

Artificial Intelligence–Based Analysis of Solar Shading

Amorotti Lucas
Retailleau Mathieu
Aurelian Andrei Panait

November 2025-February 2026



Supervised by
Ayoub Hannad
PhD

Abstract

Predicting building energy consumption traditionally relies on detailed thermal simulations, which are time consuming and require extensive modelling effort. This work evaluates the ability of artificial neural networks to predict energy performance while explicitly considering solar shading effects caused by neighbouring structures. A baseline dataset is generated for a reference building without shading. Additional datasets introduce multiple shading configurations. Two predictive models are trained and assessed using RMSE, CV(RMSE) and R^2 , and their predictions are compared with dynamic thermal simulations. A renovation scenario is finally tested to quantify whether AI models maintain accuracy when building characteristics are modified. Results demonstrate the influence of shading on prediction accuracy and discuss the capability of machine learning models to integrate physical context.

Contents

1	Introduction	3
2	Methodology	3
2.1	Reference building and simulation dataset	3
2.1.1	Baseline dataset without shading	3
2.1.2	Shading dataset	3
2.1.3	Renovation scenario	4
2.2	Preprocessing data	5
2.2.1	Transformation of the data	5
2.2.2	Training, validation and evaluation metrics	5
2.3	Network architecture and hyperparameters	5
3	Results	5
3.1	Hyperparameter search summary	5
3.2	Models performance	6
3.2.1	Baseline model (without shading inputs)	6
3.2.2	Shading-aware model	6
3.2.3	Quantitative performance metrics	6
3.2.4	Impact of shading scenarios	7
3.2.5	Model interpretability	7
3.2.6	Performance under renovation conditions	7
4	Discussion	8
4.1	Influence of shading on learning and prediction	8
4.2	Strengths and limitations of the proposed approach	8
4.3	Applicability for early design and renovation analysis	8
4.4	Sources of error and possible improvements	9
5	Conclusion	9
5.1	Key findings	9
5.2	Role of shading in predictive energy modelling	9
5.3	Future work	9

1 Introduction

The building sector accounts for a significant share of global energy use. Improving its efficiency is essential for reducing environmental impact and achieving climate objectives. Dynamic thermal simulation tools remain the main reference for performance assessment, but they require detailed physical modelling and long computation times. In contrast, machine learning approaches offer faster predictions once trained, provided that the dataset captures the relevant physical behaviour.

Solar shading is one of the parameters that strongly affects energy demand, especially in dense urban environments. Shading modifies solar gains, daylight availability and therefore heating and cooling loads. Despite its impact, shading is often poorly represented in predictive models. This study focuses on analysing how the explicit inclusion of shading parameters affects the performance of neural network models trained to predict building energy consumption.

2 Methodology

2.1 Reference building and simulation dataset

2.1.1 Baseline dataset without shading

Reference building geometry and parameters The reference building was modelled using the thermal simulation software *Pléiades*. A four storey building template with nearly identical floors was provided and slightly adapted for this study.

Figures 1 and 2 show respectively the 3D external view of the building model and the 2D ground floor plan.

Each floor consists of two rows of five adjacent rooms, each with an individual floor area of 17.04 m². The two rows are separated by a 1 m wide central corridor. The building is strictly oriented towards the south. Windows are installed on all four façades, ensuring exterior exposure for every room.




Figure 1: 3D view of the reference building model




Figure 2: Ground floor plan of the reference building

Dataset generation process To construct the baseline dataset, ten physical and operational variables were randomly sampled within predefined ranges. These variables fall into three categories: ventilation parameters, insulation properties and glazing characteristics.

For each sampled configuration, a full year dynamic thermal simulation was carried out using the AMAPOLA engine integrated in Pléiades.

The dataset was generated in three equally sized subsets:

- **First subset:** parameter values were sampled using a uniform distribution through Latin Hypercube Sampling, ensuring broad coverage of the input space.
- **Second subset:** insulation was set to zero for all simulations, while the remaining variables were sampled normally.
- **Third subset:** insulation was set to zero and ventilation was fixed at its maximum value.

Following this procedure, a total of 2000 simulations were obtained. This baseline dataset represents the heating demand of the reference building under conditions where no external shading is present.

2.1.2 Shading dataset

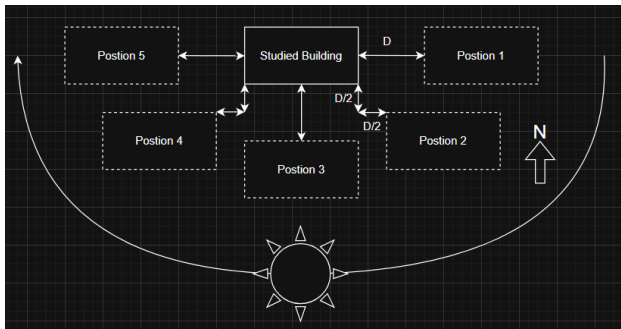
Definition of shading scenarios To investigate the impact of neighbouring structures on energy needs, a parametric study was conducted involving a single masking obstacle. This obstacle is modelled as a rectangular cuboid with identical dimensions to the reference building (20 × 10 m footprint) and a fixed height $H = 11.5$ m.

The position of the obstacle is defined relative to the reference building using two variables: the azimuth angle and the separation distance D .


- **Azimuth:** Five positions are considered relative to the South: -90° (West), -45° (South-West), 0° (South), $+45^\circ$ (South-East), and $+90^\circ$ (East).

- **Distance (D):** Three wall-to-wall distances are tested: 5 m, 10 m, and 15 m.


For the cardinal directions, D represents the perpendicular distance. For the diagonal orientations ($\pm 45^\circ$), a Manhattan distance logic is applied (offset by $D/2$ on both axes).




(a) Top-down view of the 5 angular positions



(b) Distance definitions (D) and geometry



(c) East, $D = 5$ m



(d) South-West, $D = 15$ m

Figure 3: Geometric configuration of the shading scenarios. The shading structure ($H = 11.5$ m) is identical to the reference building.

This results in a total of 15 distinct shading scenarios (5 positions \times 3 distances).

Dataset expansion strategy A specific sampling strategy was employed to isolate the effect of shading from the influence of building envelope parameters. Instead of generating new random building configurations for the shading scenarios, the exact same 2000 building configurations generated for the baseline dataset were reused.

For each of the 15 shading scenarios, the 2000 simulations were re-run with the added obstacle. This "cloning" approach ensures that any variation in energy performance is strictly attributable to the shading effect, as ventilation rates, insulation thicknesses (parameters E2 to E10), and glazing properties (E11, E19) remain strictly identical between the reference and the shaded cases.

Consequently, the total dataset comprises 32,000 simulations:

- 2,000 baseline simulations (no shading).
- 30,000 shaded simulations (15 scenarios \times 2,000 buildings).

The train/test split established in the baseline phase (1600 training / 400 testing) is preserved across all scenarios. The model is trained and tested on the same respective building IDs to prevent data leakage.

Simulation output The target variable is the annual heating load [kWh] required to maintain an indoor set-point temperature of 21°C . This value is normalized by the total floor area (754.8 m^2) to obtain kWh/m^2 . The masks are opaque and oriented at 0° (South), significantly impacting direct solar radiation.

2.1.3 Renovation scenario


Envelope and window upgrade measures The renovation scenario is introduced as a post-training robustness test. To keep the analysis consistent with the learning framework, the comparison is performed strictly through the model inputs. All geometric descriptors and window-to-wall ratios are kept unchanged, while envelope insulation and glazing properties are modified to represent an upgraded building.

Table 1: Renovation scenario: modified model input variables (before / after).

Parameters	Before	After
E2: Internal insulation thickness (external walls)	0.0	2.0
E3: External insulation thickness (external walls)	15.0	10.0
E6: External insulation thickness (staircase walls)	0.0	0.010
E7: Internal insulation thickness (roof)	0.0	2.0
E8: External insulation thickness (roof)	15.0	15.0
E10: Insulation thickness (ground floor)	15.0	2.0
E11: Window thermal conductance (North)	2.606	1.095
E19: Solar factor of glazing (North)	0.502	0.533
E31: Air renewal rate (typical floors)	1	1
E32: Air renewal rate (ground floor / corridor)	1	1
G4: Ground-floor heating enabled	1.0	1.0

The renovation scenario mainly affects envelope insulation and glazing-related inputs, while geometric, ventilation and shading parameters remain unchanged. Internal insulation layers are introduced for external walls and the roof, while external wall insulation is reduced by approximately 33%. The most pronounced change concerns ground floor insulation (E10), which decreases by nearly

87%. Window thermal conductance (E11) is reduced by about 58%, reflecting a substantial improvement in glazing performance. The solar factor of the glazing (E19) slightly increases by approximately 6%, modifying solar gains. Overall, the renovation induces large relative variations in a limited set of model input variables, providing a meaningful robustness test for the prediction models.



2.2 Preprocessing data

2.2.1 Transformation of the data

Before training the neural network, the raw dataset was preprocessed using a unified transformation pipeline. The input variables include numerical, categorical and binary descriptors of the building configuration. To ensure consistent scaling and a stable learning process, the data were processed through a `ColumnTransformer` composed of four elements:

- **Indicator transformation:** applied to the eight glazing and window-ratio parameters (*Ratio fenêtre Nord, Sud, Est, Ouest* and their corresponding ground-floor counterparts). This transformation generates binary or threshold-based indicators that help the model capture abrupt changes in façade configuration.
- **Standardization:** all remaining numerical variables were normalized using a z-score transformation to ensure a balanced contribution during optimization.
- **One-hot encoding:** the categorical variable *Type de mur* was encoded without introducing artificial ordinal relationships.
- **Passthrough:** the binary variable *Chauffage RDC* was preserved unchanged, as no scaling was required.

2.2.2 Training, validation and evaluation metrics

Each candidate MLP architecture is trained for up to 100 epochs with validation-based early stopping.

After tuning, the best model is evaluated on the test set using the root mean squared error (RMSE) and coefficient of determination (R^2).

These metrics provide a comprehensive evaluation of model accuracy and generalization ability in the no-shading scenario.

2.3 Network architecture and hyperparameters

The predictive model is a Multilayer Perceptron (MLP), whose architecture is automatically tuned using `Keras Tuner` via a `RandomSearch` strategy.

The hyperparameter search explores:

- number of hidden layers: 1 to 10,
- neurons per layer: 20 to 500,
- activation function: SELU,
- kernel initializer: `lecun_normal`,
- learning rate: 10^{-4} to $3 \cdot 10^{-2}$,
- optimizer: stochastic gradient descent with momentum.

3 Results

3.1 Hyperparameter search summary

MLP without shading inputs	MLP with shading inputs
<i>Best configuration (random search):</i>	<i>Best configuration (random search):</i>
Hidden layers 9	Hidden layers 1
Neurons/layer 80	Neurons/layer 20
Learning rate 0.001	Learning rate 0.01
Optimizer Adam	Optimizer Adam

Figure 4: Best hyperparameter configurations selected by random search for the two input scenarios. The shading-aware optimum is reported; the no-shading optimum will be added once the tuning run completes.

The random hyperparameter search yields a compact configuration for the shading-aware model, consisting of a single hidden layer with 20 neurons, trained with the Adam optimizer and a learning rate of 0.01. This outcome suggests that, once shading information is explicitly provided at the input level, the mapping from inputs to annual heating demand can be captured with limited model complexity. The same search procedure is applied

to the no-shading case; its best configuration will be reported once the tuning run is completed, enabling a direct comparison of how the optimal capacity depends on the chosen input space.

3.2 Models performance

This section compares the predictive performance of the two proposed models: (i) a baseline MLP trained without shading-related inputs, and (ii) a shading-aware MLP explicitly incorporating geometric shading parameters. Model performance is assessed using graphical indicators and standard regression metrics computed on an independent test set.

3.2.1 Baseline model (without shading inputs)

The baseline model is first evaluated to establish a reference level of performance. Figure 5 compares predicted and simulated annual heating demand on the test dataset.

Predictions are generally well aligned with the identity line, indicating a strong overall agreement between model outputs and simulation results. Nevertheless, a slight increase in dispersion can be observed for certain configurations, particularly those associated with strong solar masking. Because shading information is not explicitly provided as an input, the model must implicitly infer these effects from the remaining variables, which may limit its robustness in highly shaded situations.




Figure 5: Predicted versus simulated heating demand for the baseline model. The solid line represents the identity line.

3.2.2 Shading-aware model

The shading-aware model is evaluated using the same test dataset and performance metrics as the baseline model. Figure 6 presents the comparison between predicted and simulated heating demand.

Compared to the baseline configuration, predictions remain more tightly distributed around the identity line across the full range of values. This reduced dispersion indicates that the explicit inclusion of shading-related inputs improves the model’s ability to capture variations in solar gains induced by neighbouring obstacles, resulting in a more consistent prediction behaviour.




Figure 6: Predicted versus simulated heating demand for the shading-aware model.

3.2.3 Quantitative performance metrics

Quantitative performance indicators are summarized in Table 2. Both models achieve high predictive accuracy on the test set, as reflected by very high coefficients of determination.

Table 2: Prediction performance of both models on the test dataset.

Model	RMSE	CV(RMSE) [%]	R^2
MLP without shading	2.1047	0.96%	0.9991
MLP with shading	0.5977	0.27%	0.9999

However, the shading-aware model systematically outperforms the baseline model, with a reduction of approximately 72% in RMSE and a CV(RMSE) below 0.3%. These results confirm that incorporating shading information leads not only to slightly improved accuracy but

also to a more stable and robust prediction across different operating conditions.

Both models achieve high predictive accuracy on the test set. However, the shading-aware model consistently exhibits lower error levels and a more stable coefficient of determination, particularly for configurations where shading significantly affects heating demand.

3.2.4 Impact of shading scenarios

Several important trends can be observed.




Figure 7: Prediction error distribution across shading scenarios for both models.

First, the baseline model exhibits a strong sensitivity to the shading configuration, with error levels varying significantly across scenarios. In particular, configurations involving southern orientations at short distances lead to a sharp increase in prediction error, reaching values several times higher than in weakly shaded cases. This behaviour highlights the inability of the baseline model to correctly account for strong solar masking effects when shading information is not explicitly provided.

In contrast, the shading-aware model exhibits a much more stable error distribution across all scenarios. The mean error remains low and relatively uniform, regardless of obstacle orientation or distance. This confirms that the introduction of shading-related inputs allows the model to capture the impact of solar obstruction in a consistent manner, significantly reducing sensitivity to specific geometric configurations.

The comparison also reveals that the performance gap between the two models increases as shading becomes more pronounced. While both models perform similarly in weakly shaded cases, the baseline model rapidly degrades under strong shading conditions, whereas the shading-aware model maintains a low and nearly constant error level. This result highlights the key role of shading information in improving robustness and confirms that solar masking constitutes a major source of prediction error when not explicitly accounted for.

Overall, these results demonstrate that incorporating shading parameters does not only reduce the global prediction error, but also ensures a more uniform and physically consistent behavior across diverse shading configurations.

3.2.5 Model interpretability

To improve model interpretability, SHAP values are computed for the shading-aware model. Figure 8 presents a global summary of feature contributions to the predicted heating demand.

Envelope insulation thicknesses and glazing-related parameters are identified as the dominant contributors, in agreement with physical expectations. Shading descriptors also exhibit a significant contribution, with their impact varying across configurations. The dispersion of SHAP values highlights the nonlinear interaction between shading, envelope properties and heating demand, confirming that geometric context plays a meaningful role in the model’s predictions.

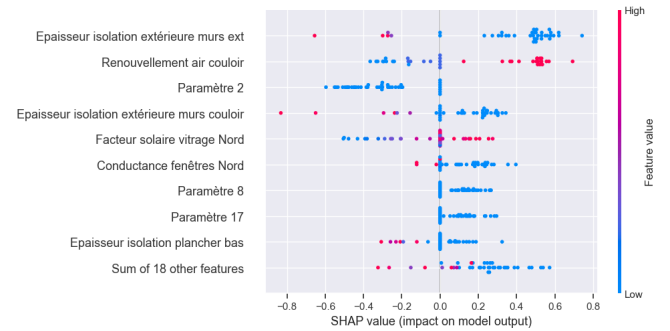


Figure 8: SHAP summary plot for the shading-aware model.

3.2.6 Performance under renovation conditions

Table 3: Model performance before and after renovation.

Model	Consumption (kWh/m^2) (before)	Consumption (kWh/m^2) (after)
MLP without shading	108.59131	109.77188
MLP with shading	106.23883	104.5243

Finally, both models are evaluated under the renovation scenario to assess their robustness to changes in envelope and glazing characteristics. Table 3 compares prediction errors before and after renovation.

While both models experience changes in prediction accuracy following renovation, the shading-aware model maintains a more stable performance level. This behaviour suggests a better generalization capability when the input space is modified, supporting the relevance of explicitly incorporating shading information even when building characteristics differ from those seen during training. It should be noted that, unlike the results reported in Table 2 (expressed in kWh/m^2), the renovation scenario errors presented in Table 3 are reported in absolute kWh. This difference in normalization explains the significantly higher RMSE values observed after renovation.

4 Discussion

4.1 Influence of shading on learning and prediction

The results clearly demonstrate that solar shading has a significant influence on the learning process and predictive accuracy of neural network models for building energy demand, as evidenced by the strong reduction in prediction error obtained when shading-related inputs are introduced (Table 2). When shading information is not explicitly included, the baseline model is forced to implicitly infer solar masking effects from unrelated envelope and operational variables, since the input space does not contain any geometric descriptors of neighbouring obstacles (Section 2.1.2). This limitation leads to increased dispersion in predictions, particularly for configurations with strong or asymmetric shading, as observed in the predicted versus simulated results for the baseline model (Figure 5). In contrast, the shading-aware model benefits from a more physically informed input space. By explicitly encoding geometric shading descriptors such as obstacle azimuth and separation distance (Section 2.1.2), the model is able to directly associate variations in solar gains with changes in heating demand. This results in a substantial reduction in prediction error (Table 2) and a more stable regression behaviour across shading configurations, as illustrated by the tighter alignment with the identity line (Figure 6) and the more uniform error distribution across shading scenarios (Figure 7). The marked improvement in RMSE and CV(RMSE) confirms that shading is not a secondary effect but a key driver of model performance when predicting energy demand in dense or obstructed environments.

4.2 Strengths and limitations of the proposed approach

One of the main strengths of the proposed approach lies in the controlled construction of the simulation dataset. By reusing identical building configurations across all shading scenarios, the study isolates the impact of solar shading from other envelope and operational parameters (Section 2.1.2). This design ensures that observed differences in predictive performance can be directly attributed to shading effects rather than to confounding variability in the input space.

Another key strength is the integration of physically meaningful descriptors into a data-driven modelling framework. The inclusion of geometric shading parameters enables the model to learn relationships that are consistent with building physics, as reflected by the improved prediction accuracy and reduced dispersion observed for the shading-aware model (Figures 6 and 5). The SHAP analysis further supports this interpretation by highlighting the contribution of shading-related features alongside envelope and glazing parameters (Figure 8).

Despite these advantages, several limitations should be acknowledged. The shading configurations considered in this study are deliberately simplified, involving a single neighbouring obstacle with fixed height and idealized geometry. Real urban environments typically include multiple surrounding buildings with varying shapes, heights and surface properties, which may lead to more complex shading interactions than those captured here.

In addition, the models are trained exclusively on simulation-generated data. While this ensures consistency and physical coherence, it may limit direct transferability to real-world buildings, where uncertainties related to construction quality, occupant behaviour and weather variability are more pronounced. Finally, although the shading-aware model demonstrates strong predictive performance within the explored parameter space, extrapolation beyond the training ranges cannot be guaranteed and should be approached with caution.

4.3 Applicability for early design and renovation analysis

The proposed modelling framework is particularly relevant for early-stage design and renovation analysis, where rapid evaluation of multiple configurations is required while detailed information about building components may still be limited. At these stages, full dynamic thermal simulations remain time-consuming and computationally expensive, making them less suitable for iterative exploration of design alternatives.

Once trained, the neural network models provide near-instantaneous predictions of annual heating demand while preserving sensitivity to key physical drivers, including envelope properties and solar shading effects. The strong agreement between predicted and simulated values observed for the shading-aware model (Figure 6) indicates that the approach can reliably approximate simulation results at a fraction of the computational cost.

The renovation scenario further illustrates the applicability of the method in decision-support contexts. Despite significant modifications to insulation and glazing parameters (Table 1), the shading-aware model maintains a stable level of predictive accuracy after renovation (Table 3). This robustness suggests that the model captures generalizable relationships between inputs and energy demand, rather than overfitting to a narrow range of building configurations.

As a result, the proposed approach can be effectively used as a screening or preliminary assessment tool to compare design and renovation options, identify the relative importance of shading effects, and prioritize configurations for more detailed simulation-based analysis. Such a workflow is particularly valuable in dense urban contexts, where solar masking plays a critical role in building energy performance.

4.4 Sources of error and possible improvements

In the proposed approach, shading effects are introduced through a finite number of predefined configurations. While this strategy allows a controlled analysis of the impact of shading, it also implies that solar obstruction is represented in a discrete manner. In practice, it is not possible to enumerate all possible shading configurations, as real urban environments exhibit a continuous range of geometric situations resulting from building positions, heights, orientations and mutual interactions.

This discretization limits the capacity of the model to generalize beyond the simulated cases. Although the neural network learns the influence of shading for the considered configurations, it cannot directly infer the effect of intermediate or unseen situations. As a consequence, the model remains case-dependent and its applicability to arbitrary urban layouts is inherently restricted.

Moreover, shading is intrinsically a time-dependent phenomenon. Its impact varies continuously throughout the day and over the year as a function of solar altitude and azimuth. By encoding shading only through static scenario labels, the model does not receive explicit information about when shading occurs, how long it lasts, or how it evolves over time. Even though this temporal variability is implicitly present in the simulation outputs, it is not explicitly available in the input space, which prevents the model from learning fine-grained solar-thermal interactions.

This limitation suggests that a more general and physically consistent approach would rely on continuous shading descriptors rather than discrete cases. Such descriptors could include geometric parameters (e.g. obstacle orientation, distance, height ratios) or radiative indicators derived from sun-path analysis, such as shadow duration or solar obstruction factors. Introducing these features would enable the model to capture the temporal dynamics of shading and to generalize more effectively to complex and realistic urban configurations.

5 Conclusion

5.1 Key findings

This study demonstrates that explicitly incorporating solar shading information into data-driven building energy models significantly improves predictive accuracy. Compared to a baseline neural network trained without shading-related inputs, the shading-aware model achieves a substantial reduction in prediction error, with RMSE decreased by approximately 72% and CV(RMSE) below 0.3.

The results further show that neglecting shading forces baseline models to implicitly infer solar masking effects

from unrelated envelope and operational variables, leading to increased dispersion in predictions for highly shaded configurations. In contrast, the inclusion of geometric shading descriptors enables the model to directly capture variations in solar gains induced by neighbouring obstacles, resulting in more consistent prediction behaviour across different shading scenarios.

Finally, the robustness analysis under renovation conditions indicates that the shading-aware model maintains stable performance when building envelope and glazing properties are modified. This suggests that the proposed approach learns generalizable relationships between physical inputs and energy demand, supporting its potential use as a fast and reliable decision-support tool in early design and renovation contexts.

5.2 Role of shading in predictive energy modelling

The results of this study highlight the central role of solar shading as a structuring variable in predictive energy modelling. Rather than acting as a secondary correction, shading directly influences solar gains and therefore constitutes a key driver of heating demand, particularly in dense or obstructed urban environments.

From a modelling perspective, the explicit inclusion of shading-related descriptors allows data-driven approaches to better reflect the underlying physical mechanisms governing building energy performance. When shading is not represented in the input space, predictive models are forced to approximate its effects indirectly through envelope or operational parameters, which can reduce robustness and increase prediction dispersion. By contrast, shading-aware input spaces enable a clearer mapping between geometry, solar exposure and energy demand.

These findings suggest that future predictive energy models, especially those intended for urban-scale applications or early-stage design support, should systematically integrate shading information alongside conventional envelope and system parameters. Doing so improves not only predictive accuracy but also model interpretability and physical consistency, reinforcing the relevance of hybrid approaches that combine geometric context with machine learning methods.

5.3 Future work

Several directions can be explored to extend the scope and applicability of the proposed approach. First, the representation of shading could be refined by considering more complex urban configurations, including multiple neighbouring obstacles with varying heights, shapes and spatial arrangements. Such extensions would allow the model to better capture the diversity and geometric complexity of real-world urban environments, in line with recent developments in urban building energy modelling (UBEM) [1, 2].

Future work could also investigate the generalization of the proposed models across different climatic conditions and building typologies. Training and evaluating the framework on datasets generated for multiple climates, or validating predictions against measured building energy consumption, would provide valuable insight into model robustness and transferability beyond the controlled simulation context.

From a methodological perspective, the integration of hybrid modelling strategies represents a promising avenue. In particular, combining data-driven approaches with simplified physical constraints or physics-informed neural networks (PINNs) could improve extrapolation capabilities while preserving computational efficiency [3, 4]. Finally, extending the framework to include cooling demand and multi-objective performance indicators would further enhance its relevance for early-stage design and renovation decision support.

References

- [1] Martina Ferrando et al. Urban building energy modeling (ubem) tools: A state-of-the-art review of bottom-up physics-based approaches. *Applied Energy*, 306:117787, 2021.
- [2] U. Ali et al. A review of urban building energy modeling (ubem): Approaches, challenges and future directions. *Energy and Buildings*, 250:111273, 2021.
- [3] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [4] George E. Karniadakis et al. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.