

Review

## Machine learning to model gentrification: A synthesis of emerging forms



Mueller Maya <sup>a</sup>, Hoque Simi <sup>a,\*</sup>, Hamil Pearsall <sup>b</sup>

<sup>a</sup> Department of Civil, Architectural, and Environmental Engineering, Drexel University, 3141 Chestnut St, PA 19104, USA

<sup>b</sup> Department of Geography and Urban Studies, Temple University, 1801 N Broad St, Philadelphia, PA 19122, USA

### ARTICLE INFO

**Keywords:**

Gentrification  
Machine learning  
Built environment  
Neighborhood change  
Computer vision

### ABSTRACT

Gentrification is a complex and context-specific process that involves changes in the built environment and social fabric of neighborhoods, often resulting in the displacement of vulnerable communities. Machine Learning (ML) has emerged as a powerful predictive tool that is capable of circumventing the methodological challenges that historically held back researchers from producing reliable forecasts of gentrification. Additionally, computer vision ML algorithms for landscape character assessment, or deep mapping, can now capture a wider range of built metrics related to gentrification-induced redevelopment. These novel ML applications promise to rapidly progress our understandings of gentrification and our capacity to translate academic findings into more productive direction for communities and stakeholders, but with this sudden development comes a steep learning curve. The current paper aims to bridge this divide by providing an overview of recent progress and an actionable template of use that is accessible for researchers across a wide array of academic fields. As a secondary point of emphasis, the review goes over Explainable Artificial Intelligence (XAI) tools for gentrification models and opens up discussion on the nuanced challenges that arise when applying black-box models to human systems. Abstract: Gentrification is a complex and context-specific process that involves changes in the built environment and social fabric of neighborhoods, often resulting in the displacement of vulnerable communities. Machine Learning (ML) has emerged as a powerful predictive tool that is capable of circumventing the methodological challenges that historically held back researchers from producing reliable forecasts of gentrification. Additionally, computer vision ML algorithms for landscape character assessment, or deep mapping, can now capture a wider range of built metrics related to gentrification-induced redevelopment. These novel ML applications promise to rapidly progress our understandings of gentrification and our capacity to translate academic findings into more productive direction for communities and stakeholders, but with this sudden development comes a steep learning curve. The current paper aims to bridge this divide by providing an overview of recent progress and an actionable template of use that is accessible for researchers across a wide array of academic fields. As a secondary point of emphasis, the review goes over Explainable Artificial Intelligence (XAI) tools for gentrification models and opens up discussion on the nuanced challenges that arise when applying black-box models to human systems.

### 1. Introduction

Gentrification is defined as the process where the social character and built landscape of a historically disinvested, inner-city neighborhood is transformed by an influx of people and capital (Glass, 1964; Smith, 1979). Gentrification is distinct from other forms of neighborhood evolution, such as in the case of incumbent upgrading where the residential housing stock is gradually renovated in situ by existing residents (Van Criekingen & Decroly, 2003, p. 2452).

Gentrification's transformation of the built environment emerges in relation to historical spatial patterns of investment and disinvestment, or

“geographies of opportunity” (Wilson, 2006). Opportunity, in this case, can refer to educational opportunity across school districts, employment opportunity, the proximity of health services and grocery stores, public transit options, and the overall quality of the neighborhood's infrastructure. Central to modern gentrification theory is the concept that structural attributes affect life quality and human outcomes, and that gentrification both flourishes in and contributes to increased inequity through the restructuring of local landscapes (Ansell, 2019; Wyly & Hammel, 2004; Zuk et al., 2018).

Information asymmetries inherent in the real estate market can exacerbate the gentrification process by creating disparities in

\* Corresponding author.

E-mail address: [sth55@drexel.edu](mailto:sth55@drexel.edu) (H. Simi).

knowledge and power between various stakeholders, such as developers, incumbent residents, potential homebuyers or renters, and local government (Krijnen, 2018). Without access to the private market data that drives the informed speculation of developers and investors, municipal planners and community groups frequently find themselves retroactively responding to gentrification's effects after development is well-advanced. For this reason, a large body of contemporary scholars work towards constructing a productive and reproducible model of the gentrification process that can extrapolate trends that reflect realities on-ground.

Research themes are multidisciplinary and encompass a wide array of topics, from examining the sociopolitical conflicts between gentrifiers and displacees (Brown-Saracino & Ghaziani, 2009; Pattillo, 2010), modeling the relationship between gentrification and broader economic trends in the housing market (Schaffer & Smith, 1986; Wyly & Hammel, 1999), identifying global patterns of gentrification outside of Anglo-sphere contexts (Lees, 2019; Slater, 2017), and finding ways to better integrate academic findings into municipal policy (Chapple, 2009; Chapple & Zuk, 2016).

For most studies, researchers are tasked to identify where and when gentrification is occurring. Quantitative models tend to focus on socio-economic shifts in the study area, such as tracking spikes in median income or rent (Atkinson, 2000; McKinnish et al., 2010), whereas more labor-intensive, mixed-methods studies integrate observation-based data in the form of community surveys or direct audits of visual neighborhood change (Chapple, 2009; Hammel & Wyly, 1996; Hwang & Sampson, 2014; Wyly & Hammel, 1998, 1999).

Many researchers have emphasized the need for better metrics, models, and validation techniques (Barton, 2016; Brown-Saracino, 2017a; Finio, 2021). Although qualitative studies can provide detailed ethnographic accounts of gentrification narratives, they are often locale-specific and time-intensive. Quantitative methods can analyze gentrification at a macro-scale with nationally available census datasets, but variables lack quality information on built characteristics and data on housing supply. The most widely applied quantitative methods are also limited in their applicability in light of gaps in the data and collinearities between variables (Royall & Wortmann, 2015). Without the capability to provide accurate forecasts and link specific urban features to the gentrification process, researchers face difficulties in communicating their theoretical insights to stakeholders in meaningful ways (Atkinson, 2008; Chapple & Zuk, 2016).

Recent advancements in Machine Learning (ML), a subset of Artificial Intelligence (AI), have opened new avenues for researchers to tackle the wicked problem of modeling gentrification (e.g., Palafox & Ortiz-Monasterio, 2020; Reades et al., 2019; Thackway et al., 2023a, 2023b; Yee & Dennett, 2022). Unlike non-learning-based methods, ML is powerful in its capacity to discern meaningful patterns from high-volume, stochastic datasets. These pioneering researchers emphasize the potential for their findings to help outline gentrification mitigation techniques, construct early warning systems for vulnerable communities, and place political pressure on stakeholders to consider the long-term, negative consequences of supporting pro-gentrification policies.

Progress in ML seems to come in leaps and bounds by the day, from the application of more computationally powerful algorithms like Extreme Gradient Boosting (Chen & Guestrin, 2016), to the ever-increasing sophistication of object recognition models (Tian, 2020). Gentrification researchers are integrating ML with a similarly rapid and eager pace, but there are few publications that reflect on the progress, untapped potential, and limitations of these novel approaches. Despite the pace of integration, there is a need for more critical evaluation and, additionally, a need to disseminate this knowledge in ways that are accessible to researchers without data science-specific backgrounds.

Although recent ML-based gentrification models have exercised a good deal of caution, the same cannot be said of other applications, such as the case with AI chatbots like Microsoft's Tay, Meta's Galactica, and ChatGPT that have been reported to perpetuate racist and sexist

stereotypes (Borji, 2023), or the proliferation of criminal profiling tools that use ML to classify suspects based on arbitrary patterns in facial proportions (Labi, 2012; Wu & Zhang, 2016). It is important to note that there is nothing inherently malicious about the structure of ML algorithms. However, due to the need for substantial amounts of input data and the lack of transparency in how the algorithms process this data to generate predictions, ML algorithms are susceptible to misuse and can easily learn biases without proper safeguards.

For modeling gentrification, a phenomenon with inexorable ties to race, class, and inequity, researchers face an ever-increasing responsibility towards applying ML in reflective and cautious ways to not do disservice to the very groups we aim to serve. Many of the ML models used in gentrification research are "black-box" in nature in that underlying model mechanisms are not readily apparent to the user and the relationships between inputs and outputs are unclear (Hamori et al., 2018). Within this context, the current paper contributes an overview of successful ML applications for gentrification modeling with an eye towards elucidating the strengths, limitations, and interpretability of these AI algorithms. By demonstrating how ML models can function within our existing theoretical frameworks, we can work towards demystifying its seemingly "magic" capacity to produce predictions from the data, promote transparency in our data and methods, and encourage a widespread usage of these tools in fields such as human geography and planning, environmental science, and other non-data science disciplines.

As a secondary point of emphasis, the current paper investigates the significance of the advent of "deep mapping," or deep learning models trained to categorize aspects of the built environment, for gentrification modeling (Ilic et al., 2019). Built characteristics have largely eluded researchers prior; Relevant census variables tend to be inconsistently available and thus unworkable for the purposes of creating a reproducible model. However, deep mapping methods provide a template for automating the ways we acquire data of built forms. Such a technique opens the possibility to implement easily reproducible methods of data collection that can be flexible for a wide range of urban environments. Lastly, the current paper proposes a visualization of how these new data collection and modeling techniques can cohere together into a single, comprehensive methodology for forecasting gentrification.

## 2. An actionable definition of gentrification

One of the greatest methodological challenges in mapping gentrification is the misidentification, or the chronic under- or over-estimation, of gentrification occurrence in a locale (Finio, 2021). By establishing a set of comprehensive metrics, researchers can better evaluate the quality of their models and pinpoint areas where a loss of information could lead to poor model performance. From this basis of thought, the paper finds it useful to break down the core characteristics that differentiate gentrification from other forms of neighborhood change before entering into an overview of the methods.

As follows, the current research defines gentrification in a neighborhood when landscape change occurs in league with other urban phenomena that are symptomatic of the presence of gentrification. From Davidson and Lees (2005), these gentrification-related indicators can be broadly grouped under three processes of neighborhood change:

1. a major investment of capital for stakeholders in a localized area, often indicated through the relative scale of development, renovation, and/or beautification efforts underway
2. landscape changes where the amenities, built features, and ease of access to these newly developed or rehabilitated locations are designed for the wealthier target demographic and their existing or anticipated demands
3. indicators of an influx of households from a high(er) socioeconomic class

There exists general consensus in the research that these criteria

serve as reliable indicators for the presence of gentrification in a neighborhood. Ideally, these three criteria are all well-represented in a comprehensive model.

Contemporary research recognizes the existence of gentrification variants that provide more dimension in the way we analyze the differences in how gentrification can manifest across context and place. For example, “new-build” gentrification involves development on brown-field sites and underperforming lots, commonly found in cases of state-led gentrification in the Global South (Lees, 2019). “Green gentrification” is primarily triggered by sustainability initiatives such as the introduction of parks or pollution remediation (Gould & Lewis, 2012; Pearsall, 2018). The redevelopment of these green amenities is often responsible for the displacement of low-income people of color and results in inequitable impacts by further distancing marginalized communities from ecosystem benefits and green space (Anguelovski, 2016; Goossens et al., 2020). With respect to varying forms of socioeconomic change, “marginal gentrification” describes the phenomena where young artists and/or students attract wealthier gentrifier groups by imbuing a trendy aesthetic to the neighborhood with their presence, inadvertently upping the marketability of the neighborhood (Owens, 2012, p. 347). For “super-gentrification,” the gentrifier group comes from the wealthiest upper echelon, a transnational elite class, residing in newly developed, luxury high-rises. This extreme form of gentrification can even result in the displacement of middle-class residents who were once gentrifiers themselves and perpetuate a spiraling up of land and property values (Butler & Lees, 2006; Lees, 2008; Rofe, 2003).

By incorporating these variants into a gentrification modeling methodology, researchers can develop more sophisticated, customized models of measurement for each gentrification variant rather than mapping gentrification under an overly broad umbrella of indicator features. For example, green gentrification may necessitate the variable of “proximity to recently (re)developed green amenity” in order to qualify as a green gentrifying area, whereas solely using traditional gentrification metrics that focus on residential structures may lead to green gentrifying tracts being overlooked in the analysis.

The following section provides an overview of common, non-ML methods for mapping gentrification occurrence (section 3) before moving into a review of ML-specific techniques (section 4).

### 3. A review of the methods for mapping gentrification

In order to map out gentrification for a study area, quantitative studies typically apply a “threshold strategy” for identifying gentrification occurrence by comparing proportional differences in population and housing characteristics before and after a given time point (Barton, 2016). Although contemporary literature prioritizes the recognition of variegated gentrification forms (Atkinson & Bridge, 2004), many existing studies tend towards recognizing a more classic typology of a gentrifying tract – i.e., a majority low-income, blue-collar inner-city neighborhood changing to a majority high-income, white-collar one. Under this diagnosis, a tract is vulnerable to gentrification given that the tract’s median income is some percentile below the municipal or national average, and a tract is gentrified given a certain percentage point increase in income- and employment-related features (Atkinson, 2000; Ellen & Ding, 2016; McKinnish et al., 2010).

For example, Atkinson (2000) quantify gentrification with proxy variables relating to household occupation. Gentrification is indicated by above-average increases in the proportion of white-collar occupations for a given ward relative to the whole of London. Based on descriptive statistics of longitudinal census data, the study finds a correspondence with these occupational shifts and decreases in socioeconomic variables related to blue-collar displacement (e.g., the proportion of working-class households, unskilled labor, renters, non-white ethnicities, the elderly, single parents, and the unemployed).

Freeman (2005) integrates more variables into their criteria for gentrification occurrence, considering only inner-city, low-income,

underinvested neighborhoods in the US that demonstrate an increase in educational attainment and housing price. The threshold is set as a proportional change relative to the median of the metropolitan area, with the variable of underinvestment quantified according to the proportion of housing built.

Threshold criteria provide an easily interpretable and implementable solution to classifying neighborhoods. Nonetheless, research suggests that simpler metrics may fail to capture the complexities within a gentrification model and thus chronically misidentify gentrification status due to that loss of information (Barton, 2016). For example, if a threshold analysis were to use median rent as their variable-of-choice, a non-gentrifying neighborhood with newly constructed affordable housing would be classified in the same way as a gentrifying neighborhood with new-build condominiums (Wyly & Hammel, 1998).

In order to elucidate how alternate methodologies can compensate for this weakness, the current paper first outlines three core limitations to standard threshold analyses. Firstly, commonly used metrics often fail to capture built development and investment, a necessary criterion for gentrification occurrence. A core feature of the gentrification experience is not only that wealthier households move into the neighborhood, but that this demographic movement heralds a restructuring of the built environment in a way that is exclusionary to existing communities. This overreliance on limited sociodemographic measures is not from negligence, but rather due to a lack of access to quality data on built characteristics, thus preventing an appropriate consideration of landscape change and a propensity to under- or over-estimate the scale of gentrification in each study area.

Threshold methods are also limited in the quantity of metrics they can incorporate. For a given neighborhood, socioeconomic change may be better captured across a wider range of metrics beyond occupational and income characteristics. As the threshold method is often implemented in tandem with a kind of “checklist” approach to identifying gentrifying places, this technique limits the capacity to integrate a wider range of indicators and lacks the ability to weight features according to their ability to distinguish between difference socioeconomic groups.

Lastly, gentrification can take on various forms beyond the standard “blue-collar to white-collar” typology. Threshold approaches can overlook more evolved stages of the gentrification lifecycle, such as with super-gentrification where the gentrifier group consists of ultra wealthy individuals who are capable of outing high-income households who were once gentrifiers themselves. Alternatively, threshold methods can overlook more nascent forms of gentrification like marginal gentrification where artists and students inadvertently up the marketability of their neighborhood, even when they lack the financial capital and income status to be quantified as gentrifiers by commonly used metrics.

In order to circumvent these limitations, certain mixed-methods approaches integrate qualitative findings to construct or validate gentrification occurrence. Often, journalistic approaches can capture on-ground truths through interviews with stakeholders, media analysis, and field observations of built changes (i.e., direct auditing).

To date, Hammel and Wyly (1996) and Wyly and Hammel (1998, 1999) provide the most extensive dataset of field observations for Boston, Chicago, Detroit, Milwaukee, Minneapolis-St. Paul, Philadelphia, Seattle and Washington D.C. The field survey observations were centered on the development or rehabilitation of residential buildings, such as the reconstruction of latticework and window frames or the installation of security systems. Along with journalistic sources (e.g., city planning documents, local press reports), these findings were drawn on to split the study area into “core” gentrified, “fringe” gentrified, and non-gentrified tracts. According to the authors, core gentrified tracts are defined as tracts with at least one rehabilitated structure on each block and with at least a third of all structures upgraded in a tract, whereas fringe gentrified tracts have at least one rehabilitated structure on the majority of blocks within a tract with at least one block having a third of all structures upgraded.

Direct auditing methods are, unfortunately, difficult to replicate due

to their time-consuming and labor-intensive nature. As a solution to this limitation, [Hwang and Sampson \(2014\)](#) apply Google Street View (GSV) images in order to effectively audit changes to the built environment using the panoramic, rotation, and zoom features of the GSV platform for different years. By eliminating the need for in-person observation, the researchers could afford to incorporate a wider array of survey metrics like marking evidence of beautification efforts (e.g., painting over graffiti, cleaning up vacant lots) and disorder (e.g., derelict buildings).

The implementation of GSV proves a promising solution to integrating landscape change into gentrification forecasts in addition to providing the validity and inter-reliability results needed to be a productive measure of neighborhood change ([Clarke et al., 2010](#); [Odgers et al., 2012](#); [Rundle et al., 2011](#)). Notably, manually parsing through GSV images is still an exhaustive process. [Hwang and Sampson \(2014\)](#), for example, worked with a sample of 2709 block faces. Recent advancements in ML and computer vision provide a solution to automating the GSV image classification process, and will be discussed in greater detail in [section 4](#).

These advancements allow for a more comprehensive and representative dataset of gentrification-indicating features, but there is still a need for a better statistical method for integrating a wider array of features and for weighting variables by their relative influence. [Wly and Hammel \(1999\)](#) apply a multivariate discriminant analysis to identify which set of independent variables are of best fit to the model, but this type of analysis requires an a priori knowledge of gentrification status and retroactively fits features to a known system state. Additionally, multivariate analyses still perform poorly in the presence of multicollinearities which run rampant when dealing with socioeconomic variables. With the rapid evolution of the gentrification process and its variability from place to place, researchers need a method that can provide a reliable map of gentrifying neighborhoods from the data at hand.

Recent research has found an easily implementable solution to this problem in the form of Principal Components Analysis (PCA) for the identification of gentrifying tracts. PCA is a popular dimensionality reduction technique that can perform optimally with a large number of interrelated variables ([Wold et al., 1987](#)). PCA functions by producing a set of principal component vectors that consist of a linear combination of the original independent variables in a way that best describes the variance of the dataset. In this way, the top principal component would ascribe a new value to each data point or tract that better describes the meaningful differences in tract characteristics.

[Bereitschaft \(2020\)](#) apply PCA on 110 urban cores of U.S. cities in order to quantify socioeconomic and demographic change that could be indicative of gentrification occurrence. The 16 variables capturing metrics like race, occupational status, education, income, age, vehicle-ownership, household size, population, and homeownership could be condensed down to four principal components that describe 67% of the variance in the dataset. PCA, on its own, is not amenable to interpretation as it can be challenging to parse through the individual effects of the original variables. However, PCA provides a useful way to tackle long and wide datasets and distinguish between summarized magnitudes of neighborhood change.

Whereas threshold and PCA methods find utility in identifying past or current gentrification status, relatively less research has been done in constructing techniques for tracking gentrification progression for future years. This subset of the literature has important implications for disseminating academic findings in a way that has immediate utility on-ground, such as providing early warning systems for vulnerable communities or assisting municipal bodies in producing targeted anti-gentrification policies.

Methodology-wise, the majority of studies on gentrification prediction rely on constructing composite indices based on variables related to gentrification susceptibility ([Bates, 2013](#); [Chapple, 2009](#); [Spinney et al., 2011](#); [Turner & Snow, 2001](#)). [Chapple \(2009\)](#) provide a measure of

gentrification susceptibility to function as an early warning toolkit for the San Francisco Bay Area municipal government. Before constructing a predictive model, the research first identifies current gentrification status as low-income tracts that had undergone above-average housing appreciation and educational attainment over the decade. The study then implements a multivariate regression to model gentrification-related variables as a function of attributes relating to demographic traits, income, transportation, housing, and locational factors. If one of the attributes is found to correlate highly with features of gentrification occurrence, then that variable becomes an indicator of future gentrification risk. For the forecasting process, each risk variable contributes an equally weighted score of 1 to the susceptibility index.

Other composite indices generally follow a similar methodology as [Chapple \(2009\)](#), just with different variable choices. For example, [Turner and Snow \(2001\)](#) construct a composite index for Washington DC with the same approach for scoring variables, this time drawing on data that described housing prices, public transit access, and presence of new coffee shops and art galleries. Given the spatiotemporal variability of gentrification occurrence, data choices should vary from place to place. For example, Los Angeles would consider access to public transit a high-risk factor whereas Portland places more emphasis on evidence of private investment in the housing sector ([Preis et al., 2020](#)). However, in the case of the early warning systems mentioned, data choices are difficult to validate for their efficacy relative to other metrics as they are often based on the personal judgements of the authors behind the study.

Of the studies that validated the accuracy of their composite metrics, many find a great discrepancy between predicted and actual gentrification status and a propensity for false positives ([Chapple & Zuk, 2016](#)). There is a persisting need in the literature for more reliable forecasting models that can perform optimally despite multicollinearities between input features and assign weights to indicators according to their contribution to the prediction.

Progress in the literature has brought us the ability to incorporate built measures into the gentrification model and the capacity to condense a large number of input features into smaller, more workable components. We are still left with limitations in our ability to identify gentrification variants i.e., modeling gentrification beyond the more rigid constructs of blue-collar to white-collar, low-income to high-income gentrification evolution. Secondly, we still need more sophisticated methodologies for predicting gentrification susceptibility ([Greene & Pettit, 2016](#)). Moreover, there is a need for less time-consuming methods for acquiring data on the built environment. The current paper will proceed with an explanation of how recent applications in ML provide solutions to these three core gaps in the literature.

#### 4. Machine learning for gentrification

Advancements in ML provide a host of solutions to persisting gaps in the literature, such as parsing out gentrification variants from the dataset ([Owens, 2012](#); [Wei & Knox, 2013](#); [Yee & Dennett, 2022](#)), producing predictive models to forecast gentrification for upcoming decades ([Alejandro & Palafox, 2019](#); [Kiely & Bastian, 2020](#); [Knorr, 2019](#); [Palafox & Ortiz-Monasterio, 2020](#); [Reades et al., 2019](#); [Thackway et al., 2023a](#); [Yee & Dennett, 2022](#)), and gathering built metrics with computer vision techniques ([Ilic et al., 2019](#); [Thackway et al., 2023b](#)).

Moreover, ML methods utilized in gentrification modeling are often amenable to stacking, where the predictions from one ML method can be coherently fed into another, allowing researchers to acknowledge multiple research questions (e.g., identifying current gentrification status, parsing out variants, gathering built metrics, building a predictive model) within a single study. The progress of ML-based gentrification research is often consistent in its tendency to expand off of the methodological structure of prior work.

The current section of this paper will describe these ML modeling techniques in a way that places emphasis on research studies that compound their findings onto the methodologies of their predecessors,

beginning with the pioneering application of a predictive gentrification model (Reades et al., 2019), to the stacking of ML models for identifying gentrification variants (Yee & Dennett, 2022), and down to the final iteration of ML studies that construct more detailed metrics of landscape change for the prediction model (Ilic et al., 2019; Thackway et al., 2023b).

For a brief summary of ML, learning-based algorithms are not so different from their task-based counterparts (e.g., non-ML regression analyses) in that they often function off of similar mathematical functions and in their ability to draw on samples to make inferences about the true system state. ML only differs in its capacity to “learn” from and improve on prior iterations of the model. All learning-based models begin with training the model on a subset of the data and testing the adapted model on the remaining data, unfamiliar to the newly trained model, until optimal performance is achieved.

ML approaches can be grouped into families of supervised, unsupervised, and semi-supervised models. Supervised models (e.g., Random Forest) learn from labeled data, where the “label” is the descriptor value for a data sample (e.g., the label for a tract could be the binary yes/no indicator of whether or not that tract gentrified).

For a simplified diagram of what the supervised ML learning process looks like in practice, see Fig. 1 where the sequential steps for data preparation, model training, model evaluation, model testing, and model explainability are presented in a top-down order. Although not a strict requirement, it is suggested that the output of low-explainability models (e.g., tree-ensembles, support vector machine) is run through an Explainable Artificial Intelligence (XAI) tool (e.g., Shapley, LIME). The recommendation holds for unsupervised and semi-supervised models.

Alternatively, unsupervised models (e.g., K-means) learn from the intrinsic patterns of unlabeled data. Where supervised models train on user-presented examples (i.e., labeled data samples), unsupervised models use the structure of the data itself to come up with a prediction. Unsupervised approaches are commonly used for clustering, dimensionality reduction, and density estimation. Semi-supervised models can

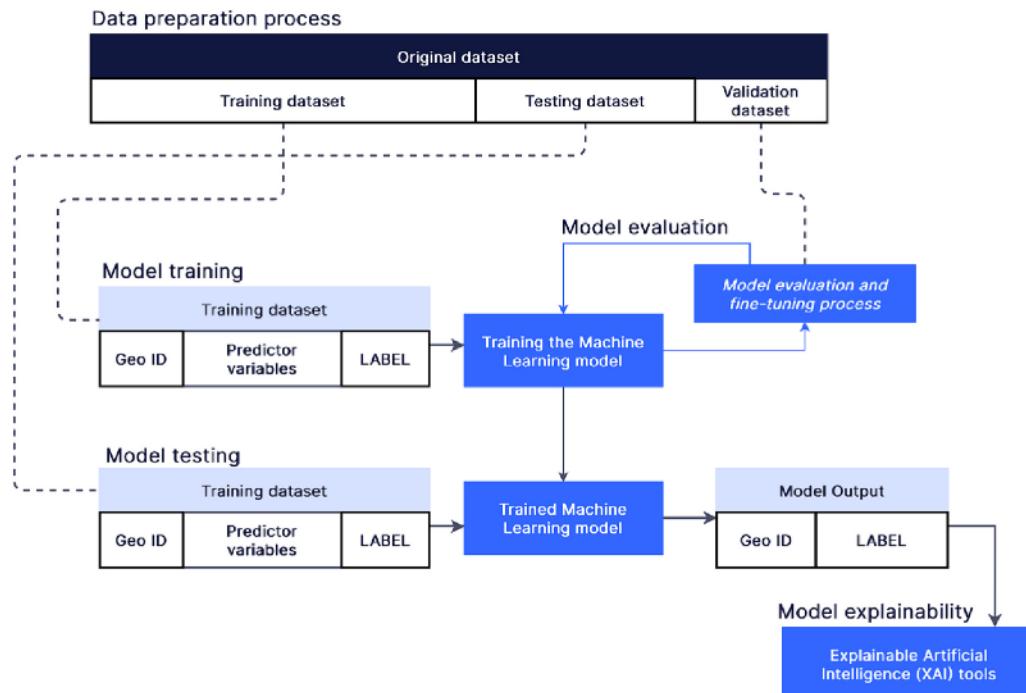
intake unlabeled and labeled data in tandem, and are often applied when labeling the entire training dataset becomes too labor intensive for the user. In the following subsection, we describe how ML models can be used to predict future gentrification occurrence. Many of the studies mentioned in the review published open-source datasets and/or methods via GitHub. Links are provided in the Supplementary Information section of the current paper.

#### 4.1. Data sources

The spatial and temporal scales of gentrification research tend to be restricted by the extent of data availability. Studies largely measure gentrification over 5- to 10-year periods corresponding to regional census data time scales. For US contexts, the census tract is a commonly used spatial unit of analysis containing around 1200 to 8000 residents per unit.

ML-based gentrification models are similar to traditional studies in the variables they incorporate. Gentrification’s influential factors, as represented in the data, can be broadly split into two camps: residential household and housing characteristics like the demographic and socio-economic descriptors of households and basic housing stock features and, secondly, the locational attributes that influence the accessibility, appeal, and marketability of urban spaces, like Point of Interest (POI) amenities, proximity to the downtown center, historic district, or transit, and qualities of the public transit system itself (Brown-Saracino, 2017b; Finio, 2022).

Unlike traditional regression analyses, many ML algorithms are capable of intaking a high number of correlated variables without running the risk of overfitting. Tree-based models like Random Forest (RF) and Gradient Boosted Machine (GBM) can methodically select features that contribute best to the final prediction and, in turn, place lower predictive weight on the more insignificant variables. This optimization process is of great utility to the research as neighborhood variables can change in predictive power over location, context, and time, even within a single study area. For example, the impact of transit



**Fig. 1.** A supervised Machine Learning (ML) model. A flowchart describing the processes for data preparation, model training and testing, model validation, and model explainability. Specific to geospatial datasets, “geo ID” refers to the geographic unit of measurement (e.g., tract). Without the “label” boxes, the figure would represent an unsupervised learning algorithm.

Source: Authors.

station access (Wardrip, 2011) and parks (Anguelovski, 2016; Rigolon & Németh, 2020) display variable influence on gentrification outcomes depending on location and feature characteristics.

The algorithmic flexibility of ML allows the model to integrate a wide range of variable domains. For example, for a Random Forest (RF) model for predicting future gentrification in London, Reades et al. (2018) incorporate 166 census variables that describe the housing market (median house price and rent), demographic descriptors (e.g., age, place of origin, gender, family structure, health), educational and occupational traits, household transportation behaviors and accessibility, income, and basic structural attributes of commercial and residential housing.

Thackway et al. (2023a) similarly cover an extensive range of features from regional census datasets like age, country of birth, dwelling type, education level, employment status, family composition, income, industry, marital status, occupation, and religion. The authors supplement their analyses with property data of housing prices and supply, commercial property sales, and additionally Point of Interest (POI) data like Airbnb count and presence of food and leisure businesses. See the Appendix for a comprehensive list of data sources for the respective publications (Reades et al., 2018; Thackway et al., 2023a).

For more detailed POI categories, Zeng et al. (2022) provide a ML-based methodology for acquiring detailed metrics on consumer amenities like food services, recreational outlets, art institutions and other regional attractions and further maps the POI data in relation to gentrification patterns with a supervised gradient boosting model.

For the research at large, built metrics are more scarcely available in census data outside of simpler descriptors like building zoning (e.g., residential, commercial), housing structure (e.g., detached vs. semi-detached designations), and building age. In Section 4.4., we provide detail on how ML models are implemented to produce data on whether a residential structure undergoes rehabilitation or is replaced by new-build development (Ilic et al., 2019; Thackway et al., 2023a).

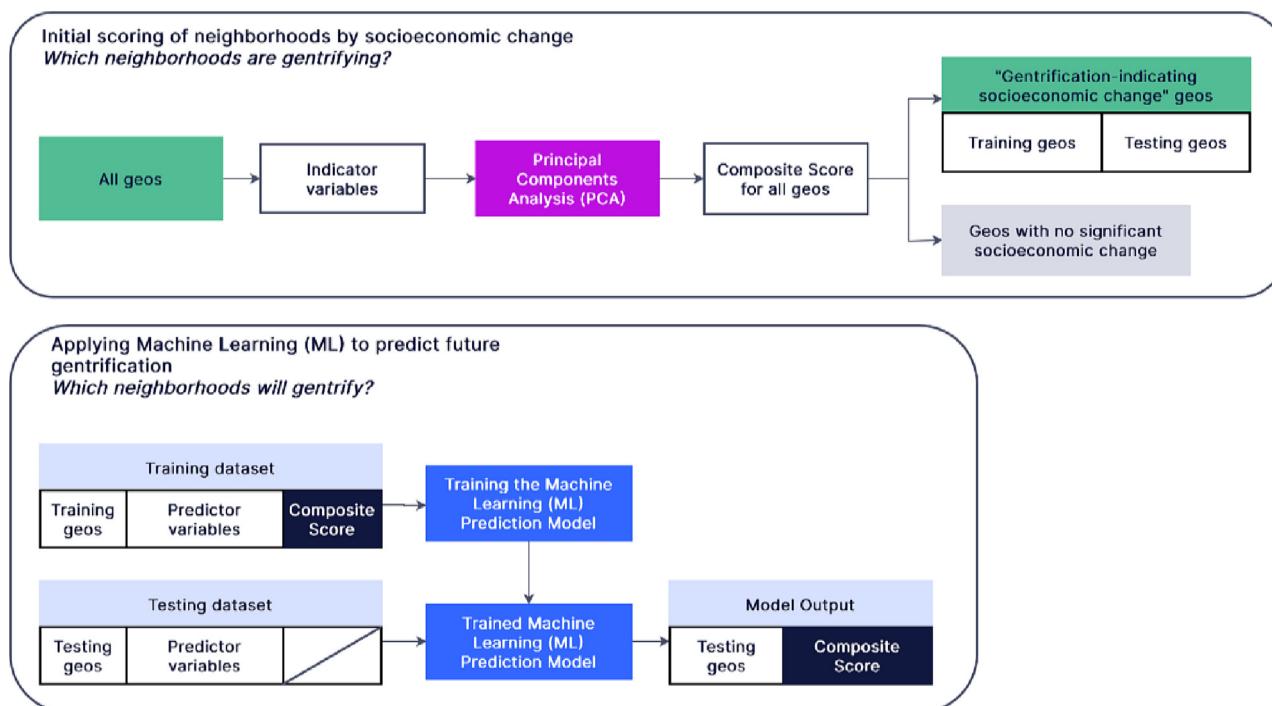
#### 4.2. A predictive machine learning model

The pioneering application of ML for gentrification analysis began with the work of Reades et al. (2019) and the implementation of the ML algorithm Random Forest (RF) in order to predict gentrification trajectories for Lower-Layer Super Output Areas (LSOAs) in London. LSOAs are approximately similar to US census tracts, the default geographic scale for gentrification-based studies.

Random Forest is a supervised machine learning model that can be adapted for classification problems for predicting categorical outcomes (e.g., whether or not a tract gentrified) or for regression problems for estimating a continuous output (e.g., the degree to which a tract underwent socioeconomic change). Due to its ability to perform well with noise and collinearities, RF is an attractive algorithm for high-dimensional, real-world problems (Breiman, 2001).

As Random Forest is a supervised learning algorithm, it is necessary to first construct a label for the training dataset. In application to the gentrification forecast, Reades et al. (2019) precede the Random Forest regression model with a training dataset where each LSOA is assigned a continuous measure indicating the degree of gentrification-related change. The label is quantified using Principal Components Analysis (PCA) to construct a single composite score from four variables of neighborhood change (e.g., median household income, median property value, proportion of higher-earning, professional employment groups, and education attainment). For more detail on PCA construction, see section 3. The results of the PCA analysis are promising and indicate a reliable composite metric, with the first principal component alone explaining 78.8% of the variance in the data. Importantly, the authors note the need for built variables for a more comprehensive composite metric.

For a visualization of the training and testing process, refer to Fig. 2 where all LSOAs (or geos) are run through a PCA in order to derive a composite score. After splitting samples into a training and testing dataset, a ML prediction model (e.g., Random Forest regression) learns how to attribute a prediction to each LSOA in the study area. The labels



**Fig. 2.** A basic gentrification prediction model. The base stack of a gentrification prediction modeling methodology where a Principal Components Analysis (PCA) and a supervised Machine Learning (ML) model like Random Forest Regression produce a gentrification forecast. Geos refer to geographical sampling unit and the Composite Index Score quantifies the degree of socioeconomic change, a proxy for gentrification.

Source: Authors, based on the methodology of Reades et al. (2019).

for the ML model are differentiated by a dark blue color.

For the testing phase, Random Forest (RF) models outperformed non-ML linear regression methods in prediction reliability. Even the initial, untuned RF model (i.e., without the optimization of model hyperparameters) could provide more accurate predictions of gentrification-indicating socioeconomic change for London LSOAs compared to the non-learning-based approach. From 2001 census data, the tuned RF model could predict 2011 PCA composite scores with high accuracy (Pearson's  $r$  of 0.99).

[Thackway et al. \(2023a\)](#) compare the performance of Random Forest (RF) against two other supervised machine learning algorithms: Gradient Boosted Machine (GBM) and an improved implementation of GBM, Extreme Gradient Boosting (XGBoost). All models are implemented to predict future gentrification for Sydney, Australia. Rather than a PCA score, the authors utilize a pre-existing municipal index on household demographics and residential and commercial housing supply. All machine learning methods outperformed task-based linear regression, but the boosting algorithms (i.e., GBM and XGBoost) demonstrated marginally higher performance to their RF counterpart in predicting gentrification growth for Sydney, Australia.

For other examples of ML modeling for gentrification prediction, see work by [Palafox and Ortiz-Monasterio \(2020\)](#) where Neural Networks are applied to predict gentrification trajectories for Mexico City. [Kiely and Bastian \(2020\)](#) provide a more economics-centric measurement of gentrification based on real estate sale price and sale probability to compare the performance of a generalized linear model (GLM), Random Forest, Gradient Boosting Machine (GBM), and artificial neural network (ANN).

#### 4.3. Gentrification variants and implementing controls

For the second “stack” of the gentrification forecasting system, we discuss the integration of clustering algorithms for the identification of gentrification variants and other forms of neighborhood change (e.g., incumbent upgrading).

A popular clustering algorithm, K-means ([Hartigan & Wong, 1979](#)) is an unsupervised learning method that partitions the dataset into  $k$  clusters where within-cluster variance is minimized. The goal of K-means is to produce clusters that are compact (i.e., samples within a cluster are similar) and distinct from other clusters. As K-means is unsupervised, the algorithm does not work with nor produce labels for each cluster. Thus, researchers are tasked to profile the clusters based on their theoretical knowledge of what each cluster may pertain to. For example, a cluster centroid characterized by features of socioeconomic ascent, but lacking features of landscape change would be inferred to pertain to an “incumbent upgrading” data cluster, in contrast to a “gentrification” cluster sharing features of both socioeconomic ascent and landscape change.

Early research has experimented with the applicability of K-means for parsing out different typologies of socioeconomic ascent. For example, [Wei and Knox \(2013\)](#) apply K-means cluster analysis and discriminant analysis to distinguish between seven clusters of neighborhood types based on combinations of race, immigrant status, and socioeconomic descriptor characteristics. [Owens \(2012\)](#) implement a PCA in order to summarize population and housing features into five to six composite indices, or linear combinations of the original variables, before running K-means to define a typology of metropolitan tracts according to factors like race, economic status, and gentrification occurrence.

[Yee and Dennett \(2022\)](#) follow directly off of the methodology of [Reades et al. \(2019\)](#) to predict future gentrification status for London LSOAs, this time stacking the model with a K-means analysis of gentrification variants. Whereas [Reades et al. \(2019\)](#) treat the PCA score for socioeconomic change as a proxy to quantify gentrification and a label for the prediction model, [Yee and Dennett \(2022\)](#) treat the PCA score as solely a measure of socioeconomic change. From the neighborhoods of

the PCA, an initial K-means analysis is run to further isolate the group of samples that demonstrate both socioeconomic shift and evidence of (re) development i.e., the gentrifying cluster. Data on (re)development is drawn from municipal-level data on new-build construction and rehabilitation of existing residential buildings.

After filtering out LSOAs that either represented socioeconomic change without built change or, alternatively, built change without socioeconomic change (i.e., incumbent upgrading), the study performs a secondary K-means to identify gentrification variants like marginal gentrification and super-gentrification.

After K-means, [Yee and Dennett \(2022\)](#) run a Random Forest (RF) Classification model where LSOAs are categorically labeled according to their gentrification status and variant type (classic gentrification, marginal gentrification, or super-gentrification). The research provides a novel methodological approach to integrating clustering techniques into the gentrification forecasting toolbox.

See [Fig. 3](#) for a visualization of how a clustering algorithm like K-means can work with a PCA and prediction model (e.g., Random Forest classification) to identify different types of neighborhood change. For simplicity, [Fig. 3](#) condenses the two disparate K-means analyses used in [Yee and Dennett \(2022\)](#) into a single process, highlighted in yellow.

K-means carries some notable limitations that are worth consideration before replicating this method with other datasets. For one, K-means works under the assumption that all clusters are approximately the same size and variance. With clusters of uneven density, the model can easily group the data incorrectly even when the clusters are apparent to the naked eye ([Holden et al., 2011](#)). With gentrification-related models where the quantity of, say, green gentrifying tracts could be much smaller than the number of non-gentrifying tracts, this limitation can lead to deteriorating model performance. With a large number of independent variables, researchers can find it difficult to visualize model performance on multiple axes and validate the accuracy of cluster assignments.

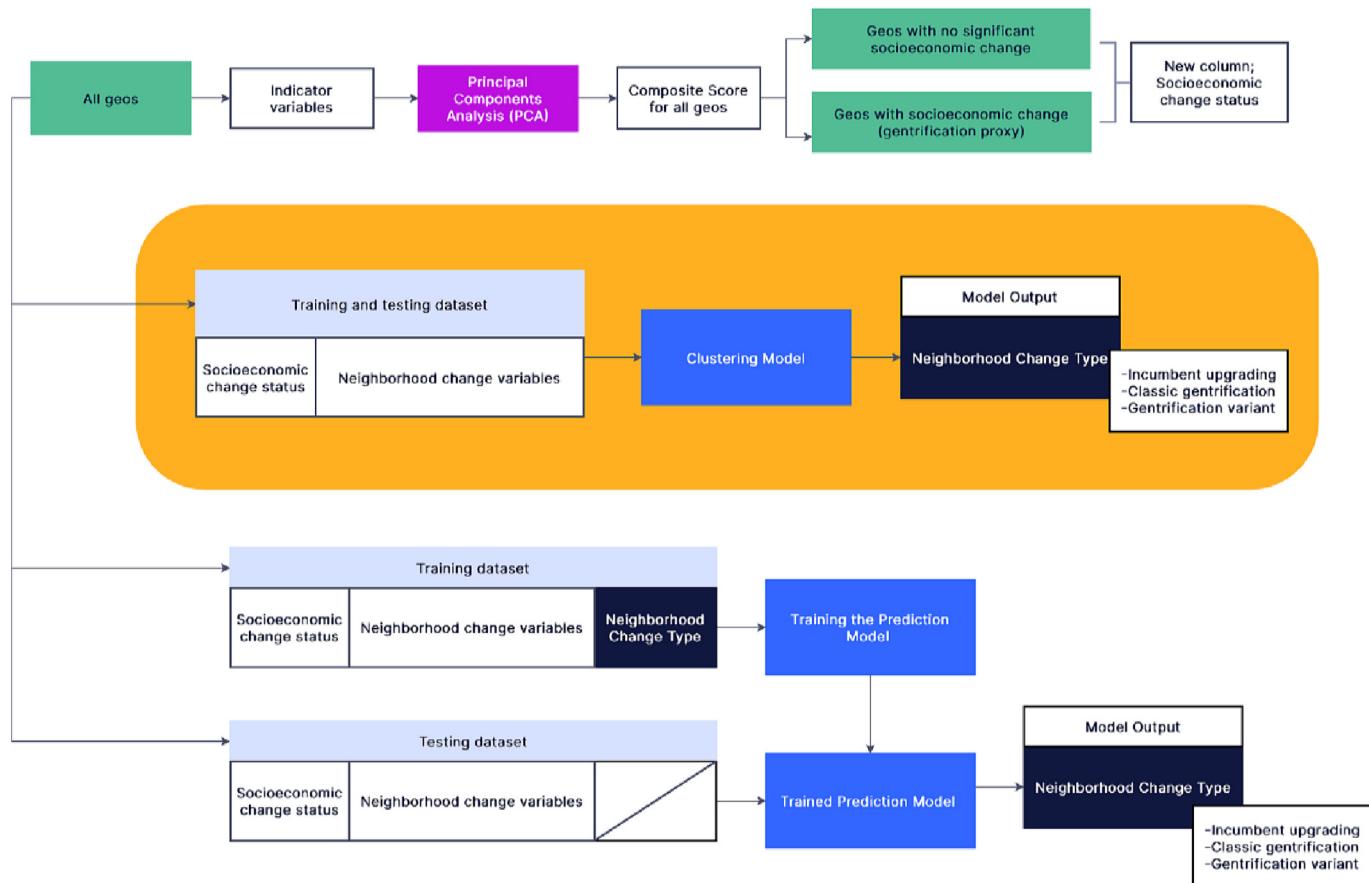
Alternate clustering algorithms may be able to circumvent these shortcomings, such as Density-Based Spatial Clustering and Application with Noise or DBSCAN ([Ester et al., 1996](#)), an algorithm that clusters samples based on density and can function optimally despite unequal cluster sizes and the presence of outliers ([Finch, 2019](#)). Ultimately, the performance of clustering algorithms can vary depending on the research application and dataset type so it can be beneficial to compare the performance of multiple methods.

The clustering-based methodology proposed by [Yee and Dennett \(2022\)](#) provides a way to control for other forms of neighborhood change and incorporate landscape change into a gentrification metric. However, the housing data applied by the study draws on municipal-level data which can be sparse and inconsistently available for different regions. The next subsection reviews how a third ML stack addresses this limitation by automating the process of data acquisition on built forms in a way that is reproducible for a variety of geographic contexts.

#### 4.4. Better built metrics with computer vision and street-level imagery

The third stack of the ML gentrification forecast methodology involves integrating more comprehensive metrics on the built environment. Mapping the built environment in an Artificial Intelligence (AI) context is referred to as “deep mapping” or “machine mapping” and presents a promising alternative to direct auditing. Deep mapping approaches are already phasing out patch-wise labeling and other traditional methods of land-use classification in their capacity to pick out fine boundaries in remote-sensing aerial images ([Kang et al., 2018](#)). Image data opens many new avenues for finally integrating built forms into the gentrification model and placing more emphasis on landscape change as an essential feature of the gentrification process.

Deep mapping incorporates two features of new urban analytics: deep learning – a subset of ML algorithms that run on neural network



**Fig. 3.** Second stack of the gentrification prediction model. The second stack of a gentrification prediction model where the methodology in Fig. 2 is supplemented by a K-means analysis (in yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.) Source: Authors, based on the methodology of [Yee and Dennett \(2022\)](#)

architectures, and computer vision – an AI application that aids in the processing, detection, classification, and interpretation of images. In practice, deep mapping consists of black-box, ML methods that draw meaningful information from images, such as Google Street View (GSV) or satellite imagery, in a way that allows researchers to automate the collection of built variables and incorporate more detailed, geospatial metrics.

Deep mapping techniques can possess different levels of model interpretability, from the more black-box models that imitate perceptual judgements of a scene (e.g., categorizing a street view as “beautiful” or “ugly”) to relatively less opaque models that learn to extract semantically-defined spatial objects in an image (e.g., road, sidewalk, tree).

For perceptual deep mapping, pioneering studies trained ML models to categorize streets according to perceived safety, drawing on crowd-sourced datasets and GSV images to produce large-scale maps for U.S. cities, such as the computer vision algorithms of Place Pulse and Street Score ([Naik et al., 2014](#); [Salesse et al., 2013](#)). The datasets and computational techniques are open-source and have been drawn on for further urban analysis, such as identifying socioeconomic characteristics that correlate with neighborhood infrastructure improvements ([Naik et al., 2017](#)).

Many of the perceptual models apply Convolutional Neural Networks (CNNs) which are versatile in their ability to learn from large sets of labeled and unlabeled images, hence the capacity to mimic human visual judgements of more esoteric qualities, such as identifying a street-level image as “safe-looking” ([Naik et al., 2014](#); [Salesse et al., 2013](#)), “scenic” ([Seresinhe et al., 2017](#)), or “walkable” ([Blečić et al., 2018](#)).

For a simplified explanation, CNN draws on multiple convolutional

layers where each layer consists of a procedure for extracting information from an image input. The initial layers extract low-level, basic features (e.g., points, edges, lines) from the image, with each additional layer extracting more and more complex feature descriptors (e.g., shapes, ridges). CNN is a learning-based algorithm in that it learns how to optimally filter out meaningful information from the image on its own without the need for hard-engineered rules ([Cheng et al., 2018](#)).

In the field of remote sensing scene classification, a nascent but growing body of research has found promise in the application of Siamese Convolutional Neural Networks (SCNN) that harnesses the strengths of two or more identical CNN models in order to compare the similarity between pairwise images ([Bertinetto et al., 2016](#); [Liu et al., 2019](#)).

[Ilic et al. \(2019\)](#) applied the first implementation of a deep mapping model to map gentrification-indicating landscape change. The study applies a Siamese Convolutional Neural Network (SCNN) that is trained to identify changes between two, sequential images. In this case, panoramic Google Street View (GSV) imagery is drawn on to construct a dataset where a single data sample of pairwise images consists of the before- and after- structural improvements for a residential building. The study defines a neighborhood as likely gentrifying when there is a high spatial density of property improvements over an area. The authors acknowledge how the lack of sociodemographic metrics could result in an over-estimation of gentrification status, as incumbent upgrading is not controlled for.

The authors adopt the protocol of [Hammel and Wyly \(1996; p. 250\)](#), where gentrification-like changes cover the replacement of an older, disinvested building with a newly constructed house, significant and structurally sound renovation of the building (e.g., reconstructed

windows and frames, steps, porches), and evidence of landscaping and beautification (e.g., porch furniture). See Fig. 4 for an example of a gentrification-indicating structural change, where a residential building is demolished in lieu of a newly-built three-story home with modernistic architecture.

Notably, the deep learning model itself does not explicitly learn how to identify semantic objects like windows, porches, or porch furniture, but rather mimics human judgements of before- and after- changes at the pixel unit. As the model does not extract any human-understandable objects, categorizations can be vulnerable to inaccuracies for less pristine image data. False positive detections occurred when there was a slight offset in the image snapshot, causing a portion of an adjacent property to appear in the image. The presence of vehicles also threw off the model.

The granular scale and black-box nature of the SCNN algorithm gives rise to dilemmas in ascertaining the breadth of model's ability to capture different built forms – is the model really targeting the right pixels (e.g., the pixels that comprise of a new window) when categorizing an image? Is there a cut-off point in terms of the *degree* of structural change that the model can identify? The low level of AI interpretability leaves much to be desired, but the high performance of the model with new, unseen testing images (95.6% accuracy) and comparisons with cross-validated Development Approval (DA) datasets suggest that the model is capable of honing in on the correct target regions of the training data for a majority of cases.

For future research, more detailed metrics on the scale of redevelopment in each image pair during manual auditing could provide insight into the mechanisms of the model and the breadth of its capacity to identify structural changes. Additionally, the research can benefit from outlining explicit protocol on the attributes in an image that define a significant, gentrification-indicative change. Quantifying and justifying the metrics that underlie the image labeling process is critical for promoting the replicability of these deep mapping models; Different researchers will inevitably label pairwise images differently and visible structural changes can vary according to the architectural specificities of a given region. Qualitative, local knowledge and community input can fill in the gaps for understanding these place-based gentrification indicators and help researchers distinguish between incumbent upgrades and capital investment from external developer groups.

Although there is no precedent, the literature may also benefit from constructing a measure of inter-class variability in (re)development intensity (i.e., the degree of renovation to a single structure) rather than solely measuring the intra-class density where all forms of property improvement are weighted equally.

From the methodology of Ilic et al. (2019); Thackway et al. (2023b)

apply a SCNN deep learning model on GSV images of residential landscape change for neighborhoods in Sydney, Australia. The research also draws on prior results of a census-based ML gentrification forecast (Thackway et al., 2023a), making it the first study to integrate predictions of socioeconomic shift with that of GSV-based built metrics. Although the landscape change data is not integrated in tandem with the sociodemographic-centric model, the study provides a solution to drawing on SCNN-derived data to ensure that the gentrifying neighborhoods are qualified according to corresponding sociodemographic shift and residential-specific built changes.

Because the manual labeling of the training data is still a labor-intensive process, the image classification process is semi-supervised in that a portion of the samples are labeled by the authors and the remaining samples are labeled by an unsupervised classification algorithm.

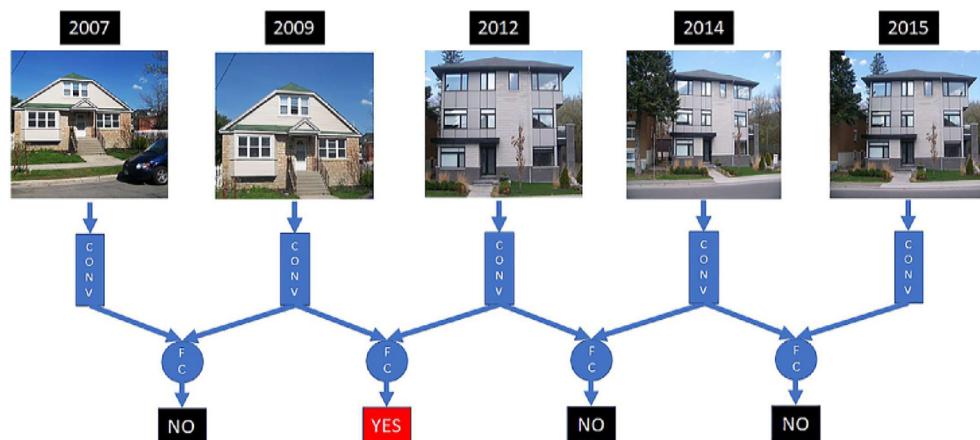
As the deep learning model can be easily thrown off by images with tree or vehicle obstructions, the paper applies a semantic segmentation algorithm (SegNet) to classify images into basic object types (e.g., building, tree, sky, car) and parse out image data unsuitable for pairwise classification (Badrinarayanan et al., 2017).

In Thackway et al. (2023b), researchers were able to draw on their familiarity with regional urban trends in Sydney, Australia and assess demographic census records to identify immigrant population turnover in Sydney's Bankstown neighborhood as a culprit of localized socio-economic changes without landscape change. These insights provide an important qualitative dimensionality to the analysis and an opportunity to distinguish between model inaccuracy and model scope.

See Fig. 5 for a proposed methodology that demonstrates how the methodologies of Reades et al. (2019), Yee and Dennett (2022), Ilic et al. (2019), and Thackway et al. (2023b) can be combined into a multi-layered gentrification forecasting system. Highlighted in yellow, Fig. 5 illustrates how the deep mapping model (e.g., Siamese Convolutional Neural Networks) can work to produce a binary label (i.e., whether or not a residential structure underwent visible changes).

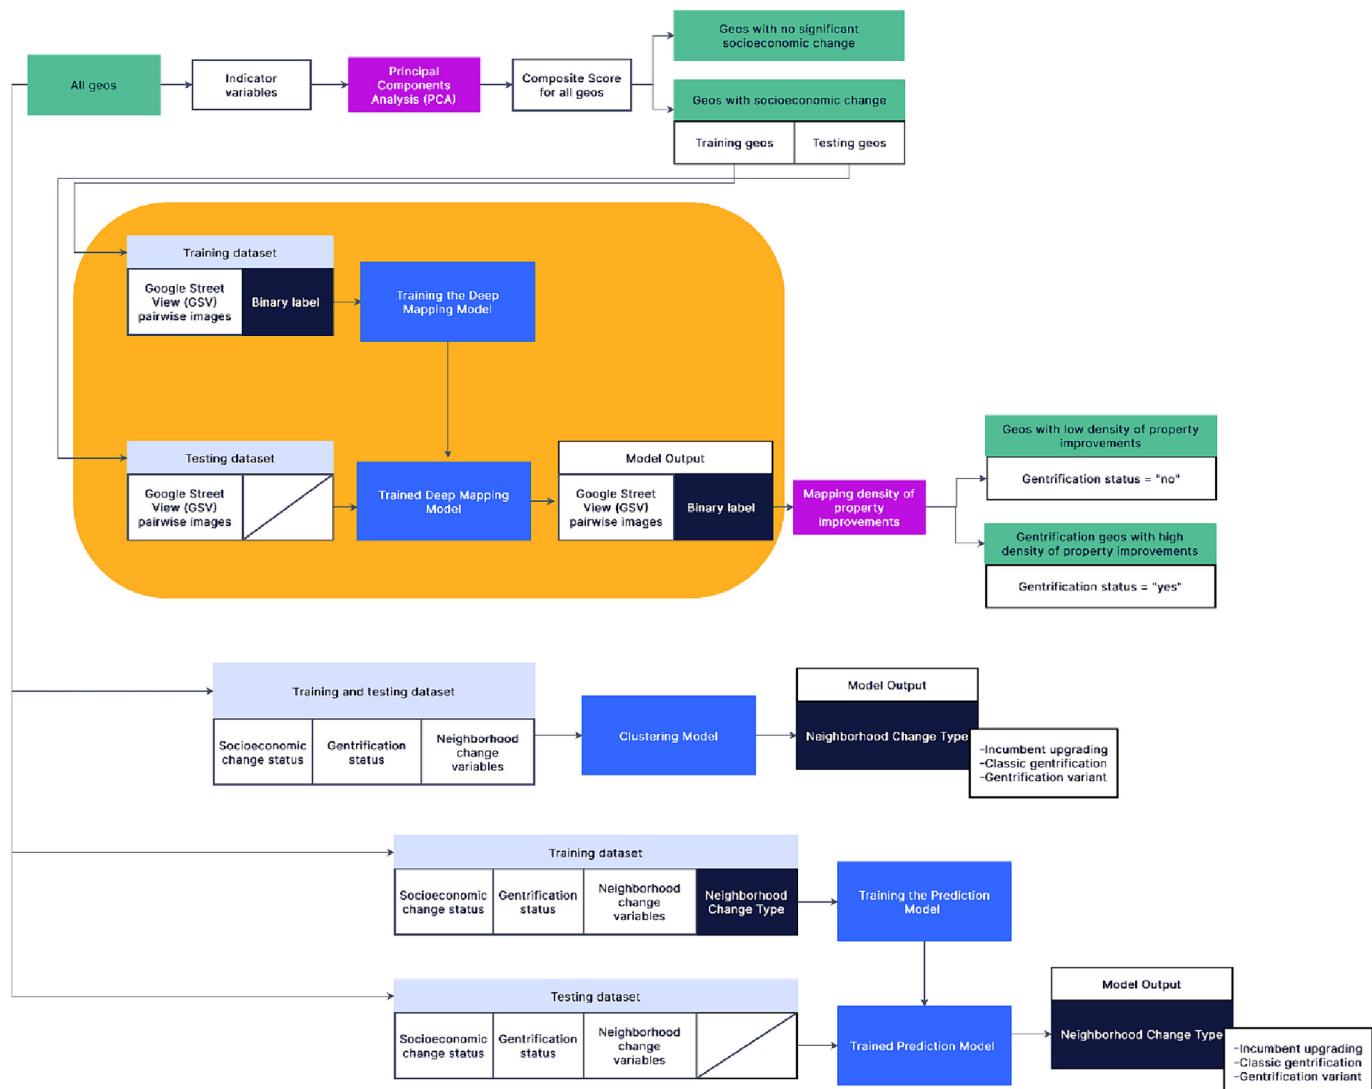
In Fig. 5, the mapped density improvements are used to construct the "Gentrification status" column input for the clustering model and can also be integrated into the "Neighborhood change variables" independently as a "density of rehabilitation" variable to identify incumbent upgrading clusters for both the clustering and the prediction model.

In Fig. 6, the methodological stacks covered in the current paper are simplified into a four-step process for further clarity. The "Data collection" phase describes how input data for subsequent steps is acquired, where a computer vision or deep mapping ML model draws out information on residential development and rehabilitation (see Fig. 5 and Section 4.4). In order to illustrate where various ML models contribute



**Fig. 4.** User pairwise images. Crowdsourced perceptions on landscape change are turned into training data. Users are asked, "Is there a property improvement?" for before- and after-images of a residential unit.

Source: (Ilic et al., 2019)



**Fig. 5.** Third, proposed stack of the gentrification prediction model. The third stack of a proposed gentrification prediction methodology incorporating the contributions of Fig. 2 and Fig. 3 with a more accurate measure of built metrics.

Source: From authors, based on work by Ilic et al. (2019) and Thackway et al. (2023b)

to the gentrification prediction, they are represented by light blue boxes. The mapped output, where a value is assigned to a geographic area (e.g., census tract), is distinguished by a darker blue color.

In step (b), “Identifying gentrification status”, relevant metrics are drawn out of the census-based household and housing data to produce a composite score. Gentrification status is attributed to neighborhoods with both socioeconomic change and evidence of built investment. As PCA is a non-ML analysis, it is represented in purple. For more detail, refer to Fig. 2, the end of Section 3, and Section 4.2.

Step (c) summarizes the process illustrated in Fig. 3 and Section 4.3 for identifying gentrification variants (e.g., green gentrification, marginal gentrification, super gentrification, classic gentrification). Lastly, step (d) produces the final gentrification prediction model (Fig. 2; Section 4.2). Steps (b) and (c) produce the requisite label for the final ML prediction model (d) by producing a map of gentrified neighborhoods (b) and further detail on their corresponding variants (c).

Instead of solely mapping the prevalence of built improvements as initial controls for identifying gentrification, the current review posits that there is utility in using the deep mapping output as a *predictive* feature for future gentrification occurrence:

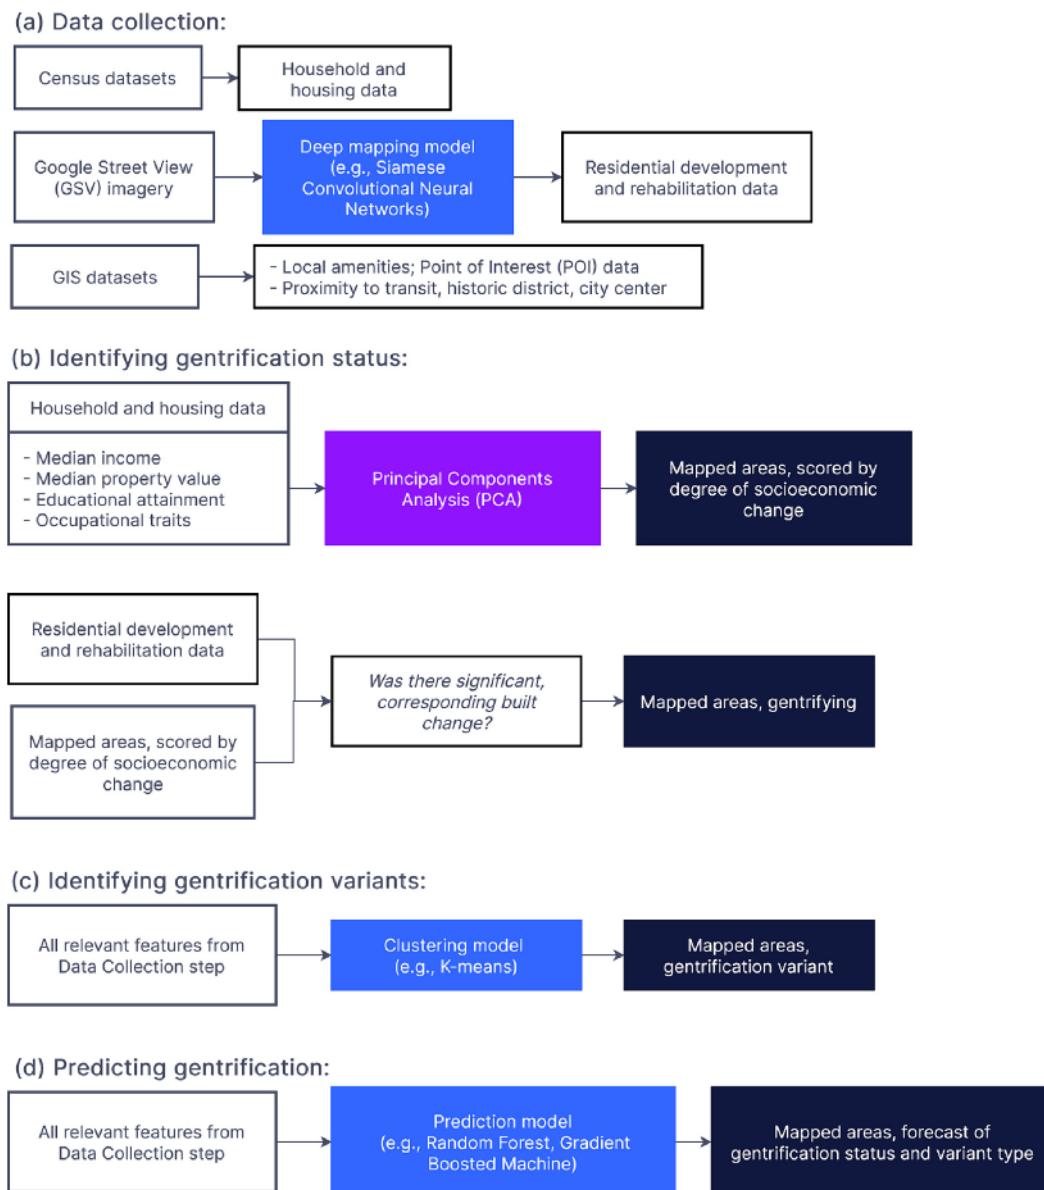
For one, there exists a major gap in the research in understanding the temporal directionality of built changes in the gentrification process. By

adding a time-based level to the redevelopment data (e.g., structural changes that occurred from 2005 to 2010, 2010–2015), we can investigate the role of built change as a predecessor, coinciding, or posthumous event on the gentrification timeline.

Mapping the temporal evolution of redevelopment gives way to new opportunities in the research: If the bulk of residential improvements occur prior to socioeconomic shift, then could we draw on visible signs of improvements as a proxy for gentrification warning systems? Do the densities of structural changes align or misalign with the magnitude of socioeconomic shift?

Although there are observational discussions on the characteristics that describe the architecture of gentrified neighborhoods (Helms, 2003), there is limited discourse on the role of structural attributes in instigating, hastening, or even preventing gentrification effects (Aoki, 1992; Miranda & Lane-McKinley, 2017).

If we can pinpoint more specific architectural attributes through auditing or more detail-oriented computer vision tools, can we assess if there are any built characteristics that make a neighborhood more vulnerable or resilient to gentrification? The research stipulates that the built environment can affect human outcomes and foster a sense of place-based community (Chitov, 2006; Mazumdar et al., 2018), and that communities with social capital can better rebound from and resist



**Fig. 6.** Summarized methodological steps for a gentrification prediction model. Steps (a)-(d) represent how multiple Machine Learning (ML) models, represented in light blue, play various roles in acquiring data, identifying gentrification variants, and predicting future gentrification. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)  
(Source: From authors)

gentrification-induced changes (Balzarini & Shlay, 2018; Bernstein & Isaac, 2023)—can we draw on ML tools to develop more gentrification-resistant neighborhoods?

Although not yet widespread in the research, there are many existing tools in the remote sensing field for producing more detailed accounts of built change and investigating the aforementioned gaps in the literature:

For more objective measures of deep mapping, neural networks have shown promising performance in identifying changes in urban greenery via street-level imagery (Lu et al., 2023). Common measures include the Green View Index (GVI) that scores relative quantities of vegetation based on the distribution of RGB pixel colors in each street-level image (Li et al., 2015), Sky View Factor (SVF) to quantify the proportion of visible sky (Zeng et al., 2018), and measures for tree height and canopy projection size (Wang et al., 2018). Deep mapping for greenery changes can be an especially helpful tool for variants like green gentrification, where gentrification can arise from the installation of green amenities.

Other deep mapping measures include the semantic classification of

buildings trained on street-level image data like GSV and OpenStreetMap to map building frontage quality and functionality (e.g., Kang et al., 2018; Law et al., 2019; Liu et al., 2017). For example, Kang et al. (2018) provide a framework for semantically classifying building instances with a CNN model to map the presence of apartments, churches, garages, industrial buildings, office buildings, retail buildings, and roofs in street-level images. These building classification maps are currently publicly available for further urban systems research at high levels of spatial resolution for the U.S. and Canada. For future research, these methodological advancements could prove an asset to the gentrification forecasting process as a more objective measure of building stock characteristics and for better understanding the role of built forms in the evolution of neighborhood change.

#### 4.5. Explainable artificial intelligence (XAI) for the predictive model

Despite its high-performance accuracy, the ML models described

above (e.g., Random Forest, XGBoost, CNN) all fall under the category of opaque, or “black-box”, algorithms in that the motivations behind the algorithm’s decisions, the underlying model structure, and the interactions between input variables are not well understood by the user (Pasquale, 2015; Rudin, 2019). With the potential to test increasingly more complex and obtuse systems of measurement, it becomes critical to integrate systems for model interpretability. This ensures that stakeholders can reliably extrapolate patterns from the research and that the model itself is grounded upon a logical, theoretical basis rather than functioning off spurious correlations in the data (Yosinski, Clune, Nguyen, Fuchs, & Lipson, 2015).

A machine learning algorithm can, in practice, predict if a household voted Republican based on a satellite image of their vehicle (Gebru et al., 2017), or estimate the Dow Jones average based on the prevalence of the word “calm” in Twitter tweets (Bollen et al., 2011), but finding happenstance patterns in the noise of data clouds (i.e., data mining) is theoretically and productively of zero substance. Such associations cannot be reliably extrapolated for future applications, and predictions, even in their immediate usage, can be easily thrown off track with minuscule alterations to the data (Smith, 2020). For this reason, recent trends in AI literature are moving to prioritize interpretability over accuracy and emphasize the need for Explainable AI (XAI) methods when dealing with human-based research questions (Angelov et al., 2021; Core et al., 2006; Guidotti et al., 2018).

Thackway et al. (2023a) provide the first implementation of an Explainable AI (XAI) analysis to better interpret the results of three models: Random Forest, Gradient Boosted Machine (GBM), and Extreme Gradient Boosting (XGBoost). Random Forest, GBM, and XGBoost all fall under the category of tree ensemble models and therefore have low levels of explainability compared to more transparent ML models like decision trees, k-nearest neighbors, and rule-based learners. In order to provide more insight to model behavior, the authors apply Shapley Additive Explanations (Shapley or SHAP analyses), an XAI tool for explaining feature relevance by ranking the predictive power of influential features (Lundberg & Lee, 2017). In practice, Shapley analyses can provide a system to ascertain whether the ML model is working with the input features in a logical way. For example, if the variables representing the change in median income, rent, and house price were found to rank highly in Shapley analyses, this would instill trust in the ML prediction as the behavior of the model corresponds to the researcher’s theoretical knowledge of gentrification occurrence. On the other hand, if a more innocuous variable, like the proportion of residents that immigrated from Canada, were found to be of high import to the model prediction, we may have reason to reconsider whether the model is taking shortcuts based on noise in the data.

Shapley analyses are versatile for a wide array of ML model types (i.e., are model agnostic) and are well-trusted by the ML community due to Shapley’s basis in game theoretic mathematics. However, Shapley is not always ideal in the case of highly correlated predictors as the Shapley explanation for an individual feature would be muddled by between-factor interactions. For tree ensemble methods in particular, the literature advises supplementing feature-oriented XAI tools like Shapley with model simplification tools (Arrieta et al., 2020).

Examples of model simplification in XAI include local surrogate methods like Local Interpretable Model-Agnostic Explanations (LIME; Dieber & Kirrane, 2020) or G-REX (König et al., 2008). Local surrogate methods function by training a transparent, toy model (e.g., Decision Trees) on a permutation of the original dataset where the labels of the training dataset comprise of the output of the original, black-box model (e.g., Random Forest). Local surrogate models like LIME are also applicable to image data where user-defined pixels are permuted from the original image data to better understand the influence of spatial objects on the image classification model.

Whereas local surrogate techniques are used to understand how the model prediction changes when data samples are perturbed by the user, global surrogates approximate the entire prediction function of the

original model with a surrogate. See Islam et al. (2022) where the global behavior of a Random Forest model is approximated with CART decision trees.

Specific to CNN-type architectures, Class Activation Maps (CAMs) and its generalizations Grad-CAM (Selvaraju et al., 2017) and Grad-CAM++ (Chattopadhyay et al., 2018) provide an XAI method where, similar to Shapley, spatial patterns in an input image are marked according to their relative importance to the CNN model prediction. CNN-specific XAI techniques have not been applied for models interpreting gentrification-related landscape change, but may be a key step in untangling how the learned perceptual decisions of the CNN model relate to discrete spatial objects in a street-view image.

For computer vision models like CNN and SCNN, feature extraction techniques can also provide a level of explainability (Naik et al., 2014; Naik et al., 2016; Naik et al., 2017; Thackway et al., 2023b) or, at the very least, ensure that the model is actually focusing on the housing unit in the image rather than on unrelated noise. By ensuring that pixels can be categorically identified in the model (e.g., sky, car, tree), we can trace back performance errors to discrete, human-understandable objects and adjust the training dataset accordingly.

## 5. Conclusion and opportunities for future research

Contemporary literature identifies gentrification as a global process defined by class succession, landscape change, and the increased segregation of housing access. From the Global North to the Global South, these three qualities are inherent in every gentrification outcome, but the specifics that describe gentrification can vary depending on geographic and sociopolitical context. Different and sometimes unexpected metrics can capture how these neighborhoods vary in meaningful ways, rendering simpler, task-based methods of measurement prone to inaccuracy and incapable of reliably extrapolating trends from the data.

A recent renaissance of Machine Learning (ML) applications provides researchers with a newfound ability to deal with the stochasticity of real-world phenomena and model the gentrification process. Predictive models such as Random Forest and Extreme Gradient Boosting (XGBoost) can draw on immense volumes of data to produce reliable forecasts of future gentrification at spatially granular scales (Reades et al., 2018; Thackway et al., 2023a). Unsupervised algorithms like K-means clustering can categorize specific gentrification variants like marginal gentrification and super-gentrification, in addition to controlling for non-gentrification-related neighborhood change (Yee & Dennett, 2022). Deep mapping techniques with Neural Network (NN) models offer efficient ways to gather data on built metrics (Ilic et al., 2019; Thackway et al., 2023b). Lastly, recent advances in Artificial Intelligence (AI) computer vision software allow researchers to incorporate built metrics that were long absent from the model (Thackway et al., 2023b).

Various opportunities exist to progress our understanding of the gentrification lifecycle with these novel ML tools. Deep mapping measures for semantic classification of building instances have yet to find wide-spread use in application to gentrification modeling but offer an avenue to integrate variables on land use, zoning, building quality, and building characteristics into analyses of landscape change.

For future research, semantic image classification can provide more objective, workable components into the gentrification model by linking specific types of built development (e.g., a new roof) to gentrification processes. Moreover, semantically classified image data could serve as inputs to a clustering algorithm to better identify gentrification variants and uncover novel patterns of landscape change for each variant. For example, researchers understand that new-build gentrification entails new-build development (Davidson & Lees, 2005), but little is known about the types of zoning (e.g., residential, commercial, mixed), floor-area ratios, and development intensity that correspond with this gentrification type.

For both semantic and perceptual-based built characteristics, these

built metrics are solely drawn on to control for alternate forms of neighborhood change, but the research has yet to integrating information on landscape change as explicit variables within a modeling or forecasting mechanism. Such information would be monumental to transcribing academic findings to workable direction for municipal bodies that seek to uncover the links between planning policy, built development, and gentrification effects.

Beyond statistical XAI tools, ML models can become more trustable when they are nested in a diversified system of measurement and analysis. Even for non-learning-based systems, methods become more robust when triangulated, such as combining quantitative, qualitative, and field-survey derived findings (Loukaitou-Sideris et al., 2017). For example, qualitative studies on gentrification may be case specific but are nonetheless highly versed in local dialectics of social structure and extensive in their coverage of neighborhood change (e.g., Betancur, 2002). By synthesizing these findings with AI-based applications, we can ascertain that model findings run parallel to the narrative threads identified in journalistic accounts of a given locale.

Pioneering ML gentrification models have set a precedent for this by weaving in discussion of a priori, contextual knowledge of the study area with the results (Ilic et al., 2019; Reades et al., 2018; Thackway et al., 2023a; Yee & Dennett, 2022). Typically focusing on a single city, like London, Sydney, or Ottawa, the researchers relate ML model findings with the narratives of specific boroughs or neighborhoods. Context-specificity and sensitivity is essential to capturing the variable ways in which gentrification can alter local demographics and landscapes (Preis et al., 2020). Although ML methods can automate aspects of the analysis and be reproduced for diverse locales, the ways we implement the model, the preparation of the training data and the choice of validation techniques, must be adapted on a case-by-case basis.

As machine learning functions for usability over interpretability, it is important for new waves of research to lay out clear targets for how model output will be utilized in real-time. Although the body of research is still in the nascent phase of testing future applicability, pioneering studies in predictive machine learning models are pointed in how they evaluate the functionality of their forecasts: to supplement early warning systems, to place greater political pressure on policy makers to instigate policies, and to, ultimately, mitigate the occurrence of gentrification by accurately identifying at-risk neighborhoods (e.g., Palafox Novack, 2019; Reades et al., 2018; Thackway et al., 2023a; Yee & Dennett, 2022). Emphasizing the intended purpose of each machine learning application helps ensure that future reproductions of the research will evaluate the productivity of study findings in furthering community-based motives. Moreover, if the aim of contemporary scholars is to move towards global reproducibility, we must construct a more robust and integrated toolbox of gentrification measurement to ensure predictions can be extrapolated for different geographic contexts, and keep the more uninterpretable, black-box algorithms in check with other, highly interpretable systems of analysis.

## Funding

This research is funded by the National Science Foundation grant #2035176. The authors gratefully acknowledge the support.

## CRDiT authorship contribution statement

**Mueller Maya:** Conceptualization, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Hoque Simi:** Funding acquisition, Project administration, Supervision, Writing – review & editing. **Hamil Pearsall:** Supervision, Writing – review & editing.

## Declaration of competing interest

The authors have declared that no competing interest exists.

## Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.compenvurbsys.2024.102119>.

## References

Alejandro, Y., & Palafox, L. (2019). Gentrification prediction using machine learning. *Advances in Soft Computing. In 18th Mexican International Conference on Artificial Intelligence, MICAI 2019, Xalapa, Mexico, October 27–November 2, 2019, Proceedings 18*.

Angelov, P. P., Soares, E. A., Jiang, R., Arnold, N. I., & Atkinson, P. M. (2021). Explainable artificial intelligence: An analytical review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 11(5). <https://doi.org/10.1002/widm.1424>

Anguelovski, I. (2016). From toxic sites to parks as (green) LULUS? New challenges of inequity, privilege, gentrification, and exclusion for urban environmental justice. *Journal of Planning Literature*, 31(1), 23–36. <https://doi.org/10.1177/0885412215610491>

Ansell, B. W. (2019). The politics of housing. *Annual Review of Political Science*, 22, 165–185. doi:10.1146/.

Aoki, K. (1992). Race, space, and place: The relation between architectural modernism, post-modernism, urban planning, and gentrification. *Fordham Urb. L.J.*, 20, 699.

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... Benjamins, R. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>

Atkinson, R. (2000). Measuring gentrification and displacement in greater London. *Urban Studies*, 37(1), 149–165. <https://doi.org/10.1080/0042098002339>

Atkinson, R. (2008). Commentary: Gentrification, segregation and the vocabulary of affluent residential choice. *Urban Studies*. <https://doi.org/10.1177/0042098008097110>

Atkinson, R., & Bridge, G. (2004). *Gentrification in a global context*. Routledge.

Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495. <https://doi.org/10.1109/TPAMI.2016.2644615>

Balzarini, J. E., & Shlay, A. B. (2018). The strength of strong ties reconsidered: Social ties and collective power in a gentrifying community. *Social Currents*, 5(1), 67–85.

Barton, M. (2016). An exploration of the importance of the strategy used to identify gentrification. *Urban Studies*, 53(1), 92–111. <https://doi.org/10.1177/0042098014561723>

Bates, L. K. (2013). Gentrification and displacement study: Implementing an equitable inclusive development strategy in the context of gentrification. In *Urban Studies and Planning Faculty Publications and Presentations*. <https://doi.org/10.15760/report-01>

Bereitschaft, B. (2020). Gentrification central: A change-based typology of the American urban core, 2000–2015. *Applied Geography*, 118. <https://doi.org/10.1016/j.apgeog.2020.102206>

Bernstein, A. G., & Isaac, C. A. (2023). Gentrification: The role of dialogue in community engagement and social cohesion. *Journal of Urban Affairs*, 45(4), 753–770.

Bertinetto, L., Valmadre, J., Henriques, J. F., Vedaldi, A., & Torr, P. H. (2016). Fully-convolutional siamese networks for object tracking. In *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part II 14*.

Betancur, J. (2002). *The Politics of Gentrification: The Case of West Town in Chicago*. <https://doi.org/10.1177/107874037006002>

Blećić, I., Cecchini, A., & Trunfio, G. A. (2018). Towards automatic assessment of perceived walkability. *Computational Science and Its Applications–ICCSA 2018. In 18th International Conference, Melbourne, VIC, Australia, July 2–5, 2018, Proceedings, Part III 18*.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1), 1–8. <https://doi.org/10.1016/j.jocs.2010.12.007>

Borji, A. (2023). A categorical archive of chatgpt failures. arXiv. <https://doi.org/10.48550/arXiv.2302.03494>. preprint arXiv:2302.03494.

Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32. <https://doi.org/10.1023/A:1010933404324>

Brown-Saracino, J. (2017a). Explicating divided approaches to gentrification and growing income inequality. *Annual Review of Sociology*, 43, 515–539. <https://doi.org/10.1146/annurev-soc-060116-053427>

Brown-Saracino, J. (2017b). Explicating divided approaches to gentrification and growing income inequality. *Annual Review of Sociology*, 43(1), 515–539. <https://doi.org/10.1146/annurev-soc-060116-053427>

Brown-Saracino, J., & Ghaziani, A. (2009). The constraints of culture: Evidence from the Chicago dyke march. *Cultural Sociology*, 3(1), 51–75. <https://doi.org/10.1177/1749975508100671>

Butler, T., & Lees, L. (2006). Super-gentrification in Bartsbury, London: Globalization and gentrifying global elites at the neighbourhood level. *Transactions of the Institute of British Geographers*, 31(4), 467–487.

Chapple, K. (2009). *Mapping susceptibility to gentrification: The early warning toolkit* (p. 43). Berkeley, CA: Center for Community Innovation.

Chapple, K., & Zuk, M. (2016). *Forewarned: The use of neighborhood early warning systems for gentrification and displacement*. JSTOR Cityscape.

Chattopadhyay, A., Sarkar, A., Howlader, P., & Balasubramanian, V. N. (2018). Grad-cam ++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE winter conference on applications of computer vision (WACV)*.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*.

Cheng, G., Yang, C., Yao, X., Guo, L., & Han, J. (2018). When deep learning meets metric learning: Remote sensing image scene classification via learning discriminative CNNs. *IEEE Transactions on Geoscience and Remote Sensing*, 56(5), 2811–2821. <https://doi.org/10.1109/TGRS.2017.2783902>

Chitov, D. (2006). Cultivating social capital on urban plots: Community gardens in new York City. *Humanity and Society*, 30(4), 437–462.

Clarke, P., Ailshire, J., Melendez, R., Bader, M., & Morenoff, J. (2010). Using Google earth to conduct a neighborhood audit: Reliability of a virtual audit instrument. *Health & Place*, 16(6), 1224–1229. <https://doi.org/10.1016/j.healthplace.2010.08.007>

Core, M. G., Lane, H. C., Van Lent, M., Gomboc, D., Solomon, S., & Rosenberg, M. (2006). *Building explainable artificial intelligence systems*. AAAI.

Davidson, M., & Lees, L. (2005). New-build 'gentrification' and London's riverside renaissance. *Environment and Planning A*, 37(7), 1165–1190. <https://doi.org/10.1068/a3739>

Dieber, J., & Kirrane, S. (2020). Why model why? Assessing the strengths and limitations of LIME. *arXiv preprint*. <https://doi.org/10.48550/arXiv.2012.00093>. arXiv: 2012.00093.

Ellen, I. G., & Ding, L. (2016). Guest Editors' introduction: Advancing our understanding of gentrification. *Cityscape*, 18(3), 3–8.

Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). *A density-based algorithm for discovering clusters in large spatial databases with noise. Kdd*.

Finch, W. (2019). A comparison of clustering methods when group sizes are unequal, outliers are present, and in the presence of noise variables. *General Linear Model Journal*, 45, 12–22.

Finio, N. (2021). Measurement and definition of gentrification in urban studies and planning. *Journal of Planning Literature*, 37(2), 249–264. <https://doi.org/10.1177/088541222111051603>

Finio, N. (2022). Measurement and definition of gentrification in urban studies and planning. *Journal of Planning Literature*, 37(2), 249–264.

Freeman, L. (2005). Displacement or succession? Residential mobility in gentrifying neighborhoods. *Urban Affairs Review*, 40(4), 463–491. <https://doi.org/10.1177/1078087404273341>

Gebru, T., Krause, J., Wang, Y., Chen, D., Deng, J., Aiden, E. L., & Fei-Fei, L. (2017). Using deep learning and Google street view to estimate the demographic makeup of neighborhoods across the United States. *Proceedings of the National Academy of Sciences*, 114(50), 13108–13113. <https://doi.org/10.1073/pnas.1700035114>

Glass, R. (1964). *London: Aspects of change*. Macgibbon & Kee.

Goossens, C., Oosterlynck, S., & Brabd, L. (2020). Livable streets? Green gentrification and the displacement of longtime residents in Ghent, Belgium. *Urban Geography*, 41 (4), 550–572. <https://doi.org/10.1080/02723638.2019.1686307>

Gould, K. A., & Lewis, T. L. (2012). The environmental injustice of green gentrification: The case of Brooklyn's Prospect Park. In *2. The world in Brooklyn: Gentrification, immigration, and ethnic politics in a global city* (pp. 113–146).

Greene, S., & Pettit, K. L. (2016). *What if cities used data to drive inclusive neighborhood change?* Washington, DC: Urban Institute.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys (CSUR)*, 51(5), 1–42. <https://doi.org/10.1145/3236009>

Hammel, D. J., & Wly, E. K. (1996). A model for identifying gentrified areas with census data. *Urban Geography*, 17(3), 248–268. <https://doi.org/10.2747/0272-3638.17.3.248>

Hamori, S., Kawai, M., Kume, T., Murakami, Y., & Watanabe, C. (2018). Ensemble learning or deep learning? Application to default risk analysis. *Journal of Risk and Financial Management*, 11(1), 12. <https://doi.org/10.3390/jrfm11010012>

Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society: Series C: Applied Statistics*, 28(1), 100–108. <https://doi.org/10.2307/2346830>

Helms, A. C. (2003). Understanding gentrification: An empirical analysis of the determinants of urban housing renovation. *Journal of Urban Economics*, 54(3), 474–498.

Holten, J. E., Finch, W. H., & Kelley, K. (2011). A comparison of two-group classification methods. *Educational and Psychological Measurement*, 71(5), 870–901. <https://doi.org/10.1177/0013164411398357>

Hwang, J., & Sampson, R. (2014). Divergent pathways of gentrification: Racial inequality and the social order of renewal in Chicago neighborhoods. *American Sociological Review*, 79(4), 726–751.

Ilic, L., Sawada, M., & Zarzelli, A. (2019). Deep mapping gentrification in a large Canadian city using deep learning and Google street view. *PLoS One*, 14(3), Article e0212814.

Islam, M. R., Ahmed, M. U., Barua, S., & Begum, S. (2022). A systematic review of explainable artificial intelligence in terms of different application domains and tasks. *Applied Sciences*, 12(3), 1353.

Kang, J., Körner, M., Wang, Y., Taubenböck, H., & Zhu, X. X. (2018). Building instance classification using street view images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 145, 44–59.

Kiely, T. J., & Bastian, N. D. (2020). The spatially conscious machine learning model. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 13(1), 31–49.

Knorr, D. C. (2019). *Using machine learning to identify and predict gentrification in Nashville, Tennessee*.

König, R., Johansson, U., & Niklasson, L. (2008). Using genetic programming to increase rule quality. In *FLAIRS Conference*.

Krijnen, M. (2018). Gentrification and the creation and formation of rent gaps: Opening up gentrification theory to global forces of urban change. *City*, 22(3), 437–446.

Labi, N. (2012). Misfortune teller. *The Atlantic*, 35(3), 191.

Law, S., Paige, B., & Russell, C. (2019). Take a look around: Using street view and satellite images to estimate house prices. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(5), 1–19.

Lees, L. (2008). Gentrification and social mixing: Towards an inclusive urban renaissance? *Urban Studies*, 45(12), 2449–2470.

Lees, L. (2019). Planetary gentrification and urban (re)development. *Urban Development Issues*, 61(1), 5–13. <https://doi.org/10.2478/udi-2019-0001>

Li, X., Zhang, C., Li, W., Ricard, R., Meng, Q., & Zhang, W. (2015). Assessing street-level urban greenery using Google street view and a modified green view index. *Urban Forestry & Urban Greening*, 14(3), 675–685.

Liu, L., Silva, E. A., Wu, C., & Wang, H. (2017). A machine learning-based method for the large-scale evaluation of the qualities of the urban environment. *Computers, Environment and Urban Systems*, 65, 113–125. <https://doi.org/10.1016/j.comenvb.2017.06.003>

Liu, X., Zhou, Y., Zhao, J., Yao, R., Liu, B., & Zheng, Y. (2019). Siamese convolutional neural networks for remote sensing scene classification. *IEEE Geoscience and Remote Sensing Letters*, 16(8), 1200–1204.

Loukaitou-Sideris, A., Gonzalez, S., & Ong, P. (2017). Triangulating neighborhood knowledge to understand neighborhood change: Methods to study gentrification. *Journal of Planning Education and Research*, 39(2), 227–242. <https://doi.org/10.1177/0739456x17730890>

Lu, Y., Ferranti, E. J. S., Chapman, L., & Pfarrang, C. (2023). Assessing urban greenery by harvesting street view data: A review. *Urban Forestry & Urban Greening*, 127917.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.

Mazumdar, S., Learnihan, V., Cochrane, T., & Davey, R. (2018). The built environment and social capital: A systematic review. *Environment and Behavior*, 50(2), 119–158.

McKinnish, T., Walsh, R., & White, T. K. (2010). Who gentrifies low-income neighborhoods? *Journal of Urban Economics*, 67(2), 180–193. <https://doi.org/10.1016/j.jue.2009.08.003>

Miranda, M., & Lane-McKinley, K. (2017). *Artwashing, or between social practice and social reproduction: A Blade of Grass*.

Naik, N., Kominers, S. D., Raskar, R., Glaeser, E. L., & Hidalgo, C. A. (2017). Computer vision uncovers predictors of physical urban change. *Proceedings of the National Academy of Sciences*, 114(29), 7571–7576. <https://doi.org/10.1073/pnas.1619003114>

Naik, N., Philpott, J., Raskar, R., & Hidalgo, C. (2014). Streetscore-predicting the perceived safety of one million streetscapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*.

Naik, N., Raskar, R., & Hidalgo, C. A. (2016). Cities are physical too: Using computer vision to measure the quality and impact of urban appearance. *American Economic Review*, 106(5), 128–132. <https://doi.org/10.1257/aer.p20161030>

Odgers, C. L., Caspi, A., Bates, C. J., Sampson, R. J., & Moffitt, T. E. (2012). Systematic social observation of children's neighborhoods using Google street view: A reliable and cost-effective method. *Journal of Child Psychology and Psychiatry*, 53(10), 1009–1017. <https://doi.org/10.1111/j.1469-7610.2012.02565.x>

Owens, A. (2012). Neighborhoods on the rise: A typology of neighborhoods experiencing socioeconomic ascent. *City & Community*, 11(4), 345–369. <https://doi.org/10.1111/j.1540-6040.2012.01412.x>

Palafox, L., & Ortiz-Monasterio, P. (2020). Predicting gentrification in Mexico city using neural networks. In *2020 international joint conference on neural networks (IJCNN)*.

Palafox Novack, L. F. (2019). *Gentrification prediction using machine learning*. OPENAIRE.

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Pattillo, M. (2010). *Black on the block: The politics of race and class in the city*. University of Chicago Press.

Pearshall, H. (2018). New directions in urban environmental/green gentrification research. In *Handbook of gentrification studies*.

Preis, B., Janakiraman, A., Bob, A., & Steil, J. (2020). Mapping gentrification and displacement pressure: An exploration of four distinct methodologies. *Urban Studies*, 58(2), 405–424. <https://doi.org/10.1177/0042098020903011>

Reades, J., De Souza, J., & Hubbard, P. (2018). Understanding urban gentrification through machine learning. *Urban Studies*, 56(5), 922–942. <https://doi.org/10.1177/0042098018789054>

Reades, J., De Souza, J., & Hubbard, P. (2019). Understanding urban gentrification through machine learning. *Urban Studies*, 56(5), 922–942.

Rigolone, A., & Németh, J. (2020). Green gentrification or 'just green enough': Do park location, size and function affect whether a place gentrifies or not? *Urban Studies*, 57 (2), 402–420.

Rofe, M. W. (2003). "I want to be global": Theorising the gentrifying class as an emergent élite global community. *Urban Studies*, 40(12), 2511–2526.

Royall, E., & Wortmann, T. (2015). Finding the state space of urban regeneration: modeling gentrification as a probabilistic process using k-means clustering and Markov models. In *Proceedings of the 2015 14th international conference on computers in urban planning and Urban Management (CUPUM)*, Cambridge, MA, USA.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.

Rundle, A. G., Bader, M. D., Richards, C. A., Neckerman, K. M., & Teitler, J. O. (2011). Using Google street view to audit neighborhood environments. *American Journal of Preventive Medicine*, 40(1), 94–100.

Salesses, P., Schechtner, K., & Hidalgo, C. A. (2013). The collaborative image of the city: Mapping the inequality of urban perception. *PLoS One*, 8(7), Article e68400.

Schaffer, R., & Smith, N. (1986). The gentrification of Harlem? *Annals of the Association of American Geographers*, 76(3), 347–365.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*.

Seresinhe, C. I., Preis, T., & Moat, H. S. (2017). Using deep learning to quantify the beauty of outdoor places. *Royal Society Open Science*, 4(7), Article 170170.

Slater, T. (2017). Planetary rent gaps. *Antipode*, 49, 114–137.

Smith, G. (2020). Data mining fool's gold. *Journal of Information Technology*, 35(3), 182–194.

Smith, N. (1979). Toward a theory of gentrification a Back to the City movement by capital, not people. *Journal of the American Planning Association*, 45(4), 538–548. <https://doi.org/10.1080/01944367908977002>

Spinney, A., Reynolds, M. A., Wulff, M., & Atkinson, R. (2011). *Gentrification and displacement: The household impacts of neighbourhood change*.

Thackway, W., Ng, M., Lee, C.-L., & Pettit, C. (2023a). Building a predictive machine learning model of gentrification in Sydney. *Cities*, 134. <https://doi.org/10.1016/j.cities.2023.104192>

Thackway, W., Ng, M., Lee, C.-L., & Pettit, C. (2023b). Implementing a deep-learning model using Google street view to combine social and physical indicators of gentrification. *Computers, Environment and Urban Systems*, 102, Article 101970.

Tian, Y. (2020). Artificial intelligence image recognition method based on convolutional neural network algorithm. *IEEE Access*, 8, 125731–125744.

Turner, M., & Snow, C. (2001). *Leading indicators of gentrification in DC neighborhoods*. DC Policy Forum.

Van Criekingen, M., & Decroly, J.-M. (2003). Revisiting the diversity of gentrification: Neighbourhood renewal processes in Brussels and Montreal. *Urban Studies*, 40(12), 2451–2468.

Wang, W., Xiao, L., Zhang, J., Yang, Y., Tian, P., Wang, H., & He, X. (2018). Potential of internet street-view images for measuring tree sizes in roadside forests. *Urban Forestry & Urban Greening*, 35, 211–220.

Wardrip, K. (2011). *Public Transit's impact on housing costs*. Insights From Housing Policy Research.

Wei, F., & Knox, P. L. (2013). Neighborhood change in metropolitan America, 1990 to 2010. *Urban Affairs Review*, 50(4), 459–489. <https://doi.org/10.1177/1078087413501640>

Wilson, W. J. (2006). *The geography of opportunity: Race and housing choice in metropolitan America*. Rowman & Littlefield.

Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 2(1–3), 37–52.

Wuly, E. K., & Hammel, D. J. (1998). Modeling the context and contingency of gentrification. *Journal of Urban Affairs*, 20(3), 303–326.

Wuly, E. K., & Hammel, D. J. (1999). *Islands of decay in seas of renewal: Housing policy and the resurgence of gentrification*.

Wuly, E. K., & Hammel, D. J. (2004). Gentrification, segregation, and discrimination in the American urban system. *Environment and Planning A*, 36(7), 1215–1241.

Yee, J., & Dennett, A. (2022). Stratifying and predicting patterns of neighbourhood change and gentrification: An urban analytics approach. *Transactions of the Institute of British Geographers*, 47(3), 770–790. <https://doi.org/10.1111/tran.12522>

Zeng, J., Yue, Y., Gao, Q., Gu, Y., & Ma, C. (2022). Identifying localized amenities for gentrification using a machine learning-based framework. *Applied Geography*, 145, Article 102748.

Zeng, L., Lu, J., Li, W., & Li, Y. (2018). A fast approach for large-scale sky view factor estimation using street view images. *Building and Environment*, 135, 74–84.

Zuk, M., Bierbaum, A. H., Chapple, K., Gorska, K., & Loukaitou-Sideris, A. (2018). Gentrification, displacement, and the role of public investment. *Journal of Planning Literature*, 33(1), 31–44.

Wu, X., & Zhang, X. (2016). Automated inference on criminality using face images. *arXiv preprint*, 2, 4038–4052. [arXiv:1611.04135](https://arxiv.org/abs/1611.04135).

Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., & Lipson, H. (2015). Understanding neural networks through deep visualization. *arXiv preprint*. [arXiv:1506.06579](https://arxiv.org/abs/1506.06579).