Fifty shades of black: Uncovering physical models from symbolic regressions for scalable building heat dynamics identification

Julien Leprince
j.j.leprince@tue.nl
Technical University of Eindhoven
Eindhoven, Brabant, Netherlands

Clayton Miller
clayton@nus.edu.sg
National University of Singapore
Singapore, Singapore

Mario Frei
mario.frei@nus.edu.sg
National University of Singapore
Singapore, Singapore

Henrik Madsen
hmad@dtu.dk
Technical University of Denmark
Lyngby, Denmark

Wim Zeiler
w.zeiler@tue.nl
Technical University of Eindhoven
Eindhoven, Brabant, Netherlands

ABSTRACT

The rapid growth of machine learning (black-box) techniques and computing capacity has started to transform many research domains, including building performance analysis. However, physical interpretation of these models remains a challenge due to their opaque nature. This paper outlines an experiment to unveil analytical expressions from an open-source machine-learning-based algorithm, i.e., symbolic regression. From 241 residential buildings in the Netherlands, 50 unique analytical expressions were produced demonstrating overall better characterization accuracies than an XGBoost baseline, while providing a powerful mean of interpretability from model structures and coefficients. These insights present a starting point for further work towards highly scalable models yielding new characterizations of residential buildings.

CCS CONCEPTS

• Applied computing → Engineering;

KEYWORDS

Buildings, Automated Model Identification, Interpretable Black-box, Symbolic Regression

ACM Reference Format:

Julien Leprince, Clayton Miller, Mario Frei, Henrik Madsen, and Wim Zeiler. 2021. Fifty shades of black: Uncovering physical models from symbolic regressions for scalable building heat dynamics identification. In *The 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '21)*, November 17–18, 2021, Coimbra, Portugal. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3486611.3491120>

1 INTRODUCTION

Modelling building thermal dynamics is an important challenge in characterizing performance towards various objectives. With applications in building retrofitting [12], demand-side management

[6], energy forecasting [13], and model predictive control [5], it has been at the center of many research publications within the past decades. Despite its momentum, the approach is still faced with the fundamental challenge of scaling across the heterogeneous building stock. Thermal dynamics modelling fits in one of three well-established categories: physics-based methods (white-box), purely data driven (black-box) and hybrid approaches (grey-box) [3]. Physics and knowledge-based methods (white-box) are known to be time-consuming and difficult to scale up. With many parameters to fix and human expertise required, they are better fitted to detailed and isolated case-study building models. Grey-box models, on the other hand, work as a hybrid approach bridging the gap between physical and statistical modeling. By exploiting physical knowledge in their models, grey-box models profit from interpretability, while exploiting the particularities of case-study data information for parameter fitting which makes them good at generalization [1]. Finally, data driven models (black-box), encompass machine learning algorithms and statistical regressions, commonly fitted from input and output time-series data of the system. They are notoriously powerful at generalizing yet struggle to produce interpretable models. And while efforts in the domain have allowed the *opening of the box* through feature importance metrics, e.g., SHAP (SHapley Additive exPlanations) values [8], physical interpretation of captured models remains an existing gap in the field.

Symbolic regression was recently put to light as a powerful black-box approach for extracting analytical equations out of data. However, when dealing with high-dimensionalities, the exponential explosion of combinations make it poor at scaling. Established data-driven building heat dynamics model identification processes (grey-box) typically require as little as 4 dimensional data measurements; inside temperature, heat signal input and outside conditions with ambient temperature and solar global irradiance [2]. Although these approaches benefit from physical knowledge included within the developed models, it becomes interesting to explore how symbolic regressions could uncover new forms of building heat dynamics.

1.1 Opening the (black) box

This paper consequently proposes to group, categorize and analyze the analytical expression outputs of symbolic regression for building heat dynamics model identification. From a case study of 241 monitored Dutch residential buildings and exploiting the paved

This work is licensed under a Creative Commons Attribution International 4.0 License.
BuildSys '21, November 17–18, 2021, Coimbra, Portugal
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9114-6/21/11.
<https://doi.org/10.1145/3486611.3491120>

path provided by grey-box approaches, we propose to uncover the relationships driving inside temperature states through interpretable black-box-produced models. Prediction accuracies of the identified models are benchmarked against a commonly employed black-box regressor within the building sector, i.e., XGBoost, as well as against a naive predictor to confirm its relative performance.

2 SYMBOLIC REGRESSION

Symbolic regression is a machine learning algorithm based on genetic programming which uses a simple tree-like representation structures to build an analytical expression from given input data and mathematical operators [10]. By iteratively mutating, performing crossovers or replications of the tree branches, multiple analytical expressions are explored to determine the best fit to the given data. The procedure produces increasingly complex analytical expressions from the given input features to predict the target output. The equation with the largest fractional drop in error metric is selected as the best model [4].

Ultimately, symbolic regression derives explicit physical relations between components of a system in an automated way. Additionally by building the symbolic expression from a tree structure, increasing orders of complexities are explored as the model develops. Thus, the algorithm only incrementally incorporates features into the model, which allows inputs with no significant impact on the target output not to be considered upfront.

3 IMPLEMENTATION

Our study considers a total of 241 homes located in the Netherlands. Anonymized measurements of the inside temperature, heat signal control inputs are collected from smart thermostats provided by a local energy distributor company at a resolution of 15 minutes. Electric and gas-meter data are also available at resolutions of 10 seconds and 1 hour respectively. Weather data are collected from the publicly available Royal Netherlands Meteorological Institute (KNMI) weather stations [7] and are paired to each building thanks to an aggregated geo-localisation process.

We consider minimum measurement periods of two months and limit the maximum times-series length to 10'000 points, which corresponds to a period of approximately 3,5 months, which is the recommended, amply sufficient, maximum input data length for the Symbolic Regression. The measurement period ranges from February 1st to the end of May 2021, which comprises the end of winter season as well as a notably cold start of spring season at the beginning of April. Electric and gas-meter data are re-sampled by average to 15 minutes intervals to match the smart-thermostat information. Available data are then filtered to obtain the most recent continuous measurement period for each building. Cumulative missing values larger than two hours are imputed and smaller gaps are filled via moving average using an eight hours window size.

We employ the open source python library *PySR* developed by Miles Cranmer et al. [4] for the symbolic learning algorithm of this study. Mathematical operators considered encompass multiplication, addition, division, cosine and sine functions. The number of iterations, or generations the regression runs for, is set to 10 and no weights are assigned to input data or operators not to influence the knowledge discovery process. Prediction accuracy of

Figure 1: Mean Squared Error (MSE) distribution of the symbolic regression (SR) output analytical functions with their respective complexities compared to naive (naive) and XGBoost (XGB) regressors.

the regressor is benchmarked against a naive predictor as well as a gradient boosting regressor, i.e., XGBoost from the scikit-learn package [11], a commonly employed black-box regressor within the building sector. The naive regressor simply predicts the step-ahead inside temperature of the considered building to be the same as the last, providing a classic reference value to compare a regressor's performance to. The XGBoost regressor is trained from a classic 20-fold *TimeSeriesSplit* function of the *sklearn* python package using the same input data as fed to the symbolic regressor.

4 RESULTS

The performance of the symbolic regression is evaluated using the distribution of the Mean Squared Error (MSE) of the fitted models. Figure 1 presents the boxplot distribution of the symbolic regressor (SR) compared to a naive (*naive*) and gradient boosted (XGB) regressor. The obtained symbolic expressions present lower MSE central tendencies and spread compared to both benchmarks. Interestingly enough, it can also be noted that the simple naive regressor seems to produce lower MSEs overall than the gradient boosted method. This result might be due to the lack of parameter tuning for this method. A number of MSE outliers for both *naive* and XGB regressors are not represented given the upper limit of MSE axis that reaches as high as 0.85 Kelvins for XGB. The distribution of obtained symbolic expression complexities from the SR exposes a predominant presence of complexities of order five, which corresponds to a typical affine expression $a \cdot T_i + b$, where coefficients a and b , variable T_i and operators $\cdot$ and $+$ each add an order 1 of complexity to the overall expression.

Figure 2: Identified symbolic expression groups and their building characteristics meta-data distribution

Figure 3: Identified symbolic expressions and their coefficients boxplot distribution

The symbolic expressions derived from the black-box SR are post-processed to uncover 50 unique analytical expressions. These equations are then grouped into similar analytical expression families which are presented along with their group size and intra-family available attributes, i.e., building characteristics in Figure 2 and symbolic coefficient values in Figure 3. While examining Figure 2, it can be noticed that all but one family constitute linear polynomial expressions of order one and two. The exception here being the $T_i + \Phi_h \cdot \sin(f \cdot T_i)$ expression comprising either a *sine* or *cosine* function, where f is a coefficient and variables T_i and Φ_h represent inside temperature and space heating input signal respectively. The two preeminent analytical families, i.e., $T_i \cdot a + \Phi_h \cdot b + c$ and the simpler $T_i \cdot a + c$, where a , b and c represent affine coefficients, evoke simple first order regressions of the inside temperature considering, or not, the space heating input signal. The building characteristics distribution per analytical group seem to suggest smaller homes to be more frequent in $T_i \cdot a + V_g \cdot d + c$, where V_g represents gas-meter measurements, along with a larger proportion of family sizes of 3. The two largest family groups appear to be mainly composed of family sizes of 2 and 4, while home type distributions cover a crushing majority of town types. Given the number of buildings grouped

per analytical family along with their available meta-data, only the top 3 groups present results that can be considered significant.

Figure 3 allows us to dive into the fitted coefficient distributions per family group. Coefficients of value 0 or 1 have here been removed not to bias the appreciation of distributions. Inside temperature related coefficients a are commonly centered around 0.97, with a negatively skewed distribution, while coefficients c are spread between values of -0.1 and 2 with positively skewed distributions centered around 0.6. Overall, larger sized groups tend to show larger coefficient distributions.

5 DISCUSSION

We here discuss how the discovered findings might bring value to the building sector by (i) creating paradigm links between discovered black-box models and established grey-box ones and (ii) uncovering physical knowledge from identified models.

5.1 Link to grey-box paradigm

Thanks to the formulation of symbolic expressions describing dynamical systems from measurements, it naturally follows that parallels can be drawn to well-established grey-box models.

While grey-box necessitates the definition of multiple models for appropriate model selection, symbolic regression inherently iteratively builds the model, thus making it far better at generalization and automation. Commonly, multiple state point estimates are included in grey-box models, corresponding to up to fifth order models. The complexities captured by these models echo quite naturally with the different thermal inertiae interacting in buildings. While these models increase in complexity, their inherent interpretability related to building physics allows a direct evaluation of estimated parameters. In opposition, analytical functions discovered by *SR*, while being interpretable, necessitate physical analysis by domain experts which can scale poorly given the variety of identified functions. It follows that natural links between a buildings' thermal properties and identified *SR* coefficients can be drawn, building on the knowledge of grey-box models. Quite concretely, a typical grey-box building model can be represented by a lumped resistance-capacity models. The below differential equation represents a 1st order model,

$$dT_i = \frac{1}{R_{ia}C_i}(T_a - T_i)dt + \frac{1}{C_i}\eta_h\Phi_h dt + \frac{1}{C_i}A_w\Phi_s dt + \epsilon$$

where the state variables T and Φ represent temperature and heat flux, estimated parameters R , C , η and A serve as heat resistance, heat capacity, appliance efficiency, and area respectively and the subscripts i , a , h , s and w relate to inside, ambient, heat, solar and window components respectively. ϵ encapsulates the measurement error, model approximations and non-recognized or modeled phenomena [9].

Linking identified *SR* functions to this formulation uncovers physical components such as building heat capacity C_i , space heating appliance efficiency η_h and thermal gains ϵ here linked to coefficients a , b and c respectively. Coefficient g may also be associated with solar window area gains A_w . The below equation explicitly links the above grey-box model to identified polynomial relationships.

$$\begin{aligned} C_i \cdot T_i(t+1) &= C_i \cdot T_i(t) + \eta_h \cdot \Phi_h(t) + A_w \cdot \Phi_s(t) + \epsilon(t) \\ C_i \cdot T_i(t+1) &= a \cdot T_i(t) + b \cdot \Phi_h(t) + g \cdot \Phi_s(t) + c \end{aligned}$$

Confirming these physical links would however require knowledge of ground truth - an interesting area of research for future studies that could help uncover direct links between measurements and building characteristics.

5.2 Model discovery

Leveraging knowledge discovery can be used as a powerful tool to build new models, enhancing white- or grey-box model identification approaches. While all discovered polynomial models possess a linear simplicity that makes their interpretation accessible, some of the more complex or non-intuitive uncovered models might just be the starting point of a new generation of models. The cosine and sine function of the inside temperature identified in $T_i + E_h \cdot \sin(f \cdot T_i)$ may here correspond to particular cyclical control strategies of the thermostat. Additionally, while no significant meta-data factor seems to separate variations of the identified polynomial functions, inside temperature patterns emerging from thermostat control and occupant behavior heat gains might unveil these model structures. This requires deeper inspections outside the scope of this paper.

6 CONCLUSION

With this work, we bring to light an automated model identification approach of building heat dynamics from data. With 241 monitored buildings, fifty unique models were uncovered and grouped into seven main families of symbolic expressions, six of which are polynomials. These results support established differential models developed with grey-box approaches, while favoring simplified symbolic complexities. It brings important perspectives to model identification in practice, e.g., for forecasting and control applications. Discovered models and coefficients may be exploited in a variety of building service applications including automated and scalable model identification and calibration for building Model Predictive Control (MPC). Building performance analytics may also leverage such findings for building characteristics benchmarking or thermostat control strategy characterisation.

ACKNOWLEDGMENTS

This work is funded by the Dutch Research Council (NWO), in the context of the call for Energy System Integration & Big Data (ESI-bida). We gratefully acknowledge the support and contribution of Eneco, with particular thanks to Dr. Kaustav Basu and Rik van der Vlist, as well as Eneco customers to this research.

REFERENCES

- [1] Abdul Afram and Farrokh Janabi-Sharifi. 2014. Review of modeling methods for HVAC systems. *Applied Thermal Engineering* 67, 1-2 (2014), 507–519. <https://doi.org/10.1016/j.applthermaleng.2014.03.055>
- [2] Peder Bacher and Henrik Madsen. 2011. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings* 43, 7 (jul 2011), 1511–1522. <https://doi.org/10.1016/j.enbuild.2011.02.005>
- [3] Mathieu Bourdeau, Xiao qiang Zhai, Elyes Nefzaoui, Xiaofeng Guo, and Patrice Chatellier. 2019. Modeling and forecasting building energy consumption: A review of data-driven techniques. *Sustainable Cities and Society* 48 (2019), 101533. <https://doi.org/10.1016/j.scs.2019.101533>
- [4] Miles Cranmer, Alvaro Sanchez-Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. 2020. Discovering Symbolic Models from Deep Learning with Inductive Biases. *arXiv:2006.11287 [cs.LG]*
- [5] Ján Drgona, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L. Vrabie, and Lieke Helsen. 2020. All you need to know about model predictive control for buildings. *Annual Reviews in Control* 50 (2020), 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
- [6] Elisa Guelpa and Vittorio Verda. 2021. Demand response and other demand side management techniques for district heating: A review. *Energy* 219 (2021), 119440. <https://doi.org/10.1016/j.energy.2020.119440>
- [7] KNMI. 2021. Hourly values of weather stations. <https://www.daggegevens.knmi.nl/klimatologie/uurgegevens>
- [8] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st international conference on neural information processing systems*. 4768–4777.
- [9] H. Madsen and J. Holst. 1995. Estimation of continuous-time models for the heat dynamics of a building. *Energy and Buildings* 22 (1995), 67–79.
- [10] B. McKay, M.J. Willis, and G.W. Barton. 1995. Using a tree structured genetic algorithm to perform symbolic regression. In *First International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications*. 487–492. <https://doi.org/10.1049/cp:19951096>
- [11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [12] Mehrdad Rabani, Habtamu B. Madessa, and Natasa Nord. 2017. A state-of-art review of retrofit interventions in buildings towards nearly zero energy level. *Energy Procedia* 134 (2017), 317–326. <https://doi.org/10.1016/j.egypro.2017.09.534> Sustainability in Energy and Buildings 2017: Proceedings of the Ninth KES International Conference, Chania, Greece, 5-7 July 2017.
- [13] Yi Wang, Qixin Chen, Tao Hong, and C. Kang. 2019. Review of Smart Meter Data Analytics: Applications, Methodologies, and Challenges. *IEEE Transactions on Smart Grid* 10 (2019), 3125–3148.