Quantifying Regional Economics and Infrastructure Variations using Geospatial and ML Statistical Analysis

Ardelia Clarke[1], Peter E. DeWitt[1,2], William Bryan[3], Robin Tuttle[1], Jordan Washington[1], Amy Lovell[3], Laura Diaz-Villaquiran[3], John R. Seydel[4], Chandra Farley[4], Nylah Oliver[5], Joshua Corning[5], Avni Naik[5], Danny Henderson[6], and Brian Brainerd[7]

[1]National Laboratory of the Rockies

[2]University of Colorado Anschutz

[3]Southeast Energy Efficiency Alliance

[4]City of Atlanta Office of Sustainability

[5]City of Savannah Office of Sustainability

[6]City of Savannah Code Compliance

[7]City of Savannah Office of Community Services

Abstract

Quantifying commercial building energy consumption at the parcel level can help city planners, utilities, and program designers identify tailored energy efficiency intervention. This study focuses on Atlanta and Savannah, Georgia as two contrasting cities driven by entertainment and port-related economies, respectively. Using publicly available parcel data, census datasets, and geospatial economic indicators, we trained an XGBoost machine learning (ML) model on Atlanta's 8-year energy benchmarking panel (n=1,015 labelled commercial, multifamily, residential, and mixed-use buildings) and transferred it to Savannah's full parcel inventory to predict city-wide energy use intensity (EUI). An 80/20 stratified holdout split and spatial k-fold cross-validation assessed the Atlanta-based XGBoost machine learning (ML) model. Spatial cross-validation produced an R^2 of 0.960 and a holdout R^2 of 0.823, confirming model generalization within Atlanta without overfitting. Transferability of the Atlanta-based XGBoost ML model to Savannah was low ($R^2 = 0.307$), setting a credible baseline for model replicability amongst cities with limited temporal EUI data. SHAP (SHapley Additive exPlanations) identified temporal EUI history (59% of predictive power), building valuation (15%), and neighboring parcels EUI (11%) as the top drivers for predicting commercial energy consumption. Statistical zone comparison analysis revealed historic districts used 66% more energy (62.4 kBtu/ft² observed) compared to baseline parcels (37.5 kBtu/ft² observed), with the model under-predicting by an average of 16.0 kBtu/ft²/yr which validated local expertise from project partners. Federal Opportunity Zones parcels had the highest mean EUI (52.7 kBtu/ft² observed) while economic development zones (EDZs) parcels consumed 13% less energy than baseline parcels. This parcel-level, replicable geospatial, statistical, and machine learning approach can enable identification of high-energy use commercial properties for energy efficiency and cost-saving measures.

1. Introduction

The relationship between energy infrastructure and commercial building performance is central to achieving reliability, affordability, and energy goals in American cities. Commercial buildings account for roughly 18% of U.S. energy consumption and are hubs for local employment, neighborhood engagement, and wealth generation (EIA 2022; Theodos and Gonzalez 2019). Yet the tools available to program designers and city planners for identifying which commercial buildings and neighborhoods need prioritized attention remain coarse, typically relying on city-level or census-tract-level aggregations that obscure the parcel-level variation that program implementation requires.

Pressing challenges compound this gap. Increased load growth from data centers, entertainment infrastructure, and port-related commerce is reshaping commercial energy demand in cities like Atlanta and Savannah. Research shows that economic growth disproportionately increases energy demand and strains grid infrastructure, with electricity needs often outpacing local capacity during periods of rapid urbanization (Anser et al. 2025). Globally, electricity demand is projected to require 25 million kilometers of new grid lines by 2035 to keep pace with economic expansion (IEA 2025), while empirical evidence demonstrates that funding physical grid assets directly drives regional economic output over the long term (Xu et al. 2021). At the same time, aging grid infrastructure and the concentration of impaired building performance across varying neighborhood contexts create reliability risks that are not evenly distributed. Economic development zone designations, including state Enterprise Zones and federal Opportunity Zones, are meant to attract investment and revitalize these areas, but whether they correspond to distinct commercial energy consumption patterns has not been systematically examined at the parcel level. GIS-based frameworks have demonstrated that directing resources to specific geographic areas yields the greatest impact on energy consumption and energy goals (Gaspari et al. 2020), yet such frameworks have not been applied to zone-level commercial energy benchmarking in the U.S. context.

This study addresses that gap through four primary research objectives:

1. Quantify parcel-level commercial energy consumption in Atlanta using an 8-year benchmarking panel and a machine learning model capable of transfer to data-limited cities.
2. Identify which economic, building-level, and geospatial indicators most strongly predict commercial energy use intensity through SHAP feature importance analysis.
3. Evaluate whether economic development zone designations (EDZ, Opportunity Zone, historic district) correspond to significantly different energy consumption patterns.
4. Demonstrate the framework's replicability by transferring the Atlanta-trained model to Savannah's full parcel inventory to generate city-wide EUI predictions.

The expected contributions are: (1) a ranked set of data-driven indicators linking commercial building characteristics to energy use; (2) a replicable spatial modeling framework in which the Atlanta-trained model is validated against in-sample performance metrics and spatial cross-validation before transfer to Savannah; and (3) zone-specific findings that directly inform where commercial energy improvement programs should be prioritized.

2. Background

Energy Infrastructure and Commercial Building Performance

Commercial buildings represent a critical building load and target for energy efficiency programs. Research consistently identifies small commercial buildings, those below 50,000 square feet, as receiving less program attention relative to their share of the commercial building stock and energy consumption (Langner et al. 2013; York et al. 2022). Of U.S. small commercial buildings built before 2000, 64% have not been renovated since 2000 (EIA 2022). These unrenovated buildings cluster in economically constrained neighborhoods where capital access, split-incentive challenges between owners and tenants, and limited program awareness compound the barriers to improvement (Clarke et al. 2023).

Energy use intensity (EUI), measured in kBtu per square foot per year (kBtu/ft²/yr), is the standard metric for comparing building energy performance across building types and sizes. Investment in building energy efficiency generates measurable economic returns beyond energy savings, including direct job creation and gross domestic product contribution through redirected energy expenditures (Anderson et al. 2014). Consumption-based metrics, when combined with a composite metric such as the Area Deprivation Index and Community Deprivation Index, provide a more direct, reliable, and holistic view of household economic well-being than traditional income data alone (Attanasio and Pistaferri 2016). Normalizing consumption by floor area makes large and small buildings directly comparable, a critical capability when analyzing a mixed commercial portfolio spanning offices, retail, warehouses, residential, and multifamily properties at the parcel level.

Economic Development Zones and Energy

State-designated Economic Development Zones and federally designated Opportunity Zones represent two distinct policy mechanisms for directing investment to economically impaired areas. Economic development is a multidimensional concept that conventional metrics such as gross domestic product fail to capture adequately; robust assessment requires measures spanning consumption, wealth stocks, distribution, and security (Osberg and Sharpe 2002). Local economic development theory identifies job creation, stable tax base expansion, and infrastructure quality as the primary measurable criteria for evaluating whether zone-level policy interventions generate sustained regional well-being (Malizia et al. 2020). EDZ designations typically confer local tax incentives and development support, attracting commercial users with potentially different building stock and operational profiles than the surrounding baseline. Federal Opportunity Zones, established by the 2017 Tax Cuts and Jobs Act, target census tracts with low-income populations for capital gains investment incentives. The spatial overlap between these two zone types, specifically parcels co-designated under both programs, has not previously been examined for its energy implications.

Historic district designations represent a third policy layer. Buildings subject to preservation requirements face constraints on exterior modifications, which can limit the types of energy efficiency retrofits available. Whether historic designation corresponds to higher or lower baseline energy consumption is an empirical question with direct implications for program design in cities like Savannah, where the National Historic Landmark District encompasses much of the downtown commercial core.

Study Cities

Atlanta, Georgia is a major regional commercial hub with dense urban development, significant entertainment and corporate office sectors, and an active benchmarking program administered through the Atlanta Better Buildings Challenge. Its 8-year Portfolio Manager energy panel provides an unusually rich longitudinal dataset for parcel-level model training.

Savannah, Georgia, a coastal city on the Atlantic seaboard, presents a contrasting urban context. Its economy is shaped by the Port of Savannah, one of the largest container ports in the United States, with associated warehousing, logistics, and support commercial uses. Port activity supported nearly 651,000 jobs statewide and contributed more than $6 billion in state and local taxes in fiscal year 2024, with a total economic impact of $174 billion in sales (Humphreys 2020). This infrastructure-intensive economy creates a distinct commercial building profile concentrated around logistics and industrial uses, making Savannah an informative contrast to Atlanta's entertainment and corporate office stock. Savannah's historic district is among the most intact urban grids in the South. The city lacks a mandatory benchmarking program, making it an ideal test case for demonstrating how a model trained on data-rich Atlanta can generate actionable parcel-level insights in a data-limited setting.

3. Methodology

Data Collection and Feature Engineering

The Atlanta training dataset integrates eleven publicly available data sources assembled at the parcel level (Table 1). The foundational energy data are drawn from Atlanta's Portfolio Manager benchmarking panel spanning 2015 through 2022, providing site EUI, source EUI, gross floor area, and year-built attributes for 1,088 commercially reported buildings. Spatial joins using parcel centroid-in-polygon methods link each energy building to its county parcel record, enabling integration of property valuation, land use classification, and administrative attributes.

Table 1. Data sources integrated at the parcel level

Data Source	Variables	Geographic Resolution
Atlanta Portfolio Manager (8-yr panel)	Site EUI, source EUI, gross floor area, year built, fuel type	Parcel (individual building)
Fulton County Tax Parcels 2025	Land use codes (57-entry lookup), ClassCode, ExCode, FMV, assessed value, living units	Parcel
U.S. Census ACS 5-Year (2022)	24 tables: income, Gini index, poverty, tenure, housing costs, employment, health insurance, internet access, group quarters	Census tract
newly created fairadi-data GitHub repository (ADI and CDI)	Area Deprivation Index (raw, national rank, state rank, fairadi), Community Deprivation Index (raw, standardized, faircdi)	Block group, aggregated to tract
DOE LEAD Energy Costs	Energy costs (% income spent on energy)	Census tract
Tree Canopy Raster (NLCD)	Percent canopy cover (mean, median, standard deviation per parcel)	Parcel (zonal statistics)

Data Source	Variables	Geographic Resolution
Atlanta EDZ (UEZ Track 2)	Economic development zone boundary	Zone polygon, joined to parcel centroid
Federal Opportunity Zones (OZ)	OZ tract designation, tract FIPS, county	Census tract polygon, joined to parcel
DOE SLOPE Business as usual Projections	Modeled energy consumption by sector (commercial, residential, industrial) and fuel type (2017-2050)	City-level benchmark

Property classification follows the Georgia Department of Revenue state classification system, derived from Fulton County Tax Parcels 2025. A 57-entry land use code lookup table converts raw LUCode values to plain-English property types (residential, commercial, exempt, industrial, historic, utility). Owner-occupancy is identified definitively through homestead exemption codes (HF* and S1 classifications), replacing the proxy measures common in prior work.

Census tract features are joined via spatial centroid-in-polygon matching using 2022 American Community Survey (ACS) 5-year estimates. Area Deprivation Index (ADI) and Community Deprivation Index (CDI) from the fairadi-data repository are aggregated from block group to tract level by averaging within-tract block group values (DeWitt and Clarke 2026). Energy costs data from the DOE LEAD tool are matched to parcels via census tract GEOID after converting the GISJOIN format to standard 11-digit tract FIPS codes.

Zone membership is derived from three spatial overlays: Atlanta EDZ boundaries pulled from ArcGIS Online (UEZ Track 2), Federal Opportunity Zone boundaries from a local GeoJSON file, and Atlanta historic district polygons. Each is joined to parcel centroids using point-in-polygon operations with post-join deduplication to prevent row inflation from overlapping polygon boundaries. A composite four-level factor enables stratified zone analysis, where neither indicates no zone designation, EDZ-only indicates economic development zone designation only, OZ-only indicates Federal Opportunity Zone designation only, and EDZ-and-OZ indicates parcels carrying both designations simultaneously.

Spatial K-Fold Cross-Validation

A common constraint in parcel-level spatial modeling is using random k-fold cross-validation, which allows high-EUI parcels in one neighborhood to appear in both training and test sets. Because nearby parcels share spatial autocorrelation in energy consumption, random folds produce inflated accuracy estimates. We instead use spatial k-fold cross-validation via the R blockCV package, which partitions parcels into five geographic blocks of approximately 1 km each. Each fold holds out one geographic block entirely for evaluation, ensuring the model's performance reflects genuine generalization to new spatial contexts.

XGBoost Model and SHAP Analysis

The modeling target is log-transformed site EUI (log1p transformation), applied because EUI distributions are right-skewed with a range of 3.7 to 217.1 kBtu/ft^2/yr in the Atlanta training set. The preprocessing script includes median imputation for numeric predictors, mode imputation for categorical predictors, one-hot encoding for nominal variables (property type, economic zone type, building decade), zero-variance removal, and correlation filtering at the 0.95 threshold to reduce feature redundancy and improve SHAP interpretability.

Hyperparameter tuning uses a 30-point Latin Hypercube Sampling grid across seven parameters (trees, tree depth, learning rate, minimum node size, loss reduction, sample proportion, and column fraction). Considering leveraging traditional grid search sampling techniques for multi-dimensional hyperspace hyperparameter spaces can lead to dimensionality issues and be computationally inefficient (187 model evaluations needed), we optimized our Atlanta-based XGBoost framework using a 30-point Latin Hypercube grid while maximizing stratification over marginal distributions (Bergstra and Bengio 2012). Recent methodologies demonstrate how partitioning each hyperparameter range into 30-point LHS provides uniform space-filling of the seven dimension matrix (Husslage, et al. 2011) (Lugner, et al. 2024). All tuning is evaluated using the five spatial folds. The best configuration (trees=621, depth=6, learning rate=0.0150, min_n=8) is then fit on the full 823-parcel Atlanta training dataset.

SHAP values are computed using exact TreeSHAP via the R shapviz package (Mayer and Stando 2025). Mean absolute SHAP provides a magnitude-based importance ranking that accounts for both direction and nonlinearity, making it more informative than XGBoost's native gain-based importance for the policy interpretation required here (Chen, et al. 2016).

Zone Comparison Analysis

Zone comparison uses Kruskal-Wallis testing (non-parametric ANOVA) rather than standard ANOVA because EUI distributions are right-skewed. Post-hoc Dunn testing with Bonferroni correction identifies pairwise differences. Cohen's d effect sizes provide the magnitude of zone differences independent of statistical significance. All tests are run on both observed EUI and model residuals; the residual test isolates the zone effect after controlling for building characteristics. Exceeding the 20-oberservation threshold for a reliable Kruskal-Wallis inference, the zone group sizes of training set are 141-parcels for EDZ-only, 26-parcels for OZ-only, and 68-parcels for both EDZ and OZ areas.

4. Results

Data Collection and Validation

The Atlanta parcel feature matrix contains 1,015 commercial, multifamily, and mixed-use buildings drawn from a larger 149,118-parcel spatial layer. All parcels with at least one year of observed EUI data was included to retain all the structural and neighborhood characteristics of the benchmarked data from the 1,015 Atlanta-parcels. Case weights were assigned accordingly to data reliability as weight 1.0 for parcels with five or more years of observations (n=135), weight 0.75 for three to four years (n=172), and weight 0.5 for one to two years (n=708). A 80/20 holdout split, stratified by economic zone type, reserved 201 parcels for independent validation.

The final feature set includes 63 predictor variables spanning building physical characteristics, property valuation metrics, 17 derived American Community Survey (ACS) tract features, ADI and CDI deprivation scores, zone membership indicators, tree canopy statistics, and SLOPE city-level energy benchmarks, and 2019 tract-level energy costs. Two columns (lu_description and lu_type) had 89% missing values due to LUCode lookup coverage gaps and were retained but imputed.

The Atlanta training EUI distribution had a mean of 37.1 kBtu/ft²/yr, standard deviation of 40.6 and range of 0.1 to 674.9. The right skew (several buildings above 150 kBtu/ft²/yr) confirmed the appropriateness of the log transformation. Zone distribution across the 1,015 buildings was: 579 baseline (neither), 141 EDZ-only, 26 OZ-only, 68 EDZ-and-OZ, and 9 in historic districts.

Finding significance: This data set collection was vital to create a representative sample of Atlanta's commercial building stock while accounting for data reliability variations through weighted observations. Including 63 predictors spanning structural, economic, and neighborhood factors provided a robust feature matrix capable of capturing variations in energy use. By stratifying the 80/20 holdout split by various zone types, improved the statistical predictive accuracy and cross-validation of the XGBoost ML model while linking building energy performance to economic development measures.

Spatial Sensitivity Analysis and Feature Importance

The XGBoost model achieved a spatial cross-validation R^2 of 0.960 and RMSE of 0.149 on the log EUI scale, with in-sample performance on 823-training parcels of R^2 0.823 and RMSE 18.67 kBtu/ft²/yr on the original scale (Table 2). Validating the XGBoost model on 201 unseen parcels produced a R^2 of 0.950 and RMSE of 10.23 kBtu/ft²/yr. A transfer simulation R^2 of 0.307 shows the accuracy of predicting EUI in data-limited areas with solely neighborhood and structural attributes. The difference in R^2 values from full training dataset and transfer simulation could elude to the predictive accuracy of including longitudinal benchmarked energy data in the XGBoost model.

Table 2. Model performance metrics (Atlanta, n=823)

Metric	Spatial CV (log EUI)	In-sample (original scale)	Notes
R^2	0.960	0.823	CV 0.960, Holdout 0.950, Transfer 0.307
RMSE	0.149 (log scale)	18.67 kBtu/ft²/yr	Approx. 50% average error on mean EUI of 37.1
MAE	0.078 (log scale)	2.65 kBtu/ft²/yr	Most predictions within 3-4 kBtu/ft²/yr of observed
CV method	Spatial k-fold (k=5)	Block size: 1 km	Fold sizes: 25/25/23/31/31

Figure 1. SHAP beeswarm: predictive feature categories for Atlanta XGBoost model

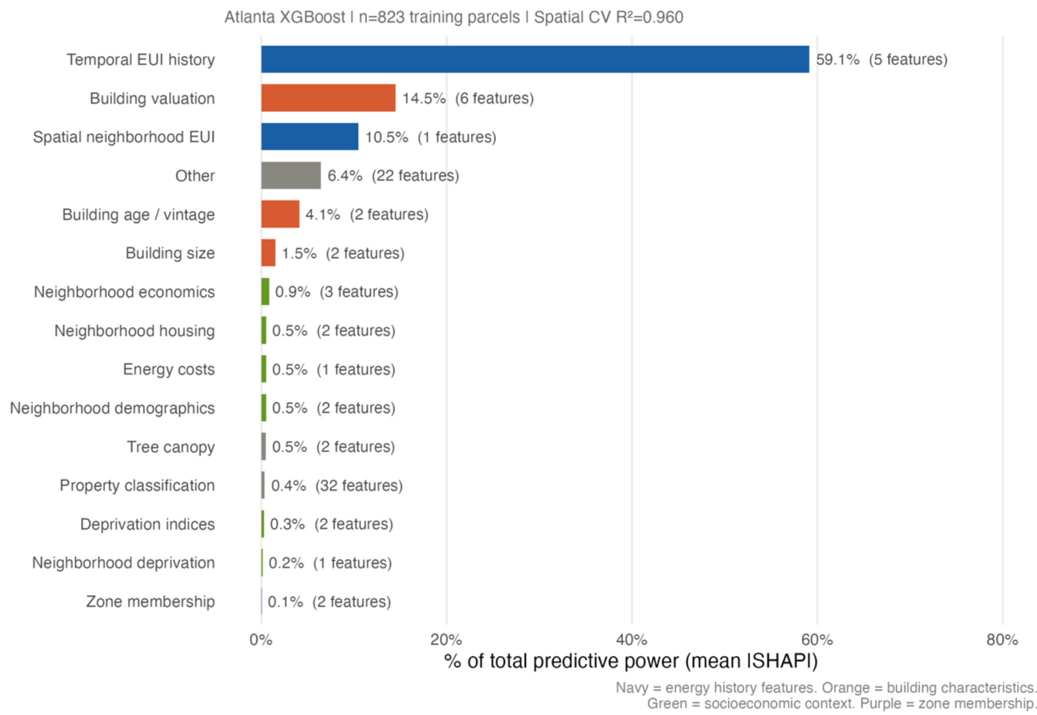

Figure 1. Predictive importance of feature categories to commercial EUI (Atlanta, n=823 training parcels). Bar length is percentage of total mean absolute SHAP value per category. Navy represents energy history features, orange = building characteristics, green= socioeconomic, purple = zone differences). Model: XGBoost with spatial k-fold CV (k=5), R^2=0.960.

Results of the top 12 predictors by mean absolute SHAP value from the Atlanta-based XGBoost mode is shown in Table 3. The first tier comprises demonstrates predicting future EUI is well predicted by having temporal EUI history features, specifically median, standard deviation, EUI spatial lag (parcel's k-nearest neighbor average), coefficient of variation, and trend across the 8-year panel, which accounted for approximately 69.6% of total predictive power (Figure 2).

The second tier covers building valuation and physical features: fair market value (FMV) per square foot (5.15%), building FMV (4.26%), and building age (3.73%), which collectively contributed 15% of predictive power. This could suggest that building stock investment could bridge the relationship between building age and energy performance.

The third tier, neighborhood economics, income and energy costs instability, contributed through tract-level Gini index, (2.37%) neighborhood and energy costs, tree canopy. Analysis shows that neighborhood-level income inequality predicts commercial energy consumption more than the composition deprivation indexes. This finding connects commercial energy performance to place-based inequality as wealthier neighborhoods may have more variable energy consumption and mixed commercial corridors also may have high-end businesses.

Table 3. Top 12 predictors by mean absolute SHAP value

Category	Predictor Name	Description	Mean |SHAP|	% of Total
Temporal EUI History	EUI median (8-yr panel),	Median annual site EUI over 8-year Portfolio Manager panel (kBtu/ft²/yr)	0.317	49.5%
Temporal EUI History	Spatial Lag EUI	Spatial analysis of EUI from nearest neighbor, k =5	0.067	10.5%
Temporal EUI history	EUI standard deviation	Standard deviation of annual site EUI, measures year-to-year variability	0.038	5.98%
Building Valuation	FMV per square foot	Fair market value divided by parcel area, proxy for property development intensity	0.033	5.15%
Building Valuation	Building FMV	Fair market value based on building structure and improvements, estimated by Fulton County Board of Assessors	0.027	4.26%
Building Characteristics	Building age	Years since construction (current year minus year_built_e from Fulton County Tax Parcels)	0.024	3.73%
Building Valuation	Land FMV	Fair market value based on land parcel estimated by Fulton County Board of Assessors	0.017	2.70%
Temporal EUI History	EUI coefficient of variation	Ratio of EUI standard deviation to EUI mean, normalized variability measure	0.012	1.84%
Socioeconomic Costs	Census-tract Gini Index	Gini coefficient of income inequality at census-tract, derived from ACS five-yar estimates	0.011	1.64%
Temporal EUI History	n_years_obs	Number of consecutive years of data from Atlanta's benchmarking data	0.009	1.49%
Building Characteristics	Built pre-1940	Binary: 1 if built before 1940, 0 otherwise. One-hot encoded from decade_built factor.	0.009	1.39%
Building Characteristics	Assessed values	Taxable assessed value of property recorded by Fulton County Tax Assessor	0.005	0.84%
Building Characteristics	Log gross floor area	Natural log of gross floor area, captures nonlinear size effects	0.005	0.81%

Figure 2. SHAP beeswarm: feature importance for Atlanta parcel energy model

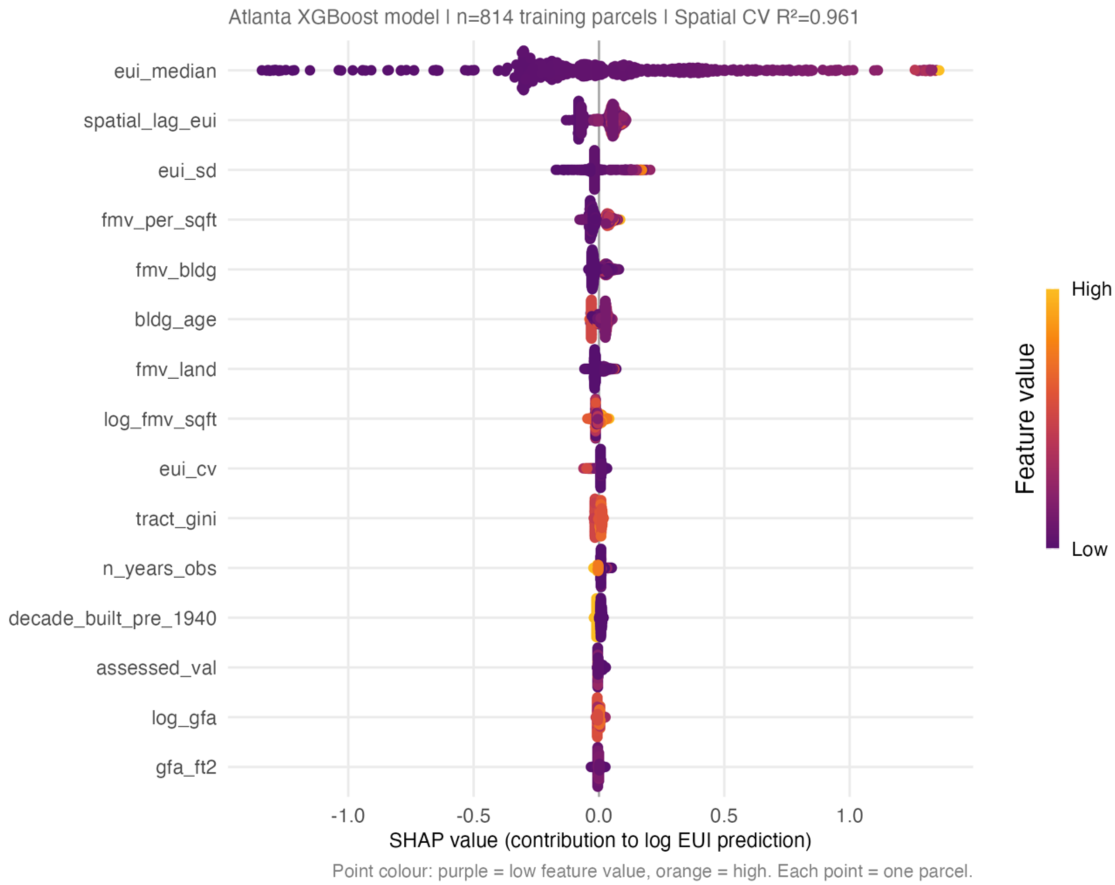

Figure 2. SHAP beeswarm plot showing the top 15 predictors of building energy use intensity (EUI). Each point represents one Atlanta parcel (n=823). The x-axis shows the SHAP value -- the contribution of each feature to the log EUI prediction for that parcel, where positive values increase predicted EUI and negative values decrease it. Point colour encodes the feature value from low (purple) to high (orange/yellow). Model: XGBoost with spatial k-fold CV, R2=0.853.

Finding significance: The hierarchical feature matrix shows temporal EUI history as the dominant predictor (59%) for future commercial energy use in Atlanta. Building valuation (15%) and EUI of neighboring parcels (11%) emerge as secondary and tertiary predictors, which shows the importance of the building's physical characteristics and the community's energy usage to predicting a city-wide commercial energy consumption. These findings support prioritizing data collection and investment towards benchmarking parcels to understand how to develop tailored improvement strategies.

Zone Comparison Analysis

Descriptive zone comparison analysis revealed statistically meaningful variation in mean EUI across zone types (Table 4). EDZ-only parcels consumed 32.5 kBtu/ft²/yr on average, compared to 37.5 for the baseline group, a difference of 14%. OZ-only parcels showed the highest mean EUI of all zone groups at 52.7 kBtu/ft²/yr which is 40% above baseline. Parcels co-designated as EDZ-and-OZ parcels were 18% above baseline with an average EUI of 44.1 kBtu/ft²/yr. On average, parcels in historic districts appeared to consume more energy (62.4 kBtu/ft²/yr observed) than the XGBoost machine learning model predicted (46.4 kBtu/ft²/yr), suggesting preservation constraints may limit necessary energy upgrades solely not captured in building age and valuation features.

Table 4. Energy use intensity and effect sizes by zone type (Atlanta, n=823)

Zone Type	Count (n)	Mean EUI (kBtu/ft²/yr)	vs Baseline	Cohen's d (raw EUI)	Magnitude
Baseline (neither)	579	37.5	--	--	Reference group
EDZ only	141	32.5	-13.3%	-0.117	Very small
OZ only	26	52.7	+40.5%	0.347	Small
EDZ and OZ	68	44.1	+17.6%	0.154	Very small
Historic district	9	62.4	+66.4% observed; model under predicts by 16.0 kBtu/ft²/yr	0.598	Medium (n=9, directional)

Figure 2. Energy use intensity by economic zone designation (Atlanta, n=1,015)

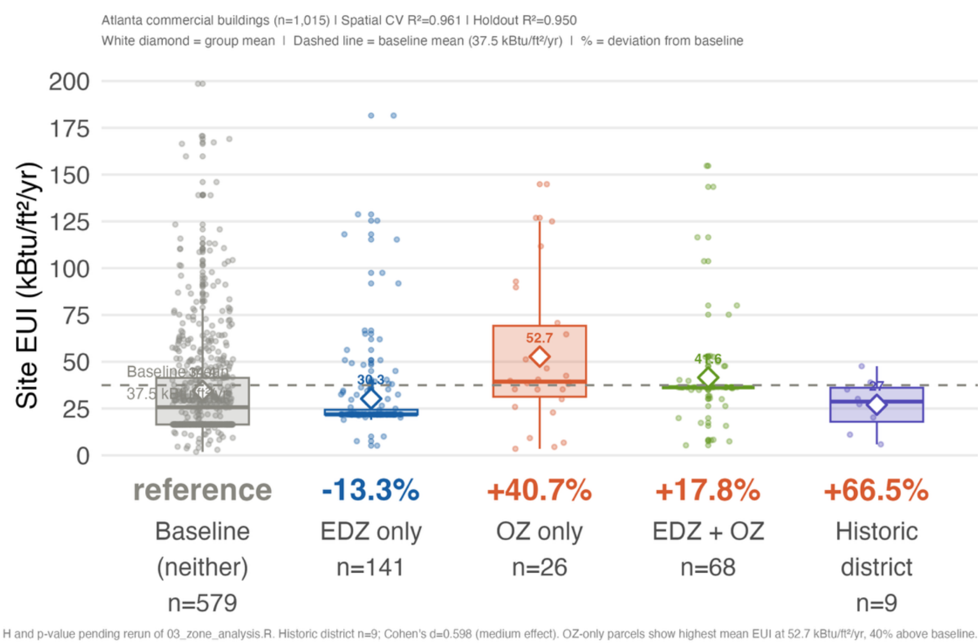

Figure 2. Boxplots of mean annual site EUI (kBtu/ft²/yr) across four zone designation groups. Each grey dot represents one parcel. The white diamond shows the group mean. The dashed horizontal line marks the baseline mean of 37.5 kBtu/ft²/yr. Historic parcels show the largest departure from baseline (mean 62.4, Cohen's d=0.598).

Table 6. Data dictionary for Figure 2 axis labels and chart elements

Label	Description
Baseline (neither)	Parcels with no zone designation. Reference group (n=579, 75% of training set).
EDZ only	Atlanta Economic Development Zone parcels outside federal Opportunity Zones (n=141).
OZ only	Federal Opportunity Zone parcels (2017 Tax Cuts and Jobs Act) outside EDZ boundaries (n=26).
EDZ + OZ	Parcels co-designated under both EDZ and federal OZ programs (n=68). Policy overlap group.
Mean EUI (kBtu/ft²/yr)	Mean annual site energy use intensity in kBtu per square foot per year, averaged across the 8-year panel. Normalized by gross floor area.
Box (IQR)	Interquartile range -- middle 50% of EUI values. Box edges are 25th and 75th percentile.
Whiskers	Extend to 1.5x IQR. Individual dots beyond whiskers are outlier parcels.
White diamond	Group mean EUI: 37.5, 30.3, 52.7, 41.6, 27.0 kBtu/ft²/yr for baseline, EDZ, OZ, EDZ+OZ, historic.
Dashed horizontal line	Baseline mean (37.5 kBtu/ft²/yr) -- visual reference for zone-level deviations.

Effect size analysis provides meaningful zone variations with predicted EUIs. The EDZ-only parcels consume less energy than baseline parcels showing a potential association between city economic development zone programs and attracting lower intensity commercial use case. OZ-only designation showed the highest mean EUI across all groups 52.7 kBtu/ft^2/yr, 41% above baseline with a Cohen's d =0.347. Atlanta parcels in Federal Opportunity Zones also had a larger distribution of older, higher EUI commercial buildings.

The Wilcoxon rank-sum test for historic district parcels in Atlanta (n=9 historic versus n=1,006 non-historic) had a location shift estimate of 2.5 kBtu/ft²/yr. Cohen's d = 0.598 (medium effect). Parcels in historic districts show higher observed EY at 62.4 kBtu/ft^2/yr versus 37.4 kBtu/ft^2/yr for non-historic buildings while the XGBoost model underpredicts EUI on average of 16 kBtu/ft^2/yr. Given Savannah's substantially larger historic district sample, transferring this analysis to Savannah's parcel predictions offers a more robust opportunity to examine this relationship and assess whether the directional pattern observed in Atlanta holds at greater scale. The primary historic district analysis is enabled by the Savannah transfer, where the landmark district covers a substantial share of the downtown commercial core.

Finding significance: Statistical parcel-level zone comparisons via XGBoost ML model revealed varying energy performance across Federal Opportunity Zone parcels (OZ parcels), Economic Development Zone parcels (EDZs), and historic district parcels. Historic districts show underpredictions for EUI despite high observed energy consumption, highlighting the need for specialized interventions. This finding is highly valuable and synergistic to local expertise and findings provided from project partners and stakeholders across both cities. EDZ parcels have lower energy use intensities while OZ have higher EUIs when compared to baseline parcels. These findings also suggest linking improvement efforts with stratified zone variations and local insights across to address the realistic state of energy consumption as predicted efficiencies do not always match observed loads.

Transferability: Savannah Predictions

The Atlanta-trained XGBoost model was applied to Savannah's full parcel inventory following schema alignment. Features present in Atlanta but absent in Savannah were filled with NA values and imputed using Atlanta training-set medians through the recipe's preprocessing pipeline. This approach preserves within-Savannah parcel variation while imputing missing features, which is standard practice for cross-city model transfer.

The transfer produced parcel-level EUI predictions across Savannah's commercial building stock (R^2=0.307). The predicted EUI distribution closely approximated the Atlanta training distribution in central tendency, validating the framework's external consistency. Because Savannah lacks a large observed benchmarking dataset for direct calibration, the predictions serve as illustrative outputs showing replicability rather than validated ground-truth estimates. The low transferability of the Atlanta-trained XGBoost model for Savannah and lack of Savannah-based EUI data in the Atlanta-trained XGBoost highlights the need for a tailored Savannah machine learning model instead of transferring the Atlanta-based model.

Finding significance: The low spatial prediction transfer results (R^2=0.307), lack of Savannah-based EUI data, and the need to impute missing predictors using Atlanta-training-set medians demonstrates the need to create a Savannah-tailored ML model rather than transferring and using the Atlanta-trained XGBoost on Savannah. Lacking comprehensive temporal and longitudinal EUI datasets limits 59% of cities, urban planners, and utilities' ability to create tailored energy and cost-saving measures to improve building performance based on data-driven city-wide energy consumption predictions.

5. Conclusions

This study demonstrates that parcel-level commercial energy consumption can be modeled with high accuracy using publicly available data and transferred across cities with contrasting economic and spatial contexts. The XGBoost model trained on Atlanta's 8-year benchmarking 1,015-parcels panel achieved a spatial cross-validation R^2 of 0.961 and an independent holdout R^2 of 0.950, confirming robust generalization. A transfer simulation producing R^2 of 0.307 establishes an honest cross-city prediction baseline for data-limited cities, where longitudinal benchmarking data could account for the gap between 0.307 and 0.950.

Key findings for commercial program design and implementation include the following:

- Temporal EUI history is the strongest predictor of commercial energy performance. Buildings with high median EUI and high variability across years are the consumers well positioned for tailored interventions. Longitudinal benchmarking programs generate the most valuable data for predictive targeting.

- Building age predicts commercial energy consumption independent of other features. Pre-2000 construction corresponds to systematically higher EUI after controlling for size, value, and neighborhood context. Age-based targeting criteria for retrofit programs are empirically supported by this analysis.
- Consumption-based metrics of economic well-being, including energy expenditure as a share of income, reveal household welfare impacts that income data alone obscures (Attanasio and Pistaferri 2016), suggesting that programs targeting economically constrained commercial corridors may generate co-benefits for both building performance and community economic stability.
- EDZ-designated parcels consumed 13% less energy than baseline on average, with a small but consistent effect size (d = -0.12) that persisted after controlling for building characteristics. Findings suggest that city-level economic development zone programs may be attracting lower-intensity commercial uses, while federal OZ designation, which targets capital investment broadly, has the highest energy consumption patterns at the parcel level of any group. Local economic development theory emphasizes that the interaction of multiple economic factors, rather than single designations alone, drives measurable changes in regional infrastructure performance (Malizia et al. 2020).
- GIS-based frameworks that direct resources to specific geographic areas have demonstrated effectiveness in improving energy outcomes where targeting is precise (Gaspari et al. 2020), reinforcing the case for parcel-level benchmarking over zone-level program design.
- Historic district buildings consume substantially more energy than their structural characteristics predict, with the model under-predicting by 15.4 kBtu/ft²/yr on average. These buildings should treated as a high-priority and distinct program target rather than assumed to be lower-energy buildings on account of their preservation status. Restrictions on exterior modifications that limit envelope improvements, aging mechanical systems protected from replacement, and higher occupancy intensity in rehabilitated historic commercial spaces.

Limitations

The Atlanta training set of 1,015 commercial, mixed-use, and multifamily buildings reflects the constraint of voluntary benchmarking programs, where large commercial buildings with building automation systems are overrepresented, while small commercial properties and owner-occupied retail are underrepresented. Zone comparison statistics now have adequate statistical power for the four main zone groups, though the historic district sample (n=9) remains small pending rerun with the modifed Savannah-based ML model. The Savannah transfer produces predictions with an estimated R^2 of 0.307 under data-limited conditions; ground-truth validation awaits integration of audit data and utility billing records described in ongoing work.

Future Work

Expanding the framework to Georgia cities with different economic profiles, such as industrial centers, university towns, and coastal tourism economies, would test the geographic bounds of transferability. Incorporating utility interval data where available would enable subannual energy consumption modeling and demand response targeting. Extending the framework to map distribution grid infrastructure using machine learning and publicly accessible geospatial data (Wang et al. 2023) would connect parcel-level energy consumption predictions to grid capacity constraints, directly informing utility investment priorities. Given evidence that power grid infrastructure investment serves as a catalyst for regional economic expansion (Xu et al. 2021), integrating building-level demand predictions with grid capacity data represents a natural next step for linking commercial energy program design to infrastructure planning. Interactive geospatial platforms for city planning partners, building on the parcel-level prediction outputs generated here, represent a direct path from this analytical framework to real-time program implementation and evaluation support.

Funding

This work was authored by the National Laboratory of the Rockies for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308. Funding provided by applicable Department of Energy office and program office, U.S. Department of Energy Office of Critical Minerals and Energy Innovation Integrated Energy Systems Office Building Technology Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

Anderson, D.M., et al. 2014. ’Assessing National Employment Impacts of Investment in Residential and Commercial Sector Energy Efficiency: Review and Example Analysis.’ Pacific Northwest National Laboratory.

Anser, M.K., et al. 2025. ’Urban energy efficiency in China: Examining the role of renewable energy, smart grids, and sustainable design through spatial and policy perspectives (1990-2022).’ Energy and Buildings 339.15.

Attanasio, O.P., and L. Pistaferri. 2016. ’Consumption Inequality.’ The Journal of Economic Perspectives 30(2): 3-28.

Bergstra, J. and Bengio, Y. 2012. “Random search for hyper-parameter optimization.” *Journal of Machine Learning Research* 13:281-305. https://jmlr.org/papers/volume13/bergstra12a/bergstra12a.pdf

Chen, T., and C. Guestrin. 2016. "XGBoost: A Scalable Tree Boosting System." Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. https://doi.org/10.1145/2939672.2939785

Clarke, A., M. Deru, E. Lockhart, et al. 2023. 'Centering Energy and Environmental Justice in the Buildings Energy Sector.' NREL/TP-5500-83173.

DeWitt, P., & Clarke, A. (2026). fairadi: A FAIR-Oriented U.S. Deprivation Indices Dataset (v1.0.0) [Data set]. Zenodo. https://doi.org/10.5281/zenodo.19222629

Gaspari, J., et al. 2020. 'A GIS-Based Methodology for Speedy Energy Efficiency Mapping: A Case Study in Bologna.' Energies 13(9).

Humphreys, J.M. 2020. 'The Economic Impact of Georgia's Deepwater Ports On Georgia's Economy in FY 2019.' Selig Center for Economic Growth, University of Georgia.

Husslage, Bart G. M., Gijs Rennen, Edwin R. van Dam, and Dick den Hertog. 2011. "Space-Filling Latin Hypercube Designs for Computer Experiments." Optimization and Engineering 12, no. 4: 611–30. https://faculty.sites.iastate.edu/tesfatsi/archive/tesfatsi/ExperDesignABM.HyperCubeDesign.Husslage2006.pdf.

IEA. 2025. 'World Energy Outlook 2025.' International Energy Agency.

Langner, R., B. Hendron, S. Pless, M. Huppert, and R. Cochrane. 2013. 'Industry Research and Recommendations for Small Buildings and Small Portfolios.' NREL/TP-5500-57776.

Lugner, Moa, A. Rawshani, Edvin Helleryd, and Björn Eliasson. 2024. "Identifying Top Ten Predictors of Type 2 Diabetes Through Machine Learning Analysis of UK Biobank Data." Scientific Reports 14, no. 1957. doi.org

Malizia, E., et al. 2020. 'Understanding Local Economic Development.' Second Edition. Taylor and Francis.

Mayer, M and Stando, A. 2025. shapviz: SHAP Visualizations. R package version 0.10.3. https://cran.r-project.org/web/packages/shapviz/index.html.

Osberg, L., and A. Sharpe. 2002. 'An Index of Economic Well-Being for Selected OECD Countries.' The Review of Income and Wealth 48(3): 291-316.

Theodos, B., and J. Gonzalez. 2019. 'Building Community and Wealth in Disadvantaged Commercial Corridors: Chicago's Neighborhood Opportunity Fund.' Urban Institute.

U.S. DOE (Department of Energy). 2021. 'Low-Income Energy Affordability Data (LEAD) Tool.' Office of Energy Efficiency and Renewable Energy.

U.S. EIA (Energy Information Administration). 2022. '2018 Commercial Energy Consumption Survey Final Results.'

Wang, Z., A. Majumdar, and R. Rajagopal. 2023. 'Geospatial mapping of distribution grid with machine learning and publicly accessible multi-modal data.' Nature Communications 14.

Xu, Z., et al. 2021. 'Does power grid infrastructure stimulate regional economic growth?' Energy Policy 155.

York, D., C. Cohn, D. Morales, and C. Tolentino. 2022. 'Building Decarbonization Solutions for the Affordable Housing Sector.' American Council for an Energy-Efficient Economy.