Forecasting Energy Demand in Large Commercial Buildings Using Support Vector Machine Regression

David Solomon[1], Rebecca Winter[2], Albert Boulanger[3], Roger Anderson[3 1], Leon Wu[3]

[1]Department of Earth and Environmental Science, Columbia College, Columbia University
[2]Department of Earth and Environmental Engineering, Fu Foundation School of Engineering and Applied Sciences,
[3]Center For Computational Learning Systems, Columbia University

As our society gains a better understanding of how humans have negatively impacted the environment, research related to reducing carbon emissions and overall energy consumption has become increasingly important. One of the simplest ways to reduce energy usage is by making current buildings less wasteful. By improving energy efficiency, this method of lowering our carbon footprint is particularly worthwhile because it reduces energy costs of operating the building, unlike many environmental initiatives that require large monetary investments. In order to improve the efficiency of the heating, ventilation, and air conditioning (HVAC) system of a Manhattan skyscraper, 345 Park Avenue, a predictive computer model was designed to forecast the amount of energy the building will consume. This model uses Support Vector Machine Regression (SVMR), a method that builds a regression based purely on historical data of the building, requiring no knowledge of its size, heating and cooling methods, or any other physical properties. SVMR employs time-delay coordinates as a representation of the past to create the feature vectors for SVM training. This pure dependence on historical data makes the model very easily applicable to different types of buildings with few model adjustments. The SVM regression model was built to predict a week of future energy usage based on past energy, temperature, and dew point temperature data.