U.S. Natural Gas & Electric Company

Company Profile

This public utility provides natural gas and electricity to 6.7 million of Michigan's 10 million residents, serving customers in all sixty-eight of the state's Lower Peninsula counties.

Products & Services

Our team used GPU accelerated deep learning to increase efficiency, upgrade performance, and significantly boost productivity.

Our team also installed and upgraded a number of Python libraries, including TensorFlow, H2O, Keras, Imblearn, Torch, XGBoost, CatBoost, and Fastai.

Pacific Data Integrators

Business Challenges

As the client was unable to reduce utility consumption in the commercial sector, they set their focus on reducing residential sector consumption. One method to achieve this was the energy-efficiency programs AC Peak Cycling, Smart Thermostat, and Peak Time Rewards. However, when customers did not sign up for the appropriate power-saving program that fit their needs:

Residential sector consumption was not effectively reduced

Potential electricity savings were not realized, costing the company potential revenue

Customers were dissatisfied with the results/benefits of the program they signed up for

Compounding the problem were some additional roadblocks:

The success of a previous energy program overshadowed the newly introduced Smart Thermostat program, believed to be a highly lucrative program in the current market

The inability to perform customer outreach via the right communication channel hampered marketing efforts and drained marketing funds.

Enrollment data was absent as some programs were in the early stages of operation. This created a scarcity of insights that could have otherwise helped create program predictions

Project Goals

The primary goal of this project was to create a program-based recommender model for 1.3 million residential Demand Response Contract Accounts. This model would predict customer preference for three energy-efficiency programs – AC Peak Cycling, Smart Thermostat, and Peak Time Rewards – and rank the programs in order of each customer's preference.

Additionally, the company wanted a channel-based recommender model to find the optimal outreach channel for customers enrolled in the AC Peak Cycling program. This predictive model would rank email, direct mail, and phone in order of customer preference, to determine the most effective mode of communicating with each customer.

Project Results

Our team aggregated the customer prediction results of three different prediction models: H2O's Gradient Boosting, LightGBM and Deep Neural Networks. Based on these, our team built robust Machine Learning-based predictive models that analyzed customer usage and demographics data to accurately predict the probability of a consumer enrolling in one of three programs. Out of 74,000 early enrollments, 94% of customers who joined the AC Peak Cycling program were successfully predicted by our team's model.

Our team also developed the ability to make accurate recommendations for which channel of communication a customer would prefer (i.e. email, direct mail, or phone). This empowered the company to more accurately budget their marketing spend for each communication channel and reach out to customers via the most effective channel.

Additional Project Results

Overcame the data insufficiencies of nascent programs with boosting and sampling techniques, providing results on par with the company's data requirements

Optimized existing code and updated modules to streamline the data pipeline

Increased productivity, operational performance, data integration capabilities, and overall data quality

Upgraded the ML server to handle any Machine Learning framework

Streamlined SQL queries and made SQL and Python aspects of the code modular for ease of handling and usability

Wrote custom-code to optimize the memory usage of datasets, thereby decreasing runtimes and increasing computational efficiency