Measuring Neighborhood Change Using Postal and Housing Choice Voucher Data: Results from a Pilot Analysis of Four Metropolitan Areas in Washington, D.C. and Ohio

Alena Stern
Urban Institute

Abstract

This article presents the results of a pilot effort to model neighborhood change in near real-time by supplementing time-lagged demographic data from the American Community Survey (ACS) with real-time U.S. Postal Service (USPS) and Housing Choice Voucher (HCV) data. The author first defines and measures three key types of neighborhood change—gentrification, decline, and inclusive growth—in the selected metropolitan statistical areas (MSAs). She then uses machine learning methods to create a model that identifies neighborhood change at the census tract level. The model identifies neighborhood change with 76 percent accuracy and 76 percent precision; that precision exceeds models trained with just ACS data or just USPS and HCV data. The model is strong at predicting neighborhood decline and less accurate at identifying gentrifying neighborhoods. These results suggest a promising application of the USPS and HCV data to model neighborhood change.

Background

Measuring neighborhood change in real time is critical to enable timely policy action to prevent displacement in gentrifying communities, intervene to mitigate community decline, and encourage inclusive growth. Many previous efforts to measure neighborhood change across jurisdictions rely on nationwide administrative datasets such as the decennial census or American Community Survey (ACS) (e.g., Bates, 2013; Bostic and Martin, 2003; Chapple, 2009; Ellen and Ding, 2016; Freeman, 2005; Steif et al., 2016; and Thomas, et al., 2020). These data are published with considerable time lags that do not allow for real-time analysis. To achieve more timely results, some studies have focused on one jurisdiction using more frequently updated local datasets such as parcel-level files or building permits that are not comparable across jurisdictions (e.g., Data Driven Detroit, 2012; Hetrick et al., 2013; Raleigh and Galster, 2015). Accordingly, methodologies are needed to identify neighborhood change in closer to real time that are applicable across jurisdictions.

The U.S. Department of Housing and Urban Development (HUD) collects two data sources that could help identify neighborhood change in real time across jurisdictions. First, HUD has received quarterly aggregate data since 2005 on total business and residential addresses and counts of addresses identified as having been "Vacant" or "No-Stat"[1] in the previous quarter, which may proxy both investment and disinvestment in neighborhoods over time (Cohen and Pettit, 2019).[2,3] Second, HUD collects real-time data on the administration of its Housing Choice Voucher (HCV) program, including the locations where vouchers are used—that may show changes in the low-income renter population and landlords' willingness to accept vouchers as a result of rising rents (Cohen and Pettit, 2019).[4] Because both of these data sets cover all neighborhoods in the United States and are updated frequently, they offer potential power to understand neighborhood change and the impact of policy in near-real time. The author assesses the explanatory power of just the USPS and HCV data alone to identify neighborhood change and the combination of the two with time-lagged demographic data from the ACS.

Data Sources

USPS Data

The HUD Aggregated United States Postal Service (USPS) Data on Address Vacancies (USPS data) provides quarterly counts at the USPS ZIP9 (or ZIP+4)[5] level of residential, business, and "other" addresses that were vacant or no-stat in the previous quarter and the count of total addresses. In addition, they report the count of addresses that have been vacant or no-stat for different time intervals (e.g. 12–24 months) and the median number of months that addresses in a given ZIP9 have been vacant or no-stat. The author aggregates the 2010–2019 quarterly ZIP9 data to the tract level for analysis using a crosswalk between the ZIP9 and census geographies provided by USPS to HUD.[6]

[1] The HUD USPS data documentation outlines several reasons that addresses can be classified as "No-Stat" including "Rural Route addresses that are vacant for 90 days or longer; Addresses for businesses or homes under construction and not yet occupied; Addresses in urban areas identified by a carrier as not likely to be active for some time." For more information, see: https://www.huduser.gov/portal/datasets/usps.html

[2] In their 2019 *Guide to Measuring Neighborhood Change to Understand and Prevent Displacement*, Mychal Cohen and Kathryn L.S. Pettit outline key indicators for neighborhood change with associated data sources, including the USPS vacancy data and HUD housing choice voucher data.

[3] The USPS reports aggregates at the ZIP+4 level of aggregation. HUD aggregates these data to the Census tract level and makes them available to governmental entities and non-profit organizations quarterly. For more information, see https://www.huduser.gov/portal/datasets/usps.html

[4] HUD aggregates these data to the Census tract level and releases them to the public on an annual basis. The data were aggregated at the block level on a quarterly basis by HUD staff and made available to the researcher for the purposes of this analysis. For more information, see https://www.huduser.gov/portal/datasets/assthsg.html

[5] For more information on ZIP Codes, see https://faq.usps.com/s/article/ZIP-Code-The-Basics

[6] Quarterly tract-level USPS data beginning in 2005 to present are available here for researchers and practitioners: https://www.huduser.gov/portal/datasets/usps.html

Housing Choice Voucher Data

The author uses aggregated data on the count of Housing Choice Voucher (HCV) tenants by tract from 2010–2019. The aggregated data include quarterly counts of tenant-based vouchers (TBV), project-based vouchers (PBVs), homeownership program vouchers (HV), and total vouchers in each tract.[7]

Resident Characteristics Data

The author uses data on resident characteristics from the 5-year American Community Survey (ACS) at the tract level for two purposes in this analysis. First, the ACS data is used to categorize neighborhoods into different neighborhood change types for this analysis. To measure actual neighborhood change from 2013–2018, changes between the 2009–2013 ACS and the 2014–2018 ACS—the most recent 5-year ACS available at the time of analysis[8]—are used. Second, the author uses the ACS to produce variables used in some of her prediction models. She uses the actual neighborhood change between 2013–2016 for these variables, measured using change between the 2009–2013 ACS and 2012–2016 ACS as the latter represents the most recent ACS 5-year data available in the classification year of 2018. One limitation of using 5-year estimates with overlapping years is that it could underestimate change between 2013–2016.[9] The author chose tract-level analysis because the margins of error of the relevant ACS variables at the block group level were too high to reliably identify neighborhood change types for the selected metropolitan statistical areas (MSAs).

Analysis Method

The objective of this analysis is to identify, as of December 2018, the type of change that each neighborhood experienced between 2013–2018. For this pilot, the author focuses on neighborhoods in four MSAs: Washington-Arlington-Alexandria, DC-VA-MD-WV; Youngstown-Warren-Boardman, OH-PA; Cleveland-Elyria, OH; and Akron, OH. The author chose these pilot MSAs to have a variety of neighborhood change types represented in places where she had connections with local experts who could help validate the neighborhood change assignments. She conducts this analysis in the following steps:

Defining Neighborhood Change Types

For the purpose of the analysis, the author defines three mutually exclusive types of neighborhood change (see Appendix A for detailed definitions):

---

[7] The author uses the counts of different types of vouchers to test whether certain voucher types are more predictive of neighborhood change (for example, whether tenant-based vouchers are more sensitive to displacement via gentrification than homeownership vouchers). Annual tract level data from 2009–2019 on counts of Housing Choice Voucher tenants are available here: https://www.huduser.gov/portal/datasets/assthsg.html#2009-2019_data; Data on current voucher tenants by tract updated quarterly are available here: https://hudgis-hud.opendata.arcgis.com/datasets/housing-choice-vouchers-by-tract

[8] The 2014–2018 5-year ACS was published in January 2020.

[9] See chapter 4 of the ACS handbook for more details: https://www.census.gov/content/dam/Census/library/publications/2019/acs/acs_aian_handbook_2019_ch04.pdf

- **Gentrifying:** The author adapts the definition of gentrification from Bates (2013). She first identifies eligible neighborhoods where gentrification and displacement risk is high at the beginning of the time period (2013) based on the average rent, home values, proportion of renters, education level, and proportion of low-income households.[10] The author then identifies eligible neighborhoods as gentrifying where the low-income population decreases (as a proxy for displacement) and the change in the proportion of homeowners, individuals with a bachelor's degree, non-low income households, and change in average rents and home values are greater than the median change in the metropolitan area.

- **Declining:** The author adapts her definition of decline from Stancil (2019) and Data Driven Detroit (2012). She first identifies eligible neighborhoods as those not in the highest quartile of rents, home values, and household incomes in 2013. The author then identifies the eligible neighborhoods as declining where the total population decreases and the change in the proportion of addresses that are vacant and households that are low-income is greater than the median change in the metropolitan area.

- **Inclusively Growing:** The author adapts her definition of inclusively growing from Stancil (2019). All neighborhoods are eligible, and she identifies inclusively growing neighborhoods as those where the change in the number of low-income and non-low-income households is positive and greater than the median change for the metropolitan area.

Because the criteria used to evaluate membership in these three neighborhood change types (referred to as "classes") are mutually exclusive, a given neighborhood can only belong to one class (see Appendix A for more detail). All neighborhoods that do not exhibit any of the three types of neighborhood change defined previously are categorized as unchanging. The author tested and refined the definitions through feedback from experts in the selected MSAs. With their valuable input, she identified the final definitions that yield the following numbers (exhibit 1) of neighborhoods in each class.

Exhibit 1

Number and Percent of Neighborhood Types in Data—Total and by MSA

Neighborhood Type	Total	Akron, OH	Cleveland, OH	Washington, DC	Youngstown, OH
Gentrifying	37 (1.6%)	3 (1.8%)	8 (1.3%)	24 (1.8%)	2 (1.3%)
Declining	188 (8.1%)	15 (8.8%)	68 (10.7%)	81 (6.0%)	24 (15.5%)
Inclusively Growing	462 (19.9%)	15 (8.8%)	112 (17.6%)	325 (23.9%)	10 (6.5%)
Unchanging	1630 (70.3%)	137 (80.6%)	447 (70.4%)	927 (68.3%)	119 (76.8%)

MSA = metropolitan statistical area. Percentages may not add to 100 percent due to rounding.
Sources: Author's calculations based on analysis of American Community Survey (ACS) and U.S. Postal Service (USPS) data

[10] The author defines low-income households as those with a household income below 80 percent of the median household income in the MSA. See Appendix A for more detail.

Predicting Neighborhood Change Type

To predict the neighborhood change type, the author trained a machine learning model in the following steps:

Data Cleaning: The author imputed missing HCV data with 0 because a missing value for a given tract indicates that no HCV tenants were in that tract for the given quarter. She also pulled the ACS data at the tract level from the Census Application Programming Interface (API). The author imputed missing estimates[11] using the median value for the county in which the tract falls.

Feature Creation: The author used USPS, HCV, and ACS data to create different variables (called features) to use in the model to classify neighborhoods by the type of change. For each of the following feature categories, the author calculated numerous features using the input variables (such as vacant addresses, total vouchers, etc.) over different periods of change (for example, change over 12 months, change over 36 months, etc.). She also included several of the raw USPS and HCV variables for the quarter of the prediction date. A full list of features created can be found in Appendix B. The author trained models using three different feature groups: first, using just the HCV and USPS data, second, using just the ACS data from 2013–2016, and finally, using all three sources.

- **Change in Variable:** Measuring change in HCV tenant and USPS address variables from different start dates to the prediction date (Quarter [Q] 4, 2018). For the ACS features, the author measured the change between 2013–2016.

The following features are calculated only for the HCV and USPS data because they rely on quarterly data:

- **Consistency of Change in Count of Addresses/Voucher Tenants:** The author calculated whether the count of addresses or voucher tenants of a given type have changed consistently in one direction (increasing or decreasing) for a given number of consecutive quarters.

- **Change in Rate of Change of Addresses/Voucher Tenants:** The author calculated the difference between the change in the count of addresses/voucher tenants of a given type in two consecutive time periods (such as the difference in the change from 2017 to 2018 and 2016 to 2017).

- **Change in Bordering Neighborhoods:** The author calculated the average change in the neighborhoods that border a given neighborhood as of Q4, 2017 as an early warning sign for change.

Split Training and Test Data: To prevent overfitting the model to the training data, such that the model would effectively learn the idiosyncrasies of the training data but generalize poorly to new data, the author split the data into two sets: a training set, which she used to select features and train the model, and a test set, which she used to evaluate the model performance. The author

[11] In the ACS, missing data (with value -666666666) indicates "either no sample observations or too few sample observations were available to compute an estimate." See https://www.census.gov/data/developers/data-sets/acs-1year/data-notes.html for more details. In the data used in this analysis, 5 percent of observations had a missing variable.

randomly selected 70 percent of the observations (n = 1,621) to be the training set and reserve the remaining 30 percent of the observations (n = 696) to be the test set.

Data Scaling: The author normalized all the variables to z-scores within the train and test set separately to ensure that no variable is given more importance in the modeling simply by virtue of having larger values and/or range.

Feature Selection: Before fitting the model, the author selected a subset of features to use for modelling that best predicts the neighborhood type. She fit a logistic regression with the L1 norm on the training data to predict neighborhood change type. The L1 norm forces many of the feature regression coefficients toward zero. The author then used only features with coefficients above a threshold to fit the model. This feature selection process reduces the likelihood of overfitting by reducing the number of dimensions used in modeling.

Model Selection: The author tested a variety of feature selection, algorithm, and hyperparameter combinations, where each combination is a different model.[12] The hyperparameters fine-tune how a given algorithm will attempt to fit the data and control different dimensions of the model, such as overfitting.

She first selected the hyperparameters that yield the best precision on the training data for each algorithm.[13] The author only considered the gentrifying, declining, and inclusively growing classes when calculating precision. She used different probability thresholds for assigning observations to each class because of the considerable imbalance among the different classes.[14] This enables her to identify those neighborhoods that are *most likely* to belong to a given class among all neighborhoods in the analysis.

Model Evaluation: The author evaluated the best model for each algorithm on the test data by using the custom precision metric described previously, other precision thresholds, and overall accuracy to identify the final model which offers the best overall performance.

Results

Exhibit 2 provides performance metrics for the best model for each feature group, which in all cases used the Gradient Boosting Classifier algorithm. The author also implemented two simple baseline classification approaches against which she assessed the performance and improvement

[12] The full set of algorithms and hyperparameters tested are available upon request.

[13] Precision is defined as the ratio of true positives to all positives, or the proportion of neighborhoods identified by the model as experiencing a type of neighborhood change that actually undergo that change.

[14] For each neighborhood, the model assigns a probability that the neighborhood belongs to each class. The author then established a cutoff for each class based on the approximate prevalence of each class in the overall data and assign the top X percent of probabilities for a given class to that predicted class, where X is equal to the cutoff. For her analysis, the author set cutoffs of 1 percent, 5 percent, and 20 percent for the gentrifying, declining, and inclusively growing classes respectively. For example, the neighborhoods with the top 1 percent of probabilities of being gentrifying according to the model are predicted as gentrifying. If a given neighborhood meets the cutoff for multiple classes, the author assigned it to the class where the probability percentile is highest. For example, if for a given neighborhood the probability of being in the declining class is the 99th percentile for all neighborhoods, and the probability of being in the inclusively growing class is the 85th percentile among all neighborhoods, the author would assign that neighborhood as declining.

offered by the models. The previous ACS baseline assigns the actual neighborhood type from 2013–2016 as the predicted neighborhood change type in 2018; the neighborhood change baseline uses a set of simple rules[15] for change in vouchers, total addresses, and active addresses between 2013–2018 to predict neighborhood change type in 2018. The full results for all the algorithms tested can be found in Appendix C.

Exhibit 2

Performance of Best Model for Each Model Approach and Baselines

Model	Training Precision—Class Threshold	Test Precision—Class Threshold	Test Precision—Top 1 Percent[a]	Test Precision—Top 2 Percent	Test Precision—Overall	Test Accuracy[b]
ACS Features Only	0.60	.62	.67	.64	.66	.76
USPS and HCV Features Only	0.37	.40	.35	.39	.58	.70
USPS, HCV, and ACS Features	0.60	.64	.72	.66	.76	.76
Previous ACS Baseline	n/a[c]	.47	.39	.50	.56	.74
Neighborhood Change Baseline	n/a	.37	.33	.37	.37	.67

ACS = American Community Survey. HCV = Housing Choice Voucher. N/a = data not available. USPS = United States Postal Service.

[a]For the top X percent precision, the author assigned the neighborhoods with the highest X percent of probabilities of belonging to a given class assigned by the model to that class for gentrifying, declining, and inclusively growing, and the author assigned the remaining observations to unchanged.

[b]Unlike the precision calculations, the test accuracy includes the model's accuracy at both identifying the presence of change (class = gentrifying, declining, or inclusively growing) and the absence of change (class = unchanged). One can see that the test accuracy is higher than the precision for both the class threshold and top 1 percent, given the large proportion of neighborhoods that are "unchanged" in the study's time period.

[c]There is no training score for the author's baseline approaches because no model is trained; instead, a set of rules is applied to the test data.

Source: Author's calculations based on model results

One can see that the model with USPS, HCV, and ACS features performs better than the models that include ACS features alone and USPS and HCV features alone, offering increases in precision of 10 percent and 18 percent, respectively—and out-performing both baselines by a considerable margin. This finding suggests that bringing in timelier USPS and HCV data can add considerable classification power to models using time-lagged ACS data. The author selected this model with USPS, HCV, and ACS features as the "final model," which she will discuss for the remainder of this article.

The confusion matrix in exhibit 3 shows the accuracy of the final model's predictions[16] for each class. Because the author focused on precision using the custom class thresholds, she expected

[15] Neighborhoods were identified as gentrifying if they experienced a decrease in vouchers, increase in active addresses, and increase in total addresses in 2013–2018; declining if they experienced a decrease in vouchers, decrease in active addresses, and stable total addresses in 2013–2018; inclusively growing if vouchers, active addresses, and total addresses all increased. For purposes of this definition, the author considered increasing to be a change that is more than one-half of a standard deviation above the median change in the metro area, decreasing to be a change that is more than one-half of a standard deviation below the median change in the metro area, and stable to be within one-half a standard deviation of the mean.

[16] Exhibits 3–5 use the class thresholds to assign the neighborhoods to classes as described.

to see some neighborhoods that actually changed identified as unchanging but hoped that those neighborhoods predicted to change actually did so.

Exhibit 3

Test Data Confusion Matrix for Final Model

		Predicted Neighborhood Change Type			
		Unchanged	Gentrifying	Declining	Inclusively Growing
Actual Neighborhood Change Type	Unchanged	420	2	9	47
	Gentrifying	9	1	0	0
	Declining	33	0	23	2
	Inclusively Growing	61	2	1	86

Source: Author's calculations based on model results

One can see that the final model has the best precision at the custom class thresholds for declining neighborhoods (70 percent), followed by inclusively growing neighborhoods (64 percent) and much worse precision for gentrifying neighborhoods (20 percent). This result may be because very few gentrifying neighborhoods were present in the data (1.5 percent), giving the model very limited data to effectively learn patterns of gentrification. Future efforts should include more metropolitan areas with significant gentrification to improve results. When one examines the model accuracy and precision by MSA (exhibit 4), they can find that the largest MSA (Washington, DC) has the best precision. An area for future analysis could be first clustering MSAs to assess whether training separate models on groups of more similar MSAs improves performance. Maps of the predicted and actual neighborhood change types and model accuracy can be found in Appendix D.

Exhibit 4

Model Performance by MSA

	Akron, OH (%)	Cleveland, OH (%)	Washington, DC (%)	Youngstown, OH (%)
Accuracy	75.5	79.8	73.9	80.5
Precision*	50.0	61.4	65.8	60

**Only includes gentrifying, declining, and inclusively growing classes.*
MSA = Metropolitan Statistical Area.
Source: Author's calculations based on model results

Looking at performance within each MSA, the author examined whether the model performs differentially for neighborhoods by race and ethnicity composition (exhibit 5). Although the author does not include race and ethnicity in the features, considerable research has shown that models can learn race and ethnicity, in some cases yielding outcomes that exacerbate existing disparities (Gillis and Speiss, 2019; Turner Lee, Resnick, and Barton, 2019).

Exhibit 5

Model Performance by Race

		Gentrifying				Declining				Inclusively Growing			
	All Tracts (%)	True Pos (%)	False Pos (%)	All Pos (%)	False Neg (%)	True Pos (%)	False Pos (%)	All Pos (%)	False Neg (%)	True Pos (%)	False Pos (%)	All Pos (%)	False Neg (%)
Black	22	18	87	58	44	41	51	43	16	27	26	27	30
White	55	38	9	23	24	45	40	44	65	44	50	46	44
Asian	7	0	0	0	5	3	1	2	5	10	7	9	8
Hispanic	11	39	3	17	25	9	5	8	11	15	13	15	16
All Other Races	3	5	1	3	3	3	3	3	4	4	4	4	4

Source: Author's calculations using 2014–2018 ACS estimates and model results

First, the author found that neighborhoods falsely identified as both gentrifying and declining have larger Black populations than those correctly identified—87 percent vs. 18 percent for gentrifying and 51 percent vs. 41 percent for declining. She also found that the false negatives for gentrifying have larger Hispanic and smaller Black populations than those neighborhoods the model identified as gentrifying (25 percent vs 17 percent Hispanic and 44 percent vs 58 percent Black). However, the author cautions against generalizing these results given the very small number of actual and predicted gentrifying neighborhoods in the test data (10 and 5 respectively). The author also found that false negatives for declining have larger White and smaller Black populations than those identified as by the model as declining (65 percent vs 44 percent White and 43 percent vs 16 percent Black). The author found that the race and ethnicity composition was more similar across inclusively growing subgroups. These racial differences could be driven by many factors, such as income disparities by race and legacies of redlining and racial segregation (Rothstein, 2017; Wilson, 2020). Such racial equity implications must be evaluated in any decision to use such models to inform future resource allocation to address neighborhood change.

For the study's top-performing model, the features that offered the most predictive[17] power (in order) are as follows in exhibit 6.

[17] This is assessed using the impurity-based feature importances from the best Gradient Boosting Classifier model. The impurity-based feature importance is computed as the (normalized) total reduction of the error criterion brought by that feature. The higher the value, the more important the feature.

Exhibit 6

Feature Importance

Feature Name	Data Source
Change in percent of residential addresses that are vacant for 3 months or less—36 months	USPS
Change in the percent of residential addresses that are active—48 months	USPS
Change in the total number of no-stat business addresses—48 months	USPS
Change in the rate of change of total vouchers—2 months	HCV
Change in the percent of residential addresses that are no-stat for 3 months or less—36 months	USPS
Change in the average number of days residential addresses are no-stat—48 months	USPS
Whether a tract is eligible for gentrification in 2013	ACS
Change in the number of individuals earning above 80 percent of AMI from 2013–2016	ACS
Change in the total number of active residential addresses in neighboring tracts from Q4 2014 to Q4 2017	USPS

ACS = American Community Survey. AMI = area median income. HCV = Housing Choice Voucher. USPS = United States Postal Service.
Source: Author's calculations based on model results

Interestingly, one can see that many of the most important variables are United States Postal Service and Housing Choice Voucher variables, despite the fact that the American Community Survey-only model performed much better than the USPS and HCV-only model. This finding also underscores the value that these variables add to the ACS variables in identifying neighborhood change.

Conclusion

This analysis provides promising early evidence of the value added by incorporating USPS and HCV data with traditionally used demographic data from the ACS to measure neighborhood change in near real-time. Future analysis could replicate this work using the public USPS and HCV data to determine whether these results hold when applying the model to a broader set of metropolitan areas and identifying neighborhood change with new years of data (such as the recently released 2019 5-year ACS). This approach also has several areas for future improvement, including incorporating other datasets to create features,[18] investigating the effect of uncertainty in the ACS estimates at the tract level on neighborhood type assignment and model performance, and including more geographically diverse MSAs with larger numbers of gentrifying tracts to improve the model's performance for classifying gentrification. Analysts should validate the resulting classifications with local experts and community members.

[18] Although the focus of this pilot was on the USPS and HCV data, previous research on neighborhood change has used many other data sources that may enhance the precision of future models such as additional administrative data sources (e.g. LEHD Origin-Destination Employment Statistics (LODES), Home Mortgage Disclosure Act (HMDA)), and private sources (e.g. Zillow). See Cohen and Pettit (2019) for further ideas.

Appendix A: Neighborhood Change Definitions

- **Gentrifying:** The author adapted her definition of gentrification from the definition advanced in Bates (2013). She first identified neighborhoods that are eligible for gentrification at the beginning of the study's time period (2013) as those that meet at least two of the three conditions:

    1. The proportion of households that rent is above the median for the metropolitan area.
    2. The proportion of individuals over 25 years old without a bachelor's degree is above the median for the metropolitan area.
    3. The proportion of households with income below 80 percent of the median household income in the metropolitan area[19] is above the median for the metropolitan area.

Eligible neighborhoods must also meet both of the following conditions:

1. Average rents are below the median for the metropolitan area.
2. Average home values are below the median for the metropolitan area.

Eligible neighborhoods are considered to be gentrifying if, at the end of the study's time period (2018), one of the following two conditions is met:

1. Change in the proportion of households that are homeowners is greater than the median change in the metropolitan area.
2. Change in the proportion of individuals over 25 years old with a bachelor's degree is greater than the median change in the metropolitan area.

Gentrifying neighborhoods must also meet all of the following four conditions:

1. Change in the proportion of households with incomes above 80 percent of the median income in the metropolitan area is greater than the median change in the metropolitan area.
2. The change in the average rent is greater than the median change in the metropolitan area.
3. The change in the average home values is greater than the median change in the metro area.
4. The number of low-income individuals living in the tract decreases.

- **Declining:** The author adapted her definition of decline from the definitions advanced in Stancil (2019) and Data Driven Detroit (2012). She first identified the neighborhoods

[19] The author used the ACS data to identify the median household income for each of the MSAs. The ACS reports the count of households in many different income buckets (e.g. less than $10,000, $10,000–$14,999, etc.). The author identified the bucket whose upper-bound is closest to 80 percent of the MSA median income without exceeding that threshold and summed the counts of households in all buckets from the lowest (less than $10,000) through that bucket to calculate the count of households with incomes below 80 percent of the median in that MSA. This calculation likely slightly underestimates households, given that the author did not count the households in the bucket into which the threshold falls.

eligible for decline at the beginning of the study's time period (2013) as those that meet the following conditions:

1. The average home value is below the 75th percentile average home value for the metropolitan area.

2. The average rent is below the 75th percentile average rent for the metropolitan area.

Eligible neighborhoods are considered to be declining if at the end of the study's time period (2018) the following three conditions are met:

1. Change in the percent of addresses that are vacant is greater than the median change in the metropolitan area.

2. The proportion of households with incomes below 80 percent of the median income in the metropolitan area changes by more than the median change in the metropolitan area.

3. The total population decreases.

- **Inclusively Growing:** The author adapted her definition of inclusively growing from the definition advanced in Stancil (2019). All neighborhoods are considered eligible, and those neighborhoods that meet the following conditions at the end of the study's time period (2018) are considered to be inclusively growing:

    1. The change in the number of households with incomes below 80 percent of the median income in the metropolitan area is positive and greater than the median change for the metropolitan area.

    2. The change in the number of households with incomes above 80 percent of the median income in the metropolitan area is positive and greater than the median change for the metropolitan area.

Appendix B: List of Features Produced for Analysis

Exhibit B1

Addresses

Feature Type	Address Type	Address Status Type	Change Duration (months)	Status Duration (months)
Change in Count/ Proportion Addresses	Residential, Business	Total,[a] Active,[b] No-Stat, Vacant	12, 36, 48	All addresses
Change in Count/ Proportion of Addresses by Duration	Residential, Business	Total, Active, No-Stat, Vacant	12, 36, 48	3 or less, 36 or more
Change in Neighbors as of December 2017	Residential, Business	Active, Vacant	12, 36	3 or less, 36 or more
Change in the Rate of Change	Residential, Business	Active, Vacant	12, 36	All addresses
Consistency of Change	Residential, Business	Active, Vacant	36, 60 (for 12-month increments)	All addresses

[a]*Proportion is not calculated for total addresses as the proportion would always be equal to 1.*
[b]*Active is the total number of addresses minus no-stat and vacant.*

Exhibit B2

Vouchers

Feature Type	Voucher Type	Change Duration (months)
Change in Count/ Proportion[a] of Vouchers	TBV, PBV, HV, Total	12, 36, 48
Change in Neighbors	TBV, PBV, HV, Total	12, 36
Change in the Rate of Change	TBV, PBV, HV, Total	12, 36
Consistency of Change	TBV, PBV, HV, Total	36, 60 (for 12-month increments)

[a]*Proportion is calculated as the ratio of the count of vouchers to total residential addresses.*
HV = housing voucher. PBV = project-based voucher. TBV = tenant-based voucher.

Change in Other Columns:

- Columns: median days business addresses vacant, median days residential addresses vacant, average days residential addresses no-stat, average days business addresses no-stat
- Change Duration (months): 12, 36, 48

Raw Columns:

- Percent vacant
- Percent vacant longer than 12 months
- Percent vacant longer than 24 months
- Tenant-Based Voucher

- Project-Based Voucher
- Housing Voucher
- Voucher
- Total active residential addresses
- Total vacant residential addresses
- Total active business addresses
- Total vacant business addresses
- Count of no-stat addresses for 3 months or less
- Count of no-stat addresses for 36 months or more
- Count of vacant addresses for 3 months or less
- Count of vacant addresses for 36 months or more

American Community Survey Variables (all change between 2013–2016):

- Average home value (absolute and percent change)
- Average rent (absolute and percent change)
- Households earning above 80 percent Area Median Income
- Households earning below 80 percent AMI
- Proportion of households earning below 80 percent AMI
- Proportion of population over 25 years old with a bachelor's degree
- Proportion of households that are owners
- Eligibility for gentrification in 2013
- Eligibility for decline in 2013

Appendix C: Full Modelling Results

Algorithm	Training Precision—Class Threshold	Test Precision—Class Threshold	Test Precision—Top 1 Percent	Test Precision—Top 2 Percent	Test Precision—Overall	Test Accuracy[a]
USPS and HCV Features Only						
Decision Tree	0.26	.37	.33	.33	.48	0.69
Random Forest	0.37	.39	.39	.38	.53	0.70
Gradient Boosted Tree	0.37	.40	.35	.39	.58	0.70
Logistic Regression	0.37	.40	.39	.33	.37	0.66
K-Nearest Neighbors	0.28	.24	.22	.21	.46	0.69
USPS, HCV, and ACS Features						
Decision Tree	0.57	.60	.44	.38	.55	0.72
Random Forest	0.61	.61	.67	.69	.65	0.77
Gradient Boosted Tree	0.60	.64	.72	.66	.76	0.76
Logistic Regression	0.60	.59	.61	.56	.59	0.71
K-Nearest Neighbors	0.38	.33	.39	.37	.55	0.69
Baseline Models						
Previous ACS Baseline	n/a	.47	.39	.50	.56	.74
All ACS Baseline	0.60	.62	.67	.64	.66	.76
Neighborhood Change Baseline	n/a	.37	.33	.37	.37	.67

ACS = American Community Survey. HCV = Housing Choice Voucher. USPS = United States Postal Service.

[a]Unlike the precision calculations, the test accuracy includes the accuracy at both identifying the presence of change (class = gentrifying, declining, or inclusively growing) and the absence of change (class = unchanged). one can see that the test accuracy is higher than precision given the large proportion of neighborhoods that are "unchanged" in the data.

Source: Author's calculations based on model results

Appendix D: Accuracy Maps

Akron, OH MSA True Neighborhood Type **Akron, OH MSA Predicted Neighborhood Type**

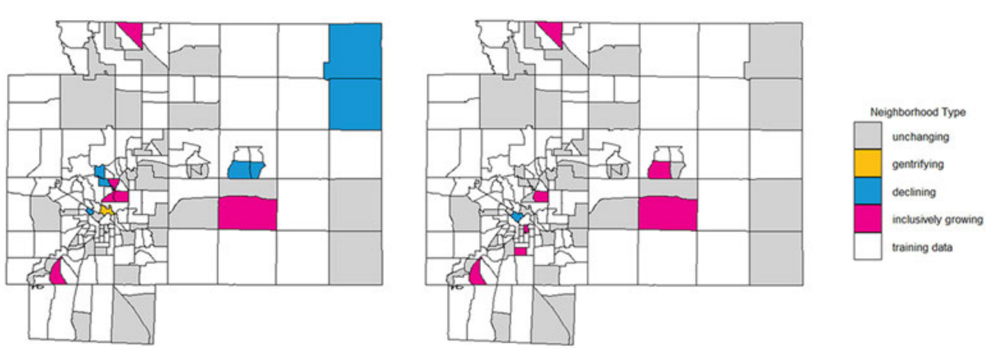

Cleveland, OH MSA True Neighborhood Type   Cleveland, OH MSA Predicted Neighborhood Type

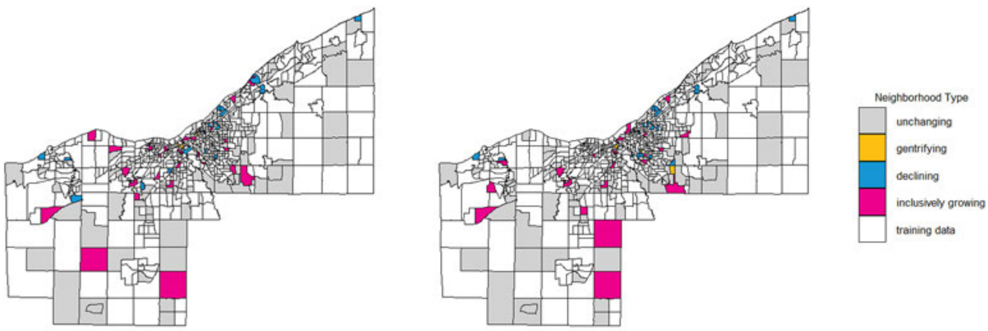

Youngstown, OH MSA True Neighborhood Type   Youngstown, OH MSA Predicted Neighborhood Type

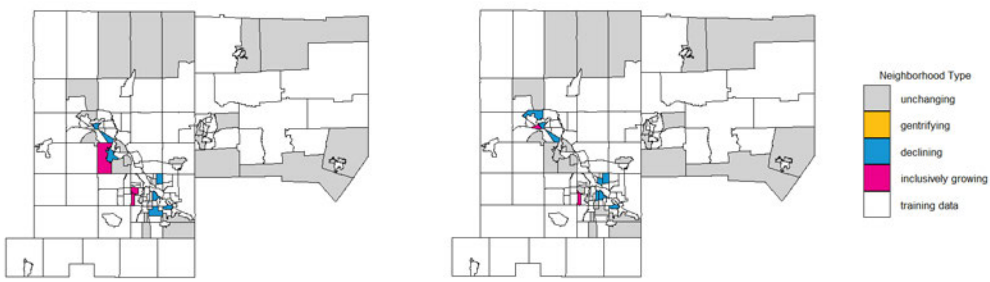

Washington, DC MSA True Neighborhood Type   Washington, DC MSA Predicted Neighborhood Type

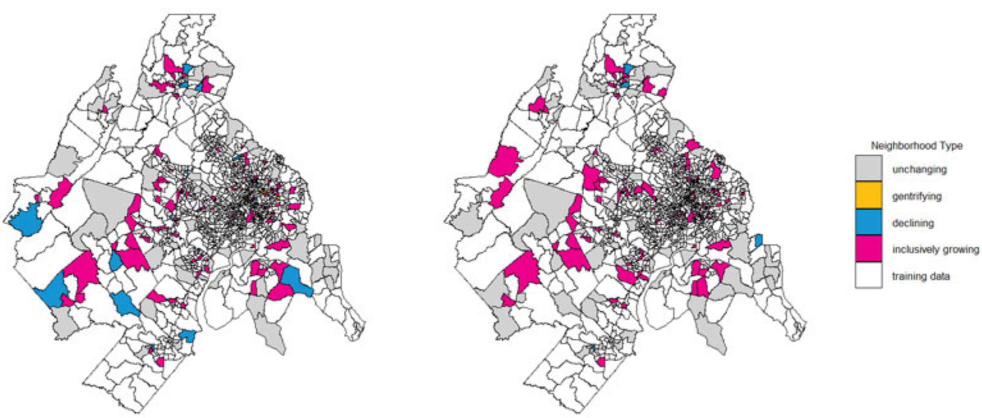

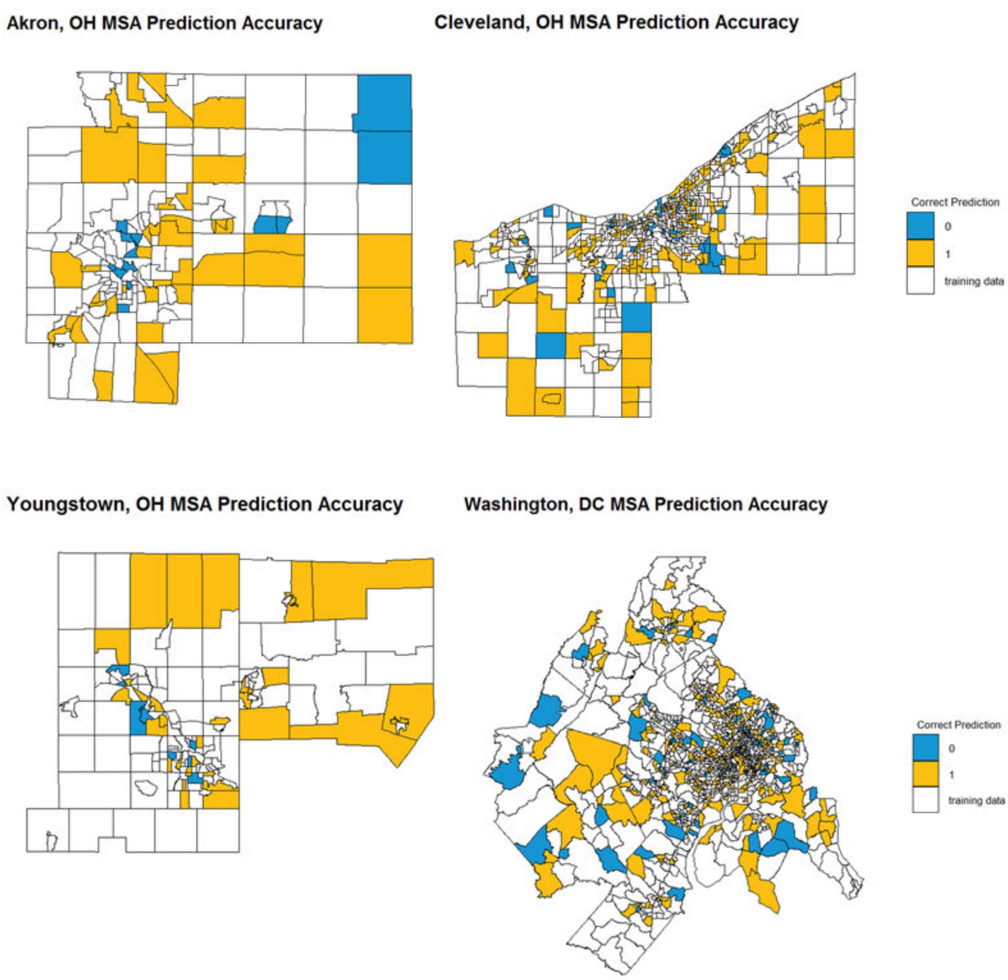

Source: Author's analysis of model results

Acknowledgements

The author would like to thank Todd Richardson, Pam Blumenthal, Ted Payne, Graham MacDonald, Mary Cunningham, and Margery Austin Turner for their expert review and guidance; John Bralich, Ian Beniston, and Sarah Lowry for providing feedback on the neighborhood change types; and John Huggins, Alex Din, Todd McNeil, Shanti Karimsetti, Katherine Winger, and Jacquie Bachand for data and technical assistance.

Author

Alena Stern is a senior data scientist at the Urban Institute. Stern is available by email at astern@urban.org.

References

Bates, Lisa K. 2013. *Gentrification and Displacement Study: Implementing an Equitable Inclusive Development Strategy in the Context of Gentrification*. Portland, OR. Portland, OR: Portland State University.

Bostic, Raphael W., and Richard W. Martin. 2003. "Black Home-Owners as a Gentrifying Force? Neighbourhood Dynamics in the Context of Minority Home-Ownership." *Urban Studies* 40 (12): 2427–2449.

Chapple, Karen. 2009. *Mapping Susceptibility to Gentrification: The Early Warning Toolkit*. Berkeley, CA: Center for Community Innovation. https://www.reimaginerpe.org/files/Gentrification-Report%284%29.pdf

Cohen, Mychal, and Kathryn L.S. Pettit. 2019. *Guide to Measuring Neighborhood Change to Understand and Prevent Displacement*. Washington, DC: Urban Institute. https://www.urban.org/sites/default/files/publication/100135/guide_to_measuring_neighborhood_change_to_understand_and_prevent_displacement.pdf

Data Driven Detroit. 2012. *CDAD Neighborhood Revitalization Strategic Framework Process: Composition and Explanation of Data Sets*. Detroit, MI: Data Driven Detroit https://datadrivendetroit.org/files/2012/03/CDAD_Strategic_Framework_Data_Explanation.pdf

Ellen, Ingrid Gould, and Lei Ding, 2016. "Advancing Our Understanding of Gentrification," *Cityscape* 18 (3): 3–8.

Freeman, Lance. 2005. "Displacement or Succession?: Residential Mobility in Gentrifying Neighborhoods," *Urban Affairs Review* 40 (4): 463–491.

Gillis, Talia, and Jann Speiss. 2019. "Big Data and Discrimination," *The University of Chicago Law Review* 86 (2): 459–487.

Hetrick, Tom, Ian Beniston, John Bralich, Jack Daugherty, Karen Schubert, Eartha Terrell, William D'Avignon, T. Sharon Woodberry, Jason Whitehead, and Karen Perkins. 2013. *Neighborhood Conditions Report*. Youngstown, OH: Youngstown Neighborhood Development Corporation. http://yndc.org/sites/default/files/NCR_EXEC_SUMMARY_140206_0.pdf

Raleigh, Erica, and George Galster. 2015. "Neighborhood Disinvestment, Abandonment, And Crime Dynamics," *Journal of Urban Affairs* 37: 367–396. https://doi.org/10.1111/juaf.12102

Rothstein, Richard. 2017. *The Color of Law: A Forgotten History of How Our Government Segregated America*. https://www.epi.org/publication/the-color-of-law-a-forgotten-history-of-how-our-government-segregated-america/

Stancil, William. 2019. *American Neighborhood Change in the 21st Century*. Minneapolis, MN: Institute on Metropolitan Opportunity—University of Minnesota Law School

Steif, Ken, Alan Mallach, Michael Fichman, and Simon Kassel. 2016. *Predicting Gentrification Using Longitudinal Census Data*. Philadelphia, PA: Urban Spatial Analysis. http://urbanspatialanalysis.com/portfolio/predicting-gentrification-using-longitudinal-census-data/

Thomas, Tim, Carson Hartmann, Anna Driscoll, Karen Chapple, Anna Cash, Renee Roy Elias, and Miriam Zuk. 2020. *The Urban Displacement Replication Project*. Berkeley, CA: The Urban Displacement Project. https://www.urbandisplacement.org/sites/default/files/udp_replication_project_methodology_10.16.2020-converted.pdf

Turner Lee, Nicole, Paul Resnick, and Genie Barton. 2019. *Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms*. Washington, DC: Brookings Institution. https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/

Wilson, Valerie. 2020. *Racial Disparities in Income and Poverty Remain Largely Unchanged Amid Strong Income Growth in 2019*. Washington, DC: Economic Policy Institute. https://www.epi.org/blog/racial-disparities-in-income-and-poverty-remain-largely-unchanged-amid-strong-income-growth-in-2019/