

A Deep Neural Network for Mass Property Valuation Using Spatial Features

(Selected Paper in the 8th ISPRS Geospatial Conference 2025, University of Tehran, Iran)

Ali Sadeghian¹ , Farid Cheraghi^{2✉}, and Hamid Mehrabi³ 

1. Department of Geomatics Engineering, Faculty of Civil and Transportation Engineering, University of Isfahan, Iran. E-mail: a.sadeghian@trn.ui.ac.ir
2. Corresponding author, Department of Geomatics Engineering, Faculty of Civil and Transportation Engineering, University of Isfahan, Iran. E-mail: f.cheraghi@cet.ui.ac.ir
3. Department of Geomatics Engineering, Faculty of Civil and Transportation Engineering, University of Isfahan, Iran. E-mail: h.mehrabi@eng.ui.ac.ir

Article Info

Article type:

Research Article

Article history:

Received 2026-01-17

Received in revised 2026-02-01

Accepted 2026-02-02

Available online 2026-05-12

Keywords:

Mass Property Valuation,
Deep Neural Network,
Multi-layer Perceptron,
Geospatial Information System,
Wales and England

ABSTRACT

This study developed an artificial neural network (ANN) model for large-scale property valuation in England and Wales, leveraging comprehensive datasets that integrate physical, environmental, spatial, and temporal features of properties. The primary goal is to enhance the accuracy of property price prediction by capturing complex nonlinear relationships among diverse property attributes.

The research utilized a Multi-Layer Perceptron (MLP) architecture with a systematic hyperparameter search to determine the optimal 10 layers configuration (512-256-128-64-32-16-8-4-2-1 nodes). Data preprocessing steps included feature encoding, standardization, and outlier removal. Spatial features, such as distances from urban infrastructure, were normalized using GIS tools. The model was trained using the Adam optimizer with mean squared error (MSE) loss and incorporated early stopping to prevent overfitting. Feature ablation studies and learning curve analyses were conducted to validate the optimal feature set (59 features) and dataset size (5.5million samples).

The final model achieved a coefficient of determination (R^2) of 0.90, with a mean absolute percentage error (MAPE) of 15.28%, demonstrating superior performance compared to existing methods like Random Forest. The root mean square error (RMSE) was approximately 45,848 units (£), and the mean absolute error (MAE) was 29,368 units (£). These metrics highlight the model's ability to capture intricate patterns in property valuation.

The study underscored the effectiveness of integrating deep learning techniques with spatially enhanced data for accurate mass property valuation. The MLP architecture proved particularly suited for handling heterogeneous and high-dimensional datasets, achieving state-of-the-art performance in urban environments characterized by complex spatial-economic relationships. This approach offers a robust alternative to traditional methods, enabling more precise and reliable property price predictions.

Cite this article: Sadeghian, A., Cheraghi, F., Mehrabi, H. (2025). A Deep Neural Network for Mass Property Valuation Using Spatial Features, *Earth Observation and Geomatics Engineering*, Volume 9, Issue 1, Pages 104-116. <http://doi.org/10.22059/eoge.2026.409770.1199>



© The Author(s).

DOI: <http://doi.org/10.22059/eoge.2026.409770.1199>

Publisher: University of Tehran.

1. Introduction

Property valuation is a process in which the value of a property — whether residential, commercial, or industrial — is determined through precise statistical analyses. This valuation is typically conducted by professional assessors for purposes such as buying and selling, securing bank loans, and investment decisions. Property valuation is based on unique characteristics of the property, its geographical location, physical attributes, market conditions, and various other factors (Blackledge, 2016).

Mass property valuation, on the other hand, refers to the process in which the values of a group of properties are estimated simultaneously using statistical methods and algorithms (International Association of Assessing Officers, 2013). This approach is particularly useful in areas with a large number of properties and is commonly used for tax assessment, urban planning, and other governmental or public purposes. Unlike individual property valuation, which is usually performed on a case-by-case basis, mass valuation requires extensive analysis and therefore involves larger datasets and more advanced statistical and computational tools (Blackledge, 2016).

Traditional property valuation systems largely rely on personal experiences of appraisers, subjective opinions, and unscientific comparisons. In these methods, valuations are heavily influenced by individual perspectives, which can lead to unrealistic estimates and price discrepancies (Droj et al., 2024). This highlights the need for developing methodologies that enable transparent, accurate, and impartial valuations independent of human bias. In this context, Automated Valuation Models (AVMs) play a key role in improving the accuracy and efficiency of the property valuation process. An AVM is a data-driven system that uses statistical and artificial intelligence techniques to estimate property prices (International Association of Assessing Officers, 2018).

Spatial data also plays a critical role in identifying the factors that influence land value. By integrating this information into valuation models, the accuracy of property value predictions improves significantly. The use of spatial analysis enables systematic and fair property valuation (Demetriou, 2017). Locational factors exert a profound influence on property worth. This need is powerfully addressed by emerging GeoAI-driven approaches, which unify international land administration standards—such as the LADM Valuation Information Model—with advanced machine learning and rich geospatial feature engineering to deliver highly accurate, transparent, and scalable valuation systems capable of supporting complex urban governance and market functions (Mete, 2025).

In this study, publicly available government datasets related to real estate transactions and Energy Performance Certificates (EPC) in England and Wales were utilized (Chi et al., 2021). This dataset includes physical building

characteristics and some energy consumption indicators, which have previously been used by researchers for large-scale property value modeling. For instance, a previous study employed spatial variables extracted from GIS along with machine learning methods such as Random Forest, XGBoost, CatBoost, and LightGBM to predict property prices (Mete & Yomralioglu, 2023). The results showed that the Random Forest model outperformed other algorithms, achieving an R^2 score of 0.85, a Mean Absolute Percentage Error (MAPE) of 19.4%, and an RMSE of 101,847 units.

Deep learning, a subfield of machine learning, has made remarkable progress in recent years by leveraging automated feature learning and data structure modeling. Factors such as the availability of large-scale datasets, advancements in processing hardware (e.g., GPUs), improvements in learning algorithms, and the development of software libraries have led to increased interest in deep neural networks in fields such as real estate valuation, urban pattern analysis, and estimation of economic variables associated with geographic spaces. These networks can better model hidden relationships between features compared to traditional methods by layer-wise processing of spatial, economic, and image-based data (Chollet, 2021).

In a study by Mimis et al. (2013), the application of Artificial Neural Networks (ANNs) was investigated for residential property valuation in Athens. The research developed a nonlinear model for predicting property prices by integrating physical data (such as floor area, building age, and floor level), environmental characteristics (including neighborhood average income and access to metro stations), and property geographical locations (using Geographic Information Systems – GIS). The dataset consisted of 3,150 residential units across a 300 km² area. Results indicated that the ANN model outperformed a Spatial Autoregressive (SAR) model, achieving an R^2 of 0.86 and a Mean Absolute Error (MAE) of €16,625, compared to an R^2 of 0.76 and an MAE of €17,384 for SAR. Moreover, while nonlinear analyses confirmed an inverse relationship between property age and price, they also revealed a nonlinear and increasing relationship between property size and price—indicating that as size increases, price rises at an accelerating rate but not linearly. This study emphasized the capability of ANNs in modeling spatial price patterns and the importance of integrating spatial data in mass property valuation.

In a study by Demetriou (2017), an Artificial Neural Network (ANN)-based Automated Valuation Model (AVM) was developed for mass land valuation in the context of land consolidation in Cyprus. The research integrated 14 spatially derived land valuation factors—spanning physical (e.g., slope, size, aspect), legal (e.g., irrigation rights), locational (e.g., distance to residential zones, road access), and economic attributes—within a GIS framework. Using a stratified sample of only 15% (73 out of 488) of land parcel values provided by a governmental Land Valuation Committee, the model was trained via a Multilayer Perceptron (MLP) architecture with one hidden layer (3

neurons) and hyperbolic tangent activation. The ANN achieved a coefficient of determination (R^2) of approximately 0.79 and a Mean Absolute Percentage Error (MAPE) of 11.7%, outperforming a linear regression baseline and matching the performance of a nonlinear regression model applied to the same dataset. Notably, quality assurance metrics—including a Coefficient of Dispersion (COD) of 11.38 and a Price-Related Differential (PRD) of 1.03—confirmed the model’s compliance with international mass appraisal standards ([International Association of Assessing Officers, 2013](#)). This study demonstrated that even with limited labeled data, a spatially informed ANN can deliver reliable, transparent, and efficient land valuations, reducing manual effort by up to 80% while enhancing consistency and fairness in rural land reallocation processes.

In a comprehensive study a self-adaptive neural network model was developed for mass appraisal and scenario-based forecasting of urban real estate values across 10 major Russian cities (including Moscow, Saint Petersburg, and Yekaterinburg) over the period 2006–2020 ([Yasnitsky et al., 2021](#)). The model integrated 15 input features, encompassing physical attributes (e.g., wall type, floor area, number of rooms), geographical coordinates, locational prestige (derived from price heatmaps), temporal indicators (year and season), and macroeconomic variables (e.g., USD/RUB exchange rate, Brent oil price, national GDP, housing construction volume, and mortgage lending). Trained on approximately 300,000 transaction records, the model employed a three-layer MLP architecture (input–3–2–output) with sigmoid activation in hidden layers and Resilient Backpropagation (RPROP) for optimization. The final model achieved a mean relative error of 6.2% on the test set and demonstrated strong generalizability across diverse urban contexts. Notably, the study enabled scenario-based forecasting—such as simulating the impact of mortgage policy changes or GDP fluctuations on property prices—highlighting the model’s capacity to capture spatio-temporal dynamics and support strategic decision-making in volatile economic environments.

Another study by [Abidoje and Chan \(2018\)](#) conducted in Nigeria compared the predictive performance of two approaches — the Hedonic Pricing Model (HPM) and Artificial Neural Network (ANN) — for residential property valuation in the Lagos real estate market. Using 321 transaction records and structural-spatial variables, the study found that ANN achieved a much lower Mean Absolute Percentage Error (MAPE) of 15.94% compared to HPM’s 38.23%. The ANN also demonstrated significant superiority in other metrics such as RMSE and MAE. These findings suggest that nonlinear methods like ANN can provide higher accuracy in mass property valuation, especially in complex markets. The study underscored the importance of adopting artificial intelligence technologies to enhance real estate valuation processes.

In a study conducted in Konya, Turkey, a mass valuation model was implemented using ANN and compared with another model employing Multiple Regression Analysis (MRA) ([Yalpir et al., 2014](#)). Both methods were integrated with spatial analysis data from the ArcGIS software and applied to 300 residential properties. Property value maps were generated for both algorithms within ArcGIS. Statistical comparison of the results showed that ANN provided more suitable outcomes for mass property valuation, and both methods yielded more accurate results when enhanced with spatial analysis data compared to when used without them.

In a recent study, Geographically, Temporally, and Characteristically Weighted Generalized Regression Neural Network (GTCW-GRNN) was proposed for housing price modelling in District 5 of Tehran, Iran. The model integrates 17 input features—including physical attributes (e.g., floor area, building age), locational proximity (e.g., distance to hospitals, schools, and fault lines), and macroeconomic indicators (notably the USD exchange rate)—to capture spatiotemporal and attribute-based heterogeneity. Trained on 7,161 validated transactions from 2017–2019, the GTCW-GRNN outperformed conventional methods such as GTWR, achieving the highest adjusted R^2 of 0.813 among all tested configurations. The study further demonstrated—through a rigorous feature ablation analysis using McNemar’s test—that the exchange rate and residential unit area were the most influential predictors, highlighting the critical role of economic context in emerging real estate markets ([Zali et al., 2025](#)).

Although previous studies have examined the application of artificial neural networks in property valuation, most of them were based on limited datasets (fewer than 5,000 samples) and fewer features (mostly less than 20 variables). Additionally, the neural network architectures used in those studies were often simpler, employing fewer hidden layers. In contrast, the present study utilizes an enriched dataset consisting of over 5 million records and 57 features — including physical, environmental, spatial, and temporal variables — enabling the extraction of more complex patterns. Furthermore, the deeper architecture of the proposed model (an MLP with 4 hidden layers and a neuron configuration of 256-128-64-32-1) offers a higher capacity for learning nonlinear relationships among diverse features. These advancements, combined with sophisticated spatial analysis integration, clearly distinguish the accuracy of the current method from previous studies. The findings not only confirm the potential of neural networks in property valuation but also represent a step forward from prior works by introducing a scalable model for complex urban environments.

The remainder of this paper is structured as follows: Section 2 describes the data preparation pipeline, the theoretical methodology, and the implementation details of the proposed deep neural network. Section 3 presents the experimental results, including learning curves and

performance metrics. Finally, Section 4 discusses the findings and concludes with implications for mass property valuation in urban environments.

2. Data and Methods

2.1. Data Preparation

The datasets used in this study were extracted from publicly available government data in England and Wales (Chi et al., 2021). This dataset includes real estate transaction records — known as Price Paid Data (PPD) — and Energy Performance Certificates (EPC), which cover residential property transactions recorded between 2011 and 2019 (Figure 1). These data are highly diverse in terms of physical, environmental, spatial, and temporal characteristics of properties (Table 2).

Initially, the real estate transaction data and energy performance certificates were integrated into a unified dataset to enable comprehensive analysis of various property features. In the next stage, the data underwent extensive preprocessing steps. Null values and zero entries across columns were identified and removed. Categorical variables with non-numeric values (e.g., Basement, 1st, Ground from the Floor Level column) were encoded into numerical representations. Additionally, variables such as date of transfer were processed to extract only the year of the transaction and converted into numeric format. As part of the outlier handling procedure, extreme values were identified and removed for price field using the 3σ standard deviation method, ensuring that only reliable and representative data points were retained for model training.

To eliminate the confounding effect of general inflation on property price trends and ensure temporal comparability of transaction values across the 2011–2019 period, all nominal prices were adjusted to real prices expressed in 2011 purchasing power. This deflation process was performed using the Consumer Price Index (CPI) published by the Office for National Statistics (ONS), with 2011 as the base year (Table 1). The adjustment followed the standard formula:

$$P_{real} = P_{nominal} \times \frac{CPI_{base}}{CPI_{current}}$$

where $P_{nominal}$ denotes the nominal transaction price in current year, and P_{real} represents the inflation-adjusted price in pounds sterling. The resulting variable—`deinflate_price`—served as the target variable for all subsequent modeling stages. This approach isolates genuine property-specific value determinants from macroeconomic price drift, thereby strengthening the validity of spatial and feature-based analyses.

Table 1: CPIs by year

Year	CPI	Adjustment Factor
2011	93	100%
2012	96	96.88%
2013	98.2	94.70%
2014	99.9	93.09%
2015	100	93.00%
2016	100.9	92.17%
2017	103.5	89.85%
2018	106.2	87.57%
2019	108.2	85.95%

Next, spatial information embedded in the dataset was utilized to derive key features such as distances from urban infrastructures (e.g., train stations, hospitals, shopping centers, universities, rivers, and seas), visibility, topography, and proximity indices. These indicators were included in the model as influential factors in determining property value. Given the critical role of geographical location in property valuation, the longitude and latitude coordinates of each property were also incorporated into the model as two independent features. Categorical variables such as property type, new or old construction status (old new), and ownership type (duration) were also encoded using appropriate techniques.

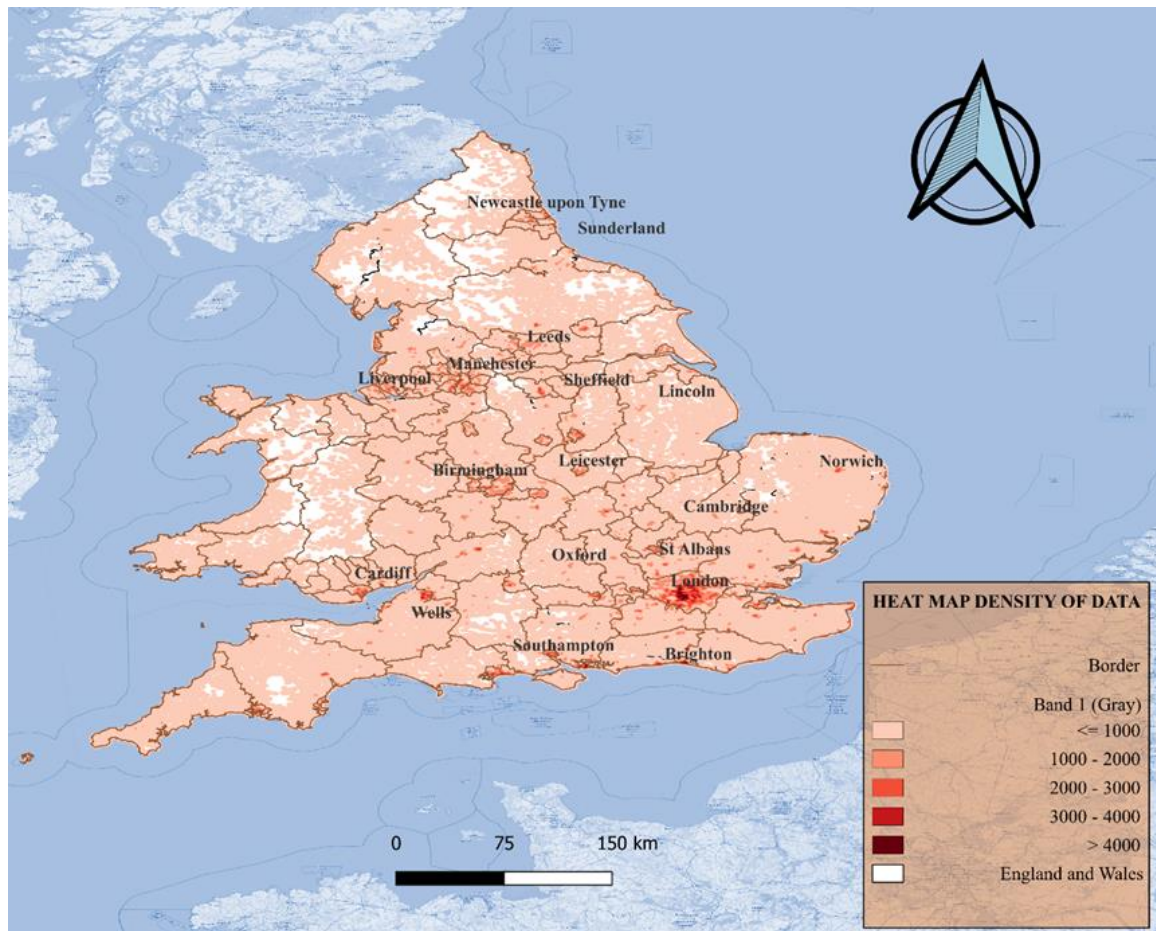


Figure 1: Transaction density map

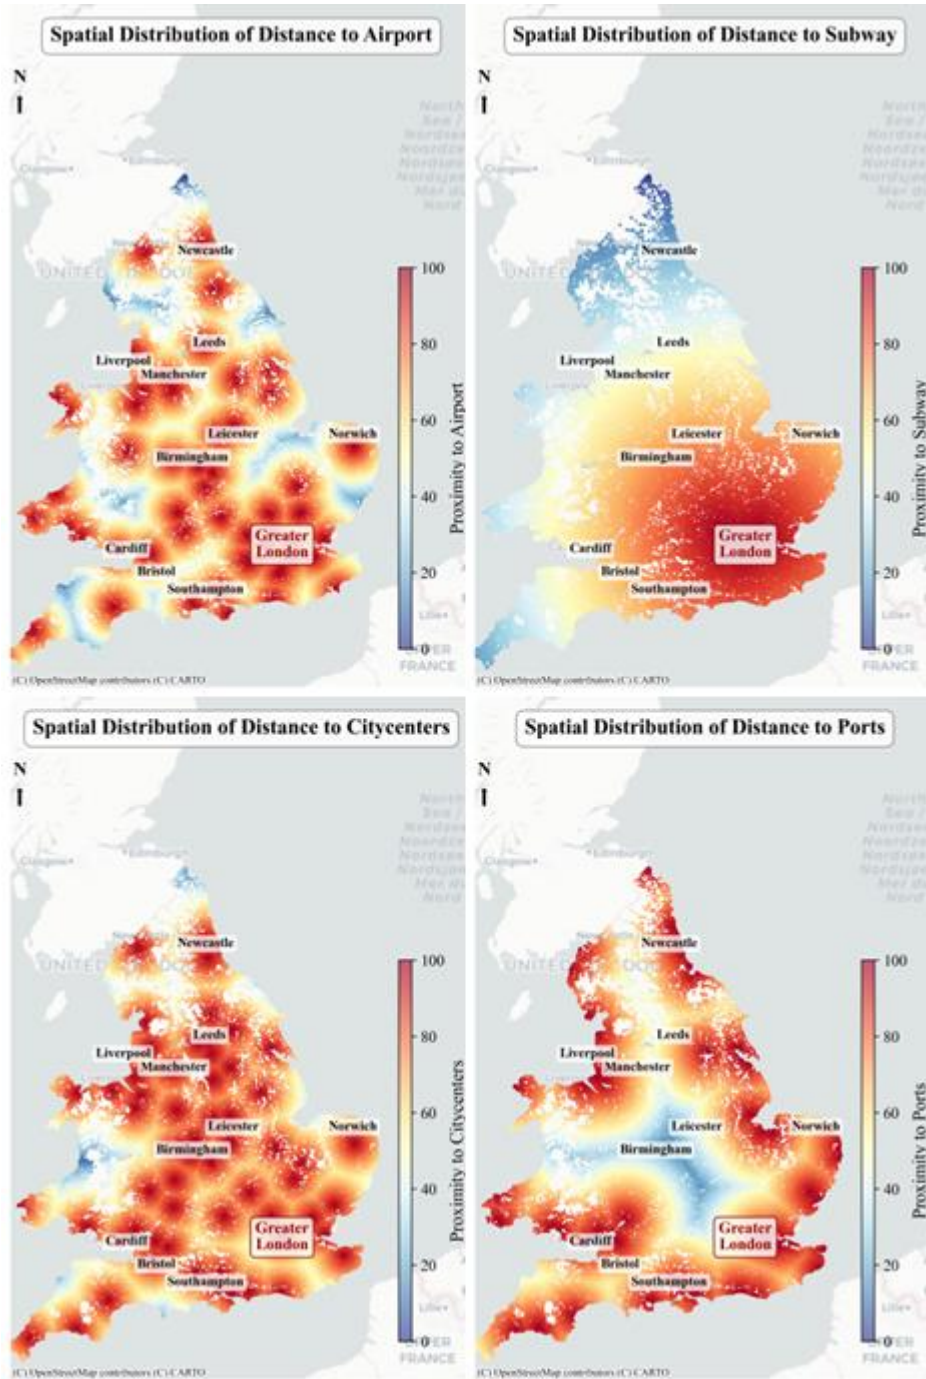


Figure 2: Proximity to (A) airport, (B) subway, (C) city centers, (D) ports normalized between 0 and 100.

During the data preparation phase, variables representing distances from urban infrastructures — commonly referred to as proximity features (e.g., subway stations, bus stops, shopping centers, etc.) — were subjected to a standardization process to ensure coherent and comparable integration with other features in the model. Initially computed as Euclidean distances in meters, these variables

were transformed into a normalized scale ranging from 0 to 100 to enhance interpretability and uniformity across scales. In this approach, complete proximity to an infrastructure (the best possible condition) was assigned a score of 100, while maximum distance (the worst condition) was assigned a score of 0. This normalization was performed using spatial

analysis tools within GIS, supported by OpenStreetMap and Ordnance Survey datasets (Table 3; Figure 2).

All features were standardised using the StandardScaler class from the Scikit-Learn library, resulting in zero mean and unit standard deviation. This step was essential due to the significant differences in scale among the features (e.g., price, area, distance, energy consumption). Ultimately, the final dataset consisted of 59 numerical features and approximately 5.5 million samples, ready for input into the neural network model.

Table 2: Classification of Features

Feature name	Variable type
Total Floor Area, Number Rooms, Environment Impact Current, Heating Cost Current, Floor Level, Extension Count, Current Energy Efficiency, Energy Consumption Current, Co2 Emissions Current, Lighting Cost Current, Hot Water Cost Current, Detached, Flats, Semi-Detached, Terraced, old property, new property, Wind Turbine Count	Physical
Airport, Main Road, Local Road, Coach Station, Vehicular Ferry, Fire Station, Hospital, Motorway Junction, Mall, Hospital, Park, Police, Primary Education, Art, Bus, City Center, Further Education, Library, Museum, Passenger Ferry, Post, Sports, Subway, Tourist, Railway, Tram, Worship, River, Sea, Secondary Education, Slope, University, Grass, Distance to London, Non State Primary, Non State Secondary, Other Health Center	Environmental
Free-Hold, Lease-Hold, Date Of Transfer	Legal
Longitude, Latitude	Coordinates

Table 3: Summary statistics of features

Feature name	Mean	Std	Min	Max
Deinflate Price	220670	146260	6296	1058000
Year	2015	2	2011	2019

Feature name	Mean	Std	Min	Max
Total Floor Area	91.68	39.51	9.7	937
Number Rooms	4.6	1.57	1	20
Environment Impact Current	56.83	13.64	1	329
Energy Consumption Current	280.85	117.09	-1365	13230
Co2 Emissions Current	4.58	2.62	-25.4	280
Lighting Cost Current	82.32	41.27	-369	56351
Heating Cost Current	745.69	448.25	-344	42928
Hot Water Cost Current	144.22	74.92	-2	1966
Extension Count	0.51	0.75	0	4
Wind Turbine Count	0.0003	0.0199	0	4
Airport	77	16	0	100
Art	85	12	0	100
Bus Station	92	7	0	100
Coach Station	88	11	0	100
Fire Station	89	8	0	100
Hospital	90	8	0	100
Library	92	8	0	100
Museum	86	11	0	100
Police Station	85	13	0	100
Post Office	94	5	0	100
Further Education	94	6	0	100
Universities	90	9	0	100
Sport Center	90	9	0	100
Worship	94	4	0	100
Primary	96	3	0	100
Secondary	95	6	0	100
Ports	70	22	0	100
Main Road	95	6	0	100
Local Road	98	3	0	100
Passenger Ferry	72	20	0	100

Feature name	Mean	Std	Min	Max
Vehicular Ferry	57	29	0	100
Motorway Junction	93	10	0	100
Railway	96	5	0	100
Subway	70	22	0	100
Tram	87	14	0	100
River	92	7	0	100
Tourist	75	17	0	100
Non State Primary	93	8	0	100
Non State Secondary	91	9	0	100
Other Health Center	96	4	0	100
Sea	76	24	0	100
Mall	91	9	0	100
Grass	96	4	0	100
City center	82	12	0	100
Distance to London	66	22	0	100
Slope	68	16	0	89
Aspect	179	107	0	360
Floor Level Number	0.2	0.8	-1	21
Longitude	-1.31560	1.327442	-5.70597	1.758335
Latitude	52.30293	1.131112	49.96456	55.78708
Detached	0.2332	0.4228	0	1
Flats	0.16	0.36	0	1
Semi-Detached	0.3	0.46	0	1
Terraced	0.31	0.46	0	1
Old Property	0.99	0.07	0	1
New Property	0.005	0.07	0	1
Freehold	0.80	0.40	0	1
Leasehold	0.21	0.40	0	1

2.2. Methodology

As discussed in the introduction, although several machine learning methods — such as Random Forest and hybrid models based on spatial data analysis using GIS software — demonstrate high capability in estimating property prices with acceptable accuracy, Artificial Neural

Networks (ANNs) were identified as a more suitable approach for this study. This superiority is particularly evident when dealing with spatially derived data and diverse physical and environmental features, as ANNs can extract hidden patterns within spatiotemporal interactions with greater flexibility. Accordingly, in this study, a Multi-Layer Perceptron (MLP) architecture was adopted as the primary modeling framework. Leveraging the capabilities of deep learning, this approach enhances predictive accuracy compared to conventional methods. In addition to overcoming limitations of traditional techniques — such as reliance on linear assumptions in regression models — this method also enables optimal integration of both spatial and non-spatial data.

The MLP consists of fully connected (dense) layers, where each neuron in a layer receives input from all neurons in the preceding layer. This structure enables the model to learn intricate, high-order interactions among physical, environmental, spatial, and temporal property attributes. The Rectified Linear Unit (ReLU) activation function defined as $f(x)=max(0, x)$, was chosen for hidden layers to mitigate the vanishing gradient problem and accommodate the right-skewed distribution of property prices (Figure 3). A linear activation was used in the output layer to support unbounded regression.




Figure 3. ReLU function

Model training was guided by the Mean Squared Error (MSE) loss function, which penalizes large prediction errors more heavily—a critical property in valuation tasks where financial consequences of outliers are significant. The Adam optimizer was selected for its adaptive learning rate mechanism, which combines momentum and RMSprop principles to stabilize convergence across heterogeneous feature scales and sparse data patterns.

To prevent overfitting and ensure generalizability, early stopping was applied based on validation set performance. Model evaluation was conducted using four complementary metrics: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). These metrics collectively assess prediction accuracy, error distribution,

relative deviation, and explained variance, aligning with international mass appraisal standards ([International Association of Assessing Officers, 2013](#)).

Root Mean Square Error (RMSE): One of the most commonly used metrics for evaluating regression models. Due to its sensitivity to large errors, RMSE serves as a robust indicator for real estate valuation tasks, where large prediction errors can lead to significant financial consequences.

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$$

where y_i = actual value of i-th sample
 \hat{y}_i = predicted value of the i-th sample
 n = total number of samples

Mean Absolute Error (MAE): Measures the average absolute difference between predicted and actual values.

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

Mean Absolute Percentage Error (MAPE): Measures error as a percentage of the actual values.

$$MAPE = \frac{100 \times \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right|}{n}$$

Coefficient of Determination (R^2): Measures the proportion of variance explained by the model. R^2 ranges between 0 and 1, with values closer to 1 indicating better model fit.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

The theoretical design prioritized two criteria:

- **Expressivity:** the ability to model complex nonlinear dependencies among diverse property features;
- **Generalizability:** robustness against overfitting when trained on large-scale, real-world datasets.

Crucially, the selection of input features was not arbitrary. Each feature was included based on established theoretical or empirical evidence of its influence on property value—such as proximity to public transport (Mete et al., 2021, p. 20), energy performance ([Fuerst et al., 2016](#)), or topographic exposure ([Demetriou, 2017](#)). The inclusion of spatially derived variables ensures that the model captures location-specific value premiums, addressing a well-documented gap in conventional AVMs.

To ensure representative distribution of spatial and temporal patterns across all subsets and mitigate potential spatial autocorrelation bias, a stratified sampling strategy was employed during dataset partitioning. Specifically, stratification was performed based on the administrative

boundaries defined by the local_authority field—representing local government districts across England and Wales—combined with the year of transaction (2011–2019). This dual stratification approach guarantees proportional representation of each spatiotemporal unit across the training (70%), validation (15%), and test (15%) subsets. Consequently, spatial clustering and temporal drift effects are minimized, ensuring that model evaluation reflects performance across the full spectrum of geographic and temporal variability present in the dataset.

To further understand the contribution of individual features to the model's predictive capability, permutation importance analysis was conducted on the test set. This technique quantifies feature importance by measuring the increase in RMSE when feature values are randomly permuted, thereby disrupting the relationship between that feature and the target variable.

2.3. Implementation

The final MLP architecture was selected through a systematic hyperparameter search over a range of candidate configurations. Specifically, we evaluated architectures with 2 to 6 hidden layers, and for each layer, the number of neurons was constrained to powers of two (i.e., 2, 4, 8, 16, 32, 64, 128, 256) to align with computational efficiency principles and common deep learning practice. All models were trained under identical conditions: Adam optimizer, MSE loss, batch size = 256, and early stopping (patience = 10 epochs) monitored on the validation set.

Among the tested configurations including 64-32-16-8-1, 64-32-16-1, 128-64-32-16-8-4-1, 128-64-32-16-1, 256-128-64-32-16-8-1, 512-256-128-64-1, etc. the architecture 256-128-64-32-1 consistently achieved the lowest validation RMSE while maintaining a minimal gap between training and validation errors, indicating strong generalization. This structure contains 205,481 trainable parameters and implements a gradual information bottleneck that compresses high-dimensional inputs while preserving valuation-relevant signals.

The model was trained on the full dataset of 5.6 million samples and 57 meaningful features, all of which were selected based on domain knowledge and statistical significance. To justify this feature count, we conducted a feature ablation study: models were trained incrementally with 10, 20, 30, 40, 50, and 57 features, ordered by descending Pearson correlation with price. As shown in Figure 4, RMSE decreased monotonically as more features were added, with the steepest improvement occurring

between 10 and 40 features. Beyond 50 features, gains plateaued, confirming that the full set of 57 features captures nearly all predictive signal available in the data.

Similarly, to justify the use of the full dataset, we trained identical models on subsamples ranging from 100 to 3 million records. As illustrated in Figure 5, RMSE declined steadily with increasing sample size, demonstrating the model’s capacity to leverage big data and learn increasingly refined spatial-economic patterns. This behaviour confirms that the model is data-hungry but not overfitting, and that the scale of our dataset is a key contributor to performance.

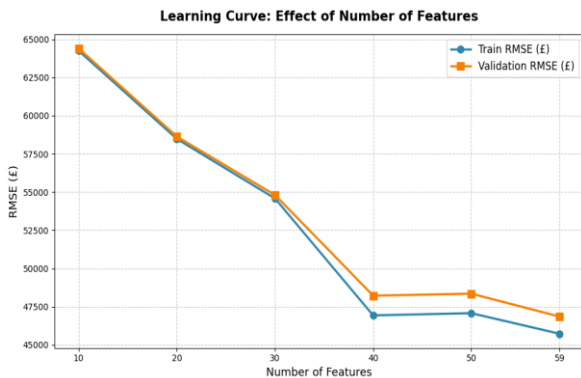


Figure 4: RMSE Using Different Number of Features

The model was implemented in Python using TensorFlow/Keras, with Scikit-Learn for preprocessing. Training was performed on GPU-accelerated hardware to handle the scale of the dataset. The same hyperparameter configuration was maintained across all experimental runs to ensure fair comparison during architecture selection and feature ablation studies.

3. Results

Learning curve analysis revealed that model performance is directly influenced by the number of features and the size of the training dataset. The learning curve based on feature count shows that increasing the number of features from 10 to 57 led to a substantial decrease in RMSE for both the training and validation sets. This highlights the important role of adding meaningful and diverse features — especially spatial and environmental ones — in improving the accuracy of property price predictions. However, beyond approximately 50 features, the rate of error reduction slowed significantly, indicating the existence of an optimal feature threshold beyond which additional features contribute little to model accuracy




Figure 5: RMSE using different numbers of training samples

Similarly, the learning curve based on sample size (Figure 6) showed that increasing the training set size from 100 to 3 million samples resulted in a continuous decline in RMSE for both the training and validation sets. This trend demonstrates that the neural network under consideration has a strong capacity for leveraging big data, and the larger the training dataset, the more accurately the model can learn complex non-linear relationships between physical, environmental, spatial, and temporal property features. This behavior indicates minimal overfitting and confirms that the model generalizes well.

Based on these findings, the final model was trained using all 57 features and over 5.6 million samples. The data were divided into three subsets: 70% for training, 15% for validation, and 15% for testing. The model was trained for a maximum of 150 epochs; however, the training process was stopped early at epoch 117 using the Early Stopping technique to prevent overfitting. Throughout the training process, all training samples passed through the network once per epoch. The number of epochs was selected based on early stopping criteria monitored via validation set error, which serves as an effective strategy for optimizing the learning process and preventing overfitting. The batch size was set to 256, meaning that gradients were computed based on 256 samples at each weight update step. This value was chosen considering the balance between learning speed and gradient estimation quality.

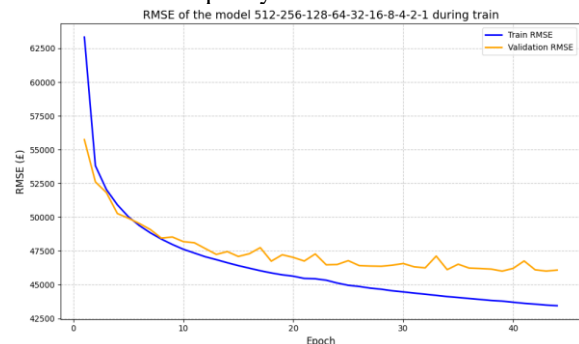


Figure 6: RMSE values across training epochs for training and validation sets

All error metrics reported below are expressed in constant 2011 pounds sterling, reflecting the deflated target variable used during model training. This standardization enables direct comparison of prediction accuracy across the entire temporal span of the dataset without distortion from general inflation trends. Model evaluation results indicated that the coefficient of determination (R^2) reached 0.90. This value suggests that the model explains 90% of the variance in observed property prices. Additionally, the Mean Absolute Percentage Error (MAPE) was approximately 15.2%, and the Root Mean Square Error (RMSE) was around 45,848 units (Pound sterling). These metrics, along with an MAE of 29,368 units (Pound sterling), collectively provide a comprehensive picture of the model's adequate predictive accuracy.

To further evaluate the model's spatial predictive capability, Figure 6 presents the spatial distribution of absolute prediction errors across the test set. The visualization reveals a predominantly uniform error pattern with lower errors (blue shades) prevalent across most regions and occasional higher errors (red shades) concentrated in specific urban centers.

Quantitative assessment of spatial autocorrelation in residuals was conducted using Moran's I statistical test, which yielded a Moran's I value of 0.1, p-value of 0.001 and z-score of 205.94. This low Moran's I value indicates minimal spatial clustering of prediction errors, confirming that the model effectively captures location-specific value determinants through the integration of spatial features. The absence of significant spatial patterns in residuals validates the model's capacity to account for complex spatial-economic relationships inherent in property valuation, thereby supporting the robustness of the DNN architecture in handling heterogeneous urban environments. This finding aligns with our theoretical framework that emphasizes the critical role of spatial feature integration in mass property valuation systems.

This level of accuracy, particularly in an urban environment characterized by complex nonlinear relationships among influential variables, reflects the notable performance of the developed model. The relatively low error values compared to actual property prices demonstrate the model's capability to manage multidimensional data and learn hidden patterns within physical, spatial, and energy-related property relationships.

Spatial Distribution of Prediction Errors (Test Set) - DNN Model

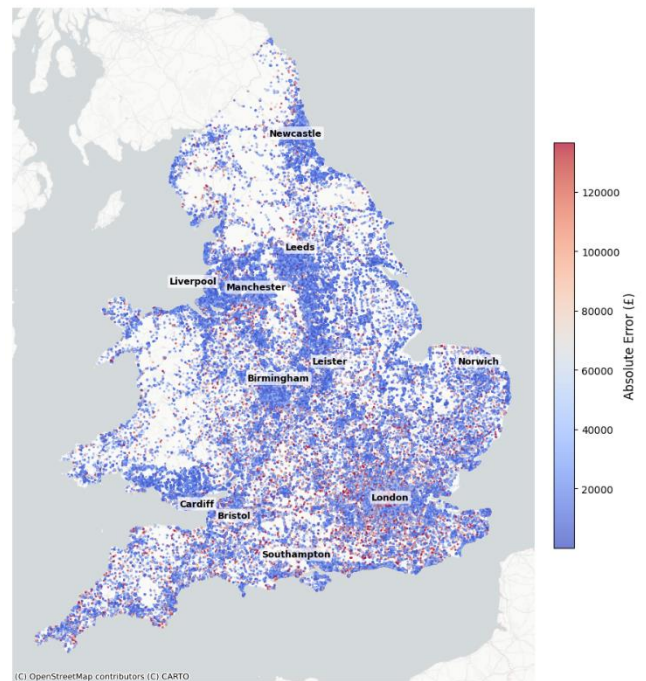


Figure 7: Spatial distribution of absolute prediction errors across the test set for the DNN model. Lower errors (blue) are prevalent across most regions, while higher errors (red) are concentrated in specific urban centers.

Finally, as shown in Figure 8, the analysis revealed that proximity to London's underground infrastructure (Subway Scaled) was the most influential feature, with a substantial RMSE increase of 57,887 units when permuted. This finding underscores the critical role of accessibility to London's public transportation network in determining property values across England and Wales. The second most important feature was distance to London (greatlondon), with a 38,780-unit RMSE increase, confirming the strong spatial gradient in property values radiating from the capital city. Other highly influential features included environmental factors (CO_2 emissions, environmental impact), proximity to major infrastructure (ports, airports, ferry terminals), and geographic coordinates (latitude, longitude), which collectively demonstrate the model's ability to capture both local and regional spatial patterns.




Figure 8: 20 features with the most impact

These results align with theoretical expectations in urban economics and validate our approach of incorporating comprehensive spatial and environmental features, while also providing actionable insights for property valuation practitioners regarding the relative importance of different location-based factors; "Reason Codes" is often a legal requirement for government tax assessment.

4. Discussion and Conclusion

In this study, an Artificial Neural Network (ANN) model was developed for large-scale mass property valuation. The proposed model leverages enriched datasets containing physical, environmental, spatial, and temporal features of properties in England and Wales, enabling the identification of complex nonlinear patterns among various property characteristics and providing high accuracy in price prediction.

Model performance evaluation based on standard metrics — including RMSE, MAE, MAPE, and R^2 — demonstrated that the designed model achieves competitive accuracy compared to existing approaches. A coefficient of determination (R^2) of approximately 0.89, a Mean Absolute Percentage Error (MAPE) of around 15.7%, a Mean Absolute Error (MAE) of 38,638 units (Pound sterling), and an RMSE of approximately 86,638 unit (Pound sterling) reflect the model's acceptable predictive capability. These results significantly outperform those reported by Mete and Yomralioglu (2023), who applied hybrid GIS-ML methods — including Random Forest — on the same datasets and achieved an R^2 of 0.85, MAPE of 19.4%, MAE of 44,888 unit (Pound sterling), and RMSE of 101,487 unit (Pound sterling). This comparison highlights the effectiveness of integrating deep neural networks with spatially enhanced

data for capturing more accurate and nuanced relationships between property attributes and their market values.

This level of accuracy is particularly significant when dealing with complex urban data characterized by high feature diversity and nonlinear relationships. When comparing the findings of this research with previous studies — especially hybrid GIS-ML models such as those by Mete and Yomralioglu (2023), which utilized similar data — it became evident that integrating artificial neural networks with spatial data enables higher precision and more accurate modeling of complex inter-dependencies between property attributes. This underscores the potential of deep learning techniques to significantly enhance the accuracy of mass property valuation in urban environments marked by high spatial-economic complexity.

The Multi-Layer Perceptron (MLP) architecture used in this study, with a configuration of 512-256-128-64-32-16-8-4-2-1, proved highly compatible with the multidimensional and heterogeneous nature of the dataset. The use of ReLU activation functions in hidden layers and a linear function in the output layer facilitated rapid convergence while improving model stability during training. Furthermore, feature standardization significantly improved the convergence speed of the model during training by ensuring uniformity across input variables with different scales.

Beyond its high predictive accuracy, one of the most significant features of this study lies in the use of extensive set of physical, environmental, spatial, and temporal features employed, all of which played a key role in enhancing model performance. These findings align with earlier studies, including those by Mimis et al. (2013) and Demetriou (2017), confirming the importance of incorporating spatial analysis-derived data into artificial intelligence models for real estate valuation.

Nonetheless, several challenges were identified throughout the study. These include the high computational demands associated with training on a large dataset (approximately 5.6 million samples), the presence of outliers, and right-skewness in the price distribution. Addressing these issues requires optimizing the training process, utilizing powerful computational resources such as GPUs, and applying preprocessing techniques like log transformation of price values.

Future research directions include several promising avenues. First, the integration of temporal dynamics—such as time-series modeling or spatio-temporal graph neural networks—could further enhance prediction accuracy by explicitly capturing market trends and price evolution over time. Second, incorporating additional data sources, such as

street-view imagery, additional land uses such as commercial properties, or socio-economic indicators at the neighborhood level, may enrich feature representation and improve model robustness. Third, extending the current framework to other geographic contexts (e.g., non-UK markets) would test its generalizability and adaptability across diverse urban and regulatory environments. Finally, exploring hybrid architectures that combine deep learning with spatial regression models (e.g., GTWR or spatial lag models) could leverage the strengths of both data-driven and theory-driven approaches for more interpretable and accurate mass appraisal systems.

Overall, this research demonstrates that the integration of artificial neural networks with spatially derived data can serve as a robust and effective tool for mass property valuation in both urban and regional settings. The proposed model has the potential to support government decision-makers, urban planners, and automated valuation models (AVMs) in conducting precise, fair, and objective property valuations independent of subjective expert judgment.

5. References

- Abidoeye, R. B., & Chan, A. P. (2018). Improving property valuation accuracy: A comparison of hedonic pricing model and artificial neural network. *Pacific Rim Property Research Journal*, 24(1), 71–83.
- Blackledge, M. (2016). *Introducing property valuation*. Routledge.
- 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*, 3, e019.
- Chollet, F. (2021). *Deep learning with Python*. Simon and Schuster.
- Demetriou, D. (2017). A spatially based artificial neural network mass valuation model for land consolidation. *Environment and Planning B: Urban Analytics and City Science*, 44(5), 864–883.
- Droj, G., Kwartnik-Pruc, A., & Droj, L. (2024). A comprehensive overview regarding the impact of GIS on property valuation. *ISPRS International Journal of Geo-Information*, 13(6), 175.
- Fuerst, F., McAllister, P., Nanda, A., & Wyatt, P. (2016). Energy performance ratings and house prices in Wales: An empirical study. *Energy Policy*, 92, 20–33.
- International Association of Assessing Officers. (2013). *Standard on Mass Appraisal of Real Property*. www.iaao.org
- International Association of Assessing Officers. (2018). *STANDARD ON Automated Valuation Models (AVMs)*. International Association of Assessing Officers. https://www.iaao.org
- Mete, M. O., Guler, D., & Yomralioglu, T. (2021). Towards a 3D real estate valuation model using BIM and GIS. *The Proceedings of the International Conference on Smart City Applications*, 945–962.
- Mete, M. O., & Yomralioglu, T. (2023). A hybrid approach for mass valuation of residential properties through geographic information systems and machine learning integration. *Geographical Analysis*, 55(4), 535–559.
- Mimis, A., Rovolis, A., & Stamou, M. (2013). Property valuation with artificial neural network: The case of Athens. *Journal of Property Research*, 30(2), 128–143.
- Yalpir, S., Durduran, S. S., Unel, F. B., & Yolcu, M. (2014). Creating a valuation map in GIS through artificial neural network methodology: A case study. *Acta Montanistica Slovaca*, 19(2).
- Yasnitsky, L. N., Yasnitsky, V. L., & Alekseev, A. O. (2021). The complex neural network model for mass appraisal and scenario forecasting of the urban real estate market value that adapts itself to space and time. *Complexity*, 2021(1), 5392170.