

# AI-Enabled Sustainability in Indian Higher Education Institutions: Use-cases, Barriers, and the KPI–Data–Duty (KDD) Framework

Anurag Tiruwa<sup>1</sup>; Shuchi Dikshit<sup>2</sup>

<sup>1</sup>Assistant Professor, Institute of Information Technology and Management, GGSIPU, Delhi

<sup>2</sup>Assistant Professor, Fortune Institute of International Business, Delhi

Publication Date: 2026/01/06

**Abstract:** Indian higher education institutions (HEIs) operate as “mini-cities” with substantial electricity, water, mobility, and material footprints. This paper synthesizes how artificial intelligence (AI) can accelerate campus sustainability through (i) resource optimization (especially building energy and water), (ii) monitoring and operational reliability (maintenance, waste, and compliance), and (iii) behavior and mobility nudges (transport and paper reduction). Using a mini-review approach, we consolidate high-impact, campus-relevant AI applications and outline the measurement logic that links interventions to auditable sustainability indicators. We then identify key adoption barriers in Indian HEIs, including limited metering and data interoperability, procurement and skills constraints, governance and privacy concerns, and the environmental footprint of AI systems themselves. To move from isolated pilots to measurable outcomes, we propose the KPI–Data–Duty (KDD) framework, which connects a small set of time-bound sustainability KPIs to minimal viable data architecture, pilot design, and a lightweight Responsible/Green AI duty checklist. The paper contributes an implementation-oriented roadmap and use-case mapping that can support HEI leaders in planning, governing, and scaling AI-enabled sustainability initiatives with accountability.

**Keywords:** AI for Sustainability; Green Campus; Smart Buildings; Higher Education Institutions; Energy Management; Water Conservation; Responsible AI; Green AI.

**How to Cite:** Anurag Tiruwa; Shuchi Dikshit (2026) AI-Enabled Sustainability in Indian Higher Education Institutions: Use-cases, Barriers, and the KPI–Data–Duty (KDD) Framework. *International Journal of Innovative Science and Research Technology*, 11(01), 13-25. <https://doi.org/10.38124/ijisrt/26jan033>

## I. INTRODUCTION

Indian higher education institutions (HEIs) function as “mini-cities” that operate energy-intensive buildings (classrooms, laboratories, libraries, hostels), water and sanitation systems, transport and mobility services, and large administrative workflows. As a result, campuses offer a concentrated setting where sustainability interventions can be designed, measured, and institutionalized at scale—contributing to national and global sustainability priorities such as the UN Sustainable Development Goals (United Nations, 2015). In the Indian context, campus sustainability initiatives have often been driven through infrastructure upgrades (e.g., efficient lighting, solar PV, water reuse systems) and compliance/rating mechanisms. For example, Indian green campus guidance emphasizes auditable performance categories and measurable outcomes across energy, water, waste management, transportation, and health and well-being—providing a useful Key Performance Indicator (KPI) structure for institutional planning and reporting (Indian Green Building Council, 2024).

Artificial intelligence (AI) can strengthen these efforts by enabling a shift from periodic audits and static interventions to continuous, data-driven optimization. Campus operations generate regular patterns and predictable loads (timetables, semester cycles, hostel occupancy, weather-driven cooling demand), making them suitable for AI-based forecasting, anomaly detection, and automated control. In particular, buildings and utilities represent a high-impact starting point because they are major contributors to campus resource consumption, and operational improvements can be measured directly through electricity and water KPIs. Evidence from systematic reviews indicates that AI-enabled HVAC control and energy management can reduce consumption by dynamically adapting to environmental conditions and occupancy patterns (Aghili et al., 2025). Beyond energy, AI approaches can support leak detection and consumption analytics for water conservation, predictive maintenance to reduce asset failures and wastage, and improved waste segregation and routing—thereby strengthening operational reliability while advancing sustainability outcomes.

However, HEIs face a dual challenge: deploying AI responsibly while ensuring that AI itself aligns with sustainability goals. Governance concerns arise because campus AI can involve sensitive data (e.g., occupancy signals, mobility traces, or administrative records), requiring clear policies on privacy, accountability, and human oversight. A recognized risk-management approach such as the NIST AI Risk Management Framework provides a structured vocabulary for documenting risks, controls, and monitoring throughout the AI lifecycle (National Institute of Standards and Technology, 2023). At the same time, “Green AI” scholarship highlights that AI systems have their own energy and carbon footprint, and argues for prioritizing efficiency and reporting compute-related costs rather than optimizing accuracy alone—an important consideration for budget-constrained HEIs (Schwartz et al., 2020).

Against this backdrop, the present paper (a) synthesizes campus-relevant AI use-cases for sustainability, (b) identifies adoption barriers in Indian HEIs, and (c) proposes a practical, implementation-oriented framework that links sustainability KPIs to data readiness, pilot design, and responsible governance—supporting HEI leaders in moving from isolated pilots to scalable, measurable sustainability outcomes (Indian Green Building Council, 2024; National Institute of Standards and Technology, 2023).

## II. LITERATURE REVIEW

### A. Campus Sustainability Measurement in India: Rating Frameworks and KPI Logic

Campus sustainability work in India is often operationalized through auditable categories and measurable performance targets embedded in green building/campus rating frameworks. The IGBC Green Campus Rating System (Version 1.0, July 2024) emphasizes measurable outcomes across energy, water, waste, transportation, and related campus systems, providing a practical structure for KPI selection and reporting (Indian Green Building Council, 2024). In parallel, GRIHA positions itself as a nationally benchmarked green building rating approach, with criteria spanning energy and water optimization, waste management, and operations & maintenance—areas that map cleanly to campus sustainability interventions (GRIHA Council, n.d.).

A key implication for AI-enabled sustainability is that successful deployments are rarely “model-first.” Instead, they follow a measurement logic: (i) define KPI baselines and targets, (ii) ensure data integrity and instrumentation, and (iii) implement iterative operational controls and verification cycles—an approach consistent with energy management thinking such as ISO 50001, which focuses on systematic improvement of energy performance (International Organization for Standardization, 2018).

### B. AI for Building Energy Management: HVAC, Occupancy Intelligence, and Plug Loads

Buildings typically dominate institutional energy footprints, making them the most evidence-rich domain for AI-enabled sustainability. Reviews of AI in HVAC and building energy management consistently report benefits from forecasting, anomaly detection, and control optimization that

adapt to weather and occupancy dynamics. A recent systematic review synthesizing post-2018 research reports that AI-driven HVAC strategies can reduce energy use in certain settings by dynamically adapting to environmental and occupancy conditions (Aghili et al., 2025).

Occupancy intelligence is central to this literature because many energy loads in academic buildings are schedule-driven yet highly variable (class timetables, lab use, hostels, events). A major review on machine-learning-based occupancy prediction highlights that occupancy-informed control can meaningfully reduce energy use while maintaining comfort, emphasizing the importance of reliable occupancy sensing or privacy-preserving proxy signals (Zhang et al., 2022). Complementing HVAC optimization, research on plug-load prediction shows that incorporating occupant count improves prediction accuracy, which is valuable for model-predictive control and operational scheduling (Wang et al., 2019).

Across this stream, two implementation lessons recur: (1) performance depends on data quality and instrumentation (sub-metering, reliable sensors, contextual features), and (2) operational adoption improves when models are interpretable and embedded into facility workflows rather than treated as “black box automation” (Aghili et al., 2025; Zhang et al., 2022). These insights motivate a KPI-linked, instrumentation-first deployment pathway.

### C. AI for Water Systems: Smart Metering, Anomaly Detection, and Leakage Analytics

Water conservation and loss reduction are increasingly addressed through AI-enabled analytics on smart meters, district/zone metering, and distribution network telemetry. A systematic review of ML-based anomaly detection in smart water metering networks (2016–2023 coverage) highlights how models are used to identify abnormal consumption patterns, leaks, and operational anomalies, while noting challenges around ground-truth labeling and generalizability across sites (Kanyama et al., 2024). Recent work continues to develop ensemble approaches for anomaly detection in smart water metering, reinforcing that AI methods support conservation when paired with adequate sensing and operational verification protocols (Kanyama et al., 2025).

For HEIs, the literature implies a practical sequencing: start with zone-level measurement + anomaly detection (to reduce losses quickly), then move toward richer optimization (e.g., predictive consumption forecasting), while maintaining low-friction operator verification to manage false alarms (Kanyama et al., 2024, 2025).

### D. AI for Waste: Segregation Support and Collection Logistics

In waste management, AI work splits into (i) segregation support (often computer vision classification) and (ii) collection/logistics optimization (routing, scheduling, bin fill prediction). A 2025 study in *Knowledge-Based Systems* presents a deep learning–enabled waste classification system and applies explainable AI methods to support interpretability, illustrating the direction of high-accuracy classification pipelines for segregation workflows (Nahiduzzaman et al.,

2025). On the logistics side, a systematic review and meta-analysis of IoT-enabled routing optimization quantifies average reductions in waste collection distance, implying emissions and cost benefits when routing analytics are integrated with operational systems (Maciel et al., 2025).

For campuses, the research suggests that model performance alone is insufficient; operational value depends on workflow redesign (bin placement, segregation incentives, collection schedules) and governance choices that avoid intrusive surveillance in public spaces (Maciel et al., 2025; Nahiduzzaman et al., 2025).

#### *E. Mobility Optimization and Behavioral Change: Routing, Demand-Responsive Transport, and Nudges*

Transport and commuting are prominent categories in campus sustainability frameworks (e.g., IGBC's transportation emphasis), and the literature offers multiple AI levers: shuttle routing, demand estimation, and demand-responsive service design (IGBC, 2024). At the campus scale, a study on university shuttle bus optimization uses ML for travel-time prediction while considering fuel cost and emissions, showing how predictive analytics can support routing decisions and monitoring (Noor et al., 2020). At the broader mobility-systems level, work on demand-responsive transport integrates mobile data analytics and routing algorithms to support flexible services that better match demand while improving efficiency (Melo et al., 2024).

Behavioral interventions complement operational optimization where user actions shape resource use (lighting/plug loads, water use, waste segregation). Evidence from student-accommodation contexts indicates that informational and competition-based nudges can reduce energy use, though effect sizes vary and persistence depends on design (Chen & Lotti, 2025). Longer-term field evidence comparing “nudges” and “boosts” in residential energy settings also suggests that intervention mechanisms and durability matter for sustained outcomes (Paunov & Grüne-Yanoff, 2023).

#### *F. Governance, Privacy, and “Sustainable AI”: Responsible AI + Green AI*

AI-enabled campus sustainability introduces a governance duality: (i) responsible processing of campus data and (ii) the environmental footprint of AI itself. On data protection, India's Digital Personal Data Protection Act, 2023 establishes obligations relevant to campuses acting as data fiduciaries—particularly around notice, purpose limitation, retention, and grievance redressal (Government of India, 2023). The DPDP Rules, 2025 provide the implementation layer and mark operationalization of compliance expectations (Ministry of Electronics and Information Technology [MeitY], 2025; Press Information Bureau, 2025).

In addition, India-specific Responsible AI guidance emphasizes principles for trust-building and alignment with societal values—highly relevant where occupancy signals, camera feeds, or mobility traces may be used (NITI Aayog, 2021a, 2021b). At the risk-governance level, the NIST AI Risk Management Framework (AI RMF 1.0) provides a widely used structure for mapping, measuring, and managing risks

across the AI lifecycle, supporting institutional controls such as transparency, accountability, security, and human oversight (National Institute of Standards and Technology, 2023).

On AI's environmental footprint, Green AI scholarship argues for reporting and valuing compute/efficiency alongside accuracy, since resource-intensive model development has both financial and carbon implications (Schwartz et al., 2020). A more recent review consolidates green AI research and reinforces efficiency-oriented approaches as central to sustainable AI practice (Bolón-Canedo & Morán-Fernández, 2024). These governance and sustainability concerns align naturally with an explicit “duty” layer in AI-for-sustainability implementations.

#### *G. AI Synthesis and Research Gap Motivating the KPI–Data–Duty Approach*

Across domains (energy, water, waste, mobility), the literature demonstrates a growing toolbox of AI techniques and a steady expansion of evidence syntheses and applied studies (Aghili et al., 2025; Kanyama et al., 2024; Maciel et al., 2025; Melo et al., 2024). However, three gaps remain especially salient for Indian HEIs:

##### ➤ *KPI-to-Intervention Traceability*

- Many studies report algorithmic performance or localized savings without a standardized, campus-administration-friendly linkage between sustainability KPIs (baseline → target → audit) and AI deployment decisions. IGBC/GRIHA provide KPI categories but do not prescribe an AI implementation pathway (GRIHA Council, n.d.; IGBC, 2024).

##### ➤ *Data Readiness as the Constraint:*

Repeatedly, outcomes depend less on model sophistication and more on instrumentation, interoperability, and operational data pipelines—suggesting a “minimum viable data stack” mindset before optimization (Zhang et al., 2022; Wang et al., 2019).

- Governance embedded early: With DPDP obligations, Responsible AI principles, NIST-aligned risk management, and Green AI concerns, campuses need a lightweight but explicit governance checklist that travels with the pilot-to-scale journey rather than being added after deployment (MeitY, 2025; NIST, 2023; NITI Aayog, 2021a; Schwartz et al., 2020).

Collectively, these gaps justify an implementation-oriented synthesis that provides a practical, auditable route from “AI idea” to “measured sustainability outcome.”

### **III. METHOD: MINI-REVIEW**

This study used a mini-review approach to identify and synthesize campus-relevant AI applications for sustainability, with emphasis on energy, water, waste, mobility, and administrative dematerialization (paper reduction). A mini-review was selected because the objective was not exhaustive coverage of all AI and sustainability research, but a focused, decision-oriented synthesis that translates the most relevant evidence into (i) a campus use-case taxonomy and (ii) implementation implications suitable for Indian higher

education institutions (HEIs), following established guidance on planning and reporting literature reviews (Grant & Booth, 2009; Snyder, 2019).

#### A. Review Scope and Unit of Analysis

The scope was defined around operational sustainability systems typical to HEIs (e.g., buildings, utilities, campus services, and mobility) rather than city-scale infrastructure. The unit of analysis was campus-scale deployment, meaning applications needed to be plausibly implementable at the level of individual buildings, hostels, labs, or the campus network (rather than purely theoretical algorithms without a deployment context).

#### B. Search Strategy and Sources

Searches were conducted using combinations of keywords structured around three blocks:

- Technology block: “AI”, “machine learning”, “deep learning”, “predictive analytics”, “anomaly detection”, “optimization”
- Context block: “campus”, “university”, “higher education”, “smart campus”, “smart building”
- Sustainability block: “energy”, “HVAC”, “building energy”, “water”, “leakage”, “waste”, “segregation”, “routing”, “mobility”, “transport”, “paper reduction”, “administrative processes”

#### ➤ Example Query Patterns Included:

- “AI” AND (“campus” OR “university”) AND (“energy” OR “HVAC” OR “smart building”)
- “machine learning” AND (“smart water” OR “smart metering”) AND (“anomaly detection” OR “leak detection”)
- “AI” AND (“waste” OR “segregation” OR “routing”) AND (“IoT” OR “collection”)
- “AI” AND (“mobility” OR “shuttle” OR “demand-responsive transport”) AND (“optimization” OR “prediction”)

To improve contextual relevance, India-specific terms were used such as “green campus”, “IGBC green campus”, “sustainability in Indian universities”, and “campus sustainability India”, with purposeful attention to practice-facing documents and standards that define KPI categories (e.g., Indian Green Building Council, 2024).

#### C. Selection Priorities and Rationale

Given the mini-review intent, priority was given to sources most likely to generalize into implementable campus guidance:

- Review papers and synthesis articles (to consolidate evidence across multiple studies and reduce over-reliance on single-case claims) (Grant & Booth, 2009; Snyder, 2019).
- Building energy/HVAC evidence because buildings are typically the most metered and controllable campus domain and have mature AI applications (e.g., Aghili et al., 2025).
- India-relevant sustainability rating frameworks and guidance to ensure the review remains KPI-compatible

with local reporting and benchmarking norms (Indian Green Building Council, 2024).

- Governance and risk frameworks to ensure proposed implementation guidance explicitly addresses transparency, accountability, and responsible deployment (National Institute of Standards and Technology, 2023).

#### D. Inclusion and Exclusion Criteria

Sources were included if they met all of the following criteria:

- Application clarity: clear description of an AI technique and the sustainability application context (e.g., HVAC control, anomaly detection in water metering, waste routing).
- Outcome relevance: reported operational outcomes, measurable impacts, or a credible pathway to KPI measurement (e.g., kWh savings, reduced downtime, leak reduction, route distance reduction).
- Campus deployability: plausibility for campus-scale deployment (instrumentation requirements and implementation context could be reasonably mapped to HEI conditions).

Sources were excluded if they:

- Focused purely on algorithmic novelty without a sustainability outcome or deployment context,
- Addressed sustainability at a scale far removed from HEIs without a clear translation pathway, or
- Lacked sufficient methodological detail to interpret claims.

#### E. Screening and Data Extraction

Retrieved sources were screened in two stages:

- Stage 1 (title/abstract screening): removal of off-topic items and non-campus-relevant studies.
- Stage 2 (full-text screening): evaluation against the inclusion criteria above, with attention to the fit between technique, instrumentation, and KPI measurability.

A structured extraction template was used to standardize interpretation across domains. For each included source, the following information was captured:

- Sustainability domain energy/ water/ waste/ mobility/ admin),
- AI technique family (forecasting, anomaly detection, classification, optimization/control),
- Required data inputs and instrumentation,
- Reported outcomes and KPI candidates,
- Implementation conditions and constraints (data quality, integration, human oversight), and
- Transferability to campus settings (what would need to be in place at an HEI).

#### F. Synthesis Approach

Evidence was synthesized using a narrative thematic synthesis consistent with mini-review practice (Grant & Booth, 2009; Snyder, 2019). Findings were coded into two complementary outputs:

- Use-case taxonomy: grouping applications by campus function and technique type (e.g., “HVAC optimization,” “water anomaly detection,” “waste routing”).

- Implementation implications: cross-cutting lessons that determine whether AI translates into measurable outcomes (e.g., instrumentation readiness, workflow integration, governance safeguards).

Rather than treating all studies as equally generalizable, the synthesis emphasized repeatable patterns (e.g., instrumentation-first logic, verification loops, integration with facilities operations) and highlighted where evidence is context-dependent.

#### G. Governance Lens for Risk and Responsibility

Since, campus sustainability AI can involve potentially sensitive signals (occupancy proxies, mobility traces, camera-adjacent data), governance considerations were integrated explicitly. Responsible AI and risk guidance was consulted to ensure that the proposed framework incorporates transparency, accountability, security, and oversight mechanisms appropriate for institutional deployment (National Institute of Standards and Technology, 2023). This governance lens was operationalized in the synthesis by identifying “duty” requirements alongside each domain (e.g., data minimization, access control, documentation, and human-in-the-loop verification).

#### H. Methodological Limitations

As a mini-review, the method prioritizes relevance and synthesis utility over exhaustive retrieval and formal meta-analysis. Therefore, results should be interpreted as a structured, implementation-oriented consolidation rather than a comprehensive mapping of all published studies. Nonetheless, the use of explicit search blocks, staged screening, and standardized extraction improves transparency and reproducibility in line with literature review guidance (Grant & Booth, 2009; Snyder, 2019).

Following guidance for transparent review reporting (Grant & Booth, 2009; Snyder, 2019), the search strategy across the defined keyword blocks (AI + campus/university + energy/HVAC/water/mobility/dematerialization, plus India-specific terms such as “IGBC green campus”) yielded 312 records. After removing duplicates, 254 unique records remained. Title and abstract screening for campus deployability and sustainability relevance excluded 186 records, resulting in 68 articles retained for full-text assessment. Full texts were then evaluated against the inclusion criteria (clear AI technique and campus context; operational outcome or KPI measurability; feasible instrumentation/deployment logic), leading to the exclusion of 35 studies (e.g., insufficient deployment context, limited KPI linkage, or non-campus scale). The final synthesis included 33 sources, which were coded into (i) a campus use-case taxonomy and (ii) cross-cutting implementation implications (instrumentation readiness, workflow integration, verification cycles), with governance and risk considerations aligned to a lifecycle risk lens (NIST, 2023) to support transparency, accountability, and oversight.

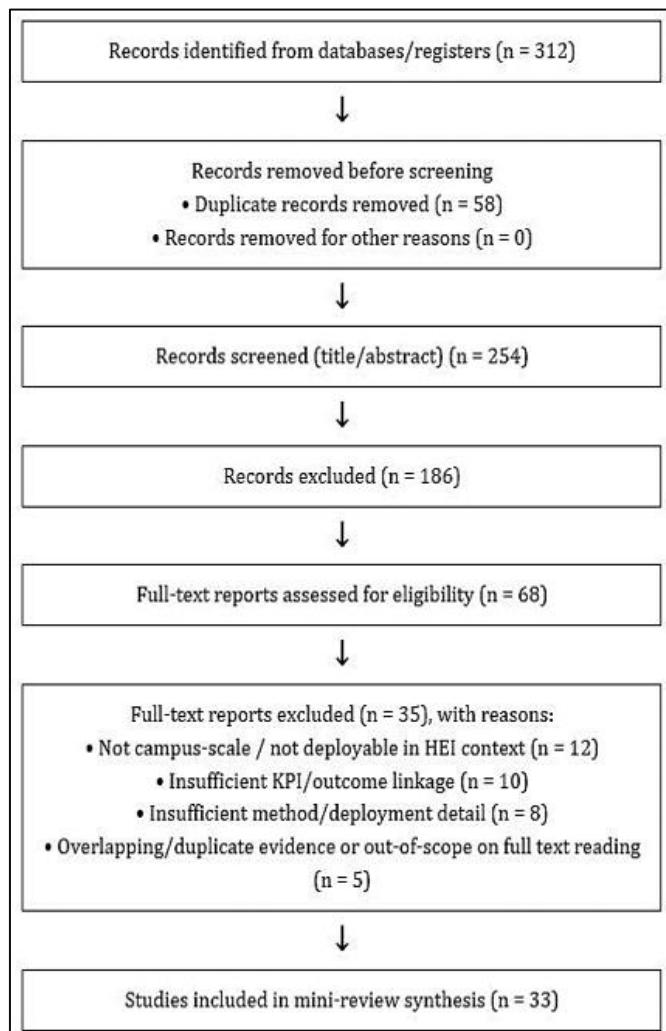


Fig 1 PRISMA Flow Diagram for the Mini-Review

## IV. AI USE-CASES FOR A SUSTAINABLE CAMPUS

### A. Resource Optimization

- Smart HVAC and thermal comfort optimization. AI-enabled HVAC control is consistently highlighted as a major lever because HVAC loads vary with occupancy, schedules, and weather. A systematic review of AI approaches for HVAC energy management reports that AI-driven control strategies can reduce energy consumption by dynamically adapting to environmental conditions and occupancy patterns (Aghili et al., 2025). For Indian HEIs, where classroom utilization fluctuates by timetable and season, even incremental optimization can yield substantial savings without compromising comfort.
- Lighting and plug-load management. Occupancy-based control (sensor-driven) combined with predictive scheduling (timetables + historical usage) can reduce waste from lights/fans/AC running in unoccupied rooms. Evidence from state-of-the-art reviews and empirical studies shows meaningful energy-saving potential from occupancy-based lighting control (de Bakker et al., 2017; de Bakker et al., 2018). In parallel, plug-load prediction and scheduling research demonstrates that occupancy/usage-aware forecasting can support predictive control and operational savings (Botman et al., 2024; Wang et al., 2019). These are “low-regret” applications

because they are explainable and KPI-linked (kWh reduction).

- Renewable integration and demand management. Campuses with rooftop solar and time-of-day tariffs can use AI forecasting (load + solar generation) to shape demand (e.g., shifting non-critical loads), improve self-consumption, and reduce peak costs. Campus microgrid research shows that demand-response-oriented energy management can deliver measurable cost savings and improve renewable utilization (Bin et al., 2022). AI-based solar and load forecasting studies further support optimized scheduling and reliability in renewable-heavy systems (Bouquet et al., 2024; Rajbhandari et al., 2024).

#### *B. Monitoring, Maintenance, and Compliance*

- Predictive maintenance for utilities and equipment. AI can detect anomalies in pumps, chillers, and electrical panels using sensor/IoT time-series data, enabling earlier interventions that reduce unplanned downtime and improve operational reliability (Es-sakali et al., 2022; Murtaza et al., 2024). The sustainability benefit is indirect but meaningful: fewer failures, reduced wastage, and longer asset life (Murtaza et al., 2024).
- Water leak detection and consumption analytics. AI-based anomaly detection on smart-meter data can identify continuous-flow and abnormal-consumption patterns (Kanyama et al., 2024; Kanyama et al., 2025). India-facing campus guidance emphasizes measurable water conservation and reduction targets, making leak analytics a strong fit for institutional KPIs and compliance reporting in campus sustainability programs (Indian Green Building Council, 2024).
- Waste segregation support and routing. Computer vision models can support waste identification/classification for segregation at source (e.g., canteens/hostels), which can reduce contamination and improve downstream recovery (Nahiduzzaman et al., 2025). For internal collection, IoT- and analytics-enabled route optimization has been synthesized to reduce collection distance on average, supporting lower fuel use and emissions (Maciel et al., 2025). Practical success often depends on workflow design and incentives in addition to model performance (Maciel et al., 2025).

Appendix 1 summarizes the identified AI use-cases mapped to campus KPIs, SDGs, data requirements, and key risks/controls.

#### **V. CONSTRAINTS ON ADOPTING AI-ENABLED SUSTAINABILITY INITIATIVES IN INDIAN HEIS**

Adoption of AI-enabled sustainability initiatives in Indian higher education institutions (HEIs) is constrained by multiple, interlinked challenges.

- First, many campuses still lack consistent metering and sensing, such as sub-metering for electricity and water, reliable occupancy sensing, and structured telemetry for key assets (HVAC plants, pumps, elevators, major plug-load zones). Where data exist, they are often incomplete, noisy, or temporally misaligned across systems, which reduces the feasibility of forecasting, anomaly detection,

and closed-loop control. Evidence from building energy management reviews shows that AI performance and realized savings are strongly contingent on data quality, coverage, and contextual features (e.g., weather, schedules, occupancy proxies), not only on model choice (Aghili et al., 2025; Zhang et al., 2022). When instrumentation is weak, AI projects risk becoming “pilot theatre”—producing dashboards and prototypes without repeatable, auditable KPI improvement.

- Second, available data are frequently siloed across facility teams, contractors, and vendors (BMS, smart meters, ticketing/CMMS, ERP/procurement, student information systems), creating interoperability and integration burdens. Smart-campus IoT deployments illustrate that multi-vendor environments routinely face integration challenges (heterogeneous protocols, inconsistent semantics), along with scalable storage/processing and reliable visualization—issues that must be addressed before analytics can be operationalized campus-wide (Domínguez-Bolaño et al., 2024). In practice, campuses often need a “minimum viable data stack” (instrumentation + integration + governance) before optimization.
- Third, procurement rigidity and limited in-house data engineering capacity can increase dependence on vendors, limiting iterative refinement and raising the risk of lock-in (long contracts, proprietary formats/APIs, restricted portability). Research on vendor lock-in in cloud migration highlights how portability and interoperability constraints can restrict switching options and create long-term dependence (Opara-Martins et al., 2016). Public ICT procurement research also documents how tendering processes can constrain agility (changes are hard after tender publication), and how large, complex procurements can unintentionally reinforce lock-in dynamics (Ghezzi & Mikkonen, 2024). For HEIs, this means that even technically sound pilots may stall at scale if contracts do not explicitly protect data access, interoperability, and exit pathways.
- Fourth, governance, privacy, and trust concerns are central because campus AI can involve potentially sensitive data (occupancy traces, Wi-Fi logs, CCTV feeds, mobility patterns, or linked student/staff identifiers). Studies on smart-campus governance show that students can experience “tensions and imaginaries” around surveillance, consent, and institutional power, which can undermine legitimacy even when the stated purpose is sustainability (Cheong & Nyaupane, 2022). In India, these concerns are sharpened by the Digital Personal Data Protection (DPDP) legal regime: the DPDP Act, 2023 sets obligations relevant to campuses acting as data fiduciaries, and the DPDP Rules, 2025 further operationalize requirements (e.g., governance, notices/consent where applicable, grievance mechanisms, safeguards, retention discipline). For implementation, this implies clear policies for purpose limitation, data minimization, retention periods, access controls, audit logs, and human oversight—especially for AI that could be repurposed beyond sustainability.
- Fifth, the sustainability of AI itself must be considered. Training and deploying models consume compute, electricity, and associated carbon, and “Green AI” argues that efficiency and compute cost should be reported and

valued alongside accuracy. Empirical work on deep learning also emphasizes that energy and carbon costs can be material and unequally distributed, with practical recommendations to reduce compute intensity. For budget-constrained academic settings, this means prioritizing lightweight models where feasible, reusing pretrained components responsibly, and logging compute/energy metrics for accountability.

Finally, change management and human factors remain decisive. Facilities teams may resist “black-box” automation, and campus stakeholders may distrust monitoring unless benefits are visible and safeguards are explicit. Explainability research shows that interpretable, instance-level explanations can improve users’ ability to assess and trust model outputs—supporting adoption in operational contexts where decisions must be justified (Ribeiro et al., 2016). Similarly, structured risk governance frameworks (e.g., NIST AI RMF 1.0) provide a practical vocabulary for documenting risks and controls across the lifecycle—helping move from pilot to scale without eroding trust.

## VI. PRACTICAL FRAMEWORK FOR HEIS

The KPI–Data–Duty (KDD) framework guides Indian HEIs in implementing AI for campus sustainability through a disciplined pilot-to-scale pathway that prioritizes measurable outcomes, data readiness, and responsible deployment.

➤ First, the framework begins with KPI selection (“KPI”): institutions define 3–5 auditable, time-bound sustainability KPIs that are meaningful for campus operations and align with recognized green campus/building categories. For example, KPI sets can be drawn from structured domains such as energy, water, waste, transportation/mobility, and operations & maintenance used in Indian rating and benchmarking systems (Indian Green Building Council [IGBC], 2024; GRIHA Council, n.d.). Each KPI is defined with (i) a clear formula and data source, (ii) a measurement frequency (daily/weekly/monthly), (iii) a KPI owner (facility lead/estate office/hostel admin), and (iv) a baseline and target window. This KPI discipline ensures AI pilots are not evaluated only on “model accuracy,” but on whether they move operational metrics that can be reported and audited.

➤ Second, KDD establishes feasibility via a minimum viable “Data” stack (“Data”), following the principle “instrument first, optimize second.” This step is a gate: if minimum instrumentation and integration are absent, the institution improves data readiness before attempting advanced AI. Building-energy and HVAC research repeatedly shows that realized energy benefits depend heavily on data quality, metering granularity, and contextual variables (weather, schedules, occupancy proxies), not just algorithm choice (Aghili et al., 2025; Zhang et al., 2022). Practically, a minimal stack typically includes:

- Sub-metering at building/zone/end-use level (HVAC plant, hostels, labs where possible)
- Occupancy proxies (aggregated counts or schedule signals) to contextualize demand variability (Zhang et al., 2022)

- Weather and tariff inputs to normalize performance and support peak/price-aware decisions
- Maintenance logs / CMMS tickets to connect anomalies to verified action and closure
- Interoperability basics (timestamps, naming conventions, a simple data dictionary, and access to APIs/exports), because multi-vendor campus systems often fail to scale without integration discipline (Domínguez-Bolaño et al., 2024).
- Where relevant (e.g., plug loads and internal gains), the minimum stack may also include occupant-count-linked signals because occupant count can materially improve prediction performance and control relevance (Wang et al., 2019).

➤ Third, KDD embeds a lightweight governance and sustainability checklist (“Duty”) to ensure AI is both responsible and resource-efficient from day one. The Duty checklist covers:

- Privacy & purpose limitation (what data is collected, why it is needed, and what it will *not* be used for), plus retention and access discipline consistent with India’s DPDP framework (MeitY, 2023, 2025).
- Transparency & accountability (who approves, who monitors drift, who handles grievances and exceptions)
- Security controls (role-based access, audit logs, vendor safeguards)
- Human oversight (AI as decision support, with override and escalation paths)
- Risk documentation alignment using a recognized lifecycle vocabulary such as NIST AI RMF 1.0 (National Institute of Standards and Technology, 2023).
- Green AI logging: track compute intensity, retraining frequency, and efficiency trade-offs—reflecting “Green AI” guidance that efficiency should be valued alongside performance (Schwartz et al., 2020) and evidence that training large models can have non-trivial energy/carbon implications (Strubell et al., 2019).
- Interpretability support so facility teams can trust and act on outputs; explainability methods are widely used to make predictions more actionable and auditable in operational settings (Ribeiro et al., 2016).

➤ Finally, KDD operationalizes learning through an 8–12 week pilot and scale decision. The pilot is implemented in a limited, well-defined zone (one academic block/hostel cluster/utility subsystem) with (i) a pre-specified evaluation plan, (ii) comparison against a seasonal/academic-calendar baseline, and (iii) a workflow that links AI outputs to action (alert → verification → ticket → fix → KPI update). Baseline–target improvement logic mirrors systematic performance improvement thinking used in energy management system practice (ISO, 2018). The framework recommends scaling only when KPI improvements are demonstrable *and* Duty controls are stable, preventing the common failure mode of scaling prototypes without governance readiness.

Appendix 2 presents the KPI–Data–Duty (KDD) framework as a concise pilot-to-scale roadmap, enabling HEI leaders to translate AI ideas into measured sustainability

outcomes with traceable KPIs, minimum viable instrumentation, and embedded responsible/Green AI safeguards.

## VII. DISCUSSION & IMPLICATIONS

For Indian HEIs, the most practical starting point for AI-enabled campus sustainability is building energy and water. These domains offer three advantages. First, outcomes are highly measurable through standard indicators (e.g., kWh, peak demand, kWh/m<sup>2</sup>; litres per day, litres per capita, anomaly/leak events) once minimum metering exists. Second, interventions are predominantly operational and system-level (controls, scheduling, maintenance response, leak repair), reducing the risk that projects are perceived as personal surveillance. Third, the evidence base is relatively mature: AI applications for HVAC and building energy management are widely studied and show that forecasting, fault detection, and control optimization can reduce energy use when data and operational integration are adequate (Aghili et al., 2025; Zhang et al., 2022). Similarly, smart-water anomaly detection literature indicates that machine learning can support leak/abnormal-consumption detection when paired with appropriate sensing and verification routines (Kanyama et al., 2024).

Campus sustainability rating frameworks provide an additional institutional lever because they act as a shared administrative language. In India, the IGBC Green Campus Rating framework explicitly foregrounds performance categories such as energy, water, waste, and transportation, which helps facility teams and administrators align on what counts as “progress” and how it should be reported (Indian Green Building Council, 2024). GRIHA’s criteria and benchmarking orientation similarly reinforce a multi-criterion, measurable approach to sustainability management (GRIHA Council, n.d.). These frameworks make it easier to translate AI work into auditable claims—reducing the likelihood that pilots remain isolated technology demonstrations.

At the same time, HEIs should avoid treating AI as a “technology add-on.” The implication of the literature is that AI value is realized only when it is embedded into a measurement-and-improvement operating cycle rather than deployed as a standalone analytics layer. This is why the KDD framing begins with KPIs and instrumentation and then links outputs to operational routines such as maintenance workflows, control actions, and verified closure (ISO, 2018; Aghili et al., 2025). The key adoption lesson here is socio-technical: energy/water AI performs best when it becomes part of how the campus runs—how anomalies become tickets, how controls are adjusted, how baselines are updated, and how results are reported.

Governance and trust are not peripheral to sustainability AI; they are often the conditions for scaling. Even “non-personal” domains can drift into privacy risk once occupancy proxies, Wi-Fi counts, access logs, or camera-adjacent data are introduced. Embedding governance early—through documentation, access controls, purpose limitation, retention rules, and human oversight—reduces reputational and compliance risk and increases stakeholder acceptance. A

useful institutional anchor is a lifecycle risk lens such as the NIST AI Risk Management Framework (AI RMF 1.0), which provides a practical vocabulary for “govern–map–measure–manage” across design, deployment, monitoring, and incident handling (NIST, 2023). From an implementation perspective, the governance implication is straightforward: HEIs should require a lightweight “duty pack” for every pilot (what data is used, why, who can access it, how long it is kept, how errors are handled), and scale only when these controls are stable.

A further implication—often overlooked in campus sustainability discussions—is that the environmental footprint of AI itself must be managed. “Green AI” argues that efficiency and compute-related costs should be reported and valued alongside accuracy and performance (Schwartz et al., 2020). Related work demonstrates that training and running models can have meaningful energy and carbon costs, making efficiency a practical governance concern rather than a purely ethical one (Strubell et al., 2019). For HEIs, operationalizing Green AI means preferring efficient models where possible, avoiding unnecessary retraining, and logging compute and energy proxies as part of sustainability reporting. Doing so also aligns with academic values of transparency and replicability—where methods are documented, claims are verifiable, and performance trade-offs are explicit (Schwartz et al., 2020).

### A. Practical Implications for HEI Leaders and Facility Teams

- Start where measurement is easiest: prioritize energy and water pilots with clear KPIs and baseline definitions (IGBC, 2024; Aghili et al., 2025; Kanyama et al., 2024).
- Build minimum viable instrumentation before “AI”: sub-metering and reliable logs are often higher ROI than sophisticated models in low-data environments (Zhang et al., 2022; ISO, 2018).
- Treat AI as decision support: integrate outputs into maintenance and operations, with clear escalation and override paths (NIST, 2023).
- Make Green AI a requirement: log compute/efficiency indicators and justify model complexity relative to marginal sustainability gains (Schwartz et al., 2020; Strubell et al., 2019).

### B. Research Implications

For researchers, the KDD logic implies that the strongest contributions will be those that connect models to auditable KPIs and document the full deployment chain (instrumentation → model → action → verification). Future work can strengthen the evidence base by reporting standardized baselines, seasonal controls, maintenance-response confounds, and durability of impact beyond pilot windows—especially in HEI settings where academic calendars create pronounced demand variability (Zhang et al., 2022; Aghili et al., 2025).

## VIII. LIMITATIONS & FUTURE WORK

This study adopted a mini-review approach that prioritizes transferable, campus-relevant AI applications and standards over exhaustive retrieval, formal risk-of-bias assessment, or meta-analysis. As such, it should be interpreted as an implementation-oriented synthesis rather than a full

systematic review, consistent with the known trade-offs of review types (Grant & Booth, 2009; Snyder, 2019). While transparency was strengthened through structured search blocks and staged screening, the process does not claim comprehensive coverage comparable to PRISMA-grade systematic reviews (Page et al., 2021).

A second limitation is context sensitivity of reported savings. Energy and water outcomes summarized from reviews can vary substantially by building typology (labs vs classrooms vs hostels), climate zone, operational schedules, occupancy volatility, and baseline efficiency. For example, the building-energy literature shows that AI impacts depend strongly on instrumentation quality, weather normalization, and occupancy dynamics, which may differ across campuses and regions (Aghili et al., 2025; Zhang et al., 2022). Similarly, anomaly detection in water systems is affected by metering granularity and ground-truth verification practices, which can vary widely across institutions (Kanyama et al., 2024). Therefore, the framework's pathway is robust, but effect sizes should not be assumed without local baselines and verification.

Third, the paper does not introduce a primary dataset or operational trial results, so claims about feasibility and ROI are based on synthesized evidence rather than demonstrated outcomes from Indian HEI deployments. This is important because campus environments involve socio-technical constraints (maintenance response capacity, procurement constraints, stakeholder trust) that can limit real-world performance even when models work in controlled studies. Future research can strengthen and validate the KDD framework in three high-value directions:

- Multi-campus validation across Indian contexts: Implement coordinated pilots across multiple HEIs (e.g., different climate zones, campus sizes, building stock ages) using shared KPI definitions and standardized reporting, to test generalizability and isolate which enabling conditions (metering maturity, governance readiness, O&M capacity) drive outcomes (Aghili et al., 2025; Zhang et al., 2022).
- Comparative evaluation of Green AI choices (energy/compute cost vs accuracy vs KPI gain): Conduct head-to-head comparisons of model families (lightweight statistical baselines, tree-based models, deep learning, hybrid control) while explicitly reporting compute/efficiency and retraining requirements. This aligns with Green AI arguments that efficiency should be valued and reported alongside performance (Schwartz et al., 2020) and with evidence that model development/training can have non-trivial energy and carbon costs (Strubell et al., 2019). More recent syntheses can be used to build standardized reporting templates for AI footprint and efficiency trade-offs (Verdecchia et al., 2023; Bolón-Canedo & Morán-Fernández, 2024).
- Governance and acceptability research on monitoring boundaries in HEIs: Empirically examine what levels of monitoring (occupancy proxies, access logs, mobility traces, camera-adjacent systems) are considered acceptable by students, staff, and administrators; what consent/notice designs increase legitimacy; and how governance mechanisms shape adoption. Smart campus work shows that perceived surveillance and data

governance tensions materially affect legitimacy and trust (Cheong & Nyaupane, 2022). These studies should be linked to practical lifecycle controls and documentation frameworks (NIST, 2023) and to evolving legal obligations under India's DPDP regime (MeitY, 2023, 2025).

Collectively, these directions would shift the contribution from a framework grounded in synthesis toward a validated, benchmarked, and governance-tested implementation model for Indian HEIs.

## IX. CONCLUSION

This article emphasises that Indian higher education institutions (HEIs) can accelerate campus sustainability outcomes with AI only when AI is treated as an operational performance program rather than a technology add-on. Across energy, water, waste and mobility, the reviewed evidence shows that AI's value is typically realized through forecasting, anomaly detection, and optimization—but the magnitude and reliability of gains depend less on “model choice” and more on measurement discipline, data readiness, and workflow integration (Aghili et al., 2025; Zhang et al., 2022; Kanyama et al., 2024). In response, the KPI–Data–Duty (KDD) framework contributes a practical pilot-to-scale roadmap that starts with 3–5 auditable KPIs, builds a minimum viable data stack (“instrument first, optimize second”), and embeds a Duty layer to ensure Responsible and Green AI practices (privacy, accountability, security, transparency, and efficiency logging) aligned with lifecycle risk governance expectations (NIST, 2023; Schwartz et al., 2020). The framework also emphasizes disciplined evaluation through an 8–12 week pilot with seasonal baselines and scale decisions tied to demonstrated KPI improvement and stable governance controls.

For practice, the implications are immediate: HEIs should begin with building energy and water, where impacts are measurable and interventions are operational, while ensuring that procurement protects interoperability and avoids lock-in. For scholarship, the paper highlights the need for future studies that report standardized baselines, multi-campus validations, and explicit Green AI trade-offs to strengthen generalizability and replicability (Schwartz et al., 2020; Strubell et al., 2019). Overall, KDD reframes AI-enabled sustainability as a trustworthy, auditable, and resource-efficient operating model—one that can help Indian campuses convert pilots into measurable environmental performance gains while maintaining legitimacy and accountability.

## REFERENCES

- [1]. Aghili, S. A., Haji Mohammad Rezaei, A., Tafazzoli, M., Khanzadi, M., & Rahbar, M. (2025). Artificial intelligence approaches to energy management in HVAC systems: A systematic review. *Buildings*, 15(7), 1008. <https://doi.org/10.3390/buildings15071008>
- [2]. Aghili, S. A., Rezaei, A. H. M., Tafazzoli, M., Khanzadi, M., & Rahbar, M. (2025). Artificial intelligence approaches to energy management in HVAC systems: A systematic review. *Buildings*, 15(7), 1008. <https://doi.org/10.3390/buildings15071008>

[3]. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). *On the dangers of stochastic parrots: Can language models be too big?* In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). <https://doi.org/10.1145/3442188.3445922>

[4]. Bin, L., Shahzad, M., Javed, H., Muqeet, H. A., Akhter, M. N., Liaqat, R., & Hussain, M. M. (2022). Scheduling and sizing of campus microgrid considering demand response and economic analysis. *Sensors*, 22(16), 6150. <https://doi.org/10.3390/s22166150>

[5]. Bolón-Canedo, V., & Morán-Fernández, L. (2024). A review of green artificial intelligence: Towards a more sustainable future. *Neurocomputing*, 599, 128096. <https://doi.org/10.1016/j.neucom.2024.128096>

[6]. Botman, L., Lago, J., Fu, X., Chia, K., Wolf, J., Kleissl, J., & De Moor, B. (2024). Building plug load mode detection, forecasting and scheduling. *Applied Energy*, 369, 123098. <https://doi.org/10.1016/j.apenergy.2024.123098>

[7]. Bouquet, P., (et al.). (2024). AI-based forecasting for optimised solar energy management and enhanced smart grid reliability. *International Journal of Production Research*. <https://doi.org/10.1080/00207543.2023.2269565>

[8]. Chen, Y., & Lotti, L. (2025). Evaluating the effectiveness of behavioural nudges in reducing energy consumption in student accommodations: A quasi-experimental approach. *UCL Open Environment Preprints*. <https://doi.org/10.14324/ucloepreprints.286.v2>

[9]. Cheong, P. H., & Nyaupane, P. (2022). Smart campus communication, Internet of Things, and data governance: Understanding student tensions and imaginaries. *Big Data & Society*, 9(1). <https://doi.org/10.1177/20539517221092656>

[10]. de Bakker, C., Aries, M., Kort, H., & Rosemann, A. (2017). Occupancy-based lighting control in open-plan office spaces: A state-of-the-art review. *Building and Environment*, 112, 308–321. <https://doi.org/10.1016/j.buildenv.2016.11.042>

[11]. de Bakker, C., Van de Voort, T., & Rosemann, A. (2018). The energy saving potential of occupancy-based lighting control strategies in open-plan offices: The influence of occupancy patterns. *Energies*, 11(1), 2. <https://doi.org/10.3390/en11010002>

[12]. Domínguez-Bolaño, T., Barral, V., Escudero, C. J., & García-Naya, J. A. (2024). An IoT system for a smart campus: Challenges and solutions illustrated over several real-world use cases. *Internet of Things*, 25, 101099. <https://doi.org/10.1016/j.iot.2024.101099>

[13]. Es-sakali, N., Cherkaoui, M., Mghazli, M. O., & Naimi, Z. (2022). Review of predictive maintenance algorithms applied to HVAC systems. *Energy Reports*, 8(Suppl. 9), 1003–1012. <https://doi.org/10.1016/j.egyr.2022.07.130>

[14]. Ghezzi, R., & Mikkonen, T. (2024). On public procurement of ICT systems: Stakeholder views and emerging tensions. In *Software Business (ICSOB 2023)* (pp. 61–76). Springer. [https://doi.org/10.1007/978-3-031-53227-6\\_5](https://doi.org/10.1007/978-3-031-53227-6_5)

[15]. Gilman, E., Tamminen, S., Yasmin, R., Ristimella, E., Peltonen, E., Harju, M., Lovén, L., Riekki, J., & Pirttikangas, S. (2020). Internet of Things for smart spaces: A university campus case study. *Sensors*, 20(13), 3716. <https://doi.org/10.3390/s20133716>

[16]. Government of India. (2023). *The Digital Personal Data Protection Act, 2023 (No. 22 of 2023)*. Ministry of Electronics and Information Technology.

[17]. Grant, M. J., & Booth, A. (2009). A typology of reviews: An analysis of 14 review types and associated methodologies. *Health Information & Libraries Journal*, 26(2), 91–108. <https://doi.org/10.1111/j.1471-1842.2009.00848.x>

[18]. GRIHA Council. (n.d.). *GRIHA rating (GRIHA V3) – criteria overview*.

[19]. Hashim, M. A. M., Tlemsani, I., & Matthews, R. (2022). Higher education strategy in digital transformation. *Education and Information Technologies*, 27, 3171–3195. <https://doi.org/10.1007/s10639-021-10739-1>

[20]. Indian Green Building Council. (2024). *IGBC Green Campus rating system (Version 1.0): Reference guide* (July 2024). Indian Green Building Council.

[21]. International Organization for Standardization. (2018). *ISO 50001:2018—Energy management systems—Requirements with guidance for use*.

[22]. Kanyama, M. N., Bhunu Shava, F., Gamundani, A. M., & Hartmann, A. (2024). Machine learning applications for anomaly detection in Smart Water Metering Networks: A systematic review. *Physics and Chemistry of the Earth, Parts A/B/C*, 134, 103558. <https://doi.org/10.1016/j.pce.2024.103558>

[23]. Kanyama, M. N., Bhunu Shava, F., Gamundani, A. M., & Hartmann, A. (2025). AI-driven anomaly detection in smart water metering systems using ensemble learning. *Water*, 17(13), 1933. <https://doi.org/10.3390/w17131933>

[24]. Kanyama, M. N., Shava, F. B., Gamundani, A. M., & Hartmann, A. (2024). Machine learning applications for anomaly detection in smart water metering networks: A systematic review. *Physics and Chemistry of the Earth, Parts A/B/C*, 134, 103558. <https://doi.org/10.1016/j.pce.2024.103558>

[25]. Kim, I., & Lee, A. J. (2024). “I know what you did last semester”: Understanding privacy expectations and preferences in the smart campus. *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. <https://doi.org/10.1145/3613904.3642174>

[26]. Komljenovic, J., Birch, K., Sellar, S., Bergviken Rensfeldt, A., Deville, J., Eaton, C., Gourlay, L., Hansen, M., Kerssens, N., Kovalainen, A., Nappert, P.-L., Noteboom, J., Parcerisa, L., Pardo-Guerra, J. P., Poutanen, S., Robertson, S., Tyfield, D., & Williamson, B. (2025). Digitalised higher education: Key developments, questions, and concerns. *Discourse: Studies in the Cultural Politics of Education*, 46(2), 276–292. <https://doi.org/10.1080/01596306.2024.2408397>

[27]. Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. [https://doi.org/10.1518/hfes.46.1.50\\_30392](https://doi.org/10.1518/hfes.46.1.50_30392)

[28]. Maciel, R. R., de Souza, A. D., Almeida, R. M. A., & Leite, J. P. R. R. (2025). The impact of IoT-enabled routing optimization on waste collection distance: A systematic review and meta-analysis. *Logistics*, 9(4), 161. <https://doi.org/10.3390/logistics9040161>

[29]. MeitY. (2025). *Digital Personal Data Protection Rules*, 2025.

[30]. Melo, S., Gomes, R. J. R., Abbasi, R., & Arantes, A. (2024). Demand-responsive transport for urban mobility: Integrating mobile data analytics to enhance public transportation systems. *Sustainability*, 16(11), 4367. <https://doi.org/10.3390/su16114367>

[31]. Ministry of Electronics and Information Technology (MeitY). (2023). *The Digital Personal Data Protection Act*, 2023.

[32]. Murtaza, A. A., Saher, A., Zafar, M. H., Moosavi, S. K. R., Aftab, M. F., & Sanfilippo, F. (2024). Paradigm shift for predictive maintenance and condition monitoring from Industry 4.0 to Industry 5.0: A systematic review, challenges and case study. *Results in Engineering*, 24, 102935. <https://doi.org/10.1016/j.rineng.2024.102935>

[33]. Nahiduzzaman, M., Ahamed, M. F., & Naznine, M. (2025). An automated waste classification system using deep learning techniques: Toward efficient waste recycling and environmental sustainability. *Knowledge-Based Systems*, 310, 113028. <https://doi.org/10.1016/j.knosys.2025.113028>

[34]. Nahiduzzaman, M., Ahamed, M. F., Naznine, M., Karim, M. J., Kibria, H. B., Ayari, M. A., Khandakar, A., Ashraf, A., Ahsan, M., & Haider, J. (2025). An automated waste classification system using deep learning techniques: Toward efficient waste recycling and environmental sustainability. *Knowledge-Based Systems*, 310, 113028. <https://doi.org/10.1016/j.knosys.2025.113028>

[35]. National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). <https://doi.org/10.6028/NIST.AI.100-1>

[36]. NITI Aayog. (2021a). *Principles for Responsible AI*.

[37]. NITI Aayog. (2021b). *Operationalizing Principles for Responsible AI (Part 2)*.

[38]. Noor, R. M., Rasyidi, M. A., & others. (2020). Campus shuttle bus route optimization using machine learning predictive analysis: A case study. *Sustainability*, 13(1), 225. <https://doi.org/10.3390/su13010225>

[39]. Opara-Martins, J., Sahandi, R., & Tian, F. (2016). Critical analysis of vendor lock-in and its impact on cloud computing migration: A business perspective. *Journal of Cloud Computing: Advances, Systems and Applications*, 5(1). <https://doi.org/10.1186/s13677-016-0054-z>

[40]. Organisation for Economic Co-operation and Development. (2025). *Digital transformation of public procurement*. OECD

[41]. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>

[42]. Paunov, Y., & Grüne-Yanoff, T. (2023). Boosting vs. nudging sustainable energy consumption: A long-term comparative field test in a residential context. *Behavioural Public Policy*. <https://doi.org/10.1017/bpp.2023.30>

[43]. Press Information Bureau. (2025, November 17). *DPDP Rules, 2025 notified: A citizen-centric framework for privacy protection and responsible data use*.

[44]. Rajbhandari, Y., (et al.). (2024). Enhanced demand side management for solar-based isolated microgrids using forecasting and optimization. *IET Smart Grid*. <https://doi.org/10.1049/stg2.12151>

[45]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. <https://doi.org/10.1145/2939672.2939778>

[46]. 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>

[47]. Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>

[48]. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*. <https://doi.org/10.18653/v1/P19-1355>

[49]. United Nations. (2015). *Transforming our world: The 2030 Agenda for Sustainable Development* (A/RES/70/1

[50]. Verdecchia, R., Sallou, J., & others. (2023). A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*. <https://doi.org/10.1002/widm.1507>

[51]. Wang, Z., Hong, T., & Piette, M. A. (2019). Predicting plug loads with occupant count data through a deep learning approach. *Energy*, 181, 29–42. <https://doi.org/10.1016/j.energy.2019.05.138>

[52]. Zhang, W., Wu, Y., & Calautit, J. K. (2022). A review on occupancy prediction through machine learning for enhancing energy efficiency, air quality and thermal comfort in the built environment. *Renewable and Sustainable Energy Reviews*, 167, 112704. <https://doi.org/10.1016/j.rser.2022.112704>

## APPENDIX 1

<b>Table 1: AI use-cases mapped to SDGs, campus KPIs, and key risks</b>					
<b>Campus domain</b>	<b>AI use-case (example)</b>	<b>Primary KPI</b>	<b>SDG linkage</b>	<b>Key data needed</b>	<b>Main risks &amp; controls</b>
<b>Energy (buildings)</b>	AI HVAC optimization (occupancy + weather adaptive control)	kWh, peak demand, comfort hours	SDG 7, 11, 13	Metering, occupancy proxy, weather	Over-automation → human override; model drift
<b>Energy (lighting/plug loads)</b>	Predictive scheduling + occupancy sensing	kWh per m <sup>2</sup>	SDG 7, 13	PIR/Wi-Fi proxy, schedules	Sensor privacy → minimize & aggregate
<b>Water</b>	Leak/anomaly detection on smart meters	L/capita, leak rate	SDG 6, 11	Water meters, zone-level logs	False alarms → thresholds + verification
<b>Waste</b>	Vision-assisted segregation + bin fill prediction	diversion %, contamination rate	SDG 11, 12	Camera/bin sensors	Surveillance concerns → avoid faces; signage
<b>Mobility</b>	Shuttle routing + ride-share matching	low-carbon commute %, utilization	SDG 11, 13	travel demand, routes	Equity → ensure access across groups
<b>Admin/paper</b>	Document classification + workflow automation	print volume, process time	SDG 12, 13	forms/docs metadata	Bias/errors → human review, audit trails
<b>Operations</b>	Predictive maintenance for pumps/chillers	downtime, wastage, energy intensity	SDG 9, 12, 13	sensors, maintenance logs	Vendor lock-in → open standards, data ownership

## APPENDIX 2

<b>Table 2: KPI–Data–Duty (KDD) Framework for Implementing AI-Enabled Campus Sustainability in Indian HEIs (Pilot-to-Scale Roadmap)</b>			
<b>Step</b>	<b>What it means</b>	<b>What to do (checklist)</b>	<b>Outputs you should document</b>
<b>Step 1: KPI (Impact clarity)</b>	Pick 3–5 auditable, time-bound sustainability metrics for the campus	<ul style="list-style-type: none"> <li>• <b>Electricity:</b> kWh per student / per m<sup>2</sup>; peak demand</li> <li>• <b>Water:</b> liters per capita; leakage rate</li> <li>• <b>Waste:</b> segregation rate; landfill diversion</li> <li>• <b>Mobility:</b> % low-carbon commuting; shuttle utilization</li> <li>• <b>Paper:</b> forms digitized; print volume Align KPIs with structured categories (transportation, water, energy, waste, health &amp; well-being)</li> </ul>	<ul style="list-style-type: none"> <li>• KPI definitions + formula</li> <li>• Baseline period selected</li> <li>• Target value + timeline</li> <li>• KPI owner (role/name)</li> </ul>
<b>Step 2: Data (Feasibility &amp; architecture)</b>	Ensure you have the minimum viable data to measure and improve the KPI	<p><b>Minimal viable data stack:</b></p> <ul style="list-style-type: none"> <li>• Sub-metering (hostels, academic blocks, chilled water plant)</li> <li>• Occupancy proxies (timetables, Wi-Fi counts, PIR sensors)</li> <li>• Weather data + tariff schedule</li> <li>• Ticketing/maintenance logs</li> </ul>	<ul style="list-style-type: none"> <li>• Data inventory (sources, frequency, quality)</li> <li>• Data access + storage plan                             <ul style="list-style-type: none"> <li>• Interoperability notes (formats/APIs)</li> </ul> </li> <li>• Data governance notes (who can access what)</li> </ul>
<b>Step 3: Duty (Responsible &amp; Green AI checks)</b>	Run responsible AI + Green AI checks before deploying	<ul style="list-style-type: none"> <li>• <b>Privacy:</b> data minimization, access controls, retention limits</li> <li>• <b>Transparency:</b> explainable rules/alerts for operators</li> <li>• <b>Accountability:</b> named owner, escalation paths, human override</li> <li>• <b>Security:</b> secure device/network layer</li> <li>• <b>Green AI:</b> track compute, choose efficient models, measure energy cost of inference</li> <li>• <b>Risk alignment:</b> use recognized AI risk framework vocabulary + documentation approach</li> </ul>	<ul style="list-style-type: none"> <li>• Responsible AI checklist completed</li> <li>• Risk register (top risks + mitigations)</li> <li>• Model/decision documentation (what, why, limits)                             <ul style="list-style-type: none"> <li>• Green AI logging plan (compute/inference energy proxy)</li> </ul> </li> </ul>
<b>Step 4: Pilot → Scale</b>	Start small, prove impact, then scale responsibly	<ul style="list-style-type: none"> <li>• Pilot <b>8–12 weeks</b> in one building/hostel</li> <li>• Compare baseline vs pilot (same season if possible)                             <ul style="list-style-type: none"> <li>• Scale only after KPI improvement is demonstrated and governance controls are stable</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Pilot report (baseline vs pilot KPI change)</li> <li>• Lessons learned + adjustments                             <ul style="list-style-type: none"> <li>• Scale plan (sites, timeline, budget, governance)</li> </ul> </li> <li>• Ongoing monitoring plan (drift, exceptions, audits)</li> </ul>