Combining learning-based COP proxies and formal optimization for residential heat pump system operation

Pedro Girón[1]
Juan Rodriguez Santiago[1]
[1] Energy Economics and System Analysis
Fraunhofer IEE
Kassel, Germany
pedro.giron@iee.fraunhofer.de

Philipp Härtel[1,2]
Jan Dobschinski[1,2]
[2] Sustainable Electrical Energy Systems
University of Kassel
Kassel, Germany

***Abstract*— Residential heat pump systems are key in achieving climate policy objectives in the building sector. Due to the variability in technical installation criteria, manufacturer requirements, local regulations and the spatial or geographical restrictions associated with buildings, these heating systems are highly heterogeneous. The operating conditions measured under standardized settings do not always align with those observed in real-world systems. This poses challenges for operational planning tools that aim to coordinate the systems' power consumption and increasingly exploit their flexibility when interacting with markets or system operators' congestion management mechanisms. We propose a predict-then-optimize approach that combines learning-based COP proxies with model-based optimization methods to use real-world operational data when making prescriptive scheduling decisions.**

***Index Terms*-- hybrid models, optimization proxy, mixed integer programming, heating systems, heat pumps**

I. Introduction

In the last years, heat pump systems have gained importance in Germany. For example, in 2023, 65% of 96800 new residential buildings in Germany were primarily using heat pumps to cover the heat demand [1]. Given the rapid adoption of heat pumps in Germany's building sector, there is an increased need to develop forecasting models and optimization techniques to maintain high operational performance [2], [3]. These models aim to improve both the performance and the efficiency in diverse operating conditions, not only of the isolated heat pump in a residential context but also as part of an interconnected system that supplies heat from an electrical network.

Prescriptive approaches based on formal optimization methods are important as they can leverage formalized domain knowledge and provide a principled approach to problem modeling with convergence guarantees. For instance, they are crucial for assessing and extracting flexibility potentials. However, formal methods face challenges when nonlinear and nonconvex effects govern the optimal operation of a system and cannot be tackled by simplification, abstraction, or aggregation. For instance, operating residential heat pump systems at the individual household level entails modulated operating points and non-linear physics in the characteristic production curve, i.e., the coefficient of performance (COP).

Data-driven approaches, particularly those based on neural networks, have emerged as promising alternatives for dealing with the heterogeneity of operating conditions. By leveraging real-world datasets, these methods can model nonlinear COP behavior more accurately than purely model-based approaches [4]. Recent studies have demonstrated that learning-based control strategies can yield efficiency gains, with some reporting reductions in electricity consumption of up to 8% for domestic hot water supply [5]. Additionally, optimization frameworks have been employed to minimize electricity consumption while maintaining thermal comfort [6]. These findings highlight the potential of data-driven techniques in enhancing operational decision-making for heat pumps.

Our work integrates predictive modeling with formal model-based optimization in a unified predict-then-optimize framework. Specifically, we employ a multi-layer neural network to predict the COP under various operating conditions measured for a specific residential heat pump system. The neural network captures intricate dependencies between physical parameters and system performance, while the optimization model incorporates these predictions to improve its prescriptive decisions. To validate the above assumptions, the process has been validated across four distinct use cases. In summary our contributions are as follows:

- We learn and build a neural network model to predict the nonlinear COP relationship of a residential heat pump system based on real-world system measurement data.
- We develop a mixed-integer linear programming model to optimize the residential heat pump operational schedule, which integrates the learning-based prediction model as an explicit model constraint.

This work is sponsored by the Federal Ministry for Economic Affairs and Climate Action (BMWK) with the project FlexLearn (Funding Code: 03SF0756).

979-8-3315-1278-1/25/$31.00 ©2025 IEEE

- We assess the implications of extensive data and their effect on the optimization of electricity consumption, especially for operating conditions of dynamic temperatures and partial load of the heat pump in four different case studies.

II. Literature Review

A. Machine Learning Models for COP Prediction

Artificial Neural Networks (ANNs) offer significant advantages in predicting the COP of heat pumps, including the ability to model complex non-linear relationships without explicit formulation and adaptability to new input data such as outdoor ambient temperature (T_{out}) and supply temperature (T_{supp}), and the resulting COP or even more input parameters [7]. However, purely data-driven models do not consider empirical operational constraints, such as minimum compressor start-up settings, which can lead to unrealistic predictions. For example, an ANN model might predict an extremely high COP for a combination of temperatures that would, in practice, result in frequent compressor on-off cycling, significantly reducing the system's actual efficiency. To address these limitations, hybrid approaches exist that combine ANN predictions with physics-based knowledge of heat pump systems. [8] developed a model that incorporates operational constraints into the network architecture, [9] demonstrated that integrating compressor speed and operational state can significantly improve the model's ability to capture the system's real behavior.

B. Optimization Models for Heat Pump Operation

The classical approach to mathematical optimization of heat pump operation uses mixed-integer linear programming (MILP) models, especially for modeling the dynamics of the thermal storage system. A common approach is to formulate the problem as a cost minimization problem subject to technical constraints on thermal comfort, equipment capacities, and operational limits. These models often incorporate time-of-use electricity pricing and thermal demand forecasts to optimize the heat pump's operation schedule.

Models for heat pump operation have focused on addressing uncertainties and integrated advanced computational techniques by considering future predictions of weather, energy prices, and demand [10]. Stochastic optimization methods incorporate probabilistic scenarios to handle uncertainties in variables like weather and energy prices, as studies optimizing ground source heat pumps and decentralized energy systems [11], [12]. Multi-objective optimization frameworks balance energy efficiency, cost reduction, and thermal comfort, often using Pareto-optimal solutions for decision-making [13]. Reinforcement Learning (RL) has emerged as a promising alternative, offering model-free control strategies that adapt dynamically to changing conditions while achieving MPC-like performance [14]. Hybrid approaches that combine physics-based modeling with data-driven techniques further improve optimization. These developments present an alternative method that can leverage machine learning models to predict unknown parameters of an optimization problem, for example. Instead of relying on solving the optimization problem within a training loop, it is possible to learn optimal solutions directly from observable features using pretrained models [15]. In the following section a method following the previous paradigm will be proposed.

III. Proposed Model

The combination of prediction and optimization approaches using pre-trained models requires considering fixed parameters, which improves decision quality while working within the constraints of pre-trained architectures. Fig. 1 shows the basic neural network scheme and its integration into the optimization framework.

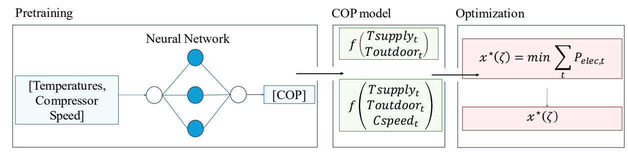

Figure 1. Neural network (pretrained) model integration as a restriction in the optimization program.

This approach focuses on adjusting specific components or introducing adaptable modules to enhance performance in different use cases or experiments presented in the next section.

A. Decision variable mapping and prediction framework

A multi-layer neural network is employed to pretrain the COP of the heat pump. The input layer is determined by the outdoor ambient temperature, supply temperature, and the compressor speed. We use two intermediate or hidden layers, whose size is varied to find the appropriate model structure. Fig. 2 shows the output layer that determines the COP.

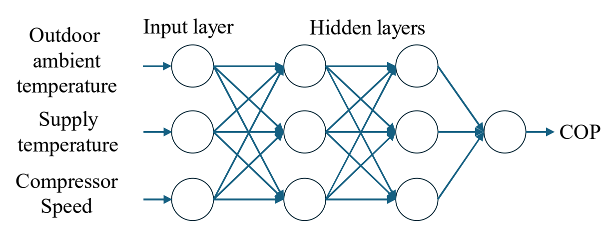

Figure 2. Standard feedforward neural network model for the correlation and data-driven prediction of the COP relationship.

For the calculation of each neuron's activation, the output processed by neuron j in each layer i is denoted as V in (1). The connections between neurons are mathematically represented by weights W. The output Y of the neuron in each layer is the sum of the outputs from layer i-1 multiplied by the respective weight connecting them to neuron ij. The respective threshold *b* is added to this sum [16].

$$V_{ij} = \sigma(Y_{ij}) = \sigma\left(\sum_k W_{ijk} V_{(i-1)k} + b_{ij}\right) \qquad (1)$$

We integrate this neuron architecture into a minimum cost-based operation optimization framework, presented below.

B. *Heat Pump Operation: Minimization of Energy Consumption Over Time*

The residential heat pump system under consideration includes the components shown in Fig. 3.

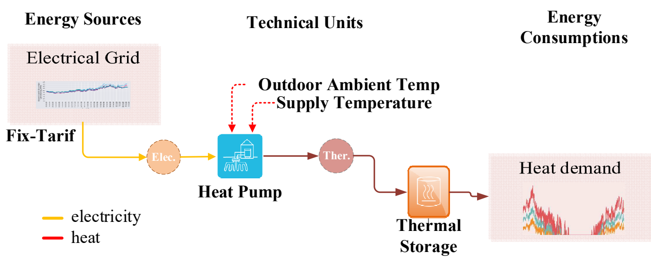

Figure 3. Interconection of a residential heating system.

Indoor comfort temperature is assumed to be maintained within the levels required by the ISO 7730:2005 [17] while optimizing electrical energy consumption P_{elec}. By adjusting the heat pump output P_{hp} as a function of the outdoor ambient temperature, the desired supply temperature and the current State of Charge (SOC) of the thermal storage to cover the thermal demand P_{demand} we formulate an optimization problem [18] for each time step t following the equation system (2) to (7)

$$min \sum_t P_{elec,t} \quad (2)$$

$$P_{hp,t} = COP_t \cdot P_{elec,t} \quad (3)$$

$$P_{out,t} \geq P_{demand,t} \quad (4)$$

$$SOC_t = SOC_{t-1} + P_{hp,t} - P_{out,t} \quad (5)$$

$$T_{min} \leq T_{supply,t} \leq T_{max} \quad (6)$$

$$COP_t = f(T_{supply,t}, T_{outdoor,t}) \quad (7)$$

It is important to note that forecast models are necessary for both supply and outdoor ambient temperature, as the COP depends on these variables [19]. Additionally, the heat demand must be accurately characterized using a robust model that effectively represents the thermal behavior of the studied buildings. A tool like TRNSYS, enables the simulation of heat demands under various operating conditions. [20].

For simplicity and understanding of the proposed optimization model, dynamics such as variable tariffs or the use of renewable energy with rooftop photovoltaic systems [21] is not considered.

IV. Experimental Setup

A. *Neural network generation and COP Model training*

Based on the measured data of outdoor ambient temperature, flow and return temperatures to the heat pump, thermal output and power consumption, and compressor speed, a COP identification is adjusted and eventually used by the optimization model.

Considering the measurements of outdoor ambient temperature and supply temperature, a first interdependence is observed in Fig. 4 illustrating their relationship for the considered system.

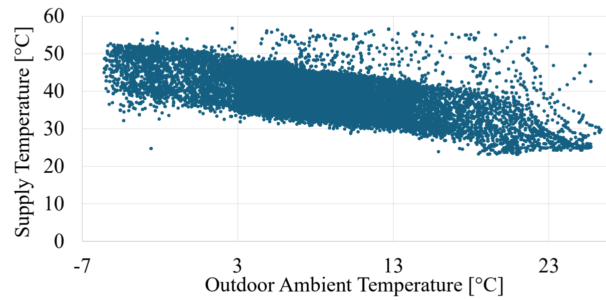

Figure 4. Outdoor temperature vs. supply temperature of a real heat pump.

The compressor speed and its various operational levels present another relevant dimension. Typically, compressor speeds can range from zero, with several intermediate levels and maximum speed depending on the requirements of the heat pump production. Fig 5 shows that for any given compressor speed, there can be different outdoor ambient temperatures.

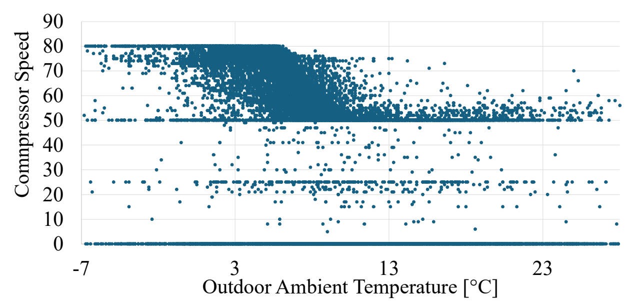

Figure 5. Outdoor ambient temperature versus compressor speed for a real heat pump, data of one year.

After defining the data for training the neural network, a preprocessing step is required. This involves a polynomial regression and percentile-based outlier detection. The pipeline includes an imputer for managing missing values and a polynomial feature transformer [22]. A polynomial regression model fitted on the transformed data [23] based on a multivariate regression calculates the intercept value (b_{cop}) and the coefficients CO and CS for the respective variables OT (outdoor ambient temperature) and ST (supply temperature) as follows (8):

$$COP_t(OT, ST) = b_{COP} + CO_{COP}OT_t + CS_{COP}ST_t \quad (8)$$

The Interquartile Range (IQR) method identifies outliers where data points falling outside the lower bound Q1 - 1.5*IQR and the upper bound Q3 + 1.5*IQR are considered anomalous [24], [25].

First, using the technical data sheet of a real heat pump, we calculate the polynomial function of (8) for two supply temperatures $COP_t(OT, 35)$ and $COP_t(OT, 55)$ to have a range of valid values, shown as lines in Fig. 6. Then we apply the IQR method explained in the previous paragraph to the COP_{real} data (green scatter plot) to obtain the COP_{clean} data (crimson red scatter plot)

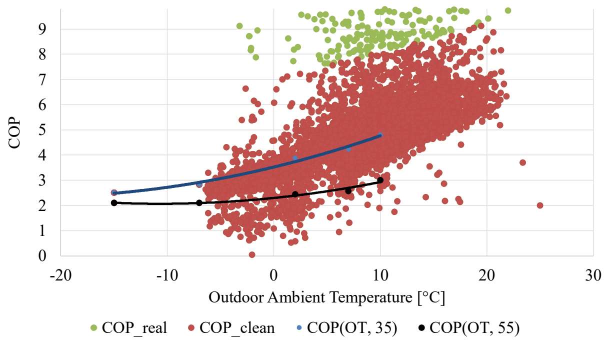

Figure 6. Polynomial regression filtering of COP combined with IQR method.

Now that COP data does not contain outliers, the training process begins with TensorFlow and Keras libraries on a python environment, defining a model that creates and trains a feedforward neural network model as described in Section II A. The model architecture consists of an input layer, two hidden layers with 10 neurons each using ReLU activation functions, and an output layer with a single neuron. Input data is normalized to improve model performance and convergence [26]. The Adam optimizer is used for training, which is an adaptive learning rate optimization algorithm that combines AdaGrad and RMSprop [27]. The model is trained for 200 epochs, which represents the number of complete passes through the training dataset [28]. The loss function [29] used is a mean squared error (MSE), calculated in (9) as the average of the squared differences between $COP_{predicted}$ and COP_{clean}.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{9}$$

B. Joint optimization considering the predictor

To consider the neural network model of the previous step, the process uses the functionality of Gurobi Machine Learning [30] [31] to formulate a trained machine learning predictor as constraint within a mathematical optimization model.

This additional restriction can be described as function $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ that can be a trained machine learning predictor, where *n* is the number of input features and *m* is the number of output variables, in our case, m = 1 as the output is only the COP. The function adds constraints to the equation system of the MILP that enforce the exact formulation of the pretrained model depending on the type of machine learning model used.

For neural networks, it represents each layer's operation, including activation functions, modifying the COP written as (7) by the set (10) to (12):

$$H_j^{(1)} = \sigma\left(\sum_{m=1}^{3} W_{2jm}x_m + b_{2j}\right) \tag{10}$$

$$H_k^{(2)} = \sigma\left(\sum_{j=1}^{n_2} W_{3kj}H_j^{(1)} + b_{3k}\right) \tag{11}$$

$$COP_{pred} = \sigma(Y_{41}) = \sigma_{out}\left(\sum_{k=1}^{n_3} W_{41k}H_k^{(2)} + b_{41}\right) \tag{12}$$

where x_m is the m-th feature of the input vector $\mathbb{x} = [x_1, x_2, \dots, x_n]$ and *n* is the total number of inputs: x_1= outdoor ambient temperature, x_2 = supply temperature, x_3= compressor speed level.

In (10), we take $H_j^{(1)}$ as the output (activation) of the first hidden layer after applying the activation function σ. $H_k^{(2)}$ in (11) is the output of the second hidden layer, which is calculated by applying σ to the linear combination of $H_j^{(1)}$. Finally (12), calculates the $COP_{predicted}$ as the output of the network, obtained after passing $H_k^{(2)}$ through the final transformation and output activation σ_{out}. The function returns an object containing information about the added constraints [32], allowing us to use the predictor inside the optimization model as:

$$COP_{pred} = COP_t = f\left(T_{supply,t}, T_{outdoor,t}, C_{speed,t}\right) \tag{13}$$

C. Operational planning problem description

With the pretrained model, the optimization algorithm can now be executed.

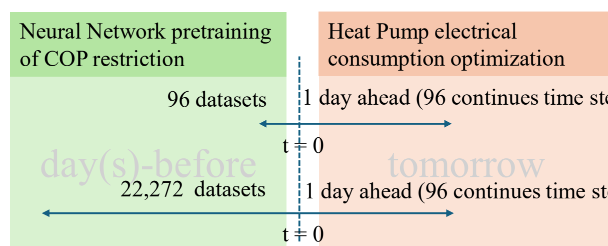

Figure 7. Connection of the COP pretrained model and electrical consumption optimization for the next day.

The study cases consider the difference between large amounts of data to train the ANN model as shown in Fig. 7, but also the impact of incorporating a third input to predict the COP. *Use cases 1* and *2* have 96 datasets as input from real measurements (emulating a day before data-reading); in case 1, the training inputs are only the temperatures, while case 2 includes the compressor speed. *Use cases 3* and *4* extend the training period to 22,272 datasets (emulating 232 days at 15-minute intervals), incorporating the same input parameters as cases 1 and 2, respectively.

All scenarios employ (2) to (6) and (13) to formulate the optimization problem as a MILP for the subsequent 96 time steps to simulate the operation for the next day. The extended training periods in *Use Cases 3* and *4* may lead to improved model accuracy and potentially better optimization results, i.e., electrical consumption minimization.

The MILP is implemented in Python 3.12 using the latest version of the Gurobi solver (12.0). The model is solved on a 12th Gen Intel(R) Core (TM) i7-1260P 2.10 GHz with 48,0 GB RAM, and Windows 11 based x64 processor.

V. RESULTS AND DISCUSSION

For the *Use Case 1* and *2,* the dataset was divided into 96 "day-before" training samples as historical dataset and 96 "tomorrow" samples as forecasted dataset, emulating 1 day in advance planning, in both cases of outdoor ambient temperature, supply temperature and compressor speed, having t=0 the 4th of December 2023 at 12:45:00.

The new 96 datasets are used by the optimization model and to train again the COP model but this time as a restriction embedded in the equation system. As shown in Fig. 8 the prediction obtained within the optimization in *Use Case 1* doesn't closely match the real behavior of the COP, using 10 neurons per layer and 2 hidden layers. The result of *Use Case 2* shows a closer representation when considering the compressor's on and off cycles.

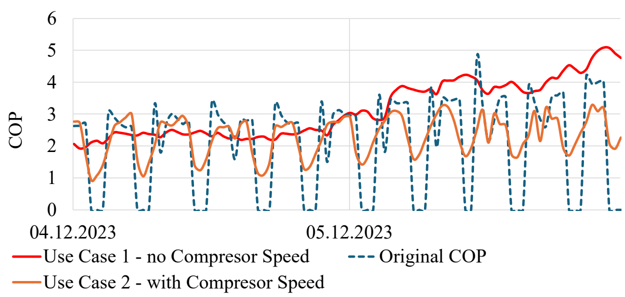

Figure 8. COP predictor as restriction of the optimization problem for a small dataset.

This approximation causes an overstimation in the electrical consumption, as seen in the Fig. 9, of almost 50% more in the worst result (*Use Case 1*)

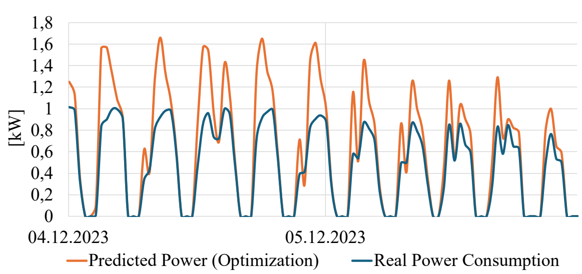

Figure 9. Optimization result of the electrical power consumption for the next day using a small dataset in *Use Case 1*.

For the *Use Cases 3* and *4*, the number of training samples (in this example, emulating 232 "days-before" with 15-minute intervals, resulting in 22,272 values per input) significantly improves the results, as it allows the model to capture drops to zero due to heat pump shutdowns.

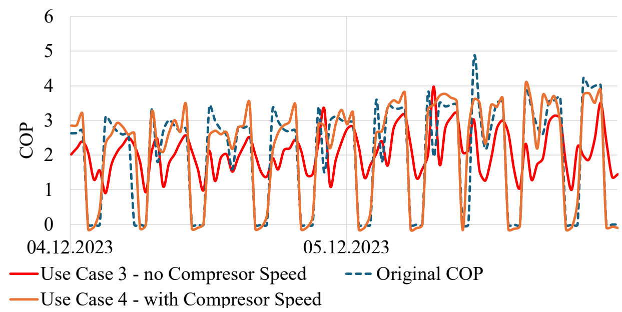

Figure 10. COP predictor considering an extended training of 22,272 dataset as restriction of the optimization problem.

Increasing the number of neurons per layer for the extended training dataset did not improve performance. Instead, negative COP values were observed for *Use Case 4* that considers the compressor speed, along with a longer convergence time as seen in Table 1.

TABLE 1. Converge time and MSE for different neurons number

Predictor	**Convergence Time (s) of the training**		**MSE**
	10 neurons per layer	*200 neurons per layer*	*10 neurons per layer*
with compressor speed	220.41	312.59	0.831
no compressor speed	230.36	287.84	2.202

VI. CONCLUSION

In conclusion, the study shows that incorporating additional variables, such as compressor speed, significantly improves the model's accuracy. Although increasing the training data volume leads to a consequent decrease in prediction error, the system representation reveals that it is crucial to identify and filter the training data carefully. This process should consider factors such as the operating points where heat is produced or the start/stop ramps of electricity consumption that can introduce outliers in the COP model. Moreover, incorporating additional decisions, such as operational ranges as binary variables in the optimization model, could be interesting to reduce errors, especially if negative values of the COP predictor are observed.

ACKNOWLEDGMENT

This research was partially supported by national funds in Germany through the Projektträger Jülich (PtJ) on behalf of the Federal Ministry for Economic Affairs and Climate Action (BMWK) within the framework of the project Verallgemeinerte Prognose von Lasten und Flexibilitätspotenzialen für den Betrieb von Wärmepumpen durch Transfer- und Multitasking-Lernen (FlexLearn). The authors gratefully acknowledge Gurobi Optimization LLC for their Gurobi Optimizer, which was instrumental in solving the optimization problems presented in this study.

REFERENCES

[1] Federal Statistical Office (Destatis), "Three in four new German buildings approved in 2023 will come with heat pump," Clean Energy Wire, 2024.

[2] Dobschinski J., Jost D., Giron P.J., Goldmaier A-K., Hollermann D.E., Rodriguez Santiago J., “Predictability of the operating behaviour of different types of heat pump systems," in Proc. 23rd Wind and Solar Integration Workshop (WIW 2024), Helsinki, Finland, Oct. 8-11, 2024. doi: https://doi.org/10.1049/icp.2024.3837

[3] M. Braun, C. Gruhl, C. A. Hans, P. Härtel, C. Scholz, B. Sick, F. Steinke, O. Stursberg, S. Wende-von Berg, “Predictions and Decision Making for Resilient Intelligent Sustainable Energy Systems,” in *2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, IEEE, pp. 1–5. doi: 10.1109/ISGTEUROPE62998.2024.10863480

[4] Fraunhofer ISE, "AI-Controlled Heat Pumps Increase Efficiency," Press Release, 2024.

[5] M. Rink et al., "More heat pumps for apartment buildings," energieforschung.de, 2024.

[6] D. Fischer et al., "Demonstration of Optimal Scheduling for a Building Heat Pump System," Energies, vol. 14, no. 23, p. 7953, 2021.

[7] A. Puttige, S. Stark, and S. Herkel, "Optimal Selection of Features for Heat Pump Models based on Artificial Neural Networks," in Proceedings of Building Simulation 2023: 18th Conference of IBPSA, Shanghai, China, 2023, pp. 1145-1152

[8] X. Luo, Y. Xia, and J. Wang, "A review on hybrid physics and data-driven modeling methods for building thermal dynamics," Renewable and Sustainable Energy Reviews, vol. 185, p. 113960, Feb. 2024.

[9] M. Thalfeldt, J. Kurnitski, and E. Latõšov, "Application of Regression and ANN Models for Heat Pumps with Field Measurements," Energies, vol. 14, no. 6, p. 1750, Mar. 2021.

[10] T. Rohrer et al., "Deep Reinforcement Learning for Heat Pump Control," arXiv:2212.12716,

[11] Z. Zhao et al., "A Stochastic Optimization Model for a Ground Source Heat Pump System with Uncertainty Quantifications," in Stanford Geothermal Workshop, Feb. 2024

[12] H. Schwarz et al., "Two-stage stochastic, large-scale optimization of a decentralized energy system – A residential quarter as case study," Karlsruhe Institute of Technology (KIT), 2024.

[13] Testi et al., "Optimal integration of electrically driven heat pumps within hybrid distributed energy systems," Stochastic Energy Systems Review, 2020.

[14] Operating heat pumps in residential complexes via reinforcement learning," PV Magazine, Jul. 2024

[15] A. Ferber, B. Wilder, B. Dilkina, and M. Tambe, "Learning Joint Models of Prediction and Optimization," arXiv preprint arXiv:2409.04898

[16] J. Hilera and V. Martínez, Redes neuronales artificiales: fundamentos, modelos y aplicaciones, no. January. 1995

[17] ISO 7730:2005, Ergonomics of the thermal environment — Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local thermal comfort criteria," International Organization for Standardization, Geneva, Switzerland, 2005.

[18] Arteconi, A., Hewitt, N. J., & Polonara, F. (2013). "Domestic demand-side management (DSM): Role of heat pumps and thermal energy storage (TES) systems." *Applied Thermal Engineering, 51(1-2)*, 155–165.

[19] Péan, T. Q., Salom, J., & Costa-Castelló, R. (2019). "Review of control strategies for improving the energy flexibility provided by heat pump systems in buildings." *Journal of Building Performance Simulation, 12(1)*, 1–22.

[20] Q. Si, Y. Peng, Q. Jin, Y. Li, and H. Cai, "Multi-Objective Optimization Research on the Integration of Renewable Energy HVAC Systems Based on TRNSYS," Buildings, vol. 13, no. 12, p. 3057, 2023. [Online]. doi: https://doi.org/10.3390/buildings13123057

[21] Fischer, D., Bernhardt, J., Madani, H., & Wittwer, C. (2017). Comparison of control approaches for variable speed air source heat pumps considering time variable electricity prices and PV. Applied Energy, 204, 93-105.

[22] A. Vaswani et al., "Attention Is All You Need," arXiv:1706.03762 [cs], Dec. 2017. [Online]. Available: https://arxiv.org/abs/2109.07127

[23] V. Kotu and B. Deshpande, "Outlier Detection Methods in Machine Learning," Towards Data Science, Sep. 15, 2021. [Online]. Available: https://towardsdatascience.com/outlier-detection-methods-in-machine-learning-1c8b7cca6cb8/

[24] S. Sharma, "How to Detect Outliers in Machine Learning," freeCodeCamp, Mar. 15, 2023. [Online]. Available: https://www.freecodecamp.org/news/how-to-detect-outliers-in-machine-learning/

[25] A. Sharma, "Outlier Detection and Treatment," Towards AI, Jan. 18, 2023. [Online]. Available: https://pub.towardsai.net/outlier-detection-and-treatment-9a9f41df0fb2

[26] S. Horvath, A. Langfelder, M. Kwak, and A. Raj, "Artificial neural network, predictor variables and sensitivity threshold lead to accurate age prediction of human tissue methylation by elastic net regression," *Scientific Reports*, vol. 11, no. 1, Jan. 2021. [Online]. Available: https://www.nature.com/articles/s41598-021-81556-2

[27] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proceedings of the 3rd International Conference on Learning Representations (ICLR), San Diego, CA, USA, 2015.

[28] S. Raschka and V. Mirjalili, "Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow 2," 3rd ed. Birmingham, UK: Packt Publishing, 2019.

[29] S. J. Russell and P. Norvig, "Artificial Intelligence: A Modern Approach," 4th ed. Hoboken, NJ: Pearson, 2021.

[30] Gurobi Optimization, LLC, "Gurobi Machine Learning Manual," 2023. [Online]. Available: https://gurobi-machinelearning.readthedocs.io/

[31] Gurobi Optimization, LLC, "Using Trained Machine Learning Predictors in Gurobi," 2023. [Online]. Available: https://cdn.gurobi.com/wp-content/uploads/Using-Trained-Machine-Learning-Predictors-in-Gurobi-slides.pdf

[32] Y. Zhu et al., "PySCIPOpt-ML: Embedding trained machine learning models into mixed-integer programming," arXiv preprint arXiv:2312.08074