BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Real Estate Price Prediction Using Machine Learning Techniques: A Review and a Proposed Explainable Framework

[1]Raj Jagdish Thakre, [2]Dr. Divya Sharma

[1]Student, [2]Assistant Professor
Dept. of Artificial Intelligence and Data Science
Ajeenkya DY Patil School of Engineering
Pune, Maharashtra, India

Abstract

Valuation of real estate can be difficult as the values can veer greatly between regions with many variables to consider. What this article introduces is the summary of several studies that have been done between 2016 and 2026 where Machine Learning (ML) strategies have been used to project real estate pricing. In realtime, location and local variables didn't receive much concentrate, and most studies have focussed on regression or tree-based models such as ensemble or neural networks. Trait selection wasn't fully carried out and involved models are hard to explain, and so, this article will aim to create a 4 part framework. 6 mixed models will be compared (Linear Regression, Random Forest, XGBoost, LightGBM, CatBoost and Graph Neural Network (GNN)). SHAP will be applied to explain the models and a LLM will be applied to supply any summary of credit values. AutoML will be applied to do quality selection. From this view, one method has been put together for the assessment of this suggested framework applying MAE and RMSE values. As, this is a further review and design paper there are no findings offered.

Keywords : Real Estate Price Prediction, House Price Prediction, Machine Learning, AutoML, Graph Neural Network, Explainable AI, SHAP, Large Language Model, Geospatial Analytics

I. Introduction

My point is it's, recent advancements have, more accurately, permitted for house prices to be usually anticipated with much more precision, at least empirically, really, stated separately, feature price prediction is typically when ML models are normally adopted to project the value of a characteristic applying historical real estate data, in real-world cases in the fields of Artificial Intelligence (AI) and ML, on the whole. [1], [2].

Most considerably, there are consistently (specifically in this setting) many factors that go into the assessing the value of a house. What's more, all these factors could, after all, maybe be the house location, size, number of bedrooms, number of amenities around, infrastructure and the market, in actual application. Taken together, mL algorithms can look through the large amounts of data to spot relationships between the factors, which helps people buy/sell, invest, banks and agencies to make regularly informed decisions; this has been largely used in real estate platforms using AI and analytics.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

More or less mainly buying a home is the biggest investment most families will, in this context, make, so money on the asset will be generally loaned by banks and investors. Against this backdrop, often applying a data driven strategy will help to make predictions quickly and consistently (in particular in this setting) (to be precise in this setting); this can only be done by having the right traits in our model and making sure our predictions are precise and easy to read as.

What's more, what this article checks out is the process of projecting house prices, which is why a literature review has been completed to reveal the weakest points in the process, on the whole; data preparations, trait selection, model comparisons and explanations have been tackled; in turn, a design has been put forward to bring everything together, all things considered. Building on this observation, key clues of this article are;

- To explain, A comparative review of fifteen studies greatly published between 2016 and 2026, to a noticeable degree.
- A list of research gaps put together from that review.
- How about this, A person framework that links AutoML characteristic selection, model comparison (including a GNN), SHAP explanation and LLM-generated reports, in many respects.
- An check plan with suitable metrics.

A comparative study of fifteen articles currently published was made happen between 2016 and 2026, and introduce a combined framework for applying AutoML for credit selection, GNN for model comparison, SHAP for explanation and LLM for report generation; this article will be currently organised with section II normally outlining background, section III outlines our key findings, section IV tackles shortcomings in research and section V provides an overview of our suggested method, in the main. Before i go a conclusion is offered. What section VI outlines is how we performed check; Section VII and VIII outline the shortcomings and future work in light of this dynamic, in the main.

A. Problem Statement

Real estate prices differ considerably as stated by geographical location, and existing ML models have little regard for real-time location and the surrounding subfield, which limits their prediction exactness, in actual application.

B. Research Objectives

- To find, via AutoML-based trait selection, the characteristics most strongly influencing characteristic prices.
- To compare six ML strategies for real estate price prediction, such as Linear Regression, Random Forest, XGBoost, LightGBM, CatBoost and Graph Neural Networks (GNN), at least in practice.
- To gather the metrics often applied to compare price prediction models, specifically MAE, RMSE and the R^2score.
- To analyse Explainable AI (XAI) techniques, in particular SHAP, for taking as ML-based price predictions.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

- To study whether Large Language Models (LLMs) can, in effect, mostly turn prediction findings into the elementary explanations and valuation summaries, to a noticeable degree, in many respects.

C. Relevance to Current Technology Trends

Machine learning. Algorithms such as XGBoost, LightGBM, CatBoost and Random Forest make more exact predictions than classic methods, all things typically considered.

AutoML. AutoML picks credits, algorithms and model factors on its own, generally shortening development time and raising model performance, at least empirically.

Graph neural networks. GNNs represent the links between neighbouring characteristics and their surroundings, that way commonly increasing prediction precision, at least in principle.

Large language models. LLMs can, in this setting, probably write straightforward explanations, market summaries and valuation reports, usually helping buyers, sellers and investors understand a prediction, broadly speaking.

Geospatial analytics. Location data, such as the distance to schools, hospitals, transport and commercial subfields, can, in effect, increase valuation exactness, at least for now.

II. Literature Review

The initial studies were uncomplicated models applying a few variables such as linear regression, broadly speaking; over time, there are fifteen studies all spanning ten years to reveal development of the schemes, in real-world cases, broadly speaking. [15]. The studies fall into a few groups, examined in detail below and summarised in Table I, at least for now; the most recent studies, from 2023 to 2026, apply ensembles, geographic data fusion, macroeconomic traits and explainable strategies [1]–[7], at least empirically, and within this framework, studies between 2017 and 2021 applied hybrid, regularised and feature-selection techniques, and direct comparisons of several ML models with each person other and with hedonic regression [8]–[14], in real-world cases.

TABLE I

SUMMARY OF REVIEWED STUDIES

Ref.	Approach	Key Contribution	Limitation
[1] (2026)	Linear regression, SVR, random forest, MLP, LightGBM; SHAP	Explainable workflow with monetary bonus/malus tables; random forest gave the best balance of accuracy and explanation	Listing prices; one city; small gain over linear model
[2] (2025)	Prototype-based learning, non-linear regression	Accurate and explainable valuation using prototype comparison	Scalability and prototype selection
[3] (2024)	Linear regression, decision tree, random forest	ML improves house price prediction when data quality is good	Data quality and feature selection

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

[4] (2024)	CNN, MLP, geographical data fusion	Geographical data improves residential valuation accuracy	Spatial processing and model complexity
[5] (2023)	Logistic regression, random forest, voting classifier	Ensemble gave better accuracy than individual models	Data preprocessing and generalisation
[6] (2023)	Extra Trees, kNN, random forest versus OLS	ML models beat the hedonic baseline in test R^2 and error	One Hong Kong district; interpretability
[7] (2023)	Decision tree, random forest, XGBoost on pooled regions	Multi-region data; interest-rate feature helped; sale-price moving averages gave suspect scores	Preprint; one market (Jeonsae)
[8] (2021)	ANN with boosted feature selection	Feature selection improved ANN accuracy and revealed important features	ANN interpretability and computational cost
[9] (2021)	Decision tree, SVM, regression	Property features improve prediction performance	Feature engineering and prediction accuracy
[10] (2021)	ANN, random forest, kNN versus hedonic model	Random forest best (R^2 0.86); non-linear feature effects	Small single-city sample; limited interpretability
[11] (2021)	SVM, random forest, gradient boosting	RF and GBM beat SVM (R^2 about 0.90 against 0.83); SVM is faster	One Hong Kong district; hard to interpret weights
[12] (2020)	Multiple linear regression, ridge, lasso, SVR	Regularised models beat basic linear regression	Data preprocessing and parameter tuning
[13] (2017)	Hybrid lasso + gradient boosting	High accuracy through effective feature engineering	Feature engineering and overfitting
[14] (2017)	RF, regression, SVM, GBT, MLP, bagging in R, RapidMiner, Weka	Random forest reported best among six models	Accuracy figures unclear; some values inconsistent
[15] (2016)	Linear regression	Historical data can be used for house price forecasting	Limited features and data coverage

Classical regression models. Mohokar et al. showed that historical data could be applied with linear regression to forecast house prices, although the credits and data coverage were restricted [15]. Checking multiple linear regression with ridge, lasso and SVR, we can conclude that they do raise prediction as opposed to the elementary linear regression. But still, you need to take care on how they're preprocessed and tuned. [12] in this broader context. Reshma et al. tested linear regression, decision tree and random

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

forest and concluded that ML helps if data quality is good; data quality and trait selection are unresolved [3]. Singh et al. adopted decision tree, SVM and regression and found that characteristic features sharpen prediction performance [9].

Hedonic regression versus ML. Hedonic regression, what which models is price as one linear function of credit characteristics, is the standard baseline for checking ML models, in actual application; yazdani connected a second hedonic model with ANN, random forest and kNN on 1,018 houses in Boulder, Colorado, in real-world situations, at least for now. Random forest achieved this best test R^2 (0.86) followed by ANN (0.85), the hedonic model (0.82) and kNN (0.75), and the study reports any non-linear association between lingering traits and price [10]; the three ML models hit test R^2 between 0.895 and 0.919, against 0.814 for OLS [6], at least in principle. Under these conditions, choy and Ho compared Extra Trees, kNN and random forest with ordinary least squares on about 25,000 housing transactions in Hong Kong.

Support vector regression. Ho et al. contrasted SVM, random forest and gradient boosting on 39,554 housing transactions from 14 estates in Hong Kong, in real-world cases; random forest and gradient boosting reached a test R^2 of about 0.90, against 0.83 for SVM, but SVM supplied reasonable predictions in much less time and the authors hint it where speed is the key issue [11], at least empirically. An RBF-kernel SVR achieved an R^2 of only 0.10, under cross-validation, proving that kernel techniques are strongly sensitive to tuning and data [1] in the Cagliari benchmark of Saiu and Mocci in parallel with this, at least in practice.

Ensemble and hybrid models. Sreeja et al. combined logistic regression and random forest in a voting classifier, which outperformed the person models, although preprocessing and generalisation were challenging [5], on the whole; XGBoost performed best and adding an interest-rate trait sharpened the metrics, in many respects. Ravikumar connected random forest, multiple regression, SVM, gradient boosted trees, a multilayer perceptron and bagging in R, RapidMiner and Weka and reported random forest as the most correct model; yet, the “exactness” percentages are unclear for a regression task, and some values are inconsistent between the table and the text of that study [14], broadly speaking. Moving averages of the sale price gave R^2 values of 0.99–1.00, which the authors assessed as reflecting hindsight bias, so they applied moving averages of the lease-to-sale price ratio rather; this source is a preprint [7], really, kong et al. pooled data from 32 regions of Seoul and Gyeonggi into 4,795 entries and contrasted decision tree, random forest and XGBoost on the Korean lump-sum lease (Jeonsae) market, broadly speaking. Along these dimensions, lu et al. combined lasso with gradient boosting and reached a high exactness with careful feature engineering; the risk is, as it stands, overfitting [13].

Neural networks and geographic data. Chanasit et al. paired an artificial neural network (ANN) with a boosted feature-selection step, and lee et al. fused geographical data with CNN and MLP models and showed geographic information raises residential valuation, at expense of spatial processing and model fine detail [4], at least in practice. Viewed through this lens, selection increased precision and spotted essential characteristics, although the ANN is hard to take as and computationally expensive [8], at least in principle.

Explainable valuation. Rodriguez-Serrano applied prototype-based learning with non-linear regression, so that a price can be explained by comparison with prototype features, in actual application, and scalability and prototype selection are the open points [2], and saiu and Mocci built the RE-VAL process on 1,153 listings from Cagliari, Italy, to a important degree, in real-world cases. They compared linear regression, SVR, random forest, MLP and LightGBM with ten-fold cross-validation and used SHAP to convert model

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

outputs to monetary bonus and malus adjustments, really, the MLP hit the lowest departure (MAE about €474 per square metre, R^2 0.59), although its margin over linear regression (0.58) and random forest (0.57) was small, and random forest supplied the best balance of exactness and explanation, to a important degree. Along these dimensions, the authors also regard missing listing credits as a possible signal and stress that SHAP findings reflect associations, not causes [1].

One pattern emerges, like the studies center on regression, ensemble and neural models, and tackle location, interpretability and generalisation across the regions only partially, in the main; more accurately, at all these shortcomings, the design in Section V is tackled, to a important degree.

III. Discussion and Findings

Simple and tree-based models are the most popular. Dataset sizes range from about a thousand data [1], [10] to about forty thousand transactions [11], and most cover only one city, district or market, in the main; against this backdrop, most of the reviewed studies are often grounded in linear, hedonic, mainly regularised, tree-based or SVR models [3], [6], [9]–[12], [15], in many respects.

Tree ensembles are generally best, but margins vary. Random forest or gradient particularly boosting out-performed SVM, OLS or the hedonic baseline in several studies [6], [10], [11], [14], at least in principle, really, in the Cagliari benchmark, still, the MLP, random forest and linear regression were, typically after all, reliably all close in R^2 (0.57–0.59) [1], so the best model hinges on the dataset, broadly speaking.

Location is underused. Location is ordinarily handled by hand-made traits, such as urban zones in [1] or distance to the city centre in [10], broadly speaking, really, none represent normally neighbouring credits as a graph, in the main. Taken together, of the fifteen studies, only Lee et al. feed geographical data directly into the model [4], broadly speaking; prices are mainly shaped mostly by location, broadly supporting the issue statement and the importance of geospatial characteristics and a GNN, in the main.

Feature selection helps but is incomplete. Taken together, kong et al. reveal the macro-economic trait can help, and traits put together from this target price can create scores near 1.00 that probably reflect leakage [7], so one automatic AutoML selection step run, within each person validation fold would be helpful, at least in principle, really, despite, a few studies still point to credit selection and feature engineering as open hurdles [3], [9], [13], broadly speaking. Selection with boosting and careful feature engineering raised findings [8], [13].

Interpretability is a gap. ANNs can (to some degree) be precise but are hard to read as [8], and Yazdani points out that random forest and ANN are less transparent than hedonic coefficients, suggesting LIME or SHAP as future work [10], to a noticeable degree; indeed, looking at it like this, only Rodriguez-Serrano [2] and Saiu and Mocci [1] concentrate on explaining prices, the latter applying SHAP to build monetary adjustment tables, at least empirically. As, Table I hints, none of the fifteen studies create natural-language valuation reports, which is the motivation for the explanation stage of the framework.

Reporting is not comparable. Studies report log prices [11], normalised prices [14], price per square metre [1] or exactness percentages, so their departures can't be compared directly; this supports routinely reporting MAE, RMSE and R^2 for all models on the same folds, on the whole.

No single model fits every need. Table II checks the primary model families, which supports widely checking several models rather than uniformly fixing one in advance, at least empirically.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

TABLE II

COMPARISON OF MODEL FAMILIES

Family	Strength	Weakness
Linear / regularised / hedonic [6], [12], [15]	Simple; readable coefficients	Limited features; misses non-linear effects
SVR / SVM [1], [11]	Fast; reasonable fit in [11]	Below RF and GBM; depends on tuning
kNN [6], [10]	Simple; non-parametric	Weaker than random forest in both studies
Ensembles [5]–[7], [11], [13]	Higher accuracy	Overfitting; generalisation
ANN / CNN / MLP [1], [4], [8], [10]	Learn complex patterns; use geographic data	Hard to interpret; costly
GNN	Uses neighbouring properties	Needs location data and graph building

IV. Research Gaps and Challenges

The review spots five gaps, and the framework is typically strong at tackling each, at least empirically.

- **G1 – Scalability and dataset diversity.** Datasets cover only a few regions and feature markets [2], [5], [6], [10], [11].
- **G2 – Feature selection.** Strange but true, noticeable characteristic and location factors are, often in effect, only partially considered, and credits received from the target can, as a matter of fact, leak information [3], [7], [9], to a important degree.
- **G3 – Generalisation.** Models systematically trained on regional data, under the varying market conditions can't be specifically trusted far beyond their training field [1], [5], [12], to a important degree, at least in practice.
- **G4 – Interpretability.** Generally involved ML and ANN models give hard-to-understand predictions [8], [10].
- **G5 – Optimisation.** Often overfitting, factor tuning and computational cost still need better usually handling [8], [11], [13].

V. Proposed Methodology

The suggested framework links data preparation, model comparison and explanation in one more person pipeline of six stages, laid out in Fig. 1, broadly speaking, and data flows through the numbered boxes in order, and the check stage computes which model yields the final prediction and explanation; this architecture is any suggested design for future execution and check, not yet built or proved genuine, broadly speaking.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

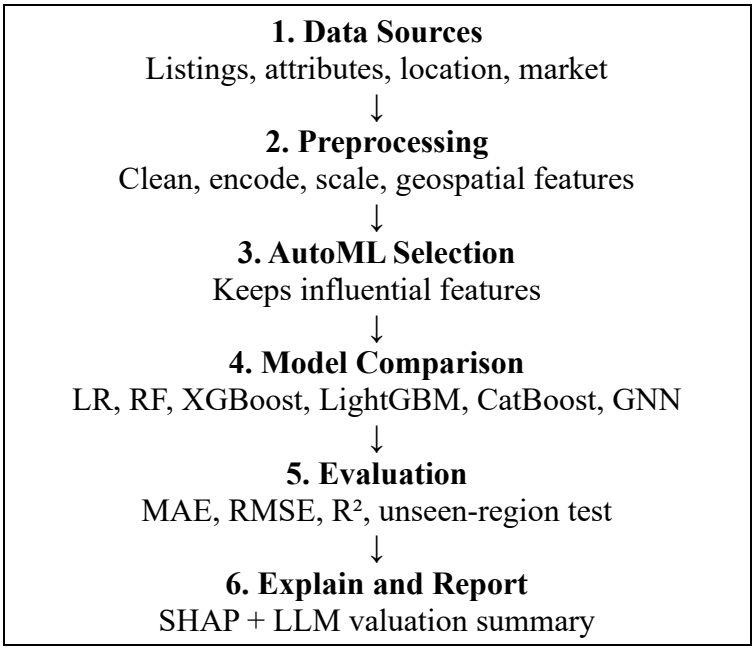


Fig. 1. Six-stage pipeline of the proposed framework, from data sources to the explained valuation.

A. Problem Formulation

What if attribute i is normally embodied by a feature vector xi covering its credits, location and surroundings, and has an gathered price yi. In parallel with this, the aim is currently to learn a function f that maps features to price so that the prediction ŷi = f(xi) is usually close to yi. This is finished by minimising the average loss over the training characteristics:

$$f^* = \arg\min_f \ (1/n) \ \sum_{i=1}^{n} \ L(y_i, f(x_i)) \qquad (1)$$

Here L is a loss such as the squared mismatch and n is the number of training features, on the whole; the vector xi carries three kinds of information, generally including credits of the characteristic, such as subfield and number of bedrooms, at least in practice, and under these conditions, its location, such as locality and coordinates; and its surroundings, such as the distance to schools, hospitals, transport and commercial fields, at least in principle.

B. Data Preparation and AutoML Feature Selection

The first three stages of Fig. 1 prepare the data, really, feature entries are pulled from public housing price datasets and, if available, from recent local listings, at least in principle. The preprocessing step removes copy and incomplete data, fills missing values, encodes categorical fields, scales numeric fields and adds geospatial characteristics such as the distance to nearby facilities, all things considered; and this way, traits which are built from the target price, such as moving averages of the sale price, are verified for leakage [7], at least in principle. In parallel with this, an AutoML step then ranks the credits and keeps the ones which most impact price, leaving the models to train on a smaller and more relevant set; the ranking is repeated within each cross-validation fold, so that no information from the test data reaches the selection, at least in principle.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

C. Model Comparison

On the picked characteristics and compared in stages 4 and 5, six models are trained, in actual application; Linear Regression is the baseline, Random Forest, XGBoost, LightGBM, and CatBoost are tree-based ensembles, and the GNN uses neighbouring credits, at least in principle. Table III lists them. Along these dimensions, their settings are tuned by the AutoML step, and every model is assessed on the same folds, making the comparison fair, and the model with the lowest validation departure is, indeed, applied for the final prediction, broadly speaking.

TABLE III

CANDIDATE MODELS IN THE COMPARISON

Model	Type	Role in the study
Linear Regression	Linear	Baseline
Random Forest	Bagged trees	Robust ensemble
XGBoost	Boosted trees	High accuracy
LightGBM	Boosted trees	Fast on large data
CatBoost	Boosted trees	Handles categories well
GNN	Graph network	Uses neighbouring properties

D. Property Graph for the GNN

Each credit in the GNN is a node which is especially connected to its nearest neighbour characteristics, as in Fig. 2. Along these dimensions, facilities, such as schools, hospitals, transport and markets, offer location setting to the node traits. In each layer, each node is principally changed from its neighbours:

$$h_v^{(l+1)} = \sigma(W \cdot AGG(\{ h_u^{(l)} : u \in N(v) \})) \quad (2)$$

The final node state is passed to a regression layer to output the price, and Algorithm 1 summarises the whole pipeline, at least in principle. Where N(v) is the set of credits which are the neighbours of characteristic v, AGG is an aggregator such as the mean, W is a learned weight matrix and σ is a non-linear function, in many respects in parallel with this, on the whole.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

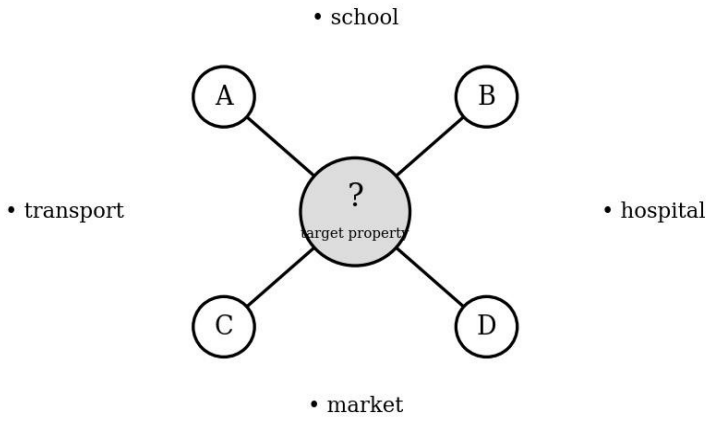


Fig. 2. Example property graph where the target property connects to four of its neighbours (A–D), with nearby facilities providing location context.

Algorithm 1. Price prediction and explanation
1: Input: property data D, query properties Q
2: clean and encode D; add location features; build graph G
3: F ← AutoML SELECT(D) {most influential features}
4: train LR, RF, XGBoost, LightGBM, CatBoost, GNN on F
5: record MAE, RMSE, R^2 and run time by k-fold CV
6: m ← model with the lowest validation error
7: for every query property x ∈ Q do
8: predict $\hat{y} \leftarrow m(x)$; explain $\varphi \leftarrow SHAP(m, x)$
9: write report R ← LLM SUMMARY(x, $\hat{y}$, leading features of φ)
10: save ($\hat{y}$, φ, R) and output R
11: end for

Fig. 3. Price prediction and explanation algorithm.

E. Explanation and Valuation Report

Once the best model is chosen, SHAP values are found out for all the predictions to outline the takeaway of credits pushing the price up or down; these values explain what the model has learned and aren't proof of causation [1], in many respects, and the projected price and the leading SHAP credits are then passed to an LLM that writes a short summary in plain language for buyers, sellers and investors, at least for now. From here, table IV reveals the fields of a valuation report. The LLM only repeats the values it's given and never computes the price, so the text stays clear with the model, and a uncomplicated check also connects the figures and factors mentioned in the summary with the model output, and any summary that doesn't match is rejected, at least for now.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

TABLE IV

FIELDS OF A VALUATION REPORT

Field	Content
Property	Location, area, bedrooms, amenities
Predicted price	Output of the selected model
Model used	Model name and validation error
Key factors	Leading SHAP features and their direction
Summary	LLM-written plain-language valuation text

VI. Evaluation Plan

A framework of this kind should be benchmarked as stated by prediction exactness, generalisation to new regions, computing cost and the quality of explanations. In this broader context, precision is measured applying three standard metrics. In them, yi represents the actual price, ŷi the forecasted price, ȳ the mean of the actual prices and n the number of test features:

$$MAE = (1/n) \sum_{i=1}^{n} | y_i - \hat{y}_i | \quad (3)$$

$$RMSE = \sqrt{[(1/n) \sum_{i=1}^{n} (y_i - \hat{y}_i)^2]} \quad (4)$$

$$R^2 = 1 - \sum_i (y_i - \hat{y}_i)^2 / \sum_i (y_i - \bar{y})^2 \quad (5)$$

What mAE outlines is the average departure relative to price units, and in price is explained by the model, r2 outlines how much of the change, at least for now; an exact description can (at least in most settings) be given since they reveal distinct facets of the difference by reporting all three, on the whole. From this view, rMSE is sensitive to extremes which is essential when a few badly valued credits are highly priced, in many respects.

Public housing datasets such as Ames and California housing can, to some extent, be applied for a preparatory benchmark, at least in practice; since, regional generalisation is a priority, testing will be done applying k-fold cross-validation together with a hold-out of entire regions which the model hasn't seen during the training, at least for now. Viewed through this lens, a time-ordered split will be added to guarantee that future prices aren't applied to anticipate past ones, if the entries span several years, and training and prediction time will (in practical terms) be recorded for each model, so that the computing cost of boosting and GNN models can be contrasted without counting on assumptions, on the whole.

VII. Limitations

The article reports a review of fifteen studies and one design, but introduces no measured findings; the models and settings will (in actual operation) need to be tested on real data before the framework can be used for valuation, in many respects, and some of the mentioned works are preprints or a student research project [7], [10], [14] and may (to a meaningful extent) not have been peer reviewed, to a important degree, on the whole. Under these conditions, the reviewed studies applied mixed datasets, currencies, price

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

transformations and metrics, so the findings aren't directly comparable and aren't applied for algorithm ranking, all things considered.

VIII. Future Work

The next step is to perform the pipeline and test it on public datasets and recent local data, including a hold-out of whole regions, broadly speaking; other directions of work are real-time location and market data feeds, transfer of models between cities, price intervals to reveal the confidence of a prediction, in real-world situations, on the whole. Along these dimensions, a web application that delivers the valuation report to the user, in actual application.

IX. Conclusion

It has been showed that Machine Learning can, in practice, be a very effective tool for housing price prediction, in the main, and yet, it was found in all these studies that refinements could (in practice) be made to given facets, including studying real time location and setting, selecting proper characteristics, interpreting models and generalising between regions, broadly speaking. A full pipeline was suggested which covers AutoML based credit selection, regression, boosting and GNN model comparison, model interpretability applying SHAP values and creating valuation summaries using an LLM. 15 distinct studies were looked into the throughout this article; under these conditions, the check was grounded in MAE, RMSE and R^2 as metrics and tests unseen regions was also offered. If adopted on a varied and high-quality set of data, such a pipeline can raise prediction and explanation.

Acknowledgment

The author thanks Dr. Divya Sharma for her mentorship and the Department of Artificial Intelligence and Data Science, Ajeenkya DY Patil School of Engineering, for the facilities, to a noticeable degree, all things often considered.

References

[1] V. Saiu and M. Mocci, "Explainable AI for urban real-estate prediction: A machine-learning framework for urban decision support," Urban Science, vol. 10, no. 6, Art. no. 315, 2026, doi: 10.3390/urbansci10060315.

[2] J. A. Rodriguez-Serrano, "Prototype-based learning for real estate valuation: A machine learning model that explains prices," Ann. Oper. Res., vol. 344, pp. 287–311, 2025.

[3] Reshma, T. Semwal, S. Pareek, and S. Kuhar, "House price prediction using machine learning approaches," in Proc. 2024 OPJU Int. Technol. Conf. (OTCON), IEEE, 2024.

[4] H. Lee, H. Han, C. Pettit, and Q. Gao, "Machine learning approach to residential valuation: A convolutional neural network model for geographic variation," Ann. Reg. Sci., vol. 72, pp. 579–599, 2024.

[5] S. P. Sreeja, C. P. Nivash, V. Asha, P. Om Prakash, B. Saju, and A. K. Singh, "Real estate price prediction using machine learning," in Proc. 3rd Int. Conf. Adv. Electr., Comput., Commun. Sustain. Technol. (ICAECT), IEEE, 2023.

[6] L. H. T. Choy and W. K. O. Ho, "The use of machine learning in real estate research," Land, vol. 12, no. 4, Art. no. 740, 2023, doi: 10.3390/land12040740.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

[7] S.-H. Kong, J. Jang, and H. Oh, "Predicting key regional real estate prices using machine learning technique: With an emphasis on Jeonsae system," SSRN preprint (not peer reviewed), 2023. [Online]. Available: https://ssrn.com/abstract=4662626
[8] K. Chanasit, E. Chuangsuwanich, A. Suchato, and P. Punyabukkana, "A real estate valuation model using boosted feature selection," IEEE Access, vol. 9, 2021.
[9] A. P. Singh, K. Rastogi, and S. Rajpoot, "House price prediction using machine learning," in Proc. 3rd Int. Conf. Adv. Comput., Commun. Control Netw. (ICACCCN), IEEE, 2021.
[10] M. Yazdani, "Machine learning, deep learning, and hedonic methods for real estate price prediction," arXiv:2110.07151 [econ.EM], Oct. 2021.
[11] W. K. O. Ho, B.-S. Tang, and S. W. Wong, "Predicting property prices with machine learning algorithms," J. Property Res., vol. 38, no. 1, pp. 48–70, 2021, doi: 10.1080/09599916.2020.1832558.
[12] J. Manasa, R. Gupta, and N. S. Narahari, "Machine learning based predicting house prices using regression techniques," in Proc. 2nd Int. Conf. Innov. Mech. Ind. Appl. (ICIMIA), IEEE, 2020.
[13] S. Lu, Z. Li, Z. Qin, X. Yang, and R. S. M. Goh, "A hybrid regression technique for house prices prediction," IEEE, 2017.
[14] A. S. Ravikumar, "Real estate price prediction using machine learning," MSc Research Project (Data Analytics), School of Computing, National College of Ireland, Dec. 2017.
[15] A. Mohokar, S. Mane, and N. Bhagat, "House price forecasting using data mining," Int. J. Comput. Appl., 2016.