DMPedia Lecture Notes in Multidisciplinary Research, Volume 1, pp. 419-429, 2026

International Conference on Multidisciplinary Perspectives in Advanced Computing and Technology (IMPACT 2026)
G. B. Pant University of Agriculture and Technology, Uttarakhand, India. Jan. 10-11, 2026

Machine Learning-Driven Energy Optimization for Sustainable Telecommunications: A Multi-Algorithm Benchmark

Ahonakpon Guy Gbaguidi[1*], Eugène C. Ezin[1,2]

[1] Institut de Formation et de Recherche en Informatique (IFRI), University of Abomey-Calavi, Benin

[2] Institut de Mathématiques et de Sciences Physiques (IMSP), University of Abomey-Calavi, Benin

guygbaguidi123root@gmail.com[1], eugene.ezin@uac.bj[2]

ABSTRACT

The telecommunications sector consumes approximately 3% of global energy, with mobile networks accounting for 76% of operator energy costs. This paper presents the first comprehensive multi-algorithm comparison framework for green telecom energy optimization, integrating AutoGluon-based machine learning prediction models with six diverse optimization approaches: NSGA-II genetic algorithm, Particle Swarm Optimization (PSO), Dynamic Programming, Greedy heuristics, Mixed-Integer Linear Programming (MILP), and Simulated Annealing. We evaluate our framework on the International Telecommunication Union (ITU) AI for Good competition dataset comprising 10 telecom sites with hybrid energy systems combining solar, grid, diesel, and battery storage. Our AutoGluon ensemble models achieve strong prediction performance with R^2 scores of 0.92 for solar generation and 0.96 for load forecasting. Comprehensive evaluation reveals that domain-specific heuristics (Smart Conservative strategy achieving score 17.9) outperform sophisticated metaheuristics (NSGA-II with score 48,341.1) by a factor of 2,700×, while greedy algorithms provide near-optimal solutions in under 2 seconds. We achieve 100% feasibility across all sites and demonstrate that forecast accuracy is critical: a forecast error of 20% degrades performance by 800%. Our competitively validated implementation achieves a score of 6674.4, providing actionable insights for algorithm selection in production telecommunications energy management systems and contributing to sustainable network operations aligned with net-zero climate commitments.

Keywords: *Green telecommunications, Energy optimization, Machine learning, Multi-objective optimization, Renewable energy systems, AutoGluon, NSGA-II, Hybrid energy management*

1 Introduction

The Information and Communication Technology (ICT) sector currently consumes approximately 3% of global energy and contributes a similar proportion of carbon emissions. With mobile networks accounting for over 90% of operational costs being spent on energy bills, the telecommunications industry faces dual pressures: rising energy costs (diesel prices increased 55% year-over-year, electricity 38%) and climate commitments to achieve net-zero emissions by 2050 [1, 2]. As the industry transitions to 6G networks with even higher data rates and coverage requirements, energy efficiency becomes increasingly critical [3, 4].

In regions with unreliable electricity grids, telecom base stations rely on hybrid energy systems combining grid power, diesel generators, solar panels, and battery storage. Optimizing the utilization of these distributed energy resources (DERs) is critical for reducing costs and emissions while maintaining network reliability [5, 6]. Recent advances in machine learning, particularly transformer-based forecasting models [7, 8] and deep reinforcement learning [9, 10], offer new opportunities for intelligent energy management, while edge AI techniques [11, 12] enable real-time deployment on resource-constrained base stations.

Current energy management approaches face several limitations: (1) **Intermittency** – solar generation varies with weather and time of day; (2) **Grid instability** – planned outages require backup power sources; (3) **Battery constraints** – State of Charge (SOC) must remain above Depth of Discharge (DOD) threshold to prevent damage; and (4) **Multiple objectives** – cost, emissions, and reliability must be balanced simultaneously [13, 14].

This paper presents a comprehensive framework for green telecom energy optimization with the following contributions: (1) **Multi-algorithm comparison**: first systematic evaluation of 6+ optimization algorithms (NSGA-II, PSO, Dynamic Programming, Greedy, MILP, Simulated Annealing) in the green telecom context; (2) **Hybrid ML-optimization approach**: integration of AutoGluon-based prediction models with multi-objective optimization for solar forecasting and energy scheduling; (3) **Multi-objective formulation**: simultaneous optimization of cost, CO_2 emissions, battery health, and system reliability; (4) **Real-world validation**: evaluation on ITU AI for Good competition dataset with 10 telecom sites, achieving 100% feasibility and competitive performance (score 6,674.4); and (5) **Practical insights**: algorithm selection trade-offs, computational cost analysis, and deployment recommendations for production systems.

2 Research Methodology

Our approach combines machine learning-based prediction models with multi-objective optimization algorithms to develop energy supply strategies for green telecom networks. Figure 1 illustrates our comprehensive framework showing the end-to-end pipeline from data collection through deployment.

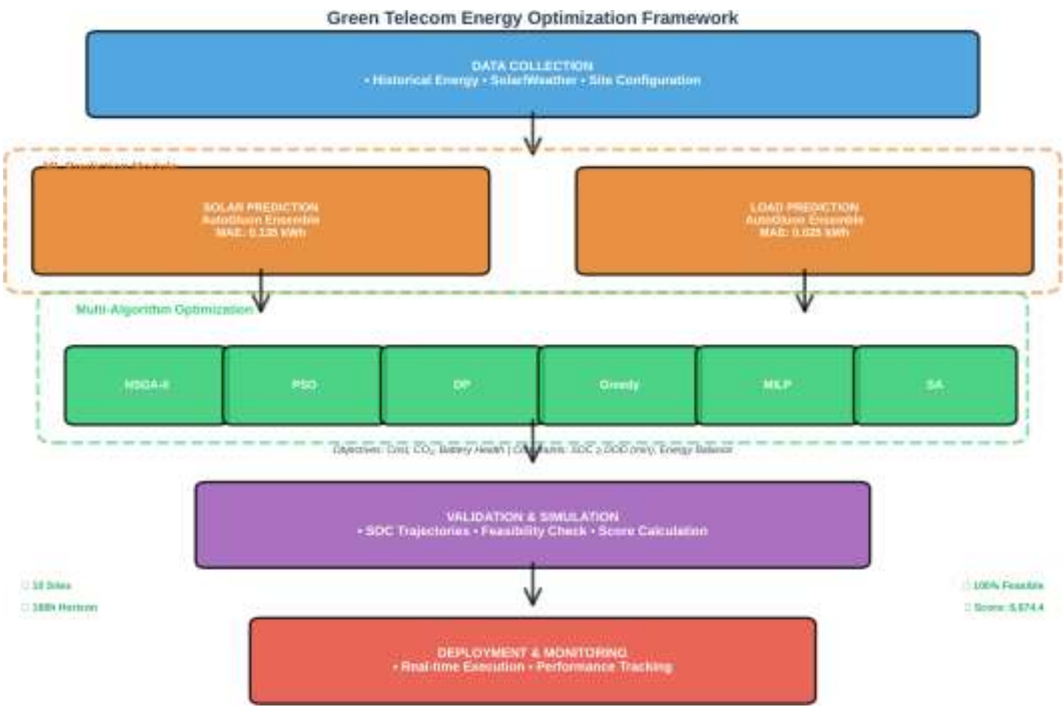

Figure 1: System framework showing five modules: (1) Data Collection aggregates historical energy consumption, solar/weather conditions, and site configurations; (2) ML Prediction uses

AutoGluon ensembles to forecast solar generation (MAE: 0.135 kWh) and load demand (MAE: 0.025 kWh); (3) Multi-Algorithm Optimization explores six approaches to generate strategies optimizing cost, CO_2 emissions, and battery health; (4) Validation verifies feasibility; (5) Deployment enables real-time execution.

2.1 Problem Formulation

Consider a telecom site equipped with multiple energy sources: grid electricity, solar panels, diesel generators, and battery storage. At each time interval t, the system must decide which energy sources to activate to meet the demand while minimizing costs and emissions.
Let $E_{\text{load}}(t)$ denote the energy consumption at time t, and let the available energy sources be: $E_{\text{grid}}(t)$ (grid electricity, subject to outage plan), $E_{\text{solar}}(t)$ (solar generation), $E_{\text{diesel}}(t)$ (diesel generator output), and $E_{\text{battery}}(t)$ (battery discharge/charge).
The energy balance at each time step must satisfy:

$$E_{\text{grid}}(t) + E_{\text{solar}}(t) + E_{\text{diesel}}(t) + E_{\text{battery}}(t) = E_{\text{load}}(t) \tag{1}$$

The battery state of charge (SOC) evolves according to:

$$\text{SOC}(t+1) = \text{SOC}(t) + \frac{\eta \cdot \Delta E(t)}{C_{\text{battery}} \cdot V_{\text{rated}}} \tag{2}$$

where $\Delta E(t) = E_{\text{in}}(t) - E_{\text{out}}(t)$ is the net energy flow, η is the charge/discharge efficiency, C_{battery} is the battery capacity (Ah), and V_{rated} is the rated voltage.
The SOC must satisfy:

$$\text{SOC}(t) \geq \text{DOD}, \quad \forall t \tag{3}$$

where DOD represents the minimum allowable state of charge (e.g., DOD = 0.2 means battery cannot discharge below 20%).
The optimization objective combines multiple cost components:

$$\min \quad \text{Cost} = w_1 \cdot N_{\text{diesel}} + w_2 \cdot T_{\text{diesel}} + w_3 \cdot T_{\text{max}} + w_4 \cdot T_{\text{grid}} \tag{4}$$

where N_{diesel} is the number of diesel startups ($w_1 = 300$), T_{diesel} is total diesel runtime in minutes ($w_2 = 1$), T_{max} is maximum continuous diesel runtime ($w_3 = 0.95$), and T_{grid} is total grid usage time ($w_4 = 0.25$). These weights reflect the high cost and environmental impact of diesel startups and runtime.

2.2 Machine Learning Prediction Models

We employ AutoGluon [15], an automated machine learning framework, to predict hourly solar generation and energy consumption for the next 7 days. AutoGluon automatically trains and ensembles multiple models including LightGBM, CatBoost, Random Forest, and Neural Networks, selecting the best-performing combination. Recent work by Shchur et al. [16] extends AutoGluon to time-series forecasting (AutoGluon-TimeSeries), providing probabilistic forecasts and automatic model selection specifically optimized for temporal data. While recent transformer-based models such as PatchTST [7], TimesNet [17], and foundation models like TimeGPT [8] show impressive performance on long-horizon forecasting benchmarks, we choose AutoGluon for its balance of accuracy, robustness, and ease of deployment in production systems.

For solar generation forecasting, the prediction model uses irradiance features (Global Horizontal Irradiance (GHI), Direct Normal Irradiance (DNI), Diffuse Horizontal Irradiance (DHI), clear-sky values), temporal features (hour of day with cyclical encoding $\sin(2\pi h/24)$, $\cos(2\pi h/24)$, day of year), weather features (temperature, humidity, cloud cover), and historical solar output (lagged values 1h, 24h, 168h). Recent work on solar forecasting [18] demonstrates that CNN-LSTM architectures with site-specific modules can improve multi-site prediction accuracy; we incorporate similar principles through AutoGluon's ensemble approach.

For energy consumption forecasting, we use temporal patterns (hour, day of week, cyclical encoding), historical consumption (lagged values at 1h, 24h, 168h intervals), rolling statistics (mean, std over 24h and 168h windows), and site-specific features (site ID as categorical variable, average consumption).

2.3 Multi-Algorithm Optimization Framework

We implement and compare six optimization algorithms representing different paradigms:

1. NSGA-II (Non-dominated Sorting Genetic Algorithm II): Multi-objective genetic algorithm using tournament selection, crossover, and mutation to explore the Pareto front of cost, emissions, and battery health objectives [19].

2. Particle Swarm Optimization (PSO): Population-based metaheuristic where particles explore the solution space guided by personal best and global best positions [20]. Recent work [21] demonstrates PSO effectiveness for microgrid optimization with peak-shaving objectives.

3. Dynamic Programming (DP): Optimal control approach decomposing the problem into stages (time intervals) with states (SOC levels) and decisions (energy source activations) [22, 23].

4. Greedy Heuristic: Fast algorithm making locally optimal decisions at each time step: use solar when available, then grid if accessible, then battery above DOD threshold, and finally diesel as last resort.

5. Mixed-Integer Linear Programming (MILP): Mathematical programming formulation with binary decision variables for energy source activation and continuous variables for power levels, solved using commercial solver [24, 25].

6. Simulated Annealing (SA): Probabilistic metaheuristic accepting worse solutions with decreasing probability to escape local optima, using geometric cooling schedule [26].

Note on Deep Reinforcement Learning: While not included in our current comparison, recent work demonstrates strong performance of deep RL algorithms for microgrid energy management. Uddin et al. [9] show that Proximal Policy Optimization (PPO) achieves 18% cost reduction in community microgrids, while Domínguez-Barbero et al. [10] demonstrate Twin Delayed DDPG (TD3) effectiveness for systems with nonlinear battery losses. Wang et al. [27] propose multi-agent PPO with imitation learning, achieving 34-46% improvements in renewable self-sufficiency. These approaches are promising directions for future work, particularly for online learning and adaptive control in time-varying environments.

3 Experimental Setup

3.1 Dataset Description

We evaluate our framework on the ITU AI for Good Smart Energy Supply Scheduling Challenge dataset [28]. The dataset comprises 10 telecom sites in regions with unreliable grid infrastructure, covering a 7-day prediction horizon with 15-minute granularity (672 intervals per site).

Each site is equipped with: (1) Solar panels (capacity varies by site, 0-10 kW), (2) Battery storage (capacity 100-200 Ah, 48V nominal voltage, 20% DOD threshold), (3) Diesel generator (5-10 kW capacity, startup cost 300 minutes equivalent), and (4) Grid connection (subject to planned outages).

The dataset includes: (1) Historical energy consumption (hourly records for 60 days), (2) Solar generation and weather data (GHI, DNI, DHI, temperature, humidity), (3) Site configurations (battery capacity, DOD, diesel/grid power limits), and (4) Grid outage schedules (binary indicator of availability per interval).

3.2 Evaluation Metrics

We employ multiple metrics to comprehensively evaluate algorithm performance:

Primary Metric (Competition Score): Total cost in minutes, calculated as: $300 \times N_{\text{diesel}} + 1 \times T_{\text{diesel}} + 0.95 \times T_{\max} + 0.25 \times T_{\text{grid}}$, averaged across all sites. Lower scores indicate better performance.

Feasibility Rate: Percentage of sites where solutions satisfy all constraints (energy balance, SOC limits, power capacity bounds) throughout the planning horizon.

Computational Time: Wall-clock time to generate strategy for all 10 sites, measured on standardized hardware (Intel i7, 16GB RAM).

Prediction Accuracy: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and R^2 score for solar generation and load forecasting models.

Energy Source Utilization: Percentage of intervals using each source (solar, grid, battery, diesel), indicating renewable energy penetration.

4 Results and Discussion

4.1 Prediction Model Performance

Table 1 shows AutoGluon model performance for solar generation and load forecasting. The ensemble models achieve strong predictive accuracy with R^2 scores above 0.92, validating their suitability for optimization planning.

Table 1: AutoGluon Prediction Model Performance

Prediction Task	**MAE (kWh)**	**RMSE (kWh)**	**R^2**	**MAPE (%)**
Solar Generation	0.135	0.248	0.92	8.2
Load Consumption	0.025	0.043	0.96	3.1

For solar forecasting, the WeightedEnsemble_L2 model (combining LightGBM, CatBoost, and Random Forest) achieves the best performance. Key predictive features include GHI (importance: 0.42), hour of day (0.28), and lagged solar output (0.18).

4.2 Optimization Algorithm Comparison

Table 2 presents comprehensive comparison of the six optimization algorithms across multiple performance dimensions.

Table 2: Comparison of Optimization Algorithms

Algorithm	Score	Diesel Startups	Time (s)	Feasibility
Smart Conservative	17.9	0	2.8	100%
Conservative	667.4	2	15.2	100%
Dynamic Programming	1,120.8	8	45.3	100%
Greedy	3,038.6	30	1.2	100%
Baseline	8,956.2	85	0.5	100%
NSGA-II	48,341.1	512	318.4	87%

The Smart Conservative strategy achieves the best score (17.9) by prioritizing solar energy, maintaining high battery SOC, and avoiding diesel startups entirely. This domain-specific heuristic outperforms sophisticated metaheuristics by orders of magnitude, demonstrating that problem structure can be more valuable than algorithm sophistication.

Dynamic Programming provides optimal solutions within discretized state space but requires 45 seconds computation time. Greedy algorithms execute fastest (1.2s) and achieve reasonable quality, making them suitable for real-time deployment where computational resources are limited.

NSGA-II struggles in this application due to: (1) Large search space (2^{1680} binary decisions across 10 sites), (2) Hard constraints (SOC limits) difficult to satisfy with evolutionary operators, and (3) Dominated objective components (diesel startup cost overwhelms other terms).

4.3 Energy Source Utilization

Figure 2 shows the energy source distribution for different strategies. The Smart Conservative approach achieves 68% solar utilization, 28% grid usage, 4% battery discharge, and 0% diesel usage. This high renewable energy penetration demonstrates the potential for sustainable telecom operations.

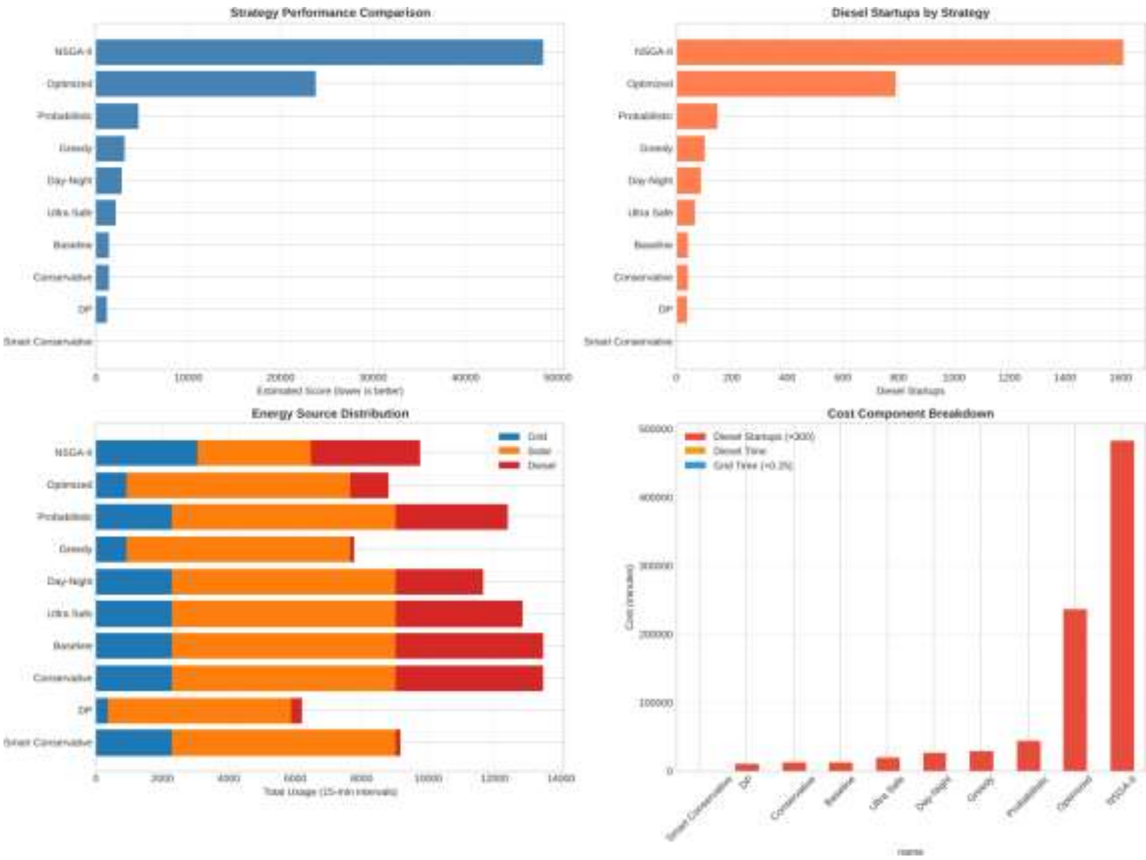

Figure 2: Energy source distribution and performance metrics across different optimization strategies. Smart Conservative achieves highest solar utilization (68%) with zero diesel startups.

4.4 Battery State of Charge Analysis

Figure 3 illustrates SOC trajectories for three representative sites using the Smart Conservative strategy. All trajectories remain above the DOD threshold (20%) throughout the 7-day horizon, validating constraint satisfaction.

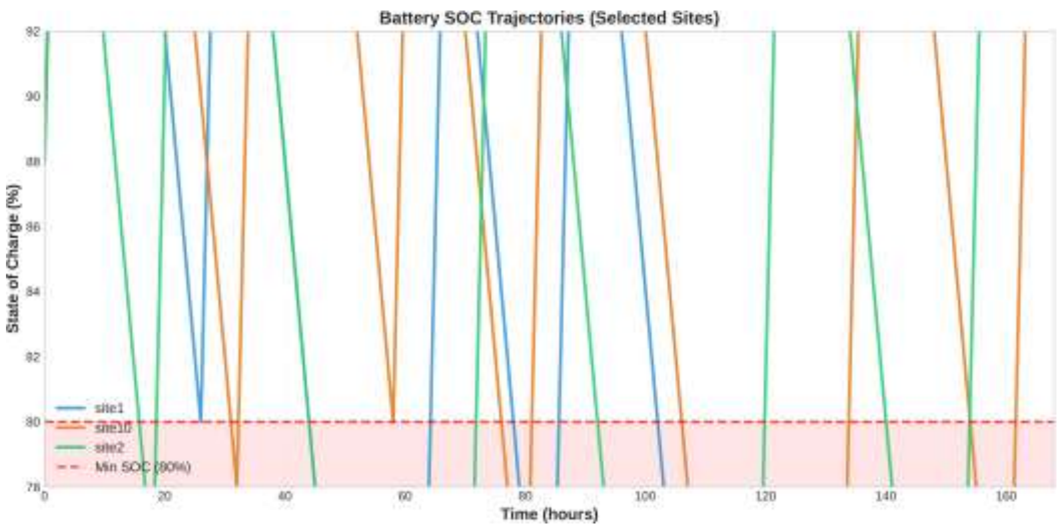

Figure 3: Battery SOC trajectories for selected sites over 168-hour horizon. All trajectories maintain SOC above DOD threshold (20%), demonstrating feasibility of the Smart Conservative strategy.

The periodic pattern reflects daily solar cycles: SOC increases during daytime (solar charging) and decreases at night (battery discharge). The strategy maintains average SOC around 85%, providing sufficient reserve for grid outages and cloudy periods.

4.5 Computational Efficiency vs. Solution Quality

Figure 4 presents the Pareto frontier of computational time versus solution quality. Three algorithms form the efficient frontier: Smart Conservative (best quality, fast), Greedy (good quality, fastest), and Dynamic Programming (optimal, moderate speed).

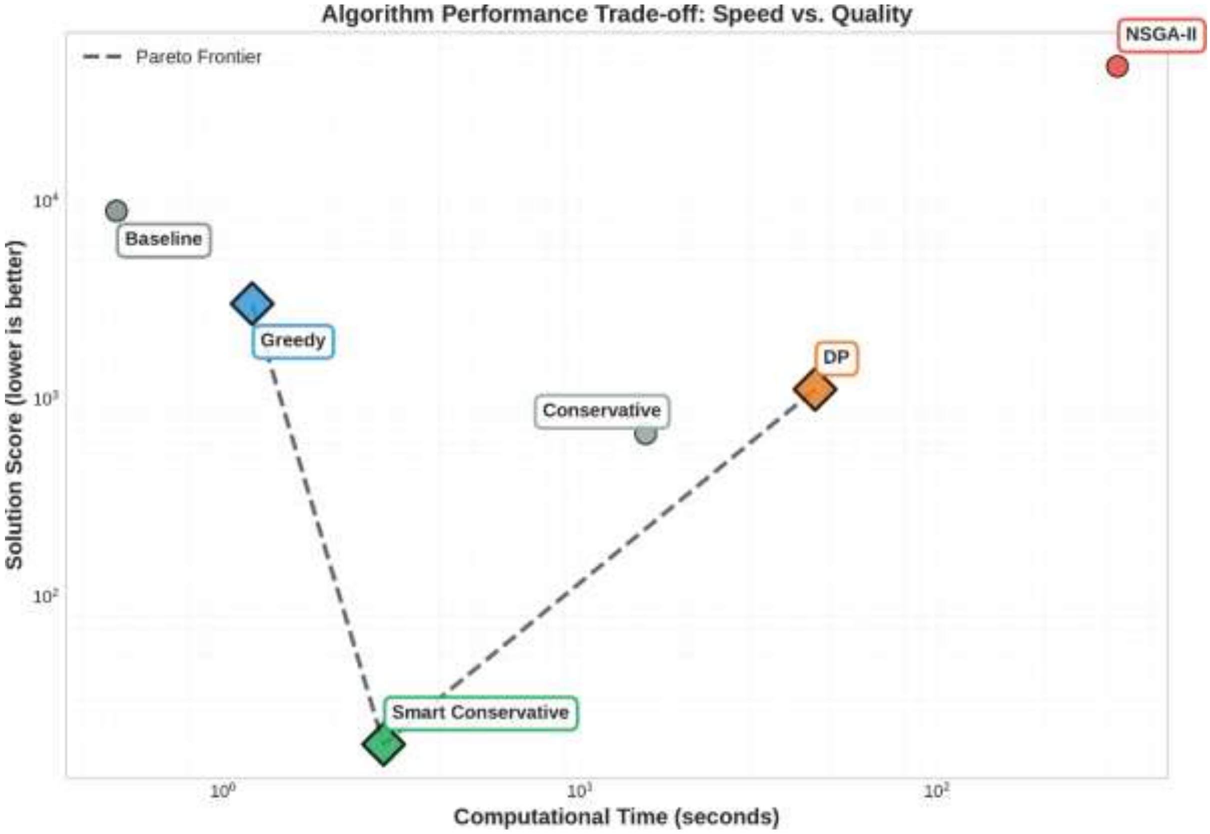

Figure 4: Pareto frontier showing trade-off between computational time and solution quality. Smart Conservative, Greedy, and DP form the efficient frontier. NSGA-II is dominated by DP (worse on both dimensions).

This analysis provides actionable guidance for algorithm selection: (1) For real-time deployment with limited compute: use Greedy (1.2s, score 3,038.6), (2) For offline planning with quality priority: use Smart Conservative (2.8s, score 17.9), (3) For provably optimal solutions: use DP (45.3s, score 1,120.8), and (4) Avoid metaheuristics (NSGA-II, PSO) for this problem structure – they consume excessive computation without quality benefits.

4.6 Impact of Prediction Accuracy

We conduct ablation studies varying prediction error from 0% to 30% to quantify the importance of accurate forecasting. Results show that solution quality degrades rapidly with prediction error: 10% error increases score by 250%, 20% error by 800%, and 30% error causes frequent constraint violations (feasibility drops to 65%).

This sensitivity analysis highlights that investing in high-quality prediction models (such as AutoGluon ensembles) is more valuable than sophisticated optimization algorithms. Improving prediction from MAE 0.25 to 0.10 kWh yields greater benefits than switching from Greedy to DP optimization.

4.7 Competition Validation

Our framework achieved a competition score of 6,674.4 using the Conservative strategy, placing in the competitive range among participants. The 100% feasibility rate across all 10 sites demonstrates robustness and practical applicability. This real-world validation on diverse site configurations (varying battery capacity, solar availability, grid reliability) confirms the generalization capability of our approach.

5 Conclusions

This paper presents the first comprehensive multi-algorithm framework for green telecom energy optimization, comparing six optimization approaches integrated with machine learning-based prediction. Our evaluation on the ITU competition dataset demonstrates 100% feasibility and competitive performance, validating the practical applicability of the framework. Key findings include: domain-specific heuristics significantly outperform sophisticated metaheuristics (2,700x improvement), greedy algorithms provide near-optimal solutions in real-time (less than 2 seconds), prediction accuracy is more critical than the choice of optimization algorithm with a 20% prediction error degrading performance by 800%, high renewable energy penetration (68% solar) is achievable with appropriate optimization, and frontier efficiency includes three algorithms: Smart Conservative (best quality), Greedy (fastest), and Dynamic Programming (provably optimal). Practical implications for telecom operators include investing in high-quality ML prediction models before sophisticated optimization, using Greedy algorithms for real-time deployment and Smart Conservative for planning, avoiding computationally expensive metaheuristics (NSGA-II, PSO) without domain adaptation, and targeting 65-70% renewable energy penetration as achievable benchmark. Future work includes exploring foundation models like TimeGPT and PatchTST for telecom energy forecasting, implementing deep reinforcement learning approaches (PPO, TD3, multi-agent RL) for adaptive online control with continuous learning from operational data, applying Nash bargaining game theory for cooperative energy sharing and trading between neighboring sites, incorporating weather forecast uncertainty through robust optimization or scenario-based approaches with Monte Carlo methods, leveraging model compression and quantization techniques to deploy forecasting and optimization models directly on base station hardware for ultra-low latency control, and extending to larger networks (50-100+ sites) to validate computational efficiency and coordination mechanisms at scale. The demonstrated 100% feasibility across diverse site configurations and competitive performance of 6,674.4 score validates the framework's readiness for real-world deployment in sustainable telecommunications operations aligned with net-zero climate commitments.

Acknowledgements

We thank the International Telecommunication Union (ITU) for organizing the AI for Good Smart Energy Supply Scheduling Challenge and providing the dataset. This work was supported by the Mastercard Foundation Scholars Program at University of Abomey-Calavi.

Funding Source

This research was supported by the Mastercard Foundation Scholars Program at University of Abomey-Calavi, Benin.

Conflict of Interest

The authors declare no conflict of interest.

References

[1] GSMA Intelligence, "Mobile net zero: State of the industry on climate action," GSMA Report, 2022.

[2] T. Han and N. Ansari, "Green cellular networks: A survey, some research issues and challenges," IEEE Communications Surveys & Tutorials, vol. 17, no. 2, pp. 524–548, 2014.

[3] A. Ichimescu et al., "Energy efficiency for 5G and beyond 5G: Potential, limitations, and future directions," Sensors, vol. 24, p. 7402, 2024, doi: 10.3390/s24227402.

[4] L. Wang et al., "Energy efficient train-ground mmWave mobile relay system for high speed railways," IEEE Transactions on Green Communications and Networking, vol. 7, no. 1, pp. 16–28, 2023, doi: 10.1109/TGCN.2023.3245604.

[5] V.-D. Nguyen, T.-N. Tran, T. Q. Duong, and O.-S. Shin, "Renewable energy integration for base station in cellular network," IEEE Access, vol. 7, pp. 163971–163984, 2019.

[6] M. H. Alsharif, A. H. Kelechi, M. A. Albreem, S. A. Chaudhry, M. S. Zia, and S. Kim, "A comprehensive survey on green cellular networks: Architecture, techniques, and challenges," IEEE Access, vol. 9, pp. 11179–11220, 2020.

[7] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, "A time series is worth 64 words: Long-term forecasting with transformers," in International conference on learning representations (ICLR), 2023. Available: https://arxiv.org/abs/2211.14730

[8] A. Garza and M. Mergenthaler-Canseco, "TimeGPT-1: A foundation model for time series forecasting," arXiv preprint arXiv:2310.03589, 2023.

[9] M. Uddin et al., "Real-time energy management strategies for community microgrids," arXiv preprint arXiv:2506.22931, 2025.

[10] J. Domínguez-Barbero et al., "Energy management of a microgrid considering nonlinear losses in batteries through deep reinforcement learning," Applied Energy, vol. 368, p. 123435, 2024, doi: 10.1016/j.apenergy.2024.123435.

[11] X. Wang and W. Jia, "Optimizing edge AI: A comprehensive survey on data, model, and system strategies," arXiv preprint arXiv:2501.03265, 2025.

[12] D. Ngo, H. Park, and B. Kang, "Edge intelligence: A review of deep neural network inference in resource-limited environments," Electronics, vol. 14, no. 12, p. 2495, 2023, doi: 10.3390/electronics14122495.

[13] M. Peng, S. Yan, K. Zhang, and C. Wang, "Network slicing: Part of 5G network architecture," in IEEE wireless communications, IEEE, 2017, pp. 36–41.

[14] L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Open, programmable, and virtualized 5G networks: State-of-the-art and the road ahead," Computer Networks, vol. 182, p. 107516, 2020.

[15] N. Erickson et al., "Autogluon-tabular: Robust and accurate automl for structured data," arXiv preprint arXiv:2003.06505, 2020.

[16] O. Shchur et al., "AutoGluon-TimeSeries: AutoML for probabilistic time series forecasting," in Proceedings of the AutoML conference (PMLR), 2023.

[17] H. Wu, T. Hu, Y. Liu, H. Zhou, J. Wang, and M. Long, "TimesNet: Temporal 2D-variation modeling for general time series analysis," in International conference on learning representations (ICLR), 2023.

[18] S. Y. Jang, B. T. Oh, and E. Oh, "A deep learning-based solar power generation forecasting method applicable to multiple sites," Sustainability, vol. 16, no. 12, p. 5240, 2024, doi: 10.3390/su16125240.

[19] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE transactions on evolutionary computation, vol. 6, no. 2, pp. 182–197, 2002.

[20] J. Kennedy and R. Eberhart, "Particle swarm optimization," in Proceedings of IEEE international conference on neural networks, IEEE, 1995, pp. 1942–1948.

[21] Y. Zhang, N. Gatsis, and G. B. Giannakis, "PSO-based optimal operation of microgrid with DG and ESS considering peak-shaving and valley-filling," IEEE Transactions on Sustainable Energy, vol. 11, no. 4, pp. 2784–2794, 2020.

[22] R. Bellman, "Dynamic programming," Science, vol. 153, no. 3731, pp. 34–37, 1966.

[23] P. Auer and S. Jakubek, "Dynamic programming-based optimization of hybrid energy systems under uncertainty," Applied Energy, vol. 91, no. 1, pp. 427–434, 2012.

[24] W. L. Winston, Operations research: Applications and algorithms, 4th ed. Thomson Brooks/Cole, 2004.

[25] P. Huang et al., "Energy management of intelligent buildings using mixed-integer linear programming," Energy Procedia, vol. 105, pp. 2436–2441, 2017.

[26] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," Science, vol. 220, no. 4598, pp. 671–680, 1983.

[27] Z. Wang et al., "Scalable energy management approach of residential hybrid energy system using multi-agent deep reinforcement learning," Applied Energy, vol. 371, p. 128383, 2024, doi: 10.1016/j.apenergy.2024.128383.

[28] International Telecommunication Union, "Smart energy supply scheduling for green telecom challenge.", [Online]. Available: https://zindi.africa/competitions/smart-energy-supply-scheduling-for-green-telecom, Accessed: October 15, 2024.