

Responsible Artificial Intelligence for Predicting Electricity Demand in Residential Communities

Leona K. Mercer

Assistant Professor, Department of Energy Analytics and Sustainable Systems
Northbridge Metropolitan University, Cedar Vale, Oregon, USA.

Corresponding Author: Leona K. Mercer

Assistant Professor, Department of Energy Analytics and Sustainable Systems
Northbridge Metropolitan University, Cedar Vale, Oregon, USA.

Article Type: Simulation-Based Research Article

Submitted: 02-January-2024 **Revised:** 09-January -2024

Accepted: 16-January -2024 **Published:** 23-January -2024

Abstract: Artificial intelligence has gained increasing relevance in residential electricity-demand forecasting due to the temporal variability, behavior-dependence, weather-responsiveness, and linear-ignorant characterizability of household electricity consumption. Alongside accuracy improvements, emerging high-capacity machine-learning approaches in residential contexts present proliferation challenges for privacy, transparency, uncertainty, model drift, data minimization, and human responsibility. In this work, we construct a community-based simulation-aided Responsibility-AI methodology for short-term electricity-demand forecasting with accuracy and governance attributes. We simulated 1,036,800 residential household-hour observations (at hourly resolution) for a synthetic community of 240 households over 180 days, and evaluated four forecasting architectures (persistence, conventional long short-term memory [LSTM], explainable LSTM, and a Responsibility-AI ensemble of recurrent modeling, gradient boosting, to reduce data input sets, to estimate uncertainty, to increase transparency, to track drift, to promote human escalation) using community-based simulated performance metrics.

The Responsibility-AI architecture performed best regarding community-average overall forecasting performance, with MAE0.278kWh, RMSE0.386kWh, MAPE7.9%, PeakΔE6.8%, and 90% PI-Coverage94.1%. Its modeled privacy-risk score was reduced to 28/100 from 72/100 under community LSTM. A secondary logistic modeling approach revealed that hot extreme temperature, missing data, and comparatively volatile loads increased the modeled

probability of high-risk forecasting events, while modeled uncertainty escalation, drift monitoring, and human review reduced the modeled probability of high-risk forecasting events. Thus, our community-based methodology affirms that forecasting accuracy is compatible with the principles of human-centered governance, which considers transparency, privacy proportionality, calibrated uncertainty, operational reliability, and accountable oversight along with accuracy.

Keywords: responsible AI; residential electricity load; demand forecasting; smart meters; explainable AI; privacy-preserving load forecasting; uncertainty quantification; energy communities.

Introduction

Prediction of residential electricity demand is playing an expanding role in modern energy management, as the deployment of distributed renewable generation, electric vehicles, smart meters, home energy-management systems, demand-response schemes, and battery storage necessitates better short-term prediction of the time and way in which electricity will be used. At the residential scale, however, energy demand can be very variable. Consumption can be affected by temperature, occupancy, family routines, appliance use, working schedules, holidays, behavioral adaptation, and other unexpected phenomena. This combination of time-dependence and human behavior variance makes residential load prediction a very natural form for artificial intelligence and machine learning.

Deep-learning methods have shown real promise in modeling non-linear household demand. Long short-term memory fields are capable of modeling relationships across several hours or days, and attention mechanisms and ensemble methods could provide further flexibility in discontinuous consumption scenarios. However, the quality of a forecast should not be judged solely by lower mean absolute error or mean absolute percentage error. With residential forecasting, data is often sourced from private homes, and precise electricity data can potentially reveal information about occupancy, derived plans or household features.

Hence, the real challenge is not only about improving the accuracy of forecasts but about creating forecasting regimes that are consistent with responsible governance of data.

The present research focuses on five interconnected problems:

- **Privacy exposure:** over-utilization of detailed household data might uncover behavioral schedules not relevant for community prediction.
- **Uncertainty hiding:** traditional point forecasts can hide how uncertain the model actually is.
- **Limited explainability:** complex models can give technically correct predictions without necessarily indicating whether demand is going up or down.
- **Distribution shift:** anomaly temperature. consumption habits changes.. missing data. new technological developments. etc. which all reduce the accuracy after deployment.
- **Automation risk:** operationally consequential predictions that are made without human intervention in situations where model confidence is low.

The focus of this work stems from the hypothesis that a technologically advanced forecasting system should provide both prediction and governance; thus, effective electricity forecasting should inherently optimize beneficial accuracy while limiting unnecessary data extraction, dominance detection, interpretability, status tracking, and responsible human administrators.

Literature Review

a) Evolution of Residential Electricity-Demand Forecasting

The traditional techniques used for electricity-demand forecasting are regression, autoregressive modeling, exponential smoothing, and other statistical techniques. These are still useful where the demand pattern is fairly predictable, but with residential energy data there tend to be nonlinear dependencies, discontinuous multiple periodicities, unexpected peaks, and other behavioral variation. Suganthi and Samuel explored the wide spectrum of energy-demand forecasting methods and confirmed the ongoing shift away from pure mathematical modeling into computational and adaptive models.

The depth of this methodological space was further extended by the rising access to smart-meter data. Wen, Zhou and Yang found deep learning applicable to residential-building load forecasting; Son and Kim showed that deep-learning architectures can adequately describe residential electricity-demand patterns; Kiprijanovska and others built a household day-ahead forecasting framework based on deep learning; all emphasizing on the significance of recurrent architectures to power consumption in the residential sector.

Several characteristics repeatedly emerge as important determinants of forecast quality:

- Lag structure analysis; historical space use; historical electricity goes demand 4.5.2. Modeling of the historical electricity price behavior 4.5.3. Modeling of the historical electricity demand behavior used in the analysis of the short-term electric load prediction.
- Variations over the course of the day and week
- Outdoor temperature and humidity
- Holidays and atypical calendar circumstances
- In general, supply pressure also includes: volatility of demand, which is influenced by consumption pattern change, and (consumption volatility and change in occupancy pattern
- Discrepancies caused by: null hypotheses that are not normally distributed alter the prediction and residual distributions;1. false fails;2. false positives.3.
- In relation to which we will have to assess the implications of what has been said for the appropriate culture for each. In addition, whether the aspects of community and geography are part of the considerations as well as the aggregation level.

Implication (1): These data imply that the predictive performance is influenced by the model structures as well as the surrounding environment.

b) Probabilistic Forecasting and Uncertainty

While point forecasting provides an expected demand value, uncertainty is inherent in actual electricity consumption. Taieb et al. were among the first to show the value of quantile-based uncertainty forecasting for smart-meter data, while Shepero et al. have pioneered probabilistic residential-load forecasts with Gaussian processes. Wang et al. have used pinball-loss guided LSTM architectures for probabilistic individual-load forecasts.

This work gets at themes and questions closely related to responsible AI, since uncertainty is not only a statistical property but also a governance signal. When a prediction interval widens or a system drops in confidence this way, a responsible system will be able to react in a different manner than a regular forecasting pipeline.

The present framework therefore treats uncertainty as:

- The degree to which a forecasting method produces consistently correct results

- An indicator of how likely a distribution change is
- A trigger for completed further diagnostic assessment
- - a human escalation threshold when operational consequences are involved

This framework starts to put probabilistic forecasting into operational governance from just describing what we do.

c) Privacy and Explainability in Residential Energy Analytics

The smart-meter data exposes privacy issues, since detailed electricity-consumption behavior might disclose privacy-sensitive details of the household activity. Beckel et al. have shown that certain household details might be predicted from the smart-meter consumption data, and Giaconi, Gündüz and Poor highlight the ongoing trade-off between information and privacy in smart metering.

Thus, protection of privacy may require more than simply removing or anonymizing names and addresses. If any fine-grained temporal information can still be associated with particular households, potential inferences may still be drawn. A good practice for forecasting is to ask whether each variable is needed for the analysis.

Provide another governance mechanism, explainability. Ribeiro, Singh, and Guestrin, previously mentioned, proposed local model explanations that explain individual predictions based on easily understood features. Later, Lundberg and Lee proposed SHAP, a unified measure that explains a prediction by defining the contribution of each variable.

For residential forecasting, explanation mechanisms can help answer practical questions such as:

- Was the predicted peak all temperature pushed?
- Was recent electricity demand relative to previous usage too high?
- Was it an anomalous observation driven by missing data?
- Is there a calendar variable that unexpectedly forecasted and possibly even dominated the forecast?
- Does the feature dependence of the model change compare to past?

The importance of these questions increases when forecasts lead to automated demand response or storage dispatch.

Research Gap

The electricity-demand forecasting literature to date has achieved significant advancements in building more accurate prediction models. However, most of these studies have employed performance measures such as MAE, RMSE, MAPE, coefficient of determination, or other similar measures of predictive-performance for the assessment of prediction models. These measures are relevant and important but do not, on their own, offer complete proof that a forecasting approach is suitable for operational use in residential dwellings.

The first significant area of unmatched accuracy-centric evaluation. A model can maintain a low mean forecasting error and be poorly calibrated during heat waves, demand peaks, missing observations and atypical community responses. A focus on average accuracy hides the potential for costly, high-risk failure modes.

Another gap is privacy proportionality. It is generally assumed in energy forecasting literature that whatever knowledge there is that can potentially improve forecast accuracy should be used. Residential data comes from private domestic environments. The existence of value in being able to predict does not necessarily mean that storing and collecting it is proportionate.

A third dimension of gap focuses on explainability and the integration with uncertainty. All the aforementioned literature have separately examined interpretable machine learning, probabilistic forecasting, privacy-preserving approach, and robustness of the models for the respective application. There are relatively fewer frameworks that incorporates all of those dimensions into a practical architecture.

A fourth gap is a lack of coverage of human escalation. Most prediction systems are automated, running predictions on an ongoing basis but little discussion is directed at deciding when to take predictions automatically, and when to escalate.

This paper helps fill these gaps by framing forecasting electric power consumption in residential sector in such a way that predictive accuracy, privacy, explanation, uncertainty, robustness, monitoring and accountable intervention are all taken into consideration.

Novelty / Proposed Contribution

Our key innovation in this paper is a Responsible-AI forecasting architecture with ethical guardrails embedded in model operation rather than bolted on after prediction.

The framework contributes in several ways:

- It measures the model accuracy in conjunction with the exposure of privacy and the calibration of uncertainty.
- It adds data minimization to the presentation of a privacy statement before the model training instead of at the end.
- Provides an interpretable means of translating uncertainty in the prediction into an operational trigger.
- It adopts the use of interpretability as part of an operational process of a routine forecast interpretation.
- Monitors drift to determine whether conditions of historical validation still remain valid for current operation.
- It differentiates routine, low-risk forecasts from the consequential or non-routine forecasts, which require human intervention.
- It assesses the errors in the peak-demand values independently of the mean forecast errors.

The proposed system thus does not consider responsible AI as one more dimension indicator. Rather, as I argued in the previous section, responsible AI is what results from the interplay of data governance, prediction models, managing the bias, explanation, monitoring, and human accountability.

Research Objectives / Research Questions

The main goal for this project is the design and on-line assessment of a simulation-based Responsible-AI framework for residential-community electricity-demand prediction.

The specific objectives are:

1. 4. To enable a comparison of the conventional and governed forecasting architectures based on known measures of forecast error.
2. To observe whether privacy aware feature reduction is possible with competitive predictive performance.
3. To check whether prediction-interval coverage can be used as a proxy for calibration of uncertainty.

4. Investigate high risk forecasting situations related to temperature, volatility, and missing observations.
5. Between the following. For each risk decomposition, the risk associate with. How we would estimate the protective associations of uncertainty escalation, human review, and drift monitoring.
6. To develop a comprehensive framework for the assessment of residential electricity demand forecasting.

The study addresses the following research questions:

Extention 1: Caused by the similar accounting (and ensemble method) to what the many LSTM ensemble forecast; we are interested in:

RQ1: Is the ensemble of Responsible-AI as comparable in predictive accuracy with classical LSTM?

RQ2: Is lack of residential data minimisation achievable in a modelled privacy risk framework without substantial loss of accuracy?

RQ3 inquires: what operating conditions lead to the highest simulated probability of predicting the system to fail?

RQ4: Does the use of uncertainty escalation, human review, and drift monitoring decrease the number of modeled high risk forecasting events?

RQ5: Either.. Is the overall assessment of residential forecasting quality derived from a multidimensional evaluation superior to that obtained by benchmarked accuracy-only benchmarks?

Methodology

a) Research Design

A simulation-based comparative research design was employed. Four different forecasting architectures were assessed in a shared environment with synthetic resid... In the same environment and with the same synthetic resid...

The research was constructed as a methodological simulation because there was no empirical utility or household data set supplied. Therefore, the numbers published in the manuscript

must not be read as performance measurements for any actual household community, electricity company or commercial forecasting system.

b) Dataset / Sample

The process of selecting and collecting data. When the original data are in the form of observations and measurements I think it is necessary to be fully explained for the importance of data collection and make the reader believed the reliability and validity of the data for further analysis.

The synthetic analytical dataset reflected 240 households logged for 180 successive days at hourly resolution. This created 1,036,800 consumer-hour observations prior to aggregation. Synthetic variables were created to simulate reasonable relationships between nearest prior electricity demand, temperature, humidity, calendar variables and short-term volatility.

Table 1: Operational Variables Used in the Simulation

Variable	Operational Definition	Measurement	Analytical Function
Hourly Demand	Simulated data of residential electricity usage	kWh	Forecast target
Temperature	Simulated outdoor temperature	°C	Exogenous predictor
Humidity	Simulated relative humidity	%	Exogenous predictor
Lagged Demand	Demanis for the last hours	kWh	Temporal predictor
Hour of Day	The time framework of the words aside from the hours times is called the hourly calendar position.	0–23	Behavioral predictor
Weekday Indicator	Weekday versus weekend. It is interesting to compare the characteristics for the weekday versus weekend days. Most indicators are lower for weekdays, but also more spread out. There is much more variety in the weekend's measures.	Binary	Calendar predictor

Holiday Indicator	Simulated public holiday condition I constructed the simulated public holiday condition, which is similar to the simulated school holiday condition in design. In order to gather the effect of public holiday, the same questionnaire in experiment one was sent to public whereas the responses received from public.	Binary	Calendar predictor
Load Volatility	Variations in demand over the short-term.	Continuous	Risk indicator
Prediction Interval	Simulated 90 predicted percentage of the actual data set.	Continuous	Uncertainty measure
Drift Flag	Distributional deviation indicator	Binary	Governance trigger
Human Review	Pre-defined escalation of high-risk forecast.	Binary	Governance control

Table 1 lists the variables that compose the synthetic residential-demand environment. The model is limited by design to omit personal identity, specific household address, appliance-level behavioral measures, employment data, demographic variables, and household variables unassociated with energy demand. This is a case of limiting the model inputs to those variables that can be justified as related to energy-demand forecasting. Lagged demand and calendar variables supply time-series information, while weather variables supply environmental drivers. Prediction intervals, drift flags, and human review indicators add robustness to the dataset by passing beyond pure point accuracy.

c) Instrumentos / Herramientas

The traditional forecasting architecture was modelled as an LSTM that could learn the time dependence of residential demand. The Responsible-AI architecture was a combination of recurrent prediction and gradient-boosting estimates in a stacked ensemble approach.

Attribute of features using SHAP was used. Model drift was modelled by modifying statistical properties of current data in comparison to training data distribution.

The analytical techniques included:

- descriptive error analysis
- MAE calculation
- RMSE calculation
- MAPE evaluation
- The variation of the percentage error for the peak-demand
- prediction-interval coverage
- privacy-risk scoring;
- The variable of the log of Bankloans is multiple binary logistic regression;
- uncertainty-based escalation;
- model-drift monitoring.

d) Experimental Procedure

The time series were split into training, validation, and testing periods in chronological order so that there would be no “leakage” from the future into an earlier stage of prediction.

All forecasting frameworks were provided with the same forecasts inputs with the exception of privacy aware feature minimization that limits information collection. Predictions were then assessed under standard and challenging conditions.

High-risk forecasts were operationally defined through one or more of the following triggers:

- Publish lengthen prediction; intervals that are longer than the width set.
- Table 3. The percent of missing observations in 5502. The percentage of missing data above 5%.
- Identify some statistical drift;
- B1B was for some reason unusually volatile community demand;
- extreme-temperature conditions;
- Community peaks (prediction above an operation threshold).

Though triggered Forecasts still entered a conceptual human-review side pathway.

Graph 1: Synthetic Residential Community Electricity-Demand Profiles




Fig 1 depicts the synthetic hourly demand profiles representing a typical residential day and a high temperature day (both include the recognized morning and evening value peaks). The high temperature situation results in a drawn out escalating afternoon load and a significantly more pronounced evening peak, which is a result of the simulated temperature-dependent residential electric demand. The graph illustrates why forecasting architecture cannot simply be based on historical averages. Environmental conditions can result in both change in the magnitude and in the time profile of loads throughout communities. These values are synthetic and are incorporated solely to illustrate the nature of the simulated demand environment.

e) Evaluation Metrics

The following metrics were used to measure: forecasting performance MAE, RMSE, MAPE, peak-demand error, prediction-interval coverage, and a simulated privacy-risk score.

We used MAE because it translates to an easily understood mean of error. RMSE is biased toward outliers, giving more penalty to large deviations, which is important in distribution-system planning, where large errors may be more impactful. MAPE allows relative comparison at different points of demand, although less stable as demand goes to zero.

Peak-demand error was examined independently as community electricity management often hinges on accurately determining the level of periods in which network tension is most extreme.

B.1 Uncertainty-Calibration should be reflected in the prediction interval coverage which was the main measure of the calibration of uncertainty. The expectation is that about 90% of the realized values will be within any nominal 90% prediction interval.

The privacy-risk score could be anywhere between 0–100, 0 being the lowest amount of information exposed and 100 the highest simulated score. It should be noted that this is a statistical construct rather than a formally tested, general privacy index.

Results


The simulation revealed significant differences in the four forecasting architectures. Persistence forecasting was found to have the poorest overall forecasting performance. Conventional LSTM significantly improved forecasting accuracy, while Explainable LSTM achieved similar accuracy while providing better uncertainty calibration and less modelling privacy leakage. The Responsible-AI Ensemble was found to have the best overall compromise between accuracy, uncertainty calibration, peak-demand accuracy, and privacy governance.

Table 2: Simulated Forecasting Performance Across Residential-Demand Architectures

Architecture	MAE (kWh)	RMSE (kWh)	MAPE	Peak-Demand Error	90% Interval Coverage	Privacy-Risk Score /100
Persistence	0.438	0.612	12.8%	15.6%	72.0%	78
LSTM	0.291	0.409	8.4%	10.1%	78.4%	72
Explainable LSTM	0.304	0.421	8.7%	9.2%	87.6%	51
Responsible-AI Ensemble	0.278	0.386	7.9%	6.8%	94.1%	28

Table 2 shows that responsible governance does not have to be associated with a forecasting penalty. The Responsible-AI architecture has the lowest MAE (0.278kWh) and RMSE (0.386kWh). Compared to persistence forecasting the the MAE has been reduced by 36.5%. Then the conventional LSTM has also achieved a high performance (reaching a MAPE of 8.4%, but an interval coverage of 78.4% still given well below the modeled 90% nominal target. Explainable LSTM optimized the coverage reaching 87.6% and increased the privacy-risk score from 72 to 51. Responsible-AI architecture has reached the 94.1% coverage but with the privacy-risk score at the far lower.28.

Graph 2: Simulated Forecasting Error Across Residential-Demand Architectures



Graph 2 presents the results of MAPE and peak-demand error for the four architectures in the four simulated forecasts. The results from persistence forecasting are the worst, showing MAPE to be 12.8% and peak-demand error at 15.6%. Conventional LSTM lowers both these measurements to 8.4% and 10.1% respectively. Explainable LSTM yields an insignificant higher MAPE of 8.7%, yet also demonstrates better peak-demand than conventional architecture at 9.2%. Responsible-AI architecture produces the lowest MAPE, measuring 7.9% and peak-demand error, measuring 6.8%. This is a crucial graph to understand, as it is not only the MAPE that is saving the disciplines here, but the peak-demand error as well. The largest gains have been shown in the peak-demand forecast, rather than only the average

forecast. This supports the govern approach as being most proficient when planning out peak operations scenarios.

Table 3: Simulated Logistic Regression Predicting a High-Risk Forecast Event

Predictor	B	SE	Odds Ratio	95% CI for OR	95% CI for OR	p-value
Extreme Temperature	0.854	0.196	2.35	1.60–3.45		<.001
Missingness >5%	1.061	0.221	2.89	1.88–4.45		<.001
High Load Volatility	0.752	0.184	2.12	1.48–3.04		<.001
Uncertainty Escalation	−0.734	0.163	0.48	0.35–0.66		<.001
Human Review	−0.916	0.199	0.40	0.27–0.59		<.001
Drift Monitoring	−0.562	0.171	0.57	0.41–0.80		.001

Table 3 lists the modeled conditions associated with high risk (above 5%) forecast events. The violation of greater than 5% missingness has the strongest positive association, with an odds ratio of 2.89. Extreme temperature has an odds ratio of 2.35, while high load volatility has an odds ratio of 2.12. These results show how issues of data quality and external conditions can affect the accuracy of forecasting, even in situations where the base model is performing well on average. Reviewing the model output indicator provides the strongest level of sensing, with an odds ratio of 0.40 in this situation. Uncertainty escalation has an odds ratio of 0.48, while drift monitoring has an odds ratio of 0.57. The simulated regression supports a governance architecture in which shortfalls detect abnormal operating conditions and trigger additional safeguards.

Discussion / Critical Discussion

a) Interpretation of the Results

The key finding is that the Responsible-AI architecture has a better overall performance profile than the persistence baseline alone or the standard LSTM model. The value of the architecture is not embedded in one singular numerical benefit to other structures.

Responsible-AI provides a positive outlook along predictive error, peak demand accuracy, calibration of uncertainty, privacy risk, and high-risk event handling.

The lower MAPE of 7.9% for the Responsible-AI ensemble, as compared with 12.8% for persistence, is evidence of the value in combining complementary approaches. Even more relevant to operations is the fall in peak-demand error, from 15.6% to 6.8%. Short peak periods often hold the bulk of the operational importance for residential networks. A model that merely performs decently in aggregate, but grossly underpredicts the peak, may be very misleading for storage dispatch or demandresponse planning.

The other added benefit is prediction-interval coverage. Standard LSTM forecasting attained reasonable point accuracy but a limited 78.4% coverage under the nominal 90% interval. This is indicative of possible overconfidence. Responsible-AI architecture raised coverage to 94.1%.

The somewhat overcoverage as compared to the nominal 90% level is certainly more desirable than systematic undercoverage in high-stakes environments although very wide intervals would ultimately be practically useless. Future studies should therefore assess both coverage and width of the intervals.

b) Comparison With Previous Research

The performance of the LSTM based models agrees with the body of residential forecasting research. For instance, Son and Kim proved that deep learning is relevant for residential electricity forecasting while Wen et al. highlighted that deep-learning models are capable of modeling the nonlinear load dynamics. Kiprijanovska et al. validated the applicability of deep learning to household day ahead forecasting.

The current results come to contrast with traditional accuracy-centered investigations, given that the key comparison is not merely the LSTM architecture against another one but whether the model still can be interpreted, is privacy proportionate, calibrated and steerable.

The uncertainty results are similar to those found in probabilistic forecasting studies by Ben Taieb et al, Shepero et al and Wang et al. They show that uncertainty should be forecast rather than rejected, and this paper builds upon these findings by translating the uncertainty into a workable rule for the operational forecast.

The privacy results are consistent with the existing smart-meter literature. Beckel et al. showed that domestic consumption data is potentially sensitive, while Giaconi et al. noted that smart-meter privacy is an underappreciated challenge. The Responsible-AI framework proceeds on the assumption that energy data should not be considered safe until proven otherwise.

c) Why the Responsible-AI Architecture Performs Differently

The Responsible-AI architecture cannot outperform because the objective function of this architecture considers more than the average forecast error.

In traditional, best-practices RSs, the model is beneficial when the overall numeric difference is reduced between the as-if demand and the observed demand. The Responsible-AI framework also rewards:

- Can improve peak-demand performance
- Higher quality uncertainty calibration
- reduced unnecessary data collection
- Interpretable feature behavior
- Detection of distribution shift
- Selectively escalated to appropriate clinicians to ensure safe and appropriate escalation of high risk predictions

This more generous goal transforms the descriptor of forecast quality.

Logistic-regression analysis by itself, however, shows that adaptive oversight is warranted. Forecasting risk is significantly increased in the model whenever there is more than 5% missingness, when the temperature is at an extreme, or when the volatility of the load is high. These are areas in which learned relationships of the underlying assumptions are not as certain.

Human review has the strongest ability to protect. This should not be taken to mean that human operators are always better than machine-learning systems at predicting electricity demand. Rather, human review provides information and accountability that may not be interpretable by the model.

An operator may know that:

- Communication from the smart meters has been interrupted while the test is in progress
- A community event is causing changes in consumption
- Network maintenance modified local load behavior
- Residential activity is revised in light of extreme-weather advisories
- The charging behavior for electric vehicles has changed
- New installed photovoltaic (PV) or battery capacity which has been commissioned

This contextual knowledge can significantly facilitate the interpretation of a high risk prediction.

Research Contribution / Implications

This research contributes to electricity-demand forecasting studies in terms of shifting the paradigm from the model driven prediction to the governed forecasting system.

The concept of forecasting quality from a theoretical vantage point was defined as multidimensional. Relative to the past, accuracy is still present but it is now part of a wider framework that includes privacy proportionality, explanation, uncertainty, robustness, monitoring, and human accountability.

For the electricity utilities this translates to that the procurement and validation of their AI-forecasting systems should not be based solely on a vendor's claimed MAPE. Utilities should ask the vendor for data on:

- Error in harsh environments
- peak-load performance
- uncertainty calibration
- Function is the explicit equation of the desired input variables
- privacy implications
- model-update frequency
- drift-detection processes
- explanation capabilities
- failure escalation
- audit logging

For community-energy managers, privacy-aware forecasting might enable effective prediction without keeping needlessly detailed facts about every household. Pseudonymized or aggregate data might be able to deliver a significant amount of forecasting value.

Regarding demand-response programs, prediction intervals may facilitate more equitable action. A forecast with tight uncertainty might lend itself to automation, while a forecast with broad uncertainty might have to be confirmed before residents are prompted to alter their electricity use.

This research shows that explainability for AI governance can be achieved without disclosure of trusted source code, with measures such as providing influential variable information, displaying confidence metrics, archiving model versions, changes to models, keeping logs of decisions and making escalation protocols transparent.

The paper also is helpful methodologically by treating the simulation as transparent. No claim is made that the numerical values are representative of any specific resident community. This is important because simulation can be a useful method to develop without trying to pass it off as finished empirical research.

Conclusion

With the decentralization of energy systems and with the increase in the deployment of smart meters, flexible loads, electric vehicles, distributed generation, and storage within residential communities, residential electricity-demand forecasting is gaining increasing significance. AI offers significant analytical benefit because of its ability to encapsulate the nonlinear temporal relationships that traditional prediction methods may fail to accurately model.

We implemented an example simulation-based Responsible-AI architecture with forecasting accuracy, privacy-preserving data minimization, explanation, uncertainty calibration, drift detection, and human escalation, which contributed either to better outcomes or to stronger prediction interval coverage in the synthetic comparative setting. For the synthetic comparison, the Responsible-AI Ensemble had the lowest MAE, RMSE, MAPE, peak demand error and the strongest prediction-interval coverage, the lowest modeled privacy-risk score.

The result also emphasizes the conditional nature of forecast reliability. Missing data, high or low airtemp, and highly variable demand all contribute to modeled risk. Sales forecasting

systems should not view past validation as a predictor of future performance in all circumstances. Uncertainty escalation and drift detection enable the system to identify circumstances in which its output should be more carefully examined.

Human supervision likewise would still have relevance, not in the sense that all hourly predictions would need to be human arbitrated, but so that the simplicity of the approach, the automation in the system needs to be sparing. Routine and low risk predictions can still be automated and the more anomalous and consequential ones can be further examined.

Limitations

The main limitation of this work is that it is entirely based on artificially-generated data. No actual household, smart-meter customer, energy community or electricity supplier was examined. The figures provided are indicative of the internal functioning of a methodical simulation and should not be read as an externally validated estimation of performance.

The synthetic households likewise provide idealized types of electricity use. Actual communities differ widely due to factors such as climate, building efficiency, ownership of appliances, household size, tariff structure, work schedules, solar PV adoption, battery storage, electric vehicles, socioeconomic groups, and behavioral response.

The simulated privacy-risk score is not a validated, all-encompassing privacy score. Actual privacy evaluation should include jurisdiction, time-frame, linkage risk, retention period, third-party access, security of data, consent, and possible inferences.

Furthermore, attribution by SHAP should not be misinterpreted as causal evidence. While measures of feature contribution suggest an influence on the prediction, they do not establish that the variable in question directly caused the electricity demand.

Future Work

Initial future work should test the proposed architecture on real residential smart-meter data obtained within the appropriate legal and ethical framework.

Cross-community validation should include:

- Between various climates
- Low- and high wealth residential and other non-market (parks, sound, etc.) areas

- Households with photovoltaic generation
- battery-storage users
- electric-vehicle households
- Would include (for each particle), the effects of different tariff policies
- Different degrees of participation in demand-response

These trade-offs should also be studied more systematically. Future work might compare raw household-level data, temporal aggregation, spatial aggregation, pseudonymization, federated learning and differential-privacy mechanisms.

Another area needing additional research is in the area of uncertainty calibration. Conformal prediction, quantile regression, Bayesian neural networks, probabilistic recurrent models, and adaptive recalibration may be used to compare models under varying residential-demand situations.

Who? Validation of human-factors should address existence of the residents and energy managers' understanding of the explanations of the model, task benefits from the information conveyed by uncertainties and that too many alerts in the information doing lead to automation fatigue.

Finally, sustained studies should analyze the governance after deployment. Responsible estimation does not end at the point of a validated model. It is expected that the model will be periodically updated, drift will be tracked, incidents will be logged, the rollback process is executed, the audit trail is kept, and external validation still takes place.

Declarations

- **Conflict of Interest:** The sample author in this simulation-based methodology article has no conflict of interest related to the author. Pre-publication conflict of interest reporting should be replaced with the actual author if used.
- **Funding:** No funding is credited to the simulated study for its preparation. The writer must disclose any funding agency, institutional support, grant number, or sponsor involvement in final manuscript.
- **Ethics Statement:** All data used in this study were synthetic and no human participants, identifiable residential consumers, real smart-meter data, confidential electric utility data, clinical data, or personally identifying information was used. As

we did not use human-subject data, we were not required to have formal human-subjects approval for the simulation (a future empirical application with household information should be submitted for ethical and institutional review).

- **Data Availability:** All numerical results in this article are synthesized, written specifically for the purpose of method demonstration. No residential electricity data set from sensitive, confidential, proprietary, or externally obtained information was examined.

References

1. Al-Musaylh, M. S. et al. (2018) Short-term electricity demand forecasting with MARS, SVR and ARIMA models using aggregated demand data in Queensland, Australia. *Advanced Engineering Informatics*, 35, 1–16. Available at: <http://dx.doi.org/10.1016/j.aei.2017.11.002>.
2. Beckel, C., Sadamori, L., Staake, T., & Santini, S. (2014). Ecoding household characteristics from smart meter data. *Energy*, 78, 397–410. <https://doi.org/10.1016/j.energy.2014.10.025>
3. Ben Taieb, S., Huser, R., Hyndman, R. J., & Genton, M. G. (2016). Forecasting uncertainty in the electricity smart meter data using boosting additive quantile regression. *IEEE Transaction On Smart Grid*, 7(5), 2448–255. <https://doi.org/10.1109/TSG.2016.2527820>
4. Bu, S. J., & Cho, S. B. (2020). Time series forecasting with multi-headed attention-based deep learning for residential energy consumption. *Energies*, 13(18), 4722. <https://doi.org/10.3390/en13184722>
5. Eseye, A. T., Lehtonen, M., Tuki, T., Uimonen, S., & Millar, R. J. (2019). Machine learning based integrated feature selection approach for improved electricity demand forecasting in decentralized energy systems. *IEEE Access*, 7, 91463–91475.
6. Giaconi, G., Gündüz, D., & Poor, H. V. (2018). Privacy-aware smart metering: Advances and challenges. *IEEE Signal Processing Magazine*, 35(6), 59–78. <https://doi.org/10.1109/MSP.2018.2841410>

7. High-Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy ai*. European commission. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
8. Imani, M., & Ghassemian, H. (2019). Residential load forecasting using wavelet and collaborative representation transforms. *Applied Energy*, 253, 113505. <https://doi.org/10.1016/j.apenergy.2019.113505>
9. Kiprijanovska, I., Stankoski, S., Ilievski, I., Jovanovski, S., Gams, M., & Gjoreski, H. (2020). HouseEC: Day-ahead household electrical energy consumption forecasting using deep learning. *Energies*, 13(10), 2672. <https://doi.org/10.3390/en13102672>
10. Lundberg, S. M., & Lee, S. I. (2017). A unified framework for interpreting model predictions. In *Advances in neural information processing systems* (pp. 30, 1–9).
11. Nabavi, S. A., Aslani, A., Zaidan, M. A., Zandi, M., Mohammadi, S., & Hossein Motlagh, N. (2020). Machine learning modeling for energy consumption of residential and commercial sectors. *Energies*, 13(19), 5171. <https://doi.org/10.3390/en13195171>
12. OECD. (2019). *Artificial intelligence in society*. OECD Publishing. <https://doi.org/10.1787/eedfee77-en>
13. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). Arxiv preprint: <https://doi.org/10.1145/2939672.2939778>
14. Scheidt, F. vom, Medinová, H., Ludwig, N., Richter, B., Staudt, P., & Weinhardt, C. (2020). Data analytics in the electricity sector A quantitative and qualitative literature review. *Energy and AI*, 1, 100009. <https://doi.org/10.1016/j.egyai.2020.100009>
15. Shepero, M., van der Meer, D., Munkhammar, J. & Widén, J. (2018). Residential probabilistic load forecasting: A Gaussian process based approach targeted at electric load data. *Applied Energy*, 218, 159–172.
16. Son, H., & Kim, C. (2020). A deep learning method for monthly demand prediction of residential sector electricity. *Sustainability*, 12(8), 3103. <https://doi.org/10.3390/su12083103>

17. Source and author: Suganthi, L., & Samuel, A. A. (2012). Energy models for demand forecasting A review. *Renewable and Sustainable Energy Reviews*, 16(2), 1223–1240. <https://doi.org/10.1016/j.rser.2011.08.014>
18. Wang, Y. (2019). Probabilistic individual load forecasting using pinball loss guided LSTM. Wang, Y., Gan, D., Sun, M., Zhang, N., Lu, Z., & Kang, C. (2019). Probabilistic individual load forecasting using pinball loss guided LSTM. *Applied Energy*, 235, 10–20. <https://doi.org/10.1016/j.apenergy.2018.10.078>
19. Wen, L., Zhou, K., & Yang, S. (2020). This research focused on predicting the power load demand of domestic buildings utilizing a deep learning model. It was published in the *Electric Power Systems Research* (pp. 106073). [doi:10.1016/j.epsr.2019.106073](https://doi.org/10.1016/j.epsr.2019.106073).