

Mapping The Price-Size Gradient: Implications for Automated Valuation Models in London's Housing Market

Wade Rogers

Laidlaw Scholars Program

September 2026

Faculty Mentor: Dr. Jon Reades

University College London

Bartlett Centre for Advanced Spatial Analysis

Abstract

Automated Valuation Models (AVMs), espoused as vehicles for modern tax reforms, increasingly rely on structural data like floor area to value homes at scale, yet existing analyses of the relationship between floor area and price paid per square meter (ppsqm) are limited by a lack of reliable public data, as well as failure to account for spatial heterogeneity. Using geographically weighted regression on a merged dataset joining Energy Performance Certificates (which, among other structural features, record floor area) with Price Paid Data recording the address, date, and price of full-market value property transactions, this paper interrogates an $n=25,000$ sample of the 2024 London housing market to better understand the spatial variation of the size-price relationship in the nation's capital. A base model predicting ppsqm with floor area alone was iteratively built up with structural and locational predictors of price as motivated by hedonic price theory, with spatial autocorrelation diagnostics and GWR run at each stage confirming the geographically-sensitive nature of the floor area-ppsqm relationship. Structural predictors (status as a flat, year of construction, etc) were found more powerful than locational predictors (distance to London's center, nearby bus stops, nearby rail stations) in attenuating the model and explaining floor area's impact on ppsqm. The direct implications for the fairness of AVMs assuming a spatial uniformity between size and price highlight the need for reliable and accessible housing data and offers a starting point for building a better understanding of the size-price gradient in the UK's capital and beyond.

Introduction

The UK's regressive 'council tax' system, which groups and taxes homes by what they would've sold for in 1991, has been widely criticized for its basis on outdated property values and deep unfairness in the modern housing market. Homes valued over £325,000 based on 1991 prices are capped at 'Band H' taxes, charged at most 3 times what 'Band A' homes which would sell for less than £40,000 in 1991 are taxed (Adam et al., 2020); in London, this has resulted in the poorest 10% of households being taxed over 6x more (8.1% of income) than the wealthiest 10% of households even after accounting for financial assistance (Murphy and Snelling, 2019).

Yet, mass valuations at a frequency reflective of a changing market—a commonly espoused vehicle to implement a more modern proportional property tax—are logistically infeasible using traditional surveyors. Consequently, the use of Automated Valuation Models (AVMs), mathematical models which factor in historical price data, property attributes, and location features to value homes en masse (GOV.UK, 2025), has been widely considered in implementing fairer tax policy. AVMs already have a foothold in policy; in 2024, Wales passed the Local Government Finance Act providing for revaluations of 1.5 million homes every 5 years beginning in 2028 to regularly update council tax bands (Welsh Government, 2025). A property's floor area is a basic input to any valuation: $\text{value} = \text{floor area} * \text{local price per square meter (ppsqm)}$. However, if the rate at which floor area translates to price is not spatially uniform, an

AVM assuming a single global relationship risks systematically misvaluing properties depending on their location and undermining the fairness it's meant to bring.

A lack of reliable, open data about internal floor area has historically hindered investigation into this relationship; existing public sources typically contain tangential observations (i.e. number of rooms) or are drawn from small samples that are difficult to generalize, like machine-learning analysis of floor plans on property portals (Özer and Jacoby, 2022). An approach devised by Chi et al. (2021) links Energy Performance Certificate (EPC) data—records of a property's structural features, namely floor area, relevant to its energy efficiency and compulsory for residential properties constructed, sold, or rented since 2007—to Price Paid Data (PPD), records of transactions sold at full-market value since 1995. EPC data currently provides the most accessible and comprehensive view of dwelling sizes, and through linkage to PPD offers a basis to interrogate size-price relationships. EPC and PPD data lack a reliable common identifier, but by using string-matching and spatially-constrained queries on addresses, Hubbard et al. (2023) achieved an 81% match rate and estimated that sub-standard homes (below 37m² internally) across the city cost 50% more per square meter than properties closer in size to London's mean 77m². Further, they found that there was a similar size premium on homes larger than average, suggesting an “effective penalty” is paid by those shopping for smaller homes.

Unaddressed by existing aggregate-level analysis is the spatial uniformity of these premiums and whether they can be attributed to other structural or locational factors that AVM might use. Accordingly, this paper samples 25,000 greater London area EPC records (valid as of 2024) following linkage with PPD data to justify and construct a geographically weighted regression model of the relationship between floor area and price paid per square meter (ppsqm) in the capital. By adjusting the model with structural and locational predictors of price, it aims to provide a truer-to-life, higher resolution image of today's size-price gradient in the UK's capital.

Methods

Data Sourcing

Since 2007, UK regulations have mandated that residences built, sold, or rented have certain structural features—including floorspace—assessed to inform an energy efficiency rating. This information is recorded in the property's EPC record, valid for 10 years upon administration, and maintained by each UK constituent country in a publicly accessible database; England's bulk EPC data is accessed through the UK Ministry of Housing, Communities and Local Government. While currently the most comprehensive dataset available, EPC coverage is incomplete; as of March 2024, the Office of National Statistics (ONS) found that 71% of English homes have EPC data, with London's figure at 73% (ONS, 2025). PPD data, which records the value, address, and date of *full market-price* transactions, has been publicly maintained in England and Wales by the UK's HM Land Registry since 1995 (ONS, 2014). As such, this paper only addresses the

purchase market since 1995; homes rented, gifted, or under shared ownership are among those with no price data and are excluded from analysis.

The datasets' only shared field is address, where entries are vulnerable to human error; assessors record streets, flat numbers, and block names differently on account of idiosyncrasy and mistakes, so genuinely matching records can fail a direct join. A further complication is resolving the nature of EPC certificates as 10-year windows into a property's history with PPD's dated instance-in-time format. This paper uses a merged database prepared by University College London's Centre for Advanced Spatial Analysis and derived from Hubbard et al.'s approach, which reconciles address input errors using fuzzy matching (Splink Python package) for street/building name and deterministic matching for discerning individual dwellings within that street/building. Because an address can accrue several EPCs over time, reflective of both logging errors and genuine changes to the property, statistical thresholds group EPCs by size regimes and collapse them into unique, dated 'phases' of development; this addresses the temporal question of deciding which EPC describes an address at the time of purchase. After filtering for outliers and errors flagged during linkage, the final analytical file consists of matches which have been confidently individuated.

Of the 93,858 linked records in the dataset's London 2024 subset, an $n=25,000$ sample was taken from the 87,267 records containing valid entries for the 9 fields used as ppsqm predictors in the final model, drawn from 2024 as the most recent complete year of data available at the time of analysis. This sample was used even for models that don't utilize all 9 fields for consistency between comparisons. Predictor field selection was motivated by hedonic price theory, which, in the context of housing, says a property's price (and, by extension, ppsqm) can be inferred by a combination of its structural features and situation among amenities (Smith, 2010); this is the core theory behind the majority of AVMs. Accordingly, 7 continuous fields—floor area, current energy efficiency, construction year, attachment style (encoded as 0 = detached, 1 = semi-detached, 2 = end-terrace, 3 = mid-terrace), distance from London's center (defined as London Waterloo Rail Station), number of bus stops within a 500 meter radius, number of 'rail nodes' within a 500 meter radius—and 2 binary fields—flat status and new build status (true for first EPC inspection prior to sale)—are used as ppsqm predictors. Continuous values were robustly scaled about their median using a 10th-90th percentile scale factor to protect against outliers. The vast majority of records contained a Unique Property Reference Number (UPRN) and all contained a postcode; address locations were largely determined using UPRN-coordinate matchings, and when necessary, postcode-coordinate matchings, both provided by the Ordnance Survey (OS). One UK postcode contains at most 100 UPRNs and averages just 15, making it a sound proxy for location (OS, 2021). Records with identical coordinates were jittered $\leq 3\text{m}$, breaking ties that would result in division by 0 while leaving spatial relationships practically untouched. Locational predictors were not included in the dataset, but calculated from the National Public Transport Access Nodes database maintained by the Department of

Transportation. The sample included addresses within a 2500m buffer from London's boundary to inform kernel-weighting of points on the outskirts, which have fewer neighbors; these only provided context to protect from edge effects and were not included in a model's final analyses. All coordinates, operated on using the geopandas library for Python (the language used for this study), were provided in the EPSG:4326 reference system and converted to EPSG:27700 for analysis and visualization.

Statistical Analytics

Ordinary Least Squares (OLS) regression, a standard statistical tool used in AVMs (Kiely and Bastian, 2019), constructs a best-fit line from aggregate data whose slope represents a single global regression coefficient per predictor. OLS was implemented using the statsmodels library and run for each model iteration. Both floor area and ppsqm are right-skewed, and the area-ppsqm relationship is expected to curve with diminishing marginal returns to space as fixed costs spread out across expanding floor areas (Feng et al., 2021); both variables were log normalized to adjust for skew and to fit area-ppsqm's non-linear relationship to a linear model. This yields the following specification:

$$\ln(\text{ppsqm})_i = \beta_0 + \beta_1 \cdot \ln(\text{floor_area})_i + \sum \beta_k \cdot X_{ki} + \varepsilon_i \quad (1)$$

where $\ln(\text{ppsqm})_i$ is the log-transformed price per square meter of property i , β_0 is the intercept, β_1 is the coefficient on log floor area, X_{ki} represents the k -th additional structural or locational predictor for property i , and ε_i is the error term. Since ppsqm is the ratio of price to floor area, the floor area coefficient represents elasticity, or the percent change in ppsqm resulting from a 1% change in floor area. For each iteration of the model, OLS yielded a ppsqm residual map alongside a global adjusted R^2 , which penalizes R^2 for the addition of weak predictors. The clustering of over and under predictions in OLS is expected to justify Geographically Weighted Regression (GWR) by illustrating that errors are not random noise but follow a spatial pattern (Oshan et al., 2019). Local and global Moran I's (significant at $p < 0.0010$, 8 nearest neighbours) and Local Indicator of Spatial Association (LISA) diagnostics (significant at $p < 0.05$), implemented through the esda library, formally probe the residual map for this. The

GWR, implemented via the mgwr library, fits a separate, locally-weighted regression to each address by considering nearby observations (based off of a selected bandwidth distance) more heavily, represented as:

$$\ln(\text{ppsqm})_i = \beta_0(u_i, v_i) + \beta_1(u_i, v_i) \cdot \ln(\text{floor_area})_i + \sum \beta_k(u_i, v_i) \cdot X_{ki} + \varepsilon_i$$

where (u_i, v_i) denotes the spatial coordinates of property i , and each coefficient is now a function of location rather than the single global value estimated by OLS above. This lets each predictor's coefficient vary across space, promising a more granular understanding of the strength of

relationships at different locations, and, specifically, how they tend to cluster (Tapia et al., 2025). A literature review by Valier (2020) found that a majority of studies (53 of 70) determined machine learning AVM models were more accurate than deterministic hedonic models like GWR; however, these black box models sacrifice explainability for stability, making them less preferred for AVM. A golden-section search tries candidate bandwidths (explicitly bounded at 300m and 2500m after empirical testing found a tendency for the search to crash outside of these bounds) to determine which scale applied across a Gaussian-weighted kernel minimizes the corrected Akaike Information Criterion (AICc). AICc was used instead of standard AIC in this fixed-bandwidth selection to protect from unstable, overfit local regressions (Oshan et al., 2019). Local coefficients were mapped with a diverging color scale, and points with local t-values statistically insignificant at a 95% confidence interval were grayed out. Predictor-vs-predictor collinearity checks were run on each iteration's OLS model to inform which predictors relate to floor area and, consequently, are likely to explain changes between models in floor area's coefficient distribution. Both OLS and GWR use an intercept column of 1's.

Mapping

Generated maps display contextualizing information such as greenspace, lakes, rivers, roads, built environment, local authority regions, and major rail stations using geographic data provided by the OS. The Middle Super Output Areas (MSOA) containing the 3 most extreme positive and negative coefficients (that are ≥ 3 km apart) are labeled; labels are provided by the ONS.

Results

Table 1: Model diagnostics across base, structural, and final

Model	Base	Structural	Final
Predictors	1	6	9
Bandwidth	402m	1050m	1300m
OLS Adj. R^2	0.001	0.032	0.280
GWR R^2	0.626	0.656	0.649
GWR Adj. R^2	0.584	0.634	0.629
AICc	1060.866	-3428.070	-3260.873
Moran's I (OLS resid.)	0.576	0.548	0.409

Table 1: Model diagnostics across base, structural, and final GWR specifications. GWR R^2 and adjusted R^2 are penalized for effective degrees of freedom ($\text{trace}(S)$); AICc is reported for the selected bandwidth at each stage.

Table 2: Predictor coefficients in OLS and GWR, final model

Predictor	OLS Coefficient	GWR Range (Min-Max)
Floor Area	-0.036	-0.375 to 0.072
Nearby Bus Transport	0.206	-0.615 to 0.271
Nearby Rail Transport	0.257	-0.404 to 0.269
Current Energy Efficiency	0.077	-0.299 to 0.121
Distance from Centre	-0.506	-1.480 to 1.272
Construction Year	-0.085	-0.215 to 0.557
Attachment	0.063	-0.188 to 0.099
Flat Status	0.103	-1.460 to 0.589
New Build Status	0.188	-1.457 to 0.755

Table 2: OLS coefficients versus GWR coefficient mean and range by predictor, final model. GWR values reflect all calibration points (n=22,083), including points not statistically significant at the 95% confidence level.

Base Model (Floor Area)

The model's first iteration used floor area as the sole predictor of ppsqm. Global OLS yields a negligible adjusted R^2 (Table 1), and its residual map (Figure 1) systematically undervalued homes (positive residual) in wealthier West London and overvalued homes (negative residual) in East London, as well as in the suburbs and beyond. Global Moran's I for the OLS residuals read 0.576 – positive and significant at $p < .0010$, confirming real spatial clustering rather than noise. This result is consistent with the Moran local scatterplot, whose points are densely concentrated in the HH and LL quadrants. The LISA cluster map placed these points at their real geography; there was strong visual correspondence between the OLS residual map and the LISA cluster map (Figure 2), also indicative of clustering. 11,084 of 22,853 locations (2,147 of the 25,000 were buffer points discarded from the final analysis) were significant at $p < 0.05$.




Figure 1: Map of base model OLS residuals in predicting log ppsqm.




Figure 2: (From left to right) Base model OLS Moran's Scatterplot, LISA cluster map, and residual quantile Chloropleth map. HH/HL indicates a point whose residual and the average residual of its 8 nearest neighbors are both high or low, respectively. HL/LH represents a mismatch between the two. ns represents statistically insignificant points at $p < 0.05$.

When GWR was applied, model fit improved substantially (Table 1). GWR's residual map was far more uniform, indicating it's accounting for spatial clustering (Figure 3). While the floor area coefficient is statistically insignificant at the 95% confidence level for the majority of addresses, there are significant (although relatively moderate) pockets of both positive and negative correlation near central London (Figure 4). Positive extremes are in the peripheral MSOAs of Keston, Pinner Green, and 'Chislehurst West & Elmstead' (though there is noticeable positive hotspot just east of King's Cross station), and negative extremes are in the peripheral MSOAs of Hounslow Heath, South Thamesmead, and 'West Thornton South & Ampere Way.'




Figure 3: Map of base model GWR residuals in predicting log ppsqm.




Figure 4: Map of base model GWR regression coefficients for floor area predicting ppsqm.

Second Model (Floor Area, Structure)

The model's second iteration added 5 predictors related to an address's structural attributes: attachment style, flat status, new build status, construction year, and energy efficiency (attributes a surveyor might use to value a home). Global OLS and Moran's I showed improvement (Table 1) while LISA diagnostics found 10,383 of 22,853 locations significant at $p < 0.05$; the structural model was confidently found to retain a high degree of spatial autocorrelation. The Moran Local Scatterplot and LISA cluster map had similar visuals to the OLS residual map, as was observed in the base model. 10,383 of 22,853 locations were significant at $p < 0.05$, a slight decline from the base model reflective of the decrease in the Global Moran's I.

GWR also improved under this model (Table 1), and floor area's coefficient distribution changed quite drastically with far more points were statistically significant at the 95% confidence level. Noticeably, large swaths of significant negative correlation emerged with extremes in the peripheral MSOAs of Enfield Wash and 'North Feltham & Hatton' and the central MSOA of 'Old Street & St. Luke's,' though statistical insignificance still blankets a vast area comprising most of the eastern greater London area. A single pocket of positive correlation remains just east

of London's King's Cross station (Figure 5). Global collinearity checks found flat status to have the strongest relationship with floor area at a value of -0.584; the 2nd strongest relationship is with construction year at -0.236.




Figure 5: Map of second model's GWR regression coefficients for floor area predicting ppsqm.

Final Model (Floor Area, Structure, Location)

The model's third and final iteration adds 3 predictors related to an address's location: distance to London's center, the number of bus stops within 500m, and the number of rail stations within 500m. Global OLS and Moran's I improved again (Table 1), and LISA diagnostics found 8,829 of 22,853 locations significant at $p < 0.05$; a diminishing yet high degree of spatial autocorrelation persisted through to the final model. The Moran Local Scatterplot and LISA cluster map had similar visuals to the OLS residual map, as was observed in the base and second model.

GWR diagnostics slightly worsened under this model (Table 1). Floor area's coefficient distribution remained largely unchanged from the second model's, though the numerical range of local elasticities did shrink very slightly while the physical regions where these elasticities were significant grew slightly in size in this model. Negative extremes were most common in the

south London MSOAs of ‘Putney Town & Wandsworth Park’ and Stockwell North, as well as the peripheral ‘North Feltham & Hatton’ (Figure 6). All traces of positive correlation have disappeared, including the area just east of King’s Cross station that was present in the first 2 models. Global collinearity checks found distance to London’s center, nearby bus stops, and nearby rail stations to correlate with floor area at coefficients .099, -0.156, and -0.112, respectively. Table 2 summarises OLS coefficients against the corresponding GWR range for all nine predictors in this model, illustrating the extent to which each relationship varies spatially beyond what a single global estimate would suggest.




Figure 6: Map of final model’s GWR regression coefficients for floor area predicting ppsqm.

Discussion

The results suggest that there is strong spatial heterogeneity in the relationship between a property’s floor area and its ppsqm, with formal diagnostics confirming a declining but persistently significant spatial autocorrelation across models (Table 1; Moran’s I: 0.576 to 0.409, $p < 0.001$). This puts into question the accuracy and, more significantly, the fairness of any AVM which assumes a single global relationship rather than a range of them (Table 2). Each model’s GWR regression coefficient map clearly defined how this relationship varies throughout London

along with hotspot regions where penalties for smaller homes are particularly steep, which the final model found specifically to be the MSOAs of 'Putney Town & Wandsworth Park,' Stockwell North, and 'North Feltham & Hatton.'

The base GWR model identified a sparse scattering of areas where the size-ppsqm relationship is statistically significant. Positive regions where ppsqm grows with floor area emerged both centrally and peripherally, weaker (coefficients approximately 0.2) and larger in central London—particularly large just east of King's Cross station, approximately the Islington area—while smaller and individually strongest (coefficients approximately 0.3) at the outer MSOAs of Keston, Pinner Green, and 'Chislehurst West & Elmstead.' Negative regions where ppsqm grows as floor area shrinks emerged moderate (coefficients approximately -0.1 to -0.2) in outer regions of London and individually strongest (coefficients approximately -0.3) in the suburbs of Hounslow Heath, South Thamesmead, and 'West Thornton South & Ampere Way.' The nature of this dual-sided relationship reflects Hubbard et al.'s aggregate-level findings that ppsqm tends to rise as floor area moves away from 'average' in either direction. Further, notably negative relationships being found in outer to suburban areas corroborate their findings that smaller homes in outer London potentially offer worse value for money than those in central London, though they attribute this more to distance from central London amenities than to an underlying size-price relationship.

Introducing structural predictors extensively adjusted the model, revealing large expanses of statistically significant negative correlation, though a sizable void persisted in most of the eastern half. Two of the most negative coefficients moved further out to the peripheral MSOAs of Enfield Wash and 'North Feltham & Hatton,' though one moved closer to the central 'Old Street & St Luke's.' Global correlation diagnostics found flat status the predictor most strongly correlated with floor area at a coefficient of -0.584, followed by construction year at a coefficient of -0.236, offering these as strong factors in the citywide attenuation of floor area's coefficient. Most every positive relationship vanished, potentially suggesting that, for these regions, structural qualities strongly influence the floor area-ppsqm relationship; these predictors' absence in the base model may have made floor area's explanatory power artificially large/significant, which their inclusion then corrected. The emergence of vast, negative regions could be attributed to structural predictors improving the model's statistical power, allowing it to detect a floor area effect that may have existed statistically insignificantly and invisibly in the base model. The importance of structural predictors is corroborated by GWR metrics, with the fall of AICc from 1060.866 to -3428.070 (a delta far exceeding 10, the conventional threshold for a decisively improved model (Burnham and Anderson, 2004) and rise of adjusted R^2 (0.584 to 0.634) tangible evidence of their benefit. Interestingly, the only region with statistically significant positive coefficients that persisted was the King's Cross-Islington area, though it shrank in both size and relationship strength (coefficients roughly halved) and only vanished with the introduction of location predictors in the final model; however, no single locational

predictor map had a correspondingly strong coefficient in this area, suggesting this may be their combined effect, a result of a broader bandwidth (1050m to 1300m). Though the exact mechanism is unknown, GWR's interpretable structure means that, with more time, this could be tested through diagnostics or systematic comparisons not possible with a black box alternative.

Introducing locational predictors otherwise worsened the model. The fall of adjusted R^2 from 0.634 to 0.629 combined with AICc's rise from -3428.070 to -3260.873 suggest these predictors contributed little explanatory power to the model after structural predictors had been accounted for. These predictors' weaker correlation with floor area (coefficients of 0.099, -0.156, and -0.112 for distance from London's center, nearby bus stops, and nearby rail stations, respectively) agree with the meager attenuation brought to the floor coefficient map that was already revealed by the structural model, though they did resolve the persistent positive pocket by King's Cross that the structural model couldn't. While still far superior to the base model, the final model offered little to an understanding of spatial variation in the floor area-ppsqm relationship once structural factors, particularly flat status, had been considered. An AVM that ignores these attributes and their share in explaining London's spatial variation in the size-price relationship risks systematically misvaluing houses depending on their location.

In tandem, these findings demonstrate that the fairness of AVMs, and consequently the tax policy they've already been set in motion to inform, can be compromised by spatially-unaware approaches that fail to capture the relationship between different attributes and price in different areas. For London, GWR revealed substantial spatial heterogeneity that existing aggregate analysis couldn't and gestured towards peripheral suburbs as where a 'small home' premium may be strongest. Additional property factors, particularly structural attributes, explained much of this observed variation. These findings are limited by a lack of reliable and accessible floor area/price data, and the EPC and PPD data used is restricted in both coverage and accuracy; there is a clear need for open housing data. This study initially used Multiscale GWR (MGWR), which enables predictors to operate on their own unique bandwidths, but quickly opted for the fixed bandwidth of GWR for computational tractability reasons; further work could employ MGWR to potentially better understand which factors matter where. This work supports a larger initiative by University College London's Centre For Advanced Spatial Analysis—the Atlas of Small Homes—which aims to extend investigation into size-price relationships to core cities throughout England and Wales; this paper's findings offer a starting point to better understand what patterns exist in both London and beyond. This study's underlying pipeline—built on Python, PySAL, and Quarto—is intended for eventual open publication for independent verification and reuse by researchers and planning authorities. This transparency is salient for public-facing applications like major tax policy reform, where those affected must understand how outcomes are determined.

References

- Adam, S., Hodge, L., Phillips, D. and Xu, X. (2020). Revaluation and reform: Bringing council tax in England into the 21st century. [online] Institute for Fiscal Studies. Available at: <https://ifs.org.uk/publications/revaluation-and-reform-bringing-council-tax-england-21st-century>.
- Burnham, K.P. and Anderson, D.R. (2004). Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33(2), pp.261–304. doi:10.1177/0049124104268644.
- Chi, B., Dennett, A., Oléron-Evans, T., & Morphet, R. (2021). A new attribute-linked residential property price dataset for England and Wales, 2011–2019. UCL Open Environment, 2. <https://doi.org/10.14324/111.444/ucloe.000019>
- Feng, S.-T., Peng, C.-W., Yang, C.-H. and Chen, P.-W. (2021). Non-linear relationships between house size and price. *International Journal of Strategic Property Management*, 25(3), pp.240–253. doi:10.3846/ijspm.2021.14607.
- GOV.UK. (2025). Valuation Office Agency: Automated valuation model. [online] Available at: <https://www.gov.uk/algorithmic-transparency-records/valuation-office-agency-automated-valuation-model>.
- Housing Analysis team. (2025). Energy efficiency of housing in England and Wales. [online] Ons.gov.uk. Available at: <https://www.ons.gov.uk/peoplepopulationandcommunity/housing/articles/energyefficiencyofhousinginenglandandwales/2025>.
- Hubbard, P., Reades, J., Walter, H. and Preston, C. (2023). Shrinking homes? The geographies of small domestic properties in London, 2010–2021. *Environment and Planning B: Urban Analytics and City Science*, [online] 51(5). doi:10.1177/23998083231208732.
- Kiely, T.J. and Bastian, N. D (2019). The spatially-conscious machine learning model. arXiv. Available at: <https://arxiv.org/abs/1902.00562v1>.

Land Registry, H. (2014). HM land registry: Price paid data. [online] GOV.UK. Available at: <https://www.gov.uk/government/collections/price-paid-data>.

Murphy, L. and Snelling, C. (2019). A poor tax reforming council tax in London final report. [online] Institute for Public Policy Research. Available at: https://ippr-org.files.svdcdn.com/production/Downloads/1559046258_a-poor-tax-reforming-council-tax-in-london-may2019.pdf?dm=1778241484.

Oshan, T., Li, Z., Kang, W., Wolf, L.J. and Fotheringham, A.S. (2019). MGWR: A python implementation of multiscale geographically weighted regression for investigating process spatial heterogeneity and scale. *ISPRS International Journal of Geo-Information*, 8(6), p.269. doi:10.3390/ijgi8060269.

Özer, S. and Jacoby, S. (2022). Dwelling size and usability in London: A study of floor plan data using machine learning. *Building Research & Information*, [online] 50(6), pp.694–708. doi:10.1080/09613218.2022.2070452.

Smith, D. (2010). Valuing housing and green spaces: Understanding local amenities, the built environment and house prices in London. [online] London: Greater London Authority. Available at: https://www.london.gov.uk/sites/default/files/gla_migrate_files_destination/GLAE-wp-42.pdf.

Tapia, J., Chavez-Garzon, N., Raúl Pezoa, Suarez-Aldunate, P. and Pilleux, M. (2025). Comparing automated valuation models for real estate assessment in the Santiago metropolitan region: A study on machine learning algorithms and hedonic pricing with spatial adjustments. *PLoS ONE*, 20(3), pp.e0318701-e0318701. doi:10.1371/journal.pone.0318701.

Valier, A. (2020). Who performs better? AVMs vs hedonic models. *Journal of Property Investment & Finance*, [online] 38(3), pp.213–225. doi:10.1108/jpif-12-2019-0157.

Welsh Government (2025). Find out about how we are reforming council tax | gov.wales. [online] GOV.WALES. Available at: <https://www.gov.wales/find-out-about-how-we-are-reforming-council-tax>.

Appendix A: Additional Coefficient Maps, Structural Model




Figure A.1: Local GWR coefficient for energy efficiency, structural model. Grey points not statistically significant at the 95% confidence level.




Figure A.2: Local GWR coefficient for construction year, structural model. Grey points not statistically significant at the 95% confidence level.




Figure A.3: Local GWR coefficient for attachment style, structural model. Grey points not statistically significant at the 95% confidence level.




Figure A.4: Local GWR coefficient for flat status, structural model. Grey points not statistically significant at the 95% confidence level.




Figure A.5: Local GWR coefficient for new build status, structural model. Grey points not statistically significant at the 95% confidence level.

Appendix B: Additional Coefficient Maps, Final Model




Figure B.1: Local GWR coefficient for distance from London's center, final model. Grey points not statistically significant at the 95% confidence level.




Figure B.2: Local GWR coefficient for bus stops within 500 meters, final model. Grey points not statistically significant at the 95% confidence level.




Figure B.3: Local GWR coefficient for rail nodes within 500 meters, final model. Grey points not statistically significant at the 95% confidence level.