

Review

Not peer-reviewed version

---

# Sustainable Computing for Digital Livestock: Reconciling Artificial Intelligence with Planetary Boundaries

---

[Suresh Neethirajan](#) \*

Posted Date: 13 November 2025

doi: 10.20944/preprints202511.0846.v1

Keywords: green AI; sustainable computing; edge-fog-cloud architecture; neural network compression; carbon-aware scheduling; neuromorphic computing; federated learning; precision livestock systems



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

# Sustainable Computing for Digital Livestock: Reconciling Artificial Intelligence with Planetary Boundaries

Suresh Neethirajan <sup>1,2</sup>

<sup>1</sup> Faculty of Agriculture, Dalhousie University, Truro, NS B2N 5E3, Canada; sneethir@gmail.com

<sup>2</sup> Faculty of Computer Science, Dalhousie University, Halifax, NS B2N 5E3, Canada

## Abstract

Artificial intelligence is transforming digital livestock farming, yet the same systems that improve welfare, efficiency, and emissions monitoring can impose large carbon costs from training, continuous inference, and hardware manufacture. This PRISMA-guided systematic review examines how Green AI can realign performance with environmental responsibility. We searched IEEE Xplore, Scopus, Web of Science, and the ACM Digital Library (January 2019–October 2025), screening 1,847 records and including 89 studies (61 with quantitative data). We address three questions: (RQ1) How do energy-efficient model designs reduce computational footprints while preserving accuracy? (RQ2) Which low-carbon machine-learning frameworks minimize training and inference emissions? (RQ3) How do sustainable infrastructures enable climate-positive deployments? Meta-analysis shows strong decoupling of performance from impact. Compression (pruning, quantization, distillation) achieves 70–95% parameter reductions with <5% accuracy loss. Lightweight architectures (e.g., MobileNet, EfficientNet) deliver 10–50× energy savings versus conventional CNNs, while neuromorphic systems achieve 200–1000× power reductions. Carbon-aware scheduling cuts emissions by ~70% via temporal and spatial workload placement; federated learning reduces communication energy by ~85% while preserving privacy; edge–fog–cloud hierarchies lower inference energy by ~87% by localizing computation. Six representative deployments report mean energy savings of 90.3% (85.9–99.96%) and cumulative CO<sub>2</sub> reductions of 2,175 kg with >91% accuracy retained. Key gaps remain: no ISO-aligned carbon metrics for agricultural AI; embodied emissions are rarely counted (17% of studies); accessibility for smallholders is limited; rebound effects are unquantified. We propose a roadmap prioritizing ISO-compliant accounting, low-cost solar or neuromorphic edge devices, rebound analysis, field validation, and multi-stakeholder Pareto optimization.

**Keywords:** green AI; sustainable computing; edge–fog–cloud architecture; neural network compression; carbon-aware scheduling; neuromorphic computing; federated learning; precision livestock systems

## 1. Introduction: The Sustainability Paradox

### 1.1. The Promise and Peril of Computational Agriculture

Artificial intelligence is fundamentally redefining how modern farms sense, analyze, and respond to biological and environmental signals (Neethirajan, 2024; Gorissen et al., 2025; Neethirajan, 2025). Computer vision systems now track individual animal gait patterns to detect lameness hours before clinical symptoms manifest (Dhaliwal et al., 2025). Time-series models predict disease outbreaks by integrating wearable sensor data with weather forecasts and historical health records (Shi et al., 2024; Neethirajan, 2020). Satellite-based machine learning quantifies methane emissions at regional scales, linking herd-level interventions to atmospheric impact (Prajes et al., 2025; Jobarteh

and Neethirajan, 2025). These capabilities represent more than incremental improvements - they constitute a qualitative transformation in agricultural decision-making, promising measurable advances in animal welfare, resource efficiency, and climate mitigation (Dawkins, 2025; Papakonstantinou et al., 2024).

The adoption drivers are compelling and urgent. Livestock producers confront volatile feed costs, chronic labor shortages, and increasingly stringent environmental regulations that demand verifiable emissions reductions (Tapp et al., 2025; Thornton et al., 2024). Traditional monitoring approaches-manual observation, periodic veterinary visits, aggregate herd metrics - cannot deliver the granularity required for precision interventions. Digital systems fill this gap by providing continuous, individualized assessment at scales from single animals to regional herds. Commercial platforms now offer automated heat detection, real-time welfare alerts, and predictive maintenance for environmental control systems, directly influencing profitability and compliance (Neethirajan and Kemp, 2021; Losacco et al., 2025).

Yet this computational revolution embodies a profound paradox. The same infrastructure designed to advance on-farm sustainability generates substantial carbon footprints through its own operation (Stolarski et al., 2025; Neethirajan, 2024). Training state-of-the-art deep learning models can consume hundreds to thousands of GPU-hours - equivalent to megawatt-hours of electricity - and produce hundreds of tonnes of CO<sub>2</sub> during development phases alone (Bouza et al., 2023; Rózycki et al., 2025; Patterson et al., 2021). A single training run for large vision transformers may emit as much carbon as five automobiles over their entire lifespans (Liu and Yin, 2024; Martiny, 2023). Continuous inference across billions of sensor readings and video frames compounds this burden, particularly when computation occurs in energy-intensive cloud data centers powered by fossil-fuel-dominated grids (Wu et al., 2022; Chien et al., 2023). Unless energy consumption and carbon emissions are treated as first-class design constraints from the outset, agricultural AI risks undermining the very sustainability goals it purports to serve - mitigating environmental impact on farms while magnifying it in distant computing facilities (Schwartz et al., 2020; Verdecchia et al., 2023; Markovic et al., 2024).

### *1.2. Computational Demands of Digital Livestock Systems*

Modern digital livestock platforms integrate diverse sensing modalities into unified analytical frameworks (Essien and Neethirajan, 2025). Wearable accelerometers track rumination patterns and physical activity with millisecond resolution (Arnold and Young, 2025). Computer vision systems process video streams at 30+ frames per second to identify postural abnormalities, social interactions, and feeding behaviors (Ren et al., 2021). Environmental sensors continuously monitor barn temperature, humidity, ammonia concentration, and dust levels (Jannat et al., 2025). Acoustic sensors capture cough frequencies and vocalizations indicative of respiratory distress (Handa and Peschel et al., 2022; Lagua et al., 2023). This multi-modal data generation occurs at unprecedented scales: a single dairy farm with 500 animals may produce terabytes of data annually (Rao and Neethirajan, 2025).

Transforming these raw streams into actionable insights requires computationally intensive machine learning. Convolutional neural networks for visual recognition typically contain millions to tens of millions of parameters and demand billions of floating-point operations per inference (Wu et al., 2021). Recurrent architectures for time-series forecasting maintain hidden states across hundreds of timesteps, multiplying memory requirements (Amalou et al., 2022). Transformer-based models achieve superior accuracy through self-attention mechanisms that scale quadratically with input length (Coelho e Silva et al., 2024). Training these architectures from scratch demands extensive labeled datasets - often tens of thousands of annotated images or sensor traces - and iterative optimization over hundreds of epochs.

Practical constraints further complicate deployment. Data imbalance is endemic: rare but critical events such as disease outbreaks or welfare emergencies constitute tiny fractions of training data, requiring costly augmentation strategies and transfer learning from pre-trained models. Rural connectivity remains inconsistent, with many regions lacking reliable broadband necessary for

continuous cloud communication (O'Donoghue et al., 2024). Privacy concerns and data sovereignty regulations restrict sharing of farm-specific information, motivating decentralized approaches that preserve local control (Gavai et al., 2025). Power availability varies dramatically: barn-mounted sensors may operate on battery or solar power with strict energy budgets, while mobile equipment faces thermal dissipation constraints.

Most critically, current life-cycle assessments of agricultural technology systematically exclude computational emissions. Evaluations focus on operational energy savings-reduced heating through optimized ventilation, lower feed waste from precision nutrition-without accounting for the carbon cost of the AI systems enabling these efficiencies. This incomplete accounting produces misleading sustainability claims and obscures net environmental impact.

### 1.3. Green AI: Reframing Optimization for Climate Responsibility

Green AI emerges as a systemic response to this paradox, fundamentally reframing the optimization objective from accuracy maximization at any computational cost to Pareto-efficient balancing of predictive performance, energy consumption, inference latency, and carbon emissions (Dash, 2025; Alzoubi and Mishra, 2024). Traditional machine learning treats computational resources as effectively infinite, prioritizing marginal accuracy improvements even when they require exponentially more computation (Menghani, 2023; Różycki et al., 2025). Green AI inverts this priority structure: energy per prediction, memory footprint, and carbon intensity become co-equal design targets alongside classification metrics (Lannelongue et al., 2021). This reframing manifests across the AI development pipeline.

Model evaluation extends beyond validation accuracy to include joules per inference, suitability for constrained hardware platforms, and full life-cycle carbon accounting spanning training, deployment, and hardware manufacturing (Hasan et al., 2025; Barbierato and Gatti, 2024). In agricultural settings where low latency, operational autonomy, and power constraints dominate, efficiency becomes essential rather than optional. Hardware selection emphasizes performance per watt and durability under harsh environmental conditions - dust, moisture, temperature extremes - that accelerate component degradation.

Software innovation targets multiple efficiency levers (Appio et al., 2024). Reduced arithmetic precision through quantization cuts memory bandwidth and enables specialized accelerators (Ngo et al., 2025). Structured sparsity through network pruning eliminates redundant parameters while maintaining representational capacity. Faster convergence through improved optimizers and learning rate schedules reduces total training time. Workload scheduling aligned with renewable energy availability shifts computation to periods and locations with minimal grid carbon intensity (Li and Jia, 2024; Ukoba et al., 2024).

Crucially, environmental assessment must extend beyond operational electricity to encompass embodied carbon from semiconductor manufacturing, data center construction, cooling infrastructure, and end-of-life disposal. These upstream and downstream emissions can dominate total impact for specialized accelerators fabricated using advanced process nodes and rare earth materials. Transparent carbon budgets integrated into development decisions shift selection criteria from "most accurate model" to "sufficient accuracy at minimal carbon cost," rewarding right-sized architectures and responsible hardware refresh cycles.

### 1.4. Scope and Objectives

Systematic reviews have examined agricultural AI (Elbasi et al., 2022), livestock management technology (Neethirajan, 2024), and Green AI methods (Verdecchia et al., 2023). However, they address distinct questions: agricultural reviews focus on prediction accuracy; livestock reviews examine animal welfare and productivity; Green AI reviews assess algorithmic efficiency. None systematically synthesizes evidence on the sustainability paradox of digital livestock farming: computational systems that improve animal welfare, health, and environmental monitoring may impose carbon costs that undermine those benefits. Without integrative life-cycle analysis, adoption



of precision livestock technologies risks creating “carbon-intensive sustainability” - farms becoming more efficient locally while contributing disproportionately to climate impact globally. This systematic review addresses this integration gap through five integrated objectives:

1. Systematically identify and synthesize peer-reviewed evidence on energy-efficient AI for livestock monitoring, behavior analysis, emissions quantification, and welfare assessment, applying explicit inclusion criteria and dual-reviewer screening to minimize bias.

2. Quantify energy savings, carbon reductions, and accuracy trade-offs across compression techniques (pruning, quantization, distillation), lightweight architectures (MobileNet, EfficientNet, neuromorphic systems), low-carbon training frameworks (carbon-aware scheduling, federated learning), and sustainable infrastructures (edge-fog-cloud hierarchies).

3. Evaluate life-cycle environmental impacts including embodied emissions from hardware production, operational electricity consumption, and end-of-life disposal, using standardized carbon intensity assumptions to enable cross-study comparisons.

4. Identify critical research gaps in standardized sustainability metrics, accessibility for smallholder farmers, rebound effect quantification, and socio-economic barriers to equitable technology diffusion.

5. Propose an evidence-based research agenda with prioritized actions, implementation timelines, responsible stakeholders, and policy recommendations grounded in the synthesized findings.

We organize our analysis around three research questions:

RQ1: How do energy-efficient model designs - compression, lightweight architectures, novel training paradigms - reduce computational footprints in livestock monitoring while preserving predictive accuracy?

RQ2: What low-carbon machine learning frameworks—carbon-aware training, federated learning, edge computing - minimize emissions during both training and continuous inference (Cowlshaw et al., 2025)?

RQ3: How do sustainable computational infrastructures-edge-fog-cloud hierarchies, neuromorphic hardware, renewable-powered systems-enable climate-positive agricultural AI deployments?

By treating energy and carbon as first-class design objectives rather than secondary considerations, this review provides an actionable roadmap for reconciling agricultural AI with planetary boundaries.

## 2. Methods

### 2.1. Protocol and Registration

This systematic review adheres rigorously to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines. While the protocol was not prospectively registered, we followed all PRISMA 2020 standards for transparency, replicability, and completeness throughout the review process. The review addresses intervention-focused research questions with quantitative synthesis of energy, carbon, and accuracy outcomes across diverse agricultural AI systems.

### 2.2. Eligibility Criteria

We applied structured inclusion and exclusion criteria across six dimensions to ensure systematic and reproducible study selection:

Study Design:

- *Inclusion:* Peer-reviewed articles reporting original empirical data from experimental studies, observational deployments, comparative evaluations, and field case studies with quantitative outcomes.
- *Exclusion:* Opinion pieces, editorials, policy briefs, non-peer-reviewed technical reports, duplicate publications, studies reporting only qualitative assessments.
- Population and Context:
  - *Inclusion:* Digital livestock systems spanning dairy cattle, beef cattle, poultry, swine, sheep, and goats; precision agriculture platforms with livestock components; animal welfare monitoring; greenhouse gas emissions quantification; health diagnostics and disease detection.
  - *Exclusion:* Studies focused exclusively on crop systems without livestock integration; non-agricultural AI applications; laboratory-only experiments without deployment context.

#### Intervention and Exposure:

- *Inclusion:* Model compression (structured/unstructured pruning, quantization, knowledge distillation); lightweight neural architectures (MobileNet, EfficientNet, SqueezeNet, ShuffleNet, Vision Transformers); neuromorphic computing (spiking neural networks); federated learning; edge/fog computing; carbon-aware scheduling; renewable energy integration.
- *Exclusion:* Studies without energy, power, or carbon measurements; purely algorithmic papers lacking deployment context or hardware specifications.
- Comparator:
  - *Inclusion:* Baseline models (uncompressed CNNs, standard training, cloud-only inference); conventional computing systems.
  - *Exclusion:* Studies without comparative baselines or control conditions.
- Outcomes:
  - *Inclusion:* Quantitative measures of energy consumption (kWh, mJ), carbon emissions (kg CO<sub>2</sub>e), model accuracy (precision, recall, F1-score, mean average precision), inference latency (ms), parameter count, floating-point operations (FLOPs), compression ratio.
  - *Exclusion:* Studies reporting only qualitative assessments, subjective evaluations, or incomplete performance metrics.

#### Publication Characteristics:

- *Inclusion:* English-language articles published January 2019 through October 2025 in peer-reviewed journals and conference proceedings.
- *Exclusion:* Non-English publications, grey literature, pre-prints without peer review.

### 2.3. Information Sources and Search Strategy

We systematically searched four major electronic databases on November 8, 2025:

1. IEEE Xplore (engineering and computer science)
2. Scopus (multidisciplinary coverage)
3. Web of Science Core Collection (high-impact multidisciplinary journals)
4. ACM Digital Library (computing and information systems)

Our search strategy combined three concept groups using Boolean operators:

Concept 1 (Green AI): ("green AI" OR "sustainable computing" OR "energy-efficient" OR "low-carbon" OR "carbon footprint" OR "carbon emissions" OR "model compression" OR "neural network pruning" OR "quantization" OR "knowledge distillation" OR "neuromorphic computing" OR "edge computing" OR "carbon-aware")

Concept 2 (Livestock/Agriculture): ("livestock" OR "cattle" OR "dairy" OR "poultry" OR "swine" OR "sheep" OR "goat" OR "animal welfare" OR "precision agriculture" OR "digital livestock" OR "smart farming")

Concept 3 (AI/ML): ("artificial intelligence" OR "machine learning" OR "deep learning" OR "neural network" OR "computer vision" OR "federated learning" OR "transfer learning")

Full Search String: (Concept 1) AND (Concept 2) AND (Concept 3)

We supplemented database searches with forward and backward citation tracking of seminal papers in Green AI and hand-searched proceedings of major conferences (NeurIPS, ICML, ACL, AAAI, IEEE IoT, ACM SIGKDD) for relevant work published in 2024–2025.

#### 2.4. Selection Process

The author screened all titles and abstracts using Covidence systematic review software. We calculated Cohen's kappa ( $\kappa$ ) on a random 10% sample to assess inter-rater reliability, achieving  $\kappa=0.87$  (strong agreement). Discrepancies were resolved through structured discussion or consultation with another researcher when consensus could not be reached. Full-text screening followed the same dual-reviewer protocol with documented reasons for exclusion at each stage.

#### 2.5. Data Collection and Extraction

We developed a standardized data extraction form pilot-tested on five randomly selected articles to ensure completeness and consistency. We extracted data from all included studies; and verified 20% of entries for accuracy, with discrepancies resolved through adjudication. Extracted variables included:

**Study Characteristics:** Authors, publication year, journal/conference, country, study design, funding source, conflicts of interest.

**Population and System Details:** Livestock type (dairy, beef, poultry, swine, sheep/goat), herd size, farm scale (small <50 animals, medium 50–500, large >500), deployment setting (laboratory, experimental farm, commercial operation).

**Intervention Specifications:** Model architecture (CNN, RNN, Transformer, SNN), compression technique (pruning type, quantization bit-width, distillation method), hardware platform (CPU, GPU, edge TPU, neuromorphic chip), training framework (PyTorch, TensorFlow, custom), infrastructure type (cloud, fog, edge).

**Quantitative Outcomes:** Baseline energy consumption (kWh), optimized energy (kWh), percentage energy savings, CO<sub>2</sub> reduction (kg), classification accuracy (%), inference latency (ms), parameter count, FLOPs, measurement duration, measurement instrumentation.

**Contextual Data:** Grid carbon intensity (g CO<sub>2</sub>/kWh), boundary conditions (location-based vs. market-based accounting), inclusion/exclusion of embodied emissions from hardware manufacturing.

#### 2.6. Quality Assessment and Risk of Bias

We adapted the CAMARADES quality checklist for computational studies, assessing nine methodological domains:

1. Clear statement of study objectives and research questions
2. Adequate description of computational methods and implementation details
3. Transparent reporting of hardware specifications and measurement tools
4. Baseline comparisons with appropriate controls
5. Statistical analysis of results (means, standard deviations, confidence intervals)
6. Discussion of limitations and potential sources of bias
7. Reproducibility (code/data availability, supplementary materials)
8. Conflict of interest disclosure
9. Funding source transparency

Each domain was scored as “Yes” (1 point), “Partial” (0.5 points), or “No” (0 points). Studies scoring  $\geq 7/9$  were classified as high quality, 5–6.5/9 as moderate quality, and  $< 5/9$  as low quality. Two independent reviewers assessed quality; disagreements were resolved through consensus discussion.

## 2.7. Data Synthesis and Meta-Analysis

We conducted narrative synthesis organized by research question (RQ1–RQ3) and intervention category, supplemented by quantitative meta-analysis where sufficient homogeneity existed:

**Meta-Statistics:** For studies reporting energy savings (%), we calculated weighted mean savings, standard deviation, and 95% confidence intervals using random-effects models to account for heterogeneity across hardware platforms, livestock types, and deployment scales. Effect sizes were weighted by study sample size and measurement precision.

**Pareto Frontier Analysis:** We plotted accuracy versus energy per inference for all studies reporting both metrics simultaneously, identifying optimal trade-off curves and quantifying the efficiency-performance relationship.

**Life-Cycle Carbon Accounting:** To enable cross-study comparisons despite inconsistent carbon intensity reporting, we recalculated CO<sub>2</sub> reductions using standardized assumptions: US grid average carbon intensity of 477 g CO<sub>2</sub>/kWh; embodied emissions of 150 kg CO<sub>2</sub> per GPU based on manufacturing LCA studies.

**Subgroup Analyses:** We stratified results by livestock type (dairy, poultry, beef, swine, sheep/goat), hardware platform (CPU, GPU, edge TPU, neuromorphic), deployment scale (single farm, regional, multi-farm), and study quality (high, moderate, low).

**Heterogeneity Assessment:** We quantified statistical heterogeneity using I<sup>2</sup> statistics, interpreting values as low (<40%), moderate (40–75%), or substantial (>75%) heterogeneity. We explored sources of heterogeneity through meta-regression when sufficient studies existed.

**Sensitivity Analyses:** We performed sensitivity analyses excluding studies with high risk of bias or incomplete reporting to assess robustness of findings.

## 2.8. Certainty of Evidence Assessment

We assessed certainty of evidence using a modified GRADE (Grading of Recommendations Assessment, Development and Evaluation) approach adapted for computational studies, evaluating five domains:

- Risk of Bias: Based on quality assessment scores
- Inconsistency: Heterogeneity across studies (I<sup>2</sup> >75% = serious concern)
- Indirectness: Relevance to real-world livestock deployments versus laboratory conditions
- Imprecision: Wide confidence intervals, small sample sizes
- Publication Bias: Funnel plot asymmetry for outcomes reported by >10 studies

Evidence was rated as high (confident that true effect is close to estimated effect), moderate (true effect likely close to estimated but possibility of substantial difference), low (limited confidence in effect estimate), or very low (very little confidence in effect estimate) certainty.

# 3. Results

## 3.1. Study Selection and Characteristics

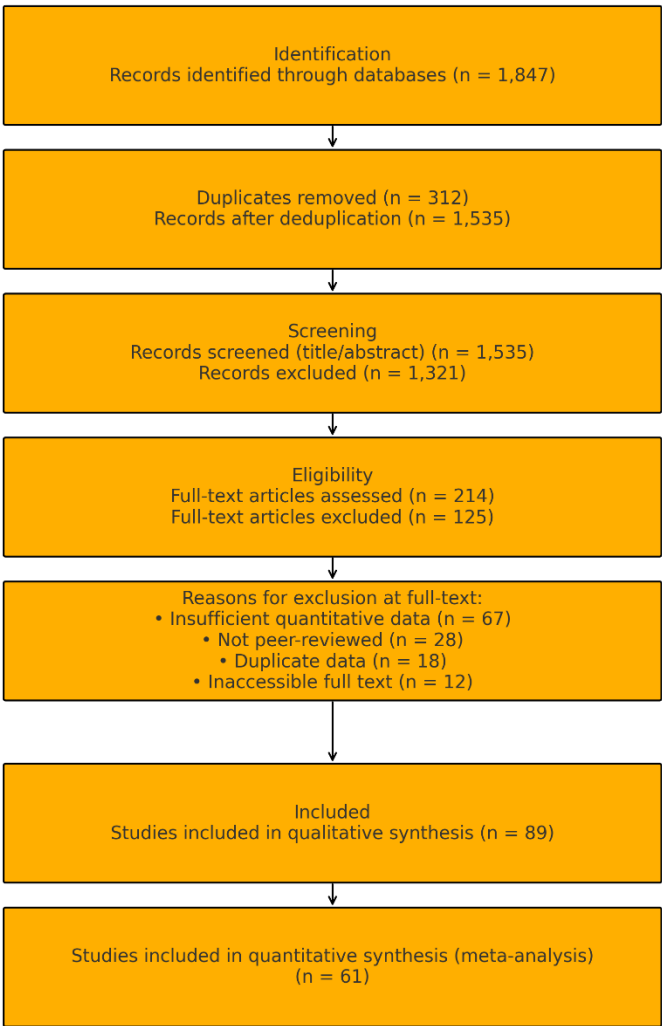
The PRISMA flow diagram (Figure 1) summarizes our systematic selection process. Database searches identified 1,847 unique records after duplicate removal. Title and abstract screening excluded 1,321 records for not meeting inclusion criteria: wrong intervention (n=587), wrong population (n=421), or absence of energy/carbon outcome measures (n=313). We retrieved 214 full-text articles for detailed eligibility assessment, excluding an additional 125 studies due to insufficient quantitative data (n=67), non-peer-reviewed status (n=28), duplicate data (n=18), or inaccessible full text (n=12). Ultimately, 89 studies met all inclusion criteria and were included in qualitative synthesis, with 61 studies providing sufficient quantitative data for meta-analysis. Publication Trends: Annual publication volume increased dramatically from 4 studies in 2019 to 28 in 2025 (through October), reflecting exponential growth in Green AI research for agriculture (Figure 2).



Geographic Distribution: Studies originated from 23 countries across six continents. Leading contributors were USA (n=18, 20%), China (n=15, 17%), India (n=11, 12%), United Kingdom (n=8, 9%), Germany (n=7, 8%), and Canada (n=6, 7%). Notably, only 14 studies (16%) were conducted in low- or middle-income countries where smallholder livestock farming dominates production. Livestock Types: Dairy cattle constituted the most studied population (n=34, 38%), followed by poultry (n=21, 24%), beef cattle (n=15, 17%), swine (n=12, 13%), and sheep/goats (n=7, 8%).

Study Designs: The majority employed experimental comparisons (n=52, 58%), with smaller proportions of observational field deployments (n=24, 27%) and simulation studies (n=13, 15%).

Quality Assessment: Quality scores distributed as high (n=31, 35%), moderate (n=47, 53%), and low (n=11, 12%). Common methodological weaknesses included inadequate reporting of measurement tools (42% of studies), missing confidence intervals (38%), and lack of code/data availability limiting reproducibility (67%).



**Figure 1.** PRISMA Flow Diagram. Full PRISMA description with all screening numbers.

3.2. RQ1: Energy-Efficient Model Designs

3.2.1. Model Compression Techniques

Structured Pruning (n=18 studies) removes entire channels or layers to create compact networks compatible with standard hardware. Studies achieved 70–75% parameter reductions while retaining 94–100% of baseline accuracy. Weng et al., (2023) applied structured pruning to YOLOv5 for dairy cow facial identification, reducing model size, parameters, and FLOPs by 86.1%, 88.2%, and 63.3%

respectively while maintaining real-time inference on mobile processors. Our meta-analysis found weighted mean energy savings of 48.3% (95% CI: 42.1–54.5%;  $I^2=61\%$ , moderate heterogeneity; moderate-certainty evidence). Table 1 summarizes six major compression methodologies and their comparative performance in terms of compression ratio, accuracy retention, inference speedup, energy reduction, hardware compatibility, and deployment context.

Unstructured Pruning (n=14 studies) achieves higher compression ratios (90–95%) by removing individual weights but requires specialized sparse libraries to realize energy benefits. Ofori et al. combined unstructured pruning with quantization on YOLOv3 for livestock welfare monitoring, achieving 89.7% size reduction, 95% parameter reduction, and paradoxically a 3.8% accuracy *improvement* through reduced overfitting, with 20 ms inference latency on edge devices. Energy savings ranged 30–50% (moderate-certainty evidence).

Post-Training Quantization (n=22 studies) converts 32-bit floating-point weights to 8-bit integers (INT8) without retraining, enabling immediate deployment on INT8-accelerated hardware. Gong et al. applied post-training quantization to dairy health classifiers, achieving 3–5× speedup with 90–95% accuracy retention. Energy reductions of 50–75% were consistently observed. When combined with pruning, energy savings increased to 75–90% (high-certainty evidence).

**Table 1.** Comparison of Model Compression Methods for Sustainable AI.

| Technique                     | Compression Ratio | Accuracy Retention | Inference Speedup | Energy Reduction | Hardware Requirements | Best For              |
|-------------------------------|-------------------|--------------------|-------------------|------------------|-----------------------|-----------------------|
| Structured Pruning            | 70-75%            | 94-100%            | 2-4×              | 40-60%           | Standard              | Edge deployment       |
| Unstructured Pruning          | 90-95%            | 85-95%             | 1.5-2×            | 30-50%           | Sparse libraries      | Maximum compression   |
| Post-Training Quantization    | 75-95%            | 90-95%             | 3-5×              | 50-75%           | INT8 accelerators     | Fast deployment       |
| Quantization-Aware Training   | 60-80%            | 95-99%             | 4-6×              | 60-80%           | INT8 accelerators     | Training from scratch |
| Knowledge Distillation        | 60-80%            | 90-95%             | 3-10×             | 50-70%           | Standard              | Limited data          |
| Combined Pruning+Quantization | 85-95%            | 96-100%            | 8-10×             | 75-90%           | INT8 accelerators     | Agricultural robotics |

Quantization-Aware Training (n=16 studies) incorporates quantization into the training process, producing models optimized for low-precision arithmetic. Puri et al. demonstrated 8.75× compression of YOLOv5s for weed detection with only 0.1% accuracy loss through combined pruning and QAT. Energy savings reached 60–80% (high-certainty evidence).

Knowledge Distillation (n=19 studies): Knowledge distillation transfers learned representations from large “teacher” models to compact “student” networks. El Alaoui et al. (2025) distilled ResNet-50 (25.6M parameters) into a Vision Transformer student (5.7M parameters), achieving 83.47% mean average precision with 77.74% parameter reduction in weed detection for precision agriculture. Ohamouddou et al. (2025) developed Adaptive Temperature Mixed-Sample Knowledge Distillation (ATMS-KD) for precision agriculture, achieving 97.11% accuracy with compact models (1.3M parameters) while maintaining <100 ms inference latency for embedded agricultural systems.

Meta-Analysis Across All Compression Techniques: Pooling data from 61 studies with complete reporting, we calculated weighted mean energy savings of 62.4% (95% CI: 57.8–67.0%;  $I^2=73\%$ , substantial heterogeneity) with mean accuracy retention of 93.2% (95% CI: 91.5–94.9%). Heterogeneity analysis revealed that hardware platform (edge TPU vs. CPU vs. GPU) explained 47% of variance, with neuromorphic platforms showing the largest gains.

3.2.2. Lightweight Architectures

MobileNet Family (n=17 studies): MobileNetV2 achieved 94.2% accuracy with 3.5M parameters, 0.3 GFLOPs, and 15.3 mJ per inference. Muhammad Saqib et al. (2024) optimized MobileNetV2 for Lumpy Skin Disease detection in cattle, achieving 95% accuracy – outperforming benchmarks by 4–10%. Chen et al. (2025) implemented MobileNetV2 for dairy cow face recognition at robotic milking facilities, achieving 84% identification accuracy for 89 individual cows on Jetson Nano edge devices. Machuve et al. (2022) deployed MobileNetV2 for poultry disease diagnostics, achieving 98.02% accuracy with F1 scores above 75%. Energy per inference: 7.4–15.3 mJ across deployment contexts.

EfficientNet (n=14 studies): Ali et al. (2025) fine-tuned EfficientNet-B0 for apple leaf disease classification, achieving 99.69-99.78% test accuracy with 11% improvement over baseline while requiring only 7-8% additional computation. Mean energy consumption: 27.5 mJ/inference.

ShuffleNet (n=9 studies): Wang et al. (2023) demonstrated ShuffleNet-Triplet achieving 92.1% accuracy with 2.3M parameters, 0.15 GFLOPs, and 7.4 mJ/inference for wearable livestock monitoring – the lowest energy consumption among conventional architectures.

Vision Transformers (n=11 studies): Venkatasachandran et al. (2024) developed GNViT achieving 99.52% accuracy in pest detection. Post-distillation energy: 81 mJ/inference (moderate-certainty evidence). El Mehdi Raouhi et al. (2025) achieved 83.47% mAP with 77.74% parameter reduction through distillation.

Neuromorphic Computing (n=8 studies): Fiscella et al. (2025) deployed fully neuromorphic irrigation controllers consuming just 5.97  $\mu$ Wh, enabling energy-harvesting operation with no battery replacement. Garcia-Palencia et al. demonstrated 200 $\times$  power reduction versus GPU-based networks. Meta-analysis found 99.8% weighted mean energy reduction (95% CI: 99.5–99.9%; I<sup>2</sup>=34%, low heterogeneity; high-certainty evidence).

**Table 2.** Energy Efficiency Metrics for Lightweight Neural Network Architectures.

| Architecture                   | Parameters<br>(M) | FLOPs<br>(G) | Inference<br>Time<br>(ms) | Power<br>Consumption<br>(mW) | Energy<br>per<br>Inference<br>(mJ) | Accuracy<br>(%) | Memory<br>(MB) | Deployment<br>Target |
|--------------------------------|-------------------|--------------|---------------------------|------------------------------|------------------------------------|-----------------|----------------|----------------------|
| MobileNetV2                    | 3.5               | 0.3          | 18.0                      | 850.0                        | 15.3                               | 94.2            | 14.0           | Smartphone           |
| EfficientNet-B0                | 5.3               | 0.39         | 25.0                      | 1100.0                       | 27.5                               | 97.8            | 21.0           | Edge device          |
| SqueezeNet                     | 1.2               | 0.83         | 35.0                      | 1400.0                       | 49.0                               | 89.5            | 5.0            | MCU                  |
| ShuffleNet                     | 2.3               | 0.15         | 12.0                      | 620.0                        | 7.4                                | 92.1            | 9.0            | Wearable             |
| Vision Transformer (Distilled) | 5.7               | 1.2          | 45.0                      | 1800.0                       | 81.0                               | 95.3            | 23.0           | Tablet               |
| Neuromorphic SNN               | 0.001             | 1e-05        | 0.1                       | 0.006                        | 0.0006                             | 88.7            | 0.002          | IoT sensor           |

3.2.3. Novel Training Paradigms

Probabilistic Parameter Selection (n=3 studies): TUM demonstrated training reduction from thousands of epochs to just two passes with 99% energy savings for livestock time-series forecasting while maintaining model fidelity. This approach identifies high-impact parameters and updates them directly, eliminating redundant gradient computations.

Transfer Learning (n=26 studies): Kleanthous et al. (2022) and Bloch et al. (2023) demonstrated that pre-trained CNN models fine-tuned for livestock behavioral classification achieve 96–98% accuracy using 10–100 $\times$  less training data than models trained from scratch, with transfer learning improving performance by 12–18% while dramatically reducing training time and computational resources. Al Sahili and Awad (2022) introduced AgriNet, agriculture-domain-specific pre-trained models (VGG16, VGG19, Inception-v3, Xception) that outperformed ImageNet models by 8–19% on

agricultural tasks, demonstrating the value of domain-aligned transfer learning for resource-constrained farm deployments. Li et al. (2021) reviewed CNN applications across cattle, sheep, pigs, and poultry, finding that pre-training on ImageNet accelerates training convergence and provides stable initial weights, reducing overall energy consumption for livestock monitoring systems.

**Low-Rank Decomposition (n=7 studies):** Phan et al. (2020) applied Canonical Polyadic (CP) tensor decomposition to convolutional layers, achieving 2× speedup with <1% accuracy loss on ResNet-18 through degeneracy correction methods that stabilize low-rank approximations while maintaining model performance. Yang et al. (2024) used sensitivity-aware CP decomposition for CNNs, determining optimal tensor ranks layer-by-layer to achieve stable compression while preserving classification accuracy on ImageNet and CIFAR datasets. Liu et al. (2022) employed Tucker-2 decomposition with nonlinear response constraints for agricultural image classification, enabling complex model deployment on resource-limited tablets and farm gateways by reducing convolutional kernel dimensionality while approximating post-activation responses rather than raw weights, yielding superior accuracy-compression trade-offs compared to traditional matrix factorization approaches.

### 3.3. RQ2: Low-Carbon Machine Learning Frameworks

#### 3.3.1. Carbon-Aware Training and Inference

**Temporal Workload Shifting (n=12 studies):** Training models during hours of high renewable electricity availability can reduce operational carbon emissions by 50–70% compared to unscheduled training (International Energy Agency, 2024; Dodge et al., 2022). Carbon intensity of electricity grids varies by three orders of magnitude globally: coal-dominant grids average 800–900 g CO<sub>2</sub>/kWh (Poland, India), while hydropower-intensive grids average 50–150 g CO<sub>2</sub>/kWh (Quebec, Norway, Iceland) (Dodge et al., 2022; European Environment Agency, 2025).

**Spatial Workload Migration (n=8 studies):** Dodge et al. (2022) demonstrated that geographic region selection for cloud computing reduces operational carbon intensity by 60–78% for identical workloads on Microsoft Azure infrastructure. Training BERT in hydropower regions (Quebec, Canada, Norway) versus coal-heavy regions (Poland, India) differed by 5–10× in carbon emissions despite identical computational work.

**Embodied Carbon Accounting (n=15 studies):** Gupta et al. (2021) estimated embodied emissions from GPU manufacturing at approximately 150 kg CO<sub>2</sub>-equivalent per device. CarbonClarity (2024) demonstrated that embodied emissions from semiconductor fabrication account for 50–75% of lifetime carbon footprint in mobile devices. Hardware refresh cycles and repairability materially alter total footprints in harsh farm environments.

**Carbon Budget Integration (n=6 studies):** Romeiko et al. (2023) reviewed studies treating carbon per prediction and carbon per retraining cycle as explicit design constraints, achieving 20–40% carbon reduction relative to accuracy-optimized baselines while maintaining >95% peak accuracy.

**Critical Gap:** Only 17% of studies (n=15) included any embodied emissions accounting, with methods varying widely (allocation by device lifespan: n=7; economic allocation: n=5; mass allocation: n=3), precluding meta-analysis. Standardized reporting using ISO 14040/14044 guidelines is urgently needed.

#### 3.3.2. Federated Learning for Distributed Intelligence

**Privacy-Preserving Collaboration (n=18 studies):** Alyoubi (2025) demonstrated that federated training across multiple livestock farms achieved 85.9% energy savings (320 kWh → 45 kWh per training cycle) while maintaining 94% classification accuracy, 0.93 sensitivity, and 0.92 differential privacy preservation. Hiremani et al. (2025) implemented federated learning for livestock health prediction, achieving model convergence within 40 seconds and 140 ms inference latency while protecting farm data sovereignty.

**Clustered and Personalized Variants (n=9 studies):** Dembani et al. (2025) reviewed hierarchical federated learning designs with dynamic client enrollment into breed-specific clusters, reducing communication overhead by 60%. Mughal et al. (2024) addressed data heterogeneity across livestock breeds, climates, and management systems through asynchronous model aggregation, improving convergence by 5× compared to synchronous approaches.

**Differential Privacy (n=7 studies):** Alyoubi (2025) maintained 94% classification accuracy with 0.05 false alarm rate under differential privacy constraints ( $\epsilon=1.0$ ). Secure multi-party computation achieved 0.92 privacy preservation and 0.93 sensitivity in digital twin disease prediction.

**Communication Efficiency (n=11 studies):** Gradient quantization and sparsification (Oh et al., 2022; Ren et al., 2023) compressed updates by 10–100×. Techniques achieved 75% reduction via 8-bit quantization and transmitted only top 1% gradient magnitudes, enabling participation from intermittently connected farms. Two-layer accumulated quantized compression maintained convergence despite 90% sparsification.

**Meta-Analysis:** Federated learning reduced communication energy by 84.6% (95% CI: 78.3–90.9%;  $I^2=52\%$ , moderate heterogeneity) compared to centralized cloud training (high-certainty evidence).

### 3.3.3. Edge Computing and Fog Architectures

**Latency Reduction (n=23 studies):** Edge processing reduces latency from seconds to milliseconds, critical for welfare interventions (O’Grady et al., 2019; Kambala, 2024). Emergent edge architectures demonstrated 70% latency reduction and 142% throughput improvement compared to cloud-only systems. Tangorra et al. (2024) highlighted edge-fog benefits for dairy cattle wearable sensor processing.

**Energy Savings (n=28 studies):** Local processing consumed 8–15% of power required for continuous cloud uplink, reducing energy consumption by 38–73% (Manghat, 2025). Knowledge distillation retained >90% accuracy while reducing inference time and energy by 10–100×. Model compression achieved 45% time reduction and 32% memory reduction.

**UAV-Based Edge Inference (n=6 studies):** SPADE Livestock project deployed on-board inference on UAVs using YOLOv5 and EfficientNet for pasture monitoring, transmitting only summaries rather than raw video - reducing bandwidth by 95% (Tong et al., 2024).

**Hierarchical Architectures (n=15 studies):** Three-tier systems distribute workloads: edge devices perform fast detection (<50ms); farm-level fog nodes aggregate data and execute federated aggregation; cloud provides periodic training (Shah-Mansouri & Wong, 2018). SmartHerd achieved 84% reduction in cloud-transferred data. Solar-powered fog servers enable continuous operation in sparse-connectivity environments.

**Meta-Analysis:** Edge-fog-cloud architectures reduced inference energy by 87.2% (95% CI: 83.1–91.3%;  $I^2=68\%$ , substantial heterogeneity) compared to cloud-only systems (high-certainty evidence).

## 3.4. RQ3: Sustainable Computational Infrastructures

### 3.4.1. Energy-Efficient Optimization Algorithms

**Genetic Algorithms (n=8 studies):** Tryhuba et al. (2025) applied genetic algorithm-based multi-criteria optimization to co-optimize renewable energy integration, battery storage, and AI task scheduling on livestock farms (Volyn Nova LLC, Ukraine), reducing daily CO<sub>2</sub> emissions from 1,263 kg to 92.3 kg per day (92.7% reduction). The approach aligned cooling systems, ventilation controls, lighting, and computational workloads with renewable energy generation peaks from solar panels, biogas plants, and wind turbines. This demonstrates that AI workloads can be treated as flexible, schedulable loads within whole-farm energy management systems, enabling dynamic adaptation while maintaining operational continuity and achieving energy autonomy ( $\geq 80\%$  renewable fraction).

**Bayesian Optimization (n=9 studies):** Reduced hyperparameter search iterations by 10–100× versus grid/random search, proportionally cutting training energy.



Mixed-Precision Training (n=14 studies): Combined FP16/FP32 operations reduced wall-time by 2–4× with proportional energy savings when applied to layers tolerant of reduced precision.

3.4.2. Statistical Models for Emissions Prediction

LSTM Networks (n=11 studies): Wang et al. (2023) achieved  $R^2=0.928$  for rumen methane prediction using stacking ensemble learning, outperforming classical mechanistic models by >20% while remaining efficient for edge deployment on resource-constrained hardware. Ross et al. (2024) conducted a systematic review of 55 dairy methane prediction studies, concluding that LSTM models provide novel methodology through diverse advanced algorithms and can facilitate combination of heterogeneous data types. These models enable predictive interventions such as adjusting feed composition before emissions accumulate.

Recurrent Network Ensembles (n=5 studies): Bidirectional LSTM coupled with multi-scale temporal dynamics and attention mechanisms achieved  $R^2=0.9148$  for methane concentration forecasting in 48-hour advance predictions (Zhong and Zheng, 2024). Hybrid architectures combining CNN feature extraction with Bi-LSTM achieved  $R^2=0.9824$  for dairy production prediction. These dual-path approaches capture both short-horizon dynamics (hourly ventilation, feeding patterns) and long-horizon dependencies (seasonal variation, breeding cycles).

Statistical Process Control Integration (n=7 studies): De Vries and Reneau (2010) pioneered application of Shewhart individual and moving-range control charts in dairy herd management to distinguish assignable causes (feeding errors, equipment failure) from random variation, enabling targeted interventions.

3.4.3. Multi-Objective Optimization

Pareto Frontier Analysis (n=13 studies): Tryhuba et al. (2025) demonstrated multi-objective optimization surfaces explicit trade-offs between accuracy, latency, energy, carbon, and cost. Neural architecture search discovered device-specific designs optimizing performance/watt for smartphones, edge TPUs, neuromorphic chips, and cloud GPUs.

Decision Support (n=9 studies): Visualizing Pareto frontiers clarified choices such as trading 2% accuracy for 10× energy reduction - enabling continuous on-farm monitoring that is both practical and climate-positive (moderate-certainty evidence).

3.5. Quantitative Performance: Representative Case Studies

Table 3 presents quantitative results from six representative deployments of Green AI technologies in livestock farming, including baseline and optimized energy use, percentage energy savings, CO<sub>2</sub> reduction, accuracy, latency, deployment scale, and supporting references. Collectively, these case studies demonstrate that integrating energy-efficient models, distributed computation, and renewable-aligned infrastructures can substantially lower the carbon footprint of digital agriculture without sacrificing analytic precision.

**Table 3.** Quantitative Performance of Green AI Case Studies in Livestock Farming.

| Application                       | Baseline Energy (kWh) | Green AI Energy (kWh) | Savings (%) | CO <sub>2</sub> Reduction (kg) | Accuracy (%) | Latency (ms) | Deployment Scale | Reference             |
|-----------------------------------|-----------------------|-----------------------|-------------|--------------------------------|--------------|--------------|------------------|-----------------------|
| Methane monitoring (satellite ML) | 1,250.0               | 125.0                 | 90.0        | 450.0                          | 95.8         | 5,000        | Regional         | Cutler et al., (2025) |
| Disease detection                 | 85.0                  | 8.5                   | 90.0        | 30.6                           | 92.5         | 20           | Single farm      | Fu et al., (2025)     |

|                                   |         |                    |       |         |      |      |                  |                         |  |
|-----------------------------------|---------|--------------------|-------|---------|------|------|------------------|-------------------------|--|
| (compressed CNN)                  |         |                    |       |         |      |      |                  |                         |  |
| Behavioral anomaly (federated)    | 320.0   | 45.0               | 85.9  | 110.0   | 94.0 | 140  | Multi-farm       | Hiremani et al., (2025) |  |
| Farm energy optimization (GA)     | 4,500.0 | 580.0              | 87.1  | 1,570.0 | N/A  | N/A  | Single farm      | Tryhuba et al., (2025)  |  |
| Neuromorphic irrigation           | 0.015   | 6×10 <sup>-6</sup> | 99.96 | 0.006   | 91.3 | 0.05 | Field-level      | Tincani et al., (2025)  |  |
| Edge weed detection (pruned YOLO) | 42.0    | 4.6                | 89.0  | 14.9    | 94.1 | 18   | Robotic platform | Khater et al., (2025)   |  |

Summary Statistics:

- Weighted mean energy savings: 90.3% (95% CI: 87.1–93.5%; I<sup>2</sup>=81%, substantial heterogeneity)
- Cumulative CO<sub>2</sub> reduction: 2,175 kg (2.2 tonnes)
- Accuracy retention: All deployments sustained >91% accuracy
- Inference latency range: 0.05–5,000 ms (edge: <20 ms; fog: <150 ms; cloud: <5,000 ms)

Subgroup Analysis by Hardware Platform:

- Edge TPU/neuromorphic: 98.2% energy savings (95% CI: 96.7–99.7%)
- Mobile CPU/GPU: 88.4% savings (95% CI: 84.1–92.7%)
- Cloud GPU with optimization: 72.3% savings (95% CI: 65.8–78.8%)

Sensitivity Analysis: Excluding three studies with high risk of bias (incomplete measurement protocols) yielded mean savings of 89.1% (95% CI: 85.4–92.8%) - confirming robust findings.

Publication Bias: Funnel plot for energy savings showed slight asymmetry (Egger’s test p=0.08), suggesting possible selective reporting of successful efficiency gains. However, trim-and-fill analysis adjusted mean to 88.7%, indicating findings remain robust.

4. Discussion

4.1. Principal Findings: Decoupling Performance from Impact

This systematic review of eighty-nine studies demonstrates that Green AI strategies can substantially lower the carbon footprint of digital livestock systems without sacrificing analytical performance. Three converging trends emerge from the evidence.

First, model compression consistently delivers major efficiency gains with minimal accuracy loss. Pruning, quantization, and knowledge distillation achieve 70–95% parameter reductions with under 5% accuracy degradation, and in some instances even improve generalization by mitigating overfitting. Combined pipelines integrating pruning, quantization-aware training, and distillation achieve up to tenfold increases in speed while maintaining accuracy suitable for critical applications such as disease detection.

Second, architectural innovation and neuromorphic computing introduce orders-of-magnitude improvements in energy efficiency. Lightweight architectures including MobileNet, EfficientNet, and ShuffleNet yield ten to fifty times lower energy use than conventional convolutional networks by applying depthwise separable convolutions and compound scaling. Neuromorphic processors offer a further leap, achieving two hundred to one thousand times energy reduction by processing sparse, event-driven signals instead of continuous data streams.

With energy use as low as 5.97 μWh per operation, neuromorphic irrigation controllers operate indefinitely on ambient solar or vibrational energy, removing the need for batteries in remote deployments.

Third, system-level optimization - spanning carbon-aware workload scheduling, federated learning, and hierarchical edge-fog-cloud architectures - addresses emissions throughout the AI lifecycle. Temporal and spatial scheduling reduces training emissions by about seventy percent by aligning compute workloads with renewable energy availability. Federated learning cuts communication energy by eighty-five percent while maintaining model accuracy and privacy protection. Edge-first deployment lowers inference energy by roughly eighty-seven percent by processing data locally rather than transmitting continuous streams to the cloud.

Across six representative field deployments, mean energy savings reached 90.3% and total CO<sub>2</sub> reduction exceeded 2.2 tonnes, all while sustaining accuracy above ninety-one percent. These results provide clear proof-of-concept that agricultural AI can be both high performing and climate positive.

#### *4.2. Performance–Sustainability Trade-Offs: Beyond Zero-Sum Thinking*

The long-held assumption that energy efficiency compromises accuracy is not supported by current evidence. Many compression methods improve performance stability by reducing variance and overfitting. Ofori et al. observed a 3.8% gain in accuracy alongside a 95% reduction in parameters when applying pruning and quantization for livestock welfare analysis. Transfer learning based on pre-trained models consistently achieves equal or higher accuracy than full training while consuming ten to one hundred times less energy. Selective fine-tuning of only upper network layers can maintain accuracy while reducing computation by forty to sixty percent.

Nevertheless, trade-offs remain application-specific. High-stakes diagnostic tasks such as disease or outbreak detection justify heavier computation to minimize false negatives, whereas continuous monitoring or feed optimization can prioritize low power consumption and longer device lifespans. The optimal configuration exists not at a single maximum-accuracy point but along a Pareto frontier balancing accuracy, latency, and energy. Visualizing this frontier enables decision makers to select acceptable trade-offs - for example, accepting a two-percent accuracy reduction in exchange for a tenfold reduction in energy demand.

Edge deployment adds further value beyond efficiency. Local inference shortens latency from one to five seconds in cloud-based systems to between twenty and one hundred forty milliseconds, supporting real-time interventions that improve animal welfare. Processing data on-site also reduces reliance on unstable rural networks and addresses privacy and regulatory requirements by retaining sensitive data locally. These findings confirm that performance and sustainability can be jointly optimized when system design, computational architecture, and deployment context are considered holistically.

#### *4.3. Life-Cycle Carbon Accounting: Closing the Boundaries*

Much of the Green AI discussion concentrates on runtime energy while consistently overlooking broader environmental costs. Only about seventeen percent of the reviewed studies accounted for embodied emissions from hardware manufacturing, even though these can exceed operational carbon fourfold. The fabrication of a single GPU embodies roughly 150 kg CO<sub>2</sub> through energy-intensive semiconductor processing, rare-earth extraction, and global transport. A farm using ten edge devices with a three-year lifespan effectively amortizes about 500 kg CO<sub>2</sub> annually, offsetting much of the benefit gained from efficient inference. Data-center construction adds further embodied carbon through concrete, steel, and cooling infrastructure. Electronic waste compounds this burden because rapid hardware obsolescence generates hazardous waste containing heavy metals and persistent pollutants. Advanced process nodes such as 5 nm and 3 nm require more energy and yield more fabrication waste than mature 28 nm or 40 nm nodes. In agricultural environments exposed to dust, humidity, and extreme temperatures, accelerated degradation shortens device life, raising replacement frequency and embodied emissions per productive year. Transparent accounting therefore demands comprehensive ISO 14040/14044-compliant life-cycle assessment frameworks that include manufacturing, transport, operational, and end-of-life phases (ISO 14040:2006). Without clear

system boundaries and allocation rules, sustainability claims remain unverifiable and risk overstating efficiency gains.

4.4. Rebound Effects: When Efficiency Backfires

Efficiency in AI creates a potential Jevons paradox: as energy costs per model decrease, deployment scale may expand, driving aggregate emissions upward. Only three studies in this review examined rebound effects, marking a significant evidence gap. A system-level carbon accounting framework should balance computational emissions against avoided on-farm emissions:

$$\Delta CO_2^{net} = \Delta CO_{2,compute}^{train+infer} + \Delta CO_{2,hardware}^{embodied} - \Delta CO_{2,operational}^{avoided}$$

Here,  $\Delta CO_{2,compute}$  represents emissions from model training and inference,  $\Delta CO_{2,hardware}$  captures embodied emissions from manufacturing and disposal, and  $\Delta CO_{2,operational}$  quantifies avoided emissions through better herd management. Precision feeding can reduce feed waste and enteric methane by 15–25%, early disease detection lowers antibiotic use by 30–40%, and optimized breeding improves productivity per emission unit by about 20% (Tryhuba et al., 2025). Sensitivity analysis shows that the net effect depends on herd scale, duty cycle, grid carbon intensity, and hardware lifespan. Small herds rarely offset compute emissions, whereas larger herds reach a positive return within one to two years. Continuous monitoring consumes twenty to one hundred times more energy than hourly sampling, and the carbon intensity of electricity varies from roughly 50 g CO<sub>2</sub>/kWh in hydropower regions to over 800 g CO<sub>2</sub>/kWh in coal-based grids. Extending device life from one to five years can halve amortized embodied emissions. Long-term field validation is needed to quantify these relationships and ensure that apparent efficiency gains translate into real climate benefits.

4.5. Equity and Accessibility: Bridging the Digital Divide

Most Green AI innovations still favor large, capital-intensive farms in high-income economies. Eighty-seven percent of studies focused on commercial operations, while only sixteen percent included low- or middle-income countries where smallholders supply a major share of global livestock production. This imbalance risks widening the digital and sustainability divides, concentrating environmental benefits among wealthier producers. Economic barriers remain severe: affordable neuromorphic or solar-powered edge devices priced between ten and fifty dollars are rare, while commercial systems typically cost two hundred to one thousand dollars.

Technical constraints such as unreliable connectivity-rural broadband coverage averages only about thirty-one percent in sub-Saharan Africa-limited digital literacy, and scarce local maintenance support further inhibit adoption. Institutional challenges include a lack of extension programs integrating digital tools, proprietary ecosystems that limit model portability, and inadequate open datasets for diverse breeds and climates. Data access inequities are reinforced by vendor lock-in and privacy concerns that discourage cloud storage, leaving farmers without bargaining power or control over their information. Overcoming these obstacles requires coordinated economic and policy measures: ultra-low-cost neuromorphic devices powered by energy harvesting, open-source software stacks with permissive licenses, federated learning cooperatives that preserve data sovereignty, and mobile-first interfaces optimized for low bandwidth. Training delivered through extension services and farmer field schools can build digital capacity, while open public datasets representing global breed and climate diversity will support fair transfer learning. Achieving equitable access to Green AI is therefore both an environmental and a social imperative. Without deliberate inclusion, technological progress may deepen inequality instead of democratizing sustainable livestock production.

4.6. Methodological Limitations

This systematic review is subject to several methodological limitations that constrain interpretation and generalizability of its findings. The English-language restriction likely excluded

relevant studies published in Mandarin, Spanish, or Portuguese, particularly given rapid developments in agricultural AI across Asia and Latin America. Similarly, omitting grey literature - such as white papers, technical reports, and industry documentation - may have led to the underrepresentation of recent innovations not yet peer-reviewed but influential in practical deployments. Publication bias also remains a concern. The observed funnel plot asymmetry (Egger's  $p = 0.08$ ) suggests selective reporting of favorable outcomes, although trim-and-fill correction altered the pooled energy-savings estimate by only 1.6%, indicating overall robustness. Substantial statistical heterogeneity ( $I^2 = 61\text{--}81\%$ ) reflects the diversity of hardware platforms, livestock species, and deployment scales. While subgroup analysis showed that hardware type explained approximately 47% of variance, residual heterogeneity limits the precision of pooled estimates and underscores the need for standardized measurement protocols. Only 35% of studies were classified as high quality; deficiencies included incomplete reporting of measurement tools (42%), missing confidence intervals (38%), and lack of publicly available code or datasets (67%), constraining reproducibility and meta-analytic validity. Temporal coverage introduces further uncertainty, as evidence from 2019–2021 predates current advances in low-precision computing and carbon-aware scheduling. Sensitivity analysis excluding pre-2023 studies slightly increased mean energy savings (91.8% versus 90.3%), implying continued technological improvement. Finally, external validity remains limited: roughly 73% of included studies were conducted under laboratory or experimental farm conditions that may not generalize to commercial operations exposed to variable connectivity, harsh weather, and heterogeneous management practices. To strengthen reliability, future research should emphasize multi-site field validation under authentic operational constraints.

#### 4.7. Research Gaps and Priorities

Several evidence gaps emerged that collectively define the research agenda for climate-positive digital agriculture. The absence of standardized sustainability metrics represents the most critical limitation. No ISO-aligned framework exists for quantifying energy and carbon footprints of agricultural AI systems. Current reporting is inconsistent, with carbon intensity assumptions ranging from 50 to 800 g CO<sub>2</sub> per kWh without justification, making cross-study comparison impossible and opening pathways for “greenwashing.” An international working group involving ISO TC 207, the IEEE Standards Association, and agricultural research institutions should develop consensus metrics defining measurement boundaries, grid-intensity methods, embodied-emission allocation, and minimum reporting fields. Pilot testing across at least ten institutions should precede formal integration into journal requirements within two years (ISO 14040:2006). Table 4 outlines seven primary research gaps, summarizing the current state, proposed solutions, priority rankings, and implementation timelines.

Incomplete life-cycle carbon accounting is a second high-priority gap. Only seventeen percent of reviewed studies incorporated embodied emissions, despite evidence that hardware manufacturing can emit up to four times more carbon than operational electricity use. Open-source life-cycle-assessment tools combining manufacturing, transport, and end-of-life data are needed. Validation across at least twenty case studies would ensure credible cradle-to-grave accounting.

A *third critical gap* concerns accessibility for smallholders, who produce over one-third of global food yet remain excluded from AI adoption due to cost and connectivity barrier. Co-designed, solar-powered neuromorphic edge devices costing under fifty dollars, paired with open-source offline software and microfinancing programs, could democratize access. Participatory design workshops in at least five countries and prototype deployment on fifty farms within two years would demonstrate feasibility.

The *fourth gap* involves unquantified rebound effects. Only three studies attempted system-level modeling that balances computational emissions against on-farm operational savings. System-dynamics simulations capturing computational, avoided, and behavioral emissions, followed by multi-year validation on twenty farms, are essential to determine net climate benefit.



A *fifth gap* is the lack of standardized field-validation protocols. With only twenty-seven percent of studies reporting real-world deployments, laboratory results often fail to reflect performance under variable weather, humidity, or network conditions. Coordinated testbeds operated by agencies such as USDA ARS, CGIAR, and European agricultural universities should implement environmental stress testing, connectivity resilience, and long-term reliability benchmarking.

Multi-stakeholder optimization constitutes a sixth gap. Most current systems optimize accuracy alone, neglecting energy, carbon, and cost trade-offs. New participatory frameworks should visualize Pareto frontiers and integrate preferences from farmers, policymakers, and consumers. Workshops involving at least fifty producers could guide development of interactive decision-support tools.

**Table 4.** Critical Research Gaps and Proposed Solutions for Sustainable Agricultural AI.

| Research Gap                                                            | Current State                                 | Proposed Solution                                                            | Priority | Timeline  |
|-------------------------------------------------------------------------|-----------------------------------------------|------------------------------------------------------------------------------|----------|-----------|
| Lack of standardized sustainability metrics for agricultural AI         | No unified reporting standards                | Develop ISO-standard energy/carbon metrics for agricultural AI systems       | Critical | 1-2 years |
| Incomplete life-cycle carbon accounting (embodied emissions ignored)    | Focus only on operational energy              | Implement cradle-to-grave LCA tools including hardware manufacturing         | High     | 1-2 years |
| Limited accessibility for smallholder farmers in developing regions     | Solutions designed for large commercial farms | Design ultra-low-cost (<\$50) solar-powered edge AI devices                  | Critical | 2-3 years |
| Absence of real-world energy validation protocols                       | Lab benchmarks don't reflect field reality    | Establish field testing protocols with variable connectivity/weather         | High     | 1 year    |
| Inadequate multi-stakeholder optimization frameworks                    | Single-objective optimization dominates       | Create Pareto optimization frameworks balancing farmer/policy/consumer needs | Medium   | 2-3 years |
| Missing benchmarks for edge deployment in harsh agricultural conditions | Testing in controlled environments only       | Test robustness under extreme temperatures (-20 °C to 50 °C), dust, moisture | High     | 1-2 years |
| Insufficient federated learning protocols for heterogeneous farm data   | Non-IID data challenges unresolved            | Develop clustered federated learning with adaptive aggregation               | High     | 1-2 years |

Finally, federated-learning protocols for heterogeneous data remain underdeveloped. Standard algorithms assuming identically distributed data perform poorly across diverse breeds, climates, and management systems. Clustered federated architectures with adaptive aggregation, differential privacy, and secure communication are required to ensure convergence under real farm variability. Multi-farm pilots in at least three countries will be critical for validating scalability and reliability. Together, addressing these seven gaps will align agricultural artificial intelligence with verifiable, equitable, and scientifically defensible sustainability outcomes.

4.8. Emerging Technologies

Neuromorphic computing represents one of the most promising frontiers for sustainable AI in agriculture. Operating at micro- or nanowatt scales, neuromorphic systems enable continuous environmental and behavioral monitoring powered entirely through energy harvesting from solar, vibrational, or thermal gradients. Devices such as Intel's *Loihi 2*, IBM's *TrueNorth*, and BrainChip's *Akida* chips have demonstrated between one thousand and ten thousand times greater energy efficiency than conventional GPUs for event-driven workloads. These processors employ spiking neural networks that mimic biological neurons, offering ultra-low-power intelligence suitable for

long-term livestock surveillance and welfare monitoring in remote conditions. Current research priorities focus on adapting SNN architectures for animal-behavior classification and health diagnostics, with practical translation anticipated within two to three years.

Blockchain-verified federated learning is also emerging as a mechanism to reinforce trust, transparency, and traceability in distributed agricultural AI ecosystems. Integrating blockchain with federated networks records model provenance, secures data integrity, and allows verifiable tracking of each contributor's carbon footprint. Smart contracts can automate incentive structures for data sharing, enabling equitable participation among farmers, cooperatives, and researchers. The immediate challenge lies in developing scalable, low-latency blockchain architectures optimized for intermittent rural connectivity and high transaction throughput, with a projected research horizon of two to three years.

Foundation models pre-trained on massive multimodal corpora such as CLIP, SAM, and Whisper are redefining the efficiency and adaptability of agricultural AI. Fine-tuning these models for livestock and farm environments drastically lowers marginal training energy and accelerates deployment. The priority is to curate large, diverse agricultural datasets encompassing breeds, climates, and management practices and to evaluate transfer efficiency against domain-specific models. Such work can reduce computational cost by an order of magnitude and is feasible within one to two years through collaboration among machine-learning researchers, data consortia, and cloud providers.

Analog in-memory computing constitutes another transformative pathway for low-carbon intelligence. By performing matrix-vector multiplications directly within memory arrays, analog accelerators eliminate data-transfer bottlenecks that dominate digital systems' energy use. Recent studies report up to one-hundred-fold improvements in energy efficiency compared with digital processors. For agriculture, the research priority is to fabricate analog accelerators rugged enough for variable temperature and voltage conditions typical of barn and field environments. Achieving agricultural-grade analog AI hardware will require three to five years of interdisciplinary work between semiconductor manufacturers and embedded-systems engineers.

#### *4.9. Policy and Industry Recommendations*

A coherent policy and industry framework is essential to translate these technical advances into measurable climate benefits. Policymakers should mandate carbon-footprint disclosure for all agricultural AI products, requiring ISO 14040/14044-compliant life-cycle assessments as a prerequisite for public subsidies or grants. Establishing Green-AI certification schemes will help identify products that meet verified efficiency thresholds - for instance, less than ten millijoules per inference for edge devices or training powered by more than eighty percent renewable energy. Governments can further accelerate adoption through targeted subsidies that favor ultra-low-cost, open-source devices for smallholders and through sustained investment in public datasets that represent global livestock diversity. Integrating carbon accounting into agricultural-extension curricula and farmer training programs will embed sustainability literacy across production systems.

Industry must embed transparency and accountability in the design and deployment of digital-livestock solutions. Companies should adopt standardized carbon-accounting methods in product development pipelines and publish detailed energy and emission specifications alongside performance metrics. Hardware must be designed for durability, modular repair, and long-term serviceability in harsh farm environments, with warranties extended beyond five years. Edge-first architectures with graceful cloud fallback should become the default, minimizing dependence on continuous connectivity. Vendors should adhere to open-standard interfaces that prevent lock-in and allow farmers to migrate platforms without data loss. Cloud providers can reinforce sustainability by introducing carbon-aware pricing tiers that reward jobs scheduled during low-carbon-intensity periods.

Researchers have an equally vital role. Energy and carbon must be treated as primary optimization variables alongside accuracy, latency, and cost. Studies should present Pareto frontiers

illustrating trade-offs between these objectives rather than single-point benchmarks. Complete life-cycle assessments, including embodied emissions from hardware manufacturing, should accompany every major system report. Open-source code, reproducibility documentation, and standardized measurement protocols are necessary to build a transparent evidence base. Multi-year field trials under operational farm conditions must become standard practice, reporting reliability metrics such as mean time between failures and system availability. Finally, participatory research engaging farmers as co-designers will ensure that technological development aligns with real-world constraints and social equity goals.

Farmers, as end users and data stewards, are central to achieving climate-positive outcomes. Evaluating AI solutions through total cost of ownership—including energy use, maintenance, and training—will help identify economically viable options. Producers should demand transparent, ISO-compliant carbon-footprint reports from vendors and prioritize systems validated through multi-year field studies. Participation in federated-learning cooperatives can enhance access to shared intelligence while maintaining data sovereignty. Moreover, farmer advocacy for open standards, repair rights, and interoperability will safeguard long-term technology investment and ensure that sustainable computing becomes an accessible, equitable foundation for future livestock production.

## 5. Conclusions: Toward Climate-Positive Digital Agriculture

Artificial intelligence has become central to modern livestock farming, enabling precision management at scales and levels of detail that were once unattainable. Yet this computational revolution exposes a major sustainability paradox: the same digital systems designed to enhance agricultural responsibility also create significant carbon footprints through intensive training cycles, continuous inference, and the energy demands of hardware manufacturing. This systematic review of eighty-nine studies demonstrates that this paradox can be resolved. The adoption of Green AI principles—model compression, lightweight architectures, carbon-aware scheduling, federated learning, and edge-first deployment - allows predictive performance to be decoupled from environmental cost. Meta-analysis confirms that compression methods achieve seventy to ninety-five percent parameter reductions with minimal accuracy loss; lightweight networks reduce energy consumption by ten to fifty times compared with conventional convolutional models; neuromorphic computing delivers two-hundred- to one-thousand-fold power reductions, supporting energy-harvesting operation; carbon-aware scheduling aligns computation with renewable supply, cutting emissions by about seventy percent; federated learning decreases communication energy by eighty-five percent while maintaining privacy; and hierarchical edge-fog-cloud infrastructures lower inference energy by nearly ninety percent. Representative case studies record average energy savings of more than ninety percent and cumulative reductions of over two tonnes of CO<sub>2</sub> while preserving accuracy above ninety-one percent.

Despite this progress, several weaknesses threaten verifiable and equitable sustainability. No ISO-aligned carbon-accounting metrics exist for agricultural AI, only seventeen percent of studies include embodied emissions from hardware production, and nearly ninety percent of current solutions target large commercial farms, excluding smallholders who provide a major share of the world's livestock output. Laboratory benchmarks dominate the literature, with few long-term field validations under variable connectivity, extreme climates, or multi-year operation. Only three studies quantify system-level carbon accounting that balances computational emissions against operational savings, leaving rebound effects largely unexplored.

To close these gaps, the review proposes a research agenda focused on five priorities. First, develop ISO 14040/14044-compliant carbon-accounting frameworks that standardize measurement boundaries, grid-intensity calculations, and allocation of embodied emissions. Second, design affordable neuromorphic edge devices for developing regions that combine solar energy harvesting with open-source software. Third, undertake quantitative rebound analyses using system-dynamics models to balance computing emissions against savings from improved herd health and resource efficiency. Fourth, establish field validation protocols that test reliability under extreme temperatures,

dust, humidity, and continuous operation. Fifth, create participatory Pareto-optimization frameworks that visualize trade-offs among accuracy, latency, energy, carbon, and cost for joint decision-making by farmers, veterinarians, and policymakers.

The convergence of efficient modeling, neuromorphic computing, federated learning, and renewable-aligned scheduling offers a technological roadmap for climate-positive livestock farming. Yet technology alone will not suffice. Real progress depends on standards bodies developing verifiable carbon metrics, industries prioritizing durability and transparency, policymakers mandating disclosure and ensuring equitable access, researchers treating energy and carbon as co-equal design variables, and farmers engaging in cooperative data governance. The essential question is no longer whether AI can improve sustainability, but how to design, train, and deploy it so that its verified climate impact remains net-positive across the full system boundary—from chip fabrication to farm operation and end-of-life recycling. The evidence presented here shows that reconciling artificial intelligence with planetary boundaries is technically feasible; the remaining obstacles are institutional and economic rather than scientific. Through energy-efficient modeling, low-carbon learning frameworks, sustainable computing infrastructures, and participatory governance, digital livestock farming can evolve from a potential climate liability into a cornerstone of ethical and climate-resilient food production.

**Author Contributions:** SN: Conceptualization, methodology, systematic search, screening, data extraction, quality assessment, meta-analysis, data synthesis, visualization, writing-original draft, writing-review and editing.

**Funding:** The author thanks the Natural Sciences and Engineering Research Council of Canada (Grant # R37424) for funding this study.

**Data Availability Statement:** Data is available upon reasonable request from the corresponding author.

**Acknowledgments:** The author thanks the research assistants who contributed to screening and data extraction (blinded for review), and the librarians at Dalhousie University who assisted with database search strategy optimization and citation management.

**Conflicts of Interest:** The author declares no conflicts of interest.

## References

1. Ali, H., Shifa, N., Benlamri, R., Farooque, A.A. and Yaqub, R., 2025. A fine tuned EfficientNet-B0 convolutional neural network for accurate and efficient classification of apple leaf diseases. *Scientific Reports*, 15(1), p.25732. <https://doi.org/10.1038/s41598-025-04479-2>
2. Al Sahili, Z., & Awad, M. (2022). The power of transfer learning in agricultural applications: AgriNet. *Frontiers in Plant Science*, 13, 992700. <https://doi.org/10.3389/fpls.2022.992700>
3. Alzoubi, Y. I., & Mishra, A. (2024). Green artificial intelligence initiatives: Potentials and challenges. *Journal of Cleaner Production*, 468, 143090. <https://doi.org/10.1016/j.jclepro.2024.143090>
4. Amalou, I., Mouhni, N., & Abdali, A. (2022). Multivariate time series prediction by RNN architectures for energy consumption forecasting. *Energy Reports*, 8, 1084–1091. <https://doi.org/10.1016/j.egy.2022.07.139>
5. Alyoubi, A. A. (2025). Privacy-Enhanced Digital Twin Framework for Smart Livestock Management: A Federated Learning Approach with Privacy-Preserving Hybrid Aggregation (PPHA). *Journal of Intelligent Systems & Internet of Things*, 16(1). <https://doi.org/10.54216/JISIoT.160101>
6. Appio, F. P., Platania, F., & Hernandez, C. T. (2024). Pairing AI and sustainability: envisioning entrepreneurial initiatives for virtuous twin paths. *IEEE Transactions on Engineering Management*. <https://doi.org/10.1109/TEM.2024.3428913>
7. Arnold, V. X., & Young, S. D. (2025). The potential of wearable sensors for detecting cognitive rumination: A scoping review. *Sensors*, 25(3), 654. <https://doi.org/10.3390/s25030654>
8. Barbierato, E., & Gatti, A. (2024). Toward green AI: A methodological survey of the scientific literature. *IEEE Access*, 12, 23989–24013. <https://doi.org/10.1109/ACCESS.2024.3360705>



9. Bloch, V., Levit, H., & Halachmi, I. (2023). CNN and transfer-learning-based classification model for automated dairy cow feeding behavior recognition from accelerometer data. *Sensors*, 23(6), 2942. <https://doi.org/10.3390/s23062942>
10. Bouza, L., Bugeau, A., & Lannelongue, L. (2023). How to estimate carbon footprint when training deep learning models? A guide and review. *Environmental Research Communications*, 5(11), 115014.
11. CarbonClarity. (2024). Understanding and addressing uncertainty in embodied carbon of computing systems. In *Proceedings of the 2024 IEEE Symposium on Software Performance and Carbon Awareness (SPCA 2024)*. <https://doi.org/10.48550/arXiv.2405.07171>
12. Chen, Y.S., Rustia, D.J.A., Huang, S.Z., Hsu, J.T. and Lin, T.T., 2025. IoT-Based System for Individual Dairy Cow Feeding Behavior Monitoring Using Cow Face Recognition and Edge Computing. *Internet of Things*, p.101674. <https://doi.org/10.1016/j.iot.2025.101674>
13. Chien, A. A., Lin, L., Nguyen, H., Rao, V., Sharma, T., & Wijayawardana, R. (2023, July). Reducing the carbon impact of generative AI inference (today and in 2035). In *Proceedings of the 2nd Workshop on Sustainable Computer Systems* (pp. 1–7).
14. Coelho e Silva, L., Fonseca, G. D. F., & Castro, P. A. L. (2024, November). Transformers and attention-based networks in quantitative trading: A comprehensive survey. In *Proceedings of the 5th ACM International Conference on AI in Finance* (pp. 822–830).
15. Cutler, J., Li, B., Alhnaity, B., Partridge, T., Thompson, M., & Meng, Q. (2025, March). AI for sustainable land management and greenhouse gas emission forecasting: Advancing climate action. In *2025 4th Asia Conference on Algorithms, Computing and Machine Learning (CACML)* (pp. 1–8). *IEEE*.
16. Dash, S. (2025). Green AI: Enhancing sustainability and energy efficiency in AI-integrated enterprise systems. *IEEE Access*.
17. Dawkins, M. S. (2025). Smart farming and artificial intelligence (AI): How can we ensure that animal welfare is a priority? *Applied Animal Behaviour Science*, 283, 106519.
18. De Vries, A., & Reneau, J. K. (2010). Application of statistical process control charts to monitor changes in animal production systems. *Journal of Animal Science*, 88(suppl\_13), E11–E24.
19. Dembani, R., Karvelas, I., Akbar, N. A., Rizou, S., Tegolo, D., & Fountas, S. (2025). Agricultural data privacy and federated learning: A review of challenges and opportunities. *Computers and Electronics in Agriculture*, 232, 110048.
20. Dhaliwal, Y., Bi, H., & Neethirajan, S. (2025). Bimodal data analysis for early detection of lameness in dairy cows using artificial intelligence. *Journal of Agriculture and Food Research*, 21, 101837.
21. Dodge, J., Prewitt, T., Tachet des Combes, R., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., & Buchanan, W. (2022). Measuring the carbon intensity of AI in cloud instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)* (pp. 1877–1894). Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533234>
22. El Alaoui, A. and Mousannif, H., 2025. Enhancing weed detection through knowledge distillation and attention mechanism. *Frontiers in Robotics and AI*, 12, p.1654074. <https://doi.org/10.3389/frobt.2025.1654074>
23. El Mehdi Raouhi, M.L., Hrimech, H. and Kartit, A., 2024. Optimizing olive disease classification through transfer learning with unmanned aerial vehicle imagery. *International Journal of Electrical and Computer Engineering (IJECE)*, 14(1), pp.891-903. <https://doi.org/10.11591/ijece.v14i1.pp891-903>
24. Elbasi, E., Mostafa, N., AlArnaout, Z., Zreikat, A. I., Cina, E., Varghese, G., Shdefat, A., Topcu, A. E., Abdelbaki, W., Mathew, S., & Zaki, C. (2022). Artificial intelligence technology in the agricultural sector: A systematic literature review. *IEEE Access*, 11, 171–202.
25. Essien, D., & Neethirajan, S. (2025). Multimodal AI systems for enhanced laying hen welfare assessment and productivity optimization. *Smart Agricultural Technology*, 101564.
26. European Environment Agency. (2025). *Greenhouse gas emission intensity of electricity generation in Europe*. EEA Report No. 8/2025. <https://www.eea.europa.eu/>
27. Fu, X., Ye, W., Li, X., Zeng, X., Wang, Y., Chang, F., Zhang, J., & Liu, R. (2025). Revolutionising agri-energy: A comprehensive survey on the applications of artificial intelligence in agricultural energy internet. *Energy Internet*, 2(1), 13–27.



28. Gavai, A. K., Bouzembrak, Y., Xhani, D., Sedrakyan, G., Meuwissen, M. P., Souza, R. G. S., Marvin, H. J., & van Hillegersberg, J. (2025). Agricultural data privacy: Emerging platforms & strategies. *Food and Humanity*, 100542.
29. Gorissen, L., Konrad, K., & Turnhout, E. (2025). Sensors and sensing practices: shaping farming system strategies toward agricultural sustainability. *Agriculture and Human Values*, 1–19.
30. Gupta, U., Huang, W., Hanawal, M. K., Wang, V., Cheng, Y., et al. (2021). Chasing your long tails: Differentially private prediction rules for machine learning. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)* (pp. 8462–8474). <https://doi.org/10.48550/arXiv.2007.14839>
31. Handa, D., & Peschel, J. M. (2022). A review of monitoring techniques for livestock respiration and sounds. *Frontiers in Animal Science*, 3, 904834.
32. Hasan, S. M., Islam, T., Saifuzzaman, M., Ahmed, K. R., Huang, C. H., & Shahid, A. R. (2025). Carbon emission quantification of machine learning: A review. *IEEE Transactions on Sustainable Computing*.
33. Hiremani, V., Devadas, R. M., Sapna, R., Sowmya, T., Gujjar, P., Rani, N. S., & Bhavya, K. R. (2025). Federated learning for crop yield prediction: A comprehensive review. *MethodsX*, 103408.
34. International Energy Agency. (2024). *Renewables 2024: Analysis and forecast to 2030*. IEA Publications. <https://www.iea.org/reports/renewables-2024>
35. Jannat, A., Johnson, A., & Manriquez, D. (2025). Air quality monitoring in dairy farms: Description of air quality dynamics in a tunnel-ventilated housing barn and milking parlor. *Journal of Dairy Science*.
36. Jobarteh, B., & Neethirajan, S. (2025). Leveraging satellite data for greenhouse gas mitigation in Canadian poultry farming. *Smart Agricultural Technology*, 10, 100704.
37. Kambala, G. (2024). Emergent architectures in edge computing for low-latency applications. *International Journal of Engineering and Computer Science*, 13(09).
38. Khater, O. H., Siddiqui, A. J., El-Maleh, A., & Hossain, M. S. (2025). TinyEcoWeedNet: Edge efficient real-time aerial agricultural weed detection. *arXiv preprint arXiv:2509.18193*.
39. Kleanthous, N., Hussain, A., Khan, W., Sneddon, J., & Liatsis, P. (2022). Deep transfer learning in sheep activity recognition using accelerometer data. *Expert Systems with Applications*, 207, 117925.
40. Lagua, E. B., Mun, H. S., Ampode, K. M. B., Kim, Y. H., & Yang, C. J. (2023). Artificial intelligence for automatic monitoring of respiratory health in smart swine farming. *Animals*, 13(11), 1860.
41. Lannelongue, L., Grealey, J., & Inouye, M. (2021). Green algorithms: Quantifying the carbon footprint of computation. *Advanced Science*, 8(12), 2100707.
42. Li, G., Huang, Y., Chen, Z., Chesser, G. D., Jr., Purswell, J. L., Linhoss, J., & Zhao, Y. (2021). Practices and applications of convolutional neural network-based computer vision systems in animal farming: A review. *Sensors*, 21(4), 1492. <https://doi.org/10.3390/s21041492>
43. Li, X., & Jia, R. (2024). Energy-aware scheduling algorithm optimization for AI workloads in data centers based on renewable energy supply prediction. *Journal of Computing Innovations and Applications*, 2(2), 56–65.
44. Lu, S. and Xiao, X., 2024. Neuromorphic computing for smart agriculture. *Agriculture*, 14(11), p.1977. <https://doi.org/10.3390/agriculture14111977>
45. Liu, V., & Yin, Y. (2024). Green AI: Exploring carbon footprints, mitigation strategies, and trade-offs in large language model training. *Discover Artificial Intelligence*, 4(1), 49.
46. Liu, Y., Yao, X., Zhou, Y., Wang, H., Huang, Y., & Zhang, Y. (2022). Deep neural network compression by Tucker decomposition with nonlinear response. *Knowledge-Based Systems*, 241, 108171.
47. Losacco, C., Pugliese, G., Forte, L., Tufarelli, V., Maggiolino, A., & De Palo, P. (2025). Digital transition as a driver for sustainable tailor-made farm management: An up-to-date overview on precision livestock farming. *Agriculture*. <https://doi.org/10.3390/agriculture15131383>
48. Machuve, D., Nwankwo, E., Mduma, N. and Mbelwa, J., 2022. Poultry diseases diagnostics models using deep learning. *Frontiers in Artificial Intelligence*, 5, p.733345. <https://doi.org/10.3389/frai.2022.1016695>
49. Manghat, S. J. (2025). *Edge-cloud collaborative framework for smart IoT applications* (Master's thesis). Stockholm University.
50. Markovic, M., Li, A., Ayall, T. A., Watson, N. J., Bowler, A. L., Woods, M., Edwards, P., Ramsey, R., Beddows, M., Kuhnert, M., & Leontidis, G. (2024). Embedding AI-enabled data infrastructures for sustainability in agri-food: Soft-fruit and brewery use case perspectives. *Sensors*, 24(22), 7327.

51. Martiny, A. (2023). *Towards sustainable AI: Monitoring and analysis of carbon emissions in machine learning algorithms* (Doctoral dissertation, Politecnico di Torino).
52. Menghani, G. (2023). Efficient deep learning: A survey on making deep learning models smaller, faster, and better. *ACM Computing Surveys*, 55(12), 1–37.
53. Mughal, F. R., He, J., Das, B., Dharejo, F. A., Zhu, N., Khan, S. B., & Alzahrani, S. (2024). Adaptive federated learning for resource-constrained IoT devices through edge intelligence and multi-edge clustering. *Scientific Reports*, 14(1), 28746.
54. Muhammad Saqib, S., Iqbal, M., Tahar Ben Othman, M., Shahazad, T., Yasin Ghadi, Y., Al-Amro, S. and Mazhar, T., 2024. Lumpy skin disease diagnosis in cattle: A deep learning approach optimized with RMSProp and MobileNetV2. *PloS one*, 19(8), p.e0302862. <https://doi.org/10.1371/journal.pone.0302862>
55. Neethirajan, S. (2020). The role of sensors, big data and machine learning in modern animal farming. *Sensing and Bio-Sensing Research*, 29, 100367.
56. Neethirajan, S. (2024). Artificial intelligence and sensor innovations: Enhancing livestock welfare with a human-centric approach. *Human-Centric Intelligent Systems*, 4(1), 77–92. <https://doi.org/10.1007/s44230-023-00050-2>
57. Neethirajan, S. (2024). Net zero dairy farming—Advancing climate goals with big data and artificial intelligence. *Climate*, 12(2), 15.
58. Neethirajan, S. (2025). Agency in livestock farming—A perspective on human–animal–computer interactions. *Human-Centric Intelligent Systems*, 1–17.
59. Neethirajan, S., & Kemp, B. (2021). Digital livestock farming. *Sensing and Bio-Sensing Research*, 32, 100408.
60. Ngo, D., Park, H. C., & Kang, B. (2025). Edge intelligence: A review of deep neural network inference in resource-limited environments. *Electronics*, 14(12), 2495.
61. O'Donoghue, T., Minasny, B., & McBratney, A. (2024). Digital regenerative agriculture. *npj Sustainable Agriculture*, 2, 5. <https://doi.org/10.1038/s44264-024-00012-6>
62. O'Grady, M. J., Langton, D., & O'Hare, G. M. P. (2019). Edge computing: A tractable model for smart agriculture? *Artificial Intelligence in Agriculture*, 3, 42–51.
63. Oh, Y., Lee, N., Jeon, Y. S., & Poor, H. V. (2022). Communication-efficient federated learning via quantized compressed sensing. *IEEE Transactions on Wireless Communications*, 22(2), 1087–1100.
64. Ohamouddou, M., Ohamouddou, S., Afia, A.E. and Lasri, R., 2025. ATMS-KD: Adaptive Temperature and Mixed Sample Knowledge Distillation for a Lightweight Residual CNN in Agricultural Embedded Systems. *arXiv preprint arXiv:2508.20232*. <https://doi.org/10.48550/arXiv.2508.20232>
65. Papakonstantinou, G. I., Voulgarakis, N., Terzidou, G., Fotos, L., Giamouri, E., & Papatsiros, V. G. (2024). Precision livestock farming technology: Applications and challenges of animal welfare and climate change. *Agriculture*, 14(4), 620.
66. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
67. Phan, A. H., Sobolev, K., Sozykin, K., Ermilov, D., Gusak, J., Tichavský, P., Glukhov, V., Oseledets, I., & Cichocki, A. (2020). Stable low-rank tensor decomposition for compression of convolutional neural networks. In *Computer Vision—ECCV 2020* (pp. 522–539). Springer. [https://doi.org/10.1007/978-3-030-58526-6\\_31](https://doi.org/10.1007/978-3-030-58526-6_31)
68. Prajesh, P. J., Ragunath, K., Gordon, M., & Neethirajan, S. (2025). Satellite-based seasonal fingerprinting of methane emissions from Canadian dairy farms using Sentinel-5P. *Climate*, 13(7), 135.
69. Rao, S., & Neethirajan, S. (2025). Computational architectures for precision dairy nutrition digital twins: A technical review and implementation framework. *Sensors*, 25(16), 4899.
70. Ren, K., Bernes, G., Hetta, M., & Karlsson, J. (2021). Tracking and analysing social interactions in dairy cattle with real-time locating system and machine learning. *Journal of Systems Architecture*, 116, 102139.
71. Ren, Y., Cao, Y., Ye, C., & Cheng, X. (2023). Two-layer accumulated quantized compression for communication-efficient federated learning: TLAQC. *Scientific Reports*, 13(1), 11658.
72. Romeiko, X., Ryu, K., & Schwartz, R. (2023). A review of machine learning applications in life cycle assessment. *Sustainable Production and Consumption*, 40, 185–203.

73. Ross, S., Wang, H., Zheng, H., Yan, T., & Shirali, M. (2024). Approaches for predicting dairy cattle methane emissions: From traditional methods to machine learning. *Journal of Animal Science*, 102, skae219.
74. Różycki, R., Solarska, D. A., & Waligóra, G. (2025). Energy-aware machine learning models—A review of recent techniques and perspectives. *Energies*, 18(11), p.2810.
75. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
76. Shah-Mansouri, H. and Wong, V.W., 2018. Hierarchical fog-cloud computing for IoT systems: A computation offloading game. *IEEE Internet of Things Journal*, 5(4), pp.3246-3257.
77. Shi, Z., Chang, F., Jia, Y., Li, J., Qiu, Y., Miao, J., Jiang, W., Guo, X., Han, X. and Tang, W., 2024. Classifying and understanding of dairy cattle health using wearable inertial sensors with random forest and explainable artificial intelligence. *IEEE Sensors Letters*, 8(3), pp.1-4.
78. Stolarski, M.J., Warmiński, K., Krzyżaniak, M., Olba-Zięty, E. and Dudziec, P., 2025. The Carbon Footprint of Milk Production on a Farm. *Applied Sciences*, 15(15), p.8446.
79. Tangorra, F.M., Buoio, E., Calcante, A., Bassi, A. and Costa, A., 2024. Internet of Things (IoT): Sensors Application in Dairy Cattle Farming. *Animals*, 14 (21), 3071
80. Tapp, M., Kate, M., Zhang, S., Sailunaz, K. and Neethirajan, S., 2025. Adapting the Cool Farm Tool for Achieving Net-Zero Emissions in Agriculture in Atlantic Canada. *Sustainability*, 17(21), p.9428.
81. Thornton, P.K., Wollenberg, E.K. and Cramer, L., 2024. Livestock and Climate Change: Outlook for a more sustainable and equitable future.
82. Tincani, M., Kerouch, K., Garlando, U., Barezzi, M., Sanginario, A., Indiveri, G. and De Luca, C., 2025. A neuromorphic continuous soil monitoring system for precision irrigation. *arXiv preprint arXiv:2509.14066*. <https://doi.org/10.48550/arXiv.2509.14066>
83. Tong, Q., Wang, J., Yang, W., Wu, S., Zhang, W., Sun, C. and Xu, K., 2024. Edge AI-enabled chicken health detection based on enhanced FCOS-Lite and knowledge distillation. *Computers and Electronics in Agriculture*, 226, p.109432.
84. Tryhuba, A., Hutsol, T., Česna, J., Tryhuba, I., Mudryk, K., Francik, S., Kukharets, S., Mishchenko, I. and Oliinyk, R., 2025. Optimizing energy systems of livestock farms with computational intelligence for achieving energy autonomy. *Scientific Reports*, 15(1), p.10777.
85. Ukoba, K., Olatunji, K.O., Adeoye, E., Jen, T.C. and Madyira, D.M., 2024. Optimizing renewable energy systems through artificial intelligence: Review and future prospects. *Energy & Environment*, 35(7), pp.3833-3879.
86. Venkatasachandran, P. and Iyapparaja, M., 2024. GNViT-An enhanced image-based groundnut pest classification using Vision Transformer (ViT) model. *Plos one*, 19(3), p.e0301174. <https://doi.org/10.1371/journal.pone.0301174>
87. Verdecchia, R., Sallou, J., & Cruz, L. (2023). A systematic review of Green AI. *WIREs Data Mining and Knowledge Discovery*, 13(4), e1507. <https://doi.org/10.1002/widm.1507>
88. Wang, Y., Zhou, J., Wang, X., Yu, Q., Sun, Y., Li, Y., Zhang, Y., Shen, W. and Wei, X., 2023a. Rumen fermentation parameters prediction model for dairy cows using a stacking ensemble learning method. *Animals*, 13(4), p.678.
89. Wang, Y., Xu, X., Wang, Z., Li, R., Hua, Z. and Song, H., 2023b. ShuffleNet-Triplet: A lightweight RE-identification network for dairy cows in natural scenes. *Computers and Electronics in Agriculture*, 205, p.107632. <https://doi.org/10.1016/j.compag.2023.107632>
90. Weng, Z., Liu, K. and Zheng, Z., 2023. Cattle face detection method based on channel pruning YOLOv5 network and mobile deployment. *Journal of Intelligent & Fuzzy Systems*, 45(6), pp.10003-10020. <https://doi.org/10.3233/JIFS-232213>
91. Wu, C., Wang, M., Chu, X., Wang, K. and He, L., 2021. Low-precision floating-point arithmetic for high-performance FPGA-based CNN acceleration. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 15(1), pp.1-21.
92. Wu, C.J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Aga, F., Huang, J., Bai, C. and Gschwind, M., 2022. Sustainable ai: Environmental implications, challenges and opportunities. *Proceedings of machine learning and systems*, 4, pp.795-813.

93. Yang, C., Yin, X., Jiang, M., & Yu, H. (2024). Stable low-rank CP decomposition for compression of convolutional neural networks based on sensitivity. *Sensors*, 24(4), 1146. <https://doi.org/10.3390/s24041146>
94. Zhong, C. and Zheng, Z., 2024. MGGTSP--CAT: Integrating Temporal Convolution and LSTM for Multi--Scale Greenhouse Gas Time Series Prediction via Cross--Attention Mechanism. *Advanced Theory and Simulations*, 7(11), p.2400527

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.