www.degruyter.com/view/j/remav

sciendo

TRAINING AND INTERPRETING MACHINE LEARNING MODELS: APPLICATION IN PROPERTY TAX ASSESSMENT

Changro Lee
Department of Real Estate
Kangwon National University
e-mail: spatialstat@naver.com

Abstract

In contrast to the outstanding performance of the machine learning approach, its adoption in industry appears to be relatively slow compared to the speed of its proliferation in a variety of business sectors. The low interpretability of a black-box-type model, such as a machine learning-based valuation model, is one reason for this. In this study, house prices in Seoul and Jeollanam Province, South Korea, were estimated using a neural network, a representative model to implement machine learning, and we attempted to interpret the resultant price estimations using an interpretability tool called a partial dependence plot. Partial dependence analysis indicated that locally optimized valuation models should be designed to enhance valuation accuracy: a land-oriented model for Seoul and a building-focused model for the Jeollanam Province. The interpretable machine learning approach is expected to catalyze the adoption of machine learning in the industry, including property valuation.

Key words: *machine learning, interpretability, neural network, partial dependence plot, house valuation*.

JEL Classification:*C13, C45, L85*.

Citation: Lee, Ch. (2022). Training and interpreting machine learning models: application in property tax assessment. *Real Estate Management and Valuation, 30*(1), 13-22.

DOI: https://doi.org/10.2478/remav-2022-0002

1. Introduction

Despite the excellent performance and increasing popularity of machine learning approaches, they have not been adopted in the business industry as fast as anticipated. The reasons for this include a lack of infrastructure required for excessive computing and the complexity of the algorithms. The former can be overcome through the recognition of the importance of infrastructure and due investments by corporate executives. The latter may be harder to overcome, because it requires an understanding and consensus of the deployed algorithms from the general public and immediate stakeholders. People tend to be reluctant or even resistant to adopt new technologies unless they understand how they work. In addition to the fear of the unknown, safety can be another reason for slow adoption in a safety-critical environment, such as medical diagnosis. The development of interpretable machine learning can mitigate these concerns.

Property valuation is a popular application area in machine learning. Valuation plays a key role in various real estate fields, including brokering, the approval of collateral loans, property portfolio management, and tax assessment. Tax assessment is the task of estimating the value of properties to calculate their property tax, and must be performed in an accurate and transparent manner. Its price estimation result must be accurate; otherwise, it leads to significant inequality in the taxpayers' burden. It also must be transparent because most countries grant taxpayers the right to appeal the estimated price, and the tax administration is responsible for explaining the price estimation process to taxpayers. In other words, property tax assessment is the area that must adopt an advanced model

to enhance the valuation accuracy and should also guarantee the transparency of the model adopted. Hence, an interpretable machine learning model would be a desirable choice for modeling property valuation.

In this study, we attempt to estimate house prices in Seoul and Jeollanam Province, South Korea, using a neural network, a popular black-box-type model for machine learning. We then interpret the price estimation results using a tool called a partial dependence plot.

Numerous studies have attempted to estimate property prices using machine learning models, but most have focused on improving the predictive accuracy. They compared the performance between valuation models and concluded that the machine learning model employed in the study contributed to an increased accuracy of valuation. In this study, we investigate the impact of relevant input variables on house price estimation and attempt to gain insight into the property tax assessment model. The study results are expected to promote more interest in interpretable machine learning and to catalyze the rapid adoption of machine learning-based valuation models in the real estate industry.

This paper reviews the literature of property valuation methodology and interpretable machine learning in the second section. The third section explains the study area, data used, and neural network chosen for house price estimation. The estimation results, interpretation of the trained neural networks, and implications for valuation modeling are provided in the fourth section. Finally, a summary of the conclusion is presented.

2. Literature review

2.1. Emerging methodology for property tax assessment

Property tax assessment determines the value of a piece of real estate and is usually conducted by certified assessors and standardized valuation models. A valuation model for property tax relies on quality data and sound modeling, and it has always received special attention from policy makers and scholars because the low performance of property assessments indicates systematic inequality of tax burdens, which is institutionalized unfairness that cannot be overcome by individual taxpayers' complaints or appeals.

Therefore, it is not surprising that a variety of property valuation methods have been proposed in the literature to enhance price estimation. The traditional method for property valuation includes a regression model. Its output is the price of a property, and the inputs are land or building attributes, such as the lot size and property age. Although the regression model is widely used for property assessment in most countries (Benjamin et al., 2004), it has demonstrated several limitations, such as the difficulty in capturing non-linear relationships between the input and the output, strict assumption of data distribution normality, and non-identifiability of a proper price functional form (Chau & Chin, 2003).

Machine learning-based models have emerged in the valuation industry owing to the boom of the big data paradigm and the advancement in artificial intelligence technologies since 2010. They have demonstrated excellent performance in price estimation, and the representative models include random forests, support vector machines, gradient boosting machines, and neural networks (Graczyk et al., 2010; Lasota et al., 2011; Abidoye & Chan, 2017; Sandbhor & Chaphalkar, 2019; Talaga et al., 2019).

However, despite its brilliant success in accurate price estimation, the machine learning approach has not been widely accepted in tax administration until now. Because machine learning models are often characterized by complex algorithms and excessive computation costs, they are generally more difficult to understand compared to traditional models, such as regression models and decision trees. In general, the higher the accuracy of the model estimation, the more complex and more difficult it is to provide explanations. In contrast, the more a machine learning model needs to explain, the simpler it should become. This problem is known as the model accuracy–model interpretation trade-off (Falk, 2019).

This low interpretability of machine learning models serves as a barrier to their adoption in several industry practices (Doshi-Velez & Kim, 2017): people, in general, would not accept even a technically superior solution if it did not explain its estimation process to a reasonable degree. Non-interpretability of models creates an obstacle to industrial acceptance and raises concerns about safety in a high-risk environment (for example, a self-driving car), estimation bias from the training data (for example, racial bias in evaluating a loan applicant), and difficulty in auditing when they are deployed

in business operations. In summary, the non-interpretability of machine learning models is the predominant reason for latency in adopting the state-of-the-art approach in many industries, including property tax administration.

2.2. Interpretable machine learning

To mitigate these black-box model characteristics, a variety of tools for machine learning interpretability have been proposed in the literature. These tools aim to clearly explain how models predict and what the influential variables are. Interpretability tools can be classified into two broad types: model-specific interpretation tools and model-agnostic tools (Molnar, 2021). The interpretation of coefficients in a regression model is a model-specific interpretation because it only works to understand the regression model. In contrast, model-agnostic tools can be used in any machine learning model and are applied after the model has been trained (post hoc). These model-agnostic tools typically work by analyzing input and output pairs, for example, permuting input variables and observing changes in output results.

Representative model-agnostic tools include permutation feature importance (Fisher et al., 2018), example-based explanation, and feature importance visualization. Popular tools such as Shapley values (Štrumbelj & Kononenko, 2014; Sundararajan & Najmi, 2020) and Shapley additive explanations (Lundberg & Lee, 2017) can be classified as permutation-based interpretation methods. Counterfactual explanations (Wachter et al., 2017), adversarial examples (Goodfellow et al., 2014), and influential instances (Koh & Liang, 2017) belong to the category of example-based explanations. The partial dependence plot (PDP) is a well-known tool for feature importance visualization.

PDP demonstrates the marginal effect of an input variable on the prediction of a machine learning model (Friedman, 2001). PDP indicates whether the relationship between an input variable and the prediction is linear, monotonic, or more complex. The partial dependence function $\hat{f}_{x_S}$ is estimated by calculating the averages in the data as follows (Molnar, 2021):

$$\hat{f}_{x_S}(x_S) = \frac{1}{n}\sum_{i=1}^{n} \hat{f}\left(x_S, x_C^{(i)}\right) \quad (1)$$

The x_S denotes an input variable for which the partial dependence function should be plotted, and x_C denotes the other input variables used in the machine learning model $\hat{f}_{x_S}$. Generally, there are only one or two variables in the set S. The variables in S are those for which the effect on the prediction must be identified. The variables x_S and x_C are combined to make up the total input variable space x. Partial dependence works by marginalizing the machine learning prediction over the distribution of the variables in set C; thus, the function shows the relationship between the variables in set S (primary variables to be investigated) and the prediction.

The computation of PDP is easy to implement and its interpretation is also clear, which clearly shows how the average prediction in the dataset changes as the j-th input variable varies. In addition, PDP can be the best way to visually describe partial dependencies (Molnar, 2021). Presenting the analysis results in a tabular format would be a very ineffective way to convey explanations for certain variables. In such a case, the interpretation of the input variables is only possible and meaningful when they are visualized. The PDP is a curve that shows the input variable and the average predicted outcome. The PDP was chosen for the analysis in this study because of these advantages.

3. Neural networks for house price estimation

3.1. Study area and dataset

This study estimates house prices in Seoul and Jeollanam Province, South Korea. Seoul is the capital of South Korea with a population of 9.6 million as of 2021 (KOSTAT, 2021). Jeollanam Province is located at the southwestern tip of the Korean Peninsula, with a population of 1.8 million as of 2021 (KOSTAT, 2021). Seoul, one of the top global cities, is well known for its high property prices and is always highlighted by real estate media. In contrast, Jeollanam Province is characterized by large-scale plains along the rivers. It produces large amounts of agricultural produce, mainly rice, wheat, barley, and potatoes (Jeollanam Province, 2019). These two areas were chosen for the analysis because they are drastically different from each other in terms of regional landscapes, and their property price-forming mechanisms are expected to differ from each other.

The data were collected from a website managed by the Ministry of Land, Infrastructure, and Transport. It is a publicly available dataset that includes the home attributes of 22,213 sample houses in Seoul and 22,106 sample houses in Jeollanam Province. The attributes include site area, building area, and county code. It also includes house sales prices, which are used as a reference for various administrative purposes, such as property taxation and compensation in the eminent domain. It was updated on January 1, 2019, and Tables 1 and 2 present the descriptive statistics of the datasets for the major variables.

Table 1

Descriptive statistics of the 22,213 houses in Seoul (2019)

Variable	Min.	Mean	Median	Max.
House price (KRW in Million)	30	561	386	27,710
Site area (m²)	13	164	142	7,160
Building area (m²)	9	259	208	2,861
County code (25 levels)	#11290: 1,467 (6.6%), #11620: 1,293 (5.8%), #11380: 1,182 (5.3%), #11230: 1,170 (5.3%), #11440: 1,123 (5.1%), #11260: 1,111 (5.0%)			
Zone (5 levels)	Residence: 20,024 (90.1%), Quasi-residence: 988 (4.4%), Commerce: 602 (2.7%), Quasi-industry: 463 (2.1%), Green belt: 136 (0.6%)			
Building structure (6 levels)	Masonry frame: 15,781 (71.0%), Reinforced concrete frame: 5,636 (25.4%), Wooden frame: 739 (3.3%), Light steel frame: 26 (0.1%), Steel frame: 18 (0.1%), Log cabin: 13 (0.1%)			
Road width (5 levels)	Wide: 549 (2.5%), Medium A: 1,312 (5.9%), Medium B: 3,524 (15.9%), Narrow: 12,949 (58.3%), Vehicle inaccessible: 3,879 (17.5%)			

Note: Only the primary levels are shown for the county code for readability. The road width in Medium A is wider than that in Medium B.

Source: own study.

Table 2

Descriptive statistics of the 22,106 houses in Jeollanam Province (2019)

Variable	Min.	Mean	Median	Max.
House price (KRW in Million)	1	35	23	876
Site area (m²)	17	371	331	4,219
Building area (m²)	13	110	88	2,226
County code (22 levels)	#46130: 2,125 (9.6%), #46150: 1,743 (7.9%), #46770: 1,484 (6.7%), #46170: 1,447 (6.5%), #46820: 1,316 (6.0%), #46110: 1,190 (5.4%)			
Zone (10 levels)	Management: 11,506 (52.0%), Residence: 6,073 (27.5%), Natural & Green: 1,643 (7.4%), Commerce: 1,395 (6.3%), Agriculture: 739 (3.3%), Nature preservation: 316 (1.4%), Quasi-residence: 290 (1.3%), Industry: 52 (0.2%), Green belt: 48 (0.2%), Quasi-industry: 44 (0.2%)			
Building structure (6 levels)	Masonry frame: 11,657 (52.7%), Wooden frame: 7,454 (33.7%), Reinforced concrete frame: 1,773 (8.0%), Light steel frame: 888 (4.0%), Log cabin: 194 (0.9%), Steel frame: 140 (0.6%)			
Road width (5 levels)	Wide: 155 (0.7%), Medium A: 587 (2.7%), Medium B: 3,317 (15.0%), Narrow: 12,079 (54.6%), Vehicle inaccessible: 5,968 (27.0%)			

Note: Only the primary levels are shown for the county code for readability. The road width in Medium A is wider than that in Medium B.

Source: own study.

As shown in Table 1, the median house price in Seoul is 386 million KRW (approximately 350,000 USD), and most areas are zoned for residence (90.1%). The county code has 25 levels, which correspond to 25 local governments in Seoul. The primary building structure is a masonry frame (71.0%), followed by a reinforced concrete frame (25.4%). Table 2 shows the same information for the Jeollanam Province. The median house price is 23 million KRW (approximately 21,000 USD), and most areas are zoned for management (52.0%). The county code has 22 levels, which are equal to the number of county governments in Jeollanam Province. The primary building structure is a masonry frame (52.7%), followed by a wooden frame (33.7%). All continuous variables in both datasets, such as

house price and site (building) area, were scaled to have a mean of zero and a standard deviation of one before being fed into the neural networks.[1]

Figure 1 presents the examples of houses included in the dataset. Figure 1 (a) depicts a typical four-story house with a reinforced concrete framework in Seoul, and Figure (b) shows a representative single-story house with a wooden framework in Jeollanam Province.



(a) Four-story house in Seoul (b) Single-story house in Jeollanam Province

Fig. 1. Examples of houses included in the dataset. *Source*: own study.

Thirteen variables were used for house valuation: county code, site area, building area, zone, width (m) of the road on which the site abuts, building structure, building usage, site shape, neighborhood characteristics, floor area ratio (FAR), site coverage ratio (SCR), number of floors, and property age. The target variable is the house price, which was log-transformed before being standardized to alleviate a skewed distribution of prices.

House values were estimated using the well-established hedonic pricing framework (Rosen, 1974) as follows:

$$Price_i = County_i + SiteArea_i + BuildingArea_i + Zone_i + RoadWidth_i + BuildingStructure_i + BuildingUsage_i + Shape_i + Neighborhood_i + FAR_i + SCR_i + Floors_i + Age_i \quad (2)$$

Where $Price_i$ denotes the price of house i, and the aforementioned 13 explanatory variables are employed. The county, zone, building structure, and road width are categorical inputs with at least five levels, as shown in Tables 1 and 2. The building usage, shape, and neighborhood are categorical inputs with two, two, and three levels, respectively.[2]

3.2. Specifying and training neural networks

Network architecture is a topology that creates a network layer by layer. That is, an input layer is first created, hidden layers are added to it, and an output layer is added at the end of the architecture design. Categorical inputs deserve special attention when specifying network architecture. In general, categorical variables cannot be directly fed into a neural network, but must first be converted into numerical representations. The simplest technique is to convert each element in a categorical variable into a separate binary dummy variable, which is often called the one-hot encoding approach. However, this approach rapidly becomes inefficient or impossible to implement when the categorical variable is highly cardinal. In addition, it treats different values of categorical variables completely independently of each other and does not consider their informative interrelations.

An excellent alternative is the embedding technique. This technique is widely used in natural language processing because words can be considered as an agglomeration of vast amounts of categorical variables. Embedding is a technique for mapping categorical values to a multi-dimensional space with fewer dimensions than the original number of levels, where values with similar function

[1] The site and building area of houses were first log-transformed before being standardized to alleviate a skewed distribution of values.

[2] Levels in building usage: residential and mixed. Levels in shape: regular and irregular. Levels in neighborhood: residential, high-class residential, and commercial.

outputs are close to each other (Guo & Berkhahn, 2016). Adopting the embedding technique avoids excessive consumption of computational resources because of the one-hot encoding of high-cardinality variables, and it intrinsically groups similar values within categorical variables together, instead of processing them completely independently of each other. This study adopted the embedding technique for high-cardinality categorical variables to enhance the learning efficiency of neural networks.

A fully connected layer neural network with a sequential architecture was used in this study. The architecture is illustrated in Figure 2. Input layers comprising 13 input variables were created, and embedding layers for the four variables (county, zone, building structure, and road width) corresponding to highly cardinal categorical variables were additionally created and added to the architecture. A suitable number of dimensions had to be determined for each embedding layer, and the prediction performance for various dimension sizes was reviewed using the usual cross-validation process. The number of embedding dimensions for each categorical variable through cross-validation is shown in Figure 2. The other three categorical variables (building usage, shape, and neighborhood) were converted using traditional one-hot encoding because they had only two or three elements within them. Then, three hidden layers were added to the architecture to include more parameters to capture minor data nuances. These three dense layers after concatenation comprised 64, 32, and 16 neurons, respectively. The output layer was a linear layer with one neuron representing a house price.

The implementation details were as follows: the Adam optimizer and Glorot initialization with a uniform distribution were used. A learning rate of 0.01 was used, and a rectified linear unit (ReLU) activation function was used for all layers, except for a linear activation function that was used for the output layer. The mean squared error (MSE) was used as the loss function. The neural network was trained for 50 epochs[3], with a batch size of 128.

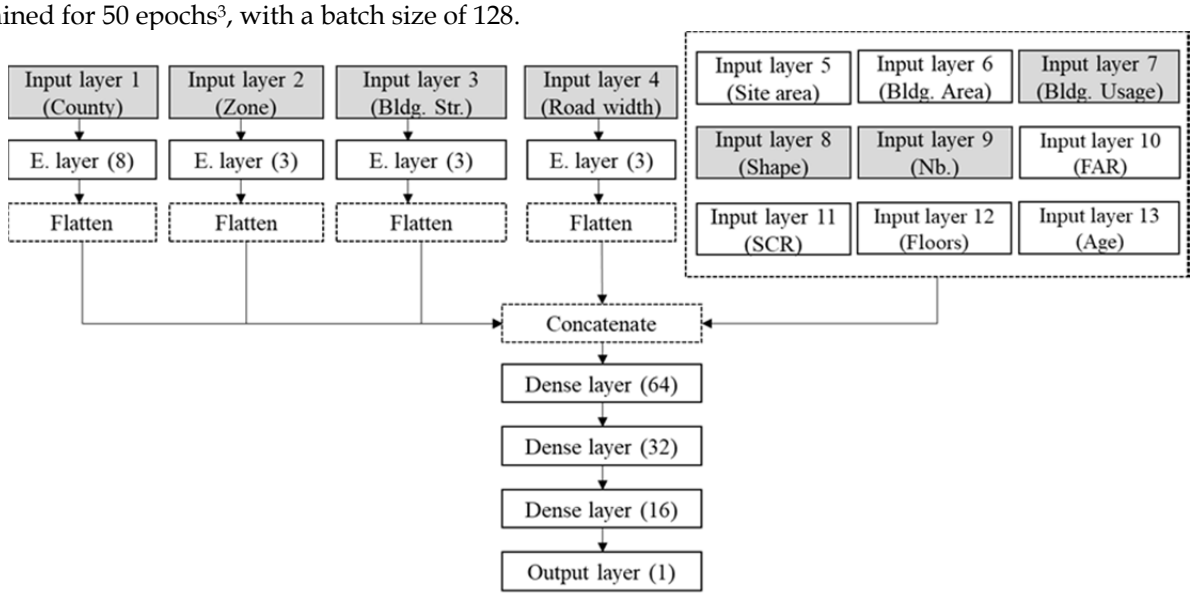

** Note: The shaded rectangles indicate categorical variables. The E. layer represents the embedding layer, and the numbers in parentheses denote the dimension sizes of each embedding layer. The numbers in the dense layers and output layer indicate the number of neurons. Bldg. Str, Bldg. Usage, Nb., FAR, and SCR stand for building structure, building usage, neighborhood, floor area ratio, and site coverage ratio, respectively.*

Fig. 2. Network architecture. *Source*: own study.

4. Results

4.1. Goodness-of-fit and PDP

The root-mean-square error (RMSE) and mean absolute percentage error (MAPE) were used to measure the goodness-of-fit as follows:

[3] The network appeared to converge sufficiently after approximately 50 epochs, and thus, this value was set for the number of training.

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{y} - y)^2} \quad (3)$$

$$MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{y-\hat{y}}{y}\right| \quad (4)$$

where$\hat{y}$ denotes the price predicted by the network, and y denotes the observed price.RMSE is a common measure of the prediction accuracy in statistics, and MAPE is a metric frequently reported by property valuation agencies (Ecker et al., 2020). MAPE indicates the prediction error in percentage and is useful for comparison across different valuation models.

The dataset (22,213 houses in Seoul and 22,106 houses in Jeollanam Province, as indicated in Tables 1 and 2) was randomly split into a training dataset (80%) and a test dataset (20%). Thus, 4,442 houses in Seoul and 4,421 houses in Jeollanam Province were reserved and used for evaluating the goodness-of-fit of neural networks. Rather than a single split, this split was repeated 100 times, because the evaluation result based on a single split cannot be representative of the true performance of the network. The network was fitted to the corresponding split dataset 100 times to provide a more comprehensive picture of the result. Table 3 presents the distribution of the predictive accuracy (RMSE and MAPE) based on these multiple splits. In Seoul, RMSE and MAPE show the range of 0.39–0.42 and 14–16, respectively. In Jeollanam Province, these metrics show higher values, indicating that the price modeling in a non-urban region can be more challenging.

Table 3

RMSE for house valuations in Seoul and Jeollanam Province

	Seoul	Jeollanam Province
RMSE	0.39–0.42	0.40–0.43
MAPE	14–16	15-18

Note: House prices in which RMSE and MAPE were measured were log-transformed and then standardized to have a mean of zero and a standard deviation of one.

Source: own study.

MAPE values near 18, as in the case of Jeollanam Province, can be considered inappropriate for practical business applications such as property tax assessment. However, most MAPEs officially reported by valuation agencies are obtained after excluding outliers. Thus, MAPE values in Jeollanam Province have potential to be reduced further if rare cases such as an extremely large-sized house were removed and appraised in a separate manner. In addition, the overall predicted prices appear to follow the observed prices reasonably well in the study areas, including Jeollanam Province, as indicated in Figure 3 (one example result of the 100 fittings). The figure presents the goodness-of-fit of the networks for the test dataset, and it can be concluded that there is no significant problem in interpreting the neural networks based on this result.

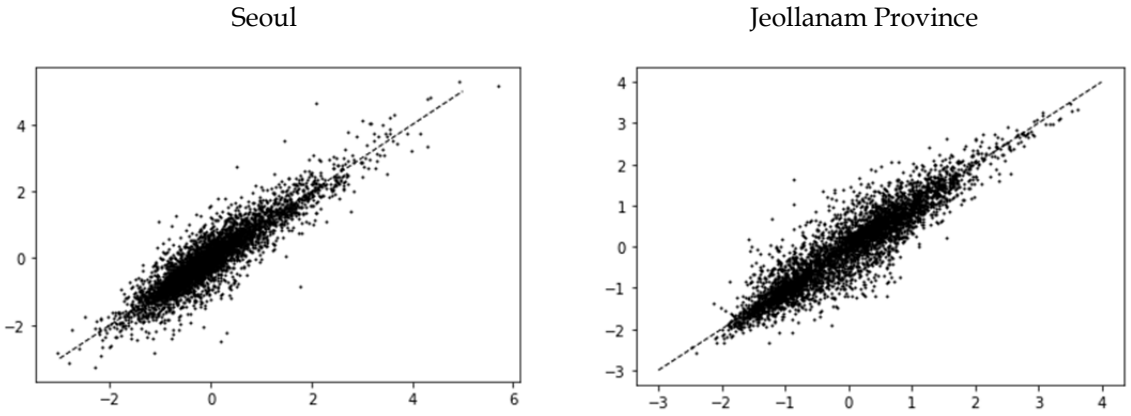

Fig. 3. Goodness-of-fit of neural networks. *Source*: own study.

Two variables of the 13 input variables produced noticeable patterns in the partial dependence analysis: site area and building area.[4] Figure 4 shows the two-way PDP for the site and building areas. The horizontal axis denotes the log-transformed site area and the vertical axis denotes the log-transformed building area. The figure shows an increase in house prices as the site area or building area increases. Rug plots were also presented along the two axes to visualize the data distribution for each variable.

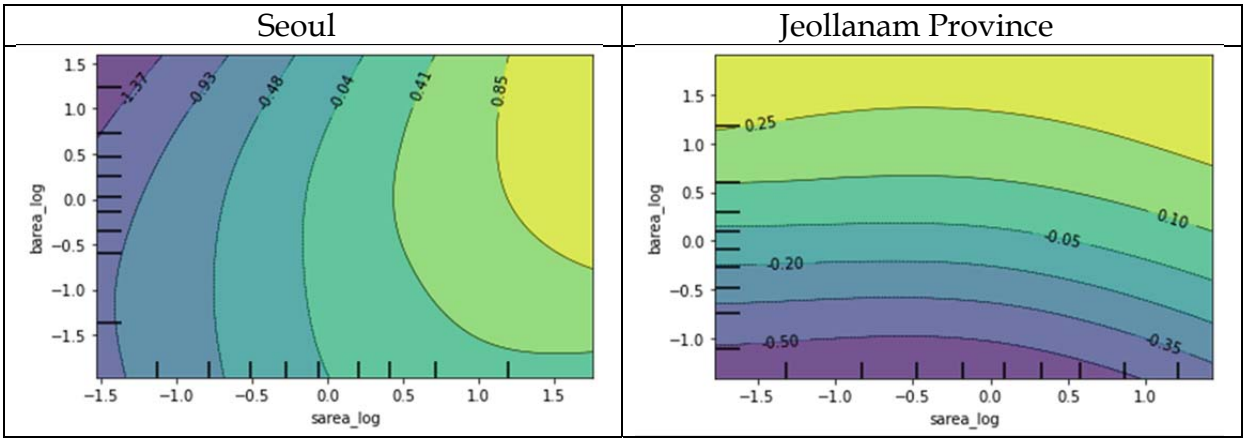

** Note: sarea_logandbarea_log indicate the log-transformed site area and log-transformed building area, respectively.*

Fig. 4. Two-way PDP for site and building areas. *Source*: own study.

As visually depicted in Figure 4, the estimated price patterns are distinct between the two regions. For Seoul, the house price is stratified by the magnitude of the site area, whereas it is stratified by the level of the building area for Jeollanam Province. Thus, the house price in Seoul appears to be determined largely by the site area, and the building area appears to have little impact on the price level. In contrast, the building area appears to be a major determinant of house prices in Jeollanam Province, and the site area does not play a significant role in estimating house prices.

Although a neural network itself is a black-box model, it is now partially possible to understand the process of price estimation in the neural network through a PDP. Both site and building areas were fed to the network as inputs for Seoul and Jeollanam Province, but the neural network learned that the site area was a primary factor for house prices in Seoul, whereas the building area was a major factor in Jeollanam Province. In other words, the neural network ignored the building area in Seoul and the site area in Jeollanam Province in estimating house prices after an appropriate period of training.

4.2. Implications for valuation modeling

A variety of valuation models are currently used for property tax assessment worldwide; however, most of these models have limitations in that the valuation is based on a single global model. That is, governments tend to employ a single model with the same input variables nationwide to estimate property prices for taxation.

However, the real estate market is always characterized as a localized market; thus, a corresponding local approach is desirable to improve the valuation accuracy. For example, most property statistics are produced at the local level instead of the global level: Case–Shiller home price indices are measured at the metropolitan area level (Lee & Park, 2022). Mortgage lenders and property owners use local property indices and local sales comparables to track asset values. Property prices in

[4] All combinations of continuous variables were examined for the analysis, and only this combination was presented in this study for the following reasons. First, the site and building areas are generally accepted primary factors for house valuation, and thus, this combination was analyzed intensively. Second, certain combinations showed common-sense pattern such as the equal contribution of the two variables concerned to the house price, which is not worthy of any further interpretation. Finally, other combinations did not show patterns that are available for interpretation.

the real estate market appreciate and depreciate in a substantially different manner depending on the region (Clapp, 2004).

As revealed in this study, the heterogeneity of price determinants across regions is apparent, and it is desirable that local valuation models be developed to reflect this regional variation. For example, site-related input variables should be further investigated and added to the model employed in Seoul. The following may be good input candidates: total site area, valid site area available for construction, and accessibility to infrastructural facilities such as roads and public open spaces. In a similar context, building-related input variables should be identified and incorporated into the model adopted in Jeollanam Province. The following may be promising input candidates: total building area, floor area for each story, and the workmanship quality of building interiors and exteriors. The tax assessment accuracy can be further enhanced by developing and deploying a locally optimized valuation model.

5. Conclusions

Property assessment for taxation has always been a major concern for policymakers and scholars because of its significant role in maintaining the equality of taxpayers' burden. Although machine learning-based valuations have received considerable attention since 2010, they have not been adopted in practice as rapidly as expected. This study attributed this to the low interpretability of machine learning models. We attempted to interpret the neural networks, the representative models to implement machine learning, by investigating the estimation results of house prices in Seoul and Jeollanam Province. We chose PDP as an interpretability tool for neural networks. By reviewing the results of the partial dependence analysis, we observed that major house price determinants vary across regions: Seoul needs to develop a more land-oriented valuation model, whereas Jeollanam Province has to use a more building-focused valuation model. Overall, the PDP analysis indicated that locally optimized valuation models should be created to enhance the predictive accuracy.

Most industries require interpretability when they deploy machine learning models in production, be it for legal reasons, owing to safety issues, or to gain insights into the underlying business. In this study, we gained insight into the house valuation model; a local variant of the standard valuation model should be developed to reflect regional heterogeneity.

The limitation of this study is that it did not attempt to utilize a variety of interpretability tools, such as Shapley values, counterfactual explanations, and adversarial examples. These interpretation methods can provide insights from different perspectives. Future studies are expected to employ methods other than that used in this study such that a richer explanation can be provided to the models concerned.

The interpretability tools for black-box models are expected to make machine learning attractive to organizations and people, particularly when explanation or transparency is demanded rigorously. Tax assessment is such a case, and interpretability tools such as PDP are expected to boost the adoption of machine learning in the industry, including property valuation.

6. References

Abidoye, R. B., & Chan, A. P. (2017). *Artificial neural network in property valuation: application framework and research trend*. Property Management.

Benjamin, J., Guttery, R., & Sirmans, C. F. (2004). Mass appraisal: An introduction to multiple regression analysis for real estate valuation. *Journal of Real Estate Practice and Education, 7*(1), 65–77. https://doi.org/10.1080/10835547.2004.12091602

Chau, K. W., & Chin, T. L. (2003). A critical review of literature on the hedonic price model. *International Journal for Housing Science and Its Applications, 27*(2), 145–165.

Clapp, J. M. (2004). A semiparametric method for estimating local house price indices. *Real Estate Economics, 32*(1), 127–160. https://doi.org/10.1111/j.1080-8620.2004.00086.x

Doshi-Velez, F., & Kim, B. (2017). *Towards a rigorous science of interpretable machine learning*. arXiv preprint arXiv:1702.08608.

Ecker, M., Isakson, H., & Kennedy, L. (2020). An Exposition of AVM Performance Metrics. *Journal of Real Estate Practice and Education, 22*(1), 22–39. https://doi.org/10.1080/15214842.2020.1757352

Falk, K. (2019). *Practical recommender systems*. Manning Publications.

Fisher, A., Rudin, C., & Dominici, F. (2018). *Model class reliance: Variable importance measures for any machine learning model class, from the "rashomon" perspective*. arXiv preprint arXiv:1801.01489, 68.

Goodfellow, I. J., Shlens, J., & Szegedy, C. (2014). *Explaining and harnessing adversarial examples.* arXiv preprint arXiv:1412.6572.

Graczyk, M., Lasota, T., Trawiński, B., & Trawiński, K. (2010). Comparison of bagging, boosting and stacking ensembles applied to real estate appraisal. In *Asian conference on intelligent information and database systems,* 340-350. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-12101-2_35

Guo, C., & Berkhahn, F. (2016). *Entity embeddings of categorical variables*. arXiv preprint arXiv:1604.06737.

Jeollnam Province. (2019). *Annual statistics in Jeollanam province*. Planning and Coordination Office of Jeollanam province.

Koh, P. W., & Liang, P. (2017). Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, 1885-1894. PMLR.

KOSTAT. (2021). *Population Census*. Korean Statistical Information Service.

Lasota, T., Łuczak, T., & Trawiński, B. (2011). Investigation of random subspace and random forest methods applied to property valuation data. In *International Conference on Computational Collective Intelligence*, 142-151. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-23935-9_14

Lee, C., & Park, K. K.-H. (2022). Forecasting trading volume in local housing markets through a time-series model and a deep learning algorithm. *Engineering, Construction, and Architectural Management*, *29*, 165–178. Advance online publication. https://doi.org/10.1108/ECAM-10-2020-0850

Lundberg, S., & Lee, S. I. (2017). *A unified approach to interpreting model predictions*. arXiv preprint arXiv:1705.07874.

Molnar, C. (2021). *Interpretable machine learning*. Lulu. com.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy, 82*(1), 34–55. https://doi.org/10.1086/260169

Sandbhor, S., & Chaphalkar, N. B. (2019). Impact of Outlier Detection on Neural Networks Based Property Value Prediction. In Information Systems Design and Intelligent Applications (pp. 481–495). Springer. https://doi.org/10.1007/978-981-13-3329-3_45

Štrumbelj, E., & Kononenko, I. (2014). Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems, 41*(3), 647–665. https://doi.org/10.1007/s10115-013-0679-x

Sundararajan, M., & Najmi, A. (2020). The many Shapley values for model explanation. In *International Conference on Machine Learning*, 9269-9278. PMLR.

Talaga, M., Piwowarczyk, M., Kutrzyński, M., Lasota, T., Telec, Z., & Trawiński, B. (2019). Apartment Valuation Models for a Big City Using Selected Spatial Attributes. In *International Conference on Computational Collective Intelligence*, 363-376. Springer, Cham. https://doi.org/10.1007/978-3-030-28377-3_30

Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech., 31*, 841. https://doi.org/10.2139/ssrn.3063289