

Green AI Strategies for Reducing Energy Consumption and Carbon Emissions in Industrial Systems: Green AI: Industrial Energy and Carbon Reduction

Vaishnavi Vadane, Abhishek P. Bangar, Sachin B. Bhosale, V. A. Karad

Department of Computer Engineering
Jaihind College of Engineering, Pune, India

Abstract — Industrial production depends on electricity, process heat, motors, pumps, compressed-air networks, refrigeration, and production schedules. Artificial Intelligence (AI) can help identify avoidable energy use and support decisions that reduce associated greenhouse-gas emissions. This review examines Green AI strategies for industrial energy and carbon reduction, including energy forecasting, predictive maintenance, process optimisation, intelligent scheduling, digital twins, renewable-energy coordination, and carbon-aware operation. It also considers the environmental cost of AI itself, such as energy used for model training and inference, data storage, networking, and hardware. The reviewed literature indicates that AI-enabled monitoring, prediction, control, scheduling, and parameter optimisation are important application areas. However, reported benefits depend on data quality, system integration, baseline selection, operator acceptance, and verification in real operating conditions. The central research gap is the lack of a consistently applied, independently verifiable framework that measures net energy and carbon benefits across the industrial system while accounting for the AI system's own footprint and comparing AI with conventional engineering alternatives. This paper organises the gaps, proposes research questions, and presents a practical framework for future work.

Keywords— Green AI, industrial energy efficiency, carbon emissions, smart manufacturing, machine learning, digital twins, predictive maintenance, carbon-aware operation.

I. INTRODUCTION

Industrial facilities consume energy through production equipment and supporting utilities. Motors, pumps, fans, boilers, furnaces, refrigeration, compressed air, lighting, and buildings all contribute to total demand. Inefficient schedules, idle running, leaks, unnecessary reheating, degraded equipment, and avoidable scrap can increase energy consumption without increasing useful output. Where electricity and fuels are associated with greenhouse-gas emissions, these losses also create an avoidable carbon burden. AI techniques can extract patterns from sensor and production data, forecast demand, detect abnormal operating conditions, estimate quality, and recommend process settings. Machine learning (ML), deep learning, reinforcement learning, optimisation methods, and digital twins have been studied for monitoring, prediction, real-time control, scheduling, and parameter optimisation. Reviews such as Keramati Feyz Abadi et al. [1] and Rojek et al. [2] describe this growing range of applications.

Green AI has two connected objectives. First, AI can be used to reduce the environmental impact of industrial operations. Second, AI should be designed and operated efficiently so that

computation does not consume excessive electricity or resources. A model that saves a limited amount of factory energy but requires substantial computing energy may provide little net benefit. Therefore, the relevant outcome is not model accuracy alone, but a measured improvement in total energy and emissions while preserving safety, quality, throughput, and reliability.

This paper is a narrative review of selected sources listed in the uploaded review paper, supplemented by relevant published review evidence. It is not a new industrial experiment and does not claim exhaustive systematic coverage. The aim is to identify the unresolved problems between promising AI methods and verifiable industrial decarbonisation, then propose a research framework to address those problems.

II. REVIEW OF RELATED LITERATURE

Research in smart manufacturing has grown alongside industrial sensors, connected equipment, and Industry 4.0 systems. Sharp, Ak, and Hedberg [3] surveyed machine-learning applications in smart manufacturing. Rojek et al. [2] reviewed AI-based demand prediction and energy management, connecting energy planning with manufacturing

operations. Keramati Feyz Abadi et al. [1] reviewed 67 journal papers and grouped energy-efficiency applications into monitoring and prediction, real-time control, scheduling, and parameter optimisation. Their work also proposed an Energy Efficiency–Digital Twin framework and highlighted limitations in validation in real industrial settings.

Lee and Lin [4] reported an AI workflow evaluated across ten factories and compared AI-based energy-saving approaches with non-AI hardware upgrades. This comparison is important because energy audits, leak repair, insulation, heat recovery, efficient motors, and equipment replacement may sometimes deliver larger or more reliable benefits than a complex AI project. The International Energy Agency [5], [6] also discusses AI's potential for energy-system optimisation while cautioning that digital infrastructure has its own energy requirements and that AI is not a substitute for wider climate action.

The systematic review by Verdecchia et al. [8] examined 98 Green AI studies and reported that much of the field focused on model-footprint monitoring, hyperparameter tuning, and benchmarking. Many studies used laboratory experiments, and the authors argued that promising findings should be transferred into industrial practice. This supports the need to connect AI-efficiency methods with real industrial energy and carbon accounting.

Across these sources, a consistent theme emerges: the technical potential is well recognised, but evidence remains fragmented across methods, industrial sectors, system boundaries, and evaluation approaches. The research gap should therefore be framed not as an absence of AI applications, but as an absence of sufficiently consistent, long-term and comparable evidence about net environmental outcomes.

III. RESEARCH GAP IDENTIFICATION

The uploaded paper identifies limited validation in real industrial settings, inconsistent data quality, model generalisation problems, the AI system's energy and hardware footprint, cybersecurity and skills barriers, and the need to compare AI with conventional engineering measures. These are the source-grounded starting points for the research-gap analysis below.

1. Gap 1: Limited Long-Term Industrial Validation

Many proposed methods are demonstrated using simulations, historical datasets, or limited pilots. Such evidence can show technical promise but may not prove sustained savings in a changing factory. Machines age, product mixes change,

ambient conditions vary, and operating teams modify processes. Short trials may miss model drift, seasonal demand, maintenance events, or rebound effects. More longitudinal field studies are needed, with documented durations, operating conditions, uncertainty, and independently checked results.

2. Gap 2: Inconsistent Baselines and Measurement

Energy savings can be overstated when production volume, product quality, operating hours, weather, or equipment utilisation changes between the baseline and the intervention period. Studies may report model accuracy or simulated energy reductions without using a common normalisation method. Research needs transparent baselines, clearly defined system boundaries, energy per unit of conforming output, total consumption, and a record of the factors used to adjust comparisons.

3. Gap 3: Incomplete Accounting of AI's Own Footprint

Industrial AI can consume energy through training, inference, data transmission, storage, cloud services, cooling, and edge devices. Hardware manufacture and replacement can also matter in a life-cycle assessment. If studies count factory savings but omit the AI infrastructure, net benefit may be exaggerated. Research needs phase-specific reporting of the AI footprint and an explicit comparison between industrial energy saved and energy used by the AI system.

4. Gap 4: Energy Savings Are Not Always Equivalent to Carbon Reductions

A reduction in kilowatt-hours does not always translate into the same reduction in CO₂-equivalent emissions. Emissions depend on the energy source, grid emissions intensity, time of consumption, fuel use, and the chosen accounting boundary. Cement and other process industries may also have emissions from chemical reactions that energy-efficiency measures alone cannot remove. Research should distinguish electricity-related emissions, direct fuel emissions, and process emissions, and document the emissions factors and accounting method used.

5. Gap 5: Fragmented Optimisation Across Machines and Plant Systems

A machine-level optimiser can reduce local consumption while creating a bottleneck or increasing heating, cooling, storage, transport, or energy demand elsewhere. Existing approaches often focus on a particular machine, process, or application. A broader plant-level approach is needed to coordinate production, shared utilities, renewable generation, storage, and grid carbon intensity while respecting delivery, quality, and safety constraints.

6. Gap 6: Weak Comparison with Non-AI Alternatives

AI is not automatically the best solution. A repaired compressed-air leak, improved insulation, efficient motor, heat recovery system, or revised operating procedure may be less expensive and more dependable. Research needs fair comparisons among AI, conventional engineering interventions, and combined approaches. Such comparisons should report implementation cost, maintenance burden, energy and carbon impact, payback, reliability, and the risk of disrupting production.

7. Gap 7: Generalisation and Model Drift

A model trained on one machine, product mix, season, or factory may perform poorly in another environment. Industrial operating conditions can change after maintenance, process upgrades, or changes in raw materials. Studies should test transferability across sites and time periods, define drift-detection procedures, and explain when retraining or retirement is necessary. Performance should be assessed for both prediction quality and environmental outcomes.

8. Gap 8: Data Interoperability and Data Quality

Industrial data may be incomplete, noisy, poorly labelled, or distributed across incompatible legacy systems. Sensors can be miscalibrated, timestamps may not align, and equipment vendors may use different naming conventions. These problems can undermine models before algorithm selection even begins. Research should evaluate reusable data schemas, data-quality checks, sensor calibration procedures, and integration methods for PLC, SCADA, meter, maintenance, and production records.

9. Gap 9: Limited Accessibility for Small and Medium-Sized Enterprises

Smaller manufacturers may not have specialist AI teams, extensive sensor networks, reliable connectivity, or capital for system integration. High-cost and data-intensive solutions may therefore remain inaccessible even when potential savings exist. More research is needed on low-cost metering, compact models, reusable software, staged deployment, and practical return-on-investment assessment for resource-constrained plants.

10. Gap 10: Safety, Explainability, and Operator Acceptance

An AI recommendation that reduces energy but violates safe process limits or reduces product quality is unacceptable. Operators need to understand the basis and limits of recommendations and retain a clear route to override unsafe or impractical actions. Research should combine explainable

models, expert knowledge, safety interlocks, human approval, and rollback procedures. Evaluation should include operator trust, usability, training needs, and the effect of human intervention.

11. Gap 11: Cybersecurity and Operational Technology Risks

Connected industrial systems can create new cybersecurity risks. Data manipulation, service interruption, unauthorised access, or unsafe control actions can undermine energy performance and plant safety. Research should assess security throughout the AI deployment lifecycle, including access controls, data integrity, network segmentation, model updates, and incident response. Energy optimisation should not be evaluated separately from the resilience of the industrial control environment.

12. Gap 12: Rebound Effects and Incomplete System Boundaries

Efficiency improvements may lower the cost of production and lead to increased output, partially offsetting absolute energy or emissions reductions. A local improvement can also shift consumption to another process or supporting utility. Research should report both absolute totals and intensity per unit of good output, clearly define system boundaries, and assess whether production changes or downstream requirements offset the expected gains.

IV. PROBLEM STATEMENT

Although AI-based techniques can support industrial energy management, the environmental value of these techniques is difficult to compare because studies use different baselines, system boundaries, time periods, and performance measures. Some focus on prediction accuracy or local energy savings, while the energy and carbon footprint of the AI infrastructure may be omitted. In addition, results from simulations or short pilots may not generalise to other facilities or remain effective over time. Consequently, industries lack a consistently applied framework for deciding when AI provides a measurable net environmental benefit compared with conventional interventions.

1. Aim

To develop a practical review-based framework for evaluating Green AI strategies in industrial systems using normalised energy measurements, transparent carbon accounting, and explicit consideration of the AI system's own footprint.

2. Objectives

- Identify the main industrial AI strategies used for energy and carbon reduction.
- Categorise the research gaps in validation, measurement, carbon accounting, integration, generalisation, affordability, and governance.
- Propose a common evaluation framework that includes industrial energy savings and AI-related energy use.
- Recommend performance indicators that cover energy, emissions, cost, product quality, throughput, safety, and reliability.
- Outline future research directions for long-term, independently verifiable industrial deployment.

3. Research Questions

- RQ1: How should industrial energy savings be normalised when production volume, product mix, and operating conditions change?
- RQ2: How can the energy and carbon footprint of AI training, inference, data movement, and hardware be included in the assessment?
- RQ3: Under what conditions does AI outperform conventional engineering interventions or provide additional value when combined with them?
- RQ4: How can industrial AI be validated across different facilities and over long operating periods?
- RQ5: What low-cost, safe, and explainable deployment practices can make Green AI accessible to small and medium-sized manufacturers?

V. PROPOSED GREEN AI EVALUATION FRAMEWORK

The framework follows six stages: establish the baseline, instrument and prepare data, identify the largest opportunity, select an efficient method, pilot under supervision, and verify/report/improve. It retains the sequence proposed in the uploaded paper and expands the evaluation requirements so that future studies can be compared more consistently.

1. Stage 1: Establish the Baseline

Define the plant boundary, reporting period, energy sources, production units, emissions factors, and relevant operating conditions. Record total electricity and fuel use, energy per unit of conforming output, production volume, scrap, operating hours, and important ambient or process conditions. Document the baseline before introducing AI.

2. Stage 2: Instrument and Prepare Data

Collect meter, machine, process, maintenance, and production records. Check sensor accuracy, missing values, outliers, timestamp alignment, data retention, and access permissions. Keep a data dictionary and record any preprocessing so the study can be reproduced.

3. Stage 3: Identify the Largest Opportunity

Use an energy audit and process knowledge to locate high-consumption equipment, persistent losses, idle operation, and flexible loads. Compare the proposed AI intervention with straightforward engineering options. Prioritise opportunities with clear potential impact and manageable risk.

4. Stage 4: Select an Efficient Method

Begin with rules, statistical forecasting, or compact ML when these are sufficient. Use complex deep-learning or reinforcement-learning methods only when they deliver additional measurable value. Record model size, training frequency, inference frequency, hardware, and deployment location.

5. Stage 5: Pilot Under Supervision

Test the system on a limited process or line. Keep safety interlocks, operator approval, rollback procedures, and predetermined success criteria. Monitor energy, emissions, quality, throughput, downtime, and false alarms. Do not allow an unconstrained algorithm to modify critical process settings.

6. Stage 6: Verify, Report, and Improve

Compare results against a normalised baseline, state uncertainty, measure AI infrastructure energy, and document changes in production conditions. Continue monitoring for model drift. Retrain, modify, or retire the model when performance or net benefit deteriorates.

VI. PROPOSED EVALUATION METRICS

A credible Green AI evaluation should report multiple indicators rather than relying on one accuracy score. The minimum set should include total industrial energy, energy intensity per unit of good output, direct fuel use, emissions in kg or tonnes of CO₂-equivalent, AI infrastructure energy, product quality, scrap, throughput, downtime, cost, and safety-related incidents.

1. Net Energy Benefit

A simple screening measure is:

$$\text{Net Energy Benefit} = \text{Industrial Energy Saved} - \text{Energy Used by the AI System (1)}$$

This is useful as an initial check, but it is not a full life-cycle assessment. The boundary should explain whether training, inference, cloud services, networking, storage, cooling, and edge-device energy are included. If hardware manufacture and replacement are material to the case, those impacts should also be considered.

2. Carbon Accounting

Carbon emissions should be calculated using documented emissions factors and a clearly stated boundary. Electricity consumption can be multiplied by an appropriate grid emissions factor; direct fuel quantities should be multiplied by fuel-specific factors. Where time-varying grid emissions data are available and reliable, the analysis can examine whether shifting flexible loads changes emissions. Location-based and market-based accounting approaches should not be mixed without explanation.

3. Normalisation and Counterfactual

Before-and-after comparisons should account for changes in production volume, product mix, operating hours, weather, and quality. Where feasible, researchers should use a control line, matched operating periods, or a documented statistical baseline to estimate what energy use would have been without the intervention. Report uncertainty and distinguish directly measured savings from simulation-based estimates.

4. Comparison Table

Table I: Comparison of Green AI Strategies, Data Inputs, and Outcomes

Strategy	Data Inputs	Outcomes
Demand forecasting	Meter, schedule, and production data	Forecast error, peak demand, and energy intensity
Predictive maintenance	Vibration, temperature, current, and service records	Fault detection, downtime, and energy waste avoided
Process optimisation	Setpoints, energy, throughput, and quality	Energy per good product and scrap
Smart scheduling	Orders, machine states, tariffs, and carbon signals	Idle time and flexible-load shifting

Strategy	Data Inputs	Outcomes
Digital twins	Process models and live data	Validated performance of alternative settings
Renewable coordination	Demand forecasts, solar output, and storage state	Renewable utilisation and emissions

VII. INDUSTRIAL APPLICATIONS AND VALIDATION PRIORITIES

1. Cement and Materials Processing

AI can forecast electricity demand, monitor grinding mills, detect equipment anomalies, and support kiln parameter optimisation. Validation must distinguish combustion emissions, electricity-related emissions, and process emissions from chemical reactions. An energy-efficiency improvement should not be described as eliminating all process emissions.

2. Food, Beverage, and Cold-Chain Facilities

AI can forecast cooling loads, detect refrigeration faults, optimise compressor sequencing, and coordinate production with storage conditions. The system must preserve hygiene, food safety, product quality, and required temperature limits. Research should report both energy performance and compliance with safety constraints.

3. Metalworking and Discrete Manufacturing

Machine tools, robotic cells, compressed air, and heat treatment can consume energy during productive and idle periods. AI-based scheduling may reduce idle running and repeated warm-up cycles. Process models can support cutting-parameter selection while balancing energy, tool wear, cycle time, and quality.

4. Chemical and Process Industries

AI can support soft sensing, fault detection, heat and steam demand prediction, and operating-condition optimisation across pumps, reactors, heat exchangers, and distillation columns. Because disturbances can create safety risks, models require validated constraints, human oversight, cybersecurity, and integration with existing process-safety systems.

5. Validation Protocol

Each industrial study should describe the site and process, equipment boundary, baseline period, data sources, model, deployment hardware, operator role, safety limits, and comparison method. It should report the duration of the pilot and whether results are measured, estimated, or simulated.

Independent review and replication across more than one operating condition would strengthen confidence in the findings.

VIII. DISCUSSION

The reviewed sources suggest that industrial AI offers a portfolio of approaches rather than a single universal solution. Monitoring and forecasting help reveal patterns; predictive maintenance targets degradation; process optimisation searches for efficient operating points; scheduling coordinates demand; digital twins support controlled testing; and renewable coordination can align demand with cleaner energy. The main unresolved issue is how to prove that these individual benefits create a net, durable improvement at the plant level.

A strong research design should compare AI with both the baseline and a credible non-AI alternative. It should measure the footprint of the AI solution and consider system-wide effects. For example, a local reduction in electricity may be offset by increased scrap, a new bottleneck, extra cooling, or additional production. Likewise, a model with excellent predictive accuracy may not change decisions enough to deliver meaningful savings. Environmental outcomes and operational constraints must therefore be evaluated together.

Future Research Directions

- Longitudinal industrial pilots: evaluate performance over seasons, maintenance cycles, and changing production mixes.
- Standardised reporting: publish common definitions for system boundaries, baselines, emissions factors, uncertainty, and AI energy use.
- Physics-informed ML: combine process knowledge with operational data to improve robustness and reduce dependence on large datasets.
- Explainable and human-centred AI: make recommendations understandable to engineers and operators, with safe override paths.
- Federated and privacy-preserving learning: explore cross-site collaboration without centralising sensitive plant data.
- Digital twins with carbon-aware optimisation: test alternative production and energy scenarios before controlled deployment.
- Efficient edge AI: measure whether compact models and local inference reduce latency, data transfer, and total footprint.
- SME-oriented solutions: develop low-cost metering, open interfaces, reusable models, and staged adoption plans.

- Comparative economics: report capital cost, operating cost, maintenance, payback, and avoided emissions alongside technical metrics.
- Independent replication: test whether reported gains can be reproduced by other facilities or research teams.

IX. CONCLUSION

Green AI can support industrial energy and carbon reduction through energy monitoring, demand forecasting, predictive maintenance, process optimisation, production scheduling, digital twins, and renewable-energy coordination. The uploaded review paper and the cited literature also show that the benefits are conditional: data quality, system integration, safety, operator acceptance, and real-world verification matter. The central research gap is the lack of a consistently applied and independently verifiable method for measuring net environmental benefit across industrial systems while accounting for the AI system's own footprint and comparing it with conventional alternatives.

Future studies should establish transparent and normalised baselines, distinguish energy savings from carbon reductions, include AI infrastructure energy, evaluate changes in quality and throughput, and monitor performance over time. The proposed six-stage framework offers a practical path from energy audits to supervised pilots and ongoing verification. AI should be treated as an enabling technology that complements proven engineering improvements, not as a standalone climate solution.

REFERENCES

1. M. M. K. Feyz Abadi, C. Liu, M. Zhang, Y. Hu, and Y. Xu, "Leveraging AI for energy-efficient manufacturing systems: Review and future perspectives," *Journal of Manufacturing Systems*, vol. 78, pp. 153–177, 2025. doi: 10.1016/j.jmsy.2024.11.017.
2. I. Rojek, D. Mikołajewski, A. Mroziński, and M. Macko, "Green energy management in manufacturing based on demand prediction by artificial intelligence—A review," *Electronics*, vol. 13, no. 16, art. 3338, 2024. doi: 10.3390/electronics13163338.
3. M. Sharp, R. Ak, and T. D. Hedberg Jr., "A survey of the advancing use and development of machine learning in smart manufacturing," *Journal of Manufacturing Systems*, vol. 48, 2018. doi: 10.1016/j.jmsy.2018.02.004.
4. D. Lee and C. Lin, "Universal artificial intelligence workflow for factory energy saving: Ten case studies,"

- Journal of Cleaner Production*, vol. 468, art. 143049, 2024. doi: 10.1016/j.jclepro.2024.143049.
5. International Energy Agency, “AI for energy optimisation and innovation,” *Energy and AI*, Paris, 2025. <https://www.iea.org/reports/energy-and-ai/ai-for-energy-optimisation-and-innovation>
 6. International Energy Agency, “AI and climate change,” *Energy and AI*, Paris, 2025. <https://www.iea.org/reports/energy-and-ai/ai-and-climate-change>
 7. IEA Industrial Energy-Related Technologies and Systems, “Digitalization, Artificial Intelligence and Related Technologies for Energy Efficiency and GHG Emissions Reduction in Industry,” IETS Task XVIII. <https://iea-industry.org/tasks/digitalization-artificial-intelligence-and-related-technologies-for-energy-efficiency-and-ghg-emissions-reduction-in-industry/>
 8. D. Verdecchia et al., “A systematic review of Green AI,” *WIREs Data Mining and Knowledge Discovery*, vol. 13, no. 4, e1507, 2023. doi: 10.1002/widm.1507.
 9. O. Oguntola, K. Boakye, and S. Simske, “Towards leveraging artificial intelligence for sustainable cement manufacturing: A systematic review of AI applications in electrical energy consumption optimization,” *Sustainability*, vol. 16, no. 11, art. 4798, 2024. doi: 10.3390/su16114798.
 10. R. X. Gao, J. Krüger, M. Merklein, H.-S. Möhring, and J. Váncza, “Artificial intelligence in manufacturing: State of the art, perspectives, and future directions,” *CIRP Annals*, vol. 73, no. 2, pp. 723–749, 2024. doi: 10.1016/j.cirp.2024.04.101.
 11. D. Verdecchia et al. (2023) is cited for Green AI lifecycle and evaluation concerns; this review uses it as contextual evidence and does not claim to reproduce its full systematic-review protocol.
 12. M. Rojahn and M. Grum, “Green AI: A systematic review and meta-analysis of its definitions, lifecycle models, hardware and measurement attempts,” *arXiv preprint*, 2025. <https://arxiv.org/abs/2511.07090>. (Additional contextual source consulted online; preprint status should be considered.)