Energy Efficiency

The role of buildings and cities

Assignments

Brightspace discussion question:

“What do you think is the biggest contributor to your own energy consumption and are there ways you could reduce it?”

Due Friday by 5pm.

First programming assignment:

Climate data visualization

Due Fri Feb 3 by midnight.

Climate change in the news

Climate change in the news

Reading Kit Climate Stripes:
https://twitter.com/ed_hawkins/status/1619456016638427136

Recap

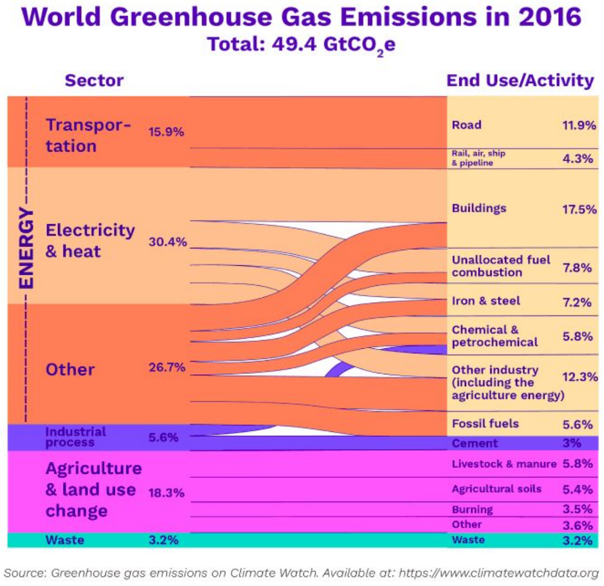

The Intergovernmental Panel on Climate Change

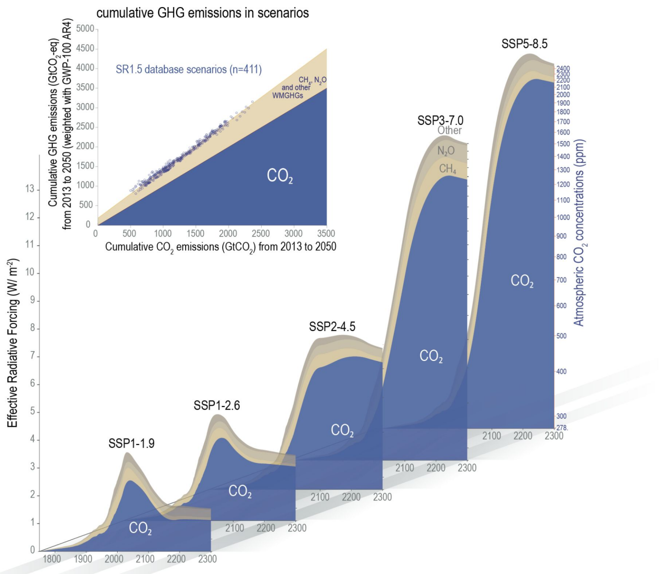

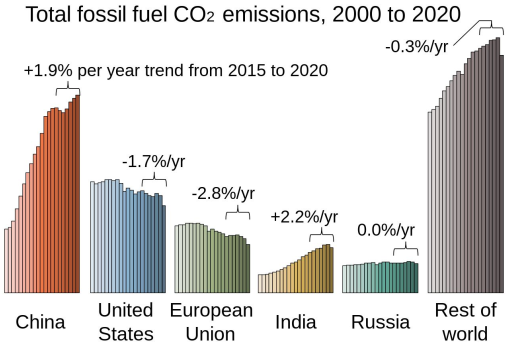

We need to rapidly stop releasing GHGs.

Energy efficiency as a way to reduce GHGs

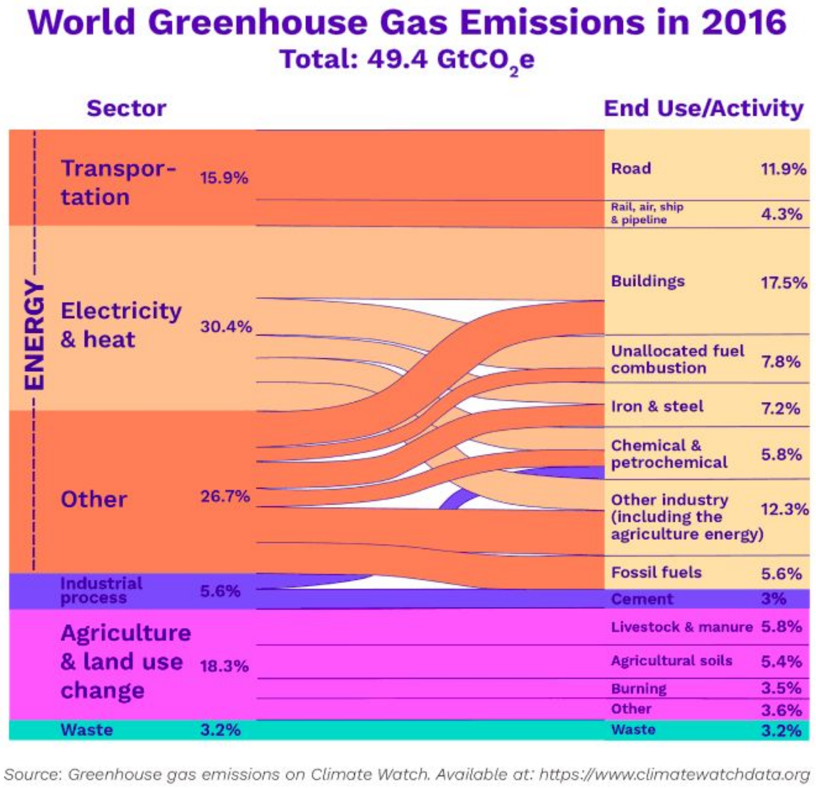

Buildings and Energy

EIA

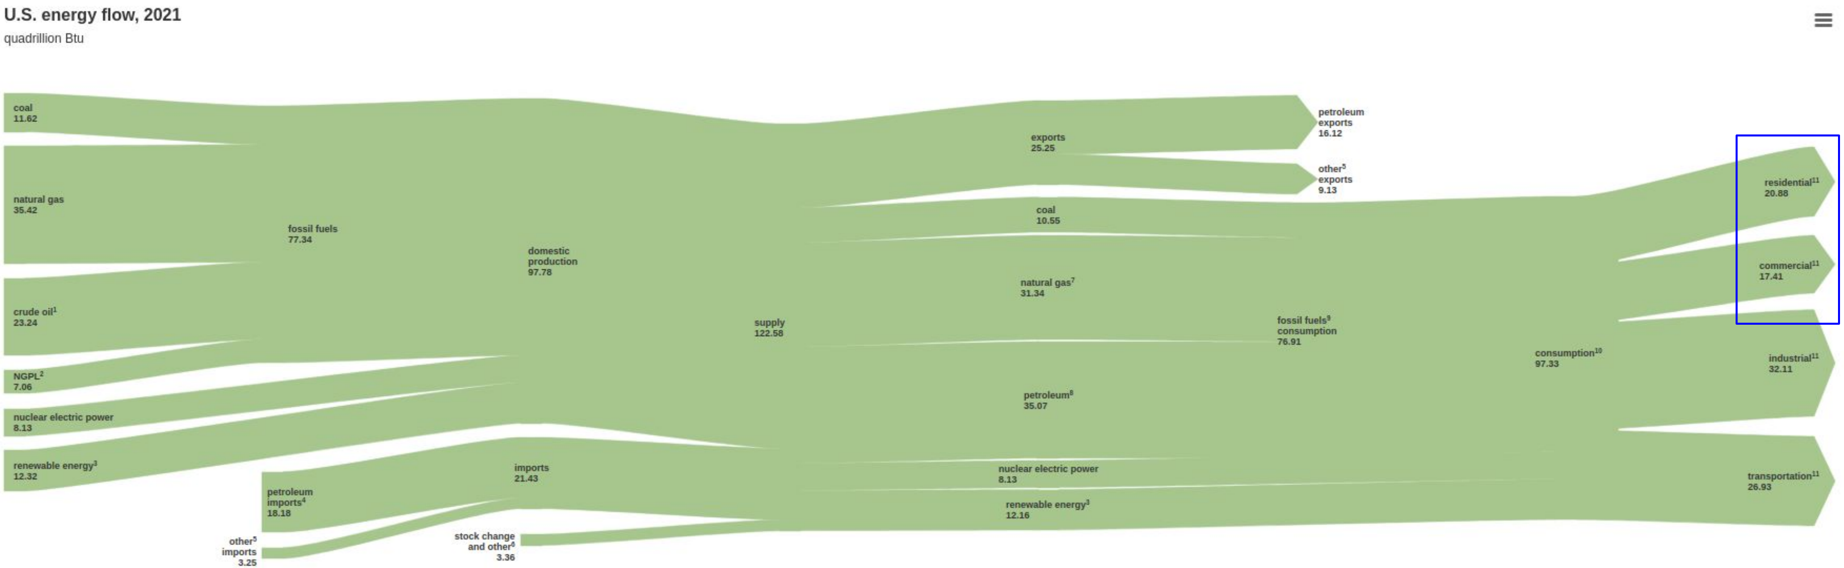

Residential and commercial building energy consumption makes up 40% of the total energy consumption in the United States.

Energy is measured in British thermal unit, defined as the amount of heat required to raise the temperature of one pound of water by one degree Fahrenheit

IPCC

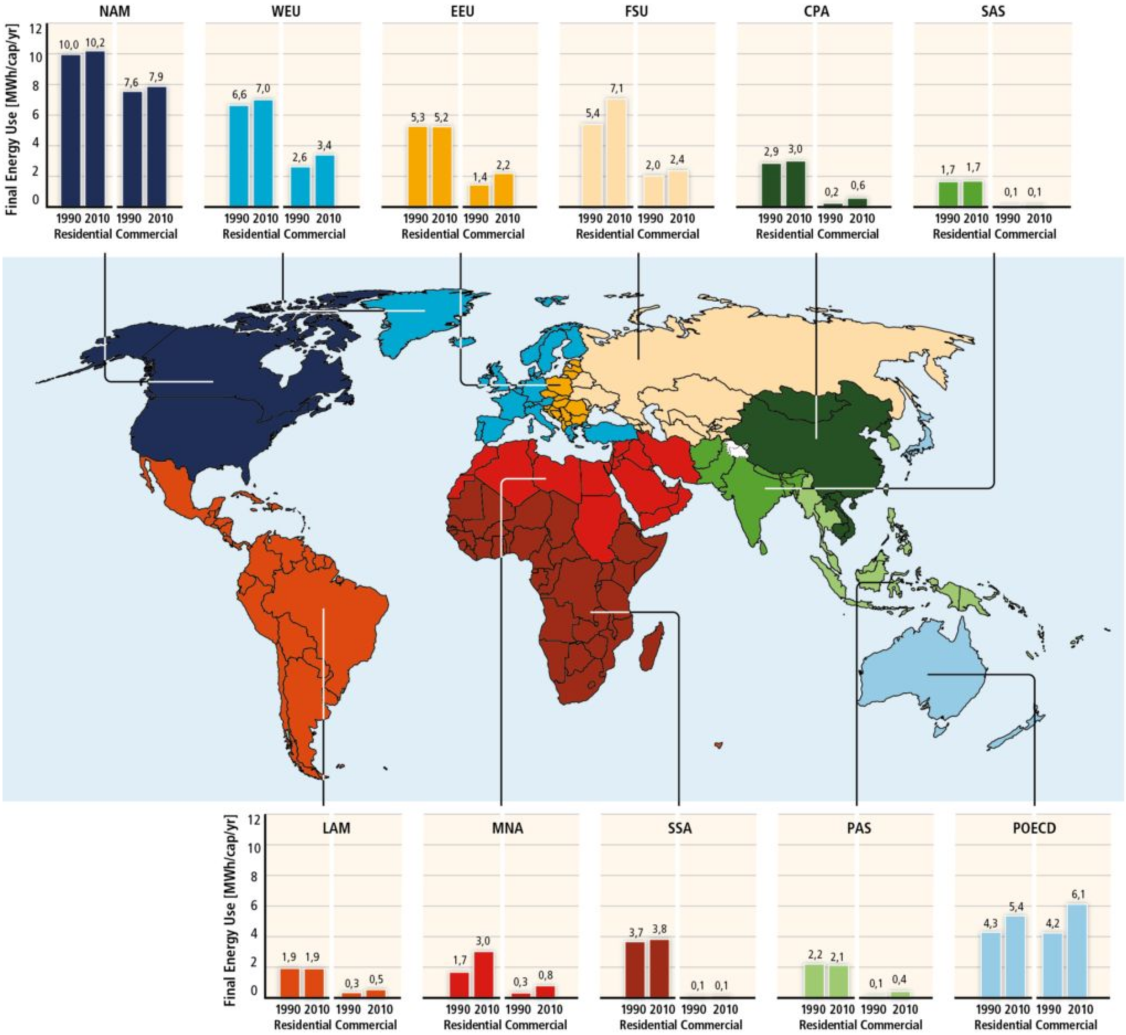

Different practices across different countries lead to different energy use intensities.

The fraction of energy used by buildings is similar in the EU (40%) but lower on average across the globe (32%).

Future Projections

IPCC

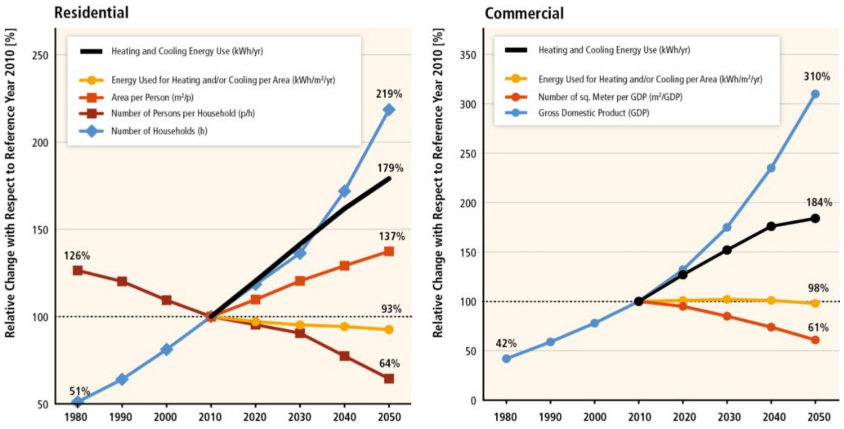

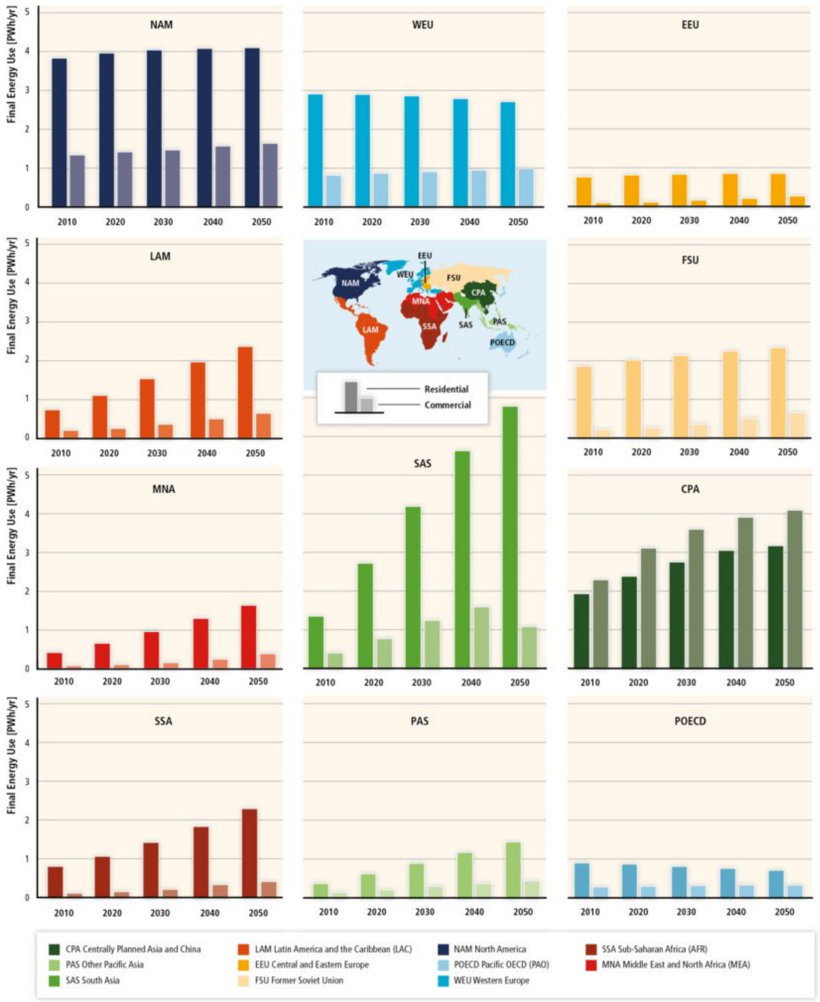

How is that energy being used?

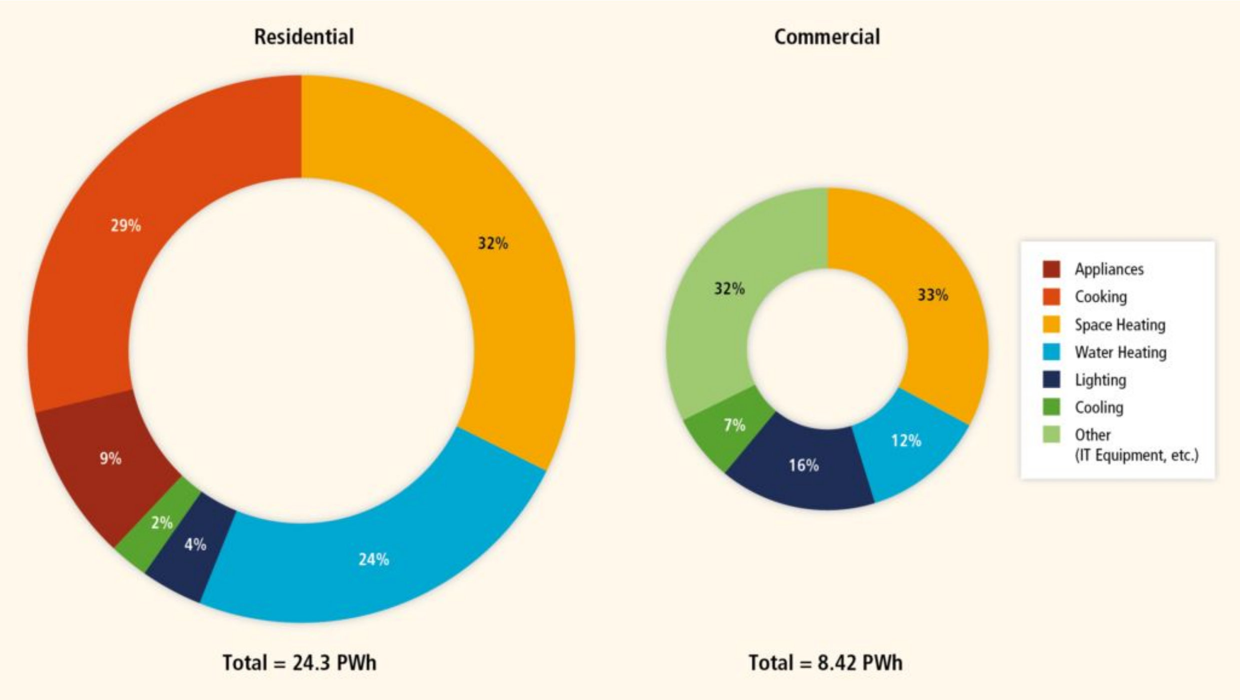

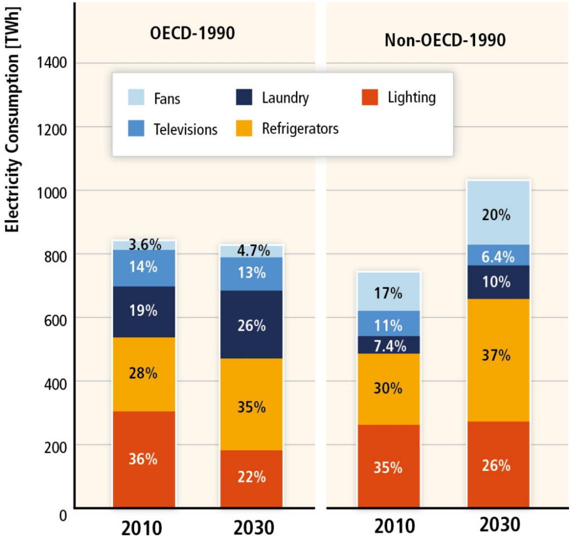

NYC Building Efficiency Grades

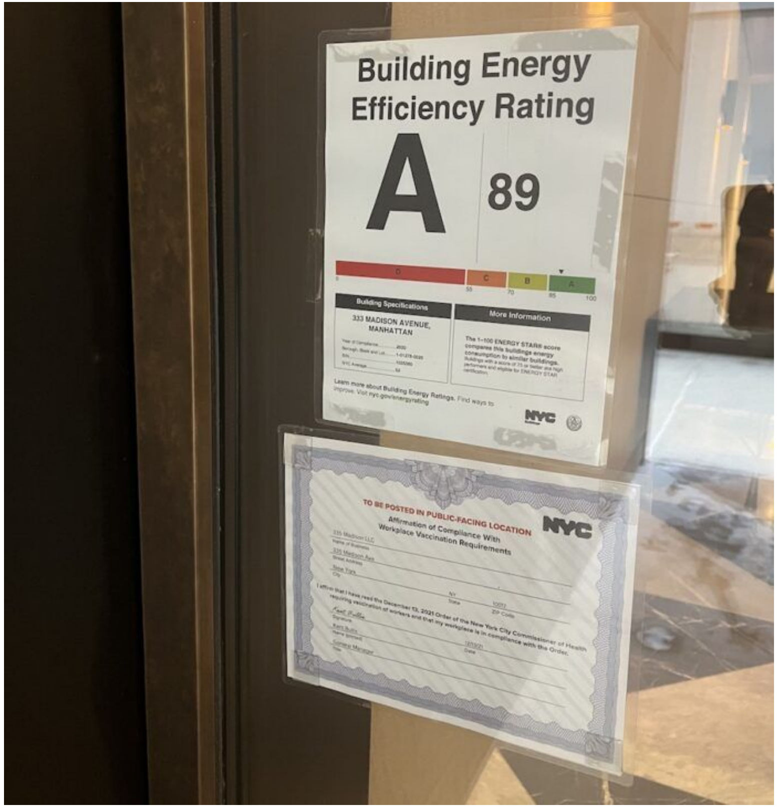

NYC Building Efficiency Grades

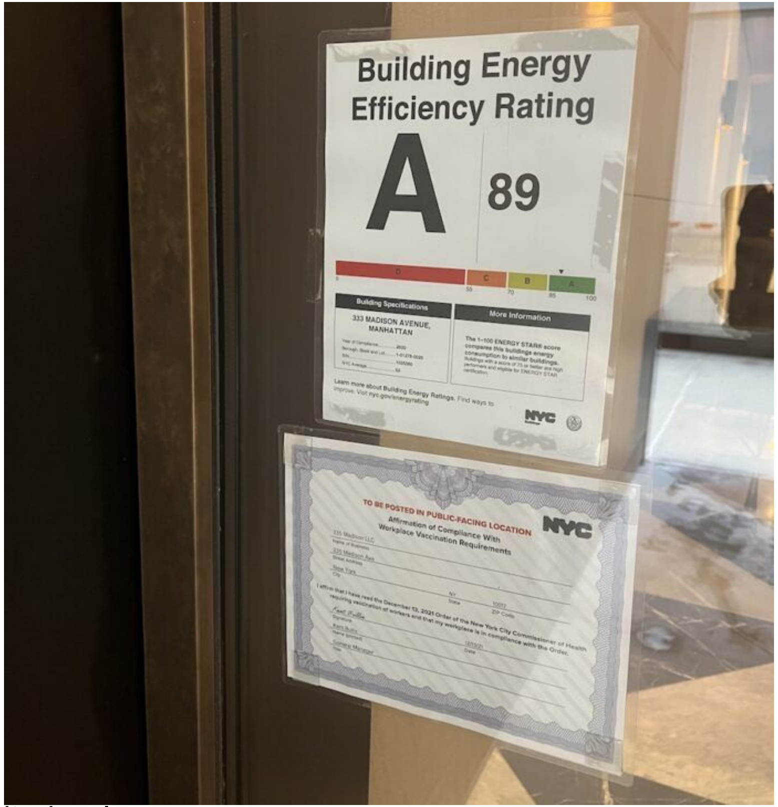

www.citysignal.com

Buildings in New York City are responsible for 70% of emissions

NYC Building Efficiency Grades

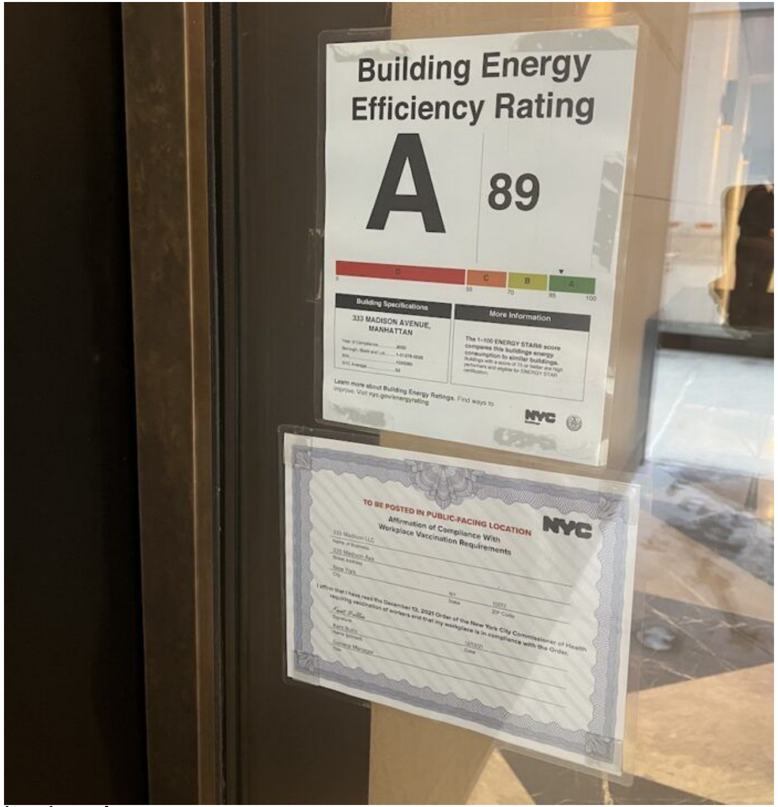

www.citysignal.com

Local Law No. 95 requires all buildings to have an efficiency rating based on Energy Star and EPA guidelines. In October 2020, an additional rule was added that larger buildings had to have these scores posted on their main entrances. That’s around 50,000 buildings and nearly two-thirds of the building area in the city.

The rankings come from Energy Star and are based on energy consumption, water consumption, and greenhouse gas emissions.

Buildings must submit energy information to the government for a 12 month period of time, and a list of fuels burned and converted on site. These criteria add together for a total score of anywhere between 1-100.

NYC Building Efficiency Grades

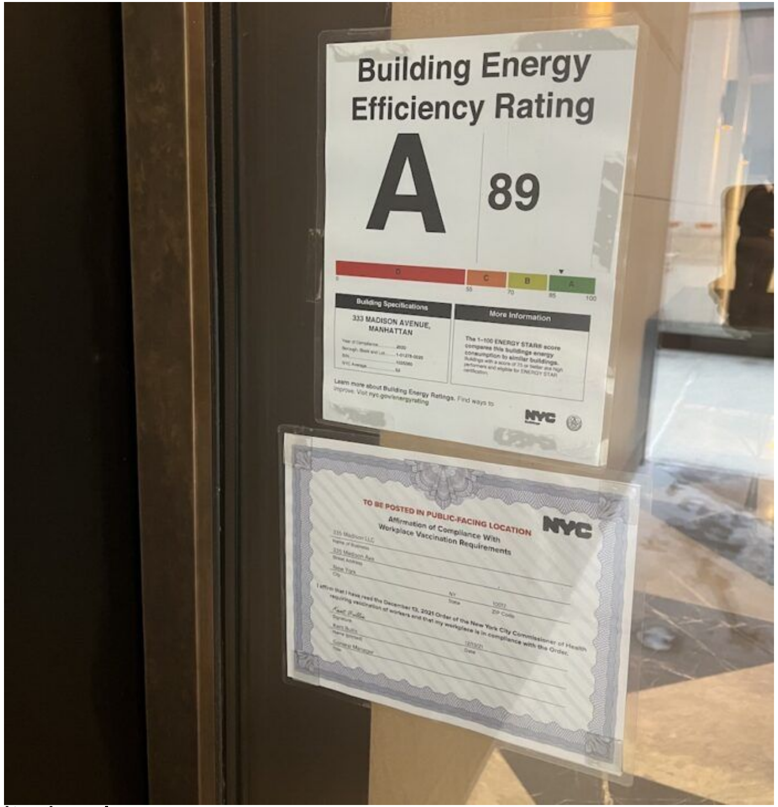

www.citysignal.com

Of the approximately 40,000 buildings that submitted reports when the law was first implemented, about half of them received a D. Thousands submitted nothing, receiving F’s. Buildings like The New York Stock Exchange and Trump Tower received some of the lowest scores, while the Flatiron Building and the Empire State Building did rather well.

In 2024 Local Law No. 97 goes into effect, which would fine buildings with lower scores. Depending on how low the scores are, the buildings could receive fines as high as hundreds of thousands of dollars. This has incentivized many buildings to change their energy consumption and distribution.

NYC Building Efficiency Grades

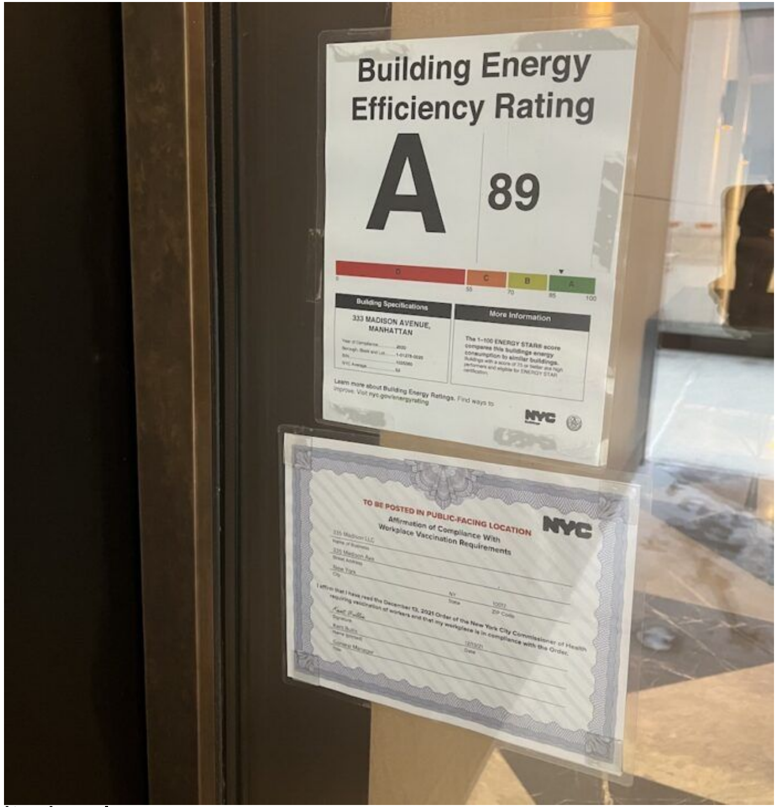

www.citysignal.com

Of the approximately 40,000 buildings that submitted reports when the law was first implemented, about half of them received a D. Thousands submitted nothing, receiving F’s. Buildings like The New York Stock Exchange and Trump Tower received some of the lowest scores, while the Flatiron Building and the Empire State Building did rather well.

In 2024 Local Law No. 97 goes into effect, which would fine buildings with lower scores. Depending on how low the scores are, the buildings could receive fines as high as hundreds of thousands of dollars. This has incentivized many buildings to change their energy consumption and distribution.

How to increase building efficiency

post-war gas low-rise

Energy Conservation Measure	Anytime/ Anywhere	Midcycle Retrofit	Refinancing/ Substantial Retrofit	Tenant Turnover	Payback (years)	Cost per SF	Energy Savings per SF
	Building Touchpoint						
Install Exhaust Fan Timers	•	•	•		5.0	$	
Install Submetering		•	•		2.0	$$$	
Install Solar/Photovoltaic			•		17.0	$$$$	
Upgrade Motors		•	•		5.5	$$	
Upgrade Lights		•	•		2.5	$	
Install Lighting Sensors		•	•		4.0	$	
Upgrade Burner			•		6.5	$$	
Upgrade Boiler			•		>20	$$$$	
Install TRVs and Zone Control		•	•		6.5	$$$	
Install Heating Controls and Thermostats		•	•		2.5	$$	
Insulate Condensate Tank	•	•	•		2.5	$	
Replace or Repair Steam Traps	•	•	•		3.5	$$	
Insulate Pipes	•	•	•		2.0	$	
Install or Upgrade Master Venting		•	•		3.0	$$	
Replace Windows and Glazing			•		>20	$$$$	
Increase Wall Insulation			•		>20	$$$$	
Increase Roof Insulation			•		>20	$$$	
Complete Air sealing	•	•	•	•	6.0	$$	
Separate DHW from Heating			•		6.5	$$$	
Install Low-Flow Showerheads	•	•	•	•	1.0	$$	
Install DHW Controls	•	•	•		0.5	$	
Install Low Flow Aerators	•	•	•	•	1.5	$$	
Insulate Pipes and Tank	•	•	•		6.0	$	

Energy Conservation Measure

- Ventilation & Cooling
- Other
- Lighting
- Heating Equipment
- Heating Distribution
- Envelope
- Domestic Hot Water

Cost per Square Foot

$	<$.05
$$	$0.05-$0.25
$$$	$0.26-$1.00
$$$$	>$1.00

Energy Savings per SF (kBtu)

- 0–3
- 3.1–8
- 8.1–12
- >12

Notes

This list of Energy Conservation Measures (ECM) is based on LL87 audit data and therefore may be incomplete. Suggested ECMs for each Building Touchpoint are representative, and not necessarily applicable to every building. Variety in specific building systems and condition of equipment must be considered in determining the appropriate packages of ECMs for individual buildings. The first step of any upgrade should be to work with a qualified service provider to develop a scope of work appropriate for your building.

Urban vs Suburban Energy Consumption

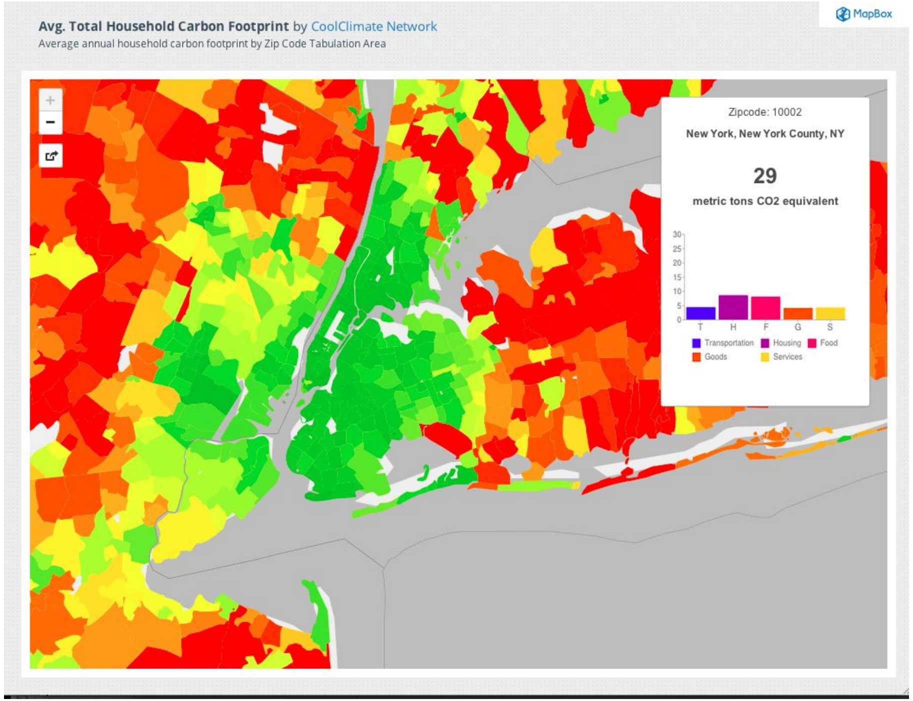

Low density suburban development is 2.0–2.5× as energy and GHG emissions intensive as high-density urban core development per capita.

Urban vs Suburban Energy Consumption

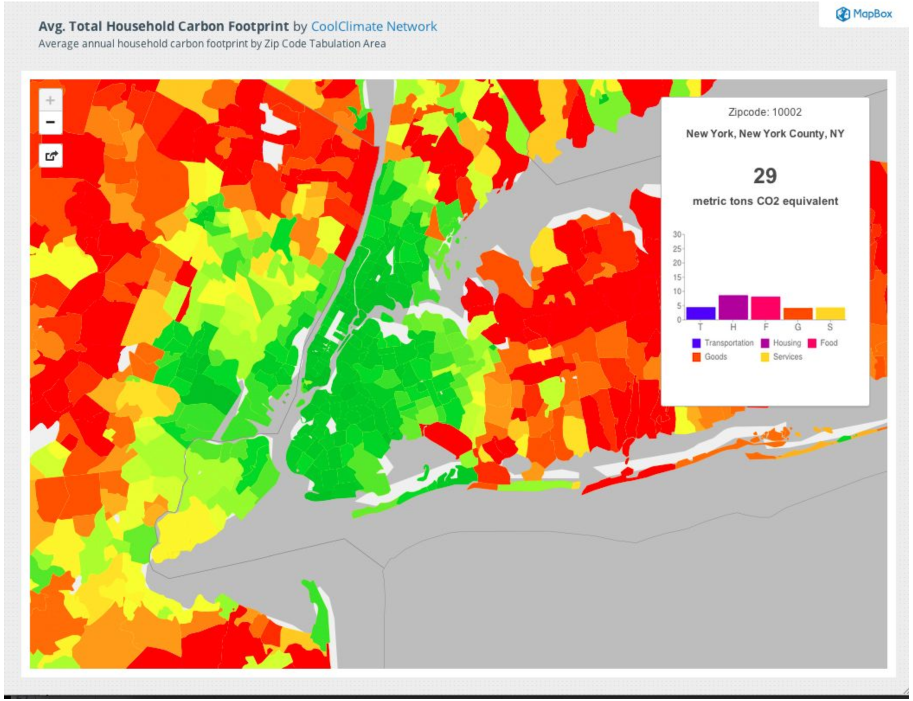

With shared resources, shared walls and generally smaller square footage, households in buildings with five or more units consume only 38 percent of the energy of households in single-family homes (Brown et al., 2005).

At a suburban density of four homes per acre, carbon dioxide emissions per household were found to be 25 percent higher than in an urban neighborhood with 20 homes per acre (Mazza, 2004).

Benefit of urban density: district heating

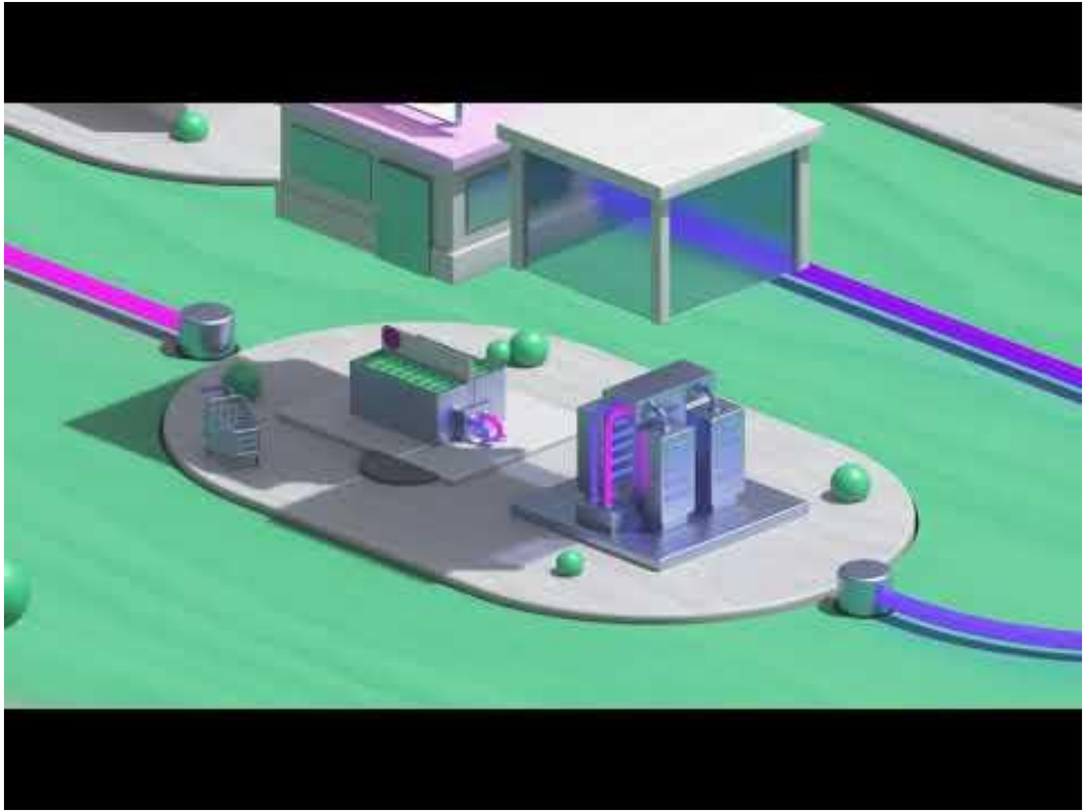

Opportunities for ML influence

District Heating

- Control—as a material flow process, a large portion of DH systems rely on improving control mechanisms at various points in the value chain. Optimal control comes down to the determination of the best control signals for the four levels of control in DH networks [11], as shown in Figure 1. Inadequate incentives and control strategies often result in wasteful energy supply margins.
- Analytics—insufficient use of analytic tools for evaluating the result from a combination of control strategies leads to lack of information about the true performance of the DH system. By improving the use of analytics, more comprehensive situation assessment can be used in planning and feedback to the control units, akin to model predictive control systems.

Open Access Article

Opportunities for Machine Learning in District Heating

by Gideon Mbiydzenyuy [1,*], Sławomir Nowaczyk [2], Håkan Knutsson [3], Dirk Vanhoudt [4,5], Jens Brage [6] and Ece Calikus [2]

https://www.mdpi.com/2076-3417/11/13/6112

Opportunities for ML influence

Automated thermostats/energy control

Journal of Building Engineering

Volume 50, 1 June 2022, 104165

Applications of reinforcement learning for building energy efficiency control: A review

Qiming Fu [a, b, 1], Zhicong Han [a, b, 1], Jianping Chen [b, c], You Lu [a, b], Hongjie Wu [a], Yunzhe Wang [a, b]

Show more

+ Add to Mendeley Share Cite

https://doi.org/10.1016/j.jobe.2022.104165

Get rights and content

https://www.sciencedirect.com/science/article/pii/S2352710222001784

4.1. Advantage

This paper reviews papers that apply RL methods to intelligent buildings. According to our research, applying RL to intelligent buildings can help reduce building energy consumption and achieve energy savings. The exploratory and exploitative nature of RL makes it unnecessary to build a complete system model, and it can demonstrate superior performance and efficiency over traditional control methods when dealing with intelligent buildings control problems with uncertain information.

4.2. Limitations

Applying RL to intelligent buildings inevitably has some problems. Paper [55] using DQN and DDPG to control refrigerant control parameters also points out that there are still many problems to be investigated when applying RL to real buildings. For example, the training time, the modeling of the environment, the difference between the simulated environment and the real environment, the cost of applying to the real environment, etc. These are all obstacles to the application of RL to building control.

4.2.1. Cost

RL itself is a process of learning by trial and error, so RL agents will perform actions during the training phase to obtain results that do not meet our expectations, or even the opposite of what we expect to achieve. If RL is deployed directly in a real building for training and learning, it will bring great inconvenience to occupants and even bring damage to the equipment, resulting in uncontrollable actual costs and easily bring security problems. For example, paper [52] applies RL to smart grids, where the security of the grid is largely and directly related to the cost issue. The paper proposes a security exploration approach to constrain the operation of active distribution grids. In the paper, a security layer is composed directly on top of the participant network of the DDPG, which predicts the change of the constrained state and thus limits the violation of the working operation of the active distribution grid.

Opportunities for ML influence

Identifying retrofit targets

ELSEVIER

Energy and Buildings

Volume 128, 15 September 2016, Pages 431-441

Applications of machine learning methods to identifying and predicting building retrofit opportunities

Daniel E. Marasco, Constantine E. Kontokosta

Show more

+ Add to Mendeley Share Cite

ELSEVIER

Applied Energy

Volume 281, 1 January 2021, 116024

A machine learning-based surrogate model to approximate optimal building retrofit solutions

Emmanouil Thrampoulidis [a, b], Georgios Mavromatidis [c], Aurelien Lucchi [d], Kristina Orehounig [b]

Show more

+ Add to Mendeley Share Cite

Highlights

- Cities want to accelerate building energy systems retrofits to reduce energy use.
- New York City energy audit data is applied to estimate energy retrofit eligibility.
- An interpretable classifier is trained to identify retrofit opportunities.
- Retrofit eligibility determined from only the most relevant building features.
- Building stakeholders can use results to rapidly identify retrofit opportunities.

Highlights

- Machine learning-based surrogate model to predict near-optimal retrofit solutions.
- Validated with a conventional building simulation-optimization model.
- Case study reveals good accuracy, ease of application and computational efficiency.
- Convenient for non-expert decision makers due to small set of input data.
- The model is scalable and applicable for retrofit analyses in wide-areas.

Opportunities for ML influence

Estimating and predicting building energy consumption

Physical Simulations:

“One way of estimating building energy consumption, in the absence of actual sensor data, is to create physical building models with a “template” of representative buildings, then run thermodynamic simulations to estimate the energy demands. These “engineering” models of building energy consumption are computationally expensive and cannot capture the wide variety of different buildings present in cities, as modeling each type of building requires very detailed input data, which is costly to collect. “

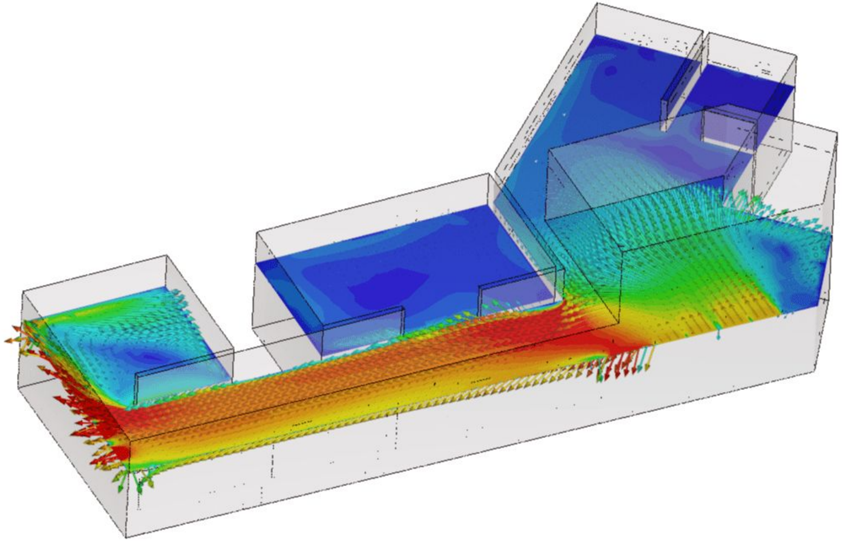

Paper Deep Dive

ELSEVIER

Applied Energy

Volume 208, 15 December 2017, Pages 889-904

Machine learning approaches for estimating commercial building energy consumption

Caleb Robinson [a], Bistra Dilkina [a], Jeffrey Hubbs [c], Wenwen Zhang [b], Subhrajit Guhathakurta [b], Marilyn A. Brown [c], Ram M. Pendyala [d]

Show more

\+ Add to Mendeley Share Cite

https://doi.org/10.1016/j.apenergy.2017.09.060

Get rights and content

https://www.sciencedirect.com/science/article/pii/S0306261917313429

Brainstorm

What kind of data would you want to have to be able to approach this problem?

What kind of methods would you apply?

How would you measure success?

If successful, how could this system be useful?

Overall goals of this work

Our two main objectives are to:

(1) train machine learning models to predict the annual major fuel, or the combination of electricity, natural gas, and fuel oil, consumption of commercial buildings from easily accessible descriptive features of buildings

(2) validate the models' ability to be applied to specific metropolitan areas. Specifically, we train and test our models using national survey data from CBECS, then use true energy consumption values from New York City's Local Law 84 (LL84) dataset to validate the ability of the nationwide CBECS-trained models to be applied accurately to a specific metropolitan area.

Commercial Buildings Energy Consumption Survey

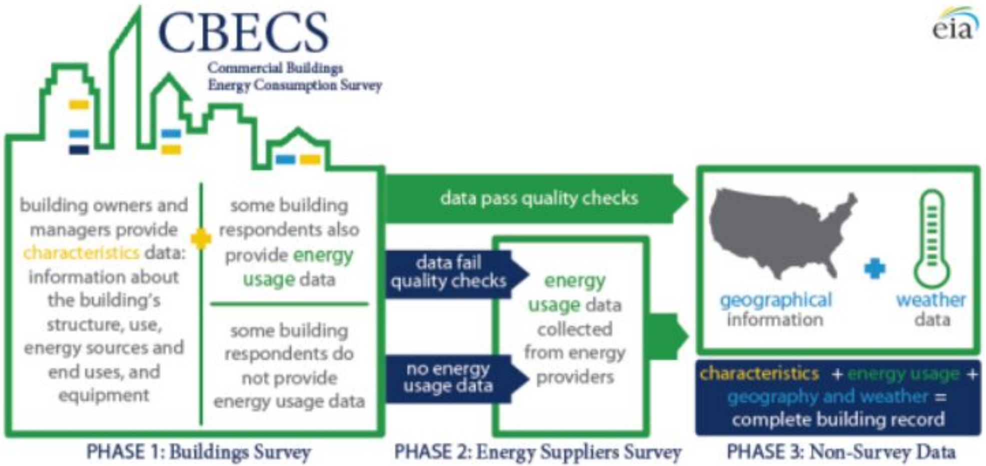

- The only independent, statistically representative source of national-level data on the characteristics and energy use of commercial buildings
- Building characteristics collected through an in-person or web survey of managers of 6,436 buildings, representing 5.9 million buildings in the United States.
- Energy usage data collected from suppliers of electricity, natural gas, fuel oil, and district heat

NYC Local Law 84 Dataset

The **New York City Benchmarking Law, known as Local Law 84** (LL84), requires buildings that are over 50,000 square feet, or lots with buildings with over 100,000 square feet combined, to ***report their annual energy and water consumption values*** in a standardized manner using the EPA's portfolio manager database.

This consumption data, along with some of the building characteristics (such as: total square feet, year built, primary building activity, and energy use intensity), have been released annually since 2011.

We augment each row in the LL84 dataset with heating degree day (HDD) and cooling degree day (CDD) features from 2015 CDD and HDD raster maps.

We further join the LL84 dataset to the New York City PLUTO dataset in order to get more information, such as the number of floors, for each building in the LL84 dataset.

After this processing, we have information on 2612 commercial buildings, which we will simply refer to as the LL84 dataset.

Data Features

The CBECS data has more features than LL84, such as, ‘Number of Employees’, ‘Number of X-ray machines’, or ‘Insulation upgraded’.

The **extended feature** set contains all features available in CBECS.

“The **common feature** set contains only the features from CBECS that are also available in the augmented LL84 data, namely: **principal building activity, square footage, number of floors, heating degree days, and cooling degree days.** ”

Building types ('Principal building activity')

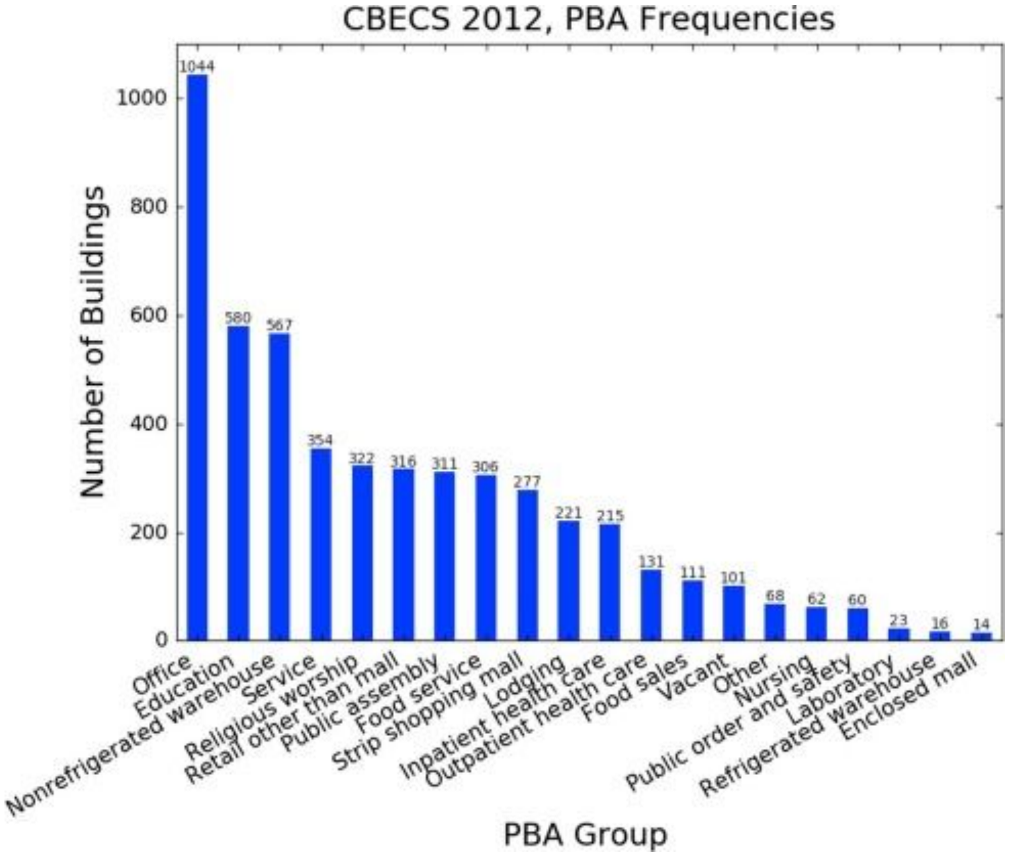

Data Features

<u>CDD and HDD:</u>

If the temperature mean is above 65°F, we subtract 65 from the mean and the result is Cooling Degree Days. If the temperature mean is below 65°F, we subtract the mean from 65 and the result is Heating Degree Days.

Heating and Cooling Degree Days

Degree days are based on the assumption that when the outside temperature is 65°F, we don't need heating or cooling to be comfortable. Degree days are the difference between the daily temperature mean, (high temperature plus low temperature divided by two) and 65°F. If the temperature mean is above 65°F, we subtract 65 from the mean and the result is ***Cooling Degree Days***. If the temperature mean is below 65°F, we subtract the mean from 65 and the result is ***Heating Degree Days***.

Example 1: The high temperature for a particular day was 90°F and the low temperature was 66°F. The temperature mean for that day was:

(90°F + 66°F) / 2 = 78°F

Because the result is *above* 65°F:

78°F - 65°F = 13 ***Cooling Degree Days***

Example 2: The high temperature for a particular day was 33°F and the low temperature was 25°F. The temperature mean for that day was:

(33°F + 25°F) / 2 = 29°F

Because the result is *below* 65°F:

65°F - 29°F = 36 ***Heating Degree Days***.

Target: ‘Annual Major Fuel Consumption’ (MFBTU)

Has a log-normal distribution, so the aim is to predict the log of the actual MFBTU

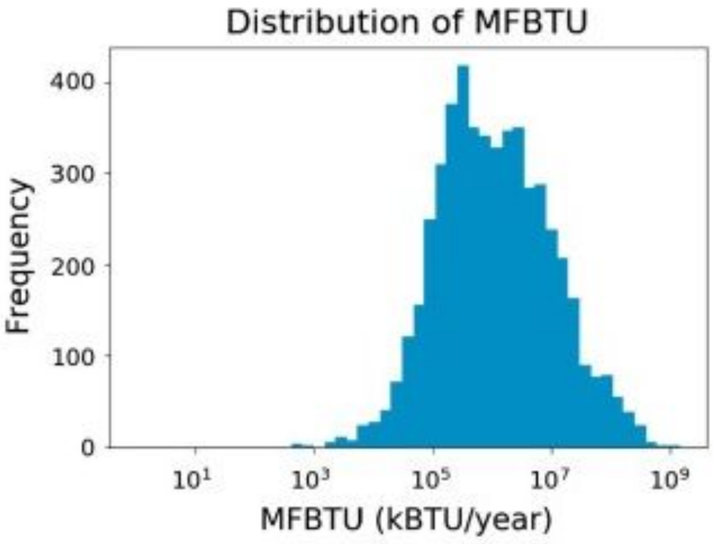

Methods

Linear regression, ridge regression, RBF kernel support vector regression (SVR), elastic net regression, linear kernel support vector regression (linear SVR), adaboost regression, bagging regression, gradient boosting regression (XGBoost), random forest regression (RF regressor), extra trees regression (ET regressor), multi-layer perceptron regression (MLP regressor), and k-nearest neighbor regression (KNN regressor)

Validation and Evaluation

To validate the models they record the cross-validated mean absolute error (mean AE), median absolute error (median AE), and the r^2 between the true and predicted values on the CBECS model.

They evaluate the generalization ability of models by using the augmented LL84 dataset.

Results

Table 2. Common features. Results of all machine learning models trained on the common feature set. The mean absolute error (mean AE), median absolute error (median AE), and the r^2 values are calculated in terms of log_{10} MFBTU values. The $10^{\text{Mean AE}}$ and $10^{\text{Median AE}}$ columns show the average number of multiples away the model's estimate is from the true value. The values following "±" are the standard deviations of each metric calculated over the 10 cross validation folds.

Results

Common features	Mean absolute error	$10^{\text{Mean AE}}$	Median absolute error	$10^{\text{Median AE}}$	r^2
XGBoost	**0.30 ± 0.01**	**1.99 ± 0.06**	**0.22 ± 0.01**	**1.66 ± 0.03**	**0.82 ± 0.02**
Bagging	0.33 ± 0.01	2.13 ± 0.07	0.24 ± 0.01	1.73 ± 0.05	0.78 ± 0.03
MLP Regressor	0.33 ± 0.01	2.16 ± 0.05	0.25 ± 0.01	1.77 ± 0.04	0.78 ± 0.02
Random Forest Regressor	0.33 ± 0.02	2.13 ± 0.07	0.24 ± 0.01	1.73 ± 0.05	0.78 ± 0.02
Extra Trees Regressor	0.34 ± 0.02	2.17 ± 0.08	0.24 ± 0.01	1.74 ± 0.05	0.76 ± 0.03
SVR	0.39 ± 0.01	2.44 ± 0.07	0.29 ± 0.01	1.95 ± 0.04	0.70 ± 0.03
KNN Regressor	0.43 ± 0.01	2.68 ± 0.08	0.32 ± 0.02	2.10 ± 0.07	0.65 ± 0.03
AdaBoost	0.43 ± 0.03	2.71 ± 0.16	0.36 ± 0.03	2.29 ± 0.17	0.68 ± 0.03
Linear SVR	0.51 ± 0.02	3.28 ± 0.15	0.40 ± 0.02	2.54 ± 0.11	0.52 ± 0.04
Linear Regression	0.52 ± 0.02	3.33 ± 0.13	0.43 ± 0.02	2.72 ± 0.12	0.53 ± 0.03
Ridge Regressor	0.52 ± 0.02	3.33 ± 0.13	0.43 ± 0.02	2.72 ± 0.12	0.53 ± 0.03
ElasticNet	0.76 ± 0.02	5.75 ± 0.32	0.67 ± 0.03	4.67 ± 0.35	0.09 ± 0.01
Lasso	0.79 ± 0.02	6.17 ± 0.35	0.69 ± 0.03	4.92 ± 0.38	0.00 ± 0.00

Table 2. Common features. Results of all machine learning models trained on the common feature set. The mean absolute error (mean AE), median absolute error (median AE), and the r^2 values are calculated in terms of log_{10} `MFBTU` values. The $10^{\text{Mean AE}}$ and $10^{\text{Median AE}}$ columns show the average number of multiples away the model's estimate is from the true value. The values following "±" are the standard deviations of each metric calculated over the 10 cross validation folds.

Results

Extended features	**Mean absolute error**	$10^{\text{Mean AE}}$	**Median absolute error**	$10^{\text{Median AE}}$	r^2
XGBoost with common features	0.30 ± 0.01	1.99 ± 0.06	0.22 ± 0.01	1.66 ± 0.03	0.82 ± 0.02
XGBoost	**0.23 ± 0.01**	**1.69 ± 0.02**	**0.17 ± 0.01**	**1.48 ± 0.03**	**0.89 ± 0.01**
Linear regression	0.24 ± 0.01	1.75 ± 0.02	0.19 ± 0.01	1.53 ± 0.04	0.88 ± 0.01
Ridge regressor	0.24 ± 0.01	1.75 ± 0.02	0.19 ± 0.01	1.53 ± 0.04	0.88 ± 0.01
SVR	0.25 ± 0.01	1.79 ± 0.04	0.19 ± 0.01	1.53 ± 0.03	0.87 ± 0.01
Bagging	0.25 ± 0.01	1.79 ± 0.04	0.18 ± 0.01	1.53 ± 0.04	0.87 ± 0.02
Random forest regressor	0.25 ± 0.01	1.79 ± 0.04	0.18 ± 0.01	1.53 ± 0.04	0.87 ± 0.01
Extra trees regressor	0.25 ± 0.01	1.79 ± 0.04	0.19 ± 0.01	1.54 ± 0.03	0.87 ± 0.01
Linear SVR	0.26 ± 0.01	1.80 ± 0.03	0.20 ± 0.01	1.58 ± 0.04	0.87 ± 0.01
AdaBoost	0.32 ± 0.01	2.07 ± 0.05	0.26 ± 0.01	1.80 ± 0.05	0.82 ± 0.01
KNN regressor	0.37 ± 0.01	2.34 ± 0.06	0.29 ± 0.01	1.93 ± 0.04	0.75 ± 0.02
MLP regressor	0.45 ± 0.02	2.82 ± 0.11	0.36 ± 0.02	2.31 ± 0.10	0.64 ± 0.03
ElasticNet	0.60 ± 0.02	4.00 ± 0.20	0.51 ± 0.02	3.26 ± 0.16	0.40 ± 0.01
Lasso	0.79 ± 0.02	6.17 ± 0.35	0.69 ± 0.03	4.92 ± 0.38	0.00 ± 0.00

Common features	**Mean absolute error**	$10^{\text{Mean AE}}$	**Median absolute error**	$10^{\text{Median AE}}$	r^2
XGBoost	**0.30 ± 0.01**	**1.99 ± 0.06**	**0.22 ± 0.01**	**1.66 ± 0.03**	**0.82 ± 0.02**
Bagging	0.33 ± 0.01	2.13 ± 0.07	0.24 ± 0.01	1.73 ± 0.05	0.78 ± 0.03
MLP Regressor	0.33 ± 0.01	2.16 ± 0.05	0.25 ± 0.01	1.77 ± 0.04	0.78 ± 0.02
Random Forest Regressor	0.33 ± 0.02	2.13 ± 0.07	0.24 ± 0.01	1.73 ± 0.05	0.78 ± 0.02
Extra Trees Regressor	0.34 ± 0.02	2.17 ± 0.08	0.24 ± 0.01	1.74 ± 0.05	0.76 ± 0.03
SVR	0.39 ± 0.01	2.44 ± 0.07	0.29 ± 0.01	1.95 ± 0.04	0.70 ± 0.03
KNN Regressor	0.43 ± 0.01	2.68 ± 0.08	0.32 ± 0.02	2.10 ± 0.07	0.65 ± 0.03
AdaBoost	0.43 ± 0.03	2.71 ± 0.16	0.36 ± 0.03	2.29 ± 0.17	0.68 ± 0.03
Linear SVR	0.51 ± 0.02	3.28 ± 0.15	0.40 ± 0.02	2.54 ± 0.11	0.52 ± 0.04
Linear Regression	0.52 ± 0.02	3.33 ± 0.13	0.43 ± 0.02	2.72 ± 0.12	0.53 ± 0.03
Ridge Regressor	0.52 ± 0.02	3.33 ± 0.13	0.43 ± 0.02	2.72 ± 0.12	0.53 ± 0.03
ElasticNet	0.76 ± 0.02	5.75 ± 0.32	0.67 ± 0.03	4.67 ± 0.35	0.09 ± 0.01
Lasso	0.79 ± 0.02	6.17 ± 0.35	0.69 ± 0.03	4.92 ± 0.38	0.00 ± 0.00

Table 2. Common features. Results of all machine learning models trained on the common feature set. The mean absolute error (mean AE), median absolute error (median AE), and the r^2 values are calculated in terms of log_{10} MFBTU values. The $10^{\text{Mean AE}}$ and $10^{\text{Median AE}}$ columns show the average number of multiples away the model's estimate is from the true value. The values following "±" are the standard deviations of each metric calculated over the 10 cross validation folds.

Feature Importance

“To aid the interpretability of our modeling process, we determine which features are the most important to the gradient boosting models

Feature importances in gradient boosting models are calculated as the amount of reduction in Gini impurity each feature causes over all splits for which that feature is present, over all of the trees that make up the model.

These values give us the relative importance of each feature included in a model, allowing us to rank the features in terms of “most useful” in the model, and compare the relative importance of features within a model.”

Extended feature set

Feature name	Feature description	Importance	Feature ranking in common set
SQFT	Square footage	0.1391	1
NWKER	Number of employees	0.0576	
WKHRS	Total hours open per week	0.0557	
ZMFBTU	Imputed major fuels consumption	0.0312	
MONUSE	Months in use	0.0299	
NGUSED	Natural gas used	0.0295	
HDD65	Heating degree days (base 65)	0.0293	3
HEATP	Percent heated	0.0278	
CDD65	Cooling degree days (base 65)	0.0224	2
NWKERC	Number of employees category	0.0221	

Feature Importance

When we train a gradient boosting regressor on *just* samples from the 'Service' class of buildings we observe that the most important feature is 'Total hours open per week', instead of 'Square footage'. This suggests that for some PBAs, the common feature set does not contain the correct signals to reproduce the MFBTU targets.

Extended feature set			
Feature name	**Feature description**	**Importance**	**Feature ranking in common set**
SQFT	Square footage	0.1391	1
NWKER	Number of employees	0.0576	
WKHRS	Total hours open per week	0.0557	
ZMFBTU	Imputed major fuels consumption	0.0312	
MONUSE	Months in use	0.0299	
NGUSED	Natural gas used	0.0295	
HDD65	Heating degree days (base 65)	0.0293	3
HEATP	Percent heated	0.0278	
CDD65	Cooling degree days (base 65)	0.0224	2
NWKERC	Number of employees category	0.0221	

How well does the model generalize to a different dataset?

Table 7. LL84 Validation. Comparison of the best external model tested on the LL84 dataset (out of sample validation result) to all machine learning models trained and tested on the LL84 dataset. The first row, 'XGBoost - CBECS', is the best external model and shows the results from applying the XGBoost model trained on all of the CBECS data, to all of the LL84 data. The remaining rows show the cross validated results on models trained and tested on the LL84 dataset. All results are shown with models using the common feature set. The mean absolute error (mean AE), median absolute error (median AE), and the r^2 values are calculated in terms of log_{10} MFBTU values. The $10^{\text{Mean AE}}$ and $10^{\text{Median AE}}$ columns show the average number of multiples away the model's estimate is from the true value.

How well does the model generalize to a different dataset?

Table 7. LL84 Validation. Comparison of the best external model tested on the LL84 dataset (out of sample validation result) to all machine learning models trained and tested on the LL84 dataset. The first row, 'XGBoost - CBECS', is the best external model and shows the results from applying the XGBoost model trained on all of the CBECS data, to all of the LL84 data. The remaining rows show the cross validated results on models trained and tested on the LL84 dataset. All results are shown with models using the common feature set. The mean absolute error (mean AE), median absolute error (median AE), and the r^2 values are calculated in terms of log_{10} MFBTU values. The $10^{\text{Mean AE}}$ and $10^{\text{Median AE}}$ columns show the average number of multiples away the model's estimate is from the true value.

	Mean absolute error	$10^{\text{Mean AE}}$	**Median absolute error**	$10^{\text{Median AE}}$	r^2
XGBoost - CBECS	0.25	1.78	0.15	1.41	0.51
XGBoost	**0.24 ± 0.02**	**1.75 ± 0.09**	**0.15 ± 0.01**	**1.40 ± 0.03**	**0.54 ± 0.09**
SVR	0.25 ± 0.02	1.77 ± 0.10	**0.15 ± 0.01**	**1.40 ± 0.03**	0.51 ± 0.11
Linear SVR	0.28 ± 0.02	1.92 ± 0.08	0.17 ± 0.00	1.50 ± 0.01	0.42 ± 0.05
MLP regressor	0.28 ± 0.04	1.92 ± 0.17	0.17 ± 0.02	1.48 ± 0.06	0.44 ± 0.13
Linear regression	0.29 ± 0.02	1.96 ± 0.10	0.19 ± 0.01	1.56 ± 0.05	0.44 ± 0.08
Ridge regressor	0.29 ± 0.02	1.96 ± 0.10	0.19 ± 0.01	1.56 ± 0.05	0.44 ± 0.08
Bagging	0.29 ± 0.02	1.95 ± 0.09	0.18 ± 0.01	1.50 ± 0.04	0.43 ± 0.08
Random forest regressor	0.29 ± 0.02	1.95 ± 0.10	0.18 ± 0.02	1.51 ± 0.05	0.43 ± 0.08
Extra trees regressor	0.30 ± 0.03	2.00 ± 0.12	0.18 ± 0.01	1.51 ± 0.05	0.39 ± 0.09
KNN regressor	0.30 ± 0.03	2.01 ± 0.15	0.19 ± 0.02	1.53 ± 0.06	0.40 ± 0.12
AdaBoost	0.42 ± 0.07	2.67 ± 0.43	0.30 ± 0.04	2.01 ± 0.20	0.14 ± 0.22
Lasso	0.45 ± 0.01	2.80 ± 0.04	0.33 ± 0.01	2.13 ± 0.06	Negative
ElasticNet	0.45 ± 0.01	2.80 ± 0.04	0.33 ± 0.01	2.13 ± 0.06	Negative

The gradient boosting model is again the best performing model, however, surprisingly, when the gradient boosting model is trained on the LL84 data, it only performs slightly better than the model that was trained on the CBECS data.

Applying the model to new locations

We applied the CBECS-trained gradient boosting regression model to the 73,388 commercial buildings in Atlanta from the CoStar real estate database.[3] We supplement the CoStar data with the 2017 heating and cooling degree day data from the Oak Ridge Climate models.

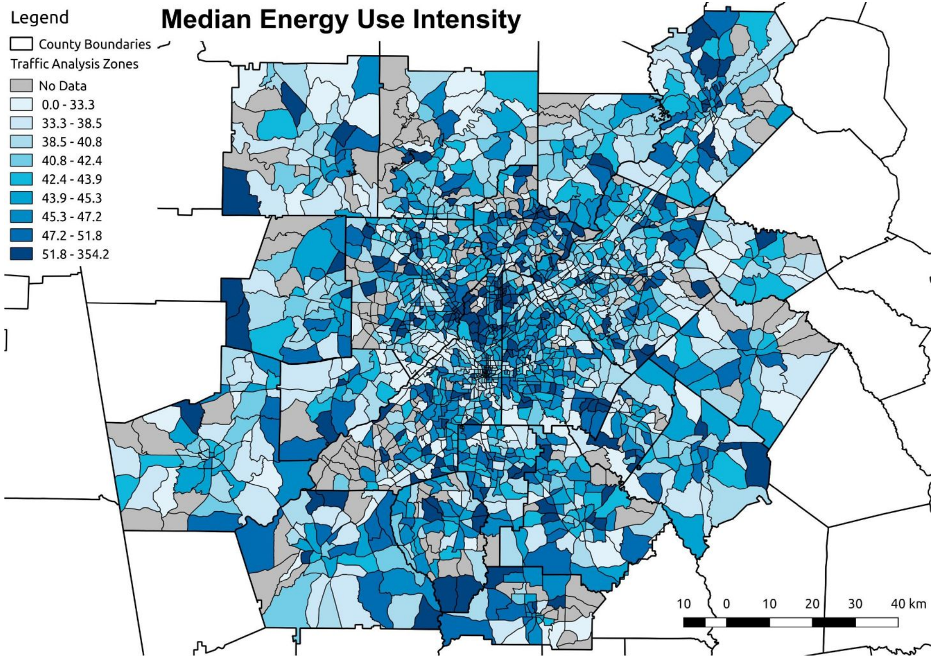

Uses and Implications of this work

Urban planners will be able to benchmark the effects of environmental or climate related policies affecting different sections of the urban region or make predictions about the outcomes of proposed policies.

Can also help city and regional planners predict the energy burdens that could result from alternative urban growth patterns and global warming scenarios.

Analysis of important features used by the machine learning models will serve to drive future data collection efforts that could help maximize the accuracy of the models.

Limitations

The error is too large for analyzing the energy consumption of any specific building. But when the models are used to make predictions for all the buildings in entire metropolitan areas (where individual prediction errors will cancel out when aggregated), as we show for Atlanta, they can offer useful insights into a city's commercial energy consumption landscape.

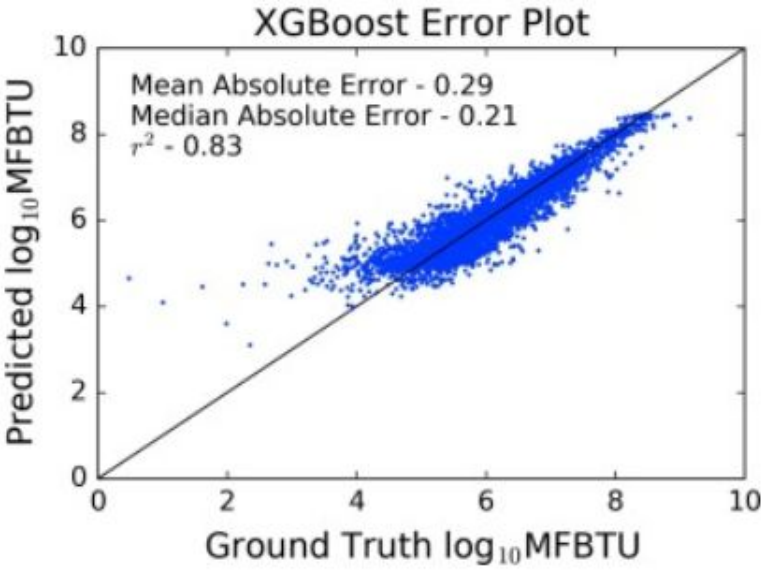

Recap

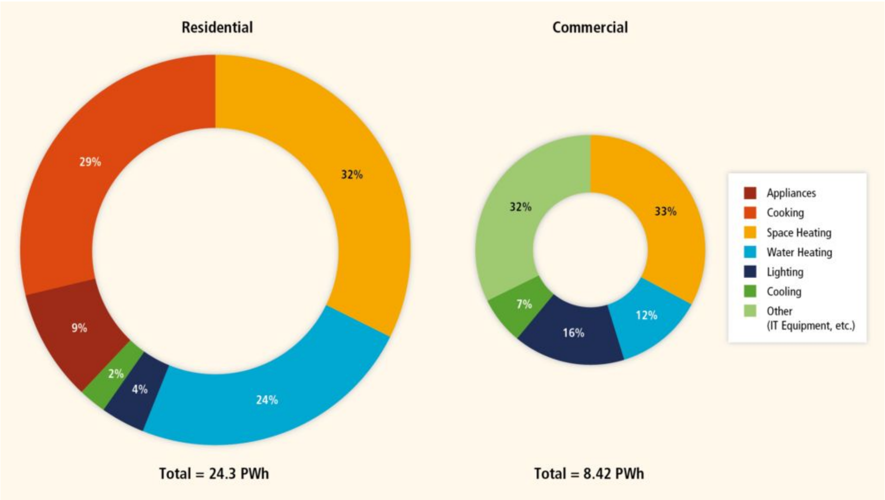

Residential
Commercial
29%
32%
9%
2%
4%
24%
33%
12%
16%
7%
Appliances
Cooking
Space Heating
Water Heating
Lighting
Cooling
Other
(IT Equipment, etc.)
Total = 24.3 PWh
Total = 8.42 PWh

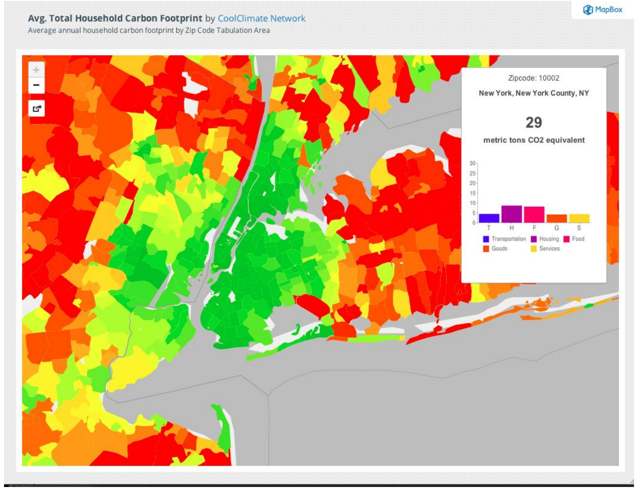

Avg. Total Household Carbon Footprint by CoolClimate Network
Average annual household carbon footprint by Zip Code Tabulation Area
MapBox
Zipcode: 10002
New York, New York County, NY
29
metric tons CO2 equivalent
Transportation
Housing
Food
Goods
Services

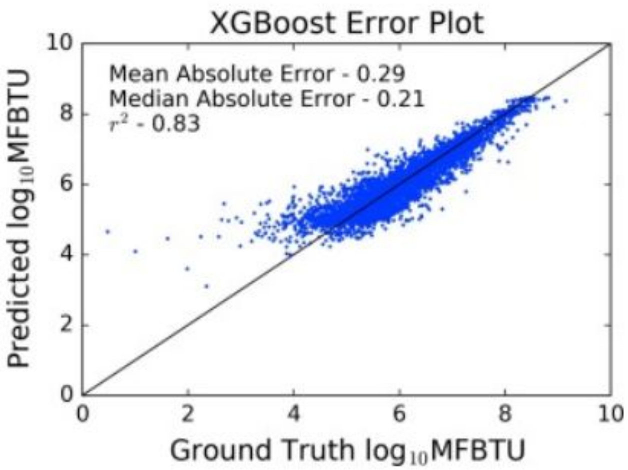

XGBoost Error Plot
Mean Absolute Error - 0.29
Median Absolute Error - 0.21
r^2 - 0.83
Predicted $\log_{10}$MFBTU
Ground Truth $\log_{10}$MFBTU