additional 5-point gain in AUC. Figure 4a shows the ROC curves for all three cases. A similar pattern emerges for the Pre-COVID (RF) and Pre-COVID (FNN) models. For both models, the AUC of ROC increased from 0.76 (only eviction data) to 0.88 (all features).

Similar to the Pre-COVID models, we observe a higher gain for the Post-COVID (XGB) model after the inclusion of the neighborhood attributes (0.73 to 0.84) and an additional gain of 5 points after including the owner attributes (0.84 to 0.90) (Figure 4b) for the ROC AUC for T_3 . We note that all the AUC differences in Pre-COVID and Post-COVID models are statistically significant at $p < 0.001$.

Performance Remains Stable for Different Models. We compare three predictive models to predict evictions where two of them are based on decision trees (Random Forest and XGBoost) and one of them is a Feedforward Neural Network. The performance remains stable for all of the models for most of the test periods. Although we observe differences in a few cases, for example, the difference in AUC of ROC between Pre-COVID (XGB) and Pre-COVID (FNN) for predicting T_2 using eviction and neighborhood features is 8 points, the difference is negligible when all of the features are included in the prediction where the AUC of ROC remains between 0.88 to

0.92 (three right most columns in the table). Most of the AUCs of PR also remain within 0.42 and 0.45 when using all features.

Models Perform Better in the Distant Future When Owner and Neighborhood Features are Used. We now turn to understanding how well models generalize over time, and we do this in two quite different contexts. In general, we would expect performance to degrade over time (that is, a model trained on data from January - October 2021 would, all else equal, be expected to perform better in predicting evictions from November 2021 - January 2022 than in predicting evictions from November 2022 - January 2023). However, with the dramatic impact of the eviction moratorium caused by COVID-19 on eviction patterns, we would expect to see a more substantial degradation in performance if the two test periods are on different sides of the eviction moratorium.

Our results in Table 2 demonstrate this clearly when only the eviction features are used. The performance of the Pre-COVID model (XGB) drops from an AUC ROC of 0.76 to 0.71 and an AUC PR of 0.38 to 0.36 when predicting T_2 rather than T_1 . The drop is not observed when using a similar “delay” but entirely in the Post-COVID timeframe. Interestingly, we do not see a substantial drop in performance when using owner and neighborhood attributes with the eviction records, showing that these attributes add robustness to the models, even during an unprecedented emergency like COVID-19, compared with simply using eviction records as features.

4 OUTREACH ROUTING POLICY

We now turn back to one of the main questions of interest: Can we leverage the generated property-level risk scores from the models to devise better routing policies for caseworkers to visit households that may be at high risk of eviction?

We describe our risk-score-based outreach policy and compare its performance in discovering evictions with commonly used policies.

4.1 Routing Path and Estimated Outreach Time

To compare outreach policies, we require an estimate of the time required to canvass a set of properties. We estimate the time it takes to canvass property B from property A in two parts: first, we estimate the time it takes to travel from A to B using a piecewise linear function; second, we add a cost proportional to the unit count of B to account for knocking on all doors upon arrival. This cost is calibrated to data provided by community partners. Further details of the implementation are provided in Appendix B.1. The geodesic distance between two properties is calculated using the GeoPy Python library [25], and the fastest route between all canvassed properties is calculated as an approximate solution to the Traveling Salesperson Problem (TSP) using the OR-tools library [28].

4.2 Routing Policies

4.2.1 Risk Score (RS) Based Outreach. We consider risk scores generated by XGBoost – the best-performing model – for T_4 . This presented the latest period (November 2022 to January 2023) available in the dataset and is most interesting to our community partners.

We categorize the risk scores into four groups (Very Low, Low, Medium, and High) for identifying properties with a greater risk of eviction. Figure 5 shows the distribution of properties in each risk group based on the number of units in the property. Most




Figure 5: Distribution of number of units that fall in a particular risk group [Low, Medium, High]. The X-axis represents the groups based on the number of units, and the Y-axis represents the proportion of properties from each category of unit sizes.

properties in T_4 (69.3%) fall into the Very Low category, where risk scores range from 0 to 0.05, reflecting a true eviction rate of 0.003. The ranges of risk scores for the Low, Medium, and High-risk groups are (0.05 – 0.2], (0.2 – 0.8] and (0.8 – 1.0]. The percentages of properties in these groups and corresponding true eviction rates are (22.08%, 7.59%, 0.98%) and (0.01, 0.09, 0.61), respectively. We might worry that the model is simply picking up on the existence of higher numbers of units in a property to assign higher risk scores, but, even within a unit size category (e.g., 100+ units), we see significant variance among the risk score groups assigned to different properties, which alleviates this concern (see Figure 5).

Neighborhood, Eviction, and Owner (based) - Targeted Outreach (NEO-T-O). We select properties in the medium and high-risk groups for targeting outreach efforts (i.e., caseworkers). There are 2,299 properties situated in St. Louis City and County in these two categories. To approximate the amount of time it would take to canvass each unit in these properties, as described above, we obtain an approximate solution of the TSP to estimate travel time, and add to that the estimated time spent “knocking on doors”.

RS Outreach Based on Eviction Records Only. We consider risk scores generated by XGB using features only from historical eviction records. We sort the properties based on these risk scores and select the top- k properties that can be visited in the same time as that spent by NEO-T-O in order to ensure fair comparison (k is estimated using incremental search).

RS Outreach Based on Eviction and Neighborhood Information. This policy is the same as the one above, except that it considers risk generated by XGB using features from both eviction records and neighborhood information.

4.2.2 Alternative Policies. We compare NEO-T-O and the other two RS outreach policies with two common outreach practices: targeting based on prior evictions (regardless of property location), and canvassing all properties within specific neighborhoods. Again, in order to ensure a fair comparison, we control for available outreach time. These policies end up looking quite different from the ones above, in that they tend to visit many more (smaller) properties, rather than concentrating resources on fewer larger properties. Thus, they spend more of their time traveling between properties and are able to visit fewer overall units. In order to separately analyze the effect of this substantive difference, we also consider variants that visit the same number of *units* as NEO-T-O.

Table 3: Cross-policy performance analysis for discovering tenant evictions. *Routing Policy* lists the risk-score-based and alternative outreach policies we consider. *Control* indicates the controlling factors in policies. *Normalized Outreach Time* shows the outreach time normalized by the time required by NEO-T-O. *Eviction Discovery Rate* reports the rate of eviction discovery per property for each policy. *Units Visited* reports the total number of units that are visited by the policy.

Routing Policy	Control	Normalized Outreach Time	Evictions Discovered	Properties Visited	Units Visited	Eviction Discovery Rate
Risk Score Based Policy						
NEO-T-O	Time	1	936	2,299	101,119	40.7%
Eviction Features	Time	1	677	1,711	101,288	39.6%
Eviction and Neighborhood Features	Time	1	919	4,000	94,556	23.0%
Alternative Policies						
Previous Eviction Count	Time	1	731	13,122	90,943	5.6%
	Unit	1.63	749	14,987	101,118	5.0%
Neighborhood Based	Time	1	863	11,460	99,532	7.5%
	Unit	1.03	870	11,905	101,116	7.3%

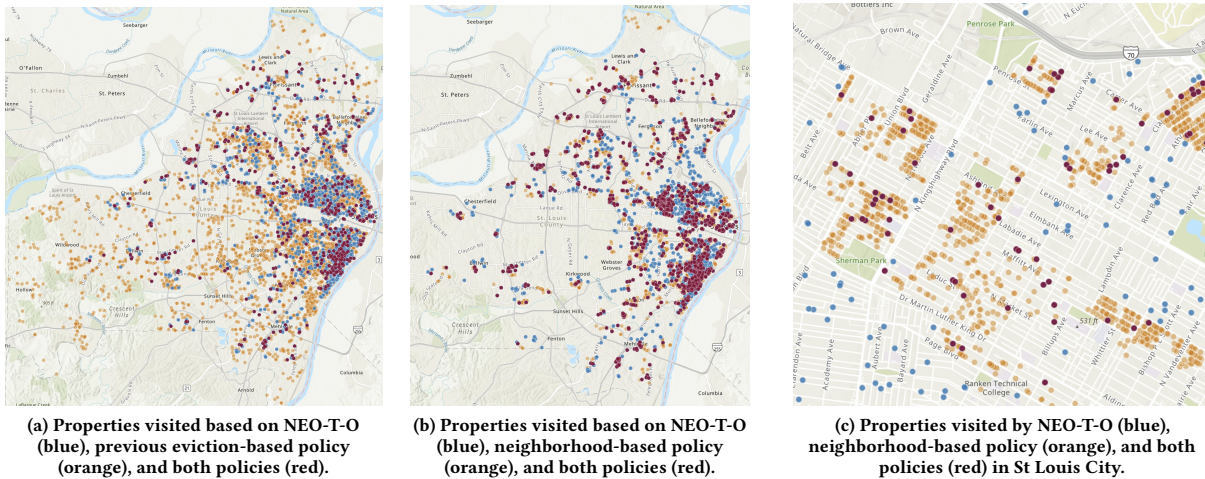


Figure 6: Canvassed properties in St. Louis City and County using different outreach policies. In all panels, blue dots represent the properties canvassed by NEO-T-O. Panel (a) shows properties selected by previous eviction-based policy (orange dots) - scattered across the map - and properties canvassed by both NEO-T-O and previous eviction-based policy (red dots). Panel (b) shows the properties selected by neighborhood-based policy (orange dots) - grouped in different areas - and properties canvassed by both NEO-T-O and neighborhood-based policy (red dots). Panel (c) shows a close-up of part of Panel (b).

Previous Eviction Count Based Policy. The properties are sorted by the number of prior evictions in a previous period, and outreach targets the “top- k ” properties in terms of prior evictions, again computing travel time by approximating the TSP solution among these k . We use observed filings in the quarter (August 2022 - October 2022) prior to period T_4 .

Neighborhood Based Policy. This routing policy visits all rental properties in certain neighborhoods. Neighborhoods – defined by census block groups – are ranked by prior eviction counts observed in the quarter before T_4 . For each neighborhood, we compute the approximately optimal route to visit all rental properties. We continue visiting neighborhoods in descending order until the time or unit constraints are met.

4.3 Outreach Routing Results and Analysis

Usefulness of Risk Scores. The headline result is that the total number of evictions discovered (936) by NEO-T-O is 8.5% higher than the neighborhood-based policy (863) and 28% higher than the previous eviction-based policy (731) when controlling for outreach time. While risk-score-based policies discover more evictions than alternative policies, they canvass a substantially lower number of properties despite both types of policies visiting similar numbers of units. Thus, RS-based policies are more useful if outreach programs want to utilize limited resources (e.g., caseworkers) to discover more evictions while canvassing a small number of properties. Both policies canvass a similar number of units, which implies that RS-based policies target large properties with a higher number of units, whereas alternative policies also include properties with fewer units. For the same reason, RS-based policies yield higher eviction discovery rates on a per-property basis than alternative policies.

Importance of Neighborhood and Owner Features. The addition of neighborhood features yields substantial marginal value to the risk scores when used in our routing policy, while the further addition of owner attributes has limited additional value beyond that. When adding neighborhood features with the eviction features, the number of discovered evictions increases to 919 from 677, yielding a 35.8% gain. However, adding owner features to the eviction and neighborhood features yields only a 1.8% gain (from 919 to 936). It is also worth noting here that the marginal benefit to the outreach strategy is better captured by the area under the ROC curve as opposed to the area under the Precision-Recall curve. This is contrary to the general conjecture in the literature that the PR curve has more value in discovery tasks with low base rates [9, 38].

Visualizing the Policies. Figure 6 shows maps of St. Louis City and County, where the dots indicate properties that would be reached by different outreach policies. We plot the properties canvassed by NEO-T-O and two categories of alternative policies when controlling for time: previous eviction-based outreach and neighborhood-based outreach.

We observe that the 13,122 properties canvassed through the eviction-based policy are scattered across the city (Figure 6a), whereas the 11,460 properties of the neighborhood-based policy are (by design) grouped in several census blocks of the map (Figure 6b). 582 properties are canvassed by both NEO-T-O and the prior eviction-based policy, covering 25% of all properties canvassed by NEO-T-O. 1,419 properties are canvassed by both NEO-T-O and the neighborhood-based policy, covering 62% of the NEO-T-O properties. This large intersection between the properties targeted by NEO-T-O and the neighborhood-based policy further explains why we do not see a large increase in the number of evictions discovered by NEO-T-O in comparison to the RS-based policy using previous eviction and neighborhood information. Figure 6c shows a closer look at the properties reached by NEO-T-O (blue), neighborhood-based policy (orange), and both (red) policies. While the orange dots cover all the properties within a neighborhood, we observe some properties near the boundary of the neighborhoods that are not canvassed by the neighborhood-based policy but have high risks of eviction (blue) according to NEO-T-O.

5 DISCUSSION

What are the benefits of using risk scores in societal service provisions such as the prevention of eviction? How should we evaluate these risk scores in real-world decision-making? This paper adds to the growing literature [43] making the case that we need to go beyond prediction accuracy and critically examine how usable the risk scores are in *decision making* – in our case, devising outreach programs and preventing evictions.

In some settings, individualized risk scoring, even when relatively accurate, has not been shown to have benefits in terms of interventions beyond those that could be obtained through less granular means [33]. We obtain a somewhat different result in a different domain. We propose a risk-score-based policy – NEO-T-O – that identifies eviction-prone properties to send a theoretical set of caseworkers to discover possible evictions and provide assistance. Our results show that such a policy is, in fact, more useful in discovering evictions than common outreach practices that focus exclusively on certain neighborhoods or properties with recent prior eviction histories. However, it is also worth noting that, consistent with existing research, a significant portion of the benefit can be achieved through the use of neighborhood socio-demographic information.

Future research should be aware of trade-offs in using risk-score-based outreach policies. NEO-T-O focuses on a small number of properties with high eviction risk. This means that NEO-T-O coverage is focused on developments with large numbers of units, which may be efficient in the short term but leave out certain populations or lead to systemic problems downstream. There is also a theoretical risk that risk-score-based policies could be prone to feedback loops [15]: caseworkers are sent to targeted properties, leading to higher rates of discovery of potential evictions, and subsequently, more caseworkers are sent to those properties, and so on.

Our results also show that risk scores that are critical for NEO-T-O are typically stable across different models and timeframes (including pre- and post-COVID timeframes). Classifiers accurately identify properties at risk of eviction up to 24 months before filing when incorporating publicly available information on past evictions, neighborhood characteristics, and owner features. The evidence suggests a relatively wide window of opportunity for tenant-landlord intervention. Moreover, the unique data linkages identify property ownership registered under various names and corporate entities across the region. The information could enable greater scrutiny of nefarious owner eviction behaviors, housing code violations, tax abatement, and other public funding that could inform strategic action by governmental and non-governmental agencies to protect tenants from problematic owners [17, 18, 36].

Our findings must be interpreted under several limitations. The results are currently limited to one city and county. This is primarily because curating such a dataset is itself a challenging task due to the presence of restrictive state and local policies and limitations of governmental infrastructures. We found that eviction itself is a relatively low-probability event and generally occurs in rental properties owned by larger corporations rather than single-owner family homes. Thus, our analysis is limited to residential rental properties with two or more units. In our experiments, the data only

represent legal eviction filings; illegal evictions in which owners displace tenants without court orders remain unobserved.

We also note that we do not have any personal information about individual tenants in the data. The eviction court orders do not contain sensitive tenant information, such as gender, race, etc., other than their names. Thus, it is difficult to evaluate the responses of the routing policies towards different social groups.

Research Ethics Statement. The research accords with the ACM Code of Ethics and Professional Conduct. We intentionally engage community partners in framing the problem, and the research team appropriately includes individuals with appropriate knowledge and skills for implementing the study design. Although the data come from publicly available information sources, we use a data-sharing agreement between community partners to ensure appropriate use. Our modeling generates property-level eviction risk scores; however, these are not currently shared with community partners, and evaluation relies on simulation to ensure privacy and avoid unintended consequences from deployment without adequate testing. Subsequent research will examine how to implement methods in ways that improve well-being and avoid harm.

Researcher Positionality Statement. The community-driven research is part of an ongoing partnership with local stakeholders working to prevent evictions and the deleterious impact on children, families, and communities. The authors include a diverse trans-disciplinary team (computer science, data science, community public health) that includes doctoral students and faculty from around the globe. Members approach the work from different perspectives. Active engagement of community partners through frequent meetings (multiple times per week when necessary) provides a strong foundation for repeatedly assessing the utility, feasibility, and ethical implications of the data and modeling. Our partners include another diverse set of experts with extensive experience in service provision, policy, and navigating unstable housing in tight affordable housing markets. The work aims to acknowledge and leverage the various perspectives into building technical and equitable solutions for keeping low-income families housed.

Adverse Impact Statement. Our study demonstrates the feasibility and potential utility of risk-score-based outreach to prevent tenant evictions. Simulation suggests AI-enabled tools could make targeting more efficient, yet deployment requires careful consideration of unintended consequences. The scarcity of housing resources forces difficult decisions regarding who to assist that must account for equity across individuals and groups. Moreover, outreach efforts that alienate rental owners could further deplete access to affordable housing. The implementation of risk-score-based outreach that prevents eviction must be carefully and collaboratively designed and evaluated – a goal of our future research.

ACKNOWLEDGMENTS

The St. Louis Regional Response Team and members of the collaborative network offer ongoing assistance in the research design, including our civic technology partner, Tolemi, who assembled the administrative records. We are grateful for support from NSF grants 2127752 and 1939677, as well as support from Amazon through an NSF/Amazon FAI Award.

REFERENCES

- [1] Binod Acharya, Dependra Bhatta, and Chandra Dhakal. 2022. The risk of eviction and the mental health outcomes among the US adults. *Preventive Medicine Reports* 29 (2022), 101981.
- [2] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. 2018. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, Willem-Jan van Hoeve (Ed.). Springer International Publishing, Cham, 35–51.
- [3] Brian Bieretz, Kimberly Burrowes, and Emily Bramhall. 2020. Getting landlords and tenants to talk: The use of mediation in eviction. *White Paper, Urban Institute* (2020).
- [4] Courtney E Boen, Lisa A Keister, Christina M Gibson-Davis, and Anneliese Luck. 2023. The buffering effect of state eviction and foreclosure policies for mental health during the COVID-19 Pandemic in the United States. *Journal of Health and Social Behavior* (2023), 00221465231175939.
- [5] Gregory Bushman and Roshanak Mehdipanah. 2022. Housing and health inequities during COVID-19: Findings from the National Household Pulse Survey. *J Epidemiol Community Health* 76, 2 (2022), 121–127.
- [6] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 134–148. <https://proceedings.mlr.press/v81/chouldechova18a.html>
- [7] Damian Collins, Esther de Vos, Joshua Evans, Meryn Severson Mason, Jalene Anderson-Baron, Victoria Cruickshank, and Kenna McDowell. 2022. “When we do evict them, it’s a last resort”: Eviction prevention in social and affordable housing. *Housing Policy Debate* 32, 3 (2022), 473–490.
- [8] Jonathan Cook and Vikram Ramadas. 2020. When to consult precision-recall curves. *The Stata Journal* 20, 1 (2020), 131–148.
- [9] Jesse Davis and Mark Goadrich. 2006. The relationship between Precision-Recall and ROC curves. In *Proceedings of the 23rd International Conference on Machine Learning (Pittsburgh, Pennsylvania, USA) (ICML '06)*. Association for Computing Machinery, New York, NY, USA, 233–240. <https://doi.org/10.1145/1143844.1143874>
- [10] Elizabeth R DeLong, David M DeLong, and Daniel L Clarke-Pearson. 1988. Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. *Biometrics* (1988), 837–845.
- [11] Matthew Desmond. 2016. *Evicted: Poverty and Profit in the American City*. Crown.
- [12] Matthew Desmond. 2022. *Unaffordable America: Poverty, housing, and eviction: American Journal of Sociology*. Taylor and Francis, United States, 389–395. <https://doi.org/10.4324/9780429299377-34>
- [13] Matthew Desmond and Rachel Tolbert Kimbro. 2015. Eviction’s fallout: Housing, hardship, and health. *Social Forces* 94, 1 (2015), 295–324.
- [14] Rachel Donnelly and Mateo P Farina. 2021. How do state policies shape experiences of household income shocks and mental health during the COVID-19 pandemic? *Social Science & Medicine* 269 (2021), 113557.
- [15] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway feedback loops in predictive policing. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 160–171. <https://proceedings.mlr.press/v81/ensign18a.html>
- [16] Patrick J Fowler, Peter S Hovmand, Katherine E Marcal, and Sanmay Das. 2019. Solving homelessness from a complex systems perspective: Insights for prevention responses. *Annual review of public health* 40 (2019), 465–486.
- [17] Philip ME Garboden and Eva Rosen. 2019. Serial filing: How landlords use the threat of eviction. *City & Community* 18, 2 (June 2019), 638–661. <https://doi.org/10.1111/cico.12387>
- [18] Meredith Greif. 2018. Regulating landlords: Unintended consequences for poor tenants. *City & Community* 17, 3 (Sept. 2018), 658–674. <https://doi.org/10.1111/cico.12321>
- [19] Ashley Gromis, Ian Fellows, James R Hendrickson, Lavar Edmonds, Lillian Leung, Adam Porton, and Matthew Desmond. 2022. Estimating eviction prevalence across the United States. *Proceedings of the National Academy of Sciences* 119, 21 (2022), e2116169119.
- [20] Gracie Himmelstein and Matthew Desmond. 2021. Eviction and health: A vicious cycle exacerbated by a pandemic. *Health Affairs Health Policy Brief* 10 (2021).
- [21] Marieke Holl, Linda Van Den Dries, and Judith RLM Wolf. 2016. Interventions to prevent tenant evictions: A systematic review. *Health & Social Care in the Community* 24, 5 (2016), 532–546.
- [22] Rebecca Ann Johnson and Simone Zhang. 2022. What is the bureaucratic counterfactual? Categorical versus algorithmic prioritization in U.S. social policy. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAcCT '22)*. Association for Computing Machinery, New York, NY, USA, 1671–1682. <https://doi.org/10.1145/3531146.3533223>

- [23] Joint Center for Housing Studies of Harvard University. 2020. The state of the nation's housing. https://www.jchs.harvard.edu/sites/default/files/reports/files/Harvard_JCHS_The_State_of_the_Nations_Housing_2020_Report_Revised_120720.pdf
- [24] Daniel Kim. 2021. Financial hardship and social assistance as determinants of mental health and food and housing insecurity during the COVID-19 pandemic in the United States. *SSM-population health* 16 (2021), 100862.
- [25] Kostya Esmukov. 2015. Geopy. <https://pypi.org/project/geopy/>. Accessed: 2023-12-06.
- [26] Amanda Kube, Sanmay Das, and Patrick J. Fowler. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 622–629. <https://doi.org/10.1609/aaai.v33i01.3301622>
- [27] Amanda Kube, Sanmay Das, Patrick J. Fowler, and Yevgeniy Vorobeychik. 2022. Just resource allocation? How algorithmic predictions and human notions of justice interact. In *Proceedings of the 23rd ACM Conference on Economics and Computation* (Boulder, CO, USA) (EC '22). Association for Computing Machinery, New York, NY, USA, 1184–1242. <https://doi.org/10.1145/3490486.3538305>
- [28] Laurent Perron and Vincent Furnon. 2023. OR-Tools. <https://developers.google.com/optimization/>. Accessed on 22 November 2023.
- [29] Kathryn M Leifheit, Craig E Pollack, Julia Raifman, Gabriel L Schwartz, Robert D Koehler, Jackie V Rodriguez Bronico, Emily A Benfer, Frederick J Zimmerman, and Sabriya L Linton. 2021. Variation in state-level eviction moratorium protections and mental health among US adults during the COVID-19 pandemic. *JAMA Network Open* 4, 12 (2021), e2139585–e2139585.
- [30] Katherine E Marçal, Patrick J Fowler, and Peter S Hovmand. 2023. Feedback dynamics of the low-income rental housing market: Exploring policy responses to COVID-19. *System Dynamics Review* 39, 4 (2023), 371–403.
- [31] Saunik Narayanan, Guanghui Yu, Chien-Ju Ho, and Ming Yin. 2023. How does value similarity affect human reliance in AI-assisted ethical decision making?. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society (AIES '23)*. Association for Computing Machinery, New York, NY, USA, 49–57. <https://doi.org/10.1145/3600211.3604709>
- [32] OpenStreetMap contributors. 2017. OpenStreetMap. <https://www.openstreetmap.org>.
- [33] Juan C. Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. 2023. Difficult lessons on social prediction from Wisconsin Public Schools. arXiv:2304.06205 [cs.CY]
- [34] Gaurab Pokharel, Sanmay Das, and Patrick Fowler. 2024. Discretionary trees: understanding street-level bureaucracy via machine learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 20 (Mar. 2024), 22303–22312. <https://doi.org/10.1609/aaai.v38i20.30236>
- [35] Dasha Pruss. 2023. Ghosting the machine: Judicial resistance to a recidivism risk assessment instrument. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 312–323. <https://doi.org/10.1145/3593013.3593999>
- [36] Eric Seymour and Joshua Akers. 2021. Building the eviction economy: Speculation, precarity, and eviction in Detroit. *Urban Affairs Review* 57, 1 (Jan. 2021), 35–69. <https://doi.org/10.1177/1078087419853388>
- [37] Marybeth Shinn and Jill Khadduri. 2020. *In the midst of plenty: Homelessness and what to do about it*. John Wiley & Sons.
- [38] Helen R Sofaer, Jennifer A Hoeting, and Catherine S Jarnevich. 2019. The area under the precision-recall curve as a performance metric for rare binary events. *Methods in Ecology and Evolution* 10, 4 (2019), 565–577.
- [39] Maryam Tabar, Wooyong Jung, Amulya Yadav, Owen Wilson Chavez, Ashley Flores, and Dongwon Lee. 2022. Forecasting the number of tenants at-risk of formal eviction: A machine learning approach to inform public policy. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, Lud De Raedt (Ed.). International Joint Conferences on Artificial Intelligence Organization, 5178–5184. <https://doi.org/10.24963/ijcai.2022/719> AI for Good.
- [40] Jack Tsai and Minda Huang. 2019. Systematic review of psychosocial factors associated with evictions. *Health & social care in the community* 27, 3 (2019), e1–e9.
- [41] U.S. Census Bureau. 2021. Household pulse survey data tables. <https://www.census.gov/programs-surveys/household-pulse-survey/data.html>
- [42] U.S. Census Bureau. 2022. Selected housing characteristics. <https://data.census.gov/table/ACSDP5Y2022.DP04?q=rent&g=050XX00US29189,29510&y=2022>. Accessed on 22 January 2024.
- [43] Bryan Wilder, Bistra Dilkina, and Milind Tambe. 2019. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 0101 (July 2019), 1658–1665. <https://doi.org/10.1609/aaai.v33i01.33011658>
- [44] Lucie Ellen Woellenstein, Fiona Cuthill, Raahat Manrai, and David Henderson. 2023. Homelessness prevention interventions for single adults at risk of or experiencing MEH: A systematic review on their effectiveness. *International Journal on Homelessness* 3, 2 (May 2023), 344–410. <https://doi.org/10.5206/ijoh.2022.2.14681>
- [45] Teng Ye, Rebecca Johnson, Samantha Fu, Jerica Copeny, Bridgit Donnelly, Alex Freeman, Mirian Lima, Joe Walsh, and Rayid Ghani. 2019. Using machine learning to help vulnerable tenants in New York City. In *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies* (Accra, Ghana) (COM-PASS '19). Association for Computing Machinery, New York, NY, USA, 248–258. <https://doi.org/10.1145/3314344.3332484>

A APPENDIX A

A.1 Model Architecture and Hyperparameters

We compare performance across three classifiers – Random Forest (RF), XGBoost (XGB), and Feedforward Neural Network (FNN) – with hyperparameters set via grid search to maximize precision.

A.1.1 Hyper-parameters for Eviction Prediction Using Random Forest.

- Measure of impurity: Gini, Entropy, Log Loss
- Max depth: [2, 5, 10, 15]
- Minimum samples in a leaf node: [1, 2, 3, 4, 5]
- n_estimators: [50, 100, 500]

A.1.2 Hyper-parameters for Eviction Prediction Using XGBoost.

- Max depth: [2, 3, 4, 5]
- Learning Rate: [0.01, 0.05, 0.1]
- n_estimators: [50, 100, 500]
- scale_pos_weight: [1, 3, 5]
- gamma: [0, 0.05, 0.1]

A.1.3 *FNN Architecture.* The FNN architecture consists of an input layer, two hidden layers, and an output layer. The input layer is designed with several neurons equal to the input features, while the first hidden layer comprises 64 neurons activated by rectified linear units (ReLU), followed by Batch Normalization for normalization and Dropout with a rate of 0.5 for regularization. The second hidden layer conducts the same processing with 32 neurons. The output layer uses a single neuron with the sigmoid activation function. This architecture is compiled using binary cross-entropy as the loss function and the Adam optimizer.

B APPENDIX B

B.1 Routing Path and Estimated Outreach Time

We use geodesic distance from the GeoPy Python library [25] to estimate the distance between properties from the latitudes and longitudes of the properties. We also evaluated OpenStreetMap [32] for calculating the distances. The Pearson correlation between the distances calculated using both methods on a random sample (100,000 pairs of properties) was 0.98. We decided to use geodesic distance as it is significantly faster than OpenStreetMap, and we have a relatively large number of property pairs for potential targeting (see subsection 4.3). With the distances, we used the OR-tools library by Google [28] to obtain an approximately optimal route to visit a set of properties. We record the estimated time (i.e., outreach time) to complete visits to the properties. We consider the following piece-wise speed function to estimate the time required to travel from one property to another. The speed gradually increases as we increase the distance. Miles (m) represents the unit for the distance and mile-per-hour (mph) for speed. For distance d , speed s :

Table 4: The combination of hyper-parameters for Random Forest Models

Model	Features	Parameters			
		Measure of impurity	Max depth	Minimum samples in a leaf node	n_estimators
Pre-COVID	Eviction	Entropy	5	2	500
	Eviction and Neighborhood	Log Loss	5	2	500
	Eviction, Neighborhood, and Owner	Entropy	5	2	500
Post-COVID	Eviction	Gini	3	2	100
	Eviction and Neighborhood	Log Loss	2	2	100
	Eviction, Neighborhood, and Owner	log Loss	5	2	100

Table 5: The combination of hyper-parameters for XGBoost Models

Model	Features	Parameters				
		Max depth	Learning Rate	n_estimators	scale_pos_weight	gamma
Pre-COVID	Eviction	3	0.05	100	1	0.1
	Eviction and Neighborhood	3	0.1	100	1	0.1
	Eviction, Neighborhood, and Owner	2	0.1	100	1	0.1
Post-COVID	Eviction	3	0.05	100	3	0.05
	Eviction and Neighborhood	3	0.05	100	1	0.05
	Eviction, Neighborhood, and Owner	3	0.05	100	5	0.05

$$s(\text{mph}) = \begin{cases} 4 & d \leq 1 \\ 15 & 1 < d \leq 3 \\ 30 & 3 < d \leq 5 \\ 55 & d > 5 \end{cases} \quad (1)$$

We estimate the actual time taken during a visit to a unit as n hours. We set n to 6 minutes or 0.1 hours. The estimate was based on data collected by community partners during door-to-door canvassing to inform tenants of COVID-19 rental relief. Canvassers knocked on an average of 35 doors during six-hour shifts, speaking with residents and/or leaving flyers with resource information. For a pair of properties p_1, p_2 with units u_1, u_2 , and distance $d_{1,2}$, the total time $t(p_1, p_2)$ to visit properties p_1 and p_2 is,

$$t(p_1, p_2) = n * u_1 + s * d_{1,2} + n * u_2 \quad (2)$$

Thus, the routing policies vary by the time needed to travel between properties.

B.2 Sensitivity Analysis on Outreach Time

We conduct a simulated experiment to assess the sensitivity of our results to assumptions about outreach time. We consider ablations of two variables, the time spent visiting a unit ($n = 3, 6, 10$ minutes, where 6 is the baseline calibrated to data reported in the main text), and speed of travel between parcels, where we consider a second speed s' , slower than the baseline in the text that could reflect speed of travel in a more dense urban setting.

Table 6: Sensitivity Analysis of estimated outreach time on discovered evictions. Discovered evictions by routing policies for different speeds and visitation time (n) when total outreach time is normalized by NEO-T-O.

Routing Policies	Speed	Time to visit unit (n)		
		3 min	6 min	10 min
Previous Eviction Count	s	481	731	824
	s'	498	701	779
Neighborhood Based	s	654	863	870
	s'	656	860	870

$$s'(\text{mph}) = \begin{cases} 2 & d \leq 1 \\ 10 & 1 < d \leq 3 \\ 20 & 3 < d \leq 5 \\ 45 & d > 5 \end{cases} \quad (3)$$

We compare our policy NEO-T-O with the Previous Eviction Count and Neighborhood Based routing policies. For each setting, the outreach time is constrained to be the same as the time required by NEO-T-O. The number of discovered evictions of 936 remains the same for NEO-T-O for all of the six settings, and changes for other policies. NEO-T-O outperforms both of the policies in all of the settings. Its relative advantage is greater for shorter visit times, and not substantially influenced by the speed of travel between parcels.