An Analysis Framework for Grid Planning and Demand-Side Strategy Evaluation

Craig Simmons, Tanushree Charan, Aditie Garg, National Laboratory of the Rockies (NLR)

ABSTRACT

Buildings represent a significant share of electricity demand, and their growing load has important implications for grid capacity and infrastructure planning. This rising electric load can lead to an increased strain on the grid often requiring costly upgrades. There is a need to assess the condition of the current state of grid capacities while planning for future capacity requirements. This project introduces an analytical framework to model demand- and supply-side scenarios, offering actionable insights to guide grid planning, enhance affordability, and inform investment strategies. The framework is developed by leveraging URBANopt, an analytical platform for high resolution district scale modeling of the distribution grid and grid edge assets including buildings. The analytical framework leverages a machine learning (ML) approach that samples ResStock or ComStock building archetypes and calibrates them using advanced metering infrastructure (AMI) load data, to accurately model the demand side building loads. This matching process enables the inference of building characteristics such as heating, cooling, and water-heating system types from AMI load shapes, allowing the framework to translate meter-level data into meaningful end-use attributes. Baseline conditions are modeled in URBANopt using the selected stock models alongside distribution grid models and load information. Based on the load modeling for the selected stock models scenarios can be created to analyze demand side management technologies such as energy efficiency, load flexibility to reduce peak load and defer or avoid grid upgrades. This simulation-based framework highlights the value of granular demand side modelling as a key resource for grid planning and can enable decision makers such as utilities, regulators de-risk technologies before incurring cost-prohibitive capital expenditures.

Introduction

Electricity demand at the grid edge is undergoing rapid change, driven by evolving building loads and other increased electrical loads that expose constraints at the distribution level (Specian et al, 2021). Utilities increasingly encounter feeders that experience both winter and summer peaks, with heterogeneous load mixes and limited visibility into how customer-level behaviors translate into coincident demand at critical nodes. These pressures challenge traditional planning practices, which rely on worst-case static load profiles, and can precipitate capital-intensive upgrades when the magnitude, timing, and spatial distribution of load growth are not well understood. Oversimplified load assumptions run the risk of over-investment in grid infrastructure without adequate consideration for dynamic building load patterns or potentially cheaper building-centric load management strategies (Schwartz, 2024).

Conventional distribution planning workflows generally rely on simplified or static load shapes, standard customer classes, and aggregated assumptions about load factors, seasonal diversity, and end-use penetration. These abstractions can obscure the hourly and sub-hourly dynamics observed in AMI data, especially where technology and distributed energy resources (DER) adoption create new coincidence patterns amongst customers and the feeder. Utilities often lack a method to connect high-resolution AMI signals with models in a scalable way, limiting the fidelity of feeder studies and the ability to test non-wires alternatives (NWAs) with confidence. Moreover, even when detailed feeder models exist, existing workflows lack the ability to integrate demand-side and supply-side resources in a unified analytical environment, hindering repeatability and the systematic exploration of scenarios.

The analysis framework presented provides a granular, combined data- and physics-based analytical approach that integrates AMI data, ResStock/ComStock physics models, and distribution system modeling to bridge the gap between customer-level behavior and grid operations. The approach supports utilities by enabling:

- Machine learning-based selection of ResStock physics-based building models to match AMI load shapes, improving the fidelity of load forecasts and scenario simulations.
- Technology adoption and demand mitigation scenario testing, including energy efficiency, load flexibility, and demand response deployment helping to quantify avoided infrastructure costs and reliability impacts.

  Potential to include additional distribution loads and generation as enabled through the URBANopt platform

The framework enables utilities to characterize current load behavior, classification of load types, quantify where and when demand flexibility might be available, and translate those findings into implications for grid reliability, customer affordability, and near-term program and infrastructure investment planning.

Approach

This paper introduces a simulation-based analytical framework for grid edge planning built on URBANopt, where AMI-matched ResStock models produce granular building load profiles that feed an OpenDSS representation of the distribution system. The method formalizes a workflow in which: (i) stock models are sampled and filtered using ML inference on expected inputs based on AMI profiles, (ii) AMI profiles are matched to candidate stock models via a multi-objective ML algorithm, and (iii) matched composite load shapes are mapped to feeder nodes for steady-state and quasi-static power flow analyses. The framework then supports policy- and investment-relevant scenarios, including energy efficiency upgrades, load shifting and curtailment strategies, targeted technology adoption pathways, and additional load integration. This allows utilities to quantify peak reduction, define reliability margins of projections, and estimate capital investment options at the locations where constraints are most acute. The framework enables actionable scenario outputs that directly inform utility investment decisions and regulatory review by quantifying the economic and reliability of demand-side management (DSM) and NWAs before committing to capital-intensive infrastructure upgrades. The framework is designed initially for utility personnel who need a transparent and repeatable method to connect AMI data with physics-based building and feeder models. By maintaining the workflow as an open-source codebase, it can also be readily adopted by consultants, regulators, and other stakeholders who benefit from a shared, extensible analytical platform.

A current limitation of the workflow is its reliance on OpenDSS for distribution system modeling, which, while flexible and open-source, is not widely used in utility planning environments. Future development could focus on enabling direct ingestion of feeder models from commonly used proprietary platforms such as CYME, improving compatibility with utility-standard workflows. Additional enhancements could also include packaging the end-to-end workflow in a form that can be more seamlessly integrated back into URBANopt or incorporated into proprietary software, allowing for broader adoption and smoother deployment within existing planning tool chains.

Framework Tools

URBANopt:

URBANopt (Urban Renewable Building and Neighborhood optimization) is an EnergyPlus and OpenStudio-based simulation platform aimed at district- and campus-scale thermal and electrical analysis. Originating as National Laboratory of the Rockies (NLR) Laboratory-Directed Research and Development (LDRD) project, URBANopt is being elevated to a more formal position within BTO's ecosystem of tools. URBANopt includes capabilities for modeling shared thermal systems, supporting integration with grid and DER analysis tools. The project is an open-source software development kit that is suitable for integration into commercial products and piloting the underlying analytics in real world projects. (U.S. Department of Energy, n.d.)

ResStock

The U.S. residential building sector stock energy model, ResStock™, is a highly granular, bottom-up model that uses multiple data sources, statistical sampling methods, and advanced building energy simulations of the residential building stock across the contiguous United States (Reyna et al. 2025). The ResStock data sets include geographically relevant building energy models that are leveraged directly in the proposed analysis framework.

OpenDSS

OpenDSS is an open-source electric power distribution system simulator (DSS) designed to support grid modelling and network analysis in a distribution network model. It enables engineers to perform complex analyses using a flexible, customizable, and easy to use platform intended specifically to meet current and future distribution system challenges and provides a foundation for understanding and integrating new technologies and resources (Electric Power Research Institute 2024).

AMI Matching

To represent the behind the meter residential demand for the distribution feeder, the AMI load profiles are aligned with physics-based building energy models from the ReStock dataset. The AMI load profiles represent the hourly electricity consumption profile for each customer associated with the distribution feeder. The load profiles do not include any behind the meter generation data associated with solar installations. While these profiles a useful in understanding gross load electricity consumption patterns, they do not contain details on the properties of associated buildings such as envelope properties, system types or occupant behavior. To bridge this gap, representative building models are matched to the observed AMI load profile. The ResStock dataset is a national set of detailed building models generated using a combination of data-driven and physics-based energy modeling methods. They provide representative characteristics for the entire US building stock (Reyna et al. 2025). A detailed methodology is developed that takes in high level inputs such as site location, the weather file and runs an optimization routine to the difference between ResStock and AMI loads. This optimization routine has been previously developed and validated for similar feeder-level studies (Speake and Parker 2024).

For this study, the AMI-load mapping methodology is applied as follows. First, for the given feeder location, relevant adjacent counties are filtered down to achieve a large enough sample of building energy models but with similar building characteristics for the AMI matching. For the selected counties (approximately 600 models), the models are simulated for the local weather conditions for the AMI year using a 2024 AMY weather file. This ensures that the simulated load profiles reflect the same weather conditions as the measured AMI consumption data. An optimization routine is then used to identify the ResStock model that best matches each AMI load profile.

In addition to the optimization approach to streamline model matches between ResStock and AMI, an inference model is developed to ensure key inputs between ResStock and AMI are aligned. The input parameters considered for the input model are:

- Presence of air conditioning
- Presence of electric heating systems
- Presence of electric water heating systems

Three logistic regression models are trained to predict the probability of the target input variable. The RessStock load profiles are used as the training dataset, since the input characteristics are known for those. The models predict the probability of each attribute for a given AMI load profile. Only high-confidence predictions are used to constrain the candidate pool of ResStock models, ensuring robustness against misclassification. This inference step reduces the search space of the optimization problem and improves physical consistency between matched models and observed loads. The objective of the optimization is to identify the ResStock building model whose hourly load profile, peak heating and peak cooling load matches the AMI electricity consumption pattern. The objective function minimizes three key errors:

- Root Mean Squared Error (RMSE) of difference in hourly load profiles: The objective function minimizes the RMSE between the hourly ResStock and AMI load profiles. This ensures the model captures the overall seasonal and temporal patterns in the electricity consumption at an hourly resolution
- Peak Heating Error: For the heating season, the top 10 daily peak loads are identified and averaged for both AMI and ResStock profiles. The peak heating error is defined as the difference between these averages.

- Peak Cooling Error: Similarly, for the cooling season, the average of the top 10 daily peaks is computed, and the difference between AMI and ResStock values defines the peak cooling error.

Minimizing the peak heating and peak cooling error ensures that the ResStock models are within the expected heating and cooling peaks for the loads associated with the distribution feeder. The optimization routine allows specifying multiple ResStock matches to find for each AMI meter and enables specifying thresholds of for the errors used in the optimization so that it is possible that a given AMI profile may not return any ResStock matches if thresholds are not met. This optimization balances between representing the accuracy of the load shape and representing the accuracy of the peak demand. Figure 1 shows two examples of matching AMI and ResStock profiles as determined by the algorithm. Note that Figure 1 depicts a single day, but the matching algorithm analyzes the entire annual hourly load.

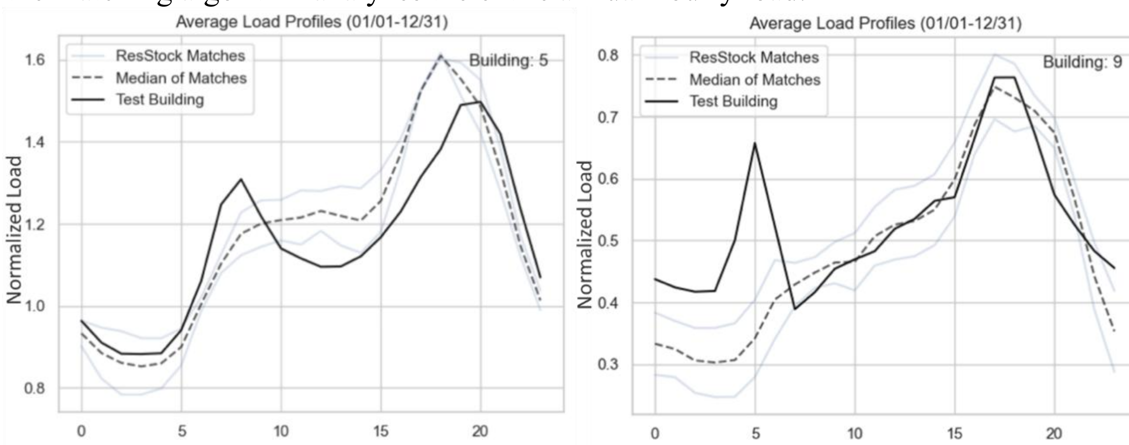

Figure 1. Two examples of a daily load profile results from the AMI matching algorithm. In both cases there are two ResStock matches for the given AMI profile and the average of the matched ResStock (dashed line) carries forward as the modeled profile for the meter in the OpenDSS model.

Results

The proposed framework is analyzed for a distribution feeder in the New York region with total 237 loads (154 residential and 83 commercial loads). The baseline feeder is modelled in OpenDSS (D. Montenegro et. al.), enabling time-series analysis of voltages, loading, and constraint violations for different load profiles. The representative feeder is highlighted in Figure below.

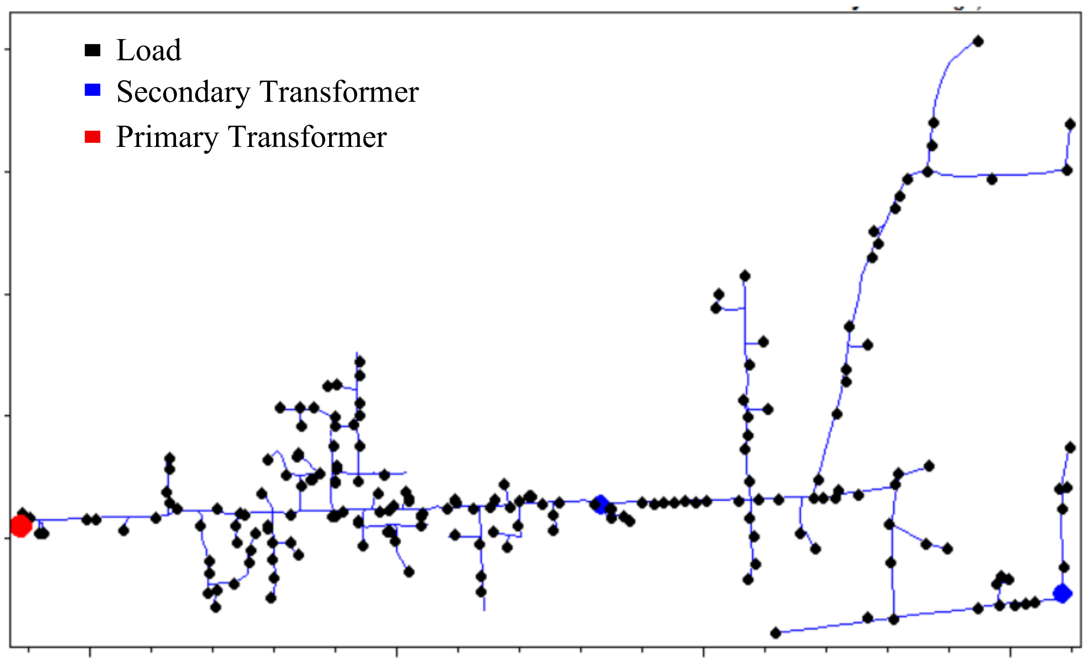

Figure 2. Feeder Line Diagram

The feeder's total rated load is 2.2 MW (1.2 MW residential and 1.0 MW nonresidential). AMI data is available for 95% of residential customers. Using the available residential AMI data (15-minute resolution) and assuming no DER on the feeder, the aggregate feeder load profile is shown in Figure . The peak load is observed in winter, with a peak of 366 kW in November 2024. This results in a rated residential load that is more than 67% higher than the observed residential peak.

The overall feeder peak (aggregated residential and commercial load) is observed in winter with a peak load of 435 kW in November 2024 and summer peak of 423 kW in July 2024.

The overall feeder voltage profile (assuming no generation on the feeder) showcases voltages to fall within acceptable 5% tolerance as highlighted in Figure 4.

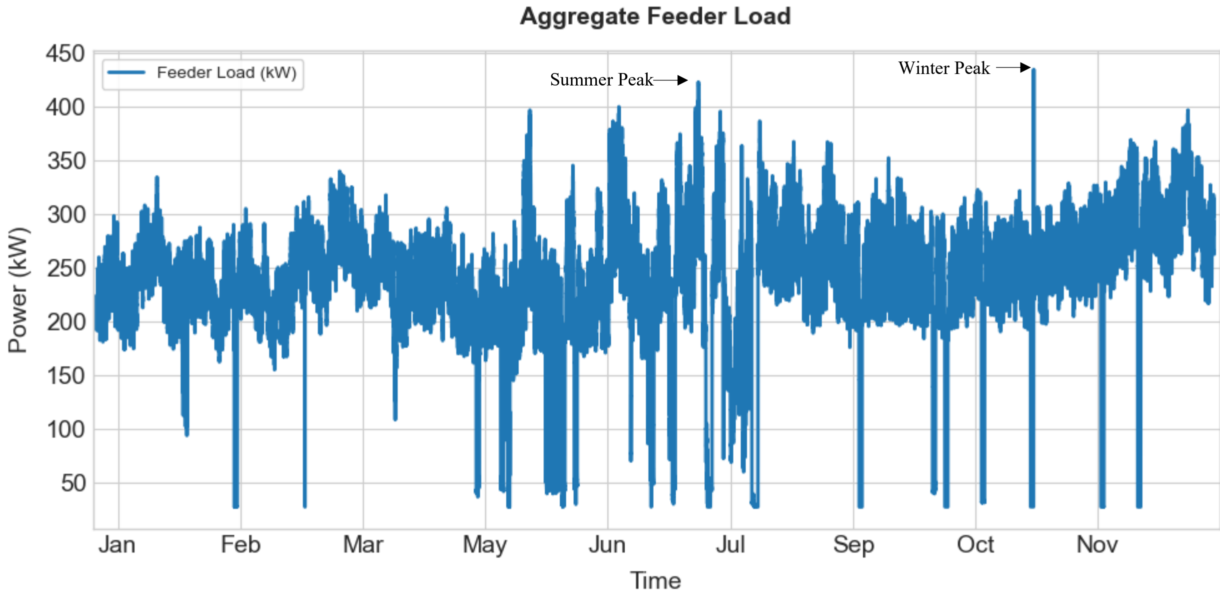

Figure 3. Feeder Aggregate AMI Load Profile

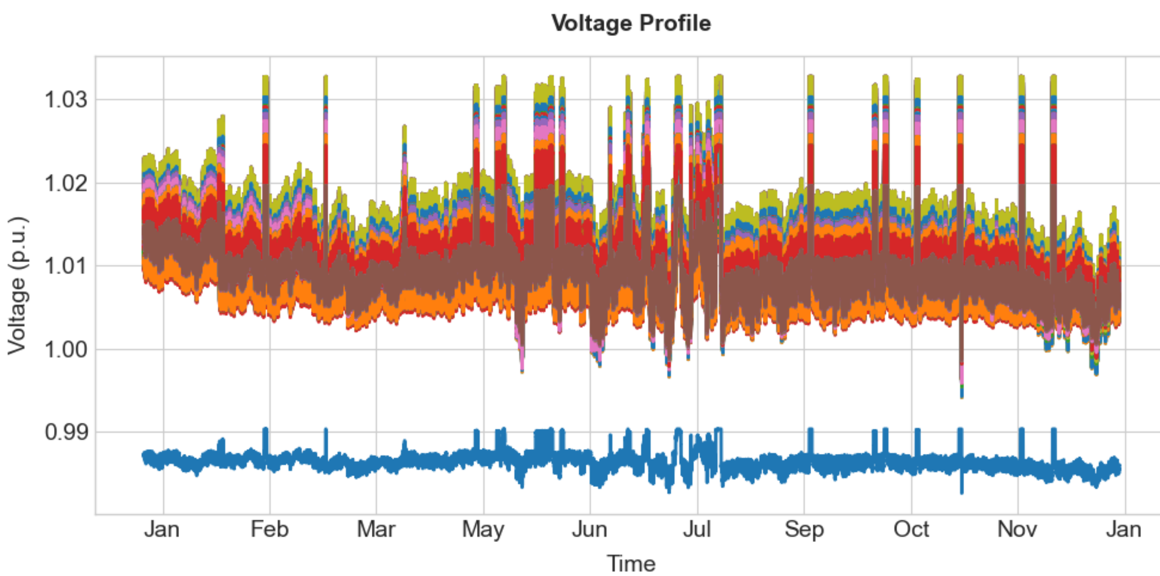

Figure 4. Feeder voltage profile for all nodes on the feeder

To map the residential AMI meter data with ResStock load models the framework utilizes weather data (AMY 2024) from a location 50 miles away from the feeder head along with filtering the ResStock datasets within the 7 counties near the feeder location. Based on these criterions the AMI profiles were matched with the representative ResStock load models. It is observed at an aggregate level the monthly CV-RMSE (Coefficient of Variance of Root Mean

Squared Error[1]) of 28.8%, and average daily CV-RMSE[2] for each month of 9.6% as highlighted in Figure 5. Higher summer CV-RMSE is expected due to increased electric cooling heterogeneity, while winter variability is lower where space heating is predominantly natural gas.

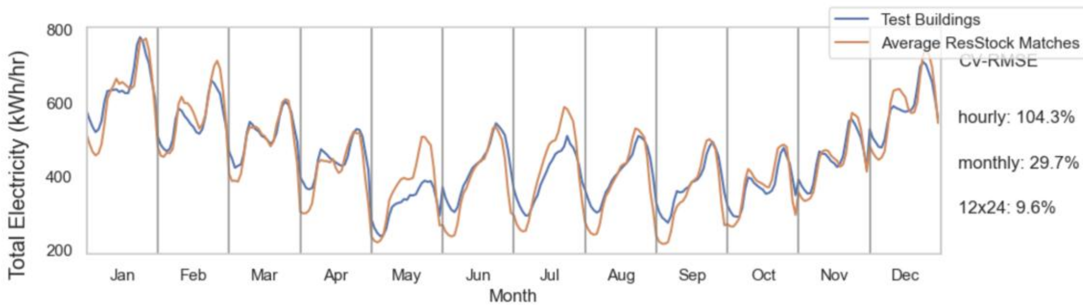

Figure 5. AMI – ResStock Matching Results. Displayed above are the hourly profile for an average day within each month.

Evaluating the matching metrics shows that monthly errors are lower and less variable than hourly errors. Hourly profiles are more sensitive to occupant behavior, making one-to-one alignment between AMI and ResStock profiles more challenging. [3] Monthly aggregation smooths these effects, yielding monthly CV-RMSE typically in the ~17–32% range, while peak cooling and heating errors remain under 10% highlighted in Figure .

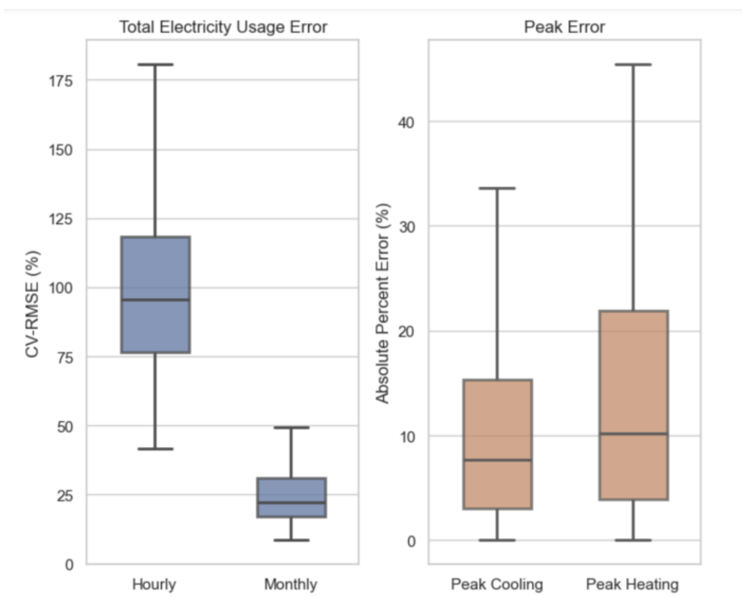

Figure 6. Error metrics AMI vs ResStock load mapping

[1] Monthly CV-RMSE = CV-RMSE applied to monthly energy consumption
[2] Average daily CV-RMSE for each month = CV-RMSE applied to each month's average 24-hour hourly load profile. This is particularly important for the utility use case where capturing seasonal and daily profiles is a specific goal.
[3] Note this analysis used a limited set of inputs related to weather and county data to match AMI data with ResStock models. The results can be improved using other matching criterion.

To evaluate building archetypes and end use usage in the residential loads the AMI interval data (15-minute resolution) are mapped to building archetypes using AI based clustering and feature mapping approach defined earlier yielding granular end-use load shapes by segment and location. This information is evaluated for various locations across the feeder including feeder head, and some key branches within the feeder. These branches are highlighted in Figure . The analysis studies three locations: feeder head that includes all the residential loads (146 loads) in the feeder, Branch 1 hosting 10 residential loads, Branch 2 hosting 11 residential loads.

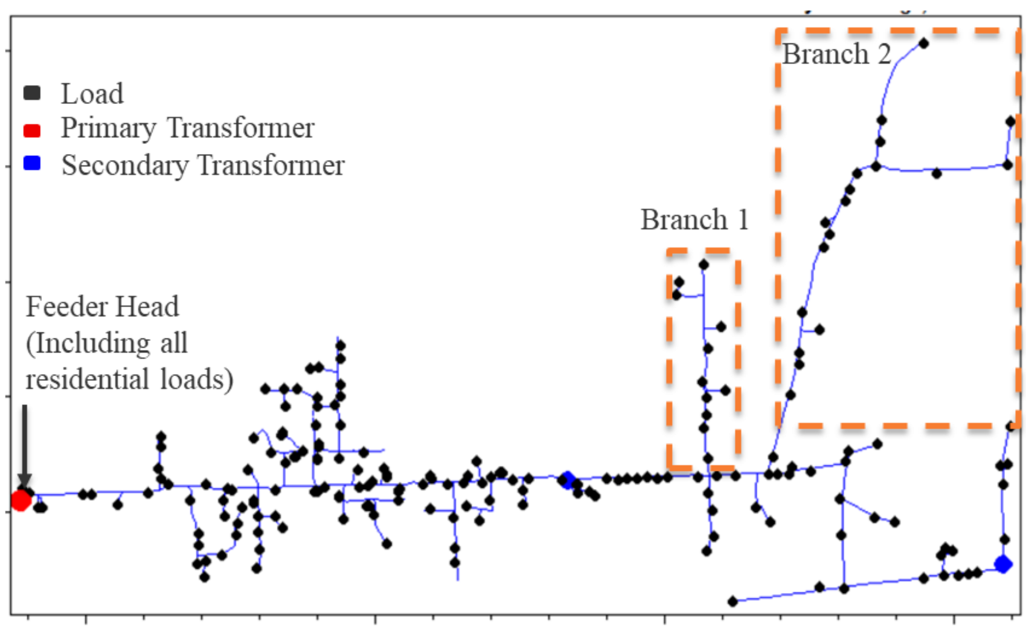

Figure 7. Branches under consideration for end use load types. Locations include: feeder head including all residential loads, Branch 1 and 2 host 10 and 11 residential loads respectively

With matched ResStock models serving as proxies for the residential units at each meter NLR is able to infer that major characteristics of the models (e.g., heating and cooling sources) are representative of the buildings on the feeder, and the project team is exploring data sources to validate this approach. As highlighted in Figure comparing the type of heating fuel distribution type across these locations it is observed that natural gas acts as the primary heating fuel serving 50% of the homes at feeder head and Branch 1 while electricity is the primary heating fuel for 50% of homes in Branch 2. A minority of other fuel oil types are observed across all cases. Thus, in this case utility programs targeting technology adoption or demand response/demand flexibility (DR/DF) implementation would vary significantly if the utility were trying to address winter distribution load conditions in Branch 1 versus Branch 2.

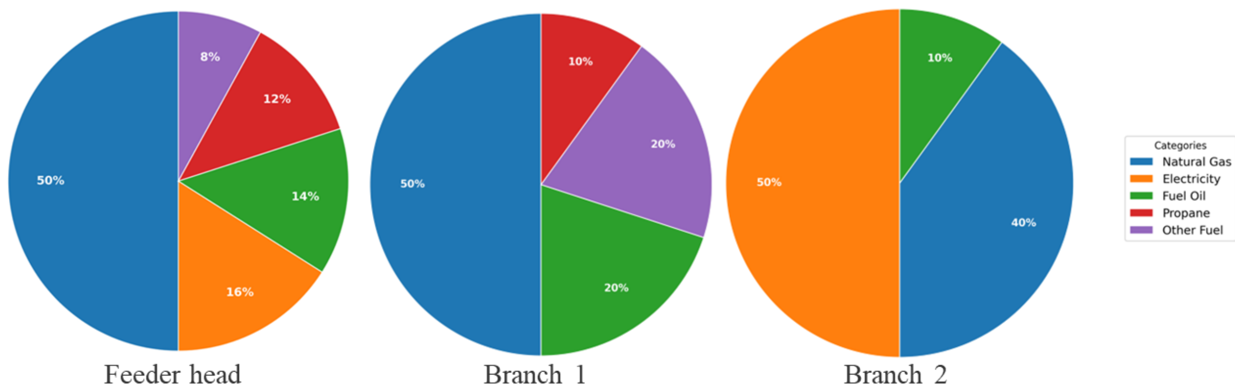

Figure 8. Heating fuel distribution across various feeder segments

Reviewing the hot water fuel type distribution highlighted in Figure reveals a similar pattern of natural gas and electricity fuels where natural gas is observed as the primary fuel for hot water load in feeder and Branch 1 while Branch 2 showcases electricity as the primary hot water fuel type. Therefore, depending on the specific locations of potential distribution stresses and technology adoption and load control objectives for the utility, the programmatic approach to relieve and control grid stress would differ greatly between Branch 1 and 2 despite being very close geographically.

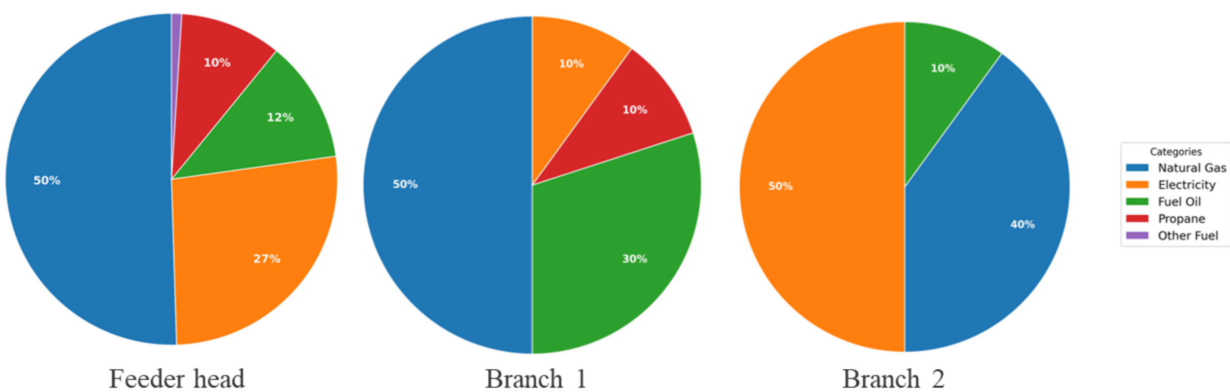

Figure 9. Hot water fuel distribution across various feeder segments

Comparing the type of cooling distribution across these locations highlighted in Figure it is observed room and central AC remain the major candidates for cooling load while a minority of the load is observed due to heat pumps. Thus, in this case planning investments could focus on heat pump adoption as the amount of penetration on this feeder seems low with an overall heat pump type cooling load noted less than 5% across the feeder. Moreover, given the major cooling is seen due to AC load a utility program related smart thermostat could be successful in this location.

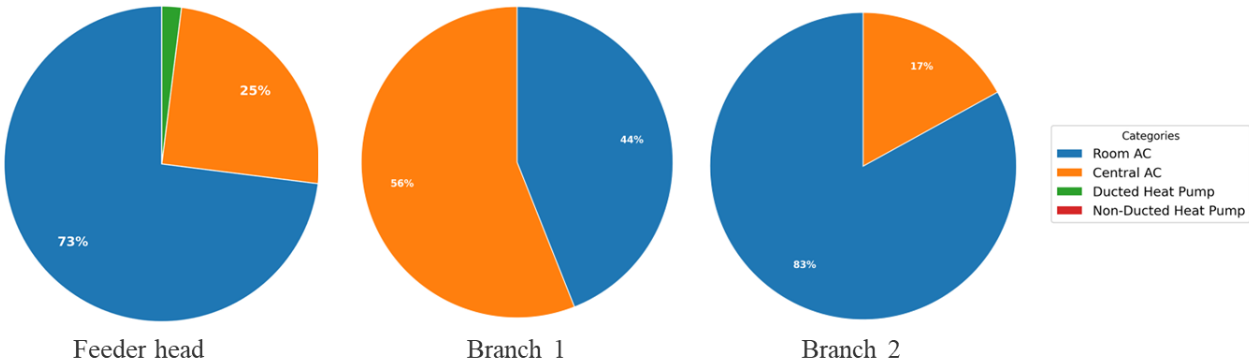

Figure 10. Cooling distribution type across various feeder segments

Reviewing these results the utility can prioritize very specific branches for technology adoption and better predict the impacts of feeder-wide DR/DF programs via smart thermostats or hot water controls. Program design can then sequence enabling controls ahead of large-scale HVAC technology adoption, target upgrades where voltage margins are tight, and set performance targets such as load reduced at peak, and avoided upgrades. This approach also streamlines measurement and verification by comparing observed AMI outcomes to archetype-based counterfactuals at the feeder head and selected branches, improving benefit attribution and cost-effectiveness. Thus, the results provide a data-driven basis to target utility programs by location and end use, enabling better-aligned planning and customer offerings.

Conclusion

This work presents a simulation-based analytical framework that connects measured feeder-level electricity demand with physics-based building stock models and detailed distribution system simulations. By integrating AMI data with ResStock building models through an optimization routine that minimizes hourly RMSE and peak-season errors, the framework provides a calibrated representation of building loads that reflects local weather, building characteristics, and customer behavior. These AMI-matched archetype models are then combined with validated OpenDSS feeder representations to quantify current operating conditions and to evaluate demand-side scenario outcomes at both feeder and branch scales.

The results demonstrate that the proposed methodology captures key spatial and temporal variations in load, enabling utilities to identify where localized demand is driving winter or summer peaks and where targeted interventions (e.g., heat pump water heaters, smart thermostats) may yield the greatest operational benefits. By translating AMI data into building-level characteristics and feeder-specific load patterns, the framework offers a defensible means of identifying the most impactful efficiency measures, demand-flexibility strategies, and other building-centric interventions that can mitigate grid stress. These insights provide a quantitative foundation for utility planning, supporting more informed decisions about program design, load management strategies, and the prioritization of capacity upgrades.

The analysis further shows that the approach can distinguish how end-use characteristics differ across branches within the same feeder, revealing location-dependent opportunities for targeted mitigation strategies. Such granular insights help utilities avoid oversimplified planning assumptions based on system-wide averages and instead focus on where interventions yield the highest value. In doing so, the framework improves the transparency and precision of planning by connecting measured load data to physics-based counterfactuals and enabling scenario-level comparisons grounded in observed customer behavior.

Future work can enhance these capabilities by expanding the integration of demand side load types, quantifying coincident peak contributions at finer spatial resolution, and evaluating the bounds of feasible demand flexibility across representative building archetypes and feeder branches. Additionally, the approach can be extended to examine incremental hosting capacity with load growth scenarios before thermal or voltage limits are reached. These extensions would further strengthen the technical basis for location-specific planning and enhance the utility of AMI-matched stock models in grid-modernization studies. Overall, the results illustrate how a combined data- and physics-driven modeling workflow can provide actionable, evidence-based insights to support grid planning, reduce uncertainty, and improve the cost-effectiveness of distribution-level investment decisions.

References

D. Montenegro, R. Dugan, J. Taylor and M. McGranaghan, "Open-source software projects for advancing the power systems analysis," *2022 Open Source Modelling and Simulation of Energy Systems (OSMSES)*, Aachen, Germany, 2022, pp. 1-6, doi: 10.1109/OSMSES54027.2022.9768968.

U.S. Department of Energy. "URBANopt." Accessed March 18, 2026. https://www.energy.gov/cmei/buildings/urbanopt.

Reyna, J., Fontanini, A., Parker, A., Present, E., Liu, L., Adhikari, R., Bianchi, C., Brossman, J., Chintala, R., Clark, K., Harris, C., Horowitz, S., Lou, Y., Maguire, J., Merket, N., Moore, N., Munankarmi, P., Phoung, S., Robertson, J., Sandoval, N., Speake, A., Stenger, K., White, P., Wilson, E. (2025). *ResStock Technical Reference Documentation, Active Development Version* [Technical report] National Renewable Energy Laboratory (NREL). Retrieved from https://NatLabRockies.github.io/ResStock.github.io/assets/trd/ResStockTechnicalReferenceGuide_2025_1.pdf

Electric Power Research Institute. (2024). *OpenDSS* [Software documentation]. Retrieved March 17, 2026, from https://www.epri.com/pages/sa/opendss

Speake, Andrew, and Andrew Parker. 2024. A Framework for Identifying Building Energy Models of Localized Utility Service Areas Using Smart Meter Data. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-90443. https://www.nrel.gov/docs/fy25osti/90443.pdf

Specian, M., C. Cohn, and D. York. 2021. Demand-Side Solutions to Winter Peaks and Constraints. Washington, DC: ACEEE. www.aceee.org/research-report/u2101.

Schwartz, Lisa C. “Integrated Distribution System Planning.” Presentation, Office of Electricity Grid Controls and Communications Program Review, Lawrence Berkeley National Laboratory, October 30, 2024. https://www.pnnl.gov/sites/default/files/media/file/Lisa%20Schwartz%20Berkeley%20Lab%20IDSP%20peer%20review%2020241030.pdf