CityTFT: Temporal Fusion Transformer for Urban Building Energy Modeling

Ting-Yu Dai
The University of Texas at Austin
Austin, TX 78712
funnyengineer@utexas.edu

Dev Niyogi
The University of Texas at Austin
Austin, TX 78712
dev.niyogi@jsg.utexas.edu

Zoltan Nagy
The University of Texas at Austin
Austin, TX 78712
nagy@utexas.edu

Abstract

Urban Building Energy Modeling (UBEM) is an emerging method to investigate urban design and energy systems against the increasing energy demand at urban and neighborhood levels. However, current UBEM methods are mostly physic-based and time-consuming in multiple climate change scenarios. This work proposes CityTFT, a data-driven UBEM framework, to accurately model the energy demands in urban environments. With the empowerment of the underlying TFT framework and an augmented loss function, CityTFT could predict heating and cooling triggers in unseen climate dynamics with an F1 score of 99.98 % while RMSE of loads of 13.57 kWh.

1 Introduction

Urbanization is one of the greatest challenges of modern society. Almost one-third of global greenhouse gas emissions come from buildings and 70% of energy is consumed by urban. As of the latest available data in 2021, the global population stood at over 7.8 billion individuals [1]. Projections indicate that in 2050 [2], the world's population may surpass 9.7 billion, reflecting a substantial demographic expansion over the intervening years. To address the long-term challenge posed by urbanization effectively, urban building energy modeling emerges as an imperative and requisite approach in academic research and urban planning endeavors.

The fundamental purpose of urban building energy modeling is to simulate and analyze the intricate dynamics of energy consumption within urban environments. Compared to building energy modeling, UBEM simulates while considering building height, surface coverage, and spatial arrangement, probing their interactions and discerning their collective influence on energy dynamics at the urban scale. Swan and Ugursal [3] set a tone in reviewing urban modeling of the residential sector, subdividing the modeling methodologies into top-down and bottom-up approaches. Specifically, bottom-up physic-based UBEM methods have garnered attention in recent times [4]. Robinson et al. developed CitySim to assist urban settlements with sustainable planning and also simulate the energy use of a few buildings ranging from tens to thousands. Those UBEM methods are reasonably accurate in simulating the performance of almost any building combination systems [6]. However, customized urban projects, and optimization problems involving many UBEM evaluations, are time-consuming and labor-intensive, and based on the scalability, the simulation runtime can be exponentially high if a broad set of design variations is analyzed.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

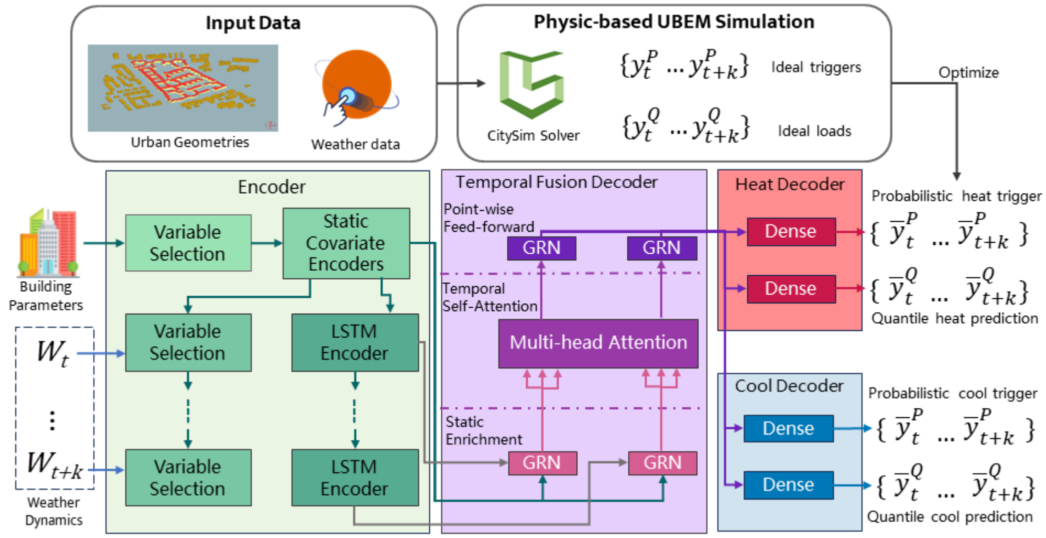

Figure 1: Overview of CityTFT. CitySim and CityTFT apply the same weather input while CityTFT uses simplified building parameters.

Here, we propose a surrogate data-driven approach to accelerate the simulation process in UBEM. Compared to similar previous works [7–9], we based our work on the extensively used forecasting model, Temporal Fusion Transformer (TFT) [10]. We extract the static covariate encoder and variable selection network from the TFT structure while adding a small neural network to model the probability of triggering heating and cooling needs. The major improvements are: 1.) Sequential input has been applied in a transformer-based model to improve the temporal accuracy. 2.) A training strategy that models from weather dynamics and urban interactions to energy demands simultaneously with a customized loss function. 3.) Improved generalizability for the proposed surrogate model. Urban planners or Energy Sectors could benefit from urban-level energy demand estimations by the proposed approach to enhance the decision-making process.

2 Methods

Data Prepration As a surrogate model, this work aims to duplicate the ability of physic-based UBEM models, CitySim [5], to simulate urban building energy. Two major inputs of UBEM are urban geometries and weather data. This study utilized the university campus geometries and the calibrated building parameters from [11] which contains 114 buildings in a wide variety. 13 variables are extracted from the building geometries to represent the static building covariates. The detailed transformation method can be found in the appendix. The other requirement of CitySim is weather information. This study collects 21 Typical Meteorological Year (TMY) files in different climate zones globally by Meteonorm [12] to obtain fruitful weather dynamics. 12 environmental variables are served in climate files as the temporal covariates while one additional variable, hour of year, is added to improve the temporal identification. Those variables are listed in Table 3. The ideal hourly heating and cooling demands are simulated through CitySim solver for each building which concludes a dataset with roughly 17 million samples. The sequence length in our training process is 24. All the building variables, weather data, and electricity loads have been normalized for preprocessing.

Temporal Fusion Transformer The Temporal Fusion Transformer (TFT) is composed of key elements: a variable selection network, static covariate encoder, and temporal processing, facilitating precise energy demand modeling. The variable section network is applied to both static building properties and temporal weather covariates to provide instance-wise variable selection. Linear transformations are applied to transform each variable into a d_{model} dimension vector to match the subsequent paper for skip connections. Gated Residual Networks (GRN) are utilized for static enrichment and point-wise feed-forward processes while self-attention layers are utilized for temporal feature processing. To decode the transformed latent correctly, the decoder of TFT first applied GRN to enrich temporal signals with static building latents. Secondly, another self-attention network is applied to assemble static-enriched features into a single matrix. The attention mechanism here is to help that each temporal dimension can only attend to features preceding it.

Different from the original TFT, this work aims to predict the energy demands in the same temporal period i.e. sequential modeling, unlike the purpose of the forecasting mission in the original paper. Therefore, this work aborts the future encoder part from the model and applies the original observed temporal features to proceed with the static enrichment and self-attention processing. The static enrichment is applied to embedded weather representations, and an interpretable attention mechanism builds the attention matrix in each time step of the enriched features. We also add one more linear layer with sigmoid activation for probabilistic projection. Heating and Cooling loads are divided into two variables while cooling loads are negative and heating loads are positive.

Probabilistic Loss: Will it trigger and how much it will cost Since about 30 and 50 percent of the heating and cooling loads are zero, this work proposes a probabilistic-based loss to handle the imbalanced data. The loss equation is summarized in Eq. 1 where t is whether heating/cooling systems are triggered, y^P is the probabilistic output produced by the TFT model, a is the actual heating/cooling loads, and y^Q is the quantile projections by the TFT model. Instead of computing deterministic energy consumption predictions, a probabilistic output of whether the heating/cooling system will be triggered is computed first. Eq. 2 calculates the loss between the output and the triggering probability from target loads. On the other side, the quantile loss in Eq. 3 only optimizes with the projections that are paired with the non-zero loads while q represents the target quantile. The quantile loss aims to optimize the range of likely target values instead of a deterministic output, and the probability loss focuses on preventing our networks from predicting smaller values while optimizing with void loads.

$$l = l_{prob}(t, y^P) + l_{quantile}(a, y^Q) \tag{1}$$

$$l_{prob}(t, y^P) = \frac{1}{N}\sum_{n=1}^{N}\{-w_n(y_n^P \cdot \log(t_n)) + (1 - y_n^P) \cdot \log(1 - t_n))\} \tag{2}$$

$$l_{quantile}(a, y^Q) = \frac{1}{N}\sum_{n=1}^{N}\begin{cases} q \cdot (a_n - y_n^Q) & , y_n^Q \leq a_n \\ (1-q) \cdot (y_n^Q - a_n) & , y_n^Q > a_n \end{cases} \tag{3}$$

3 Result & Discussion

Model projections will be compared with the conjectural heating and cooling demands by CitySim to quantify the effectiveness of our proposed methods. 4 additional simulations by CitySim are conducted with weather files in different climate zones and the same campus environment to produce conceptual heating and cooling demands. To assess the effectiveness of the proposed method, a classic recurrent neural network and a transformer are trained to compare with the TFT model. All models are trained for 400 epochs using AdamW Optimizer in learning rate as 1e-4.

Ability to simulate unseen climate dynamics The probabilistic and deterministic outputs from the TFT model are compared with the CitySim simulation results. Extra 4 simulations are conducted to evaluate the ability of our CityTFT model in unseen climate dynamics. The usage of climate files could be referred to Table 6. The prediction from CityTFT is made by observing probabilistic and quantile output at the same time. If the probability at a certain time step is over the threshold, then the prediction at the same time step is filled with the median projection in the quantile output. Fig. 2 demonstrates the predicted loads and ideal loads by CityTFT and CitySim. Each of the three models demonstrates a proficient ability to estimate both heating and cooling demands concluded in Table. 1. In particular, CityTFT exhibits the highest level of proficiency in accurately predicting energy demands, whereas the other two models occasionally display tendencies to underestimate or overestimate the regressed values.

Ability to anticipate the trigger of heating and cooling. To evaluate the probabilistic prediction, we compare the F1 score of three different models while classifying the CitySim results into a binary class as zero and non-zero. Table 1 indicates that all three models have a great ability to predict whether heating and cooling are triggered. We could also see that the scores in cooling

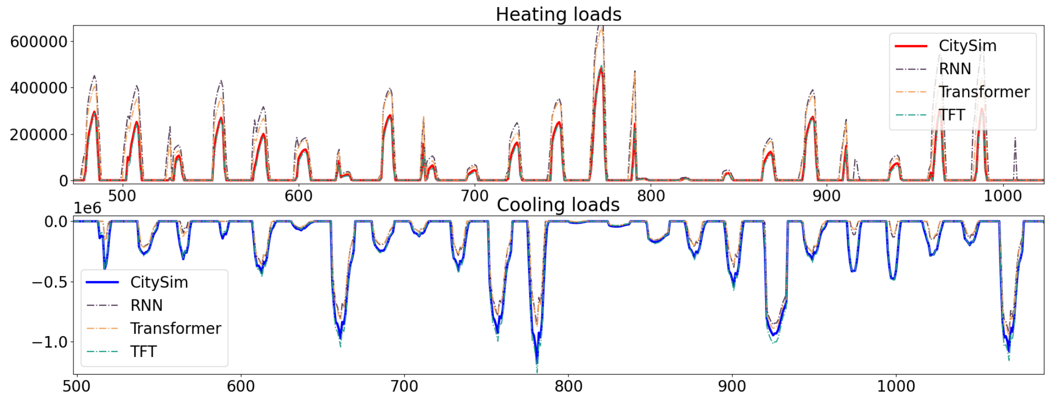

Figure 2: Comparison between ideal and predicted heating/cooling loads

loads are normally lower than those in heating loads shown in Appendix A.4. The reason is that, in physic-based UBEM tools, triggering heating systems is more straightforward which related to weather dynamics. In contrast, the activation of cooling systems is intricately affected by weather, solar heating, and human comfort examination.

However, even though the transformer model is predicting slightly worse than the other two, it shows comparable output in Fig. 2 and lower MAPE values in Table 1. We could observe that the predictions from the transformer are even closer than those from RNN. RNN, on the other hand, could better predict the triggering probability while the regressed projection is worse. Those findings could be observed in the demonstration in Fig. 2. The debate between the self-attention and recurrent mechanisms is worth investigating more deeply. The findings suggest that employing a hybrid structure, such as CityTFT, can yield exceptional performance in this task when contrasted with the exclusive utilization of either RNN or self-attention models.

	F1 score (%)	Non-zero RMSE (kWh)	Total RMSE (kWh)	Non-zero MAPE (%)
RNN	91.91	114.06	75.91	136.89
Transformer	91.33	118.43	79.74	113.65
TFT	**99.98**	**21.34**	**13.57**	**11.62**

Table 1: F1 score, RMSE with total loads and non-zero loads Comparison

4 Conclusion

This work proposes CityTFT, a temporal fusion transformer, to model the urban building energy in unseen climate dynamics. This reduces the barrier for individuals to possess a meticulously designed building geometry for simulating energy demand, thereby accelerating the more precise quantification of energy modeling within the context of climate change. This permits decision-makers to conduct a thorough exploration of energy usage across multiple buildings and climate zones during the design phase, with reduced reliance on intensive simulation efforts. From a technical perspective, several facets merit additional assessment. A thorough examination should be made to evaluate CityTFT on climate change data like CMIP6. An ablation analysis involving building parameters and weather data could be additionally integrated to assess the lower limit of CityTFT's performance. The community could realize substantial benefits in the event that CityTFT could operate comparably on more concise input from publicly available data sources, such as satellite observations or open street map data.

References

[1] Department of Economic United Nations and Population Division Social Affairs. World population prospects 2022. 2022. URL https://population.un.org/wpp/. Accessed on September 12, 2021.

[2] United Nations Department of Economic and Population Division Social Affairs. World population prospects 2022: Summary of results. 2022. URL https://www.un.org/development/desa/pd/sites/www.un.org.development.desa.pd/files/wpp2022_summary_of_results.pdf.

[3] Lukas G Swan and V Ismet Ugursal. Modeling of end-use energy consumption in the residential sector: A review of modeling techniques. *Renewable and sustainable energy reviews*, 13(8): 1819–1835, 2009.

[4] Romain Nouvel, Kai-Holger Brassel, Marcel Bruse, Eric Duminil, Volker Coors, Ursula Eicker, and Darren Robinson. Simstadt, a new workflow-driven urban energy simulation platform for citygml city models. In *Proceedings of International Conference CISBAT 2015 Future Buildings and Districts Sustainability from Nano to Urban Scale*, number CONF, pages 889–894. LESO-PB, EPFL, 2015.

[5] Darren Robinson, Frédéric Haldi, Philippe Leroux, Diane Perez, Adil Rasheed, and Urs Wilke. Citysim: Comprehensive micro-simulation of resource flows for sustainable urban planning. In *Proceedings of the Eleventh International IBPSA Conference*, number CONF, pages 1083–1090, 2009.

[6] Jan LM Hensen and Roberto Lamberts. *Building performance simulation for design and operation*. Routledge, 2012.

[7] Paul Westermann, Matthias Welzel, and Ralph Evins. Using a deep temporal convolutional network as a building energy surrogate model that spans multiple climate zones. *Applied Energy*, 278:115563, 2020.

[8] Paul Westermann and Ralph Evins. Using bayesian deep learning approaches for uncertainty-aware building energy surrogate models. *Energy and AI*, 3:100039, 2021.

[9] Jose Vazquez-Canteli, Aysegul Dilsiz Demir, Julien Brown, and Zoltan Nagy. Deep neural networks as surrogate models for urban energy simulations. In *Journal of Physics: Conference Series*, volume 1343, page 012002. IOP Publishing, 2019.

[10] Bryan Lim, Sercan Ö Arık, Nicolas Loeff, and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4): 1748–1764, 2021.

[11] Aysegul Demir Dilsiz, Kingsley E Nweye, Allen J Wu, Jérôme H Kämpf, Filip Biljecki, and Zoltan Nagy. How spatio-temporal resolution impacts urban energy calibration. *Energy and Buildings*, 292:113175, 2023.

[12] Jan Remund, Stefan Müller, Michael Schmutz, and Pascal Graf. Meteonorm version 8. *METEOTEST (www. meteotest. com)*, 2020.

A Appendix

A.1 Input Variables in CitySim and CityTFT

Table 2 concludes the processed input for CityTFT. All the parameters in this table could be traced back to the geometries or some material paramters in the original CitySim XML file input.

A.2 Distribution of loads: With or Without zero

Fig. 3 indicate the distributions of heating and cooling load with and without zero loads. The upper two histograms are for heating loads while the lower two are suggesting cooling loads. Both heating and cooling distributions show less skew after trimming out zero values. In general, the model should be relatively easy to fit with the target distribution when the skewness is lower. The y-scale is also different in that the first bar in loads with zero demonstrates dominant height as the major distribution. This implies that when only employing regression loss to optimize the model, there is a large likelihood of the model becoming overly responsive to input variables with minimal values, consequently leading to the generation of predictions that systematically underestimate the overall energy demands.

Building Variables
Building height
Building perimeter
Wall glazing ratio
Footprint area
Heating setpoint temperature
Cooling setpoint temperature
Average walls U-value
Roof U-value
First floor U-value
Average windows U-value (W/m2k)
Average walls short-wave reflectance
Wall short-wave reflectance
Roof short-wave reflectance

Table 2: Building Properties used in the training dataset

Weather Variables
Day of month
Month
Hour
Diffuse radiation horizontal
Beam
Temperature
Surface temperature
Wind speed
Wind direction
Relative humidity
Precipitation
Cloud cover fraction

Table 3: Used Weather variables in TMY filest

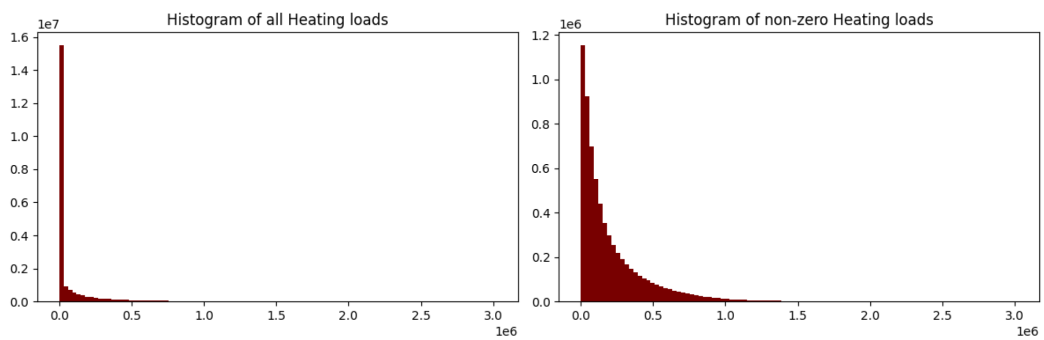

(a) Distribution of heating loads

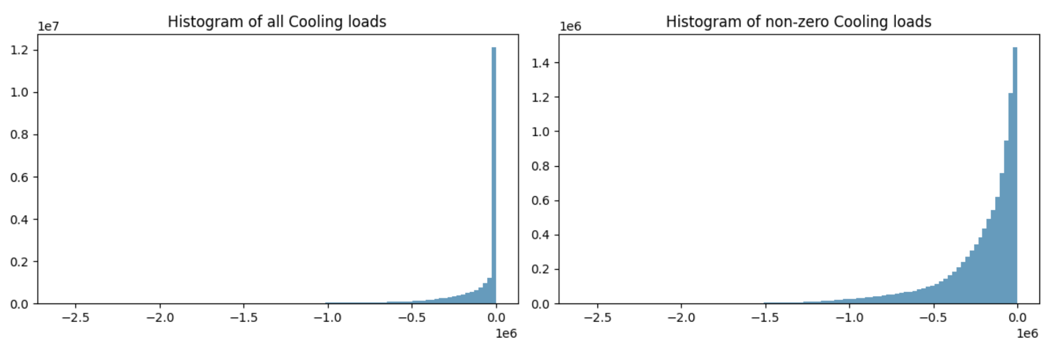

(b) Distribution of cooling loads

Figure 3: Distribution comparison between all loads and non-zero loads.

A.3 Climate data collection & Splitting datasets

Table 6 lists all the climate files used in this study.

A.4 Result metrics with only heating or cooling loads

While calculating the overall performance in the previous section, separate metrics for heating and cooling loads are also computed and shown in Table 4 and 5. Each of those models performs better in predicting heating loads than cooling, no matter in predicting triggering or in quantile regression.

	F1 score (%)	Non-zero RMSE (kWh)	Total RMSE (kWh)	Non-zero MAPE (%)
RNN	93.22	110.55	64.95	287.37
Transformer	93.05	95.19	61.30	205.88
TFT	**99.99**	**18.24**	**9.46**	**12.92**

Table 4: Heat only result

	F1 score (%)	Non-zero RMSE (kWh)	Total RMSE (kWh)	Non-zero MAPE (%)
RNN	90.10	115.78	85.48	61.86
Transformer	88.96	12.845	94.66	67.66
TFT	**99.97**	**22.73**	**16.70**	**10.97**

Table 5: Cool only result

Location	Lat	Lon	Elevation	Koppen Climate Zone	Climate Full Name	Train/Val/Test
Bogota/El-Dorado	4.7	-74.133	2547.0	C	Oceanic Climate (Warm Summer)	Train
Colombo	6.9	79.867	7.0	A	Tropical Monsoon Climate	Val
Dublin Airport	53.433	-6.233	82.0	C	Oceanic Climate (Warm Summer)	Val
Jacksonville Airp. FL	30.5	-81.7	9.0	C	Subtropical Humid Climate (Hot Summer)	Train
Key West FL	24.55	-81.75	1.0	A	Tropical Wet And Dry Climate (Winter Dry Season)	Train
Kinloss	57.65	-3.567	5.0	C	Oceanic Climate (Warm Summer)	Train
Kota Bahru	6.167	102.283	5.0	A	Tropical Monsoon Climate	Test
Marseille	43.433	5.217	6.0	C	Mediterranean Climate (Hot Summer)	Train
Matamoros Intl	25.76	-97.53	8.0	B	Hot Semi-Arid Climate (Steppe)	Train
Medford/Jackson Co.	42.367	-122.867	405.0	C	Mediterranean Climate (Warm Summer)	Train
Mobile AL	30.683	-88.25	67.0	C	Subtropical Humid Climate (Hot Summer)	Train
Ocotal	13.617	-86.467	612.0	A	Tropical Monsoon Climate	Train
Penang/Bayan L.	5.3	100.267	3.0	A	Tropical Monsoon Climate	Train
Portland OR	45.567	-122.717	12.0	C	Mediterranean Climate (Warm Summer)	Train
Redding CA	40.5	-122.3	153.0	C	Mediterranean Climate (Hot Summer)	Test
Reykjavik	64.133	-21.9	66.0	C	Subpolar Oceanic Climate (Cool Summer)	Train
Sacramento Airp. CA	38.517	-121.5	7.0	C	Mediterranean Climate (Hot Summer)	Test
Tucson AZ	32.117	-110.933	779.0	B	Hot Semi-Arid Climate (Steppe)	Test
Tunis	36.833	10.233	3.0	C	Mediterranean Climate (Hot Summer)	Train
Victoria Airp. TX	28.85	-96.917	32.0	C	Subtropical Humid Climate (Hot Summer)	Val
West Palmbeach Airp.	26.683	-80.117	5.0	A	Tropical Wet And Dry Climate (Winter Dry Season)	Train

Table 6: Location information of TMY climate file for CitySim and CityTFT