A neural network approach to predicting urban building energy consumption

Rohan Aras, Alex Nutkiewicz, Max O'Krepki
{rohana, alexer, maxo}@stanford.edu

OVERVIEW

Cities account for over 75% of all primary energy usage and over 80% of greenhouse gas emissions, with a large majority of each coming from the built environment [1]. Because 90% of urban buildings are estimated to be energy inefficient and 30% of a building's energy consumption is wasted, buildings represent a great opportunity to enhance the sustainability of cities [2].

Building energy modeling (BEM) is a tool widely used in building design and construction to understand and predict the energy performance of buildings. However, these tools are largely inaccurate when predicting energy consumption because they are unable to account for a building's "urban context" - both surrounding buildings and its urban microclimate.

This project explores multilayer perceptrons (MLP) and ResNets to integrate deep learning into an urban building energy simulation to more accurately predict the energy consumption of urban buildings.

Energy Simulation Model

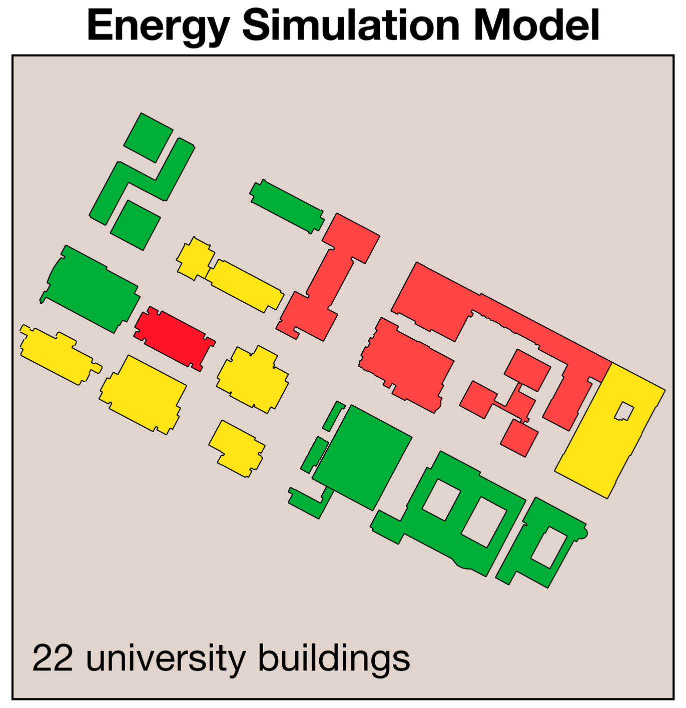

Network Model

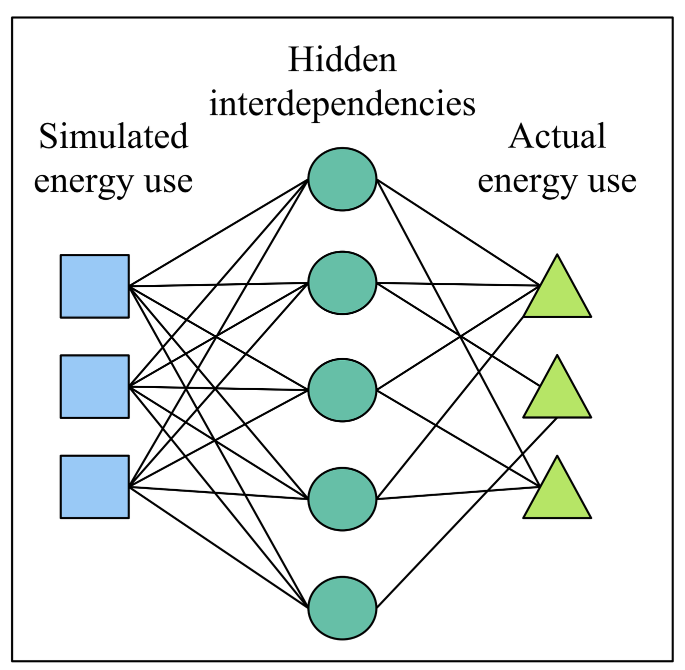

DATA & FEATURE ENGINEERING

Our project studies a dense grouping of 22 buildings at a university in southern California. We draw data from three sources:

- EnergyPlus: building energy simulations (kWh/interval)
- NREL: Weather data
- University: metered building energy use (kWh/interval)

The energy data comes in the form of 15-minute interval time series, where each feature represents one of the 22 buildings - each with two years of interval data (~70,080 observations). While the outputs for both the MLP and ResNet are the same (22 features), the inputs vary.

`Only for the MLP do we include the hourly weather data as part of the input space, resulting in 25 features. And for the ResNet model, we one-hot encode *day of the week* and *month* into our input space, resulting in 41 initial input features.

METHODS

MULTILAYER PERCEPTRON (MLP)

Our baseline of comparison is a basic MLP that consists of three fully connected layers, containing a hidden layer with 24 neurons - one with and one without using weather data as an input. These models are formalized as:

$$y = \mathbf{W} \cdot \mathbf{x} + \mathbf{b}$$

For our implementation, we experimented with both ReLU and logistic activation functions as well as the number of hidden layers within the model. We used a learning rate of 0.01, determined by cross-validation, and MSE as the model's loss function.

RESIDUAL NETWORK (ResNet)

Deep residual networks (ResNets) are multi-layer neural networks in which each layer consists of a residual module f_i and an identity mapping connection that can bypass f_i [3], [4]. Because layers in a residual network can consist of multiple convolutional layers, they are often referred to as "residual blocks." Residual blocks attempt to learn

$$H(x) = x + F(x)$$

where $F(x)$ is a series of convolution layers with batch normalization and leaky ReLU activation, and x is the original input.

For our implementation we shape our input data into a signal of 96 time steps by 41 channels (buildings + one-hot vectors). This input is fed into a 1D-convolutional layer that outputs 64 channels. This output is passed through K residual blocks, each containing two 1D convolutions, pre-activated as in [3]. A diagram of this can be seen in Figure 2. The output of these blocks is passed through a final 1D convolution that outputs 22 channels. This output is added back to the building channels of the original inputs for the final output of the model.

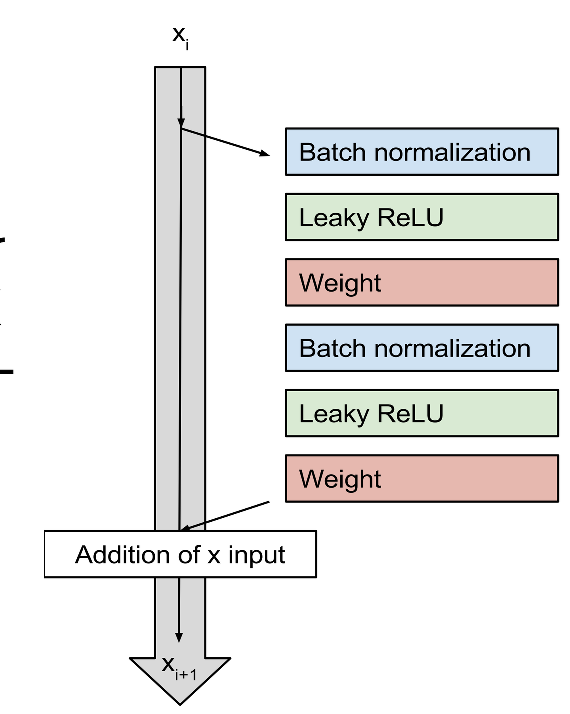

Figure 2. Pre-activation function

Figure 3. The relationship between buildings and input dataset on ResNet.

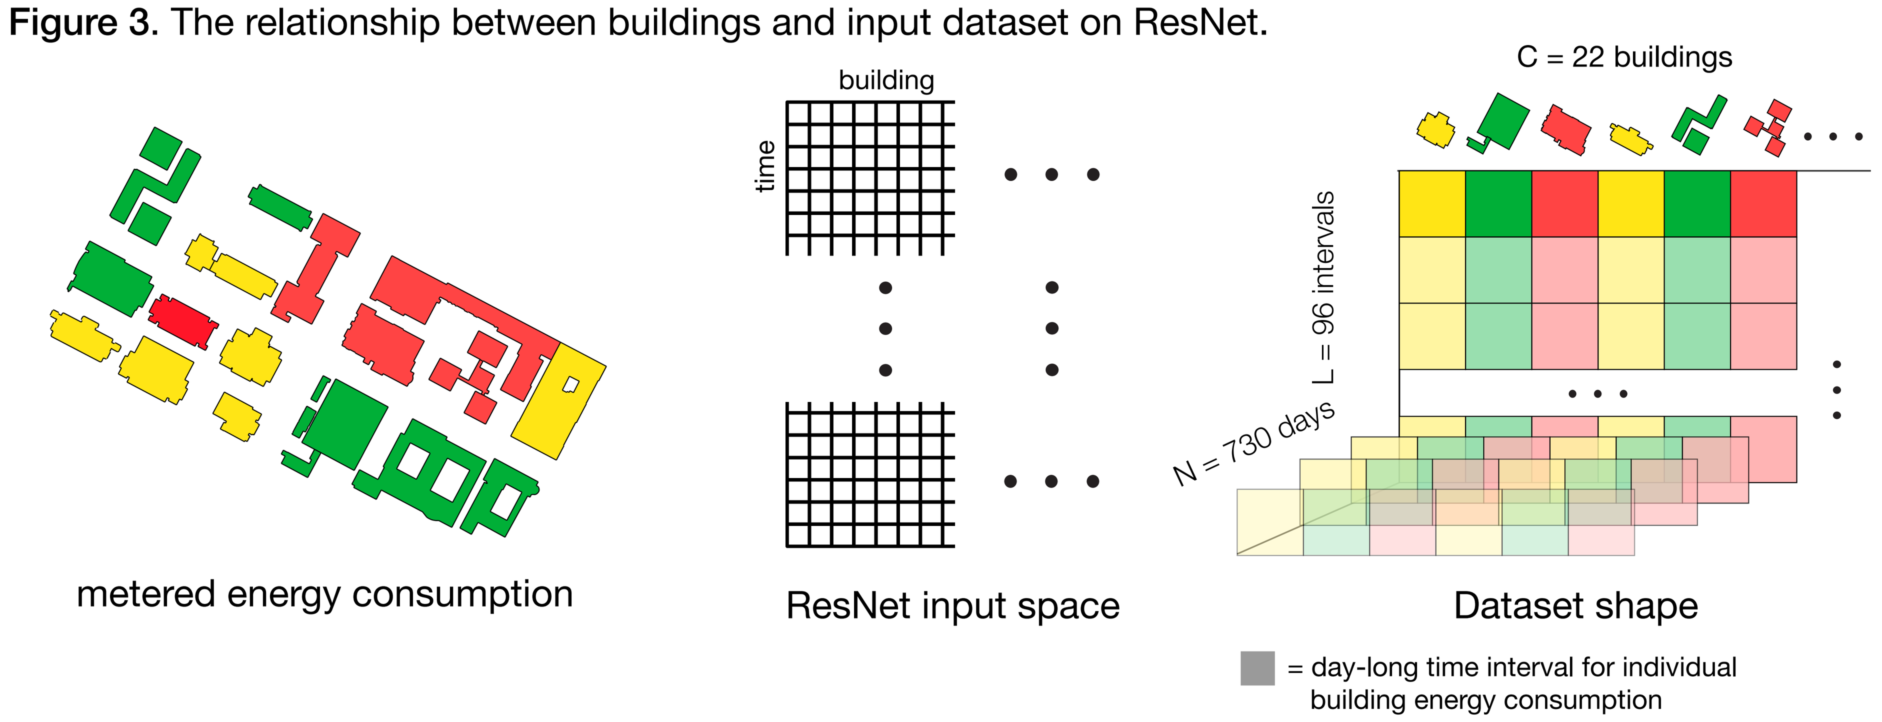

RESULTS

Below shows the performance of each model, evaluated with MSE.

Evaluated Model	# Samples	MSE
Simulation data, no neural network model	N/A	38.500
MLP, single layer, ReLU, without weather	730	36.419
MLP, single layer, ReLU, with weather	730	36.418
MLP, three layers, logistic, with weather	730	36.418
ResNet, 4 residual blocks, 64 output channels	70,080	29.391
ResNet, 8 residual blocks, 64 output channels	70,080	30.600
ResNet, 4 residual blocks, 128 output channels	70,080	32.602
ResNet, 8 residual blocks, 128 output channels	70,080	27.820

Costs of Evaluated ResNet Models

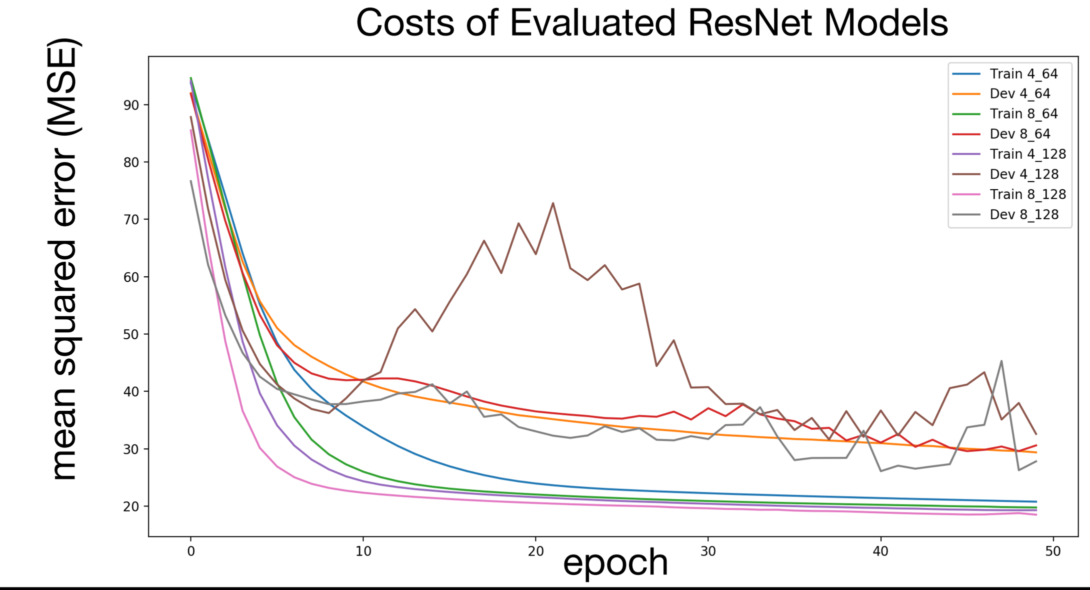

DISUCSSION & FUTURE WORK

The above results show that overall, the ResNet significantly outperformed the MLP, as expected. Because of the residual network's local convolutions, we believe it did the better job of correctly understanding the spatial-temporal relationships between urban context and energy usage. We were surprised that weather had little impact on the performance of the MLP and plan to explore this more for our final deliverable. And finally, while adding additional residual blocks and output channels improves train set performance, it also makes the dev set performance more unstable.

Our future work on this project aims to reduce the mean squared error further and evaluate it against other simulation-based approaches to predicting energy consumption. We would also like to evaluate these errors at different time intervals (e.g., hourly, yearly) and spatial scales (e.g., individual building scale, city-wide scale) to further test its prediction ability.

REFERENCES

[1] United States Energy Information Administration, "Annual Energy Outlook 2017," 2017.
[2] Breaking Energy, "The Top Ten Ways We Waste Energy in Water and Buildings," 2011.
[3] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," Lect. Notes Comput. Sci., vol. 9908 LNCS, pp. 630–645, 2016.
[4] S. Wu, S. Zhong, and Y. Liu, "Deep residual learning for image steganalysis," Multimed. Tools Appl., pp. 1–17, 2017.