Context-Aware Urban Energy Efficiency Optimization Using Hybrid Physical Models

Benjamin Choi
Department of Electrical Engineering
Stanford University
benchoi@stanford.edu

Alex Nutkiewicz
Department of Civil and Environmental Engineering
Stanford University
alexer@stanford.edu

Rishee Jain
Department of Civil and Environmental Engineering
Stanford University
rishee.jain@stanford.edu

Abstract

Buildings produce more U.S. greenhouse gas emissions through electricity generation than any other economic sector [1]. To improve the energy efficiency of buildings, engineers often rely on physics-based building simulations to predict the impacts of retrofits in individual buildings. In dense urban areas, these models suffer from inaccuracy due to imprecise parameterization or external, unmodeled urban context factors such as inter-building effects and urban microclimates [2, 3, 4]. In a case study of approximately 30 buildings in Sacramento, California, we demonstrate how our hybrid physics-driven deep learning framework can use these external factors advantageously to identify a more optimal energy efficiency retrofit installation strategy and achieve significant savings in both energy and cost.

1 Introduction

Urban areas dominate contributions to anthropogenic climate change, with cities accounting for over 75 percent of greenhouse gas emissions and two-thirds of global energy use [5]. A prime opportunity to reduce urban greenhouse gas emissions is in the buildings sector, which is responsible for over 30 percent of greenhouse gases from energy generation [5]. When deployed strategically, building energy retrofits are widely accepted as a cost-effective way to achieve energy savings in buildings [6].

Engineers often rely on physics-based simulation tools to understand building energy use and to project possible energy savings under a variety of different retrofits. These methods tend to focus on prediction tasks for individual buildings, requiring a separate model for each building, where the time and resources required to produce each model makes it difficult to model an entire urban area that municipal policymakers require to make decisions. Additionally, these models are often prone to high prediction errors because of a large number of required modeling assumptions [7, 8]. However, emerging sophisticated time series prediction models using supervised machine learning methods have been applied to accurately predict energy consumption in buildings [9, 10, 11, 12, 13]. Fully data-driven approaches to predicting impacts of energy retrofits are difficult to use in practice due to the lack of ground truth labels from buildings with comparable retrofits. We draw inspiration from the emergence of models that integrate physical simulations with deep learning [14]. Other work [15, 16] has used a combination of physical simulations and machine learning to match buildings with potentially effective retrofits, but these models do not consider the impacts of urban context (e.g., mutual shading and reflection, building-to-building heat transfer) on retrofit efficacy, which has been previously shown to significantly influence building energy performance [17].

Tackling Climate Change with Machine Learning Workshop at NeurIPS 2020

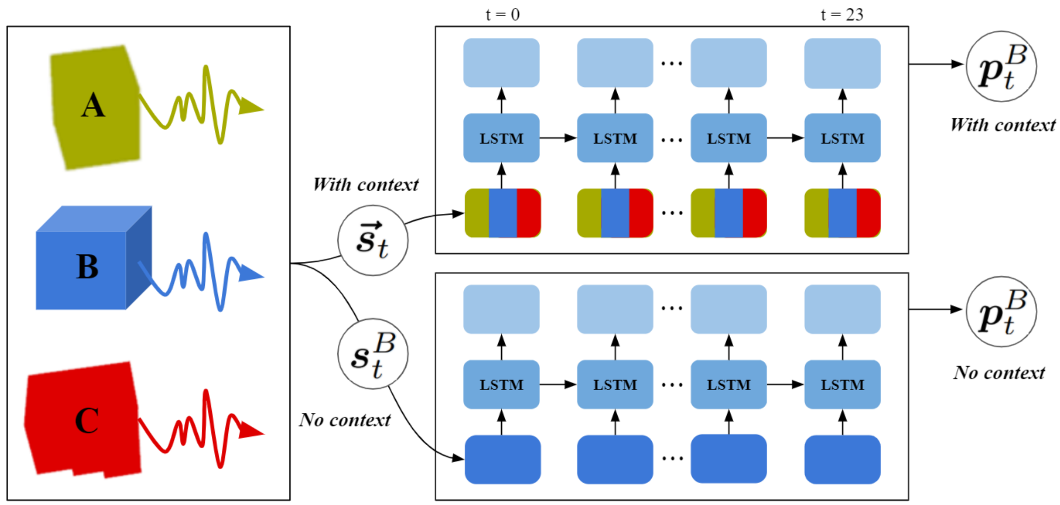

Figure 1: From left to right, 2.5D geometry and other building parameters are used by EnergyPlus to simulate energy consumption. This is the input to either the no context or with context model. The model then predicts energy consumption for the target building (in this case, building B).

Here, we extend the hybrid data-driven urban energy simulation (DUE-S) model presented in [18] by adapting it to time series prediction tasks and enabling quantification of potential retrofits. We use this new model to predict urban context-aware energy usage estimates given static physical parameters for each building in an urban area. We also isolate the effect of context to highlight its utility in retrofit optimization. A key advantage of our approach is the automatic identification of buildings that can, with strategic retrofitting, achieve a "multiplier effect" by improving energy savings in other buildings via inter-building effects. Our contributions are twofold: the application of a hybrid deep learning approach to 1) account for urban context to simultaneously predict retrofit installation effects in multiple buildings and 2) derive a retrofit strategy that minimizes installation demands.

2 Methodology

2.1 Data

Ground truth labels for our baseline model training consist of hourly electricity consumption data for 2016-2018 from $N = 29$ buildings in Sacramento, California. Buildings range in size and purpose (office, retail, warehouse, apartment). To avoid data leakage from random shuffling of sequential time series data [19], we split the first two years into consecutive sequences of training and validation sets with lengths in a ratio of 3:1 respectively. The last year (2018) is reserved for testing. For each ground truth label, a corresponding simulated energy consumption estimate is generated using EnergyPlus – a physics-based building energy simulation engine [20]. This model uses historical weather data and typical building-level information on materials, operating schedules, and mechanical systems to produce hourly predictions of energy usage by non-retrofitted buildings in our study period. These are based on archetypal building characteristics from DOE commercial reference buildings [21]. We also generate "retrofit" simulations by varying physical building parameters based on retrofit recommendations often adopted by commercial or mixed-use buildings [22]. We apply three different retrofits: a window retrofit, a lighting retrofit, and a "full" retrofit that combines the two. The window retrofit modulates thermal transmittance (U-value) and the lighting retrofit modulates lighting power density.

2.2 Context-Aware Time Series Prediction

The first step in our hybrid physical modeling approach is the training of a baseline deep learning model that aims to capture discrepancies between simulation predictions and observed energy consumption associated with urban context effects. We train and validate our model via hindcasting – using historical simulations as inputs to predict historical observed energy usage. To capture

	Building	Urban
Hourly	23.2%	9.9%
Daily	14.8%	5.6%
Monthly	9.8%	1.4%

Table 1: Mean average percent error of *with context* model across spatiotemporal scales

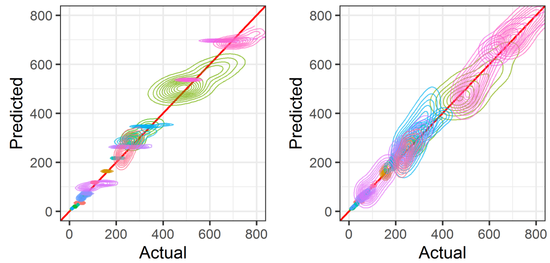

Figure 2: (a) Density plot of predictions and actual values for *no context* model. (b) Density plot for *with context* model. Horizontal features absent, indicating advantage of reduced model bias.

dependencies across multiple timesteps, we adopt a model suitable for a many-to-many prediction task. We input the 24 most recent simulation timesteps $\boldsymbol{S} = [\vec{\boldsymbol{s}}_{t-23} \cdots \vec{\boldsymbol{s}}_t]$ (one full day) to generate predictions $\boldsymbol{P}$. Although we use just one model to predict energy consumption for all 29 buildings, we optimize our model from a single loss curve to avoid difficulties in model training diagnosis stemming from multiple optimization objectives and loss curves. We accomplish this by one-hot encoding the target building for prediction as a model input. Therefore, for each target building b_k, the model output at each timestep is a scalar prediction $\boldsymbol{p}_t^k$ instead of the length N output vector $\vec{\boldsymbol{p}}_t$. Optimization is performed using an Adam optimizer with a learning rate of 1e-3, chosen for its adaptive learning rate and invariance to scaling of the objective function.

A variant of the recurrent neural network (RNN) called the long-short term memory (LSTM) model [23] is chosen for its resistance to exploding and vanishing gradients during training and adaptability to many-to-many prediction tasks. We also evaluate more complex models using convolutional neural networks (CNN-LSTM) and LSTM autoencoders [24] but in practice find that a vanilla LSTM is an appropriate level of complexity. All subsequent experiments therefore use 2 sequential LSTM 64 unit layers and 2 time-distributed fully connected layers configured for a many-to-many prediction task.

Urban context is implicitly modeled by inputting all simulation outputs $\vec{\boldsymbol{s}}_t$ regardless of the target building b_k. The model is then free to learn the time series characteristics of different retrofits, as well as complex interdependencies between physics-based simulations of multiple buildings. Using a single model for prediction of multiple buildings is complicated by the wide range of energy use observed across buildings of different sizes. For each building, we seek an absolute error value roughly proportional to its average energy use, so we choose mean absolute error loss in our training loop for its intermediacy between mean absolute percentage error and mean squared error. The former optimizes for low-energy buildings with proportionally larger noise, while the latter penalizes the relatively larger errors in high-energy buildings disproportionately.

We isolate the effect of urban context by training an additional model "without context." Concretely, instead of making simulation outputs of all buildings accessible to the model in each prediction, only the simulation output $\boldsymbol{s}_t^k$ corresponding to the target building b_k is given. After *with context* and *no context* models are trained, their outputs are compared to characterize the effects of urban context.

2.3 Using Context to Optimize Retrofit Strategy

Next, we create a retrofit installation plan for our 29-building case study that uses urban context to optimize energy savings and reduce costs. Ideally, we can identify a retrofitted subset of buildings that maximizes savings while minimizing retrofits. Using our model trained on a baseline dataset of non-retrofitted simulations and ground truth observations, we can simulate retrofitting different building subsets and output new energy consumption predictions. In a set of n elements, there are 2^n possible subsets, making a brute force approach to constructing and testing every subset of 29 buildings intractable. In contrast, evaluating buildings individually for their retrofit potential ignores aggregate inter-building effects that may arise when multiple buildings are retrofitted.

We design a greedy optimization algorithm that aims to maximize cumulative energy savings across buildings. For each building candidate, our model predicts the marginal change in energy savings

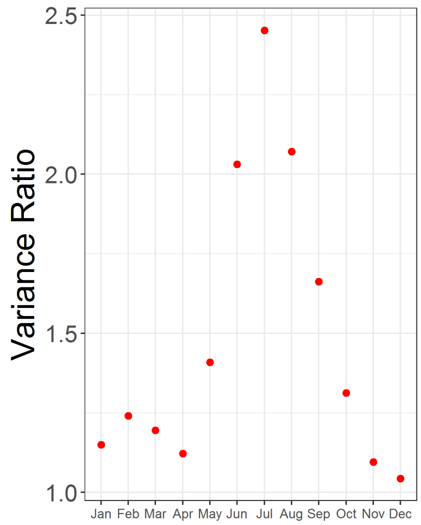

Figure 3: Ratio of monthly variance in the *with context* and *no context* models.

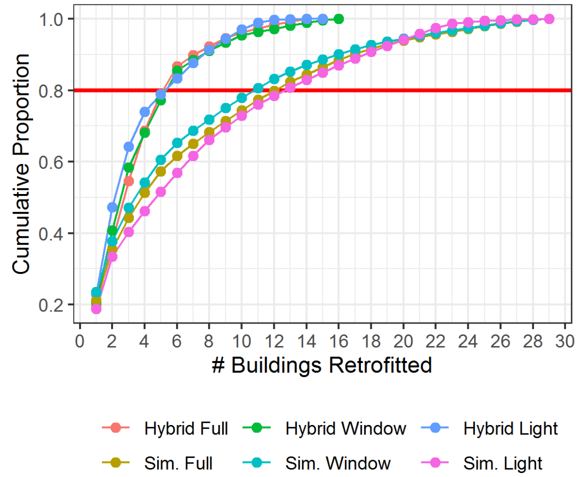

Figure 4: Cumulative proportion of maximum possible energy savings for each retrofit and model combination.

achieved from its addition to the retrofitted subset. The candidate with the greatest marginal savings is permanently added to the subset. The algorithm continues until all there are no building candidates left that result in marginal energy savings. We find the global maximum in energy savings for building subsets of size $n \leq 4$ through brute force and find that our method converges on an equivalent optimal subset of buildings. Future work may identify strategies to avoid converging on local optima.

We compare this performance to a naive procedure in which buildings are simply selected sequentially according to decreasing projected savings in the energy usage simulation alone.

3 Results

Context-aware models improve accuracy at multiple scales: Our model achieves 23.2% mean average percentage error on energy usage prediction at an hourly, individual building resolution. Performance at less granular spatiotemporal scales can be viewed in Table 1. In comparison, *no context* models suffer from approximately 5% greater error across spatiotemporal scales. Regions of horizontal density in Figure 2(a) convey narrow prediction ranges in *no context* models, suggesting an excessively biased model due to missing urban context. In contrast, *with context* models are able to capture this lost variation, as shown in Figure 2(b). In this case, our hybrid physical model demonstrates consistent performance across all buildings in our case study despite highly variable energy footprints and usage patterns. Figure 3 emphasizes how the *with context* model is able to capture greater seasonal variance arising from urban context effects.

Context can be used to reduce retrofit installation requirements: The maximum projected savings across all buildings is 14.4% in simulation and 13.8% for the hybrid physical model. While comparable in magnitude, these savings are achieved through very different means. According to Figure 4, in the case of a full retrofit, 80% of these savings are achieved by retrofitting 11 buildings in the naive simulation case. In contrast, our hybrid physical model suggests that only 6 buildings need to be retrofitted to break the 80% savings mark across all retrofit types tested. Among these 6 buildings, savings are on average 7.9 times greater than when isolated in simulation, emphasizing how improved modeling capacity of our model can help expose latent "multiplier" effects attributed to urban context. This should be regarded as an upper bound for multipliers attributed to urban context as it is possible that other latent variables are partially responsible for predicted savings.

While not necessarily a global optimum, our results suggest that policymakers may greatly benefit from using hybrid physical models to inform retrofit strategies that maximize energy savings while reducing the costs and logistical challenges associated with retrofitting a large number of buildings.

4 Conclusion and Future Work

The dominance of buildings among contributors to greenhouse gases greatly augments the potential importance of retrofits in mitigating climate change. While simulation and deep learning have individually made significant progress in characterizing and highlighting buildings with a current *deficit* in energy efficiency, we demonstrate the potential of hybrid physical models in identifying buildings where multiplier effects may yield a *surplus* of energy efficiency through careful modeling of the urban context. Further validation on other buildings, cities, and retrofit types is required to gain confidence in the usefulness of our approach to wide variety of urban environments worldwide.

Appendix

A 3D model of the buildings used in our study area can be found in Figure 5. Specific details on the retrofit parameters used in our experiments are in Figure 6.

We also further characterize the effect of geographic proximity on retrofit influence via an additional "block" experiment. Instead of retrofitting all buildings in our case study or retrofitting buildings in stepwise increments, we choose a "block" of proximal buildings to retrofit. We then evaluate deviation in energy use of non-retrofitted buildings from the baseline. A slight tendency for retrofitted buildings to exert less influence on distant buildings is observed in Figure 7, but a larger dataset is necessary to substantiate any claims.

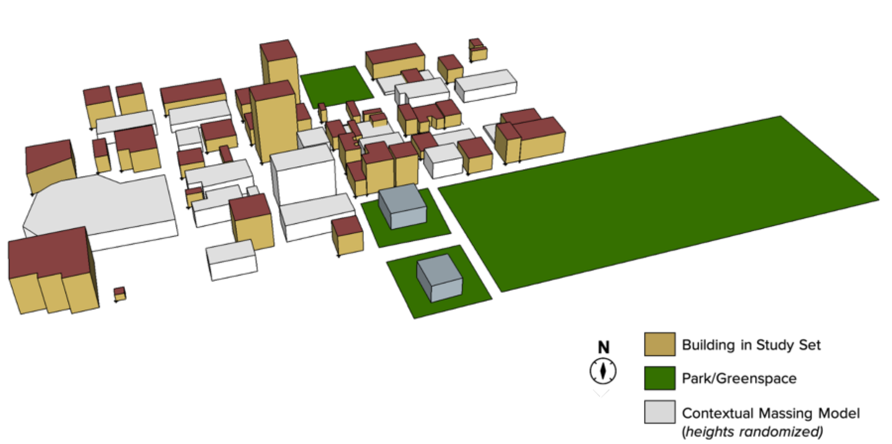

Figure 5: Buildings in our Sacramento, California study area.

Retrofit Type	Baseline Scenario	Retrofit Scenario
"Lighting"	DOE Commercial Reference Buildings (dependent on specific building)	Switch to LED bulbs (~27% decrease in LPD)
"Windows"		Updated to ASHRAE 90.1-2010 standard materials
"Full"		Include both lighting and window retrofits

Figure 6: Detailed information on retrofits used in experimentation.

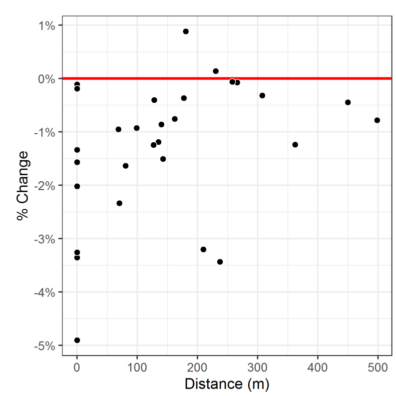

Figure 7: % Change in energy use (negative is savings) vs. Distance to the nearest retrofitted building.

References

[1] *Inventory of U.S. Greenhouse Gas Emissions and Sinks 1990-2018*. U.S. Environmental Protection Agency, 2018.

[2] Ermanno Lo Cascio, Zhenjun Ma, Davide Borelli, and Corrado Schenone. Residential building retrofit through numerical simulation: A case study. *Energy Procedia*, 111:91–100, March 2017.

[3] Yixing Chen, Tianzhen Hong, and Mary Ann Piette. Automatic generation and simulation of urban building energy models based on city datasets for city-scale building retrofit analysis. *Applied Energy*, 205:323–335, November 2017.

[4] Jun Wei Chuah, Anand Raghunathan, and Niraj K. Jha. ROBESim: A retrofit-oriented building energy simulator based on EnergyPlus. *Energy and Buildings*, 66:88–103, November 2013.

[5] UNEP (2014). *Climate Finance for Cities and Buildings - A Handbook for Local Governments*. UNEP Division of Technology, Industry and Economics (DTIE).

[6] O Lucon, D Ürge-Vorsatz, A Zain Ahmed, H Akbari, P Bertoldi, LF Cabeza, N Eyre, A Gadgil, LDD Harvey, Y Jiang, et al. Buildings climate change 2014: Mitigation of climate change. contribution of working group iii to the fifth assessment report of the intergovernmental panel on climate change. In *Intergovernmental Panel on Climate Change, Cambridge*. 2014.

[7] Aneta Strzalka, Jürgen Bogdahn, Volker Coors, and Ursula Eicker. 3d city modeling for urban scale heating energy demand forecasting. *HVAC&R Research*, 17(4):526–539, 2011.

[8] Daniel Coakley. *Calibration of detailed building energy simulation models to measured data using uncertainty analysis*. PhD thesis, 05 2014.

[9] Rishee K Jain, Kevin M Smith, Patricia J Culligan, and John E Taylor. Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy. *Applied Energy*, 123:168–178, 2014.

[10] Geoffrey K.F. Tso and Kelvin K.W. Yau. Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks. *Energy*, 32(9):1761–1768, September 2007.

[11] Salah Bouktif, Ali Fiaz, Ali Ouni, and Mohamed Adel Serhani. Optimal Deep Learning LSTM Model for Electric Load Forecasting using Feature Selection and Genetic Algorithm: Comparison with Machine Learning Approaches. *Energies; Basel*, 11(7):1636, 2018. Num Pages: 1636 Place: Basel, Switzerland, Basel Publisher: MDPI AG.

[12] Chengdong Li, Zixiang Ding, Dongbin Zhao, Jianqiang Yi, and Guiqing Zhang. Building Energy Consumption Prediction: An Extreme Deep Learning Approach. *Energies*, 10(10):1525, October 2017. Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.

[13] Daniel L. Marino, Kasun Amarasinghe, and Milos Manic. Building energy load forecasting using Deep Neural Networks. In *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, pages 7046–7051, October 2016.

[14] Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, et al. Deep learning and process understanding for data-driven earth system science. *Nature*, 566(7743):195–204, 2019.

[15] Gaby M Baasch and Ralph Evins. Targeting buildings for energy retrofit using recurrent neural networks with multivariate time series.

[16] Usman Ali, Mohammad Haris Shamsi, Mark Bohacek, Cathal Hoare, Karl Purcell, Eleni Mangina, and James O'Donnell. A data-driven approach to optimize urban scale energy retrofit decisions for residential buildings. *Applied Energy*, 267:114861, June 2020.

[17] Anna Laura Pisello, John E. Taylor, Xiaoqi Xu, and Franco Cotana. Inter-building effect: Simulating the impact of a network of buildings on the accuracy of building energy performance predictions. *Building and Environment*, 58:37–45, December 2012.

[18] Alex Nutkiewicz, Zheng Yang, and Rishee K. Jain. Data-driven urban energy simulation (due-s): A framework for integrating engineering simulation and machine learning methods in a multi-scale urban energy modeling workflow. *Applied Energy*, 225:1176 – 1189, 2018.

[19] Christoph Bergmeir and José M. Benítez. On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191:192–213, May 2012.

[20] Drury B. Crawley, Linda K. Lawrie, Frederick C. Winkelmann, W.F. Buhl, Y.Joe Huang, Curtis O. Pedersen, Richard K. Strand, Richard J. Liesen, Daniel E. Fisher, Michael J. Witte, and Jason Glazer. EnergyPlus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4):319–331, April 2001.

[21] Michael Deru, Kristin Field, Daniel Studer, Kyle Benne, Brent Griffith, Paul Torcellini, Bing Liu, Mark Halverson, Dave Winiarski, Michael Rosenberg, et al. Us department of energy commercial reference building models of the national building stock. 2011.

[22] Guopeng Liu, Bing Liu, Weimin Wang, Jian Zhang, Rahul A. Athalye, Dave Moser, Eliot Crowe, Nick Bengtson, Mark Effinger, Lia Webster, and Mike Hatten. Advanced energy retrofit guide office buildings. 9 2011.

[23] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[24] Tae-Young Kim and Sung-Bae Cho. Predicting residential energy consumption using CNN-LSTM neural networks. *Energy*, 182:72–81, September 2019.