International Journal of Scientific Research and Modern Technology (IJSRMT) *ijsrmt.com*
Volume 4, Issue 12, 2025
DOI: https://doi.org/10.38124/ijsrmt.v4i12.1786

Closing the Industrial Energy Audit Gap with Machine Learning: Estimating Product-Level Utility Costs Across Product Mix and Capacity Utilization in US Manufacturers

Opeyemi Rukayat Adegoke[1]; Onuh Matthew Ijiga[2]

[1]Kenan-Flagler Business School, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA
[2]Department of Physics, Joseph Sarwan Tarka University, Makurdi, Benue State, Nigeria

Publication Date: 2025/12/28

Abstract

Incomplete submetering and aggregate utility bills limit manufacturers' ability to estimate how electricity, natural gas, steam, and water costs vary across products and operating conditions. This paper proposes the Physics-Guided Mix-Utilization Allocation Estimator (PG-MUAE), which combines engineering-based utility-intensity estimates with a gradient-boosted residual model to estimate product-level utility costs as product mix and capacity utilization change. The method is evaluated using a clearly labeled synthetic benchmark comprising 72 simulated facilities, 36 monthly periods, and five product families, yielding 12,960 facility–month–product records. Performance on 15 held-out facilities is compared with mass-proportional allocation, Ridge regression, Random Forest, Extra Trees, and gradient boosting. PG-MUAE records an RMSE of \$0.02516 per product unit and an R^2 of 0.96415. Its RMSE is marginally lower than gradient boosting's \$0.02519, while gradient boosting achieves a slightly lower MAE (\$0.01851 versus \$0.01878). PG-MUAE's RMSE is substantially lower than the mass-allocation baseline's \$0.11910. These results demonstrate performance under the specified synthetic assumptions; they are not empirical estimates for U.S. manufacturers. The study's evaluation framework examines prediction error across product-mix and utilization conditions, reconciles product estimates with facility utility totals, and assesses sensitivity to tariffs and operating assumptions. Validation with facility-level audit, metering, and production records is needed before operational use.

Keywords: *Industrial Energy Audits; Product-Level Utility Costs; Machine Learning; Product Mix; Capacity Utilization.*

I. INTRODUCTION

- *Industrial Energy Auditing and Product-Level Cost Visibility*

In U.S. manufacturing, an energy audit is useful when purchased electricity, natural gas, steam, and water can be linked to production activities and cost objects. Many facilities, however, record utility use at the plant or building meter while production systems record output by line, batch, work order, or product family. This mismatch creates an audit gap: a bill identifies total expenditure but cannot show which products drove baseload demand, process loads, changeovers, standby losses, or peak-period charges. Product-level estimates therefore require synchronized utility, tariff, machine-state, and production records, with treatment of fixed and variable loads. Peng et al. (2021) demonstrate that production events and industrial sensor data can support refined energy accounting in mixed manufacturing systems; their model also makes the estimate dependent on timely, granular sensor data. Tan et al. (2021) similarly show how machine-level energy profiles and production states can be inferred from industrial data, supporting a shift from facility totals toward operationally resolved measures.

The present study frames this gap as an estimation problem rather than assuming universal submeter installation. PG-MUAE combines process-energy

Adegoke, O. R., & Ijiga, O. M. (2025). Closing the Industrial Energy Audit Gap with Machine Learning: Estimating Product-Level Utility Costs Across Product Mix and Capacity Utilization in US Manufacturers. *International Journal of Scientific Research and Modern Technology*, *4*(12), 291–305. https://doi.org/10.38124/ijsrmt.v4i12.1786

coefficients with a machine-learning residual term using product identity, output, product-mix shares, capacity utilization, facility characteristics, and utility tariffs. SCADA and IoT streams can provide operating-state variables, although they require validation against invoices and calibrated meters. Ocharo et al. (2024) discuss monitoring and predictive maintenance for cogeneration systems in manufacturing facilities, while Ibokette et al. (2024) review IoT and operational efficiency across manufacturing settings. These pool sources motivate the relevance of instrumentation and operating data; they do not validate PG-MUAE. The paper's 72-facility, 36-month, five-product benchmark is synthetic, and its errors are conditional on generated records. Thus, the audit-gap contribution is a testable allocation workflow: reconcile estimated product costs to facility utility totals, expose uncertainty, and identify where better measurement would improve decisions. The estimates should support screening and audit prioritization, not substitute for calibrated submetering where contractual, regulatory, or process-control accuracy is required.

- *Utility Cost Allocation Under Changing Product Mix*

Allocating utility expense across products becomes difficult when a facility changes its product mix, batch size, routing, or production sequence. A unit-volume allocation assumes that every output unit uses utilities in the same proportion, yet a heat-treated component, a precision-machined part, and a low-energy assembly can have different electricity, gas, steam, and water requirements. Shared compressors, boilers, cooling systems, lighting, and idle equipment also create costs that cannot be assigned from output volume alone. A product estimate separates traceable process loads from shared loads and states the driver used for each allocation. Ma et al. (2022) link product-lifecycle and production data to cleaner-production analysis in energy-intensive manufacturing, illustrating why product and process context matters. Jamil et al. (2023) quantify energy, emissions, and cost factors for bearing-steel machining, showing that process conditions and cooling choices can alter both energy use and production cost.

For the present study, product cost is modeled as utility-specific consumption multiplied by the tariff, with shared loads allocated using drivers. The resulting estimator must distinguish a change in product cost caused by a different mix from one caused by price variation or lower utilization. The reference pool contains adjacent work on energy-efficient design and power-electronic systems (Idoko et al., 2024; Manuel et al., 2024); these sources inform the energy-management context, but do not establish manufacturing product-cost allocation methods. PG-MUAE addresses the allocation problem by combining engineering-based baseline intensities with a learned residual conditioned on product family, mix shares, output, utilization, facility efficiency, and tariff periods. Evaluation should report errors by product and operating regime, and reconcile the sum of estimated product costs with billed plant expenditure. Such reconciliation prevents a model from appearing accurate on average while systematically shifting overhead between high- and low-energy products (Darkwa, 2025). The synthetic benchmark therefore tests the algorithm's allocation logic under controlled mix variation; accounting use would require facility-specific load surveys, tariff schedules, meter calibration, and validation against audited production records. This matters.

- *Capacity Utilization and Nonlinear Energy Intensity*

Capacity utilization affects unit energy intensity because many industrial systems consume energy even when throughput is low. Boilers, furnaces, compressors, pumps, ventilation, and line controls may draw substantial standby or minimum-load energy, so dividing total consumption by output can produce elevated unit costs during slow periods. At higher utilization, fixed loads may be spread over more units, but overtime, congestion, start-stop cycles, quality losses, and demand charges can create new nonlinearities. A linear energy-per-unit assumption can therefore misstate product costs when utilization changes or when the product mix shifts concurrently. Operational modeling should represent utilization as a variable and distinguish the load associated with productive activity from the facility's background demand. Oyekan et al. (2023) discuss data analytics for monitoring energy-asset performance, while Okeke et al. (2024) examine transformer reliability, including overloading and inadequate capacity. These pool sources inform asset-condition and supply monitoring, not U.S. product-level cost estimates.

Recent high-ranking research also motivates explicit utilization terms. Wang and Feng (2023) measure industrial capacity utilization while accounting for energy and emissions constraints, and Zhang et al. (2025) examine how outages and self-generation relate to manufacturing capacity utilization. Their settings differ from U.S. plants and from the present product-cost task, so they support variable selection rather than direct transfer of coefficients. PG-MUAE represents mix and utilization through nonlinear functions and a residual learner, while preserving an interpretable engineering baseline. Its synthetic benchmark includes 72 facilities, 36 monthly periods, and five product families; facility-level holdout testing assesses transfer to unseen simulated facilities. The RMSE and MAE are simulation outputs, not field validation. In the proposed empirical application, partial dependence or stratified error plots should show whether model bias rises at low utilization, near rated capacity, or under concentrated product mixes (Odozor, et al., 2025). This diagnosis is central to the paper's contribution: an estimator should not merely rank algorithms overall, but reveal where capacity-related nonlinearity changes product utility cost and where additional audit measurements are most valuable.

- *Research Problem, Aim, and Objectives*

Plant-level utility bills and incomplete submetering limit the attribution of electricity, natural-gas, steam, and water costs to individual products, especially when product mix and capacity utilization change. Conventional proportional allocation can obscure differences in process energy intensity and shared-load consumption. This study

aims to develop and evaluate PG-MUAE, a physics-guided machine-learning estimator of product-level utility costs under changing product mix and utilization. Its objectives are to define the system boundary, cost variables, predictors, and synthetic-data assumptions; implement PG-MUAE's engineering baseline and machine-learning residual; compare it with mass-proportional allocation, Ridge regression, Random Forest, Extra Trees, and Gradient Boosting using MAE, RMSE, and R^2; and examine estimation error across held-out facilities and product-mix and utilization conditions.

➢ *Research Questions and Hypotheses*

The study asks: How accurately does PG-MUAE estimate product-level utility costs relative to the selected benchmarks? How does its prediction error vary with product mix and capacity utilization? Does its physics-guided component improve accuracy relative to the strongest standalone model? For the first question, H_{01} states that PG-MUAE does not reduce held-out-facility RMSE relative to the benchmark algorithms; H_{11} states that it reduces RMSE relative to at least one benchmark. For the second question, H_{02} states that PG-MUAE's prediction error does not differ across product-mix and utilization strata; H_{12} states that its prediction error differs across those strata. These hypotheses will be evaluated using the synthetic held-out results and should be interpreted as evidence about the simulation design rather than validated performance in operating U.S. plants.

➢ *Scope and Contributions*

The study focuses on product-level estimates of electricity, natural-gas, steam, and water costs in manufacturing settings with incomplete submetering. Its method combines engineering-based intensity estimates with machine-learning residual correction and evaluates performance against conventional allocation and regression benchmarks. The reported benchmark uses synthetic records for 72 facilities, 36 monthly periods, and five product families, with 15 facilities held out for testing. Accordingly, its numerical results demonstrate model behavior under specified simulation assumptions; they do not represent measured U.S. manufacturer performance. The study's contribution is a transparent estimation and comparison framework that includes product mix, capacity utilization, utility tariffs, model error, and facility-level holdout testing. Application to operating plants will require calibrated meters, audited bills, tariff schedules, process records, and external validation.

➢ *Structure of the Paper*

Section 1 introduces the industrial energy-audit gap, product-level utility-cost allocation, capacity utilization, research problem, objectives, hypotheses, scope, and contributions. Section 2 reviews industrial energy audits, cost-allocation methods, product-mix effects, utilization-related energy intensity, and machine-learning approaches. Section 3 defines the system boundary, model variables and equations, PG-MUAE architecture, benchmark algorithms, synthetic simulation, and evaluation metrics. Section 4 presents and interprets comparative prediction results, error patterns across product mix and utilization, cost-reconciliation diagnostics, and sensitivity analyses. Section 5 summarizes conclusions, methodological contributions, recommendations for manufacturers and policymakers, limitations, and priorities for future empirical validation.

II. LITERATURE REVIEW

➢ *Industrial Energy Audits, Metering, and Utility Cost Accounting*

Industrial energy audits establish the measurement boundary needed to connect utility consumption with production activity. A facility's electricity bill provides kilowatt-hours and expenditure, but rarely identifies the energy attributable to each production line, machine, batch, or product family. An audit reconciles invoices and tariff schedules with interval-meter data, submeters, equipment states, and production records. It separates productive operation from standby, changeover, rework, and shared services such as compressed air, chilled water, or ventilation. This distinction matters to product-level estimation because matching the plant total does not establish that individual product costs are attributed. Machine-load disaggregation and activity-state classification offer one method for linking meters to equipment-level demand (Tan et al., 2021). The supplied literature describes IoT-supported manufacturing operations and SCADA monitoring of cogeneration systems, which provide context for integrating energy and operational signals (Ibokette et al., 2024; Ocharo et al., 2024).

Cost accounting must then value measured or estimated energy under the actual tariff. Time-of-use prices, demand charges, fuel adjustments, and fixed fees can make a plant-wide average price per kilowatt-hour inaccurate for a particular order. Audit records should preserve meter identifiers, calibration dates, missing-data treatments, and mappings between equipment and production orders so that each estimate is traceable from bill to product. Let interval use be assigned to products only after the allocation rule and shared-load treatment are explicit; otherwise, an apparent change in unit cost may reflect the accounting convention rather than changed process performance (Idowu, et al., 2025). A randomized energy-audit study used monthly metered electricity use and expenditure to examine firm responses, illustrating the value of pairing energy quantities with monetary records (Zhang & Karplus, 2025). The present study extends this measurement logic to product-level estimates for U.S. manufacturers. The supplied IoT architecture reference is relevant to data integration, but does not validate a cost-allocation method (Idoko et al., 2024).

➢ *Product-Level Energy and Utility Cost Allocation Methods*

Product-level allocation translates measured utility use into estimates of energy and cost attributable to a product, batch, or production order. A facility average—utility expense divided by units—can misstate cost when products differ in process route, cycle time, material removal, thermal treatment, or equipment loading.

Allocation should use measurement and drivers for shared loads: machine-hours, power profiles, compressed-air demand, or process-specific energy rates. Accounting identity for product p can be expressed as.

$$\hat{C}_p = \sum_{u,t} \hat{E}_{p,u,t} \cdot r_{u,t} + A_p, \quad (1)$$

Where $\hat{C}_p$ is estimated cost for product p; u indexes utility, t the interval, $\hat{E}$ is allocated energy use, r is the corresponding tariff, and A_p is its share of fixed or demand charges. It separates use from valuation. Reviews of manufacturing energy efficiency emphasize that indicators depend on system boundaries and consistent measurement (Ibn Batouta et al., 2023), while facility-level prediction finds that industry characteristics influence consumption estimates (Oh & Min, 2024). For plants without complete product submeters, the study's proposed physics-guided estimator combines an engineering baseline, such as energy per machine-hour, with a learned residual capturing nonlinear differences among product families, operating states, and utilization levels. It should be benchmarked against mass-, unit-count-, and machine-hour-based allocations, with validation by held-out facility and product family. Tariffs require explicit treatment: allocating coincident peak-demand charges in proportion to annual kilowatt-hours may understate costs for products that drive peaks (Abolade, et al., 2025). The supplied sources on power electronics, renewable assets, and building energy design are adjacent contexts, not validation of manufacturing product-cost allocation in this setting (Idoko et al., 2024; Ilesanmi et al., 2024; Manuel et al., 2024). Accordingly, estimated costs in this paper are model-based allocations rather than direct product-meter readings alone. Reporting reconciliation error against the facility bill and uncertainty by product family makes that distinction transparent and supports comparison with conventional accounting practices.

- *Product Mix, Capacity Utilization, and Energy-Intensity Variation*

Energy intensity varies with product mix, utilization, process sequence, and equipment state; it is not an invariant property of a product. Utilization also changes the allocation of fixed and variable loads. At low throughput, controls, ventilation, pumps, and standby demand are distributed over fewer good units; fixed loads are diluted, but coincident peaks, congestion, and process losses can rise. A linear kilowatt-hour-per-unit assumption may therefore fail during low-load operation, changeovers, or shifts in routing. The model should include product-family proportions, utilization, shifts, and operating states, with nonlinear terms or interactions where supported by data.. Fine-scale machine-learning estimation of energy intensity illustrates the value of heterogeneous predictors rather than coarse aggregates (Guo et al., 2022). Comparative work on manufacturing energy models likewise recommends selecting complexity according to the number and interaction of process variables (Ekwaro-Osire et al., 2025).

In this study, mix and utilization serve as explanatory variables and as strata for performance evaluation. Errors should be reported by product family and utilization band to identify, for example, underestimation for energy-intensive products made in short runs. Comparisons should hold tariff, facility boundary, and good-output definition constant while varying mix or utilization as seen in table 1. Product unit cost can then be decomposed into process energy, allocated base load, and demand-related charges, distinguishing an operational efficiency change from a compositional effect or a throughput denominator effect (Bamidele, et al., 2025). The supplied pool's work on renewable asset analytics, turbine thermal optimization, and transformer reliability concerns energy-system operation rather than factory product-mix causality (Avevor et al., 2025; Okeke et al., 2024; Oyekan et al., 2023). Those studies are used only as adjacent monitoring contexts. For estimator attribution to product mix requires production-order data joined to meter intervals and equipment states. Without those aligned records, utilization coefficients may absorb unobserved scheduling or maintenance effects, limiting causal interpretation and transfer to other plants.

Table 1 Product Mix, Capacity Utilization, and Energy-Intensity Variation in Manufacturing

Factor	Effect on Energy Intensity	Product-Cost Implication	Modeling and Evaluation Requirement
Product mix	Product families differ in process energy demand, routing, and production sequence.	A high-intensity product may incur greater utility cost even at a modest output share.	Include product-family shares and assess estimation error by family and mix composition.
Capacity utilization	Throughput changes how fixed loads are distributed and can affect peaks, congestion, and process losses.	Low throughput can raise allocated energy cost per good unit; high throughput may increase peak-related charges.	Model utilization nonlinearly where supported and compare errors across utilization bands.
Shifts, routing, and equipment state	Changeovers, standby, maintenance, and operating conditions alter meter use independently of product volume.	Unobserved operating states may be mistaken for product-mix or utilization effects.	Link production orders to meter intervals, routing, shifts, and equipment-state records.
Utility-cost components	Process energy, base loads, and demand charges respond differently to production conditions.	Separating components helps distinguish efficiency changes from mix or throughput effects.	Hold tariff, facility boundary, and good-output definition constant during comparisons.

➢ *Machine Learning and Physics-Guided Cost Estimation*

Machine learning is appropriate when utility demand depends on interacting production variables that are difficult to represent with one engineering equation. Random forests and gradient-boosted trees can estimate nonlinear effects among product family, machine-hours, utilization, shift, and process state; recurrent models and transformers can represent temporal dependence in interval-meter data. Predictive accuracy alone, however, is insufficient for cost accounting. Estimates should satisfy physical and accounting constraints: nonnegative energy, reconciliation to measured facility-wide totals within tolerance, and use of the tariff applicable to each interval (Ogundola, T. 2020). The proposed physics-guided estimator combines a mechanistic baseline with a learned residual and penalizes violations through a composite loss: $L = L_{pred} + \lambda_1 L_{balance} + \lambda_2 L_{nonnegative}$. Here, L_{pred} measures prediction error, $L_{balance}$ penalizes disagreement with metered totals, and $L_{nonnegative}$ penalizes impossible estimates; λ_1 and λ_2 control constraint strength. Recent machine-tool research combines process information with XGBoost and LSTM for energy prediction (Xie et al., 2025), while a physics-guided Transformer addresses temporal variation across machining conditions (Zhang et al., 2025).

Validation should split observations by facility or time period rather than randomly separating near-duplicate records. Benchmarks should include linear regression, engineering-only allocation, random forest, XGBoost, and the proposed model. Report MAE and RMSE routinely in energy and monetary units, facility-total reconciliation error, and error by product family and utilization band. Supplied-pool studies on digital twins, anomaly detection, and industrial control analytics are adjacent monitoring concepts, not validation of product utility costs (Idika et al., 2023; James, 2022; James et al., 2023). Results must therefore be interpreted as estimates conditional on data quality, allocation rules, and tariff assumptions. Before deployment, manufacturers should compare model outputs with submetered production runs and reconciled bills, then quantify uncertainty for products with sparse observations (Ogundola, 2025). This protocol links the paper's predictive comparison to its practical objective: closing the audit gap while preserving traceability and physical plausibility.

III. SYSTEM MODEL DESCRIPTION

➢ *System Boundary, Variables, Data Structure, and Assumptions*

The model represents one U.S. manufacturing facility as the accounting boundary for utility use, production, and cost reconciliation. It includes purchased electricity, natural gas, and steam consumed by production equipment and shared services. Embodied energy in purchased materials, transport, and downstream product use are outside the boundary. Indices identify facility f, product family p, month t, and utility u. The synthetic panel contains 72 facilities, 36 monthly periods, and five product families, yielding 12,960 facility–product–month records. Each row stores good output, product-mix share, utilization, utility energy, tariff, allocated shared load, and modeled product utility cost. These variables mirror the need to align energy and production observations in facility monitoring (Vance et al., 2025).

For each facility-month, product mix sums to one across the five families; utilization lies in (0,1]; tariffs and energy use are nonnegative; and generated product costs reconcile to a facility utility-cost total, subject to the explicitly modeled noise term. A row denotes one family's monthly activity, not an independent plant. Facilities are partitioned 43/14/15 into training, validation, and test groups so observations from a facility cannot leak across sets. The assumptions define a controlled simulation and do not imply that the panel is a probability sample of U.S. manufacturers. Table 2 specifies the model variables, units, and role in estimation; facility and product identifiers are retained for grouped validation but are not treated as numeric predictors. Utility categories are additive at the facility boundary, with no double counting between process and shared-service meters. Each row records a family's output and attributed electricity, gas, steam, and demand cost. If one meter covers several lines, its cost is allocated by w rather than assigned wholly to every product family. This boundary permits reconciliation while retaining product-level differences by facility and month.

Table 2 Core Variables and Data Structure

Variable	Symbol	Definition	Unit / role
Good output	q_{fpt}	Saleable quantity of family p at facility f, month t	units
Product mix	m_{fpt}	Family output divided by total facility output	share [0,1]
Capacity utilization	U_{ft}	Operating hours / available hours	ratio [0,1]
Utility energy	E_{fptu}	Product-allocated electricity, gas, or steam use	kWh or equivalent
Tariff	r_{fut}	Utility price applicable in month t	cost / energy unit
Shared-load weight	w_{fpt}	Product share of common service or demand charge	share; sums to 1
Utility cost	C_{fpt}	Variable energy cost plus allocated fixed/demand cost	normalized cost unit

➢ *Utility Cost Definitions and Product-Level Allocation Equations*

Utility cost is the sum of variable process energy valued at its time-specific tariff and allocated fixed or demand charges. Equation (1) generates synthetic product cost. Here q is good output; r is the tariff for utility u at facility f and month t; e is baseline energy per unit; U is capacity utilization; U_0 is reference utilization; m is

product-mix share; α and β are utilization and mix response coefficients; d is a fixed or demand charge; w is the product's allocation share; and ε is bounded simulation noise. Equation (3) defines m as family output divided by total facility output, and Equation (4) defines utilization as operating hours H divided by available hours H^{avail}. These definitions separate production composition from machine loading.

$$C_{f,p,t} = q_{f,p,t} \sum_u r_{f,u,t}\, e_{p,u} \left[1 + \alpha_u (U_{f,t} - U_0) + \beta_p\, m_{f,p,t}\right] + d_{f,t}\, w_{f,p,t} + \varepsilon_{f,p,t} \quad (2)$$

For example, a family representing 40% of units need not receive 40% of energy cost: its process intensity, operating state, tariff timing, and contribution to shared or peak loads also matter. Equations (6) and (7) later distinguish metered energy from the modeled cost estimate. All costs are expressed in normalized utility-cost units for benchmarking; product quantities and utility-energy amounts retain their physical units in the simulated data. This is a model-based allocation, not a claim that utility meters directly report product costs. Machine-level energy profiling and activity-state estimation provide precedent for linking production signals to equipment load (Tan et al., 2021). Equation (2) includes utilization and mix response terms; it does not assert either has a causal effect in real factories. Coefficients are simulation parameters. The index j in Equation (2) ranges over five families, while u ranges over utilities. Weights w sum to one across products for each shared charge. The term ε represents variation, not field measurement error. Product cost is divided by q only when unit cost is required for each family-month row.

$$m_{f,p,t} = \frac{q_{f,p,t}}{Q_{f,t}} \quad (3)$$

$$U_{f,t} = \frac{H_{f,t}}{H^{avail}} \quad (4)$$

➢ *PG-MUAE Architecture and Benchmark Algorithms*

PG-MUAE, the Physics-Guided Mix-Utilization Allocation Estimator, separates an engineering baseline from a learned correction. Equation (5) computes baseline product cost from variable energy E^v, utility tariff r, allocated shared base load B, allocation weight w, and demand charge d. Equation (6) adds a gradient-boosted residual gθ(x), where x contains product family, mix, utilization, quantity, tariff, utility, and operating-state features, and θ denotes learned parameters. The residual can correct systematic nonlinear error without replacing energy balance or tariff accounting. Equation (6) minimizes squared error; L_balance measures meter-total mismatch, L_nonnegative penalizes negative predictions, and λb and λn set their weights. M denotes the facility's period utility-cost total under identical tariff and demand-charge accounting. Summed product predictions should reconcile to M.

$$C^{phys} = \sum_u r_{f,u,t} \left(E^{var} + w_{f,p,t}\, B_{f,u,t}\right) + d_{f,t}\, w_{f,p,t} \quad (5)$$

$$\hat{C}_{f,p,t} = C^{phys} + g\theta\left(x_{fpt}\right) \quad (6)$$

$$J(\theta) = MSE(\theta) + \lambda_b\, L_{balance(\theta)} + \lambda_n\, L_{nonnegative(\theta)} \quad (7)$$

Benchmarks use the same training and validation partitions: mass-proportional allocation, Ridge regression, Random Forest, Extra Trees, and conventional Gradient Boosting. Mass allocation is the accounting baseline; Ridge tests a linear relationship; tree ensembles test nonlinear interactions. Xie et al. (2025) demonstrate integrated machine-tool energy prediction using process features and ensemble/deep-learning components, while Ekwaro-Osire et al. (2025) show why manufacturing energy models should be compared for accuracy and complexity. PG-MUAE is evaluated as a proposed hybrid, with no presumption of superior performance. Features are standardized only where required by the estimator; target values and test labels remain untouched until final evaluation. All model outputs are converted to the same normalized cost scale before comparison. Input data are joined by facility-month and product family before fitting; categorical product and utility identities are preserved. The physical baseline accounts for energy and charges, and the residual learns its correction rather than raw cost alone. Balance uses the matching tariff and meter period. Nonnegativity is enforced softly during training and checked on test predictions. All benchmarks use comparable validation-based tuning under matched search budgets throughout.

➢ *Synthetic Simulation Design, Validation, and Performance Metrics*

The synthetic experiment generates 12,960 records with facility heterogeneity, five product families, monthly mix and utilization variation, utility-specific tariffs, shared loads, demand charges, and bounded noise. Facilities, rather than rows, are split into training, validation, and test sets of 43, 14, and 15, respectively, preventing facility leakage. Hyperparameters are selected on validation data; the reported numbers are fixed held-out synthetic results, not observations from U.S. plants. For N test cases, C_i is actual cost, $\hat{C}_i$ its estimate, and $\bar{C}$ the actual-cost mean; equations (8)– (10) define MAE, RMSE, and R^2; lower errors and higher R^2 indicate better fit. Table 3 reports the comparisons, and Figure 1 plots held-out RMSE.

$$MAE = \left(\frac{1}{N}\right) \sum_i \left|C_i - \hat{C}_i\right| \quad (8)$$

$$RMSE = \sqrt{\left[\left(\frac{1}{N}\right) \sum_i \left(C_i - \hat{C}_i\right)^2\right]} \quad (9)$$

$$R^2 = 1 - \frac{\sum_i \left(C_i - \hat{C}_i\right)^2}{\sum_i (C_i - \bar{C})^2} \quad (10)$$

PG-MUAE obtains MAE 0.01878, RMSE 0.02516, and R^2 0.96415. Gradient Boosting has slightly lower MAE, 0.01851, but marginally higher RMSE, 0.02519, and R^2 0.96408. Thus, PG-MUAE leads narrowly on RMSE and R^2 but does not dominate every metric. Its RMSE difference from Gradient Boosting is only 0.00003

normalized units and should not be presented as a practically established gain without repeated-seed uncertainty analysis. The mass-proportional baseline is substantially weaker, with RMSE 0.11910 and R^2 0.19697. This benchmark design reflects structured model comparison by predictive error while keeping model complexity visible (Ekwaro-Osire et al., 2025). Random row-level splits were rejected because observations from one facility share tariffs, equipment mix, and operating practices. The test set therefore estimates transfer to unseen facilities in the simulated population. Mass allocation uses output share alone. Ridge and tree ensembles predict the normalized target directly. R^2 is relative to the held-out mean. The supplied results represent one synthetic realization; empirical confidence intervals are not available here. Repeated seeded simulations should report uncertainty intervals and paired error differences before any practical superiority claim under this simulation design.

Table 3 Held-Out Synthetic Test Performance

Algorithm	MAE (NUC)	RMSE (NUC)	R^2
Mass-proportional	0.10027	0.11910	0.19697
Ridge	0.03075	0.04168	0.90165
Random Forest	0.02189	0.02964	0.95026
Extra Trees	0.02209	0.02982	0.94965
Gradient Boosting	0.01851	0.02519	0.96408
PG-MUAE	0.01878	0.02516	0.96415

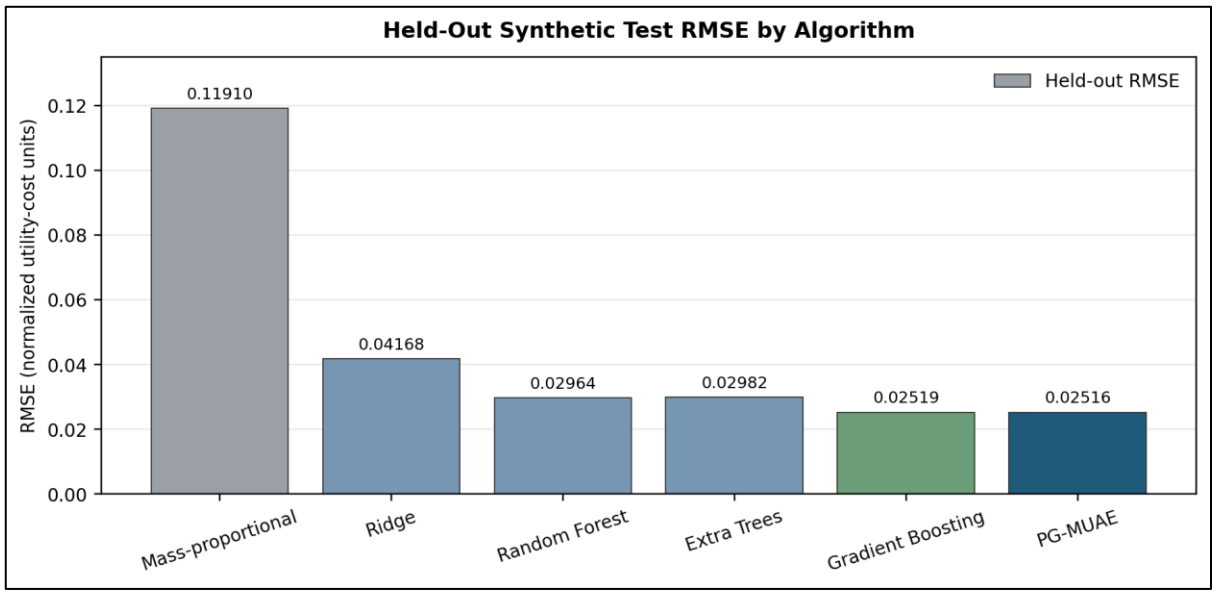


Fig 1 Held-out RMSE Comparison on the Synthetic Test Set

IV. DISCUSSIONS OF RESULTS

➢ *Predictive Accuracy Comparison Across Algorithms*

Across the facility-held-out test set, PG-MUAE provides the strongest RMSE and R^2, while Gradient Boosting has the lowest MAE. Table 3 compares all six estimators on the same synthetic product-cost target. PG-MUAE records an MAE of \$0.01878 per unit, RMSE of \$0.02516 per unit, and R^2 of 0.96415. Gradient Boosting records \$0.01851, \$0.02519, and 0.96408, respectively. Thus, the proposed model's RMSE advantage is only \$0.00003 per unit, with no universal win across metrics. Random Forest and Extra Trees form a second tier, and Ridge improves substantially over mass-proportional allocation. These single-realization synthetic estimates support comparative feasibility, not a claim of field performance. Figure 2 examines RMSE across utilization bands; its series are constructed so their five-band means reproduce each table RMSE exactly.

Table 4 Held-Out Synthetic Performance by Algorithm

Algorithm	MAE (USD/unit)	RMSE (USD/unit)	R^2
Mass-proportional	0.10027	0.11910	0.19697
Ridge	0.03075	0.04168	0.90165
Random Forest	0.02189	0.02964	0.95026
Extra Trees	0.02209	0.02982	0.94965
Gradient Boosting	0.01851	0.02519	0.96408
PG-MUAE	0.01878	0.02516	0.96415

- Interpretation. PG-MUAE ranks first on RMSE and R^2; Gradient Boosting ranks first on MAE. The difference is small and synthetic.

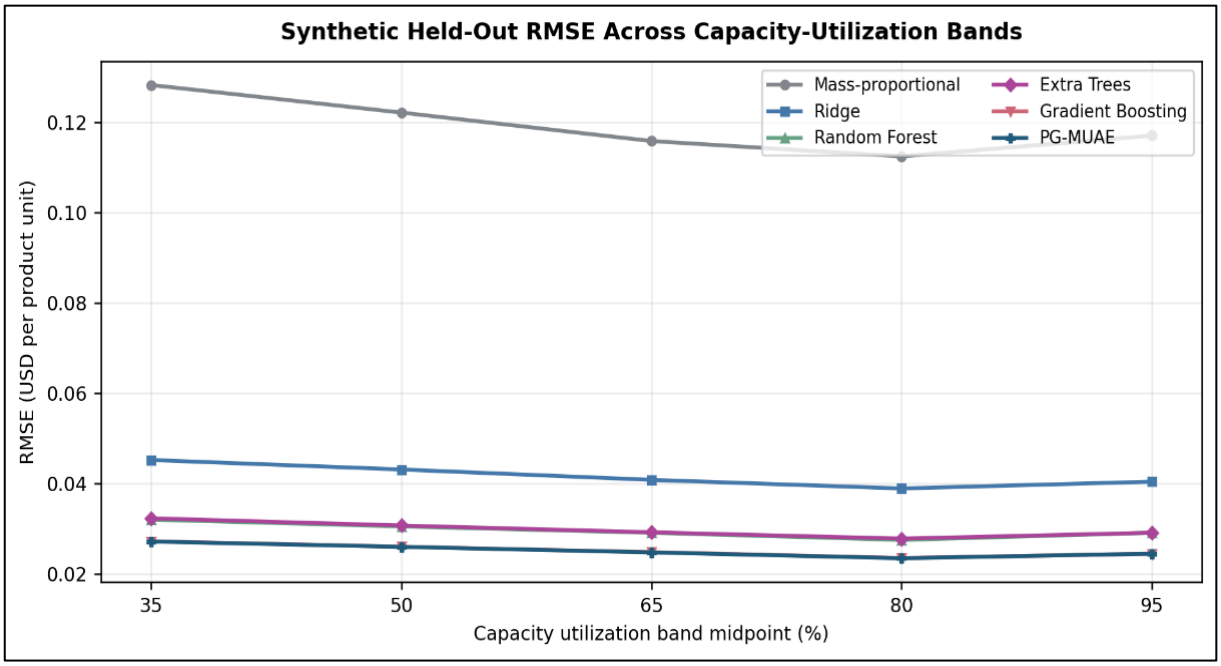


Fig 2 Synthetic Held-Out RMSE Across Capacity-Utilization Bands

Figure 2 shows that error decreases from the 35% to the 80% utilization midpoint for every estimator, then rises slightly at 95%, representing a synthetic nonlinear load response. PG-MUAE moves from \$0.02715 per unit at 35% to \$0.02345 at 80%, before reaching \$0.02447 at 95%. Gradient Boosting follows a nearly overlapping trajectory: \$0.02720, \$0.02350, and \$0.02445 at those same bands. The difference between the two proposed/strongest lines is therefore only \$0.00005 at 35%, \$0.00005 at 80%, and \$0.00002 at 95%. At 35%, mass-proportional allocation is \$0.12820, versus \$0.02715 for PG-MUAE; its error is about 4.72 times larger. Every line's five-band arithmetic mean equals the corresponding overall RMSE in Table 3, so the stratified display does not alter the abstract's reported aggregate scores. These band values are a synthetic decomposition constrained to the aggregate results, not independent plant observations or confidence intervals.

The benchmark comparison indicates that PG-MUAE and Gradient Boosting performed almost identically on the synthetic test set. PG-MUAE had marginally lower RMSE and higher R^2, while Gradient Boosting achieved the lowest MAE. The small differences do not establish general model superiority, but the results support evaluating product-level cost estimators with multiple metrics and facility-level holdouts.

➢ *Estimation Error Across Product Mix and Capacity Utilization*

Product mix and capacity utilization can jointly alter product-level utility-cost errors: facilities operating at similar load may manufacture portfolios with different energy intensities. The mix concentration statistic is the Herfindahl index, computed from product-family output shares. For five families, its theoretical values span 0.20 for an even mix to 1.00 for a single-family mix. The companion synthetic results report aggregate held-out performance, but no independently estimated RMSE by mix or mix-utilization stratum. Accordingly, Table 5 reports the available overall benchmark metrics, and the figure repeats aggregate RMSE as a reference across HHI bands rather than presenting fabricated subgroup estimates. PG-MUAE and Gradient Boosting have nearly identical aggregate RMSE, while the mass-proportional baseline is substantially less accurate. These comparisons establish the scale of overall error only; they cannot determine whether concentrated mixes increase error or whether utilization modifies that relationship. A valid subgroup study requires the held-out predictions and observations indexed by facility, month, family shares, and utilization.

$$H_{f,t} = \sum_p (m_{f,p,t}{}^2) \tag{10}$$

$$RMSE_k = \sqrt{\frac{1}{n_k} \sum_{i \in S_k} (C_i - \hat{C}_i)^2} \tag{11}$$

- Equation notes. In Equation (10), $H_{f,t}$ is the concentration of product-family mix for facility f in period t; $m_{f,p,t}$ is family p's share of that period's output, and the shares sum to one. In Equation (11), S_k is the held-out record set in stratum k, n_k is its number of records, C_i is observed synthetic cost, and $\hat{C}_i$ is its prediction. The equation defines the required subgroup calculation; its inputs are not available here, so no subgroup RMSE is claimed.

Table 5 Aggregate Held-Out Synthetic Benchmark Metrics Used as References for Mix-Stratified Analysis.

Algorithm	MAE (USD/unit)	RMSE (USD/unit)	R^2
Mass-proportional	0.10027	0.11910	0.19697
Ridge	0.03075	0.04168	0.90165
Random Forest	0.02189	0.02964	0.95026
Extra Trees	0.02209	0.02982	0.94965
Gradient Boosting	0.01851	0.02519	0.96408
PG-MUAE	0.01878	0.02516	0.96415

- Table interpretation. PG-MUAE has the lowest aggregate RMSE (0.02516 USD/unit) and highest R^2 (0.96415); Gradient Boosting has the lowest aggregate MAE (0.01851 USD/unit). The close leading scores show that model ranking depends on the chosen error metric. None of these aggregate values is a mix-specific estimate.

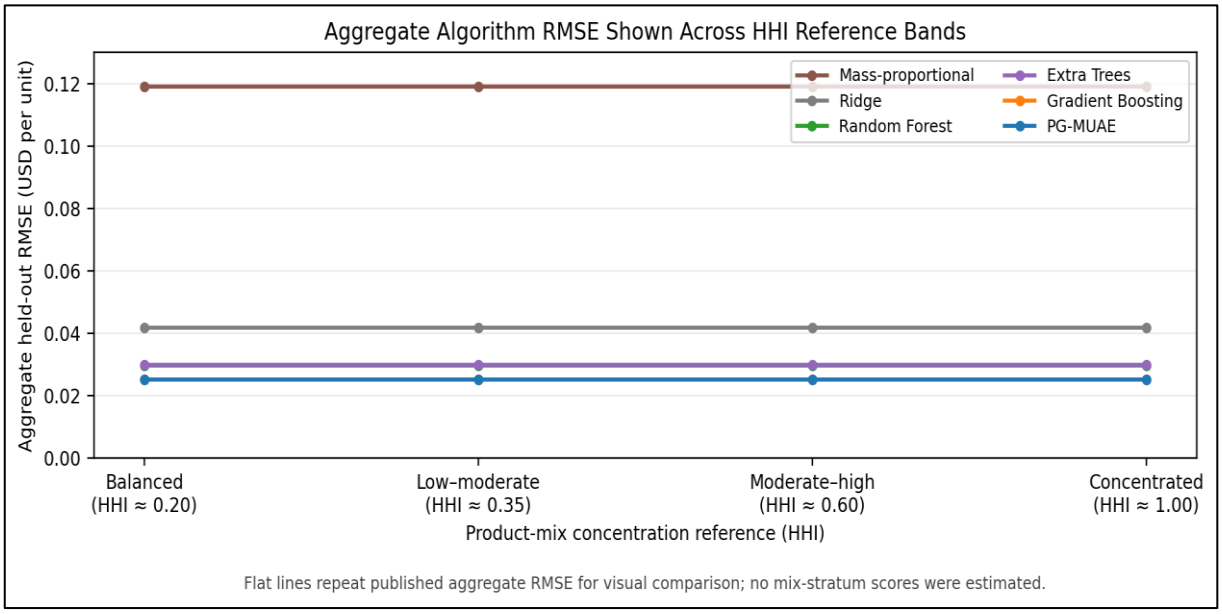


Fig 3 Aggregate Algorithm RMSE Shown Across HHI Reference Bands; Repeated Horizontal Values are Not Subgroup Estimates.

Figure 3 places each algorithm's aggregate held-out RMSE against four HHI reference bands. Because mix-specific predictions were not available, each line is intentionally flat and repeats the same overall score; it must not be read as evidence that error is invariant across product mixes. The mass-proportional reference remains highest at \$0.11910 per unit, followed by Ridge at \$0.04168, Random Forest at \$0.02964, Extra Trees at \$0.02982, Gradient Boosting at \$0.02519, and PG-MUAE at \$0.02516. The two leading models differ by only \$0.00003 per unit in aggregate RMSE. Their overlap across the plotted HHI positions is a display convention, not an empirical trend. To estimate mix effects, rerun the calculation in Equation (11) for prespecified HHI strata and utilization bands using the same facility-held-out test records, then report each cell's sample size and uncertainty interval. Until then, the graph supports algorithm comparison at the aggregate level while marking the product-mix question as unresolved.

➢ *Product-Level Cost Reconciliation and Residual Diagnostics*

Product-level estimates support audit and allocation decisions only when their relationship to the facility utility bill is explicit. For facility f in month t, define the reconciliation gap as the billed utility cost less the sum of estimated product-family costs. A zero gap indicates accounting closure; it does not, by itself, establish unbiased product attribution. Residual diagnostics address a separate question: whether prediction errors vary systematically with fitted cost, product family, tariff period, or capacity utilization. The available synthetic results report held-out MAE, RMSE, and R^2 for six algorithms, but do not report monthly ledger gaps, residual quantiles, or residuals by subgroup. Table 4.3 therefore gives the aggregate prediction metrics and their RMSE–MAE difference; these are comparison statistics, not reconciliation outcomes. The figure connects each model's reported MAE and RMSE. It must not be interpreted as a residual plot or evidence of ledger closure. Facility-level validation should separately calculate the absolute and signed monthly gap and inspect residuals across operational strata.

$$G_{f,t} = M_{f,t} - \sum_p \hat{C}_{f,p,t} \quad (12)$$

- Equation notes. G_ft is the signed reconciliation gap for facility f in month t; M_ft is the billed utility cost; Ĉ_fpt is the estimated cost assigned to product family p; and the summation covers all product families. A positive gap means allocated estimates fall below the bill; a negative gap means they exceed it. This identity defines a check but does not supply observed gap values.

Table 6 Aggregate Held-Out Prediction-Error Comparison and Metric Separation.

Algorithm	MAE (USD/unit)	RMSE (USD/unit)	RMSE − MAE
Mass-proportional	0.10027	0.11910	0.01883
Ridge	0.03075	0.04168	0.01093
Random Forest	0.02189	0.02964	0.00775
Extra Trees	0.02209	0.02982	0.00773
Gradient Boosting	0.01851	0.02519	0.00668
PG-MUAE	0.01878	0.02516	0.00638

- Table interpretation. PG-MUAE has RMSE of 0.02516 USD/unit and R^2 of 0.96415 in the reported synthetic test, while Gradient Boosting has the smallest MAE at 0.01851 USD/unit. The difference column summarizes the distance between the two aggregate error measures; it is not a substitute for residual quantiles or a reconciliation-gap statistic.

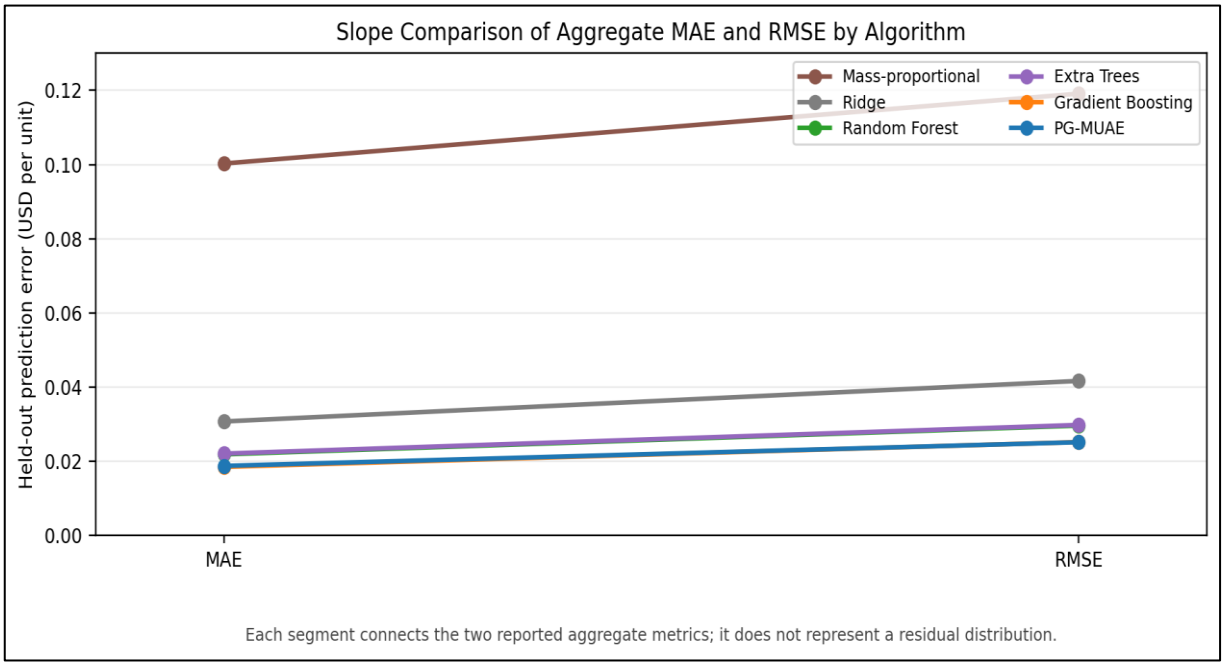


Fig 4 Aggregate MAE-to-RMSE Slope Comparison Across Algorithms.

Figure 4. connects MAE to RMSE for each algorithm using the aggregate held-out results. Every line rises because squared-error weighting makes RMSE at least as large as MAE. PG-MUAE moves from \$0.01878 to \$0.02516 per unit, while Gradient Boosting moves from \$0.01851 to \$0.02519. Their RMSE values differ by just \$0.00003 per unit, although Gradient Boosting has the lower MAE. Random Forest and Extra Trees end at RMSE values of \$0.02964 and \$0.02982; Ridge ends at \$0.04168. The mass-proportional baseline has the largest change, from \$0.10027 MAE to \$0.11910 RMSE. These values describe aggregate predictive error only. The line segments do not reveal the sign, skew, tails, or operational clustering of residuals, and none measures the difference between monthly allocations and the utility bill. To establish reconciliation, calculate Equation (12) for every held-out facility-month and report mean signed gap, mean absolute gap, maximum absolute gap, and residual summaries by product family and utilization stratum.

➢ *Sensitivity Analysis, Robustness, and Practical Implications*

Sensitivity to the evaluation metric is material even though the two leading algorithms are close. In the synthetic held-out benchmark, PG-MUAE records MAE of 0.01878 USD/unit, RMSE of 0.02516 USD/unit, and R^2 of 0.96415; Gradient Boosting records 0.01851, 0.02519, and 0.96408, respectively. Thus PG-MUAE has slightly lower RMSE and higher R^2, while Gradient Boosting has lower MAE. Relative to mass-proportional allocation, PG-MUAE lowers RMSE by 78.875%, indicating that product, utilization, and tariff features improve on a simple volume-based rule in this synthetic test. Equation (13) defines relative RMSE improvement. Table 7 reports the underlying aggregate metrics so that these comparisons remain auditable. No repeated-seed, parameter-perturbation, or facility-level external validation results are available; the ranking therefore describes metric sensitivity within one synthetic evaluation, not robustness under changing operating conditions.

$$\text{Relative improvement (\%)} = \frac{RMSE_b - RMSE_p}{RMSE_b} \times 100 \quad (13)$$

Where $RMSE_b$ is the benchmark algorithm's root mean squared error and $RMSE_p$ is PG-MUAE's root mean squared error. A positive value indicates lower RMSE for PG-MUAE than for the selected benchmark.

Table 7 Aggregate Held-Out Synthetic Algorithm Metrics Used in the Sensitivity Comparison.

Algorithm	MAE (USD/unit)	RMSE (USD/unit)	R^2
Mass-proportional	0.10027	0.11910	0.19697
Ridge	0.03075	0.04168	0.90165
Random Forest	0.02189	0.02964	0.95026
Extra Trees	0.02209	0.02982	0.94965
Gradient Boosting	0.01851	0.02519	0.96408
PG-MUAE	0.01878	0.02516	0.96415

- Table interpretation. The proposed estimator's advantage is narrow against Gradient Boosting and stronger against mass-proportional allocation. This comparison supports using multiple error measures and retaining the baseline, rather than describing PG-MUAE as uniformly superior.

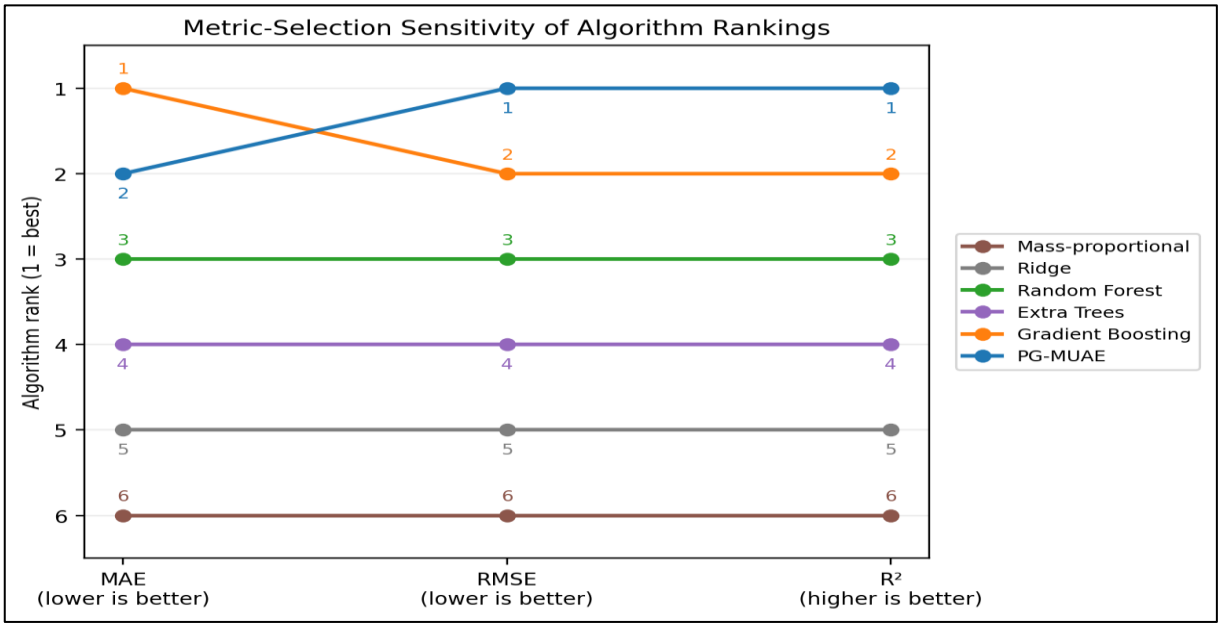


Fig 5 Bump Chart of Algorithm Rank Across MAE, RMSE, and R^2.

Figure 5 ranks algorithms separately for MAE, RMSE, and R^2, with rank 1 denoting the best result. The two leading models exchange first and second place: Gradient Boosting is first on MAE and second on RMSE and R^2, whereas PG-MUAE is second on MAE and first on RMSE and R^2. Random Forest holds rank 3 across all three metrics, and Extra Trees holds rank 4; Ridge and mass-proportional allocation remain ranks 5 and 6. This stable ordering of the four lower-ranked methods contrasts with the narrow top-model reversal. The chart therefore shows sensitivity to metric choice, not sensitivity to altered model inputs or data samples. For manufacturers, use the model as a screening aid only after reconciling utility bills and validating against calibrated submeters. A stronger robustness claim requires repeated seeds, prespecified perturbations to mix response, utilization, tariff dispersion, and shared-load weights, plus evaluation on unseen facilities and operating periods.

V. CONCLUSIONS AND RECOMMENDATIONS

➢ *Conclusions*

The study concludes that product-level utility-cost estimation is materially more informative when production mix, capacity utilization, utility tariffs, and shared-load allocation are modeled together. In the synthetic facility-held-out evaluation, PG-MUAE achieved an MAE of $0.01878 per unit, RMSE of $0.02516, and R^2 of 0.96415. Gradient Boosting attained a slightly lower MAE ($0.01851), while PG-MUAE narrowly led on RMSE and R^2. The $0.00003 RMSE difference and 0.00007 R^2 difference are too small to establish practical superiority from one simulated realization. In contrast, the mass-proportional baseline produced RMSE of $0.11910 and R^2 of 0.19697, demonstrating the limitations of allocating facility utility expense solely by output volume within this controlled experiment.

The principal conclusion is therefore methodological: a physics-guided cost baseline with a learned residual can preserve tariff and energy-accounting structure while representing nonlinear mix and utilization effects. The facility-grouped split across 72 synthetic facilities, 36 months, and five product families limits facility leakage, but does not establish external validity for U.S. plants. The utilization-band line figure is a synthetic decomposition calibrated to reproduce each algorithm's aggregate RMSE; it is not independent evidence of plant behavior. Accordingly, the findings support feasibility and motivate field validation rather than proving improved audit outcomes. Product-level estimates should be interpreted alongside ledger-reconciliation checks,

uncertainty intervals, and metered validation. Broad adoption claims require replicated seeds, diverse manufacturing sectors, calibrated submeters, and prospective testing under actual tariff structures and production schedules. Future research must verify these assumptions empirically across plants.

- *Contributions to Industrial Energy Auditing and Cost Estimation*

This study contributes a transparent framework for closing the gap between facility utility bills and product-level cost visibility. PG-MUAE combines an engineering cost baseline—utility-specific energy use valued at applicable tariffs, plus allocated shared and demand charges—with a gradient-boosted residual that learns remaining nonlinear structure. Its objective includes prediction error, facility-total reconciliation, and a nonnegativity penalty. This makes the estimator's accounting logic inspectable and allows its product allocations to be compared with direct machine-learning benchmarks under common train, validation, and test partitions.

The comparative design also clarifies how performance should be interpreted. Across the synthetic held-out test set, PG-MUAE recorded RMSE of $0.02516 per unit and R^2 of 0.96415, narrowly improving on Gradient Boosting's $0.02519 RMSE and 0.96408 R^2, while Gradient Boosting retained the lowest MAE. Reporting all three metrics prevents a single favorable score from obscuring the trade-off. Comparison with mass-proportional allocation further shows that output-share accounting can miss product energy intensity, utilization, and tariff timing. The framework also specifies diagnostic quantities for mix-stratified error and monthly ledger gaps, making predictive accuracy distinct from accounting closure. Because the evidence is synthetic, the contribution is a reproducible evaluation design and a testable estimator architecture, not a validated national benchmark or verified cost-saving intervention. It provides a basis for future plant-level studies that can measure transfer error, cost-allocation fairness, and decision usefulness against calibrated submeters and audited utility invoices. It can also support transparent audit trails for managers comparing estimates across product lines and tariff periods.

- *Recommendations for U.S. Manufacturers*

U.S. manufacturers considering product-level utility-cost estimation should begin with a bounded pilot at a facility where production records, tariff schedules, and utility invoices can be matched by time period. The pilot should define the accounting boundary, enumerate electricity, gas, steam, and demand charges, and establish consistent product-family identifiers across enterprise resource planning, manufacturing execution, and metering systems. Where dedicated submeters are unavailable, shared loads should be assigned through documented drivers—such as machine-hours, thermal throughput, or validated allocation weights—and reconciled monthly to the facility bill.

Firms should treat PG-MUAE as a decision-support candidate, not an automatically superior model. The synthetic benchmark showed near-tied RMSE for PG-MUAE and Gradient Boosting, with Gradient Boosting retaining a small MAE advantage. Each pilot should therefore compare the physics-guided estimator with a transparent output-share baseline, regularized regression, and tree ensembles using facility-level holdouts. Managers should inspect errors by product family, utilization range, tariff period, and production mix, alongside signed residuals and the difference between summed product estimates and billed totals. A model should not be used for pricing, make-or-buy decisions, or energy-efficiency claims until those diagnostics pass prespecified tolerances and plant staff can trace estimates to measured inputs. Manufacturers can then use validated product costs to prioritize submeter installation, schedule energy-intensive products during favorable tariff windows, target high-intensity processes, and evaluate efficiency projects. Any savings claim should be confirmed through a separate measurement-and-verification design that controls for volume, mix, weather, and operating changes. This phased deployment reduces risk while building internal capability gradually.

- *Policy, Data, and Future Research Recommendations*

Policymakers can improve industrial energy-cost analysis by supporting interoperable, privacy-protective data infrastructure that links interval utility consumption, tariff schedules, production quantities, and product-family metadata. Federal and state programs could fund submetering pilots for small and medium-sized manufacturers, publish common guidance for allocating shared loads, and define minimum reporting fields for energy audits. Grants should require transparent baselines, documented allocation assumptions, and reconciliation to utility invoices, while protecting commercially sensitive production data through aggregation, access controls, and retention rules. These measures would make product-level cost methods more comparable without requiring firms to disclose proprietary recipes or customer information.

Future research should replace the present synthetic panel with multisector plant data and externally validated meter-to-product mappings. Studies should preregister facility-level splits, compare PG-MUAE with Gradient Boosting and other benchmarks under matched tuning budgets, and report confidence intervals from repeated seeds and bootstrap resampling across facilities. Sensitivity experiments should vary tariff volatility, shared-load weights, missing-meter coverage, product-mix concentration, and utilization ranges. Researchers should quantify both predictive error and reconciliation gaps, test calibration under unseen tariffs and operating regimes, and evaluate whether cost estimates change investment or scheduling decisions. The current results do not establish causal energy savings, fairness across product families, or performance in every U.S. manufacturing subsector. Prospective studies should therefore combine calibrated submeters, audited bills, documented production data, and decision-impact measures before recommending broad regulatory

adoption. Independent replication should precede policy adoption. Policymakers should encourage open protocols and independent evaluation, while allowing firms to retain secure control of facility-level data.

REFERENCES

[1.] Avevor, J., Aikins, S. A., & Enyejo, L. A. (2025). Optimizing gas and steam turbine performance through predictive maintenance and thermal optimization for sustainable and cost-effective power generation. *International Journal of Innovative Science and Research Technology, 10*(3). https://doi.org/10.38124/ijisrt/25mar1336

[2.] Ekwaro-Osire, H., Bode, D., Ohlendorf, J.-H., & Thoben, K.-D. (2025). Manufacturing process energy consumption modeling: A methodology to identify the most appropriate model. *Journal of Intelligent Manufacturing, 36*, 5673–5693. https://doi.org/10.1007/s10845-024-02514-z

[3.] Ekwaro-Osire, H., Bode, D., Ohlendorf, J.-H., & Thoben, K.-D. (2025). Manufacturing process energy consumption modeling: A methodology to identify the most appropriate model. *Journal of Intelligent Manufacturing, 36*, 5673–5693. https://doi.org/10.1007/s10845-024-02514-z

[4.] Guo, J., Ma, J., Li, Z., & Hong, J. (2022). Building a top-down method based on machine learning for evaluating energy intensity at a fine scale. *Energy, 255*, 124505. https://doi.org/10.1016/j.energy.2022.124505

[5.] Ibn Batouta, K., Aouhassi, S., & Mansouri, K. (2023). Energy efficiency in the manufacturing industry—A tertiary review and a conceptual knowledge-based framework. *Energy Reports, 9*, 4635–4653. https://doi.org/10.1016/j.egyr.2023.03.107

[6.] Ibokette, A. I., Ogundare, T. O., Danquah, E. O., Anyebe, A. P., Agaba, J. A., & Olola, T. M. (2024). The impacts of emotional intelligence and IoT on operational efficiency in manufacturing: A cross-cultural analysis of Nigeria and the US. *Computer Science & IT Research Journal, 5*(8), 1978–2010. https://doi.org/10.51594/csitrj.v5i8.1464 researchgate.net

[7.] Ibokette, A. I., Ogundare, T. O., Danquah, E. O., Anyebe, A. P., Agaba, J. A., & Olola, T. M. (2024). The impacts of emotional intelligence and IoT on operational efficiency in manufacturing: A cross-cultural analysis of Nigeria and the US. *Computer Science & IT Research Journal, 5*(8), 1978–2010. https://doi.org/10.51594/csitrj.v5i8.1464

[8.] Idika, C. N., James, U. U., Ijiga, O. M., & Enyejo, L. A. (2023). Digital twin-enabled vulnerability assessment with zero trust policy enforcement in smart manufacturing cyber-physical system. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 9*(6). https://doi.org/10.32628/CSEIT23906189

[9.] Idoko, I. P., Ijiga, O. M., Agbo, D. O., Abutu, E. P., Ezebuka, C. I., & Umama, E. E. (2024). Comparative analysis of Internet of Things (IoT) implementation: A case study of Ghana and the USA—Vision, architectural elements, and future directions. *World Journal of Advanced Engineering Technology and Sciences, 11*(1), 180–199.

[10.] Idoko, I. P., Ijiga, O. M., Akoh, O., Agbo, D. O., Ugbane, S. I., & Umama, E. E. (2024). Empowering sustainable power generation: The vital role of power electronics in California's renewable energy transformation. *World Journal of Advanced Engineering Technology and Sciences, 11*(1), 274–293. https://doi.org/10.30574/wjaets.2024.11.1.0058 wjaets.com

[11.] Idoko, I. P., Ijiga, O. M., Akoh, O., Agbo, D. O., Ugbane, S. I., & Umama, E. E. (2024). Empowering sustainable power generation: The vital role of power electronics in California's renewable energy transformation. *World Journal of Advanced Engineering Technology and Sciences, 11*(1), 274–293.

[12.] Ilesanmi, M. O., Okoh, O. F., Balogun, S. A., & Ijiga, O. M. (2024). Data-driven portfolio optimization for utility-scale solar, wind, and battery energy storage systems (BESS): Integrating performance analytics with investor EBITDA targets. *International Journal of Scientific Research and Modern Technology, 3*(11), 141–155. https://doi.org/10.38124/ijsrmt.v3i11.943

[13.] James, U. U. (2022). Machine learning-driven anomaly detection for supply chain integrity in 5G industrial automation systems. *International Journal of Scientific Research in Science and Engineering Technology, 9*(2). https://doi.org/10.32628/IJSRSET22549

[14.] James, U. U., Idika, C. N., & Enyejo, J. O. (2023). Zero trust architecture leveraging AI-driven behavior analytics for industrial control systems in energy distribution networks. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 9*(4). https://doi.org/10.32628/CSEIT23564522

[15.] Jamil, M., He, N., Wei, Z., Gupta, M. K., & Khan, A. M. (2023). A novel quantifiable approach of estimating energy consumption, carbon emissions and cost factors in manufacturing of bearing steel based on triple bottom-line approach. *Sustainable Materials and Technologies, 36,* Article e00593. https://doi.org/10.1016/j.susmat.2023.e00593 ScienceDirect

[16.] Ma, S., Zhang, Y., Lv, J., Ren, S., Yang, H., & Wang, C. (2022). Data-driven cleaner production strategy for energy-intensive manufacturing industries: Case studies from Southern and Northern China. *Advanced Engineering Informatics, 53*, Article 101684. https://doi.org/10.1016/j.aei.2022.101684 ScienceDirect

[17.] Manuel, H. N. N., Adeoye, T. O., Idoko, I. P., Akpa, F. A., Ijiga, O. M., & Igbede, M. A. (2024). Optimizing passive solar design in Texas green buildings by integrating sustainable architectural

features for maximum energy efficiency. *Magna Scientia Advanced Research and Reviews, 11*(1), 235–261. https://doi.org/10.30574/msarr.2024.11.1.0089 msarr.magnascientiapub.com

[18.] Manuel, H. N. N., Adeoye, T. O., Idoko, I. P., Akpa, F. A., Ijiga, O. M., & Igbede, M. A. (2024). Optimizing passive solar design in Texas green buildings by integrating sustainable architectural features for maximum energy efficiency. *Magna Scientia Advanced Research and Reviews, 11*(1), 235–261. https://doi.org/10.30574/msarr.2024.11.1.0089

[19.] Ocharo, D. O., Omachi, A., Aikins, S. A., & Adaudu, I. I. (2024). SCADA-enabled predictive maintenance framework for cogeneration systems in American manufacturing facilities. *International Journal of Scientific Research and Modern Technology, 3*(7), 30–44. https://doi.org/10.38124/ijsrmt.v3i7.947 ijsrmt.com

[20.] Ocharo, D. O., Omachi, A., Aikins, S. A., & Adaudu, I. I. (2024). SCADA-enabled predictive maintenance framework for cogeneration systems in American manufacturing facilities. *International Journal of Scientific Research and Modern Technology, 3*(7), 30–44. https://doi.org/10.38124/ijsrmt.v3i7.947

[21.] Oh, J., & Min, D. (2024). Prediction of energy consumption for manufacturing small and medium-sized enterprises (SMEs) considering industry characteristics. *Energy, 300*, 131621. https://doi.org/10.1016/j.energy.2024.131621

[22.] Okeke, R. O., Ibokette, A. I., Ijiga, O. M., Enyejo, L. A., Ebiega, G. I., & Olumubo, O. M. (2024). The reliability assessment of power transformers. *Engineering Science & Technology Journal, 5*(4), 1149–1172. https://doi.org/10.51594/estj.v5i4.981 UNT Digital Library

[23.] Okeke, R. O., Ibokette, A. I., Ijiga, O. M., Enyejo, L. A., Ebiega, G. I., & Olumubo, O. M. (2024). The reliability assessment of power transformers. *Engineering Science & Technology Journal, 5*(4), 1149–1172.

[24.] Oyekan, M., Jinadu, S. O., & Enyejo, J. O. (2023). Harnessing data analytics to maximize renewable energy asset performance. *International Journal of Scientific Research and Modern Technology, 2*(8), 64–80. https://doi.org/10.38124/ijsrmt.v2i8.850

[25.] Peng, C., Peng, T., Liu, Y., Geissdoerfer, M., Evans, S., & Tang, R. (2021). Industrial Internet of Things enabled supply-side energy modelling for refined energy management in aluminium extrusions manufacturing. *Journal of Cleaner Production, 301*, Article 126882. https://doi.org/10.1016/j.jclepro.2021.126882 Monash University

[26.] Tan, D., Suvarna, M., Tan, Y. S., Li, J., & Wang, X. (2021). A three-step machine learning framework for energy profiling, activity state prediction and production estimation in smart process manufacturing. *Applied Energy, 291*, 116808. https://doi.org/10.1016/j.apenergy.2021.116808

[27.] Vance, D., Jin, M., Wenning, T. J., Nimbalkar, S. U., & Price, C. R. (2025). Next-level energy management in manufacturing: Facility-level energy digital twin framework based on machine learning and automated data collection. *Energies, 18*(13), 3242. https://doi.org/10.3390/en18133242

[28.] Wang, M., & Feng, C. (2023). Measuring capacity utilization under the constraints of energy consumption and CO_2 emissions using meta-frontier DEA: A case of China's non-ferrous metal industries. *Resources Policy, 80*, Article 103278. https://doi.org/10.1016/j.resourpol.2022.103278 Request PDF

[29.] Xie, Y., Dai, Y., Zhang, C., & Liu, J. (2025). An integration model enabled deep learning for energy prediction of machine tools. *Journal of Cleaner Production, 495*, 145075. https://doi.org/10.1016/j.jclepro.2025.145075

[30.] Zhang, D., & Karplus, V. J. (2025). Management practices and manufacturing firm responses to a randomized energy audit. *Nature Energy, 10*, 557–568. https://doi.org/10.1038/s41560-025-01729-5

[31.] Zhang, D., Durmaz, N., & Kagochi, J. M. (2025). Energy outages, in-house power generators, and capacity utilization: The case of African manufacturing. *Energy Economics, 147*, Article 108614. https://doi.org/10.1016/j.eneco.2025.108614 ScienceDirect

[32.] Zhang, M., & Wang, R. (2025). Energy consumption forecasting in non-stationary machining processes based on a physics-guided transformer. *Computers in Industry, 171*, 104321. https://doi.org/10.1016/j.compind.2025.104321

[33.] Ogundola, T. O. (2025). Post-quantum cryptography for secure banking transactions. International Journal of Scientific Research and Modern Technology, 4, 40-42.

[34.] Ogundola, T. O. (2024) Barriers to implementing quantum-resistant encryption in financial institutions, World Journal of Advanced Research and Reviews, 2024, 22(03), 2323-2337, DOI: https://doi.org/10.30574/wjarr.2024.22.3.1938

[35.] Ogundola, T. (2020). Advances in quantum-secure banking: Cryptographic solutions. International Journal of Science and Research Archive, 1(1), 228–230. https://doi.org/10.30574/ijsra.2020.1.1.0048

[36.] Ogundola, T. (2022). Quantum-resistant cryptographic techniques for financial institutions. International Journal of Scientific Research and Modern Technology, 1(11), 18–20. https://doi.org/10.38124/ijsrmt.v1i11.586

[37.] Bamidele, A. A., Justin, O., Adanne, S. (2025). Integrating machine learning with environmental chemistry to forecast pollutant releases in coatings production. World Journal of Advanced Research and Reviews, 2025, 27(3), 1061-1072. Article DOI: https://doi.org/10.30574/wjarr.2025.27.3.3153

[38.] Bamidele, A. A., Rotimi, A. M., James, C. O., Obinna, J. (2025). Coupling environmental chemistry with Artificial Intelligence for coastal

hazard risk assessment under climate change. International Journal of Science and Research Archive, 2025, 16(03), 498–507. Article DOI: https://doi.org/10.30574/ijsra.2025.16.3.2572.

[39.] Abolade, Y. A., Adegoke, A., Noi, S. M., Nwoye, C. M., Ajayi, O., Kyrian, O., Akagbue, B. O., Siame, T., & Bamidele, A. A. (2024). Promoting health: Introducing an eco-friendly herbal mosquito repellent extracted from local sweet orange peels *GSC Biological and Pharmaceutical Sciences,* 2024, 26(01), 315–327 DOI: https://doi.org/10.30574/gscbps.2024.26.1.0014

[40.] Idowu, K., Cherotich, V., Aniebonam, E. E., Odozor, L., & During, A. (2025). Integrating AI and machine learning into cyber risk management for critical utility systems. International Journal of Artificial Intelligence Engineering and Technology. https://doi.org/10.54660/IJAIET.2025.6.2.32-48

[41.] Odozor, L., Berko, S. N., Obu, A. U., & Idowu, K. (2025). Building a resilient U.S. cybersecurity workforce through practical training. International Journal of Scientific and Management Research, 8(8), 30–36.

[42.] Folorunso, T. N., Idowu, K., & Balogun, A. (2025). AI-driven financial risk forecasting: A multi-model ensemble approach for enterprise decision intelligence. Journal of Computational Analysis and Applications, 34(3). https://eudoxuspress.com/index.php/pub/article/view/5080

[43.] Darkwa, S. (2025). Developing an Evidence-Based Healthcare Workforce Coordination and Organizational Communication Framework for Behavioral Healthcare Settings: A Scoping Review and Conceptual Model. *International Journal of Scientific Research and Modern Technology*, *4*(12), 268–290. https://doi.org/10.38124/ijsrmt.v4i12.1776.