Automated valuation models

Exploration of a machine learning approach

1 Introduction

Automated valuation models (AVMs) provide efficient means for local government to determine fair and equitable property taxes, for mortgage providers to limit risks, and for asset owners to make complex investment decisions.[1]

Traditionally, AVMs have been econometric models, such as linear regression models. However, recent advancements in the field of machine learning (ML) have opened up a new, and in many fields successful toolbox, providing additional methods for the same data, as well as approaches to access new sources of information and to create new variables.

An important distinction between traditionally applied methods and more recently introduced techniques lies in the structure definition of a model. Econometric models require a model specification – transformation of variables, selection of functional form, interaction effects, and distributional assumptions – prior to estimating parameter values, whereas most ML algorithms determine the model's structure and parameter values simultaneously (Athey, 2018).

This fundamental difference has theoretical consequences that are naturally reflected in practical applications. The main goal of this paper is therefore to discuss how ML algorithms compare to econometric models for residential real estate valuation in theory and to show what these theoretical differences mean in practice.

In Section 2 we discuss the position of ML algorithms within the landscape of AVMs by comparing econometric models and ML algorithms from a theoretical perspective. Section 3 shows two different ML applications within residential property valuation to highlight the advantages and disadvantages of ML algorithms. Finally, Section 4 concludes and provides routes for future research.

David Kroon and Marc Francke

2 Automated Valuation Models

Various parties, such as governments or mortgage providers, are interested in the value of a property at any given moment in time. An indication of a property's market value[2] is only revealed at the time of a transaction. AVMs aim to find the value of any property at any given moment in time, based on a (limited) set of transactions, in what is essentially a direct sales comparison approach. From a technical perspective, the valuation problem requires a (linear) regression model, in real estate finance and urban economics literature called a hedonic price model, see Malpezzi (2002) for an extensive overview of hedonic price models. At the core lies the assumption that a property's value can be broken down into the combined values of its characteristics.

Econometric models and ML algorithms can be used to provide solutions for this problem. Both approaches fit a function, such that we predict a property's value y on the basis of property characteristics X. The main distinction between the two is that econometric models require the user to specify in advance any interactions between characteristics' contributions to a property's value (Mullainathan and Spiess, 2017). ML algorithms aim at finding these relations empirically from the data.

In the next two paragraphs, we consider econometric models and ML algorithms more in-depth. We then point out some of the advantages and disadvantages that both types of methods introduce. While we discuss AVMs in the context of residential real estate, the differences between econometric and ML methods apply more generally.

2.1 Econometric models

Econometric models require an econometrician to define a model structure prior to parameter estimation, so its functional form, included features, and their transformations and interactions must be 'hardwired'. Practically, it is not feasible simply to try all possible transformations and combinations of features. Mathematically, it can even be impossible to do so, considering for example a linear regression model with more features than observations. Consequently, model structures are generally defined based on the econometrician's understanding of the modeled processes, and the model performance is compared for only a limited number of competing models.

Given a model's structure definition and the assumed statistical distribution of the error term[3], posterior distributions of the parameters β can be derived from observed data (sale prices y and features X). From the posterior distributions one can provide point estimates, such as expected or most likely values, typically denoted by $\hat{\beta}$, and corresponding credible intervals, for each individual parameter. It allows one to perform statistical inference and to make statements on the statistical significance of parameters, for example, what the effect is of wind turbines on the value of nearby properties (Dröes and Koster, 2016). As such, Mullainathan and Spiess (2017) argue that econometric models revolve around correctly estimating parameters β.

Predictive distributions of property values can be derived for individual properties. Expected or most likely values (predicted values), as well as credible intervals (depending on the features of the property), can be computed from these predictive distributions.

2.2 Machine learning algorithms

While econometric models focus on inference on the parameters of interest β, because the functional form and included features are predefined, ML algorithms simultaneously search for a functional form and parameter values. Please note that for most ML algorithms, functional forms and parameter values are not observed, or cannot (easily) be interpreted. Principally, the ML calibration procedure is driven by minimizing a distance measure between observed values y and predicted values $\hat{y}$, e.g. by minimizing the mean of absolute percentage errors (MAPE[4]) between predicted and observed values. Mullainathan and Spiess (2017) therefore argue that ML algorithms revolve around finding $\hat{y}$.

Due to their flexible calibration process, ML algorithms are able to adapt very well to a given data set, and are therefore prone to over-fitting. Adding regularization to a model can mitigate over-fitting by penalizing a model's complexity. However, complexity reduction might also prevent the algorithm from finding variable interactions that do require a high level of complexity. Domain knowledge is often required to successfully engineer features that contain all information, such that the algorithm can extract it.

It should be noted that ML algorithms do not operate completely without structure. Each algorithm requires a set of algorithm-specific hyperparameters that delimit the space within which it is free to calibrate a model. As these hyperparameters are specified by the user, they are usually 'tuned' to select the best-performing model.

Repeated calibrations for the same data set result in different model structures and parameter values. The meaning that can be ascribed to parameters in econometric models can therefore not be attached to parameters in ML algorithms.[5]
Apart from predicting property values, ML algorithms can be used to extract additional information from previously inaccessible sources, such as images and text. These features can

subsequently be used in a machine learning algorithm or econometric model, see Francke and Kroon (2021) for some examples.

2.3 Pros and cons

The nature of econometric models and ML algorithms affects which method is most appropriate under which circumstances. In this section, we highlight the strengths and weaknesses of both types of approaches in relation to property valuation, summarized in Table 2.1.

Variable dependency on space and time

Variables that determine property value can be grouped into property characteristics and spatial features. Both the features and their influence can vary over time and space. In the past decades, hedonic price models have been developed that explicitly handle space and time (see Gelfand et al., 1998; Pace et al., 1998; Anselin and Lozano-Gracia, 2009; Francke and van de Minne, 2020). However, their proper application requires in-depth statistical knowledge and expertise. In contrast, ML algorithms typically assume identically and independently distributed (i.i.d.) error terms (Varian, 2014). Spatial and temporal data violate this assumption. Solutions generally involve incorporating elements from econometric models, although recently ML algorithms have been applied to construct price indices (Dutra Calainho et al., 2020).

Ease of implementation

The process of defining a model's structure, estimating parameter values, and testing model assumptions, can be cumbersome and takes a high degree of expertise of both modeling and the modeled processes. Adding additional variables to an existing model requires repeating the process. ML algorithms demand no pre-defined model structure, except hyperparameter settings. Combined with efficient algorithmic implementations in publicly available software packages, they allow direct calibration of models and easy addition of variables.

Performance

Reliable AVMs predict with an average ratio of predictions to observed sale prices of one, with

TABLE 2.1 ▸ ECONOMETRIC VERSUS MACHINE LEARNING MODELS

Aspect	Econometric models	Machine learning models
Variable dependency on space and time	Tailored solutions exist, but do require in-depth statistical knowledge.	Dependency violates basic assumption of most 'off the shelf' algorithms. Solutions incorporate elements from econometric models.
Ease of implementation	The model estimation process can be cumbersome and requires expertise. Addition of variables involves repeating the process.	Most state-of-the-art algorithms are available in widely used software packages. Addition of variables takes no additional effort.
Performance	Advanced models, e.g. state-space models, are among the most accurate AVMs. Require a limited amount of data.	Off the shelf algorithms easily perform better than simple econometric models, and similar to advanced econometric models. Require relatively many data.
Explainability of predictions	Estimated parameters inherently offer tools to relate predicted values to property characteristics.	Model selection inconsistency prohibits attaching explanatory power to a model's parameters. Existing methods are complex and lack standardization.
Credible intervals	Standard tools exist to give credible intervals for individual predictions.	Existing methods are complex and lack standardization.

minimal dispersion around this average, and with equal accuracy for expensive and cheap properties.[6] ML algorithms can perform similar to advanced econometric models, and easily outperform simple econometric models, such as linear regression models. However, ML algorithms need many observations to learn a structure from the data. In Section 3.2.3 we analyze the impact of sample size on machine learning performance.

Explainability

In many applications it is important to be able to attribute differences in values between properties to differences in their characteristics. Because of their primary goal to estimate model parameters that quantify the predefined model structure, econometric models inherently provide the tools for explainability. On the other hand, ML algorithms' combined search for parameter values and model structure results in complex nonlinear models that vary for each repeated calibration of the model. A substantial body of research is concerned with explainable ML (see for example, Molnar, 2019; Arrieta et al., 2020), but the gap with econometric models remains in terms of complexity and standardization of methods.

Credible intervals

Besides a property value, often credible intervals are needed, depending on the specific property characteristics, for example for mortgage lending in case of a high loan-to-value ratio. For econometric models credible intervals can easily be calculated for both predicted values and parameters, while they are not standard output for ML algorithms, and require additional analysis.

3 Machine Learning for property valuation

After discussing advantages and disadvantages of ML algorithms compared to econometric models in the previous section, this section reflects on advantages and disadvantages of ML algorithms in practice, specifically in the context of valuating residential properties. We discuss two separate applications. Firstly, we discuss a hybrid model, in which we include a ML algorithm within an econometric model, combining the relative strengths of both model types, as described in Table 2.1. Secondly, we present the main results of a pilot study, where we use a stand-alone ML algorithm to test its flexibility in terms of adding variables, ease of implementation and performance, while attempting to mitigate the challenges regarding variable dependency on time.

3.1 Hybrid econometric-machine learning models

In this application, we combine the flexibility of ML algorithms to find complex relations between a property's characteristics with the ability of econometric models to model price trends and credible intervals.

The base of our AVM is a hedonic price model, specified as a so-called hierarchical trend model (HTM). The HTM relates the log transaction price $p_{i,t}$ of property i at time t (in months) to:

- The level of the common trend μ at time t in the housing market considered.
- The level of the market segment trend at time t, where the market segment trend is in deviation from the common trend. The market segment is for example defined by district and house type. The district and house type price trends are denoted by λ and θ.
- Property characteristics $x_{i,t}$. They enter the specification in a nonlinear way. Moreover, the corresponding coefficients β_t are allowed to vary over time. The impact of the property characteristics for property i and time t is denoted by $f(x_{i,t}, \beta_t)$.
- Time invariant neighborhood effect φ where neighborhood is more granular than district.

The HTM (Francke and De Vos, 2000; Francke and Vos, 2004; Francke, 2008) is provided by (3.1)

$$p_{i,t} = \mu_t + d^{\lambda}_{i,t}\lambda_t + d^{\theta}_{i,t}\theta_t + f(x_{i,t}, \beta_t) + d^{\phi}_{i,t}\phi + \varepsilon_{i,t} \quad (3.1)$$

where the d's are selection row vectors (containing zeros and a one) to choose per property the appropriate district, house type and neighborhood, and ε is an error term. The common trend μ_t is modeled as a local linear trend model, the market segment trends λ_t, θ_t and the coefficients of the property characteristics β_t as random walks and the neighborhood effects φ as random effects.

We have investigated replacing the a priori (rigid) specified function $f(x_{i,t}, \beta_t)$ by more flexible ML algorithms, namely a random forest and a neural network, based on sale prices in the period January 2009 to September 2016 in the Netherlands (for details we refer to Ceyhan, 2017). As such, the hybrid model combines the strengths of structured econometric models and flexible machine ML algorithms. It accurately deals with market segment specific price trends, and gives credible intervals for individual price predictions, which are both hard to do with only ML algorithms. It is also able to freely find the complex relations between the property characteristics (such as house and plot size, construction year and house type) and the price. Both ML variants proved to be an improvement over the functional form $f(x_{i,t}, \beta_t)$, as measured by the MAPE in an out-of-sample (a random 20% of the total sample) prediction test. The hybrid models with a random forest and a neural network yielded decreases in MAPE of 1.6% and 2.2%, respectively.

While the hybrid models outperformed the HTM in terms of prediction accuracy, their level of explainability decreased, a pitfall we described in Section 2.3. Especially on the individual level, the utilization of black-box models obscures the influence of the property characteristics on the transaction price. At the same time, the hybrid models did outperform a stand-alone random forest and neural network, showing the relevance of properly accounting for the market segment specific price trends.

3.2 Combining data sources in ML algorithms

In this application, we combine the flexibility of ML algorithms to find complex relations between a property's characteristics with the ability of econometric models to model price trends and credible intervals.

In contrast with the previous model, the application in this subsection concerns a stand-alone ML algorithm. Our main objective in this pilot study is to explore the main advantages of ML: The flexibility to add variables, and the prediction accuracy increase that additional variables might yield, see Section 2.3. Some variables have nationwide coverage, but most municipalities gather additional information that thus covers only a subset of transactions in our data set. In cooperation with three partnerships of municipalities (main cities: Amsterdam, Enschede and Oss, in total 25 municipalities), our investigation focused on integrating these different data sources in our ML algorithm.

A secondary objective is to appropriately consider the price trend, described in Table 2.1 as a challenge for ML algorithms.

3.2.1 Data

We made use of transaction data and property characteristics from Cadastre, Land Registry and Mapping Agency (Kadaster), and enriched these data with information from listing websites. The national data set includes sales in the period March 2012 up to February 2020. The base set of variables in our data set pertains to property characteristics, such as year of construction, plot and floor area, and house type.

We expanded this base data set with a wide range of publicly available resources, providing spatial features. Some of these are available from Statistics Netherlands (CBS) on a neighborhood level, whereas others are sourced from various maps that a range of Dutch government institutes make available. The resulting data sources include socio-economic features (such as population density), physical environmental features (such as levels of noise intensity related to airplanes, trains and cars) and functional environmental features (such as proximities to services and amenities).

Besides the aforementioned nationally available information sources, we have a number of sources at our disposal, supplied by the cooperating municipalities. The sample with extra information accounts for roughly 10% of all observations. The additional variables include maintenance condition, appearance and quality of the structure (ratings in five categories), the presence and size of annexes, garages, sheds, and so on.

3.2.2 Light Gradient Boosting Machine

To freely explore the possibilities of flexible specifications, we concentrated on a stand-alone ML algorithm. The applied one is the Light Gradient Boosting Machine (LightGBM), which is widely accepted by the ML community as a well-performing algorithm in terms of speed and accuracy for this type of data (Ke et al., 2017).[7] In short, LightGBM is an efficient implementation of the gradient boosting technique that combines a number of regression trees, such that each additional tree predicts the residual of the previous trees combined.

Given that the algorithm operates under the assumption of i.i.d. error terms, it does not incorporate temporal dependencies between sale prices. This necessitates the explicit modeling of a time trend component. As a simple solution, we model market segments trends by monthly averages and train the algorithm on corrected sale prices.

3.2.3 Model performance

National sample and a limited set of characteristics

The algorithm, trained on nationally available data, has an out-of-sample (a random 15% of the total sample) MAPE of 9.3%, see Figure 3.1 for the frequency distribution of the ratio of model value to sale price, and additional performance statistics. Leaving out any groups of features gives decreases in performance, most notably the physical and functional environmental features. Replacing neighborhood characteristics by neighborhood dummy variables, does not improve results, showing that most neighborhood-specific information is contained in our data set. An advantage of including neighborhood features is the potential of attributing feature importance more specifically.

Regional sample and an extensive set of characteristics

ML algorithms tend to perform better with larger cross-sectional data sets, which implies using the national sample. However, the additional information for a subset of municipalities, could distinguish properties that look identical considering only the nationally available information. Merging both sources would introduce many missing values, which are not straightforward to deal with in econometric models.

Harnessing the flexibility of ML algorithms, we have been able to leverage both the sample size of the national sample (more rows), and additional information (more columns) for a subset of municipalities, within a single model, see Figure 3.2.

FIGURE 3.1 ▸ OUT-OF-SAMPLE PERFORMANCE NATIONAL MODEL

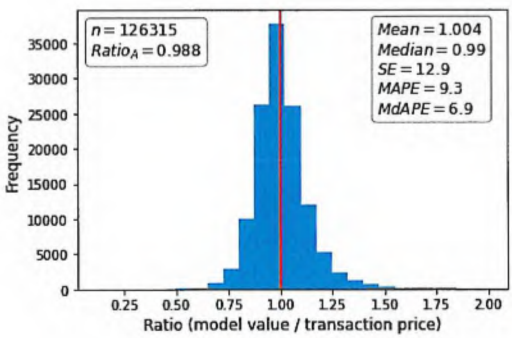

FIGURE 3.2 ▸ REGIONALLY EXTENSIVE VERSUS NATIONALLY LIMITED AVAILABLE DATA

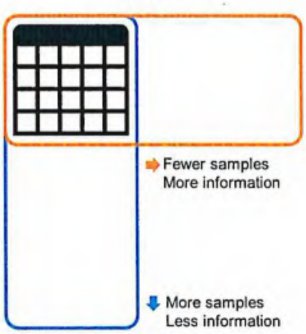

TABLE 3.1 ▸ OUT-OF-SAMPLE PERFORMANCE STATISTICS FOR VARIOUS IMPLEMENTATIONS OF THE ML ALGORITHM

Region	Training set	Extra variables	Ratios Average	Ratios Weighted	MAPE	MdAPE	SE
Amsterdam	Amsterdam	No	1,007	0,983	9	6,7	11,9
(n = 4817)	National	No	1,005	0,986	8,7	6,5	11,7
	Amsterdam	Yes	1,006	0,989	8,5	6,4	11,4
	National	Yes	1,004	0,986	8,3	6,2	11,2
Enschede	Enschede	No	1,006	0,985	10	7,6	13,4
(n = 3504)	National	No	1,006	0,986	9,5	7,2	13,2
	Enschede	Yes	1,005	0,985	9,5	7,3	12,8
	National	Yes	1,005	0,994	8,9	6,9	12,6
Oss	Oss	No	1,007	0,993	9,5	7	13
(n = 2061)	National	No	1,004	0,995	8,8	6,6	12,5
	Oss	Yes	1,004	0,992	8,8	6,4	12,1
	National	Yes	1,003	0,992	8,1	6,3	11,2

Notes: Statistics are provided on the ratio of the predicted value to the observed value. The weighted ratio is the sum of all predicted values by the sum of all observed values. MAPE stands for mean absolute percentage error, MdAPE for median absolute percentage error, and SE for standard error. The extra variables include maintenance condition, appearance and quality of the structure (ratings in five categories), the presence and size of annexes, garages, sheds, and so on.

Table 3.1 presents for the municipality partnerships with Amsterdam, Enschede, and Oss out-of-sample (a random 15% of the sample in these municipality partnerships) performance statistics: average ratios and dispersion measures. We evaluate the performance for different implementations of the ML algorithm: (i) a national or municipal training set, and (ii) without or with additional variables. A couple of conclusions can be drawn:

(i) All models have average ratios close to one, the weighted ratios are somewhat lower, around 0.99. This means that there is no systematic under- or overvaluation, and absence of inequities;

(ii) The ML algorithms trained on the municipal samples with the limited characteristics set perform worst, based on the dispersion statistics;

(iii) The ML algorithms trained on the national sample with the limited characteristics set perform better (compared to the previous one);

(iv) The ML algorithms trained on the municipal samples with the extended characteristics set perform better, in most measures (compared to the previous one);

(v) The ML algorithms trained on the national sample with the extended characteristics set perform best;

(vi) The added value of training on the national sample is the smallest for the largest city (MAPE in Amsterdam from 8.5 to 8.3), and the largest for the smallest city (MAPE in Oss from 8.8 to 8.1).

These results have great implications, namely that we are able to add any variable, regardless of whether the source is available for all observations within our data set. It implies that we can help clients improve the accuracy of their valuations by incorporating any proprietary data that they might have, and at the same time exploit information from the national sample.

Price trend

Additionally, we have evaluated the performance of the model, trained on the national sample, on a test set that is out-of-sample in time, by predicting next month's sale prices, on a rolling window basis, for the months March 2019 up to February 2020.[8] For example, we predict March 2019 sale prices using sales up to February 2019. The model performs as well for predictions in the test set consisting of a random 15% of the total sample, as it does for predictions for a month ahead in time, so outside the training set's sample period. This effectively shows that, for our purposes, we are able to mitigate the challenge listed in Table 2.1, by substituting a simple econometric model.

4 Conclusions and Future Research

The landscape of AVMs for residential properties has traditionally been dominated by econometric models. Developments in recent years in the field of machine learning have brought algorithmic modeling to the mainstream. ML algorithms are a valuable addition to the toolbox, due to the flexibility they offer in terms of finding relationships between variables, adding variables, ease of implementation, and even extracting information from previously inaccessible sources, such as text and images (Francke and Kroon, 2021).

From our research we have found to indeed be able to exploit the flexibility that algorithmic modeling offers:

- ML algorithms are relatively easy to implement, and require less statistical knowledge and experience compared to econometric analysis.
- The out-of-sample performance of machine learning algorithms is at least similar to the performance of advanced econometric models, and has the potential to improve further.
- ML algorithms can handle large numbers of variables, without overfitting the data.
- ML algorithms can easily deal with additional features for only a subset of observations.
- ML algorithms can be combined with econometric models to take at least partial and sometimes full advantage of both types of models.

ML algorithms are easy to apply compared to advanced econometric models, and most likely outperform econometric models in terms of prediction performance, when large numbers of observations and features are available. When numbers of observations are relatively small, econometric models, imposing a priori structure on the data, are likely beneficial for model fit.

At the same time, most stand-alone ML algorithms struggle to incorporate temporal and spatial dependencies, as methods are typically focused on cross-sectional data, assuming identically and independently distributed error terms. Moreover, the lack of explainability and the absence of credible intervals for individual predictions, are important obstacles to use ML algorithms as sole solution in real estate valuation. Perhaps unsurprisingly, explainable machine learning and the computation of credible intervals are the main topics of our ongoing research in the area of machine learning.

About the Authors

David Kroon MSc is research engineer at Ortec Finance, where he works on applying Machine Learning methods to financial modeling.

Prof. dr. Marc Francke is professor of real estate valuation at the University of Amsterdam and head of real estate research at Ortec Finance.

FOOTNOTES

1 An extended version of this article can be found in Francke and Kroon (2021).

2 The estimated amount for which the property should exchange on the date of valuation between a willing buyer and a willing seller in an arm's length transaction after proper marketing wherein the parties had each acted knowledgeably, prudently and without being under compulsion, TEGOVA (The European Group of Valuer Associations) Blue book.

3 And (uninformative) prior distributions for β. We discuss econometric models from a Bayesian perspective.

4 MAPE $= \frac{1}{n}\sum_{i=1}^{n} |\frac{M_i - P_i}{P_i}|$, where P is transaction price, M predicted value, and n the number of observations.

5 The lack of interpretation applies even more to ensemble ML methods, such as boosting and random forests.

6 See the Standard on Ratio Studies (2013), from the International Association of Assessing Officers.

7 See Ceyhan (2017) and Beimer and Francke (2019) for the application of other ML algorithms on sale prices in the Netherlands.

8 We correct for the general price change by the mean monthly sale price.

REFERENCES

- Anselin, L. and Lozano-Gracia, N. (2009). Spatial hedonic models. In Mills, T. C. and Patterson, K., editors, *Palgrave Handbook of Econometrics, Volume 2, Applied Econometrics*, pages 1213– 1250. Palgrave MacMillan.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, S., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., GilLopez, S., Molina, D., Benjamins, R., Chatila, R., and Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58:82 – 115.
- Athey, S. (2018). The impact of machine learning on economics. In *Economics of Artificial Intelligence*. University of Chicago Press.
- Beimer, J. and Francke, M. K. (2019). Out-of-sample house price prediction by hedonic price models and machine learning algorithms. *Real Estate Research Quarterly*, 2019(2):13–20.
- Ceyhan, C. (2017). Evaluation of machine learning techniques for house valuations. Technical report, Erasmus University Rotterdam, Econometrics & Management Science, Master thesis.
- Dröes, M. I. and Koster, H. R. A. (2016). Renewable energy and negative externalities: The effect of wind turbines on house prices. *Journal of Urban Economics*, 96:121–141.
- Dutra Calainho, F., van de Minne, A. M., and Francke, M. K. (2020). A machine learning approach to price indices: Applications in real estate. Technical report, SSRN: https://ssrn.com/abstract=3689632.
- Francke, M. K. (2008). The hierarchical trend model. In Kauko, T. and d'Amato, M., editors, *Mass Appraisal Methods. An International Perspective for Property Valuers*, pages 164–180. Wiley-Blackwell RICS Research.
- Francke, M. K. and De Vos, A. F. (2000). Efficient computation of hierarchical trends. *Journal of Business and Economic Statistics*, 18:51–57.
- Francke, M. K. and Kroon, D. P. (2021). Machine learning in real estate valuation. Technical report, Ortec Finance, R&D Labs.
- Francke, M. K. and van de Minne, A. M. (2020). Modeling unobserved heterogeneity in hedonic price models. *Real Estate Economics*, https://doi.org/10.1111/1540-6229.12320.
- Francke, M. K. and Vos, G. A. (2004). The hierarchical trend model for property valuation and local price indices. *Journal of Real Estate Finance and Economics*, 28:179–208.
- Gelfand, A. E., Ghosh, S. K., Knight, J. R., and Sirmans, C. F. (1998). Spatio-temporal modeling of residential sales data. *Journal of Business and Economic Statistics*, 16:312–321.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 30, pages 3146–3154. Curran Associates, Inc.
- Malpezzi, S. (2002). Hedonic pricing models and house price indexes: A select review. In Gibb, K. and O'Sullivan, A., editors, *Housing Economics and Public Policy: Essays in Honour of Duncan Maclennan*, pages 67–89. Oxford: Blackwell Publishing, U.K.
- Molnar, C. (2019). *Interpretable Machine Learning. A Guide for Making Black Box Models Explainable*. Leanpub, https://christophm.github.io/interpretable-ml-book/.

- Mullainathan, S. and Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2):87–106.
- Pace, R. K., Barry, R., Clapp, J. M., and Rodriquez, M. (1998). Spatiotemporal autoregressive models of neighborhood effects. *Journal of Real Estate Finance and Economics*, 17(1):15–33.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28.