Gentrification from the Sky: Using Remote Sensing and Machine Learning for Urban Change Detection

Javier Alfaro, Sanja Šćepanović, Stephen Law, and Daniele Quercia

Abstract Gentrification is an urban phenomenon marked by socioeconomic shifts that can displace long-term residents and increase inequality. Accurate measurement is essential for effective urban planning and equitable development. Traditional reliance on census data is costly, slow, and lacks the spatial and temporal resolution needed to detect neighborhood-level changes in real time. This study addresses these challenges by combining open satellite imagery with machine learning techniques to quantify gentrification more effectively. By analyzing high-resolution imagery, we detect physical changes—such as shifts in building density, rooftop materials, and green spaces—that are linked to gentrification but often overlooked by census-based approaches. Using the Greater London Area as a case study, our method improves measurement accuracy by up to 8%, achieving a balanced accuracy of 77% across 4,085 neighborhoods. Even a small improvement in accuracy can enable better identification of at-risk neighborhoods, helping policymakers intervene before displacement pressures become irreversible.

Keywords Open data · Satellite Imagery · Gentrification · Machine Learning

J. Alfaro (✉)
London School of Economics and Political Science, London, UK
e-mail: j.a.alfaro-garcia@lse.ac.uk

S. Šćepanović
Nokia Bell Labs, Cambridge, UK

University of Oxford, Oxford, UK
e-mail: sanja.scepanovic@nokia-bell-labs.com

S. Law
University College London, London, UK
e-mail: stephen.law@ucl.ac.uk

D. Quercia
Nokia Bell Labs, Cambridge, UK

King's College London, London, UK

Politecnico di Torino, Turin, Italy
e-mail: quercia@cantab.net

R. Goodspeed et al. (eds.), *Digital-Era Urban Transformations*, The Urban Book Series, https://doi.org/10.1007/978-3-031-98300-9_12

1 Introduction

In today's rapidly changing urban landscapes, the dynamics of gentrification have become increasingly critical to understand. Gentrification is characterized by socio-economic shifts that often lead to the displacement of long-term residents, alteration of neighborhood character, and rising social inequality. These changes pose significant challenges for urban planners, policymakers, and local communities. As cities evolve, accurately measuring gentrification is essential to mitigate its negative impacts and promote equitable development. Urban planning and policy decisions rely on timely and accurate data to understand these shifts, and without effective methods of quantification, interventions may fall short of addressing the needs of affected communities.

Traditionally, gentrification has been measured through census data, which focuses on socioeconomic indicators such as income, education, and housing characteristics. However, relying solely on census data presents several limitations. First, the process of collecting and analyzing census data is both time-consuming and costly, often resulting in delays that make it difficult to track the rapid evolution of neighbourhoods. Additionally, census data lacks the spatial and temporal resolution necessary to detect fine-grained changes at the neighbourhood level. As gentrification is often a gradual and localized process, subtle transformations—such as changes in building density, rooftop materials, urban greenery, or new infrastructure developments—can go unnoticed. These indicators, which reflect the physical aspects of urban change, are critical for a complete understanding of how neighbourhoods evolve over time. As a result, conventional methods are ill-suited for capturing the full picture of gentrification as it unfolds in real-time.

To address these challenges, this study introduces a novel approach that leverages open satellite imagery data and machine learning techniques to more effectively quantify the change of physical gentrification. By analyzing high-resolution satellite images, our method captures the physical changes in the urban environment that traditional datasets often overlook. These changes include modifications in building structures, rooftop materials, and green spaces, providing a more detailed and immediate picture of neighborhood transformations. Through the integration of satellite-derived features and machine learning models, we improve the accuracy of gentrification measurement by 8% over the baseline model, achieving a balanced accuracy of 77%. While this represents a measurable improvement, it also highlights the potential of linking physical changes observed in satellite imagery to socio-economic shifts, offering a complementary approach for studying gentrification dynamics.

The combination of physical and socio-economic data not only enhances the granularity and timeliness of gentrification measurement but also provides actionable insights for urban planners and policymakers. By applying this method to specific case studies within planning policy areas, this study demonstrates how satellite imagery and machine learning can contribute to a more equitable approach to managing urban development. We conclude with a discussion of the potential

applications of this model, particularly in shaping policy and fostering community-driven interventions in rapidly gentrifying areas.

2 Literature Review and Related Work

2.1 Gentrification in London

First introduced by Glass (1960), the term *gentrification* describes the social and physical transformation of economically disadvantaged neighborhoods due to the influx of wealthier households. This process is marked by socioeconomic shifts, such as changes in demographic composition, alongside visible physical upgrades in local infrastructure, both of which signal rising economic status at the neighborhood level.

British studies on gentrification have traditionally favoured qualitative approaches—media analysis, interviews, and ethnography—over quantitative methods (Watt, 2008). This is partly due to concerns that secondary data may overlook the underlying class dynamics driving urban change. However, studies like Barton (2016), which used bivariate analysis, have uncovered gentrification in areas not typically recognized by qualitative assessments, suggesting that a broader, data-driven perspective can reveal more widespread transformations.

In London, research has often focused on emblematic neighbourhoods like Hackney, Hoxton, Stratford, and Brixton, where the social and cultural markers of gentrification are most visible (Harris, 2012; Butler et al., 2013; Benson and Jackson, 2017). This focus on well-known areas can obscure the subtler, yet significant changes occurring across the city. To gain a more comprehensive understanding of gentrification, it is essential to apply quantitative and multivariate analysis to a wider array of neighborhoods, identifying potential "gentrification frontiers".

This study employs machine learning (ML) techniques to analyze urban change across 4,835 neighbourhoods in London from 2011 to 2021. By leveraging deep learning to extract latent features from satellite imagery, we aim to uncover hidden aspects of urban dynamics and relate them to gentrification. Integrating these insights with London Planning Policy geospatial data provides a more precise understanding of neighbourhood changes and offers actionable guidance for urban planning and policy decisions.

2.2 Machine Learning for Gentrification

Machine learning (ML) has been increasingly applied to predict urban changes, including gentrification. One of the most notable studies by Reades et al. (2019) used a Random Forest (RF) model to predict gentrification in London, combining

socio-economic indicators into a gentrification index using Principal Component Analysis (PCA). While this multidimensional approach helped capture the complexity of gentrification, it introduced challenges such as computational intensity and potential statistical errors.

Building on these advancements, researchers have explored alternative data sources to improve ML models. Jain et al. (2021) showed that Airbnb data, including listings and user reviews, can track neighborhood changes in real time. Similarly, Glaeser et al. (2018) used Yelp data on business activity to provide timely insights into urban transitions. Both studies address census limitations, offering real-time, and more granular information on gentrification.

Several studies have built upon this, exploring the advantages of ML over traditional regression models. Algorithms like RF and Gradient Boosted Machines (GBM) excel in their predictive accuracy and in prioritizing the most important features. For instance, Thackway et al. (2023a) used a wide range of census features and property data, including housing prices, Airbnb density, and proximity to points of interest like food and leisure establishments, to predict gentrification. These models offer flexibility in accounting for the varying predictive power of neighborhood characteristics across different contexts, such as proximity to transit or parks (Wardrip, 2011).

However, existing ML approaches still rely heavily on traditional demographic and socioeconomic data, similar to earlier studies, which often depend on infrequent tabular census data. These approaches overlook the potential of remote sensing, where open satellite data, available at weekly frequencies and continually improving in resolution, offer a more timely and detailed alternative for monitoring urban changes.

2.3 Urban Change Detection with Deep Learning

A key advancement in forecasting gentrification lies in the use of deep learning to analyze the built environment, often termed *deep mapping* or *machine mapping* (Ilic et al., 2019; Salesses et al., 2013). This approach, leveraging satellite, aerial, or Google Street View (GSV) imagery, automates the collection of detailed geospatial metrics, surpassing traditional land-use classification methods. By capturing subtle landscape changes, deep learning helps illuminate the evolving urban fabric crucial to understanding gentrification.

Computer vision techniques such as Convolutional Neural Networks (CNNs) are one of the core methodologies, which automatically learns image features—ranging from edges and textures to complex patterns—without predefined rules. This makes CNNs particularly powerful for analyzing urban imagery, as they can discern structural changes indicative of neighborhood transformations (Kang et al., 2018).

Recent studies, such as Ilic et al. (2019) and Thackway et al. (2023b), employed Siamese Convolutional Neural Networks (SCNNs) to detect gentrification by analyzing sequential GSV images. Ilic et al. (2019)‘s study in Canada identified

structural improvements in buildings over time, using these changes as proxies for gentrification. However, they noted a limitation: the absence of sociodemographic data could lead to overestimations, as not all property upgrades signify gentrification. Thackway et al. (2023b) extended this work in Sydney by integrating socioeconomic data into their predictions, creating a more robust model that accounted for both physical and demographic shifts.

Deep learning methodologies, particularly those using satellite imagery, have also been applied to urban vitality and housing price predictions (Law et al., 2019; Šćepanović et al., 2021a, b). These techniques excel at extracting nuanced details from high-dimensional data, offering new insights into urban dynamics. Building on a rich body of research in change detection (Bergamasco et al., 2022; Codegoni et al., 2022; Yin et al., 2023; Corley et al., 2024), our approach applies these advancements to enhance the spatial and temporal granularity for gentrification quantification, enhancing urban planning and decision-making through more precise, data-driven insights.

3 Data and Methods

This section describes the main data sources used for this study, the first dataset comprises census data from the UK and is used to profile neighborhoods in London. It comes from the Office for National Statistics (ONS, 2024). The second dataset entails remote-sensing imagery from the European Space Agency Copernicus Sentinel-2 mission (Copernicus Sentinel-2, 2021). A more detailed description of their processing is presented in the following sections.

3.1 Gentrification Score

To predict gentrification, we generate a score that captures changes in neighborhood socioeconomic conditions and create features from satellite data to predict such changes. We define neighborhoods as a Lower Layer Super Output Area (LSOA) in London because these represent granular administrative divisions for which governments provide socioeconomic data.

The LSOA contains between 1,000 and 3,000 inhabitants living in between 400 and 1,200 households: a geography small enough that even modest changes in the makeup of an area should show up, but large enough that the sample size of each is statistically robust and more privacy preserving. Whilst data are available on a finer scale (e.g., Output Areas) and a coarser scale (e.g., wards, Middle Layer Super Output Areas), LSOAs exemplify the characteristics of spatial proximity and social homogeneity which relate to *neighborhood effects* (Reades et al., 2019).

Inspired by Jain et al. (2021), we construct the gentrification score for each neighborhood (LSOA) based on changes in four socioeconomic measures: age,

Table 1 Socioeconomic measures for UK, London LSOAs

Measure	Definition	Data source 2011	Data source 2021
Age	Percent aged between 25	2011 ONS	2021 ONS
	and 34		
Education	Percent NVQ Level 4 or	2011 ONS	2021 ONS
Housing	Above Barriers to Housing and	IMD 2010	IMD 2019[a]
	Services Domain		
Income	Income Deprivation Domain	IMD 2010	IMD 2019

[a]The ONS collects data for the Indices of Multiple Deprivation (IMD) two years prior to release

education, housing, and income Table 1. These measures are grouped into two time periods: 2011 and 2021, then between these, the score is computed. First, socioeconomic measures are aggregated to create a neighborhood index for each LSOA. Values are standardized using percentile ranks, as raw Indices of Multiple Deprivation (IMD) values are not comparable across years (ONS, 2015). The neighborhood index is computed for the time windows:

$$\text{Neighborhood index}_{t1} = 1/4\text{x}\left(\text{age}_{t1}\,\text{education}_{t1}\,\text{housing}_{t1}\,\text{income}_{t1}\right) \quad (1)$$

Lower index values indicate more disadvantaged neighborhoods in terms of older, uneducated, poorer residents, or worse housing access.[1] Disadvantaged neighborhoods are defined as those in the bottom 50th percentile in 2011. Each neighborhood index is standardized using its percentile, and the gentrification score is aligned with literature:

$$\text{Gentrification score}_{t2} = \text{neighborhood index}_{t2} - \text{neighborhood Index}_{t1} \quad (2)$$

Here, t_2 is 2021, and t_1 is 2011. A higher gentrification score reflects a greater influx of younger, educated, wealthier residents, or improved housing access, whereas negative scores indicate disadvantaged neighborhoods with no signs of gentrification. These gentrification scores, derived from official census data, serve as the ground truth for evaluating our models. Although census data are collected at decadal intervals, our approach uses these labels to model gentrification trends, providing insights into long-term urban transformation. The spatial distribution of these scores is shown in Fig. 1.

In this study, we focus only on disadvantaged neighbourhoods because the concept of gentrification is usually discussed primarily in the context of less developed neighbourhoods. Even though affluent spaces may experience some change in

[1]According to the Office for National Statistics (ONS), 'Barriers to Housing and Services' refers to the physical and financial accessibility of housing and local services, captured through two subdomains: (1) 'geographical barriers', relating to the proximity of essential local services, and (2) 'wider barriers', which include factors such as affordability and homelessness.

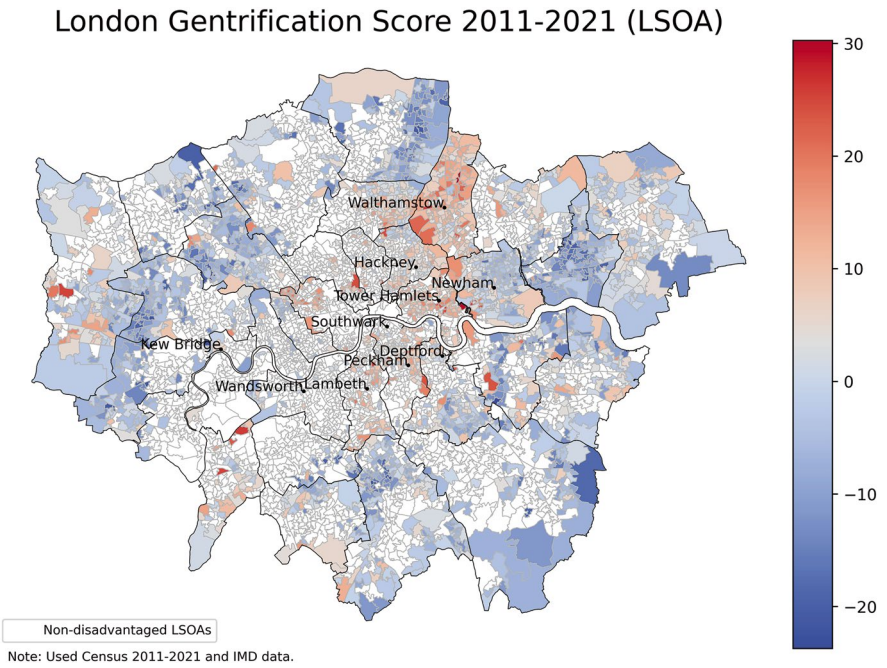

Fig. 1 Spatial distribution of gentrification scores, serving as the ground truth, calculated from changes in census data variables: age, education, housing, and income

socioeconomic measures, this is usually not considered to be gentrification (Zuk et al., 2018).

3.2 Satellite Imagery

In 2014, the European Space Agency (ESA) launched the first of the Sentinel satellites as part of the Copernicus Program (Torres et al., 2012), which aims at opening access to Earth Observation (EO) data to the public. Thanks to this investment, we can freely access accurate and timely data from the radar (Sentinel-1, Sentinel-3), altimeter (Sentinel-3), and optical (Sentinel-2) sensors (Šćepanović et al., 2021a, b).

For this study, we processed Sentinel-2 images for the years 2016 and 2021[2] through the (WASDI platform, 2023) by calculating average values during the summer season (June 1st to August 31st), as this period has fewer clouds over London, reducing obstructions in detecting physical changes. We focused on all 11 bands,

[2]The Sentinel-2 satellite was launched in 2015, limiting available data to June 2015 onwards. We focus on 2016 as the starting point, due to incomplete data for the 2015 composite.

capturing relevant environmental factors. To ensure data consistency and quality, we resampled the bands to a uniform resolution of 10 m, applied cloud masks to exclude affected pixels, and computed pixel-per-pixel averages over time under cloud-free conditions (Scepanovic et al., 2023). The resulting composite images represent the typical environmental characteristics for each season.

3.3 Deep Learning Feature Extractors

The full images comprise a dimension of (B, H, W) where B is the number of bands, H and W are the height and width of the image, the raw image size for the composite is (11, 4754, 6050), i.e., 11 bands each 4754 × 6050, each pixel represents parcels of land of 10 × 10 m^2. In order to process these images, we resized them to create imagelets of 256 × 256 pixels. This is also a standard practice in studies that process high resolution (HR) and very high resolution (VHR) satellite imagery (Corley et al., 2024; Fang et al., 2022; Ščepanović et al., 2021a, b; Yin et al., 2023).

To extract meaningful features from the satellite images, we applied a pipeline consisting of several deep learning techniques followed by image thresholding for change detection purposes. Given an image $i \in I$, where I is the set of imagelets, the goal was to extract a vector of visual components $v_{ib} \in V$ where i corresponds to each image, and b refers to each band. Some studies focus only on four bands (Papadomanolaki et al., 2019), but in order to fully leverage the whole dataset, the aforementioned strategy was employed.

In our experiments, we compare 6 methodologies, starting from a basic normalized image difference and then compiling their results. These are described as follows:

Simple-Diff (Sinergise, 2017) standardizes the image in t_2 with standard deviation and mean values from the image in t_1 to avoid detecting changes in brightness and contrast, and then computes the absolute difference between t_1 and the corrected t_2 image. Change is detected by using Multi-Otsu thresholding on this difference (Liao et al., 2001).

Res-Net demonstrates the use of Res-Net blocks (He et al., 2016) to extract deep image features for change detection with bi-temporal remote-sensed imagery.

FC-SiamDiff (Daudt et al., 2018) is a fully convolutional (FC) Siamese approach, which, instead of concatenating both features from the encoding streams, calculates the absolute value of their difference.

CGNet (Han et al., 2023) is a change-guiding network that improves upon the U-Net architecture by incorporating an attention-based Change Guide Module, designed to capture long-range dependencies across pixels.

Bi-Temporal Siamese Model (Papadomanolaki et al., 2019) uses a simple Siamese U-Net model that first extracts spatial features from multi-date inputs, which are subsequently used for change detection. The LSTM block was not included in this study.

TinyCD (Codegoni et al., 2022) is another Siamese UNet change detection architecture, which uses an EfficientNet (Tan and Le, 2019) encoder backbone to extract convolutional features and feed them to a custom attention-based decoder network.

These methods aim to perform Change Detection (CD) in bitemporal images I_1 and I_2 acquired at times t_1 and t_2, where, in our case, t_1 is 2016 and t_2 is 2021. Among them, Simple Differencing follows a traditional differencing approach, ResNet serves as a feature extractor, and the remaining models operate as end-to-end change detection methods.

3.3.1 Bi-temporal Siamese Network

Assume that a set of N unlabeled samples $X = X_n$, $n = 1, ..., N$ extracted from I_1 is available. These methods use X to train from scratch the convolutional autoencoders (CAE) in an unsupervised way. Once the CAE is trained, we process I_1 and I_2, and extract bitemporal deep feature maps of the inputs from the models. These features are compared and fused to detect changed π_c and unchanged π_{nc} pixels, where π_c includes all the relevant changes that occurred in the image, while π_{nc} represents no change. The Siamese approach, which extracts features from bitemporal images, is shown in Fig. 2 as a demonstration of how the aforementioned methods process the data.

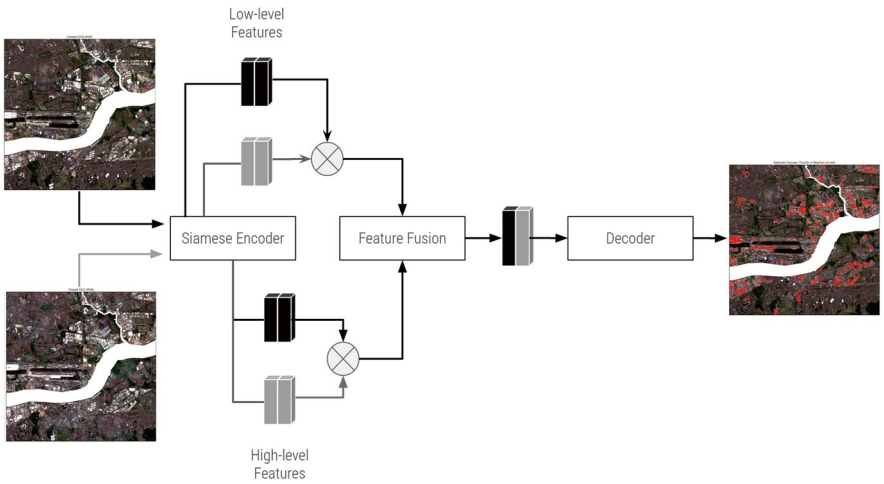

Fig. 2 Bitemporal Siamese Network Architecture for Urban Change Detection. The black and gray blocks represent satellite images captured at different times t_1 and t_2. Both high-level and low-level features are extracted from these images using a Siamese encoder. The feature maps from each image are fused and passed through a decoder, which detects urban changes, highlighted in red on the right. This process is applied independently to each spectral band of the satellite imagery

3.4 Implementation Details

The models are implemented in PyTorch (Paszke et al., 2017) and run on a single NVIDIA A100 GPU. We use the Adam optimizer (Kingma & Ba, 2014) with linear weight decay and a base learning rate of 0.001 for a single epoch across all models, except for the customized Res-Net, which runs for 100 epochs. The batch size is set to 8 for all experiments.[3]

For change detection, we employ a one-shot learning approach with Siamese neural networks, effective in learning object similarity (Nath, 2021). However, this was not applied to SimpleDiff, which relies on traditional differencing, or ResNet, which required additional training steps for meaningful feature extraction. The other models converged in a single training step, justifying the use of one-shot learning. This approach remains relatively unexplored in gentrification studies.

The models are trained using Mean Squared Error (MSE) loss following (Bergamasco et al., 2021, 2022). The goal is to minimize the reconstruction error between the original image and the predicted output over a batch $B \in X$ of patches:

$$MSE = \frac{1}{n}\sum_{n=1}^{B}\left(X_n - X_n^{'}\right)^2 \tag{3}$$

To extract feature maps, the model processes I_1 and I_2 separately. The encoder's initial layers capture simple features like edges and lines, while the decoder layers extract more spatial context about the changes. We apply bilinear interpolation to the retrieved features to ensure the maps have the same spatial dimensions for each patch. The image is then reconstructed by applying a no-data mask from I_1.

The main idea is to associate the label π_c, π_{nc} with each pixel, and to do so at the most reliable level. To that end, we apply a threshold τ retrieved using the (Otsu, 1979) algorithm to the output features φ from each model. The pixel of the final change map CM in position i, j is assigned to π_c or π_{nc} according to the class detected in the pixel position i, j.[4]

$$CM\left(i,j\right) \in \{\pi_{nc}, \text{if } \theta\left(i,j\right) \leq \tau_h\, \pi_c, \text{if} \theta\left(i,j\right) > \tau^4{}_h \tag{4}$$

3.5 Changes at the LSOA Level

A raster image is generated for each change map and band. To contextualize this process into our neighbourhoods, we translate these outputs into vector features. To handle large datasets efficiently, this workflow includes raster-to-vector conversion, spatial joining, and parallel processing.

[3] The code implementation can be found in the following GitHub repository: https://github. com/jaroxciv/Gentrification-from-the-Sky.

[4] Scikit-image's implementation of the Multi-Otsu algorithm outputs three classes by default. The subscript h indicates the highest threshold level was selected for the change detection π.

The extracted change pixel counts are then aggregated for each LSOA. The total change counts for each band are normalized by the area of the respective neighborhood to calculate the percentage of change pixels, which provides a measure of the intensity of the change within each area. This results in a matrix Φ_{nb} where each row represents an LSOA neighborhood (n), and each column corresponds to a specific extracted satellite band feature (b), which is used as a predictor of neighborhood gentrification.

4 Modeling

4.1 Setup

We cast this as a binary classification task whose dependent variable is the gentrification score computed in Sect. 3.1.

The first step is the binarisation of the gentrification score based on percentiles used to classify neighborhoods as either gentrified or not gentrified. This process involves setting a threshold at the top and bottom 25th percentiles of the gentrification scores. The resulting binary classification assigns a value of 1 to neighborhoods with gentrification scores above the top 25th percentile, and a value of 0 to areas with scores below the bottom 25th percentile.

Let G_n represent the gentrification score for neighborhood n, $Q_{75}G$ be the 75th percentile (top 25%) of the gentrification scores, and $Q_{25}G$ be the 25th percentile (bottom 25%) of the gentrification scores. The binary gentrification label GB_n for neighborhood n is defined as:

$$GB_n \in \{1, if\ G_n > Q_{75}(G) 0, if\ G_n < Q_{25}(G) dropped, if Q_{25}(G) \leq G_n \leq Q_{75}(G) \quad (5)$$

The second step is to select a proper baseline model for these binary labels. To do so, we run four independent classification models using housing, population churn,[5] income, and Gini coefficient as predictors (The Centre for Labour and Social Studies, 2021; Suss, 2023). The results show that the Logistic Regression model and LinearSVC provides the best results to predict our binary gentrification target with a value of 0.72 for both balanced accuracy and F1-score. Additionally, for each of the models, since both the house and satellite feature vectors were skewed, we used the Yeo-Johnson power transformation Weisberg (2001), a method previously used to predict gentrification in London (Yee & Dennett, 2022).

[5] The CDRC Residential Mobility Index was used in a study to predict gentrification.

4.2 Classification Methods

Given the produced neighborhood-level feature vectors Φ_{nb} (calculated in sub-Sect. 3.5 and their corresponding labels GB_n (calculated in Equation 5), we aim to predict $\hat{y}$, using a logistic regression of the form:

$$\hat{y}_n^{\alpha} = w_0^{\alpha} + w_1^{\alpha} x_1 + \ldots + w_n^{\alpha} x_n, \quad n = 1, \ldots, N \tag{6}$$

$$\hat{y}_n^{\Omega} = w_0^{\Omega} + \sum_{n=1}^{N} w_n^{\Omega} x_n + \sum_{n=1}^{N} w_n^{\Omega} \Phi_{nb}, \quad n = 1, \ldots, N \tag{7}$$

where the coefficients w are learned by their respective model, y_n is equivalent to GB_n, and N is the total number of neighborhoods.

In addition to the **Logistic Regression**, which minimizes the following log-loss function:

$$\sum_{(x,y)\in D} \left[-y log\left(\hat{y}\right) - \left(1-y\right) \log\left(1-\hat{y}\right) \right]$$

we also experimented with **Linear Support Vector Classification (SVC)**. The Linear SVC objective function is given by

$$\frac{1}{2}\|w\|^2 + C\sum_{n}^{N} \max\left(0, 1 - y_n \bullet \hat{y}_n\right)$$

where $\|w\|$ is the norm of the weight vector (reflecting model complexity), and C is a regularization parameter that controls the penalty on misclassifications. We chose to experiment with these two models because Logistic Regression provides a well-calibrated, interpretable baseline for binary classification, while Linear SVC is a more flexible alternative that can optimize decision boundaries more effectively in some cases (Ghosh et al., 2019).

Finally, we also tested an ensemble of decision trees with extreme gradient boosting, equivalent to **XGBoost Classification**, which minimizes the log loss function. On small structured datasets, decision tree-based methods are considered to be the best-performing methods (Chen & Guestrin, 2016; Shwartz-Ziv & Armon, 2022).

5 Results

5.1 Evaluation Metrics

To evaluate the performance of each classification method in predicting our binary gentrification labels (Equation 7), we use standard classification metrics: accuracy (balanced), F1-score (weighted) (Chase Lipton et al., 2014), and ROC-AUC (Hand & Till, 2001). These metrics are calculated based on the counts of true positives *tp*,

false positives *fp*, false negatives *fn*, and true negatives *tn*, typically represented in a confusion matrix. Under these terms, F1 can be expressed as a function of counts of true positives, false positives, and false negatives:

$$F1 = \frac{2tp}{2tp + fp + fn} \tag{8}$$

One last component to include, **specificity** *tn/(tn + fp)*, also known as true negative rate, helps us in defining balanced accuracy in a shorter way:

$$Balanced\ Accuracy = \frac{Sensitivity + Specificity}{2} \tag{9}$$

Balanced accuracy is crucial in this context to mitigate the bias towards the more frequent class (ungentrified areas), ensuring that both gentrified and ungentrified neighborhoods are given equal consideration during evaluation (Brodersen et al., 2010).

5.2 Gentrification Classifier Results

We compare the selected models from Sect. 4.2 using our evaluation metrics as a reference. For each model, we perform a Grid-Search ten-fold cross-validation to find the best possible baselines and optimal augmented models. After that, we apply a 30-fold cross-validation 5 times (repeated k-folds), and record the average scores and errors. The objective is to obtain results that are robust against a high variance derived from a low sample size (Šćepanović et al., 2021a, b).

The results are shown in Fig. 3. The first observation is that adding the satellite features yields improved scores in predicting gentrification. For the F1-score, this results in a 9% increase for the best-performing model (XGBoost), and an average improvement of 7% across all models. For ROC-AUC, the average improvement is 8%. Additionally, the scores between the Logistic Regression and the LinearSVC are comparable. These results suggest that a non-linear model is capable of capturing variation better. A plausible reason is that the Support Vector Machine algorithms are more geometrically motivated (Bennett & Bredensteiner, 2000), whilst the decision tree algorithm is better at discriminating the labels.

The table in Fig. 3 shows a numeric comparison of the findings, suggesting that the best model also has the lowest standard error in terms of ROC–AUC.

A heatmap analysis Fig. 4 highlights two key insights: **linear models' consistency** and **tree-based models' specialization**. Linear models show steady, reliable performance across different methods, while tree-based models yield more varied outcomes but offer flexibility through hyperparameter tuning. The result also shows that, among the different change detection methods, *TinyCD* and *Res-Net* appear to achieve the best overall performance (F1 = 0.78).

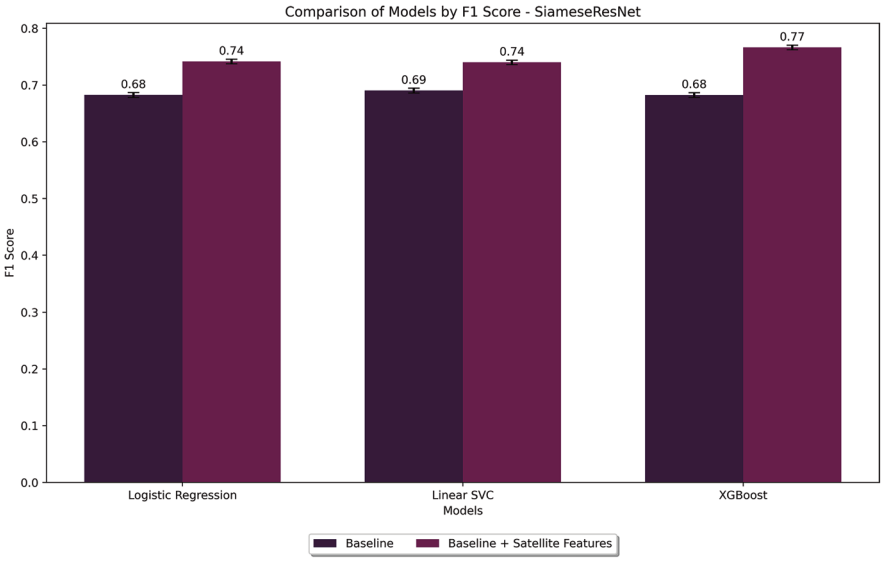

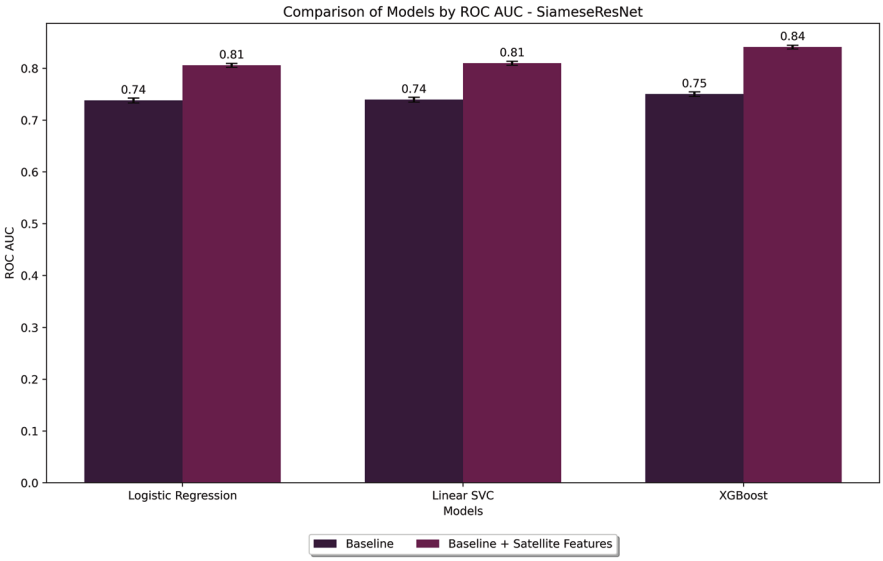

Method	Balanced Accuracy	F1 Score	ROC AUC
Logistic Regression	0.743 ± 0.051	0.742 ± 0.050	0.806 ± 0.047
Linear SVC	0.743 ± 0.050	0.740 ± 0.048	0.810 ± 0.046
XGBoost	**0.767 ± 0.049**	**0.766 ± 0.049**	**0.841 ± 0.043**

Fig. 3 Predicting gentrification from satellite features. The baseline model uses house median predictors, whilst the enhanced model adds an additional term $\sum_{n=1}^{N} w_n^{\Omega} \phi_{nb}$ from Equation 7. Balanced accuracy is omitted since it is equal to the F1-score results. Error bars represent the standard error of the mean across 30 folds

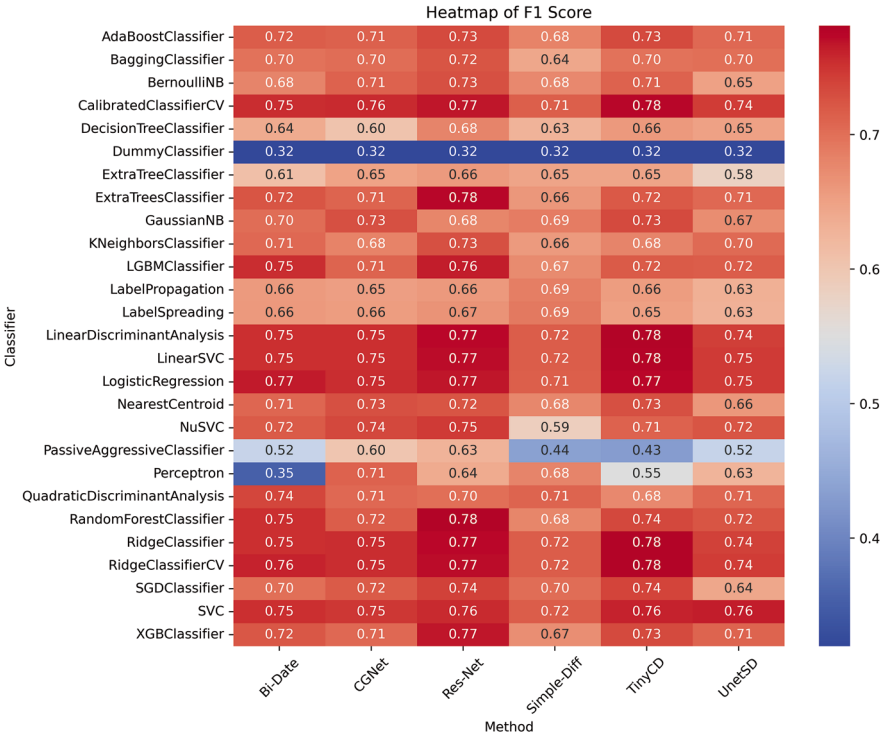

Fig. 4 The heatmap visualizes model performance, where the darker red indicates higher F1 scores, and the darker blue represents lower scores. Both linear and tree-based methods perform strongly, with tree-based models showing slightly more variation. XGBoost is ultimately chosen due to its high scores and flexibility in hyperparameter tuning, offering potential for further optimization

Additionally, overlaying detected changes on the gentrification score map in Fig. 5 reveals that different bands capture distinct spatial features. For instance, Band 4 highlights urban areas and soil changes, while Band 8 is more sensitive to vegetation and biomass variations. Changes are detected in areas with both high and low gentrification scores, suggesting the models are identifying transformations even in stagnant or highly gentrified areas.

5.3 Spatial Patterns

We first examine the evaluation results across neighborhoods, focusing on town centers and the presence of underground stations.

Newham has four main centers; in two of them, the model correctly predicts gentrification, while in the other two, it does not (Fig. 6). The figure clearly shows that the gentrified area has better transport connectivity, with four underground stations, as well as access to both the Overground and the Docklands Light Railway

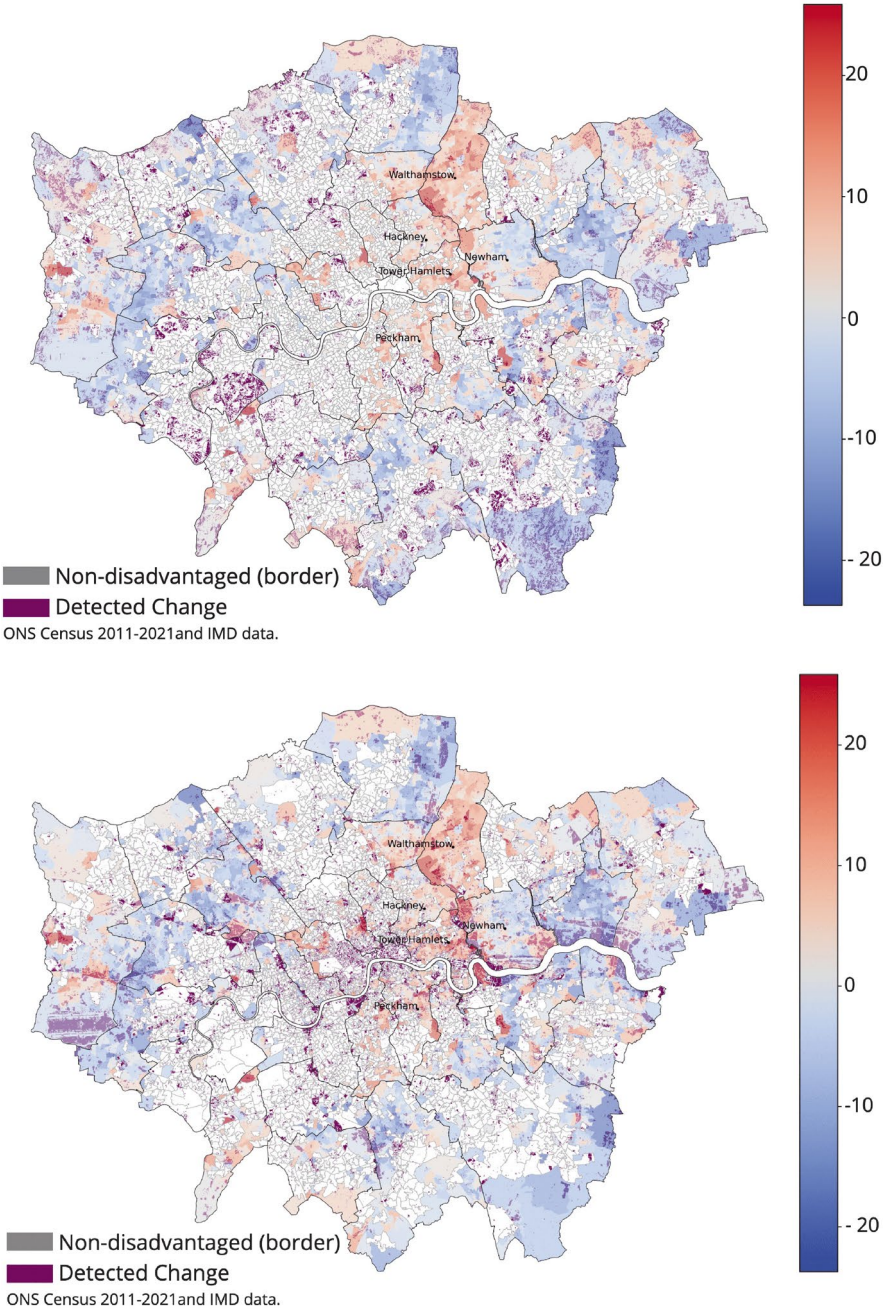

Fig. 5 London Gentrification Score with Detected Changes—LSOA. The top panel shows the detected urban changes from Band 8 (NIR), while the bottom illustrates changes from Band 4 (Red). Purple patches highlight areas of urban change as detected by our deep learning methods. Overlaying them reveals the spatial relationship between physical urban transformation and gentrification levels

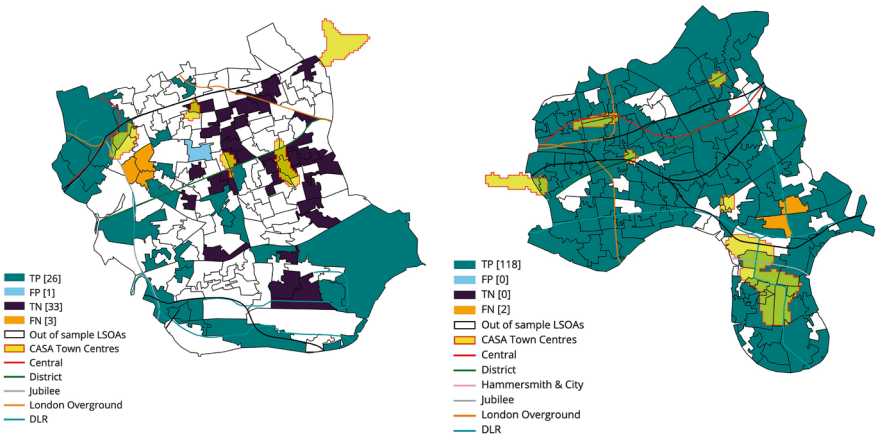

Fig. 6 Classification metrics (TP, FP, TN, FN) by Borough: Newham and Tower Hamlets

(DLR). In contrast, the center in the false negative area has only a single connection to the District Line. These regions have an average population of 31,528 and relatively high house prices, around £400,000.

The situation is similar in Tower Hamlets, which has a high level of connectivity with five underground stations, as well as access to both the Overground and the Docklands Light Railway (DLR). This borough contains six town centers, with average house prices ranging from £300,000 to £885,000. These centers have a lower average population of 22,358. Our model largely predicts gentrification correctly in this borough, aligning with previous academic studies and gray literature (Reades et al., 2019; The Centre for Labour and Social Studies, 2021).

5.4 Urban Change Detection

The previous maps offer a clear visual representation of the model's performance, showcasing areas where it successfully detects urban changes and areas where its predictions need improvement. However, understanding these changes goes beyond identifying where they are likely to occur—it also involves validating the accuracy and broader implications of these predictions.

The literature on urban change detection often uses well-labeled datasets like the Onera Satellite Change Detection (OSCD) dataset (Daudt et al., 2018), with models trained on them (Chen & Shi, 2020; Ji et al., 2019; Liu et al., 2022). Lacking such data, we validate our model using Conservation Areas—zones with strict regulations limiting physical and structural modifications to the built environment (City of London Corporation, 2024). By comparing satellite-derived change maps from

Equation 4 with Conservation Areas as a proxy for non-change, we can assess the accuracy of the change detection model. False positives may arise within these protected zones, while changes outside likely represent real urban transformations, as shown in Fig. 7.

The overlaid map validates our deep learning approach for detecting urban changes in Tower Hamlets, showing that 99% of the changes occurred outside conservation zones, where development is more permissible. Across all of London, 41,230 pixels (1.72%) of detected changes in Band 4 fell within conservation areas, further supporting the model's accuracy.

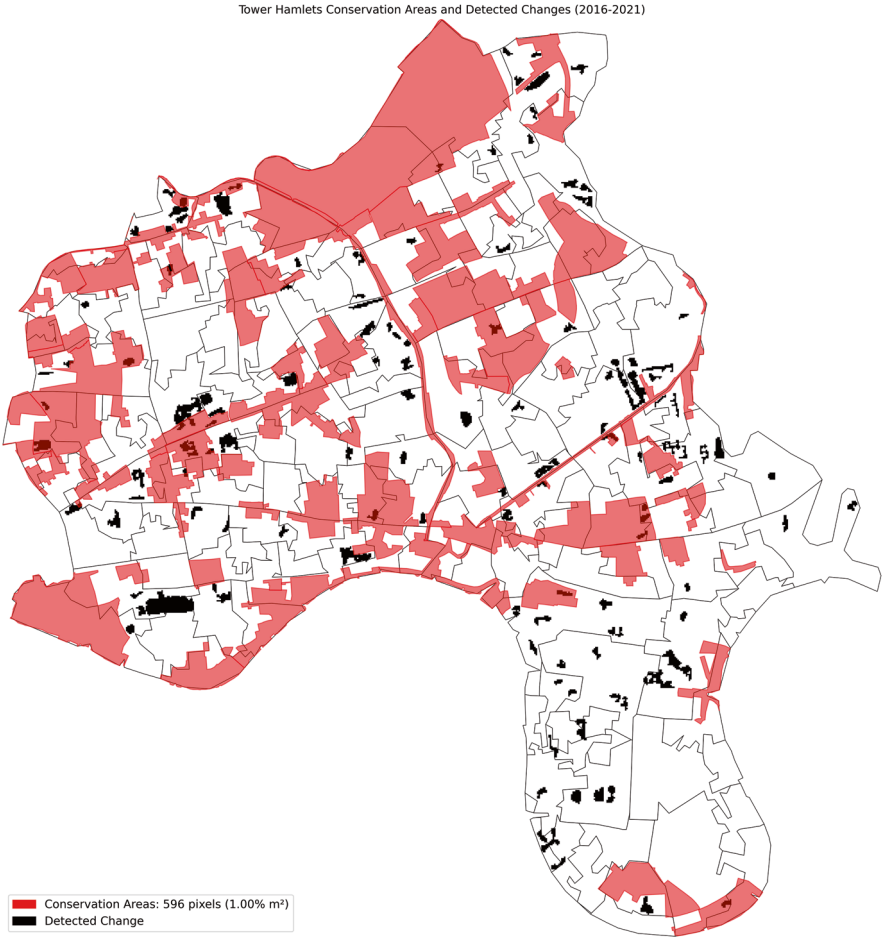

Fig. 7 Changes detected in Tower Hamlets within conservation areas (in red), using Band 4 of the TinyCD algorithm. Detected changes (in black) mostly occur outside these protected zones, where development is more likely. Changes within conservation areas may represent false positives due to strict regulations, while those outside likely reflect genuine urban transformations. This comparison provides insight into how well conservation policies prevent significant developments

Coverage across individual bands varies, from 0.22% (Band 2) to 42.39% (Band 8), with Bands 6 to 9 showing over 35% coverage within conservation zones, indicating sensitivity to certain urban changes. In contrast, Bands 2, 3, and 4 show more selective detection. On average, 19.55% of changes occur within conservation areas, reflecting the variability in how different bands capture urban transformations.

5.5 Ablation Study

The methods outlined in 3.3 define change by setting a threshold on the satellite imagery data. The choice of threshold can significantly impact the evaluation metrics as it influences the sensitivity of the model to detect true changes.

To investigate the impact of thresholding, we conducted an ablation study by defining a range of thresholds based on the minimum and maximum values identified for each spectral band. Changes were then computed at multiple points within this range.[6] Finally, we regressed only the satellite features on gentrification, accumulating accuracy and F1 scores at each threshold.

The ablation study reveals distinct curves for different models, suggesting variability not only at the threshold level but also in model specification (Fig. 8). The

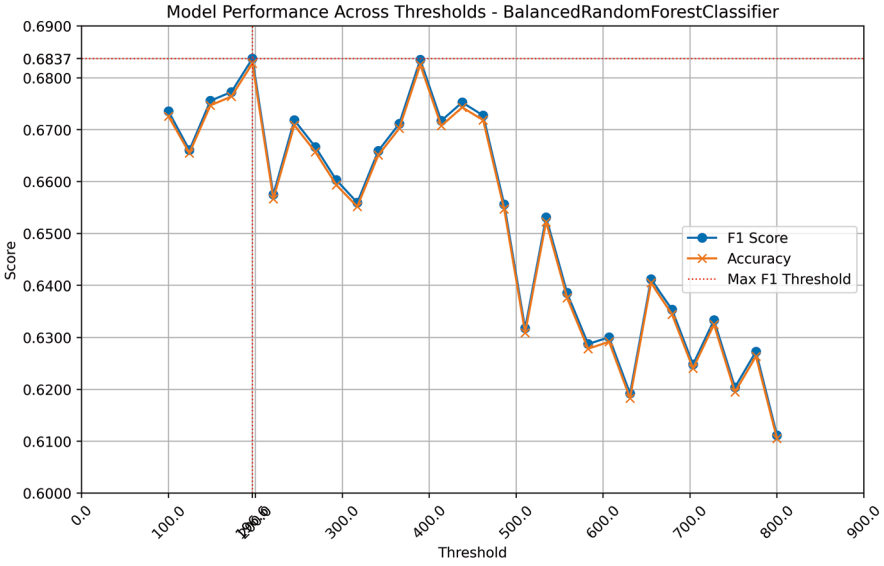

Fig. 8 Performance metrics across thresholds. Predictive model performance is sensitive to the choice of threshold, with an optimal range required to balance accuracy and F1 score effectively

[6] We use the simple-diff approach described in Sect. 3.3, which is both efficient in detecting changes and effectively demonstrates the robustness check.

study highlights two key peaks in performance: one around a threshold value of 200, which temporarily declines, and another around 400, where accuracy and F1 score reach a similar level. However, beyond 400, performance declines significantly. It is important to note that these threshold values may vary depending on the method used, further emphasizing the need for method-specific calibration to maintain predictive accuracy and reliability.

6 Discussion

6.1 Planning Policy

The observed spatial patterns in gentrification probabilities reveal associations with factors like transport connectivity and proximity to town centers, suggesting that accessibility and urban density are closely related to areas experiencing gentrification. However, these spatial patterns raise further questions about how policy decisions, infrastructure investments, and urban planning initiatives might correlate with these urban changes.

To explore these relationships, we analyze the UK government's official Planning Policy layers, which provide insights into land use designations, conservation areas, and transport infrastructure. By overlaying these policy layers with the detected urban changes from deep learning techniques, we aim to understand the associations between policy frameworks, infrastructure developments, and socio-economic shifts. Visualizing these interactions, as shown in Fig. 9, offers insights into how planning policies coincide with urban transformations, supporting arguments like those from Meltzer (2016), who noted that gentrifying neighbourhoods tend to experience employment growth and industrial restructuring.

Several key observations arise from analyzing the map. Local Industrial Locations account for the highest percentage of detected changes, with 12.48% of the total area, suggesting significant activity and potential redevelopment in these zones. The London Legacy Development Corporation (LLDC) within Newham follows with 9.90%, likely driven by ongoing large-scale projects, particularly in the context of post-Olympic regeneration (The Centre for Labour and Social Studies, 2021). Strategic Sites and Employment Hubs collectively represent around 17% of the detected changes.

Interestingly, a weaker relationship appears between detected changes and predicted gentrification probability in Local Industrial Locations. These areas, often smaller and concentrated, may lack the space for large-scale redevelopment, which could explain their low aggregated gentrification probability of 0.42 (Hubscher, 2021). The limited room for transformative projects likely constrains significant socio-economic shifts.

In contrast, larger zones like the Legacy Development Corporation LLDC within Newham and Strategic Sites exhibit higher gentrification probabilities, likely due to

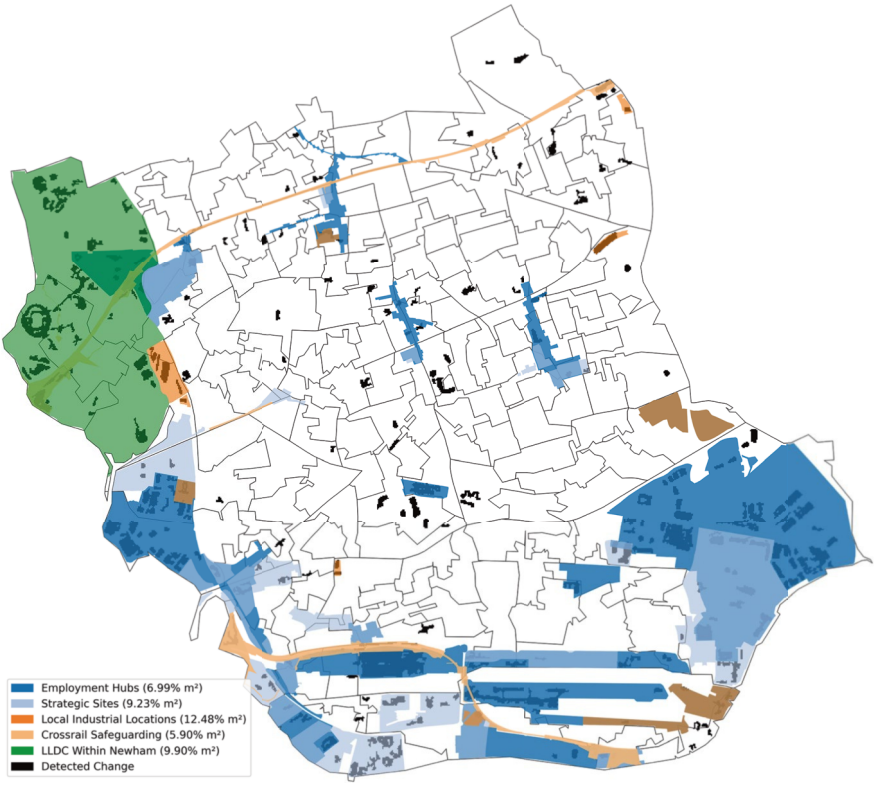

Policy Layer	Avg Probability
LLDC Within Newham	0.707
Strategic Sites	0.554
Employment Hubs	0.502
Crossrail Safeguarding	0.484
Local Industrial Locations	0.417

These probabilities are calculated using the XGBoost model across all neighborhoods.

Fig. 9 Detected changes within policy areas in Newham (top) and average probability of gentrification by policy layer (bottom). The legend shows the percentages of area in square meters corresponding to detected changes in the specified policy areas

their capacity for extensive infrastructure projects (Hubscher, 2021). These larger areas can more easily absorb spillover effects from development, stimulating gentrification across broader regions. This highlights how the scale and concentration of policy areas are associated with both physical changes and their socio-economic impacts.

While these findings offer valuable insights into the interaction between policy areas and detected changes, it is important to approach these results with caution. These results should be interpreted as indicative rather than definitive, highlighting potential areas of change but requiring further validation against ground truths and additional data sources.

7 Conclusion

Based on our analysis and results, we propose a method to evaluate urban change detection and gentrification probabilities using Sentinel-2 imagery and a deep learning framework. Our approach began with a baseline model, having the median house prices for 2011, 2016, and 2021 as predictors, and by incorporating satellite features, we were able to increase the predictive power by up to 8%. These enhanced probabilities were then used to uncover spatial patterns that, when paired with policy planning data, revealed deeper insights into the interplay of economic drivers and urban change. Our results are significant for two reasons. First, they demonstrate that satellite features, both alone (with 64% predictive power) and combined with baseline socioeconomic features like median housing price (77%), enable faster and large-scale gentrification assessments using globally available data. Second, the high temporal and spatial frequency of satellite data allows not only for detecting but also for tracking gentrification and its subtle interactions with the urban landscape.

Larger policy-supported areas, such as expansive policy zones or development sites, tend to have a higher potential for changes and, consequently, higher gentrification probabilities. The spread of new developments within these large zones often leads to spillover effects, where the impact of gentrification extends beyond the immediate vicinity of the changes.

Conversely, smaller, more concentrated development areas, like compact industrial zones or specialized policy areas, may exhibit a high percentage of detected changes due to their focused nature. However, their limited surface area and compactness constrain the scope of redevelopment and the potential for widespread gentrification. As a result, these areas may show lower aggregated probabilities of gentrification despite significant physical changes. This dynamic illustrates how the scale and density of a development zone (region) influence both the extent of urban change detected and the socio-economic outcomes that follow.

7.1 Practical Implications

The findings of this study offer several practical applications that can significantly influence urban planning and policy-making, especially in areas undergoing rapid change or facing potential gentrification pressures.

Enhanced Decision-Making through Probability Maps: The probability maps generated in this study can serve as powerful tools for urban planners and policymakers. By visualizing the likelihood of gentrification across different policy areas, these maps can help prioritize interventions in neighbourhoods at higher risk of socio-economic shifts. This targeted approach allows for more effective allocation of resources, ensuring that areas most in need of support receive timely attention.

Open Public Access Platforms for Community Engagement: Making these findings accessible to the public through open platforms and dashboards can empower communities by providing them with valuable information about their neighbourhoods. Residents can use this data to engage in informed discussions with policymakers, advocate for their needs, and participate more actively in the urban planning process. This transparency fosters a collaborative environment where both community members and decision-makers can work together towards equitable urban development.

Near Real-Time Monitoring of Urban Change: The use of open satellite data to detect changes within policy areas, rather than relying on private platform data such as Airbnb Jain et al. (2021) or Yelp Glaeser et al. (2018), offers an opportunity to develop a near real-time monitoring system for urban environments. By continuously updating the detected changes and probabilities of gentrification, cities can respond more swiftly to emerging trends, adapting policies as needed to mitigate negative impacts or capitalize on positive developments.

Guiding Policy Interventions in Vulnerable Areas: The variability in gentrification probabilities across different policy areas suggests that tailored interventions are necessary. Policymakers can use these insights to design specific strategies that address the unique challenges and opportunities within each area, ensuring that policies are both context-sensitive and effective in promoting sustainable urban growth.

7.2 Limitations and Future Work

Our approach for predicting and validating the gentrification prediction from satellite data comes with three main limitations:

Lack of Ground Truth Labels: Since ground-truth data on gentrification are unavailable, we follow previous work Jain et al. (2021) in assuming that socio-economic changes serve as a reliable proxy. Additionally, we infer the accuracy of our unsupervised gentrification detection method by examining its relationship with conservation areas. However, without explicit validation data, these relationships cannot be definitively confirmed.

Need for Human Annotations: To enhance the robustness of our model, integrating human annotations could be highly beneficial. Expert reviewers could validate a subset of the detected changes, providing ground truth labels that the

model can then use for retraining. This process would improve the accuracy and reliability of the model, ensuring it is better at distinguishing between true changes and noise.

Band Overfitting: While we have analyzed all available spectral bands, some such as Band 10 (used for cirrus cloud detection) might lead to "overdetection", potentially introducing noise and bias into the classification. Being wary of this issue can help refine the model's overall performance.

While our approach has demonstrated the potential of using Sentinel-2 imagery and deep learning to analyze urban change and gentrification probabilities, there are several avenues for future work that could further enhance the accuracy and utility of our method. First, improving the deep learning feature extraction process could lead to more precise detection of urban changes. This could be coupled with refined validation techniques that incorporate expert manual labelling, ensuring that the detected changes align closely with on-the-ground realities. Additionally, addressing potential multicollinearity issues within the spectral bands through the use of indices could further refine the models' predictions, making them more robust.

One limitation of our study is the use of medium-resolution imagery (10 m), which, while effective, could be improved with Synthetic Aperture Radar (SAR) data, such as Sentinel-1, which is unaffected by cloud cover, or very high-resolution satellite imagery, providing a more detailed view of urban dynamics. Additionally, integrating explainable AI techniques (Höhl et al., 2024) could enhance transparency in the feature extraction process, offering clearer insights into how specific urban features influence gentrification predictions. These advancements would not only refine the technical capabilities of our model but also create a more precise and reliable tool for urban planners and policymakers.

Acknowledgements The content of this chapter is primarily based on the first author's Masters Thesis, supported by the Consumer Data Research Centre (CDRC) and Nokia Bell Labs. The funding support from these institutions is gratefully acknowledged.

References

Barton, M. (2016). An exploration of the importance of the strategy used to identify gentrification. *Urban Studies, 53*(1), 92–111.

Bennett, K. P. & Bredensteiner, E. J. (2000). Duality and geometry in svm classifiers. In *ICML* (Vol. 2000, pp. 57–64). Citeseer.

Benson, M., & Jackson, E. (2017). Making the middle classes on shifting ground? r esidential status, performativity and middle-class subjectivities in contemporary l ondon. *The British Journal of Sociology, 68*(2), 215–233.

Bergamasco, L., Martinatti, L., Bovolo, F., & Bruzzone, L. (2021). An unsupervised change detection technique based on a super-resolution convolutional autoencoder. In *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS* (pp. 3337–3340). https://doi.org/10.1109/IGARSS47720.2021.9553859

Bergamasco, L., Saha, S., Bovolo, F., & Bruzzone, L. (2022). Unsupervised change detection using convolutional-autoencoder multiresolution features. *IEEE Transactions on Geoscience and Remote Sensing, 60*, 1–19. https://doi.org/10.1109/TGRS.2022.3140404

Brodersen, K. H., Ong, C. S., Stephan, K. L., & Buhmann, J. M. (2010). The balanced accuracy and its posterior distribution. In *2010 20th International Conference on Pattern Recognition* (pp. 3121–3124). https://doi.org/10.1109/ICPR.2010.764

Butler, T., Hamnett, C., & Ramsden, M. J. (2013). Gentrification, education and exclusionary displacement in e ast l ondon. *International Journal of Urban and Regional Research, 37*(2), 556–575.

The Centre for Labour and Social Studies. Pushed to the margins: A quantitative analysis of gentrification in london in the 2010s. Technical report, The Centre for Labour and Social Studies, United Kingdom, 2021. Retrieved from https://policycommons.net/artifacts/2650208/pushed-to-the-margins/3673081/ on 08 Aug 2024. CID: 20.500.12592/1wdtb3.

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, KDD '16* (pp. 785–794). Association for Computing Machinery. ISBN 9781450342322. https://doi.org/10.1145/2939672.2939785

Chen, H., & Shi, Z. (2020). A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *Remote Sensing, 12*(10). https://doi.org/10.3390/rs12101662. ISSN 2072-4292.

City of London Corporation. (2024). *Conservation areas*. https://www.cityoflondon.gov.uk/services/planning/historic-environment/conservation-areas. Accessed: 2024-08-11.

Codegoni, A., Lombardi, G., & Ferrari, A. (2022. ISSN 0941-0643). Tinycd: a (not so) deep learning model for change detection. *Neural Comput. Appl., 35*(11), 8471–8486. https://doi.org/10.1007/s00521-022-08122-3

Copernicus Sentinel-2. (2021). Msi level-2a boa reflectance product. collection 1. *European Space Agency*. https://sentinels.copernicus.eu/web/sentinel/sentinel-data-access/sentinel-products/sentinel-2-data-products/collection-1-level-2a. Processed by ESA

Isaac Corley, Caleb Robinson, and Anthony Ortiz. A change detection reality check. *arXiv preprint arXiv:2402.06994*, 2024. https://arxiv.org/abs/2402. 06994.

Daudt, R. C., Le Saux, B., & Boulch, A. (2018). Fully convolutional siamese networks for change detection. In *2018 25th IEEE International Conference on Image Processing (ICIP)* (pp. 4063–4067). IEEE.

Fang, S., Li, K., Shao, J., & Li, Z. (2022). Snunet-cd: A densely connected siamese network for change detection of vhr images. *IEEE Geoscience and Remote Sensing Letters, 19*, 1–5. https://doi.org/10.1109/LGRS.2021.3056416

Ghosh, S., Dasgupta, A., & Swetapadma, A. (2019). A study on support vector machine based linear and non-linear pattern classification. In *2019 International Conference on Intelligent Sustainable Systems (ICISS)* (pp. 24–28). https://doi.org/10.1109/ISS1.2019.8908018

Glaeser, E. L., Kim, H., & Luca, M. (2018). Nowcasting gentrification: Using yelp data to quantify neighborhood change. *AEA Papers and Proceedings, 108*, 77–82. https://doi.org/10.1257/pandp.20181034

Han, C., Chen, W., Guo, H., Meiqi, H., Li, J., & Chen, H. (2023). Change guiding network: Incorporating change prior to guide change detection in remote sensing imagery. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 16*, 8395–8407. https://doi.org/10.1109/JSTARS.2023.3310208

Hand, D. J., & Till, R. J. (2001). A simple generalisation of the area under the roc curve for multiple class classification problems. *Machine Learning, 45*, 171–186. https://api.semanticscholar.org/CorpusID:43144161

Harris, A. (2012). Art and gentrification: pursuing the urban pastoral in hoxton, london. *Transactions of the Institute of British Geographers, 37*(2), 226–241.

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90

Höhl, A., Obadic, I., Torres, M. Á. F., Najjar, H., Oliveira, D., Akata, Z., Dengel, A., & Zhu, X. X. (2024). Opening the black-box: A systematic review on explainable ai in remote sensing. *arXiv preprint arXiv:2402.13791*.

Hubscher, M. (2021. ISSN 2071-1050). Megaprojects, gentrification, and tourism. a systematic review on intertwined phenomena. *Sustainability, 13*(22). https://doi.org/10.3390/su132212827

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), 1–21. https://doi.org/10.1371/journal.pone.0212814

Jain, S., Proserpio, D., Quattrone, G., & Quercia, D. (2021). Nowcasting gentrification using airbnb data. *Proceedings of the ACM on Human-Computer Interaction, 5*(CSCW1), 1–21.

Ji, S., Wei, S., & Meng, L. (2019). Fully convolutional networks for multisource building extraction from an open aerial and satellite imagery data set. *IEEE Transactions on Geoscience and Remote Sensing, 57*(1), 574–586. https://doi.org/10.1109/TGRS.2018.2858817

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. https://doi.org/10.1016/j.isprsjprs.2018.02.006. ISSN 0924-2716. Deep Learning RS Data.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014

Law, S., Paige, B., & Russell, C. (2019. ISSN 2157-6904). Take a look around: Using street view and satellite images to estimate house prices. *ACM Transactions on Intelligent Systems and Technology, 10*(5). https://doi.org/10.1145/3342240

Liao, P.-S., Chen, T.-S., Chung, P.-C., et al. (2001). A fast algorithm for multilevel thresholding. *Journal of Information Science and Engineering, 17*(5), 713–727.

Lipton, Z. C., Elkan, C., & Narayanaswamy, B. (2014). Thresholding classifiers to maximize f1 score. *arXiv e-prints*, arXiv–1402.

Liu, M., Chai, Z., Deng, H., & Liu, R. (2022). A cnn-transformer network with multiscale context aggregation for fine-grained cropland change detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 15*, 4297–4306. https://doi.org/10.1109/JSTARS.2022.3177235

Meltzer, R. (2016). Gentrification and small business: Threat or opportunity? *Cityscape, 18*(3), 57–86. ISSN 1936007X. http://www.jstor.org/stable/26328273

Nath, S. (2021). Style change detection using siamese neural networks. In *CLEF (Working Notes)* (pp. 2073–2082).

ONS. (2015). *The english indices of deprivation technical report*. Technical report, Office of National Statistics. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/464485/English_Indices_of_Deprivation_2015_-_Technical-Report.pdf. Accessed: 2024-08-07

ONS. *Census statistics*. https://www.ons.gov.uk, 2024. Retrieved January 2024.

Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics, 9*(1), 62–66. https://doi.org/10.1109/TSMC.1979.4310076

Papadomanolaki, M., Verma, S., Vakalopoulou, M., Gupta, S., & Karantzalos, K. (2019). Detecting urban changes with recurrent neural networks from multitemporal sentinel-2 data. In *IGARSS 2019–2019 IEEE International Geoscience and Remote Sensing Symposium* (pp. 214–217). IEEE.

Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., & Lerer, A. (2017). Automatic differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*. https://openreview.net/forum?id=BJJsrmfCZ

Reades, J., De Souza, J., & Hubbard, P. (2019). Understanding urban gentrification through machine learning. *Urban Studies, 56*(5), 922–942. https://doi.org/10.1177/0042098018789054

Ruth Lazarus Glass. (1960). *London: aspects of change* (Vol. 3). MacGibbon & Kee.

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

Šćepanović, S., Joglekar, S., Law, S., & Quercia, D. (2021a). Jane jacobs in the sky: Predicting urban vitality with open satellite data. *Proc. ACM Hum.- Comput. Interact., 5*(CSCW1). https://doi.org/10.1145/3449257

Šćepanović, S., Antropov, O., Laurila, P., Rauste, Y., Ignatenko, V., & Praks, J. (2021b). Wide-area land cover mapping with sentinel-1 imagery using deep learning semantic segmentation models. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 14*, 10357–10374. https://doi.org/10.1109/JSTARS.2021.3116094

Scepanovic, S., Obadic, I., Joglekar, S., Giustarini, L., Nattero, C., Quercia, D., & Zhu, X. (2023). Medsat: A public health dataset for england featuring medical prescriptions and satellite imagery. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, & S. Levine (Eds.), *Advances in Neural Information Processing Systems* (Vol. 36, pp. 77810–77851) Curran Associates, Inc.

Shwartz-Ziv, R., & Armon, A. (2022). Tabular data: Deep learning is not all you need. *Information Fusion, 81*, 84–90.

Sinergise. (2017. https://medium.com/sentinel-hub/next-mission-automated-detection-of-land-changes-8e988dce55ff). *Next mission: Automated detection of land changes*. Sentinel Hub Blog.

Suss, J. H. (2023). Measuring local, salient economic inequality in the UK. *Environment and Planning A: Economy and Space, 55*(7), 1714–1737.

Tan, M., & Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning* (pp. 6105–6114). PMLR.

Thackway, W., Ng, M., Lee, C.-L., & Pettit, C. (2023a. ISSN 0264-2751). Building a predictive machine learning model of gentrification in sydney. *Cities, 134*, 104192. https://doi.org/10.1016/j.cities.2023.104192

Thackway, W., Ng, M., Lee, C.-L., & Pettit, C. (2023b. ISSN 0198-9715). Implementing a deep-learning model using google street view to combine social and physical indicators of gentrification. *Computers, Environment and Urban Systems, 102*, 101970. https://doi.org/10.1016/j.compenvurbsys.2023.101970

Torres, R., Snoeij, P., Geudtner, D., Bibby, D., Davidson, M., Attema, E., Potin, P., Rommen, B. Ö., Floury, N., Brown, M., Traver, I. N., Deghaye, P., Duesmann, B., Rosich, B., Miranda, N., Bruno, C., L'Abbate, M., Croci, R., Pietropaolo, A., Huchler, M., & Rostan, F. (2012). Gmes sentinel-1 mission. *Remote Sensing of Environment, 120*, 9–24. https://doi.org/10.1016/j.rse.2011.05.028. ISSN 0034-4257. The Sentinel Missions—New Opportunities for Science.

Wardrip, K. (2011). *Public transit's impact on housing costs: A review of literature*. Center for Housing Policy. https://scholar.google.com/scholar_lookup?title=Public%20Transit%E2%80%99s%20impact%20on%20housing%20costs&author=K.%20Wardrip&publication_year=2011

WASDI platform. *Earth observation tech for everyone*. https://www.wasdi.cloud, 2023. Online; accessed 3-June-2023.

Watt, P. (2008). The only class in town? gentrification and the middle-class colonization of the city and the urban imagination. *International Journal of Urban and Regional Research, 32*(1), 206–211.

Weisberg, S. (2001). *Yeo-johnson power transformations*. Department of Applied Statistics, University of Minnesota. Retrieved 1 June 2003.

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.

Yin, H., Ma, C., Weng, L., Xia, M., & Lin, H. (2023. ISSN 2072-4292). Bitemporal remote sensing image change detection network based on siamese-attention feedback architecture. *Remote Sensing, 15*(17). https://doi.org/10.3390/rs15174186

Zuk, M., Bierbaum, A. H., Chapple, K., Gorska, K., & Loukaitou-Sideris, A.- s. (2018). Gentrification, displacement, and the role of public investment. *Journal of Planning Literature, 33*(1), 31–44. https://doi.org/10.1177/0885412217716439