International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(5s) 2026 pp. 118-132
https://doi.org/10.70917/ijcisim-2026-2680
Article

A Predictive Analytics Framework for Data-Driven Sustainability in Reducing Energy Consumption and Carbon Footprint Across Urban Infrastructure

Evha Rozario[1], Shuchita Shahnaz[2], Foysal Mahmud[3], Shuchona Malek Orthi[4], Gazi Touhidul Alam[5], Ashraful Islam[6], Subha Shamarukh[7*]

[1]International American University, Los Angeles, California, USA, Email: evhaaccabd@gmail.com, ORCID: https://orcid.org/0009-0006-3077-2011
[2]Bangladesh University of Professionals, Mirpur, Dhaka, Bangladesh, Email: shuchi.mingu@gmail.com, ORCID: https://orcid.org/0009-0007-0725-394X
[3]Westcliff University, Irvine, California, USA, Email: f.mahmud.130@westcliff.edu, ORCID: https://orcid.org/0009-0002-1059-0166
[4]Westcliff University, Irvine, California, USA, Email: s.orthi.339@westcliff.edu, ORCID: https://orcid.org/0009-0007-5397-4561
[5]Trine University, Detroit, Michigan, USA, Email: touhid.one@gmail.com, ORCID: https://orcid.org/0009-0007-0400-3654
[6]Washington University of Science and Technology, Alexandria, VA, USA, Email: ashralam.student@wust.edu, ORCID: https://orcid.org/0009-0001-8067-7331
[7]University of Rochester, Rochester, New York, USA, Email: shamarukhsubha@gmail.com, ORCID: https://orcid.org/0009-0000-2170-1541

Abstract: Rapid urbanization has intensified the energy demand and carbon intensity of the built environment, positioning cities at the center of global decarbonization efforts. Although predictive analytics, machine learning, and ubiquitous sensing now generate unprecedented volumes of urban energy data, their translation into measurable reductions in consumption and emissions remain fragmented, opaque, and weakly connected to decision-making. This paper develops an integrative predictive analytics framework that operationalizes data-driven sustainability across urban infrastructure. Grounded in a structured synthesis of recent scholarship on urban building energy modelling, smart-grid analytics, digital twins, and explainable artificial intelligence, the framework is organized as six interdependent layers: data acquisition and sensing, data integration and governance, predictive modelling, explainability and trust, optimization and decision support, and a continuous feedback loop linking analytical outputs to policy and operational action. The framework treats interpretability not as an optional refinement but as a precondition for institutional uptake, embedded carbon accounting within the analytical pipeline rather than appending it downstream and explicitly aligns analytical objectives with the Sustainable Development Goals, particularly those concerning affordable clean energy, sustainable cities, resilient infrastructure, and climate action. A proposed validation design is specified, comprising candidate data sources, a comparative modelling pipeline, an evaluation protocol, and a staged deployment with outcome feedback, so that the conceptual contribution can be empirically tested without recourse to fabricated results. The paper contributes a coherent theoretical scaffolding that connects technical prediction to governance, clarifies persistent barriers related to data quality, interoperability, privacy, and trust, and articulates a research agenda for evidence-based, transparent, and equitable urban decarbonization.

Keywords: Data-Driven Sustainability; Predictive Analytics; Urban Infrastructure; Energy Consumption; Carbon Footprint; Machine Learning; Explainable Artificial Intelligence; Smart Cities

1. Introduction

Cities have become the decisive arena for confronting climate change. As the share of the world population living in urban areas continues to rise, the infrastructure that sustains urban life, including buildings, transport networks, water systems, and electricity grids, accounts for a disproportionate and growing fraction of global energy use and greenhouse-gas emissions. The building sector alone is consistently estimated to be responsible for close to forty percent of energy-related carbon dioxide emissions, a figure that reflects both the sheer scale of the urban stock and the inefficiency embedded in its daily operation (Bibri & Krogstie, 2024; Qiao et al., 2024). Because the trajectory of global emissions will be determined in large part by how urban infrastructure is built, operated, and retrofitted over the coming decades, reducing the energy consumption and associated carbon footprint of that infrastructure is not merely a technical aspiration but a structural requirement for meeting national and international decarbonization commitments. The challenge is rendered more acute by the long service lives of urban assets, which lock in patterns of consumption for decades, and by the spatial concentration of demand, which magnifies the consequences of inefficient operation.

Two parallel developments have transformed the conditions under which this challenge can be addressed. The first is the proliferation of sensing and metering technologies, which has turned the contemporary city into a continuous generator of high-resolution data describing how energy is produced, distributed, and consumed. Smart meters, building-management systems, Internet-of-Things devices, and remote-sensing platforms now capture temporal and spatial patterns of demand at a granularity that was unattainable only a decade ago, and the bidirectional communication characteristic of modern smart grids allows these measurements to be coupled with control (Mikati et al., 2025). The second development is the maturation of predictive analytics and machine learning, which provides the methodological apparatus to convert these data into forecasts, diagnoses, and prescriptions. Methods ranging from regularized regression and tree-based ensembles to recurrent neural networks have demonstrated strong performance in anticipating building and city-scale energy demand, and they increasingly outperform the physics-based simulation approaches that long dominated the field when monitored data are abundant (Breiman, 2001; Chen & Guestrin, 2016; Hochreiter & Schmidhuber, 1997; Qiao et al., 2024).

Despite this convergence of data and method, a persistent gap separates analytical capability from sustainability outcomes. Much of the existing scholarship concentrates on improving predictive accuracy for narrowly defined tasks, such as forecasting the energy use intensity of a particular building typology or the short-term load of a single feeder, while comparatively little attention is paid to how these predictions are integrated into the governance processes through which energy and emissions are actually managed. Models are frequently developed in isolation from the data-quality, interoperability, and institutional constraints that determine whether their outputs can be trusted and acted upon in practice. The opacity of high-performing models compounds the difficulty, because the decision-makers in public administration and infrastructure operation who must authorize interventions are understandably reluctant to act on recommendations they cannot interpret, contest, or defend before elected officials and citizens (Bibri & Krogstie, 2024; Doshi-Velez & Kim, 2017; Lundberg & Lee, 2017). The result is a proliferation of accurate but unused models, and a corresponding shortfall in realized decarbonization.

This paper responds to that gap by proposing an integrative predictive analytics framework for data-driven sustainability across urban infrastructure. Rather than introducing a single algorithm or improving a single benchmark, the contribution is architectural: it specifies how data acquisition, integration, modelling, explanation, optimization, and feedback can be assembled into a coherent system whose explicit and overriding purpose is the reduction of energy consumption and carbon footprint. The framework is conceptual and methodological in nature. It is grounded in a structured synthesis of recent literature and is accompanied by a detailed validation design, so that its components can be empirically tested on real data rather than asserted through unsubstantiated results. This stance is deliberate; in a field where the credibility of analytics depends on the integrity of evidence, the framework is offered as a testable proposition rather than a claim of demonstrated performance.

The study is guided by three research questions. First, how can heterogeneous urban energy data be organized into a pipeline that supports accurate and interpretable prediction of consumption and emissions across diverse asset types and temporal scales? Second, how can the analytical outputs of such a pipeline be linked to optimization and decision support in a manner that is transparent, auditable, and accountable to the institutions responsible for urban infrastructure? Third, how can the resulting framework be aligned with established sustainability objectives so that statistical performance translates into measurable progress toward decarbonization?

Further consideration motivates the analytical orientation of this paper. Historically, the management of urban energy has been largely reactive, responding to demand as it materializes and to inefficiency as it is discovered. Predictive analytics enables a shift from this reactive posture to an anticipatory one, in which demand is forecast before it occurs, inefficiency is diagnosed before it accumulates, and interventions are evaluated before they are committed. This shift is not merely a matter of efficiency but of effectiveness, because many of the most consequential decisions affecting urban emissions, such as retrofit programs and capacity planning, are made years in advance and cannot be corrected quickly once implemented. An analytics framework that supports anticipatory decision-making therefore addresses the temporal structure of the problem more faithfully than reactive monitoring alone, and it is this anticipatory capability, harnessed in service of decarbonization, that the present framework is designed to provide.

The paper makes three contributions corresponding to these questions. It articulates a layered architecture that connects sensing to governance through interpretation and feedback; it positions explainability and carbon accounting as constitutive elements of the pipeline rather than optional additions; and it specifies a rigorous validation design that enables empirical testing without fabricated evidence. The remainder of the paper is structured as follows. Section 2 reviews the relevant literature on urban energy prediction, carbon assessment, smart-city data infrastructures, and explainable analytics, and identifies the barriers and the gap the framework addresses. Section 3 presents the conceptual framework and its layered architecture, supported by figures. Section 4 specifies the proposed methodology and validation design. Section 5 discusses theoretical, practical, and policy implications. Section 6 considers limitations and future research, and Section 7 concludes.

2. Literature Review

Literature relevant to data-driven urban decarbonization spans several communities that have, until recently, developed largely in parallel. This section synthesizes four strands, namely predictive analytics for energy, carbon footprint assessment, smart-city data infrastructures, and explainable analytics, and then examines the cross-cutting barriers that impede their integration. The synthesis establishes both the maturity of the individual components and the fragmentation that the proposed framework is designed to resolve.

2.1 Predictive analytics for urban and building energy

The data-driven prediction of energy consumption has become one of the most active research frontiers in sustainable urban development. Tree-based ensemble methods, in particular the random forest and gradient-boosted variants, have repeatedly been shown to deliver robust accuracy across diverse building types because they accommodate non-linear interactions among heterogeneous predictors without demanding restrictive distributional assumptions, and because they are comparatively resilient to the noise and missingness that characterize real measurement data (Breiman, 2001; Chen & Guestrin, 2016). Gradient boosting in particular has emerged as a frequent best performer in comparative studies, owing to its capacity to model subtle feature interactions while controlling overfitting through regularization.

At the scale of urban neighbourhoods, hybrid strategies that couple physics-based urban building energy modelling with machine learning have proven effective for overcoming the chronic scarcity of monitored data. In such designs, a simulation engine is used to generate a training corpus that captures the physical relationships between building form and energy demand, and a learned model is then deployed for rapid large-scale prediction across the remaining stock. This approach combines physical fidelity of simulation with the computational efficiency of machine learning, and gradient boosting frequently emerges as the best-performing estimator while requiring only a compact set of the most informative features (Qiao et al., 2024). The practical significance of this hybridization is considerable, because it extends data-driven prediction to cities and neighborhoods that lack the dense monitoring infrastructure on which purely empirical methods depend.

Temporal dependence in energy demand has motivated the adoption of recurrent architectures, most notably the long short-term memory network, which captures the sequential structure of load profiles and the lagged influence of weather and occupancy on consumption (Hochreiter & Schmidhuber, 1997). Comparative studies of educational and public buildings have evaluated long short-term memory networks alongside random forest and gradient-boosting regressors, reporting that no single algorithm dominates across all contexts and that model selection must be tailored to the temporal and structural characteristics of each building stock (Elhabyb et al., 2024). This finding is consequential for framework design, because it implies that a credible analytics system cannot privilege one estimator a priori but must support disciplined model comparison and selection as a first-class function. A framework that hard-codes a single algorithm is liable to underperform whenever the characteristics of the asset portfolio depart from those for which that algorithm was originally validated.

2.2 Carbon footprint estimation and assessment

A recurring theme across this body of work concerns the data and feature requirements on which predictive accuracy depends. The performance of even the most capable estimator is bounded by the quality and informativeness of its inputs, and studies consistently report that careful feature engineering, drawing on building attributes, occupancy, and meteorological conditions, contributes as much to predictive success as the choice of algorithm (Elhabyb et al., 2024; Qiao et al., 2024). The identification of the most influential features also serves a dual purpose, since reducing a model to its most informative inputs lowers complexity and improves both interpretability and computational tractability without materially degrading accuracy (Qiao et al., 2024). This interdependence between feature selection, parsimony, and interpretability foreshadows the explainability concerns discussed below and reinforces the case for treating data preparation as a substantive component of the analytics system rather than a preliminary chore.

Predicting energy consumption is a necessary but insufficient step toward decarbonization; the analytical output must be translated into carbon terms through emission factors that reflect the composition of the supplying grid and the operational characteristics of the asset. Literature increasingly couples consumption prediction with carbon assessment, recognizing that the same kilowatt-hour of electricity carries markedly different emissions depending on regional generation mixes and the time of day at which it is drawn (Bibri & Krogstie, 2024). The temporal dimension is especially important as grids decarbonize unevenly, since shifting demand toward periods of low-carbon supply can reduce emissions even when total consumption is unchanged.

Embedding carbon accounting within the analytical pipeline, rather than treating it as a downstream calculation appended after prediction, allows interventions to be prioritized by their marginal abatement potential and cost. This integration is decisive when retrofit budgets are constrained and decision-makers must justify the allocation of scarce resources, because it permits a transparent ranking of candidate measures according to the emissions they avert per unit of expenditure. The literature on building decarbonization consistently emphasizes that envelope improvements and the upgrading of heating, ventilation, and cooling systems carry substantial abatement potential, and an analytics framework that quantifies this potential at the level of the individual asset provides the evidence base for such prioritization (Bibri & Krogstie, 2024). Carbon-aware analytics therefore reframes the objective from minimizing energy use in the abstract to minimizing emissions subject to cost and feasibility constraints, which more faithfully represents the decisions that infrastructure managers actually face.

2.3 Smart-city data infrastructures: sensing, smart grids, and digital twins

The analytical methods described above are only as effective as the data infrastructures that feed them. Internet-of-Things sensing and smart-grid telemetry provide the continuous, bidirectional flow of measurements that make real-time monitoring and demand response feasible, while simultaneously introducing acute challenges of scalability, interoperability, and security (Mikati et al., 2025). The heterogeneity of these data streams, which originate from devices and systems that were never designed to interoperate, is a recurring obstacle, and the volume and velocity of the resulting data place substantial demands on storage, transmission, and processing. Recent work integrating distributed-ledger mechanisms with Internet-of-Things energy management illustrates both the promise and the cost of these architectures, demonstrating that demand-response optimization and decentralized coordination can reduce energy costs and enable peer-to-peer energy trading, while exposing the system to new computational burdens and governance questions concerning the energy intensity of the coordinating mechanism itself (Mikati et al., 2025).

Digital twins have emerged as an organizing concept for binding these data streams into a coherent, queryable representation of urban infrastructure. By maintaining a continuously updated virtual counterpart of the physical city, a digital twin enables planners to simulate the consequences of interventions, such as envelope retrofits, electrification, or grid reconfiguration, before committing to them, and to monitor environmental performance against sustainability targets in something approaching real time (Bibri & Krogstie, 2024). The convergence of artificial intelligence, the Internet of Things, and digital-twin technology is widely regarded as the computational substrate of the sustainable smart city, providing the means to move from descriptive monitoring to predictive and prescriptive management. Yet the literature notes that these technologies have often been studied in isolation, with research treating urban artificial intelligence, the artificial intelligence of things, and urban digital twins as separate domains, leaving their synergistic integration conspicuously underdeveloped (Bibri & Krogstie, 2024). This observation directly motivates an architectural response that specifies how the components fit together.

2.4 Explainability and trust in sustainability analytics

A further characteristic of modern smart-city data infrastructures bears on framework design, namely the increasing importance of real-time and near-real-time analytics. The value of much urban energy data decays rapidly, since a forecast of imminent demand or a detection of anomalous consumption is useful only if it is produced in time to inform action. This temporal constraint has motivated architectures to distribute computation toward the network edge, processing data closer to its source to reduce latency and to limit the transmission of sensitive raw measurements (Mikati et al., 2025). The implication for the proposed framework is that the modelling and decision-support layers must be capable of operating within the time horizons that operational decisions impose, and that the data and governance layers must be engineered for throughput as well as accuracy. Real-time operation also intensifies the demand for interpretability, because decisions taken quickly and at scale leave little opportunity for the manual scrutiny that might otherwise compensate for opacity.

As predictive models grow more capable, their internal complexity increasingly obstructs the transparency that public decision-making requires. Explainable artificial intelligence has therefore moved from a peripheral concern to a central requirement of credible sustainability analytics, and the field has developed a rigorous vocabulary for distinguishing the goals, methods, and evaluation of interpretability (Doshi-Velez & Kim, 2017). Two broad families of post-hoc explanation are now in widespread use. Local surrogate methods approximate the behavior of a complex model in the neighborhood of an individual prediction, allowing analysts to understand why a specific building or district received a particular forecast (Ribeiro et al., 2016). Additive feature-attribution methods, of which the Shapley-value approach is the most theoretically grounded, decompose individual predictions into the contributions of constituent features in a manner that satisfies desirable consistency properties, and they support both local explanation for individual assets and global explanation across an entire portfolio (Lundberg & Lee, 2017).

Applied work in energy forecasting confirms that these interpretability techniques can be combined with high-performing models to deliver both accuracy and transparency, thereby supporting auditability and regulatory compliance in settings where unexplained recommendations would be unacceptable (El-kenawy et al., 2025). The recurring conclusion across this literature is that the absence of explanation, rather than any deficit of accuracy, is frequently the binding constraint on the adoption of analytics in policy and operational settings. Where officials must account publicly for the basis of their decisions, an opaque model that cannot be interrogated is effectively unusable regardless of its statistical performance. This conclusion has a direct architectural implication, namely that interpretation must be designed into the analytics pipeline as a dedicated function rather than retrofitted after deployment, and it is on this basis that the proposed framework elevates explainability to the status of a distinct layer.

2.5 Cross-cutting barriers to integration

Beyond the individual strands, the literature identifies a set of cross-cutting barriers that impede the assembly of these components into operational systems. Data quality and availability constitute the most pervasive obstacle, since the validity of every downstream inference depends on the coverage, resolution, and reliability of measurement, and many cities, particularly smaller or resource-constrained ones, lack the dense monitoring infrastructure that data-driven methods presuppose (Qiao et al., 2024). Interoperability across heterogeneous systems is a second barrier, as is the protection of privacy and security in the handling of fine-grained energy data, which can reveal sensitive patterns of occupancy and behavior if inadequately governed (Mikati et al., 2025). A fourth barrier is institutional, encompassing the absence of governance frameworks that connect analytical outputs to decision authority, and the deficit of trust that opaque models engender. The framework developed in the next section is structured precisely so that each of these barriers is addressed by a dedicated element of the architecture rather than left to ad hoc resolution.

2.6 Synthesis and research gap

Taken together, the literature establishes that the individual components required for data-driven urban decarbonization are mature. Accurate predictive estimators exist and have been validated across building types and temporal scales; carbon-accounting methods are available and can be coupled with prediction; sensing, smart-grid, and digital-twin infrastructures are operational in leading cities; and explanation techniques can render even complex models intelligible to non-specialist decision-makers. What remains underdeveloped is the integration of these components into a single framework whose explicit objective is the reduction of energy consumption and carbon footprint, and whose architecture connects prediction to governance through interpretation, optimization, and feedback. The prevailing fragmentation, in which models are optimized for accuracy in isolation from the institutional and data-governance conditions of their use, and in which interpretability and carbon accounting are treated as afterthoughts, constitutes the gap this paper addresses. The framework developed in the following section is designed to close it by specifying not merely which techniques to use but how they cohere into a system oriented toward measurable decarbonization.

3. The Proposed Predictive Analytics Framework

This section presents the framework. It begins with the design principles that govern its construction, proceeds to the layered architecture that constitutes its core, addresses the data foundation on which the architecture rests, and concludes with the explicit alignment of analytical objectives to sustainability goals. Throughout, the emphasis is on how the elements cohere into a system rather than on the elements in isolation.

3.1 Design principles

The framework is built on four principles derived from the preceding synthesis. The first is integration: technical prediction is treated as one element of a larger socio-technical system rather than an end in itself, so that the value of a model is judged by its contribution to decarbonization outcomes rather than by its accuracy in isolation. The second is interpretability by design, meaning that explanation is embedded in the architecture as a dedicated function rather than appended after deployment, because the evidence indicates that institutional uptake depends on it (Doshi-Velez & Kim, 2017; Lundberg & Lee, 2017). The third is actionability: every analytical output is oriented toward a decision, whether an operational adjustment, a retrofit prioritization, or a policy intervention, and outputs that cannot be connected to action are regarded as incomplete. The fourth is alignment, whereby analytical objectives are explicitly mapped to recognized sustainability goals so that performance is measured in terms of decarbonization and societal benefit rather than statistical accuracy alone (United Nations, 2015). These principles are not merely rhetorical; each is instantiated in a specific feature of the architecture described below.

3.2 Layered architecture

The framework comprises six interdependent layers, depicted in Figure 1. Each layer transforms the outputs of the previous one and supplies the inputs of the next, while a governance and feedback mechanism closes the loop between analytical insight and infrastructural action. The vertical arrangement reflects the logical flow from raw measurement to realized outcome, and the feedback path signifies that the framework is a continuous cycle rather than a linear sequence terminating in a single prediction.

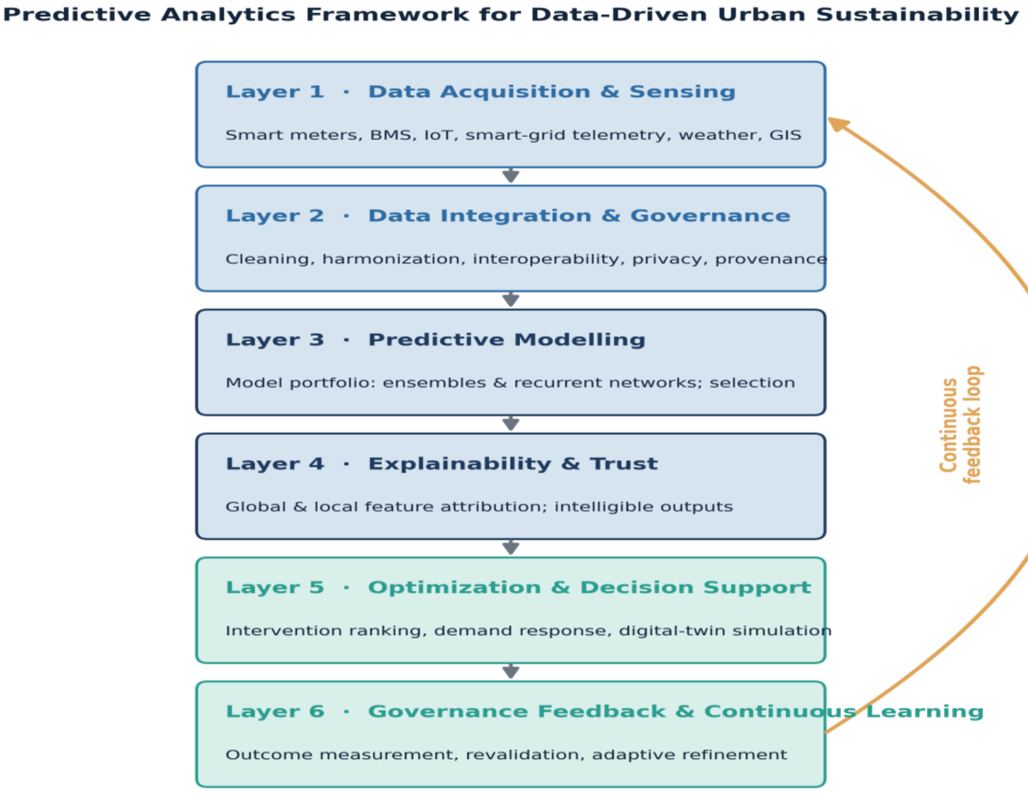

Figure 1: The six-layer architecture of the proposed predictive analytics framework, showing the flow from data acquisition to governance feedback and the continuous learning loop.

Layer 1: Data acquisition and sensing. The foundational layer ingests heterogeneous data from smart meters, building-management systems, Internet-of-Things sensors, smart-grid telemetry, meteorological services, and geospatial sources. Its principal function is to assure coverage, resolution, and reliability, since the validity of every downstream inference depends on the quality of measurement at this stage. The layer must contend with the practical realities of intermittent connectivity, sensor drift, and gaps in coverage, and it should incorporate mechanisms for monitoring data freshness and completeness so that deficiencies are detected before they propagate (Mikati et al., 2025).

Layer 2: Data integration and governance. Raw streams are cleaned, harmonized, and reconciled against a common spatial and temporal reference, with explicit attention to interoperability across systems that were never designed to communicate. This layer also enforces data governance, encompassing privacy protection, access control, provenance tracking, and systematic quality assurance, which are prerequisites for trustworthy analytics in public infrastructure. Because fine-grained energy data can expose sensitive behavioral patterns, the governance functions situated here are not administrative overhead but a substantive safeguard that conditions the social license under which the entire system operates (Mikati et al., 2025).

Layer 3: Predictive modelling. The integrated data supports the training, comparison, and selection of predictive models for energy consumption and, through emission factors, carbon footprint. Consistent with the evidence that no single estimator dominates, the layer is designed to accommodate a portfolio of methods, including regularized linear baselines, tree-based ensembles for cross-sectional prediction, and recurrent networks for temporal forecasting, with selection governed by validated performance on held-out data rather than by convention or convenience (Breiman, 2001; Chen & Guestrin, 2016; Elhabyb et al., 2024; Hochreiter & Schmidhuber, 1997). Treating model selection as a first-class function protects the framework against the brittleness that follows from committing prematurely to a single algorithm.

Layer 4: Explainability and trust. Predictions are passed through feature-attribution analysis that quantify the contribution of each driver, both globally across the asset portfolio and locally for individual buildings or districts. Global attribution reveals the systemic determinants of consumption and emissions, informing policy, while local attribution explains particular forecasts, informing operational decisions and enabling stakeholders to scrutinize and contest recommendations. This layer converts opaque outputs into intelligible explanations and is treated as indispensable rather than optional, in direct response to evidence that the intelligibility of outputs, not their accuracy, is frequently the binding constraint on adoption (Doshi-Velez & Kim, 2017; El-kenawy et al., 2025; Lundberg & Lee, 2017; Ribeiro et al., 2016).

Layer 5: Optimization and decision support. Interpreted predictions feed optimization routines that rank interventions by their abatement potential, cost, and feasibility, and that support operational decisions such as demand response and load scheduling. A digital-twin representation can serve as the environment in which candidate interventions are simulated before implementation, allowing the consequences of a decision to be examined in advance and reducing the risk of costly missteps (Bibri & Krogstie, 2024; Mikati et al., 2025). The output of this layer is not a prediction but a prioritized and justified set of actions.

Layer 6: Governance feedback and continuous learning. Implemented interventions generate new measurements that re-enter the pipeline, allowing models to be revalidated against realized outcomes and recommendations to be refined in light of what actually occurred. This closed loop distinguishes the framework from static predictive studies, embedding analytics within an adaptive cycle of monitoring, action, and learning that is essential for sustained decarbonization and that guards against the gradual degradation of model performance as conditions change.

Table 1

Layers of the proposed predictive analytics framework and their core functions

Layer	Core function	Principal output
1. Data acquisition and sensing	Capture heterogeneous urban energy data at adequate coverage and resolution	Raw measurement streams
2. Data integration and governance	Clean, harmonize, and govern data for interoperability, privacy, and quality	Curated, analysis-ready dataset

3. Predictive modelling	Train, compare, and select consumption and carbon models	Validated predictions
4. Explainability and trust	Attribute predictions to drivers globally and locally	Interpretable explanations
5. Optimization and decision support	Rank and simulate interventions; support operations	Prioritized action set
6. Governance feedback	Revalidate against outcomes; enable continuous learning	Refined models and policy signals

3.3 The data foundation and its convergence on the pipeline

The credibility of the framework rests on the data foundation established by the first two layers, and the manner in which heterogeneous sources converge on the analytics pipeline merits explicit treatment. Figure 2 depicts this convergence, in which distinct measurement streams, each with its own format, resolution, and reliability characteristics, are reconciled within the integration and governance layer before being passed to modelling, explanation, and decision support. The diagram makes visible a point that is easily overlooked in studies focused on a single data source, namely that the integration layer performs the indispensable work of transforming a collection of incompatible feeds into a coherent, analysis-ready representation. Without this reconciliation, the downstream models would be trained on inconsistent or misaligned inputs, and their outputs would inherit and amplify the defects of the raw data.

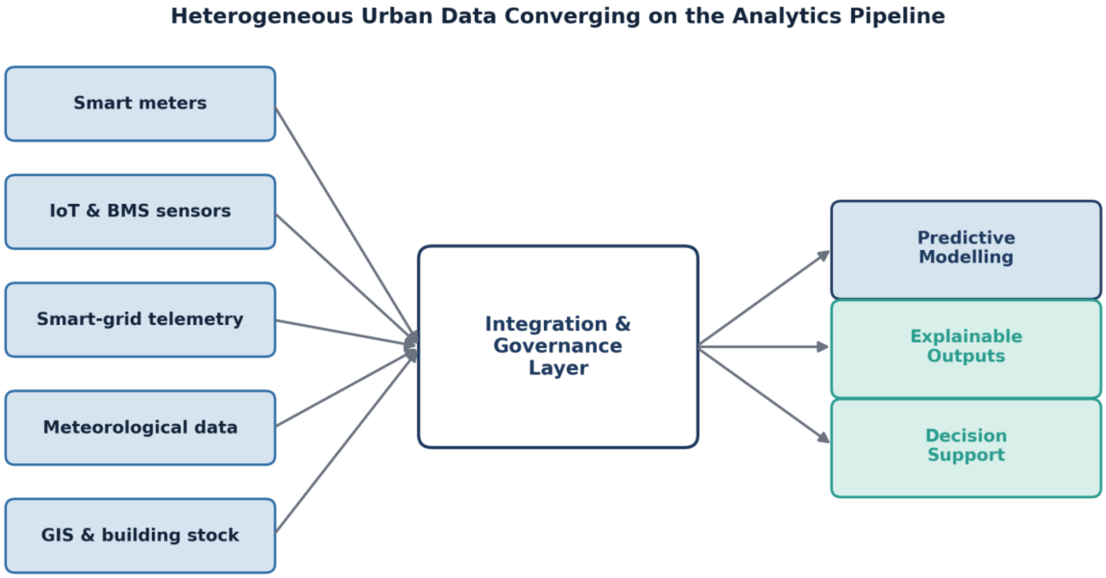

Figure 2: Convergence of heterogeneous urban data sources on the integration and governance layer, which supplies the predictive, explanatory, and decision-support functions of the pipeline.

The variety of sources also underscores the importance of the governance functions situated in the second layer. Smart-meter and sensor data are continuous and high-frequency, meteorological data are exogenous and forecastable, geospatial and building-stock data are largely static, and grid emission factors are time-varying and externally determined. Harmonizing these along common spatial and temporal references, while maintaining provenance and protecting privacy, is a non-trivial undertaking whose success or failure determines the validity of everything that follows. The framework therefore treats data integration not as preprocessing but as a substantive layer with its own governance obligations.

3.4 Alignment with the Sustainable Development Goals

The framework is deliberately aligned with the 2030 Agenda for Sustainable Development, and this alignment is summarized in Figure 3 (United Nations, 2015). Its core objective of reducing energy consumption and carbon footprint advances the goal of affordable and clean energy by improving efficiency and facilitating the integration of renewable supply, and the goal of climate action by lowering emissions across the urban stock. Because the framework operates at the scale of urban infrastructure and is intended to inform its planning and operation, it contributes directly to the goal of sustainable cities and communities, while its reliance on advanced analytics and digital infrastructure connects it to the goal of industry, innovation, and infrastructure.

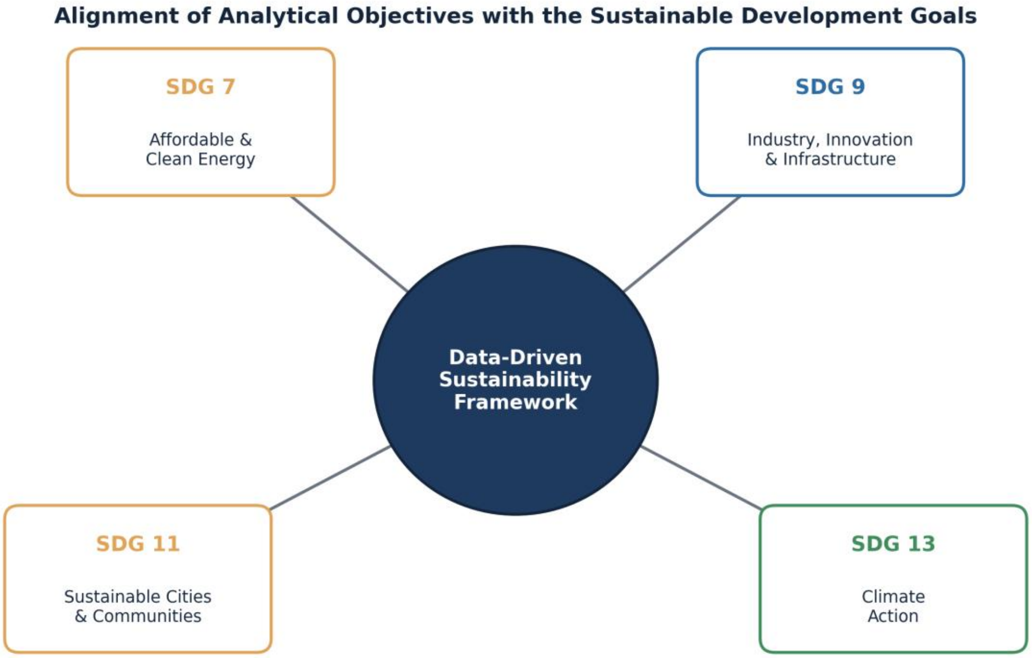

Figure 3: Alignment of the framework's analytical objectives with the Sustainable Development Goals most directly served, namely affordable and clean energy, industry and infrastructure, sustainable cities, and climate action.

This explicit mapping serves more than a rhetorical purpose. By framing analytical objectives in terms of recognized societal goals, it provides a vocabulary through which technical results can be communicated to policymakers and the public, and it supplies criteria against which the framework's success should ultimately be judged. The emphasis on transparent, accountable decision-making, embodied in the explainability layer, further supports the institutional dimensions of sustainable governance, ensuring that the pursuit of decarbonization does not come at the expense of the legitimacy on which public action depends.

3.5 Cross-cutting concerns: privacy, security, and computational cost

Three concerns cut across the layered architecture and merit explicit treatment, because they cannot be confined to any single layer. The first is privacy. Fine-grained energy data are capable of revealing intimate details of occupancy and behavior, and the framework therefore embeds privacy protection within the data-governance layer while requiring that explanation and decision support operate on appropriately aggregated or safeguarded representations. The second is security. The bidirectional connectivity that enables monitoring and control also expands the attack surface of the system, and the integrity of measurement, model, and recommendation must be protected against manipulation that could otherwise translate directly into compromised decisions (Mikati et al., 2025). The third is computational cost, which is itself a sustainability consideration. The energy consumed by sensing, transmission, storage, and model training is not negligible, and a framework whose purpose is decarbonization must account for its own footprint, favoring parsimonious models and efficient architectures where these do not materially compromise performance. Treating these concerns as cross-cutting rather than incidental ensures that the pursuit of analytical capability does not undermine the very objectives the framework is intended to serve.

4. Proposed Methodology and Validation Design

This section specifies how the framework can be operationalized and empirically evaluated. It is presented as a research design rather than a report of completed experiments; no results are claimed here, and the protocol is intended to be executed on real data by the investigator. This stance preserves scientific integrity while providing a concrete and reproducible pathway to validation. The overall workflow is shown in Figure 4.

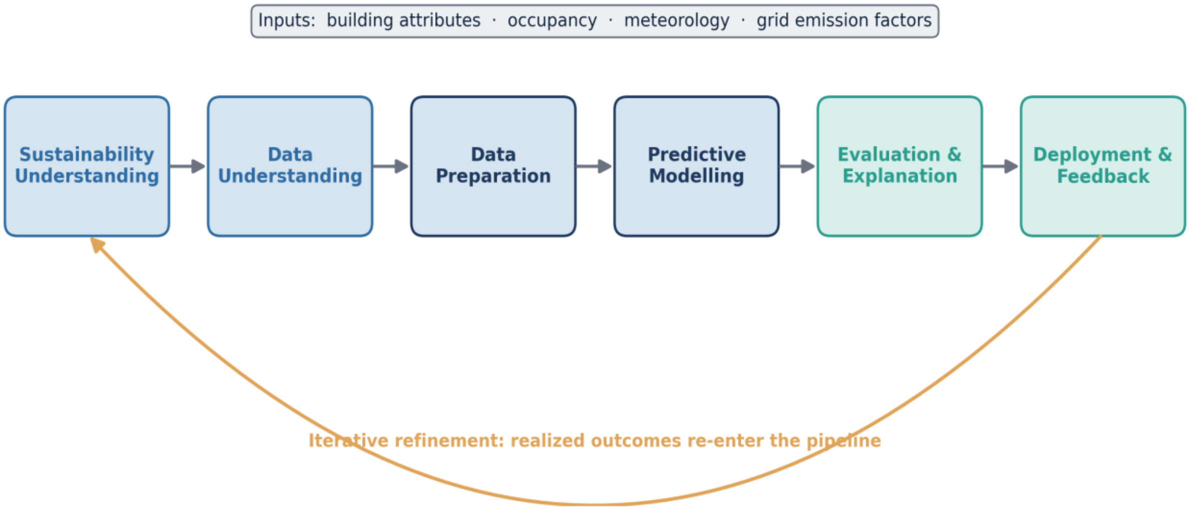

Figure 4: The proposed analytics workflow aligned with an established cross-industry process for data mining, from sustainability understanding through deployment, with realized outcomes feeding back into the pipeline.

4.1 Overall research design

The framework should be evaluated through a staged, mixed methods design that combines quantitative model assessment with qualitative evaluation of decision usefulness. A widely used process model for analytics projects, comprising business understanding, data understanding, data preparation, modelling, evaluation, and deployment, provides a disciplined structure for empirical study and ensures that each layer of the framework is tested in sequence (Wirth & Hipp, 2000). In the present context the initial phase is recast as sustainability understanding, in which the decarbonization objectives, the relevant decision-makers, and the criteria for success are defined before any modelling is undertaken. This reframing keeps the investigation oriented toward outcomes rather than accuracy from the outset, and it ensures that the subsequent technical work is disciplined by a clear statement of the decisions the analytics are meant to inform.

4.2 Data sources and variables

The empirical study should assemble a dataset spanning a defined portfolio of urban buildings or districts over a multi-year period sufficient to capture seasonal and inter-annual variation. Candidate predictors include building characteristics such as floor area, vintage, function, and envelope properties; operational variables such as occupancy and equipment schedules; meteorological variables such as temperature, humidity, and solar radiation; and grid variables such as time-varying emission factors used to convert predicted consumption into carbon terms. The target variables are energy consumption, expressed as energy use intensity to permit comparison across assets of different sizes, and the derived carbon footprint. Publicly documented urban energy datasets and municipal monitoring archives are appropriate sources, and their provenance, licensing, and known limitations should be recorded as part of the data-governance layer so that the conditions of validity are transparent. Where monitored data are sparse, the hybrid simulation-and-learning strategy discussed earlier offers a means of generating an adequate training corpus (Qiao et al., 2024).

4.3 Modelling pipeline

Following data preparation, a portfolio of estimators should be trained and compared under identical conditions to ensure a fair assessment. A reasonable portfolio includes a regularized linear baseline that establishes a lower bound on acceptable performance, tree-based ensembles such as random forest and gradient boosting for cross-sectional prediction, and a recurrent network for the temporal forecasting task (Breiman, 2001; Chen & Guestrin, 2016; Hochreiter & Schmidhuber, 1997). Hyperparameters should be tuned through cross-validation, and particular care must be taken to prevent temporal leakage by respecting chronological order in the validation splits, since the naive random partitioning of time-series data inflates apparent performance and produces models that fail in deployment. Feature-attribution analysis should be applied to the selected model so that the explainability layer is evaluated alongside predictive performance rather than as an afterthought, and the stability of attributions across resampling should itself be examined as evidence of the reliability of the explanations (Lundberg & Lee, 2017; Ribeiro et al., 2016).

4.4 Evaluation metrics

Predictive performance should be assessed with established regression metrics whose definitions are standard and are stated here only to fix the evaluation protocol. The mean absolute error and root mean squared error quantify average and variance-sensitive deviations in the original units, with the latter penalizing large errors more heavily; the mean absolute percentage error expresses error in relative terms to permit comparison across assets of different scale; and the coefficient of determination summarizes the proportion of variance explained. Carbon-related performance should be reported as the accuracy of predicted emissions against measured or independently estimated values, since accurate energy prediction does not guarantee accurate emissions prediction when emission factors are uncertain. Beyond accuracy, the evaluation should assess the stability of feature attributions and the usefulness of explanations to the intended decision-makers, the latter through structured qualitative appraisal, because an explanation that is technically valid but unintelligible to its audience fails the purpose for which the explainability layer exists. Table 2 summarizes the protocol.

Table 2

Evaluation metrics specified for the proposed validation protocol

Metric	Purpose
Mean absolute error (MAE)	Average magnitude of prediction error in original units
Root mean squared error (RMSE)	Error measure that penalizes larger deviations
Mean absolute percentage error (MAPE)	Scale-independent error enabling cross-asset comparison
Coefficient of determination (R^2)	Proportion of variance in the target explained by the model
Carbon prediction accuracy	Agreement between predicted and reference emissions
Explanation stability and usefulness	Consistency of attributions and value to decision-makers

4.5 Staged deployment and feedback

Validation should culminate in a controlled deployment in which the framework's recommendations are implemented for a subset of the portfolio while a comparable subset continues under existing practice, providing a basis for attributing observed changes to the intervention rather than to extraneous factors. Realized energy and emissions outcomes are then measured and fed back into the pipeline, testing the governance-feedback layer and providing evidence of impact under operational conditions rather than in simulation alone. This staged approach allows the framework to be assessed not only on predictive fidelity but on its capacity to produce measurable reductions in consumption and carbon footprint, which is the criterion that ultimately matters. It also surfaces the practical frictions of deployment, including the responses of operators and occupants to the recommendations, that are invisible in purely computational evaluation and that frequently determine whether analytical promise is realized in practice.

4.6 Threats to validity and reproducibility

A credible validation design must anticipate the threats that could undermine its conclusions. Internal validity is jeopardized chiefly by temporal leakage and by confounding in the deployment phase, the former addressed through chronologically respectful validation splits and the latter through the use of a comparable control subset against which intervention effects can be assessed. External validity is constrained by the specificity of any single portfolio, since a framework validated on one building stock or climatic context may not transfer directly to another; this threat is mitigated, though not eliminated, by validating across diverse contexts and by reporting the conditions under which performance was obtained. Construct validity depends on whether the chosen metrics genuinely capture the objective of decarbonization, which is why the protocol pairs energy-accuracy metrics with explicit carbon-accuracy metrics and with qualitative assessment of decision usefulness. Finally, reproducibility requires that data provenance, preprocessing steps, model configurations, and evaluation procedures be documented in full, so that the study can be independently repeated and its claims independently scrutinized. Addressing these threats systematically is what distinguishes a rigorous validation from a demonstration, and it is the standard to which any empirical application of the framework should be held.

5. Discussion

5.1 Theoretical implications

The principal theoretical contribution of this paper is the reframing of urban energy analytics as an integrated socio-technical system rather than a collection of discrete predictive tasks. By specifying how data acquisition, integration, modelling, explanation, optimization, and feedback interlock, the framework offers a scaffolding that connects the technical literature on prediction to the governance literature on decision-making, two bodies of work that have largely proceeded in isolation. It also advances the argument that interpretability is constitutive of, rather than supplementary to, credible sustainability analytics, since the binding constraint on adoption is repeatedly shown to be the intelligibility of model outputs rather than their accuracy (Doshi-Velez & Kim, 2017; El-kenawy et al., 2025; Lundberg & Lee, 2017). A further implication is that carbon accounting belongs within the analytical pipeline rather than downstream of it, because the objective of decarbonization is most faithfully represented as the minimization of emissions subject to constraints rather than the minimization of energy use in the abstract.

5.2 Practical and policy implications

For practitioners, the framework provides a blueprint for assembling analytics capabilities in a sequence that protects against common failure modes, most notably the temptation to invest in sophisticated models before securing data quality and governance. The layered structure makes it explicit that the value of advanced modelling is contingent on the integrity of the foundation beneath it, and that resources devoted to prediction will be wasted if the data and governance layers are neglected. For policymakers, the explicit alignment with the Sustainable Development Goals supplies a means of translating analytical results into the language of public commitment, and the governance-feedback layer offers a mechanism for demonstrating accountability by measuring interventions against realized outcomes rather than projected ones (United Nations, 2015). The emphasis on explanation is particularly consequential in public administration, where decisions must be justified to elected officials and citizens, and where unexplained algorithmic recommendations are unlikely to survive scrutiny or to command the legitimacy that sustained action requires.

5.3 Relation to existing approaches

The framework does not displace existing predictive methods; it situates them within a coherent whole. Tree-based ensembles, recurrent networks, digital twins, and feature-attribution techniques each retain their established roles, but they are coordinated toward a shared decarbonization objective rather than pursued in isolation for incremental gains in accuracy (Bibri & Krogstie, 2024; Breiman, 2001; Chen & Guestrin, 2016; Hochreiter & Schmidhuber, 1997). In this sense the contribution is complementary to methodological literature and directly responsive to the recurring observation that the field's principal weakness is fragmentation rather than any deficiency in its constituent techniques. The framework's value lies less in any single component than in the discipline it imposes on their combination and in the orientation, it gives to the system as a whole.

5.4 Generalizability and transferability

A question central to the practical value of any framework concerns the extent to which it transfers across settings. The architecture proposed here is deliberately specified at a level of abstraction that is independent of any particular city, building stock, or technology, so that its layers and their relationships remain applicable even as the specific instruments populating them vary. The estimators in the modelling layer, the attribution methods in the explainability layer, and the simulation environments in the decision-support layer are all interchangeable, which allows the framework to be instantiated with whatever methods are best suited to local conditions and available data. This abstraction is a strength, because it confers generality, but it also implies that transferability must be demonstrated rather than assumed. The hybrid simulation-and-learning strategy is especially pertinent to transferability, since it offers a route to instantiating the framework in data-poor cities that would otherwise be excluded from data-driven approaches (Qiao et al., 2024). Establishing the conditions under which the framework transfers successfully, and identifying where local adaptation is required, is among the most important tasks for the empirical research agenda set out below.

6. Limitations and Future Research

Several limitations qualify the contribution and define an agenda for further work. As a conceptual and methodological framework, its validity ultimately depends on empirical testing through the proposed design, and the results of that testing may reveal context-specific constraints not anticipated here. Data availability and quality remain a pervasive obstacle, particularly in smaller or resource-constrained cities where monitoring infrastructure is sparse, and the framework's performance in such settings warrants dedicated investigation; the hybrid simulation-and-learning strategy may partially mitigate this constraint but does not eliminate it (Qiao et al., 2024). Privacy and security considerations attending fine-grained energy data require careful and continuing treatment, since the same granularity that enables accurate prediction can expose sensitive patterns of occupancy and behavior if governance is inadequate (Mikati et al., 2025). The framework is also articulated primarily with reference to the building sector, which, while the largest single source of urban emissions, is not the only one.

Future research should pursue four directions. The first is empirical validation across multiple and diverse urban contexts to establish the generalizability of the framework and to characterize the conditions under which it performs well or poorly. The second is the extension of the approach from buildings to other infrastructure systems, including transport, water, and waste, so that cross-sector interactions and the trade-offs among them can be modelled within a common analytical environment. Third is the deeper integration of explanation with optimization, so that interventions are not only ranked but accompanied by intelligible justifications that decision-makers can act upon with confidence. The fourth is the systematic incorporation of equity considerations, ensuring that the benefits of data-driven decarbonization are distributed fairly across communities and that analytical systems do not inadvertently entrench existing disparities in service quality or environmental exposure. Progress along these directions would move the framework from a coherent proposition toward a validated and broadly applicable instrument of urban climate policy.

7. Conclusion

Urban infrastructure stands at the center of the global effort to reduce energy consumption and carbon emissions, and the convergence of ubiquitous sensing with mature predictive analytics has created an unprecedented opportunity to act on that responsibility. Yet the opportunity has been only partially realized, because analytical capability has too often been developed in isolation from the governance processes through which energy and emissions are managed, and because the opacity of advanced models has impeded their adoption in public decision-making. This paper has proposed an integrative predictive analytics framework that addresses these shortcomings by organizing data acquisition, integration, modelling, explanation, optimization, and feedback into a coherent system aligned with recognized sustainability objectives. By treating interpretability as a precondition for action, embedding carbon accounting within the analytical pipeline, and closing the loop between prediction and outcome, the framework connects technical performance to measurable decarbonization. The accompanying validation design provides a rigorous and reproducible pathway for empirical testing without recourse to fabricated evidence. Realizing the framework's potential will require sustained attention to data quality, interoperability, privacy, and equity, but it offers a principled foundation for transparent, accountable, and effective data-driven sustainability across the cities that will determine the trajectory of global emissions.

REFERENCES

1. Bibri, S. E., & Krogstie, J. (2024). The synergistic interplay of artificial intelligence and digital twin in environmentally planning sustainable smart cities: A comprehensive systematic review. Environmental Science and Ecotechnology, 20, 100433. https://doi.org/10.1016/j.ese.2024.100433

2. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
3. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939785
4. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv. https://doi.org/10.48550/arXiv.1702.08608
5. El-kenawy, E.-S. M., et al. (2025). Explainable AI for intelligent green energy forecasting: Deep learning with iHow optimization algorithm (iHOW). Scientific Reports, 15(1), 41158. https://doi.org/10.1038/s41598-025-26167-x
6. Elhabyb, K., Baina, A., Bellafkih, M., & Deifalla, A. F. (2024). Machine learning algorithms for predicting energy consumption in educational buildings. International Journal of Energy Research, 2024, Article 6812425. https://doi.org/10.1155/2024/6812425
7. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
8. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30, pp. 4765–4774). Curran Associates.
9. Mikati, A., et al. (2025). Integrating IoT and blockchain for smart urban energy management: Enhancing sustainability through real-time monitoring and optimization. Cluster Computing, 28(15), 960. https://doi.org/10.1007/s10586-025-05677-3
10. Qiao, R., et al. (2024). Predicting building energy consumption in urban neighborhoods using machine learning algorithms. Frontiers of Urban and Rural Planning, 2(1), 6. https://doi.org/10.1007/s44243-024-00032-3
11. 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 (pp. 1135–1144). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939778
12. United Nations. (2015). Transforming our world: The 2030 Agenda for Sustainable Development (A/RES/70/1). United Nations.
13. Wirth, R., & Hipp, J. (2000). CRISP-DM: Towards a standard process model for data mining. In Proceedings of the 4th International Conference on the Practical Applications of Knowledge Discovery and Data Mining (pp. 29–39).