Using a Convolutional Neural Network to Determine the Thermal Characteristics of a Building

Liam Jowett-Lockwood[1,2], Ralph Evins[1,2]
[1]Energy in Cities group, Department of Civil Engineering, University of Victoria, BC, Canada
[2]Institute for Integrated Energy Systems, University of Victoria, BC, Canada

Abstract

Surrogate models are machine learning models that are trained using detailed simulation input and output data and can practically provide instant results. They use a forward implementation strategy and normally do not use building simulation outputs to identify inputs. This paper introduces inverse surrogate modeling of a building by using a convolutional neural network that uses temperature data to estimate the building characteristics, such as the wall insulation conductivity and infiltration flow rate. The training data resembles thousands of 10-minute interval temperature time series along with their respective building parameters. The first proof of principle uses synthetic data from a building energy model to train and test the neural network. This will later predict actual building parameters using real temperature data from an existing building. Findings demonstrate that combining temperature data with a convolutional neural network could serve as a cost-effective method of building parameter identification.

Introduction

Many cities in the Canadian province of British Columbia contain a significant number of homes built close to or over half a century ago. The 2016 Canadian census states that based on a sample of 25% of occupied private dwellings province-wide, over 43% of homes were built prior to 1981 and over 14% were built prior to 1961 (Statistics Canada (2016)). This number is greater for older cities, such as the city of Victoria which has percentages of 56% and 24% for homes built prior to 1981 and 1961 respectively. Due to the age of these homes, information regarding their parameters (e.g., wall insulation U-value, electric equipment energy) may not be recorded and therefore not readily available for the current owner or potential buyer. Furthermore, some parameters, such as the infiltration flow rate, require tests that are expensive or are intrusive for building occupants (Edwards et al. (2017)). Lack of knowledge on key parameters makes it difficult to prioritize retrofit measures and estimate their benefits.

Buildings are a significant contributor to climate change with a contribution of one-third of global GHG emissions (Bamdad et al. (2021)). A significant portion of these emissions is related to meeting the operational energy demand composed of heating, cooling and electricity use. Effective retrofit strategies do exist, however, that can substantially influence the required energy demand. Due to variability in home construction, climate and deterioration, consideration should be applied to determine whether retrofit strategies are an effective means of reducing emissions and providing long-term cost savings. Not applying sufficient consideration may result in some unfortunate scenarios, where completed retrofits produce little benefit and the required extent of retrofits needed may be too uneconomical to justify (Urge-Vorsatz et al. (2013)). Even if some retrofit strategies are not viable, homes may have sizeable energy usage in other areas, including high equipment and lighting energy usage. Determining whether or not performing retrofits are suitable or if other options should be taken can be challenging, especially if information is unknown.

One approach to determine a building's parameter information would be to use a building energy simulation model to replicate the building energy performance. However, modifying parameter values and producing multiple simulation runs to match the energy performance to that of the actual building has two immediate issues. Firstly, there is the case where the energy performance is matched by multiple different sets of parameter values. This would suggest that some parameters should be found via other means to remove or at least reduce this occurrence. Using outputs of a different form may also help alleviate this issue (e.g., using a time series of energy usage instead of a single aggregate value). Secondly, completing multiple building simulation model runs is computationally intensive and requires significant time. For this point, a potential alternative would be the use of a Surrogate Model (SM).

A SM is a machine learning model designed to replicate the output(s) produced from a high-fidelity model (Westermann and Evins (2019)). For a building energy simulation model, the SM is trained on the model's inputs and outputs, such that the model can accurately predict outputs from unseen inputs. For synthetic data, the process often involves the development of a building simulation model to produce extensive training data. While the development of a building simulation model for the purposes of training a SM may seem to make use of the SM nonsensical, ideally a trained SM can be applied to other projects/geometries as well. Research has shown that when developed effectively, a SM can be deployed to achieve strong prediction performance where multiple building designs are trained for and the location is held constant (Lawrence et al. (2021)).

For buildings, SMs have several different uses including parameter sensitivity analysis, early design exploration, and optimization (Westermann and Evins (2019)). Furthermore, the broad field of surrogate modelling is composed of many different types of machine learning models including linear regression, Gaussian process and Artificial Neural Network (ANN) models. In particular, a Convolutional Neural Network (CNN) is an ANN with the unique implementation of convolutional kernels. This allows for a degree of invariance in the model as kernel weights are reused at each section of input. The usage of convolutional kernels allows CNNs to more easily handle potentially extensive time series inputs, such as building energy usage or temperature recordings, in comparison to other machine learning models (Westermann et al. (2020)). This would thereby suggest it as a strong candidate as a SM for projects involving inputs of this scale.

SMs for buildings typically use building specific parameters as model inputs and energy use or simulated results as outputs. Creating an inverse building SM would involve a machine learning model that uses typical building simulation outputs as inputs with the intention of predicting building parameters. With the effective ability of CNNs to accurately learn from extensive time-series inputs, there is a potential to treat building time series information, such as hourly energy usage or zone temperature, as an inverse SM input. For existing buildings, time series outputs may be known from data from recent years or easily collected sensor data. These can be used to determine unknown building parameters through the use of an Inverse Surrogate Model (ISM).

This paper presents the development and evaluation of a building ISM to facilitate the determination of specific building parameters. The ISM is based on a deep CNN because of the strong performance of CNNs with time series inputs. Trained on multi-zone building temperature data, this CNN ISM attempts to identify building parameters that may influence the temperature of different rooms over time. This study currently only examines the effectiveness of the model on synthetic data, however it is intended to eventually extend the research to include real data. This would resemble the model being training on temperature outputs from a building simulation model designed on a real structure and will then predict the actual building parameters of the same structure using sensor data.

This paper is organized as follows: Section 2 describes the main elements of the overarching methodology to formulate trained inverse building surrogate models. Section 3 describes and discusses the findings and adjustments made to the methods in analyzing results. Section 4 discuss the planned future work of incorporating real data. Section 5 concludes the paper.

Methodology

The process demonstrating the overall workflow is shown in Figure 1. The two most significant portions of this process was the development of the synthetic data generation and ISM. Data pre-processing is applied to the data prior to training to improve performance. Various error metrics are used to evaluate synthetic data, however they will be omitted from eventual real data evaluation due to the lack of reference data.

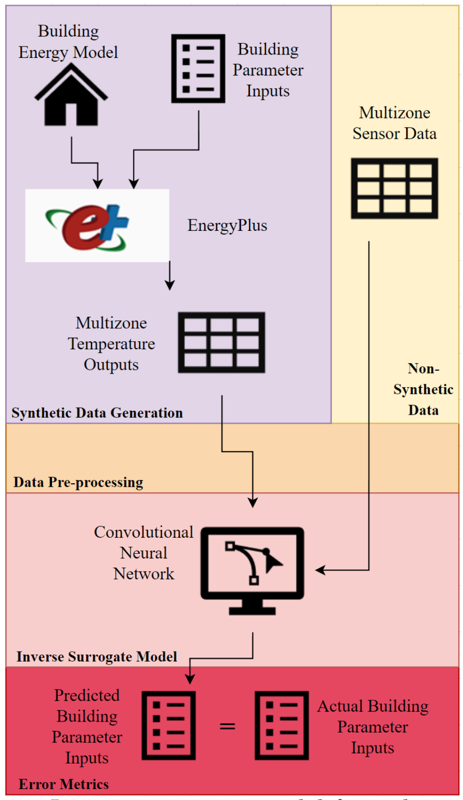

Figure 1: Inverse surrogate model formulation workflow to predict building parameter inputs.

For predicting synthetic data, the process can almost be considered circular, as varied parameter inputs are initially fed into the building simulation model and the final output of the ISM are predictions of the original inputs.

Synthetic Data Generation

In order for the trained ISM to be best suited for handling real data, the building simulation model was designed to reflect a specific existing residential building in Victoria, British Columbia. The building simulation software EnergyPlus was used for the simulation modeling. Within the simulation model, each of the 14 rooms of the home are treated as separate thermal zones from which temperature values are produced. Out of these rooms, 7 are present on the first floor, 5 on the second floor and 1 each for the basement and attic. Some of these rooms include hall and entry ways. For the privacy of the building owner and occupants, an exact floor plan and room layout have been omitted here. As some rooms are not fully separated from others (e.g. lack of a door), full interior partitions/walls are placed in the software to separate them.

In regards to HVAC, the EnergyPlus Ideal Air Loads HVAC system object was used in place of modelling mechanical systems. This object is typically used when the designer wishes to omit the implementation of a complicated HVAC system in the model while still wanting to examine building performance. The object resembles an ideal HVAC system that applies or removes heat and moisture at 100% efficiency to provide a supply air stream (USDOE (2022)). The unit is not connected to a central air system, rather each system object supplies air to a zone in necessary quantities to heat or cool in order to reach specified setpoints.

To produce training data for the ISM, various parameters were varied for each simulation run, see Table 1. The parameters and ranges were chosen based on previous research (Westermann et al. (2020)), their importance for a home retrofit and potential variability among households. As the considered building is an old construction (built in 1930s), values for many of the actual parameters are unknown. Engineering judgement was applied to develop assumptions on the properties of the constructions, however the extent of these assumptions is a likely contributor to a poor representation of the actual building. Randomized parameter value selection was conducted by using Latin Hypercube Sampling (LHS), which is a strong and widely used random sampling method because of its ability to greatly lower variance amongst different applications (Shields and Zhang (2015)).

Notably, both the Zone Ventilation Flow Rate and Zone Infiltration Flow Rate refer to the airflow in and out of the home. EnergyPlus also provides methods of transferring airflow around the home via mixing objects or its AirflowNetwork model, however these methods were not included in the model due to added complexity and how flow amounts would very significantly between different zones in the actual home because of opening sizes. Whereas the Zone Ventilation Flow Rate is defined as the purposeful flow of outside

Table 1: List of the varied building parameters for ISM training.

Parameter	Units	Range
North Axis	°	0-360
Wall Insulation Conductivity	W/mK	0.01-0.1
Roof Insulation Conductivity	W/mK	0.01-0.1
Foundation Concrete Thickness	m	0.05-0.2
Window Conductivity	W/mK	0.01-0.05
Lighting Energy	W/m^2	6-10
Electric Equipment Energy	W/m^2	5.5-9.5
Personnel	People/m^2	0.015-0.035
Zone Ventilation Flow Rate	L/sPerson	5-10
Zone Infiltration Flow Rate	L/m^{2}s	0.1-0.5

air to a thermal zone to provide cooling, the Zone Infiltration Flow Rate is the airflow into the zone through unintended openings, such as cracks around windows (USDOE (2022)). A significant difference between the two is that, for the Zone Infiltration Flow Rate, the flow per exterior surface area is what is specifically varied whereas the flow rate per person is varied for the Zone Ventilation Flow Rate. Calculations do not take into account location of envelope leakage. Furthermore, EnergyPlus does provide the option to use various formulas/models for airflow for both parameters, which would include the influence of additional factors, such as wind speed. However, for simplicity and the unknowns of the building itself, these have been omitted

It is also noteworthy that, for some of the material parameters, conductivity or thickness were selected instead of U-value, so that the variation could be more easily implemented into the simulation model. However, the variation in U-value can still be obtained if need be.

In total 4000 simulation runs were completed, taking roughly a minute for each. Temperature values for each of the 14 zones were extracted with a Python script and organized in a CSV file resulting in 4000 separate data files. Temperature values are recorded in 10-minute intervals for a year which resulted in a total of 52560 values for each zone. Additionally included in the files was a column of binary values for which values of 1 indicated that heating was present and 0 indicating otherwise. For the purposes of training, this additional column is treated as another zone making the total 15 zones.

As each data file constitutes one simulation run, the decision to use only 4000 simulations was based on the sheer size of each file (15x52560 values). Out of the total 4000 files, the training, validation and testing sets were composed of 2560, 640 and 800 files respectively. The determination of these quantities was done by using a train-test split of 0.2 and a train-validation split of 0.2 as well. While the CNN may perform stronger with a larger amount of simulation data to train on, lowering the amount of data saves time and computation effort as less simulation runs are completed, and less data is used for ISM training.

While 52560 values were produced to encompass the entire year, training and testing occurred on 30-day intervals to determine if different time periods were impactful for results, hence only a portion of each file was used. The months chosen were March, June, September and December. The decision of choosing which months to use was based on the distance apart from each other and the time frame of the preliminary sensor data gathering to be eventually used for parameter prediction of the real structure.

Data Pre-processing

Prior to training the SM, pre-processing was applied to the data. This involved standardizing the dataset to resemble standard normally distributed data. Standardization re-scales the feature values into a more suitable range, such that the influence of differences between the values of data is less significant (P.Angelov and Gu (2019)). For this study Scikit-Learn's StandardScaler (Pedregosa et al. (2011)) functions were used to compute preprocessing for both the training and testing data. This involved centering the data by removing the mean value from each feature and scaling such that mean and variance values of 0 and 1 are achieved. To conduct the preprocessing it was required to load the entire data sets in memory. For the outputs (varied parameter values), these were simply loaded in as a single CSV file. For the inputs however, as these spanned multiple data files (one for each parameter variation) composed of multiple time series, the process involved loading each file and then modifying the two-dimensional matrix into a one-dimensional array as shown in Figure 2.

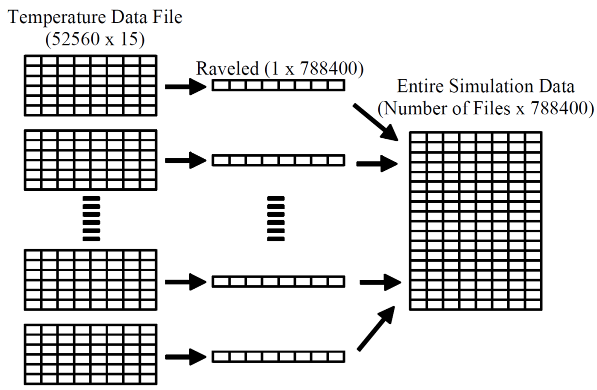

Figure 2: Pre-processing methodology for ISM input data where multiple files are loaded into memory as one.

It should be noted that the mean and standard deviation used for centering and scaling are calculated only from the samples in the training (and validation) set. The same mean and standard deviation values are then used on standardizing the testing set. Including the non-standardized testing set in the calculation of the mean and standard deviation is a form of data leakage and is a common mistake in machine learning, which can lead to false optimistic results.

Convolutional Neural Network Structure

The CNN structure is heavily based on the ResNet model architecture deployed by Westermann et al. (2020). The reader is referred to their research for a comprehensive breakdown of the various design decisions and underlying model features. The overview of the model structure is presented in Figure 3.

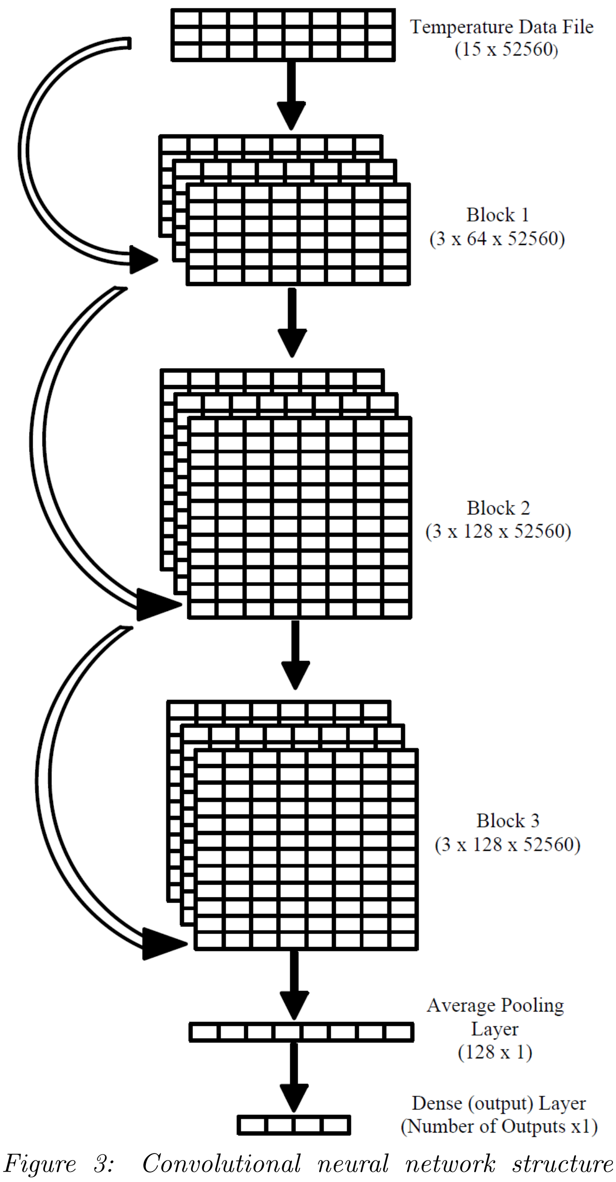

Figure 3: Convolutional neural network structure where straight arrows indicate a direct connection to subsequent blocks and curved arrows for residual shortcuts.

Each of the three blocks is composed of a one-dimensional convolutional, batch normalization and activation layer. The activation layer computes a Rectified Linear Unit (ReLU) activation. A residual shortcut is also included, thereby connecting the last layer of the previous block with the following block, in an effort remove the occurrence of the vanishing gradient. The main difference between the model developed by Westermann et al. (2020) is the removal of the inclusion of additional features after the final dense layer. In their research, the ANN is continued with a concatenation of additional features to the dense layer which is followed by an additional dense layer and eventual output. This is unnecessary for this project as no features are trained for outside of the multiple inputted time series.

Error Metrics

The purpose of error metrics is to determine the effectiveness of the CNN in being able to predict parameters. After being trained, the ISM will predict building parameter values from the temperature time series data for each simulation run. The error metrics themselves are equations whose values serve as indicators of performance. The equations accept both the predicted building parameter values and actual as inputs. Assuming the building energy model is an accurate representation of the actual building, parameters that the ISM can predict with high accuracy with synthetic data would ideally achieve higher prediction accuracy than those with poorer performance when real data is used.

Various error metrics were used in this study. These match some of the metrics previously used by Westermann et al. (2020). The metrics consisted of the Coefficient of Determination (R^2), Mean Absolute Percentage Error (MAPE) and the Root Mean Squared Percentage Error (RMSPE), which are defined below:

$$R^2(y,\dot{y}) = 1 - \frac{\sum_{i=1}^{n}(y_i - \dot{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y}_i)^2} \tag{1}$$

$$MAPE(y,\dot{y}) = 100(\frac{1}{n}\sum_{i=1}^{n}\frac{|y_i - \dot{y}_i|}{y_i}) \tag{2}$$

$$RMSPE(y,\dot{y}) = 100\sqrt{\frac{1}{n}\sum_{i=1}^{n}(\frac{y_i - \dot{y}_i}{y_i})^2} \tag{3}$$

Here, y is defined as the vector of actual values, $\dot{y}_i$ is the vector of predicted values and $\bar{y}$ is the mean of actual values. Actual values refer to recorded measurements of indoor air temperature. Combining these error metrics yields promising insights into the model performance. In their literature review, Østergård et al. (2018) showed how R^2 values are often presented for a variety of different building SMs and allow an easy comparison between them despite the diversity of the problems. Therefore, it can be considered a metric of high importance as higher values indicate the closeness of fit and overall performance. The other errors are relative metrics where the RMSPE is more drastically affected by small, large errors than MAPE (Westermann et al. (2020)).

Zone Quantity Impact

Preliminary sensor data was conducted to eventually be used to predict the home's parameters. Unfortunately, the data gathering encountered an issue where some sensors' batteries were depleted faster than anticipated resulting in severe data gaps for many of the sensors. Time periods still remained where a few sensors were able to provide a sufficient amount of data. To anticipate the required reduction in the number of zones it was insightful to examine the impact of training the model for a smaller number of thermal zones. These zones included two on the first floor, one on the second and the attic. The binary heating column was also removed to investigate performance without it; therefore only 4 zones were used.

Results and Discussion

Training was originally completed with 100 epochs, which took roughly 3 minutes for each, while using a default Batch Size of 32. Early stopping was also included to prevent overfitting. Early results suggested that training for 1 parameter output typically yielded better results than training for all parameters at once (Table 2). Exceptions do remain (such as the R^2 for Zone Infiltration Flow Rate), but error metric scores more often tended to show a notable decrease when trained for 1 parameter instead of all.

Table 2: R^2 error metric results for training on all parameters at once and individually for the month of December.

	Train		Test	
Parameter	**Group**	**Individual**	**Group**	**Individual**
North Axis	0.795	-	0.755	-
Wall Insulation Conductivity	0.983	0.996	0.981	0.995
Roof Insulation Conductivity	0.976	0.984	0.969	0.979
Foundation Concrete Thickness	0.268	-	0.194	-
Window Conductivity	0.000	-	-0.004	-
Lighting Energy	0.426	-	0.462	-
Electric Equipment Energy	0.482	0.555	0.497	0.439
Personnel	0.560	-	0.532	-
Zone Ventilation Flow Rate	0.487	-	0.409	-
Zone Infiltration Flow Rate	0.972	0.958	0.968	0.959

Further investigation should be conducted by varying the number of epochs and batch size to examine the significance of the differences. Due to these perceived difference in errors, 4 parameters (highlighted in Table 2) were examined for the rest of the project while the others remained varied but not trained for, thereby creating noise to an extent. The choice of these 4 parameters was to capture those with the highest potential performance in terms of error metrics and one with poorer prediction accuracy.

Results on Synthetic Data

One disadvantage of switching from training on all outputs at once to only one is that it involves the creation of multiple ISMs (one for each time period and parameter). This slows down training significantly as each ISM requires its own training to be completed. To account for this disadvantage, the 4 parameters were trained at once instead of separately. The number of epochs was also increased from 100 to 150, with early stopping still enabled.

While the initial results in Table 2 do not seem quite promising, the overall analysis identifies the sensitivities of individual parameters on temperature. Henceforth, the ISM can easily identify highly influential parameters (such as Wall Insulation Conductivity) compared to those that only make minor adjustments in temperature. This is perhaps unsurprising, as thermal insulation and infiltration are expected to be a significant influence on energy demand (Zhiqiang John Zhai (2019)), hence it can be expected that variations among values can cause larger temperature fluctuations (i.e., more energy would be needed to provide a comfortable temperature which is most impacted by these parameters). This would thereby imply that the method may be best suited for determining envelope characteristics. Results for the ISMs for each of the 4 parameters and their respective time periods are shown in Table 3.

Table 3 indicates strong performance for ISMs to determine exact building parameter values. Both the insulation conductivities and infiltration flow rate showed high performance overall with the R^2 frequently above 95%. Even when trained individually, Electric Equipment Energy (a previous poor performing parameter), still displayed disappointing performance. It should not be understated though that the ISMs are still able to provide a somewhat accurate, though not ideal, estimate for it. This would suggest that an ISM, depending on the user's requirements, may still be adequate in providing at least an idea of where the building parameter value lies.

An interesting observation can be made that the R^2 values in Table 3 tend to be slightly lower than those in Table 2. While this difference is minimal, it remains surprising. A possible explanation is that while the train and test sets were taken from the same pool of data, the distribution of the data into the two sets was different (i.e. some data used for training with 15 parameters were used for testing in the case of 4 parameters). This may suggest that the training set used when training for 15 parameters was better suited than the set with 4 parameters. Larger data sets will be explored in the future to determine if improvements can occur. Furthermore, the improvement of reducing the number of parameters may not be seen, as the discrepancy between the R^2 values for the Electric Equipment Energy parameter suggests that overfitting might be occurring. As early stopping is being applied, a possibility exists that training is being ended too early in an attempt to limit the overfitting of this parameter, while the other parameters still require longer training periods.

Lastly, it is worth mentioning the performance of the North Axis parameter shown in Table 2. One concern of applying the error metrics to this parameter is that the values are in degrees of a circle (e.g., 0-360°). Therefore, as an example, while values of 5° and 355° may not actually be significantly far apart, the error metric calculates the difference as being 350° thereby producing significant error. Future work will explore better methods of performance evaluation for this parameter.

Comparison between 4 and 15 zones

The comparison of results between ISMs trained on 15 zones and those trained on 4 zones are shown in Table 4.

Unsurprisingly, the ISMs perform worse when only trained on 4 sets of temperature data instead of 15. Differences are mostly minimal for the R^2 values for the insulation conductivities, however a large rise in both the RMSPE and MAPE values are notable. Additionally, the Zone Infiltration Flow Rate performs significantly worse overall with the R^2 going below 90%. Interestingly, the Electric Equipment Energy parameter achieves a slightly higher performance with reduced zones, however the R^2 is still much poorer in comparison to the other parameters. The reduction in zones would appear to remove key informative locations, thereby reducing the ability of the model to train effectively. This is especially notable for the Roof Insulation Conductivity parameter as zones on the main floor of the building (not directly touching the roof) are much less influenced by the conductivity of the roof insulation. One of the remaining 4 zones is the attic, hence the ISM performance is still solid as it is heavily influenced by the parameter, however this still demonstrates that caution needs to be applied when selecting a subset of zones to train on.

Future Work

Future work largely pertains to eventually being able to perform predictions with real data. Significant improvements can still be made to the current building simulation model. While EnergyPlus has shortcomings when replicating real building energy simulation performance, an inadequate representation is often much more a result of poor input data (Edwards et al. (2017)). For this reason, it is imperative that the building simulation model is designed to mimic the current home as accurately as possible. In our case, inaccuracies can be largely attributed both to the model inputs that were not varied (such as setpoints) and additional input (e.g., the weather file). Temperature setpoints should be adjusted such that

Table 3: *Error metric results for the synthetic data train and test sets. Colouring is applied separately for each error metric. An upper bound of 30% was only used for colouring and both RMSPE and MAPE can go higher.*

Training	Wall Insulation Conductivity			Roof Insulation Conductivity			Electric Equipment Energy			Zone Infiltration Flow Rate		
	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE
March	0.976	10.263	7.074	0.968	11.729	7.965	0.598	10.183	8.18	0.982	6.264	4.573
June	0.95	13.029	9.572	0.969	11.844	8.127	0.444	12.832	10.027	0.976	9.413	6.171
Sept	0.96	15.182	10.566	0.967	16.879	10.163	0.563	10.528	8.518	0.978	10.152	6.598
Dec	0.973	11.926	7.904	0.942	14.142	10.293	0.553	11.324	8.878	0.977	9.224	6.244
Testing	Wall Insulation Conductivity			Roof Insulation Conductivity			Electric Equipment Energy			Zone Infiltration Flow Rate		
	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE	R^2	RMSPE	MAPE
March	0.974	10.796	7.383	0.968	13.213	8.888	0.477	11.187	9.01	0.977	6.679	4.999
June	0.949	13.375	9.593	0.97	12.173	8.205	0.361	13.327	10.399	0.974	9.732	6.259
Sept	0.958	13.842	10.008	0.962	19.352	11.501	0.462	11.441	9.097	0.976	9.975	6.713
Dec	0.97	11.781	7.97	0.945	13.379	10.134	0.447	12.185	9.494	0.976	8.931	6.205
Legend	R^2	0										1
	RMSPE & MAPE	0										30

Table 4: *Comparison between results of the previous 15 zone ISMs and 4 zone ISMs.*

Parameters	15 Zone Test Set	4 Zone Test Set
Wall Insulation Conductivity	R^2: 0.970, RMSPE: 11.781, MAPE: 7.97	R^2: 0.940, RMSPE: 19.901, MAPE: 12.094
Roof Insulation Conductivity	R^2: 0.945, RMSPE: 13.379, MAPE: 10.134	R^2: 0.935, RMSPE: 27.172, MAPE: 15.167
Electric Equipment Energy	R^2: 0.447, RMSPE: 12.185, MAPE: 9.494	R^2: 0.458, RMSPE: 11.588, MAPE: 9.310
Zone Infiltration Flow Rate	R^2: 0.976, RMSPE: 8.931, MAPE: 6.205	R^2: 0.892, RMSPE: 15.184, MAPE: 11.364

the minimum and maximum zone temperatures are an accurate reflection of those currently in the home. Furthermore, the weather file currently used for simulations does not contain the exact weather features of the time period when real data was recorded. Lastly, as mentioned previously, larger train and test sets will be explored as well as better methods of evaluation for the North Axis parameter

While future work is promising, it still unfortunately remains difficult for the ISM to estimate the exact values of the home's inner components. Additionally, as the exact parameter values are unknown, a comparison to the actual cannot be made to what was predicted. However being able to predict in expected bounds of the parameters (i.e. those in Table 1) could still suggest the method is an effective method of parameter prediction.

Conclusions

This paper presents the development and use of a Inverse Surrogate Model designed to accept building temperature time series data as input and accurately predict building parameter values. ISMs could serve as a cost effective alternative method of building parameter identification. This would assist home owners with relevant information to help determine if a retrofit is justifiable. When trained and tested on synthetic data, the model showed high performance with the insulation conductivity parameters and infiltration flow rate. Due to limitations when trying to collect real data with physical sensors, the model was retrained with only 4 zones worth of data. This produced similar results, however there was a notable decrease in performance among the conductivity and infiltration flow rate parameters.

Predicting with real sensor data was not completed at this stage of the project. As the results of some of the ISMs on synthetic data are promising however, future work is planned with a focus of achieving strong performance on real data. This will include creating a more appropriate weather file and performing slight model adjustments and editing temperature setpoints for each individual zone/room to ensure that the temperatures reach a more correct default.

References

Bamdad, K., M. E. Cholette, S. Omrani, and J. Bell (2021). Future energy-optimised buildings — addressing the impact of climate change on buildings. *Energy and Buildings 232*.

Edwards, R. E., J. New, B. Cui, and J. Dong (2017). Constructing large scale surrogate models from big data and artificial intelligence. *Applied energy 202*, 685–699.

Lawrence, C. R., R. Richman, M. Kordjamshidi, and C. Skarupa (2021). Application of surrogate modelling to improve the thermal performance of single-family homes through archetype development. *Energy and Buildings 237*.

P.Angelov, P. and X. Gu (2019). *Empirical Approach to Machine Learning*. Springer.

Pedregosa, F., 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*, 2825–2830.

Shields, M. D. and J. Zhang (2015). The generalization of latin hypercube sampling. *Reliability Engineering & System Safety 148*, 96–108.

Statistics Canada (2016). Census profile, 2016 census - victoria [census metropolitan area], british columbia and british columbia [province]. `https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/prof/details/page.cfm?Lang=E&Geo1=CMACA&Code1=935&Geo2=PR&Code2=59&Data=Count&SearchText=victoria&SearchType=Begins&SearchPR=01&B1=All&TABID=1`. Accessed: 2022-02-10.

Urge-Vorsatz, D., K. Petrichenko, M. Staniec, and J. Eom (2013). Energy use in buildings in a long-term perspective. *Current Opinion in Environmental Sustainability 5*(2), 141–151.

United States Department of Energy (2022, March). *EnergyPlus Version 22.1.0 Documentation: Input Output Reference*.

Westermann, P. and R. Evins (2019). Surrogate modelling for sustainable building design – a review. *Energy and Buildings 198*, 170–186.

Westermann, P., M. Welzel, and R. Evins (2020). Using a deep temporal convolutional network as a building energy surrogate model that spans multiple climate zones. *Applied Energy 278*.

Zhiqiang John Zhai, J. M. H. (2019). Implications of climate changes to building energy and design. *Sustainable Cities and Society 44*, 511–519.

Østergård, T., R. L. Jensen, and S. E. Maagaard (2018). A comparison of six metamodeling techniques applied to building performance simulations. *Applied Energy 211*, 89–103.