

Machine Learning-Based Residential Property Price Prediction: A Multi-City Study on Romanian Real Estate Data

Cristina Florentina Mihai
University of Craiova
Craiova, Romania
mihai.cristina.n8i@student.ucv.ro

Paul Stefan Popescu
University of Craiova
Craiova, Romania
stefan.popescu@edu.ucv.ro

Marian Cristian Mihaescu
University of Craiova
Craiova, Romania
cristian.mihaescu@edu.ucv.ro

ABSTRACT

This paper presents a machine learning-based approach for residential property price prediction across five Romanian cities: Bucharest, Brasov, Constanta, Craiova, and Timisoara. A dataset of 1,921 real estate listings was collected from multiple sources, including public notary market reports and Romanian real estate platforms. Four regression algorithms were evaluated: Random Forest, XGBoost, HistGradientBoosting, and a Multi-Layer Perceptron (MLP). To address the limited number of samples in certain cities, a Gaussian noise augmentation technique was applied, increasing the dataset size fourfold. Results show significant improvements after augmentation, with R^2 scores ranging from 0.839 (Timisoara) to 0.948 (Brasov) and Mean Absolute Errors between 5,145€ and 9,839€. The models are integrated into a web platform built with ASP.NET Core MVC, enabling automatic price estimation for newly listed properties.

Author Keywords

Real estate price prediction; machine learning; random forest; XGBoost; data augmentation; Romanian housing market; ASP.NET Core.

ACM Classification Keywords

H.5.m. Information interfaces and presentation: Miscellaneous; I.2.6. Learning; J.4. Social and behavioral sciences.

General Terms

Algorithms; Measurement; Performance

DOI: 10.37789/icusi.2026.5

INTRODUCTION

Accurate property price estimation is a critical challenge in modern real estate markets, where prices are influenced by a complex combination of factors including location, surface area, number of rooms, and proximity to urban centers. In Romania, the residential real estate market has experienced significant growth over the past decade, with prices varying considerably between cities and neighborhoods. Despite this growth, automated valuation tools tailored to the Romanian market remain scarce, and

most existing platforms rely on manual price suggestions provided by sellers.

The Romanian residential real estate market has undergone substantial transformation in recent years. According to public notary data, apartment transaction volumes have increased significantly in major cities, with Bucharest, Timisoara, and Constanta recording the highest activity levels among the cities analyzed in this study. Price disparities between cities are considerable - apartments in Bucharest trade at median prices nearly double those in smaller cities such as Brasov or Constanta. These factors make accurate automated price estimation particularly valuable for both buyers and sellers navigating an increasingly complex market.

Machine learning techniques have demonstrated strong predictive capabilities in real estate valuation tasks across various countries [1, 2]. Methods such as Random Forest and gradient boosting have consistently outperformed traditional regression approaches, particularly when dealing with heterogeneous datasets and non-linear relationships between features and target variables [3].

This paper addresses the lack of localized prediction tools by presenting a multi-city study conducted on Romanian residential property data. The main contributions of this work are:

- (1) the construction and analysis of a real estate dataset covering five Romanian cities, collected from heterogeneous sources;
- (2) the application of a Gaussian noise augmentation strategy to address class imbalance and limited sample sizes;
- (3) a comparative evaluation of four machine learning algorithms across all cities; and
- (4) the integration of the best-performing models into a web-based real estate platform.

The remainder of the paper is organized as follows: Section 2 reviews related work, Section 3 describes the platform, Section 4 the dataset, Section 5 the methodology, Section 6 the results, and Section 7 concludes.

RELATED WORK

Real estate price prediction has been extensively studied using machine learning techniques. Géron [1] provides a comprehensive overview of supervised learning methods applicable to regression tasks, including ensemble methods that have proven effective in property valuation. Breiman [2] introduced Random Forest as a robust ensemble learning method, demonstrating strong performance on high-dimensional regression problems with non-linear feature interactions - characteristics commonly found in real estate datasets.

Gradient boosting methods have also shown strong results in property price prediction. Chen and Guestrin [3] proposed XGBoost, an optimized gradient boosting framework that has become widely used in structured data regression tasks due to its efficiency and accuracy. Hastie et al. [4] provide a theoretical foundation for ensemble and regularization methods applied to regression problems, including regularization techniques that help prevent overfitting on small datasets.

Several studies have investigated the application of machine learning to residential price prediction in various geographic contexts. Čeh et al. [5] demonstrated that Random Forest outperforms traditional mass appraisal methods for apartment price prediction, achieving higher accuracy on urban real estate datasets. Antipov and Pokryshevskaya [6] applied ensemble methods to the Russian real estate market and found that gradient boosting models consistently produce lower prediction errors compared to linear regression and support vector machines. Truong et al. [7] conducted a comparative study of multiple regression algorithms on Vietnamese housing data, concluding that XGBoost achieves the best balance between accuracy and computational efficiency for structured tabular data.

Feature engineering plays a critical role in real estate price prediction. Geographic features such as distance from the city center, proximity to public transport, and neighborhood characteristics have been shown to significantly improve model accuracy [5, 7]. In this work, distance from the city center was computed using OpenStreetMap geocoding, providing a continuous spatial feature that captures the effect of location on property value across different urban contexts.

Data augmentation techniques for tabular regression tasks have received less attention than in image or text domains. However, perturbation-based methods similar to the approach used in this work have been shown to improve model generalization when training data is scarce [1]. To the best of our knowledge, this is the first study to apply Gaussian noise augmentation specifically to Romanian real estate data across multiple cities.

Beyond the algorithms adopted in this study, a broader family of methods has been explored for property valuation.

Classical approaches build on the hedonic price framework of Rosen [8], which models a property's price as a function of its individual attributes and remains the theoretical basis for most feature-driven valuation models. Subsequent work has emphasized the spatial nature of housing data: Bourassa et al. [9] compared several techniques for incorporating spatial dependence and found that modeling location explicitly improves out-of-sample accuracy, which motivates the geographic features used in this work. Computational-intelligence methods have likewise been applied to mass appraisal, with Kontrimas and Verikas [10] reporting that neural networks and support vector machines outperform classical linear regression on heterogeneous transaction data.

Within the gradient-boosting family, several optimized implementations have followed XGBoost. LightGBM [11] introduces histogram-based learning and leaf-wise tree growth for faster training on large datasets, while CatBoost [12] adds ordered boosting and native handling of categorical variables, both of which suit the mixed numerical and categorical attributes typical of real estate listings. As such models grow more complex, interpretability becomes important for practical adoption; SHAP values [13] offer a unified, game-theoretic approach for attributing a prediction to individual input features, providing a principled way to explain per-property estimates to prospective buyers and sellers.

While the studies discussed above focus almost exclusively on predictive accuracy, the usability of the interface through which price estimates are delivered is an equally important factor for the adoption of real estate tools. Heuristic evaluation, introduced by Nielsen and Molich [14], is a widely adopted inspection method in which a small number of evaluators examine an interface against a set of recognized usability principles. Nielsen's ten usability heuristics [15] remain among the most frequently used criteria for assessing interactive systems, including web and e-commerce applications, because they reveal a large proportion of usability problems at a low cost. Building on this body of work, the present paper complements its predictive evaluation with a heuristic assessment of the proposed platform, reported in Section 6.

PLATFORM OVERVIEW

The prediction models are integrated into a web-based real estate platform developed using ASP.NET Core MVC, following the Model-View-Controller architectural pattern. The platform allows users to browse, publish, and manage property listings across the five cities covered by the study. A dedicated price estimation module is embedded directly into the listing creation form, providing sellers with an automatic price suggestion based on the property's characteristics.

The machine learning service is implemented in Python using scikit-learn and exposed as a REST API via Flask, running on port 5001. When a user completes the listing form with details such as city, property type, number of rooms, surface area, and distance from the city center, the web application sends an asynchronous JSON request to the Flask service, which returns the estimated price along with a confidence interval. Separate models are trained and deployed for each city, allowing the system to account for local market differences. Figure 1 shows the overall system architecture.

The price estimation widget, illustrated in Figure 2, is embedded directly in the listing creation form. As the user fills in property details, the estimated price is updated in real time. If no trained model is available for the selected city, an informative message is displayed instead. The platform also includes user authentication via ASP.NET Core Identity, a private messaging system, a favorites module, and a property promotion feature.




Figure 1. Architecture of the system



Figure 2. Price estimation module

DATASET

The dataset used for training and evaluating the prediction models was constructed by integrating data from multiple sources, depending on the city. For Brasov, data was collected from market studies and public reports published by the Romanian National Union of Public Notaries (UNNPR). For the remaining four cities - Bucharest, Constanta, Craiova, and Timisoara - data was gathered through web scraping from publicly available Romanian real estate platforms. The study focused exclusively on the apartment sales segment, which represents one of the most active segments of the Romanian residential market. The data collection pipeline is illustrated in Figure 3.




Figure 3. Data collection pipeline

The final dataset contains 1,921 instances distributed across five cities, as shown in Table 1. Each record contains 12 attributes covering property characteristics, location data, and pricing information.

CITY	INSTANCES
Brasov	406
Bucharest	366
Constanta	480
Craiova	114
Timisoara	555
TOTAL	1921

Table 1. Dataset distribution per city

Table 2 provides a detailed description of each attribute. The dataset includes both property-level features - such as surface area, number of rooms, and construction year - and location-based features, including geographic coordinates and distance from the city center. The zone attribute, extracted through address parsing, captures neighborhood-level information that may influence property value. Price information is available in both EUR and RON, with per-square-meter values included to facilitate cross-city comparisons.

Attribute	Description	Type
Nr. rooms	Number of rooms	Numeric
Address	City, street	Text
Zone	Neighborhood extracted from address	Text
Price (€)	Sale price in EUR	Numeric
Price (€/m ²)	Price per square meter in EUR	Numeric
Surface (m ²)	Useful surface area	Numeric
Construction year	Year the building was constructed	Numeric
Latitude	Geographic coordinate	Numeric
Longitude	Geographic coordinate	Numeric
Distance (km)	Distance from city center	Numeric

Table 2. Dataset attribute description

Exploratory Data Analysis

Descriptive statistics reveal considerable price variation across the five markets. Bucharest records the highest median price (114,000€) and the widest price range (40,000€–315,000€), reflecting the heterogeneity of the capital's real estate market. Brasov presents the lowest median price (67,600€) and the most compact distribution. Timisoara and Craiova occupy intermediate positions with similar median values around 100,000–110,000€. All cities exhibit right-skewed price distributions, a behavior typical of real estate markets, where the majority of transactions are concentrated at lower price ranges with a long tail toward premium properties, as illustrated in Figure 4.




Figure 4. Apartment price distribution per city

Coefficient of variation analysis further confirms the heterogeneity of each market. Timisoara (CV=40.1%) and Craiova (CV=42.0%) present the most uniform price distributions, suggesting stable and predictable markets. Bucharest (CV=51.4%) and Constanta (CV=51.5%) show higher variability, with Constanta's dispersion likely influenced by the mixed residential and touristic character of the city.

Correlation Analysis

Pearson correlation analysis was conducted between each independent variable and the target price to identify the most relevant predictors. Surface area emerged as the strongest predictor across all cities, with correlation coefficients of $r = 0.88$ in both Constanta and Craiova, $r = 0.78$ in Timisoara, $r = 0.75$ in Brasov, and $r = 0.66$ in Bucharest. The lower correlation in Bucharest suggests that factors such as neighborhood and zone play a more significant role in the capital.

Number of rooms also shows strong correlation with price in most cities ($r = 0.65$ – 0.79), with the exception of Brasov ($r = -0.09$), where price appears to be driven mainly by surface area and building characteristics rather than by room count.

Distance from the city center presents a moderate negative correlation in Craiova ($r = -0.25$) and Bucharest ($r = -0.44$), confirming that proximity to the center increases property value in these markets. In Constanta, the relationship becomes slightly positive ($r = 0.09$), reflecting the presence of premium residential zones located outside the city center. These differences across cities justify the decision to train separate models per city rather than a single unified model.

Overall, the correlation analysis confirms that surface area and number of rooms are the most reliable predictors for apartment price estimation across the Romanian markets analyzed, while distance from the city center provides additional but context-dependent predictive value. Figure 5 illustrates the correlation structure obtained for the Craiova subset.




Figure 5. Correlation matrix – Craiova

METHODOLOGY

Data Preprocessing

Prior to model training, several preprocessing steps were applied to ensure data quality. Listings with unrealistic prices were removed, keeping only properties in the range of 15,000–400,000€ for apartments. Outliers were identified and removed using the Interquartile Range (IQR) method with a coefficient of 2, applied as follows:

$$Threshold = [Q1 - 2 \times IQR, Q3 + 2 \times IQR]$$

Only properties with surface areas between 20 m² and 1,000 m² were retained. Missing values for distance from the city center were imputed using the city median, except for Bucharest, where distances were estimated based on the administrative sector extracted from the listing address. The city was divided into six administrative sectors, each assigned a representative median distance value based on its geographic position relative to the city center: Sector 1 (3.2 km), Sector 2 (4.1 km), Sector 3 (3.8 km), Sector 4 (4.5 km), Sector 5 (5.1 km), and Sector 6 (4.8 km). The sector was automatically extracted from the listing address using string parsing.

Categorical variables derived from address processing were encoded using Label Encoding. Given the right-skewed distribution of prices, a logarithmic transformation was applied to the target variable before training:

$$y_{transformed} = \ln(1 + Price)$$

Predictions were converted back to the original scale using the inverse transformation.

Data Augmentation

After preprocessing, some cities - particularly Bucharest (329 instances) and Craiova (107 instances) - had limited sample sizes. To improve model generalization and reduce overfitting risk, a Gaussian noise augmentation strategy was applied. For each original record, three synthetic instances were generated by adding small perturbations ($\pm 1\%$) to continuous features such as surface area and distance from the city center. This resulted in a fourfold increase in dataset size. A fixed random seed (seed=42) was used to ensure reproducibility of the augmented samples. The dataset sizes before and after augmentation are summarized in Table 3. The counts in Table 3 are those left after outlier removal, hence the difference from Table 1.

CITY	AFTER PREPROCESSING	AFTER AUGMENTATION
Brasov	388	1 552
Bucharest	329	1 316
Constanta	445	1 780
Craiova	107	428
Timisoara	534	2 136

Table 3. Dataset size before and after augmentation

Model Training and Evaluation

Four regression algorithms were trained and evaluated: Random Forest, XGBoost, HistGradientBoosting, and a Multi-Layer Perceptron (MLP). A separate model was trained for each city to account for local market differences. Data was split into 80% training and 20% test sets. Performance was evaluated using three metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). For the MLP model, features were standardized using StandardScaler prior to training. A fixed random seed was applied across all models to ensure reproducibility of results.

Hyperparameter tuning was performed for each algorithm using cross-validation on the training set. For Random Forest, the number of estimators was set to 100 and the maximum depth was determined through grid search. XGBoost was configured with a learning rate of 0.1 and early stopping to prevent overfitting. HistGradientBoosting used its default configuration with a maximum of 100 iterations. The MLP architecture consisted of two hidden layers with 64 and 32 neurons respectively, using ReLU activation functions and the Adam optimizer. All models were saved using joblib and deployed as part of the Flask API for integration with the web platform.

RESULTS

This section presents the experimental results obtained by training and evaluating the four regression models across all five cities. Results are reported in two stages: first, baseline performance on the original dataset without augmentation, and then final performance after applying the Gaussian noise augmentation strategy.

Table 4 presents the performance of the models before applying data augmentation, serving as a baseline for comparison.

City	MAE (€)	RMSE (€)	R^2
Brasov	8 696	11 791	0.903
Craiova	9 656	20 681	0.878
Constanta	16 404	21 176	0.740
Timisoara	15 977	23 126	0.645
Bucharest	22 568	28 958	0.622

Table 4. Model performance before data augmentation

After applying Gaussian noise augmentation, all models showed significant improvements in predictive accuracy. Table 5 presents the final results obtained by the best-performing model per city.

City	MAE (€)	RMSE (€)	R^2
Brasov	5 145	7 812	0.948
Bucharest	8 908	14 362	0.925
Constanta	7 390	12 500	0.914
Craiova	7 222	11 350	0.885
Timisoara	9 839	14 821	0.839

Table 5. Final model performance after data augmentation




Figure 6. Median price heatmap by city and number of rooms

The most significant improvement was observed for Bucharest, where R^2 increased from 0.622 to 0.925 and MAE dropped from 22,568€ to 8,908€. This confirms that Gaussian noise augmentation is particularly effective for cities with limited training data. Constanta and Timisoara also benefited considerably, with R^2 gains of 0.174 and 0.194 respectively.

Across all cities, Brasov achieved the best overall performance ($R^2=0.948$, MAE=5,145€), likely due to the higher quality and consistency of the UNNPR source data compared to web-scraped listings. Bucharest, despite having the lowest R^2 among augmented models, still achieves strong predictive accuracy, with the higher MAE explained by the greater price dispersion in the capital's market.

Random Forest and XGBoost consistently outperformed HistGradientBoosting and MLP across most cities, confirming the suitability of ensemble tree-based methods for structured real estate data. The MLP model showed competitive results only for cities with larger datasets after augmentation, such as Timisoara and Constanta.

Impact of Room Count on Price

Analysis of price distribution by number of rooms reveals a consistent upward trend across all cities. In Bucharest, median prices range from approximately 75,000€ for studio apartments to over 300,000€ for five-room properties, representing the steepest price gradient among all cities. In Timișoara and Constanța, prices follow a similar progressive pattern, with five-room apartments converging around 250,000–260,000€. Craiova and Brașov present more moderate gradients, with Brașov remaining the most affordable market across all room categories. These findings confirm that room count is a relevant predictor for apartment price estimation across all cities analyzed. Room count is also strongly correlated with surface area ($r = 0.82-0.86$); both features were retained, as they contribute complementary information. Figure 6 summarises the median price per city and room count.




Figure 7. R² score comparison across all algorithms and cities

Algorithm Comparison

Figure 7 reports the R² obtained by each of the four algorithms in each city. Random Forest achieved the best or second-best R² score in four out of five cities, confirming its robustness for structured real estate data. XGBoost performed comparably to Random Forest in most cases, with particularly strong results in Constanța (R²=0.92) and Craiova (R²=0.88). HistGradientBoosting showed competitive performance in cities with larger datasets after augmentation, such as Timișoara and Constanța. The MLP model produced the weakest results overall, particularly in Bucharest (R²=0.68), where the limited original sample size and high price variability made it difficult for the neural network to generalize effectively even after augmentation.

Usability Evaluation

In addition to the predictive performance of the estimation models, the usability of the platform was assessed through a heuristic evaluation following Nielsen's ten usability heuristics [14, 15]. The main user flows – registration and authentication, browsing and filtering listings, creating and managing a listing, and requesting a price estimate – were inspected against each heuristic, and conformance was rated on a five-point scale (1 = not satisfied, 5 = fully satisfied). The results are summarized in Table 6. The inspection was performed by a single expert evaluator, the second author, who was not involved in the implementation of the platform; no inter-rater agreement could therefore be computed.

Nielsen heuristic	Score (1–5)	Main observation
Visibility of system status	5	Live price estimate and status messages
Match between system and the real world	5	Familiar Romanian real-estate terminology
User control and freedom	4	Listings and favorites can be edited or removed
Consistency and standards	5	Uniform Bootstrap-based layout across pages
Error prevention	3	Field validation present; few confirmation dialogs
Recognition rather than recall	4	Visible filters, dropdowns and saved favorites
Flexibility and efficiency of use	3	No shortcuts or saved searches for expert users
Aesthetic and minimalist design	4	Clean, uncluttered listing cards
Help users recognize and recover from errors	3	Error messages shown but partly generic
Help and documentation	3	Contact form available; no dedicated help section

Table 6. Heuristic evaluation of the platform based on Nielsen's usability principles

The platform obtained an average conformance score of 3.9 out of 5. The highest ratings were recorded for visibility of system status, match between the system and the real world, and consistency and standards, reflecting the real-time price feedback, the use of familiar Romanian real-estate terminology, and the uniform Bootstrap-based interface. The lowest scores corresponded to error prevention, flexibility and efficiency of use, error recovery, and help and documentation. These results indicate concrete directions for improvement, namely confirmation dialogs for destructive actions, keyboard shortcuts and saved searches for frequent users, more specific error messages, and a dedicated help section. Overall, the evaluation confirms that the platform provides a usable interface, while identifying a small set of targeted refinements for future iterations.

Limitations

Several limitations should be considered. The Gaussian noise augmentation was applied before the train/test partition, so synthetic samples derived from the same

original records may occur in both partitions; the figures in Table 5 should therefore be read as an upper bound, while Table 4 provides the conservative estimate. All results come from a single split, without repeated runs or significance testing, so small differences between algorithms are not statistically supported. The comparison also excludes LightGBM, CatBoost and a SHAP-based interpretability analysis, and the heuristic evaluation involved a single evaluator. Addressing these points is the priority for the extended version of this work.

CONCLUSION

This paper presented a machine learning-based approach for residential property price prediction across five Romanian cities. A dataset of 1,921 real estate listings was constructed from heterogeneous sources, including public notary reports and web-scraped listings, and enriched with geographic features computed via the OpenStreetMap API.

Exploratory data analysis revealed significant differences between the five markets in terms of price distribution, variability, and feature correlations, confirming that a per-city modeling approach is more appropriate than a single unified model. Gaussian noise augmentation proved effective in improving model performance, particularly for cities with limited training data such as Bucharest and Craiova, where R^2 scores increased by 0.303 and 0.007 respectively after augmentation.

Among the four algorithms evaluated, Random Forest and XGBoost consistently achieved the strongest results, with R^2 scores ranging from 0.839 to 0.948 across all cities. The best overall performance was obtained for Brasov ($R^2=0.948$, MAE=5,145€), benefiting from the higher quality of the UNNPR source data.

The trained models were integrated into a web-based real estate platform built with ASP.NET Core MVC and a Python Flask API, enabling automatic price estimation for newly listed properties. Future work will focus on expanding the dataset with additional cities and property types, incorporating neighborhood-level features, and exploring more advanced augmentation strategies to further improve prediction accuracy in data-scarce markets.

REFERENCES

1. Géron, A. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*. O'Reilly Media, 2019.
2. Breiman, L. Random Forests. *Machine Learning* 45, 1 (2001), 5–32.
3. Chen, T. and Guestrin, C. XGBoost: A Scalable Tree Boosting System. In *Proc. KDD 2016*, ACM Press (2016), 785–794.
4. Hastie, T., Tibshirani, R., and Friedman, J. *The Elements of Statistical Learning*. Springer, New York, 2009.
5. Čeh, M., Kilibarda, M., Lisec, A., and Bajat, B. Estimating the Performance of Random Forest versus Multiple Regression for Predicting Prices of the Apartments. *ISPRS International Journal of Geo-Information* 7, 5 (2018), 168.
6. Antipov, E. and Pokryshevskaya, E. Mass appraisal of residential apartments: An application of Random Forest for valuation. *Expert Systems with Applications* 39, 2 (2012), 1508–1516.
7. Truong, Q., Nguyen, M., Dang, H., and Mei, B. Housing Price Prediction via Improved Machine Learning Techniques. *Procedia Computer Science* 174 (2020), 433–442.
8. Rosen, S. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy* 82, 1 (1974), 34–55.
9. Bourassa, S., Cantoni, E., and Hoesli, M. Predicting House Prices with Spatial Dependence: A Comparison of Alternative Methods. *Journal of Real Estate Research* 32, 2 (2010), 139–160.
10. Kontrimas, V. and Verikas, A. The Mass Appraisal of the Real Estate by Computational Intelligence. *Applied Soft Computing* 11, 1 (2011), 443–448.
11. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems* 30 (NIPS 2017), 3146–3154.
12. Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., and Gulin, A. CatBoost: Unbiased Boosting with Categorical Features. In *Advances in Neural Information Processing Systems* 31 (NeurIPS 2018), 6639–6649.
13. Lundberg, S. M. and Lee, S.-I. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems* 30 (NIPS 2017), 4765–4774.
14. Nielsen, J. and Molich, R. Heuristic Evaluation of User Interfaces. In *Proc. CHI '90*, ACM Press (1990), 249–256.
15. Nielsen, J. *Usability Engineering*. Morgan Kaufmann, San Francisco, 1993.