NEURAL INFORMATION PROCESSING SYSTEMS

Bridging the Temporal Gap: From Historical Monthly Invoices to Granular Hourly Energy Forecasting for Sustainable Operations

Prathamesh Pawar*[1] Alec Hewitt[1] William Schuerman[1] Seyma Gunes[1] Will Sorenson[1]
[1]Amazon Web Services

amazon | science

Bridging the Energy Data Gap

A novel machine learning approach transforming utility invoices into granular hourly energy forecasts for sustainable operations.

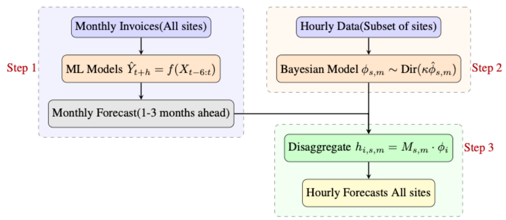

Figure 1. Three-step framework for monthly to hourly energy forecasting.

Our framework transforms monthly utility invoices into accurate hourly energy forecasts. By learning hourly patterns from 50% of 351 industrial sites with hourly meters, we accurately predict hourly consumption for the remaining 50% using only their monthly invoices, achieving 30% better accuracy than baseline methods for 1-3 month ahead predictions. This enables carbon accounting, demand response, and renewable integration - without requiring comprehensive metering infrastructure.

Key Innovations

Our approach solves a fundamental challenge in **energy forecasting: generating high-resolution predictions from limited data.** Key innovations include:

- Novel temporal bridging methodology combining sparse hourly data with widespread monthly invoices
- Bayesian disaggregation framework for going from monthly to hourly resolution
- Adaptive clustering to improve facility-specific load prediction
- Scalable solution requiring only existing utility invoices

Climate Impact

Our methodology enables immediate climate action through granular energy insights, supporting the urgent transition to sustainable operations.

- Helps grid operators **predict stability impacts of new industrial loads** through accurate hourly consumption forecasts
- Provides **real-time carbon accounting** vs. 3-month delayed insights
- **Unlocks precise solar/storage sizing** by forecasting facility's hourly energy patterns without expensive metering

This innovation democratizes access to hourly energy forecasting, making grid-scale decarbonization planning accessible to facilities worldwide using only existing utility data. With demonstrated accuracy, our approach removes a critical barrier to climate action while requiring minimal infrastructure investment.

Methodology

Our framework uses machine learning to turn monthly utility bills into hourly energy forecasts. We do this by learning hourly usage patterns from facilities that already have hourly meters, then applying these patterns to facilities with only monthly bills. Our Bayesian disaggregation approach incorporates day-of-week variations and facility-specific characteristics through cluster-aware template mixing, enabling robust predictions across diverse industrial settings.

1. **Step 1**: Match facilities that have both monthly bills and hourly meters (175 sites) to create a training dataset of monthly-to-hourly relationships
2. **Step 2**: Apply ML models (TabPFN) to forecast 1-3 months ahead of monthly totals for all facilities
3. **Step 3**: Use Bayesian disaggregation to convert these monthly forecasts into hourly predictions by learning patterns from similar facilities

Technical Validation

Our framework involves two key prediction tasks where we evaluated performance and compared different approaches:

Monthly Forecasting (Step 1) We compared different ML models for predicting monthly totals:

Model	Horizon	MAE	RMSE	MdAPE	Model	Horizon	MAE	RMSE	MdAPE
Baseline	1	50	440	8.0	LightGBM[8]	1	69	481	9.2
	2	99	921	13.7		2	69.2	479	9.1
	3	128	1173	18.0		3	68.9	484	9.2
TabPFN2[9]	1	68	524	8.4	AutoGluon[10]	1	73	509	10
	2	65	480	8.3		2	76	607	8.9
	3	65	462	9.7		3	78	615	8.5

Table 1. Monthly Forecasting Performance Comparison. Values show Mean Absolute Error (MAE) in MWh.

TabPFN achieves stable 65-68 MWh error across all horizons, while baseline degrades sharply from 50 to 128 MWh, demonstrating superior long-range forecasting capability.

Hourly Disaggregation (Step 2-3) We compared our method against standard approaches for converting monthly data to hourly predictions:

- **Uniform Distribution**: Spreads monthly total equally across all hours
- **Template Scaling**: Uses average historical patterns but lacks facility-specific adaptation
- **Cluster-Based Grouping**: Groups similar facilities and applies shared patterns
- **Our Bayesian Method**: Automatically learns and adapts to facility-specific patterns

Figure 2 demonstrates performance across three representative facilities. Our Bayesian approach achieves competitive performance with cluster-based methods while substantially outperforming template and uniform baselines. For monthly consumption forecasting 1-3 months ahead, our Enhanced Dirichlet Bayesian method reduces mean absolute error by 28-30% compared to baseline methods across all horizons (Table 2).

Model	Baseline			Forecast			Actual
	H1	H2	H3	H1	H2	H3	Observed
Uniform	178.8	185.7	189.3	174.2	176.7	178.3	157.6
Template-Scaled	152.4	160.6	164.9	159.2	160.1	159.1	145.6
Cluster-Based	132.2	143.6	150.7	131.4	133.5	134.0	112.7
Dirichlet-Bayesian	**130.5**	**142.5**	**150.2**	**129.7**	**132.1**	**133.1**	**109.0**

Table 2: Performance comparison of disaggregation models across forecasting scenarios and horizons. Values shown are MAE.

Results

When converting monthly bills to hourly patterns across 176/351 facilities, our Bayesian hourly disaggregation method combined with ML monthly forecasting achieves the lowest normalized MAE compared to other approaches (28~30% reduction over uniform baseline).

Our method successfully captures daily peaks and valleys, weekday-weekend variations, and facility-specific characteristics that simpler approaches miss. This accurate pattern capture across diverse facility types is crucial for applications like demand response program participation and renewable energy integration planning.

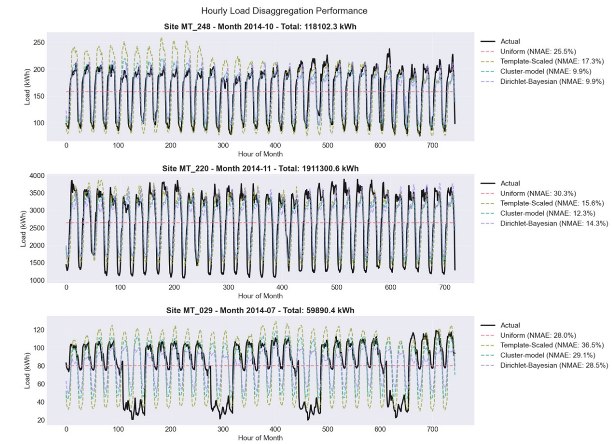

Figure 2. Hourly load disaggregation performance comparing actual consumption (black) against four methods for three representative facilities.

What's next?

Building on this demonstration with Portuguese industrial facilities, here are the next steps:

- Extending the method to handle sub-hourly predictions for more granular grid balancing
- Incorporating external factors like weather patterns and local grid conditions
- Developing automated facility clustering techniques to improve pattern learning
- Deployment in diverse geographic regions with limited metering infrastructure

References

[1] PJM Interconnection, "PJM Manual 11: Energy & Ancillary Services Market Operations, Revision 129," PJM Interconnection, Feb. 22, 2024.
[2] Bu, F., Dehghanpour, K., Yuan, Y., Wang, Z., and Guo, Y. (2021). Disaggregating Customer-Level Behind-the-Meter PV Generation Using Smart Meter Data and Solar Exemplars. IEEE Transactions on Power Systems, 36(6):5417-5427. DOI: 10.1109/TPWRS.2021.3074614.
[3] Trindade, A. (2015). ElectricityLoadDiagrams20112014 [Dataset]. UCI Machine Learning Repository. https://doi.org/10.24432/C58C86.
[4] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). Light-GBM: a highly efficient gradient boosting decision tree. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17), pages 3149–3157, Long Beach, California, USA.
[5] Hollmann, N., Müller, S., Eggensperger, K., and Hutter, F. (2023). TabPFN: A Transformer That Solves Small Tabular Classification Problems in a Second. arXiv preprint arXiv:2207.01848. https://arxiv.org/abs/2207.01848.
[6] Erickson, N., Mueller, J., Shirkov, A., Zhang, H., Larroy, P., Li, M., and Smola, A. (2020). AutoGluon-Tabular: Robust and Accurate AutoML for Structured Data. arXiv preprint arXiv:2003.06505. https://arxiv.org/abs/2003.06505.
[7] Lamagna, M., Nastasi, B., Groppi, D. et al. Hourly energy profile determination technique from monthly energy bills. Build. Simul. 13, 1235–1248 (2020). https://doi.org/10.1007/s12273-020-0698-y.

*Corresponding author: prathavp@amazon.com