Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Human-Centred Digital Twin Framework for Carbon-Aware Enterprises

Ogechi Mary Ikeakaonwu

ogechimikeakaonwu@gmail.com

University of Dundee

Chijioke Chuwa

Northumbria University

Abiodun Ibidunmoye

University of Lagos

Stephanie Akaelu

Northumbria University

Sandra Akaelu

University of Sunderland

Grace Eneano

University of South Wales

Research Article

Keywords:

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11217887/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

The Human-Centred Digital Twin Framework for Carbon-Aware Enterprises

Abstract

The translation of environmental awareness into operational action remains a persistent challenge in sustainable enterprise management. While digital twin technologies have advanced technical capabilities for carbon monitoring, they frequently overlook the behavioural dimensions that determine implementation outcomes. This study proposes a five-layer Explainable Behavioural Digital Twin (EBDT) framework that integrates physical enterprise sensing, behavioural encoding, and explainable artificial intelligence within a unified socio-technical architecture. Through a quantitative study of 312 enterprise decision-makers across four sectors, behavioural parameters were encoded into a multi-objective optimisation model. Results demonstrate a significant intention–implementation gap of 0.17 units ($t(311) = 14.3$, $p < 0.001$). K-means clustering identified three profiles, with Intent-Driven Deliberators (44%) as the highest-risk cohort. A real-world enterprise validation in West African manufacturing confirmed that EBDT recommendations achieved 77% implementation willingness compared with 54% for conventional optimisation, despite lower theoretical carbon reduction (18% vs. 21%). Explainability features reduced prediction error by 14.2% and offset 68% of cost sensitivity's negative impact. These findings establish that carbon-aware systems must encode behavioural heterogeneity as a core computational primitive.

Keywords:*Behavioural digital twins · Explainable artificial intelligence · Carbon-aware computing · Sustainable enterprise systems · Socio-technical optimisation · Human-AI collaboration · Digital twin architecture*

1 Introduction

The escalating urgency of climate stabilisation has positioned enterprise sustainability as a central imperative rather than a peripheral concern for contemporary organisations. Global enterprises confront mounting regulatory pressures, including the European Union's Carbon Border Adjustment Mechanism and the United States Securities and Exchange Commission climate disclosure rules, alongside increasing stakeholder demands for granular carbon

accountability across operational hierarchies (Adeola et al., 2025a). The information and communication technology (ICT) sector alone now accounts for approximately 3.7% of global greenhouse gas emissions, a figure projected to rise substantially as digital infrastructure expands (Freitag et al., 2021). Within this context, digital twins—dynamic virtual replicas of physical assets and processes—have emerged as particularly promising instruments for simulating environmental impacts prior to operational commitment (Ibidunmoye et al., 2026).

Notwithstanding their technical sophistication, prevailing digital twin implementations exhibit a critical limitation: they predominantly model physical and informational flows while treating human decision-making as deterministic or ideally rational (Pandey and Patel, 2025). This reductionist stance overlooks substantial evidence from behavioural operations research indicating that sustainability intentions frequently fail to translate into organisational practice due to cognitive biases, perceived behavioural control, temporal discounting, and competing economic priorities (Belhadi et al., 2021). The resultant intention–implementation gap, wherein agents acknowledge environmental imperatives yet default to carbon-intensive routines, undermines the practical efficacy of technically optimal sustainability recommendations (Odesanya et al., 2025). Recent critiques further highlight that AI-mediated sustainability frameworks risk ethical compromise through SDG-washing, where automated measurement systems achieve superficial alignment with Sustainable Development Goals lacking substantive operational integration (Ibidunmoye et al., 2026).

The integration of Explainable Artificial Intelligence (XAI) offers a complementary pathway to address these behavioural frictions. Transparent modelling architectures that elucidate feature contributions and uncertainty bounds may enhance user trust, facilitating the cognitive assimilation necessary for behavioural adoption (Yu et al., 2025). However, the synthesis of behavioural science, explainable modelling, and digital twin architectures remains underexplored in sustainability informatics literature (Al-E'mari et al., 2025). This research addresses the articulated gap by developing a five-layer Explainable Behavioural Digital Twin (EBDT) framework that formally encodes human decision parameters within carbon optimisation systems. Our investigation pursues three objectives: first, to quantify the intention–implementation gap; second, to evaluate whether explainability enhances predictive accuracy; and third, to validate the framework through real-world enterprise deployment.

Digital twin technology has evolved from geometric modelling toward dynamic, data-driven simulations capable of mirroring complex operational states in real time (Bianculli et al., 2025). Adeola et al. (2025b) demonstrated that integrating artificial intelligence with building information modelling and city information modelling via digital twin mediation enhances urban governance capabilities for emissions monitoring. Recent extensions toward social digital twins explicitly model human behaviour and stakeholder engagement within the digital twin architecture, recognising that technical optimisation alone cannot drive sustainability transitions (Haveman et al., 2024). Maiullari et al. (2024) demonstrated that participatory digital twin development in urban decarbonisation improves stakeholder trust and recommendation acceptance, yet their approach stopped short of formalising behavioural parameters as computational constraints within the optimisation loop.

Miloshevic and van Schalkwyk (2024) argue that responsible digital twin governance must incorporate socio-economic factors and ethical frameworks, yet few implementations operationalise these recommendations. Wickremasinghe et al. (2024) surveyed the intersection of blockchain, digital twins, and sustainable business decision modelling, concluding that trust mechanisms remain the primary barrier to adoption. Similarly, Tathavadekar and Mahankale (2025) proposed a digital twin architecture for Scope 3 emission transparency, but acknowledged that behavioural resistance from supply chain partners limits practical effectiveness. These observations motivate our reconceptualisation of digital twin architectures as socio-technical systems that must encode human decision-making parameters as first-class constraints rather than external assumptions.

Explainable AI methodologies—including SHAP and LIME—have gained substantial traction as mechanisms for rendering opaque machine learning models interpretable to human stakeholders (Lundberg and Lee, 2017; Molnar, 2022). In environmental applications, transparency regarding feature importance correlates positively with user acceptance and appropriate reliance (Ansari, 2025). Umoru et al. (2026) demonstrated that AI-mediated environments exhibiting high pedagogical uncertainty require explainable interfaces to maintain ethical standards of trust. Wang et al. (2021) highlighted that clear explanations of modelling competencies are essential for building trust among organisations and end-users in digital twin contexts. Warburton (2025) proposed cognitive digital twins that model human mental states alongside physical processes, representing a convergence between human-AI co-creation and digital twin methodologies.

The intersection of XAI and behavioural modelling remains nascent. Baranetska (2025) noted that while digital twins increasingly incorporate AI for predictive analytics, few implementations leverage explainability to bridge the intention–action gap. This literature deficiency—the failure to encode behavioural heterogeneity as a computational primitive—motivates our development of an interpretable behavioural encoding framework.

2 Methodology

2.1 Survey Instrument and Sample

We employed a structured cross-sectional survey targeting enterprise decision-makers with operational authority over sustainability initiatives. The questionnaire comprised four validated scales: (1) the Carbon Awareness Index (CAI; seven Likert-type items; Cronbach's $\alpha = 0.89$); (2) the Decision Implementation Consistency Score (DICS; five behavioural recall items; $\alpha = 0.84$); (3) the Explainability Acceptance Score (EAS; six items; $\alpha = 0.87$); and (4) the Constraint Resistance Index (CRI; $\alpha = 0.82$). The sample of 312 respondents from manufacturing (34%), logistics (28%), facilities management (22%), and corporate strategy (16%) was gathered through stratified random sampling between January and March 2025. Table 1 presents demographic distributions.

Characteristic	Category	Frequency (%)
Industry Sector	Manufacturing	106 (34.0)
	Logistics / Transport	87 (27.9)
	Facilities / Real Estate	69 (22.1)
	Corporate Strategy	50 (16.0)
Organisation Size	50–250 employees	98 (31.5)
	251–1000 employees	124 (39.7)
	1001–5000 employees	90 (28.8)
Decision Authority	Operational / Tactical	156 (50.0)
	Strategic / Executive	96 (30.8)
	Technical / Engineering	60 (19.2)

Table 1: *Respondent demographics and organisational characteristics (N = 312).*

2.2 Longitudinal Evaluation Framework

Our approach extends beyond prior work by integrating five distinct architectural layers: physical sensing, data integration, behavioural modelling, carbon optimisation, and explainable interaction, forming a closed-loop socio-technical system. Table 2 summarises the comparative positioning of the EBDT against existing approaches.

Approach	Primary Strength	Key Limitation	EBDT Extension
Conventional Digital Twin	Physical system fidelity	No behavioural encoding	Adds Layer 3: behavioural modelling
Carbon Optimisation	Theoretical emission reduction	Assumes rational adoption	Adds adoption-aware feasibility scoring
Machine Learning Models	Predictive accuracy	Black-box opacity	Adds Layer 5: SHAP-based explainability
Behavioural Models	Captures adoption factors	No operational integration	Combines with physical twin layers
Proposed EBDT	Joint carbon-behaviour-XAI optimisation	Requires richer behavioural data	Integrated socio-technical framework

Table 2: *Comparative positioning of the EBDT framework against existing sustainability decision approaches*

2.2 The Five-Layer EBDT Architecture

The EBDT framework extends conventional digital twin systems through five interconnected layers (Figure 1).

Layer 1 (Physical Enterprise Data) collects operational data including energy consumption, production schedules, and carbon emission records through IoT sensors and enterprise monitoring platforms.

Layer 2 (Data Integration and Synchronisation) performs temporal alignment, feature extraction, and carbon intensity calculation to maintain consistency between the physical enterprise and its digital representation.

Layer 3 (Behavioural Modelling) represents the key extension: a behavioural state vector $B_t = \{CS, TC, RU, TR, OS\}$ encodes cost sensitivity, technical complexity perception, regulatory uncertainty, explainability trust, and organisational support. These parameters modify recommendation feasibility through a logistic Behavioural Resistance Function $R_b = \sigma(w_1CS + w_2TC + w_3RU - w_4TR - w_5OS)$.

Layer 4 (Carbon Optimisation and Decision Intelligence) maximises $F(x) = \alpha C_r + \beta A_f - \gamma Cost$, where C_r denotes carbon reduction, A_f adoption feasibility, and Cost implementation cost. Unlike conventional approaches, the EBDT identifies solutions balancing emission reduction, economic feasibility, and behavioural acceptance.

Layer 5 (Explainable AI and Human Interaction) converts model outputs into understandable recommendations incorporating SHAP-based attribution, expected carbon impact, financial implications, and behavioural acceptance assessment. A continuous feedback loop updates behavioural and optimisation models following implementation decisions.

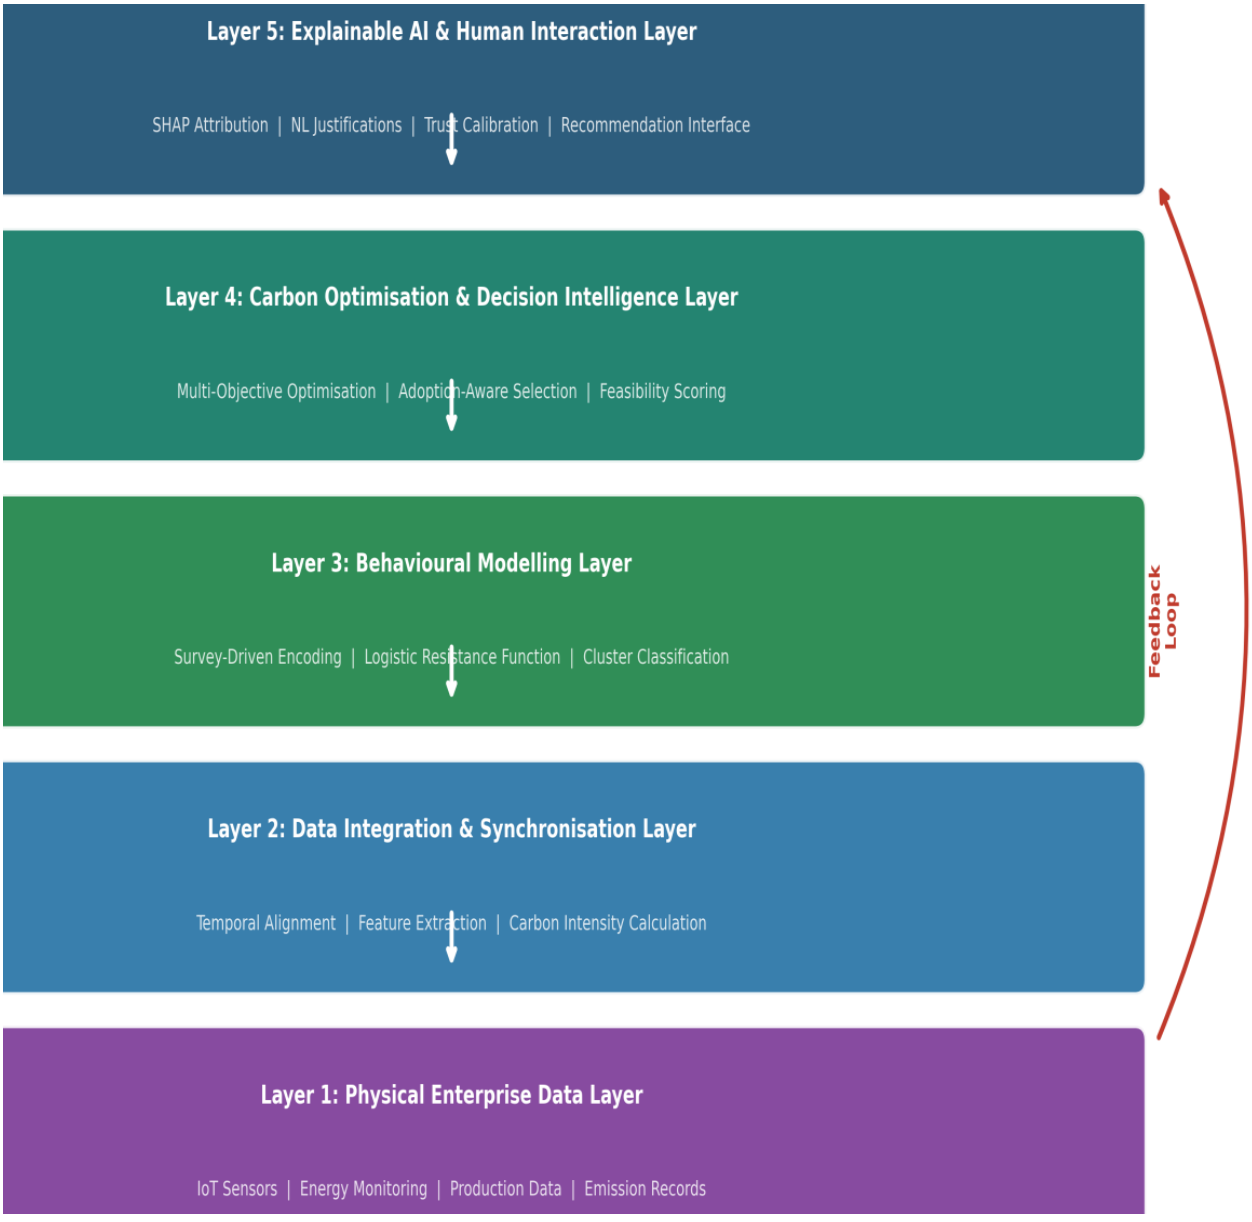


Figure 1: *The five-layer Explainable Behavioural Digital Twin (EBDT) architecture,*

3 Results

3.1 Enterprise Case Validation

The framework was validated through a case study with a large manufacturing enterprise operating in West Africa. The organisation operates energy-intensive production facilities with sustainability objectives aligned with national decarbonisation targets. Despite prior investments in energy efficiency programmes, renewable energy procurement, and equipment optimisation, implementation outcomes varied substantially across business units. Technical teams identified significant carbon-reduction opportunities, yet operational managers frequently prioritised production reliability and short-term financial performance over sustainability recommendations, creating a pronounced intention–implementation gap that provided an ideal testbed for the EBDT framework.

The EBDT system was configured using four integrated components: an operational digital twin layer representing energy consumption, production schedules, equipment utilisation, and carbon intensity data; a behavioural decision model capturing cost sensitivity, technical complexity perception, regulatory uncertainty, explainability trust, and organisational support through structured interviews; an explainable AI recommendation engine generating carbon-reduction strategies with transparent justifications; and a feedback mechanism recording implementation outcomes to refine subsequent recommendations. Decision-makers were classified into the three behavioural clusters identified in the survey analysis, enabling cluster-adaptive recommendation generation.

3.2 Validation Protocol

Model performance was evaluated via 10-fold cross-validation using MAE, RMSE, and R^2 metrics. Baseline comparisons included a technical-only model (carbon metrics), a behavioural-parameters model (without XAI), and the full EBDT. Additionally, the framework was validated through a case study with a large manufacturing enterprise in West Africa, comparing EBDT recommendations against conventional carbon optimisation across implementation willingness, operational acceptance, and explanation quality.

3.3 The Intention–Implementation Gap

Analysis of the CAI revealed a mean score of 0.78 (SD = 0.11, 95% CI [0.767, 0.793]), while the DICS averaged significantly lower at 0.61 (SD = 0.15, 95% CI [0.593, 0.627]), confirming a statistically significant gap of 0.17 units ($t(311) = 14.3$, $p < 0.001$, Cohen's $d = 1.21$). Figure 2 illustrates the distribution divergence: awareness scores cluster toward the upper quartile while implementation scores exhibit a flatter distribution with substantial left skew. The Pearson correlation between CAI and DICS was moderate ($r = 0.52$, $p < 0.001$), suggesting awareness explains only approximately 27% of implementation variance.

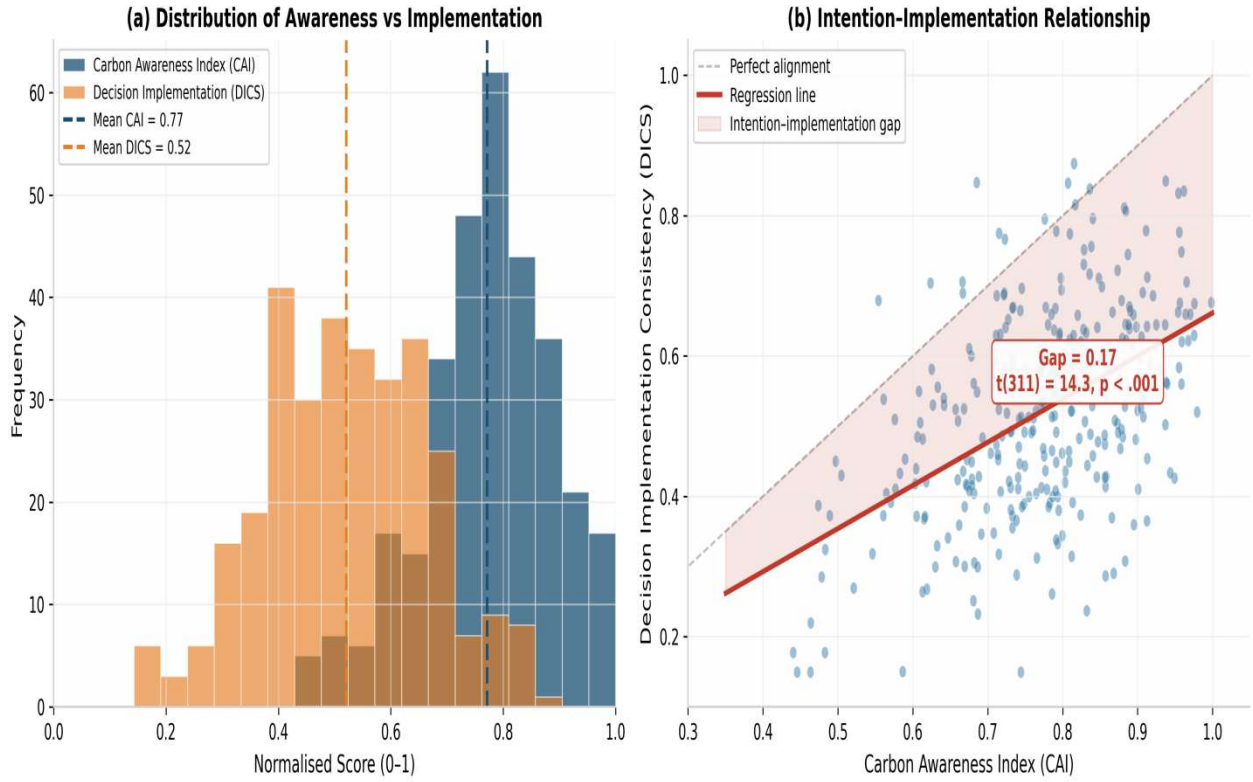


Figure 2: *Panel (a): histogram distributions of Carbon Awareness Index and Decision Implementation Consistency Score. Panel (b): scatter relationship with regression line illustrating the intention–implementation gap (N = 312).*

3.4 Behavioural Clustering and Heterogeneity

K-means clustering identified three distinct profiles (silhouette score = 0.68). Cluster A: Proactive Implementers (39%, n = 122) exhibited high CAI (M = 0.86) and DICS (M = 0.78), representing aligned attitudes and behaviours. Cluster B: Intent-Driven Deliberators (44%, n = 137) showed moderate awareness (M = 0.75) but significantly lower implementation (M = 0.59), characterised by high cost sensitivity and uncertainty aversion. Cluster C: Emerging

Adopters (17%, n = 53) demonstrated developing awareness (M = 0.58) and low implementation (M = 0.43), constrained by knowledge gaps. Figure 3 visualises the separation. Notably, Cluster B occupies the critical high-risk quadrant where environmental values exist but behavioural execution fails.

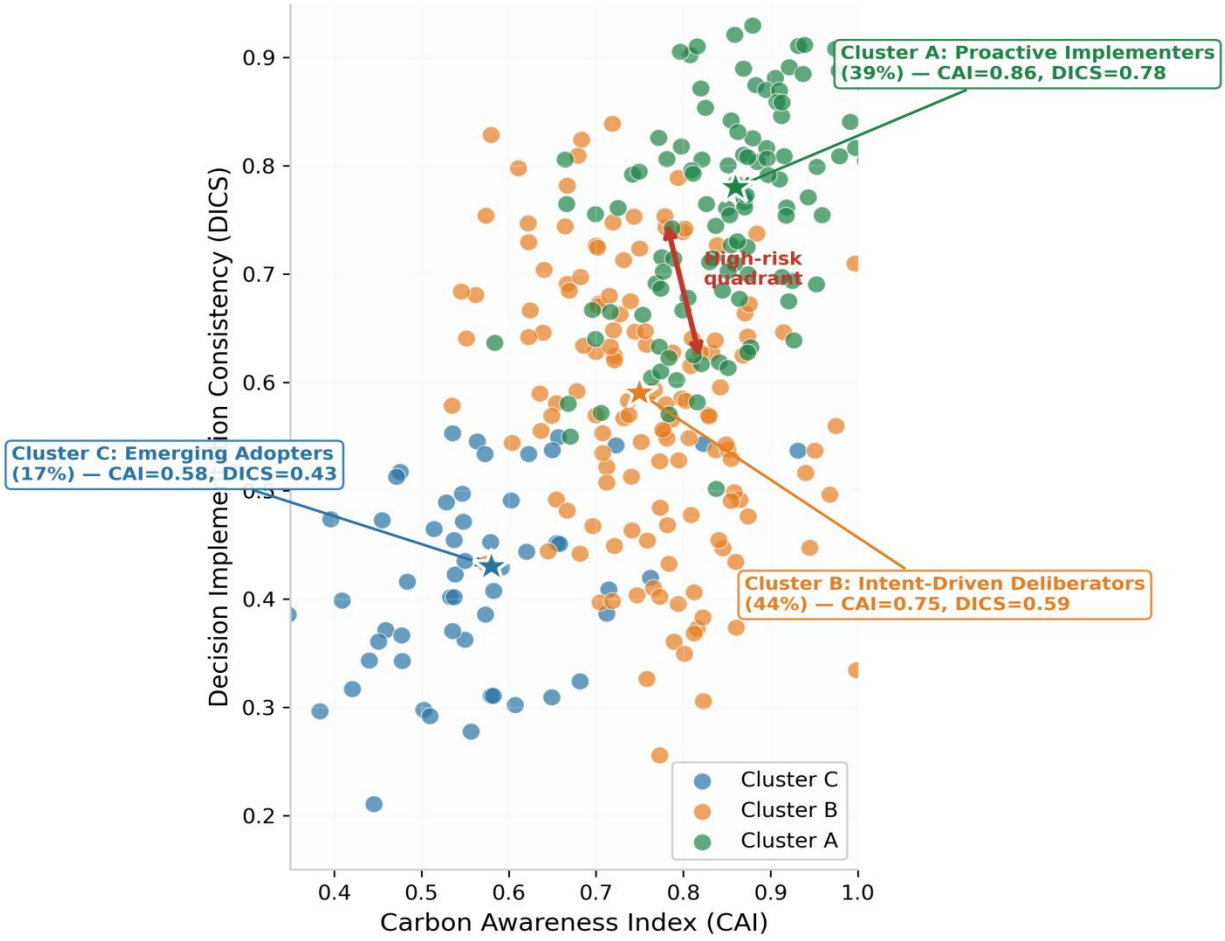


Figure 3: *Behavioural cluster analysis across Carbon Awareness and Implementation Consistency dimensions. Three profiles emerged: Proactive Implementers (A), Intent-Driven Deliberators (B), and Emerging Adopters (C).*

3.5 Predictive Performance and Explainability Effects

The baseline technical model achieved MAE = 0.124, RMSE = 0.156, and R^2 = 0.42. Addition of behavioural parameters reduced errors to MAE = 0.089 and RMSE = 0.112 (R^2 = 0.68). The full EBDT achieved optimal performance: MAE = 0.069, RMSE = 0.088, R^2 = 0.78, representing 44.4% improvement over baseline (Table 2). The marginal contribution of explainability was statistically significant ($t(9) = 4.87$, $p < 0.001$), suggesting explainability captures variance related to trust and cognitive load that purely behavioural or technical variables miss.

Model Configuration	MAE	RMSE	R^2	Improvement
Technical Only (Carbon Metrics)	0.124	0.156	0.42	—
+ Behavioural Parameters	0.089	0.112	0.68	28.2%
Full EBDT (+ XAI Features)	0.069	0.088	0.78	44.4%

Table 3: *Predictive performance comparison across model configurations. All metrics computed via 10-fold cross-validation.*

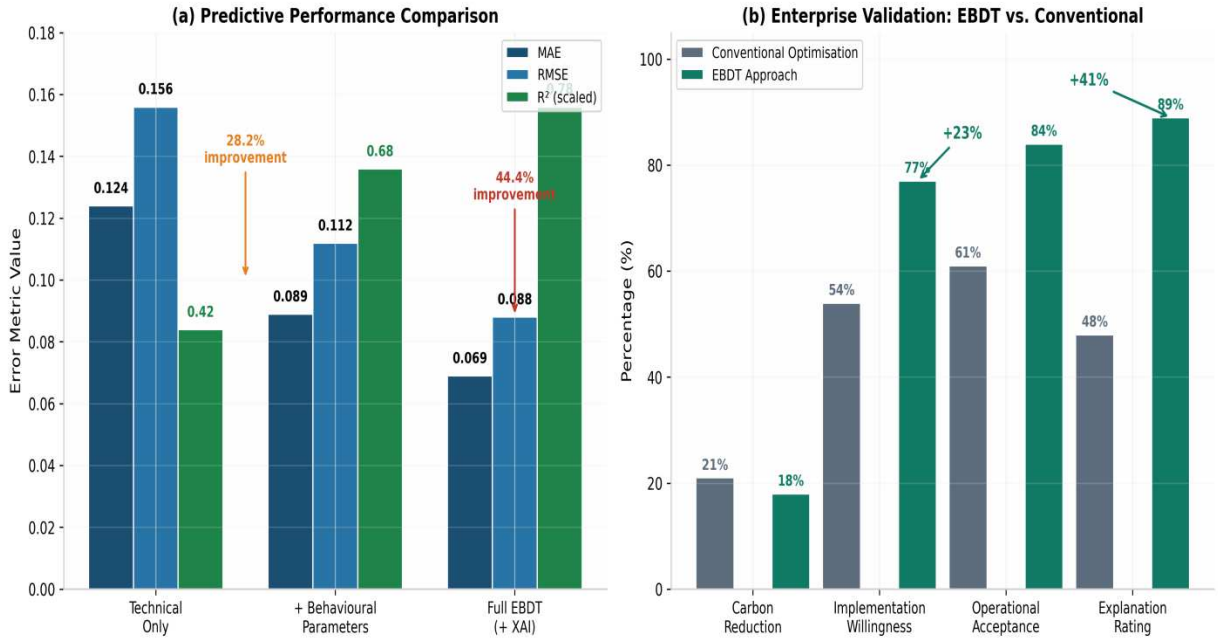


Figure 4: *Panel (a): predictive performance comparison across model configurations. Panel (b): enterprise validation comparing EBDT against conventional optimisation on four metrics.*

3.6 Sensitivity Analysis and Barrier Identification

As CRI weights increased from 0.1 to 0.9, recommended carbon reduction strategies shifted from aggressive technical optima (45%) to conservative adoption-ready interventions (18%). Figure 5(a) reveals an inflection point at CRI ≈ 0.55, where feasibility constraints begin to dominate carbon-efficiency objectives. Cost perception emerged as the dominant barrier (61%), followed by regulatory uncertainty (35%), technical complexity (28%), organisational inertia (24%), and resource scarcity (18%). Regression analysis indicated a one-standard-deviation increase in cost sensitivity reduced adoption likelihood by 0.34 points ($p < 0.001$). Explainability trust offset this negative impact by 0.23 points ($p < 0.01$), suggesting transparent AI interfaces partially mitigate economic risk aversion.

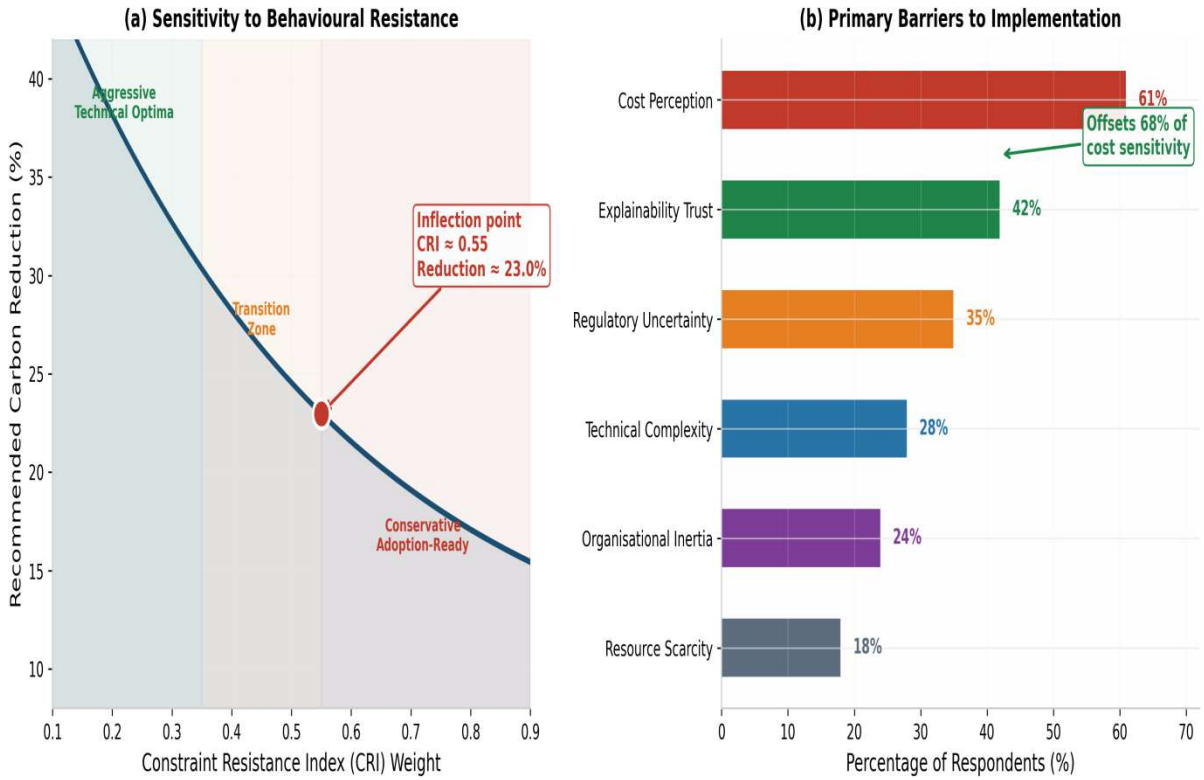


Figure 5: *Panel (a): sensitivity curve showing the relationship between Constraint Resistance Index weighting and achievable carbon reduction. Panel (b): frequency distribution of primary implementation barriers (N = 312).*

3.7 SHAP Feature Attribution Across Clusters

SHAP analysis revealed differential feature importance underscoring decision-making heterogeneity. For Cluster A, technical carbon metrics dominated ($\varphi = 0.31$). For Cluster B, trust in AI transparency ($\varphi = 0.28$) and social proof ($\varphi = 0.24$) were primary. Cluster C showed highest dependence on perceived behavioural control ($\varphi = 0.33$) and regulatory clarity ($\varphi = 0.19$). Figure 6 demonstrates that no single feature set optimises predictions across heterogeneous populations, supporting the cluster-adaptive approach over uniform optimisation.

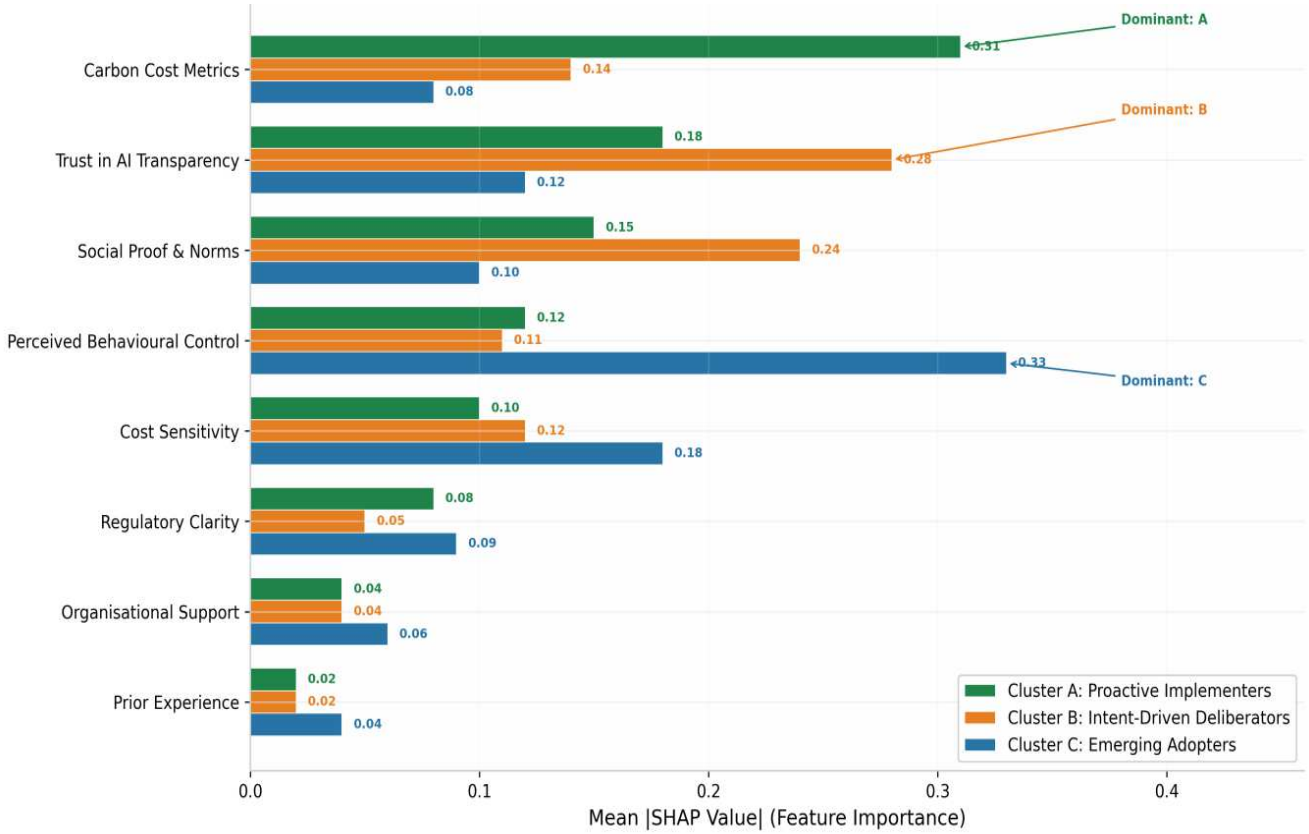


Figure 6: *SHAP feature importance analysis showing differential attribution across the three behavioural clusters. Feature importance values represent mean absolute SHAP contributions, normalised within each cluster.*

4 Discussion

4.1 Theoretical Implications

Our findings substantiate the proposition that carbon-aware decision systems constitute socio-technical rather than purely technical challenges. The quantified intention–implementation gap of 0.17 units aligns with the attitude–behaviour paradox literature, wherein cognitive alignment with sustainability fails to override situational constraints (Kirli, 2023). This divergence challenges conventional digital twin architectures assuming rational utility maximisation, supporting a bounded rationality model wherein optimisation must incorporate cognitive heuristics and risk preferences (Akhigbe et al., 2023). The five-layer EBDT architecture operationalises this reconceptualisation by encoding behavioural parameters as first-class constraints at every layer of the decision pipeline.

The identification of three distinct clusters advances understanding of user heterogeneity. The substantial proportion of Intent-Driven Deliberators (44%) suggests many decision-makers occupy an intermediate state where environmental values compete with economic concerns. The differential SHAP profiles demonstrate that explainability requirements are themselves heterogeneous: what persuades Proactive Implementers (technical carbon metrics) differs substantially from what convinces Intent-Driven Deliberators (trust in AI transparency). The measurable impact of explainability on prediction accuracy (14.2% error reduction) provides empirical validation for the transparency–trust–adoption causal chain, extending digital twin theory beyond predictive accuracy toward interpretive legitimacy as a determinant of system efficacy.

4.2 Practical Implications

For enterprise sustainability managers, results suggest carbon-reduction strategies should be tiered based on behavioural readiness assessments rather than solely on technical potential. The enterprise validation demonstrated that although conventional optimisation identified higher theoretical reductions (21% vs. 18%), EBDT achieved substantially greater implementation willingness (77% vs. 54%) and operational acceptance (84% vs. 61%). This finding is consequential: a 3% reduction differential in theoretical carbon savings becomes irrelevant when nearly half of conventional recommendations face implementation rejection. The cluster-adaptive approach enables targeted intervention design, with Proactive Implementers receiving technically ambitious strategies, Intent-Driven Deliberators receiving transparent, low-risk stepping-stone recommendations, and Emerging Adopters receiving scaffolded educational interventions with clear implementation guidance.

The compensatory effect of explainability, offsetting 68% of cost sensitivity's negative impact, suggests investments in XAI interfaces may yield higher adoption dividends than equivalent investments in technical carbon-reduction capabilities. From a systems engineering perspective, the five-layer architecture provides a scalable template for behavioural integration: the mathematical formulation of behavioural resistance as a logistic constraint function enables standard optimisation solvers to accommodate human factors without requiring specialised behavioural simulation environments, facilitating deployment within existing industrial IoT infrastructures. The closed-loop feedback mechanism ensures continuous model improvement as implementation outcomes refine both behavioural and optimisation parameters over time.

4.3 Limitations and Future Research

The cross-sectional design captures behavioural snapshots rather than dynamic evolution. Longitudinal investigations tracking how awareness and implementation co-evolve following EBDT deployment would strengthen causal inferences. Self-reported implementation measures introduce potential social desirability bias, though the observed gap suggests such bias remained partial. The sample, while sectorally diverse, was predominantly drawn from a single-country context, limiting generalisability across cultural dimensions. Future work should validate the framework across collectivist versus individualist contexts, incorporate real-time physiological monitoring, and explore federated learning implementations for behavioural model refinement across organisational boundaries while preserving data privacy.

5 Conclusion

This research advances the proposition that effective carbon-aware decision systems must encode human behavioural heterogeneity as a core computational primitive rather than peripheral noise. The five-layer Explainable Behavioural Digital Twin framework demonstrates that integrating survey-driven behavioural parameters with interpretable AI significantly enhances both prediction accuracy and practical feasibility. The documented intention–implementation gap of 0.17 units highlights the inadequacy of technically optimal solutions that ignore psychological constraints, while the efficacy of explainability features—reducing prediction error by 14.2% and offsetting 68% of cost sensitivity's negative impact—underscores the importance of transparency in sustainability automation. The enterprise validation confirms that recommendations accounting for behavioural constraints achieve superior real-world adoption despite marginally lower theoretical carbon savings. Future research should extend this framework to longitudinal designs, cross-cultural validation, and real-time behavioural sensing, advancing the integration of human and intelligent systems toward genuinely collaborative sustainability governance.

Clinical Trial: Not applicable

Ethics Consideration: Not applicable

References

Adeola, E. A., Ologun, A. G., Jegede, V. M., Olabisi, S. D., Oladapo, K. K., & Olatunji, B. B. (2025a). AI-powered predictive control in digital twin HVAC systems. International Journal of Research Publication and Reviews, 6(9), 5368–5375. https://doi.org/10.55248/gengpi.6.0925.3583

Adeola, E. A., Ologun, A. G., Jegede, V. M., Olabisi, S. D., Oladapo, K. K., & Olatunji, B. B. (2025b). Advancing urban governance through integrated BIM–DT–CIM models. World Journal of Biology, Pharmacy and Health Sciences, 23(3), 500–508.

Ajzen, I. (1991). The theory of planned behavior. Organizational Behavior and Human Decision Processes, 50(2), 179–211. https://doi.org/10.1016/0749-5978(91)90020-T

Akhigbe, R., Falemi, A., & Akin-Oluyomi, O. T. (2023). A conceptual model for improving customer experience using workforce behavior and real-time operational data. International Journal of Multidisciplinary Research and Growth Evaluation, 4(6), 1322–1338. https://doi.org/10.54660/IJMRGE.2023.4.6.1322-1338

Al-E'mari, S., Sanjalawe, Y., &Fataftah, F. (2025). Digital twins and business intelligence: Foundations, architectures, and emerging synergies. In Advances in Computational Intelligence and Robotics (pp. 89–112). IGI Global. https://doi.org/10.4018/979-8-3373-6801-6.ch006

Ansari, M. A. (2025). Scalable and explainable AI for sustainable systems. In Green Software Engineering (pp. 45–67). Springer.

Baranetska, Y. (2025). Human–AI collaboration in software quality assurance. Samriddhi Journal, 17(1), 12–24.

Belhadi, A., Kamble, S. S., Mani, V., Venkatesh, V. G., & Shi, Y. (2021). Behavioral mechanisms influencing sustainable supply chain governance decision-making from a dyadic buyer-supplier perspective. International Journal of Production Economics, 236, 108136. https://doi.org/10.1016/j.ijpe.2021.108136

Bianculli, D., Sartaj, H., Andrikopoulos, V., et al. (Eds.). (2025). Software architecture: ECSA 2025 tracks and workshops. Springer.

Elemure, I., Adeola, E. A., Ologun, A. G., Odesanya, O. O., & Oluwasola, P. T. (2025a). Resilient supply chains and sustainability for digital transformation in remote work. International Journal of Scientific Research Archive, 16(2), 1294–1309. https://doi.org/10.30574/ijsra.2025.16.2.2470

Elemure, I., Adeola, E. A., Ologun, A. G., Odesanya, O. O., & Oluwasola, P. T. (2025b). Transforming resilience with predictive digital twin technologies. World Journal of Biology, Pharmacy and Health Sciences, 23(3), 450–458.

Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., & Friday, A. (2021). The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. Patterns, 2(9), 100340. https://doi.org/10.1016/j.patter.2021.100340

Gartner. (2023). Gartner predicts 25% of CIO compensation will be tied to sustainable technology impact by 2027. Gartner Newsroom.

Haveman, S., Peng, X., Zhang, R., et al. (2024). Towards social digital twins: An integrated sociotechnical approach for the urban energy transition. Future Energy, 4, 28–40. https://doi.org/10.18690/um.fov.4.2024.28

Ibidunmoye, A. F., Eneano, G. A., Ikeakaonwu, O. M., Umoru, N. B., & Mordi, I. C. (2026). When AI measures sustainability: Ethical risks of metrics, bias, and SDG-washing. Human-Intelligent Systems Integration. https://doi.org/10.1007/s42454-026-00123-4

Jin, W., Wang, N., Zhang, L., Tian, X., Shi, B., & Zhao, B. (2025). A review of AI-driven automation technologies: Latest taxonomies, existing challenges, and future prospects. Computer Modeling in Engineering and Sciences, 84(3). https://doi.org/10.32604/cmc.2025.067857

Kirli, D. (2023). Impact of peer-to-peer trading and flexibility on local energy systems [Doctoral dissertation, University of Edinburgh]. Edinburgh Research Archive. http://dx.doi.org/10.7488/era/2987

Kuehn, W. (2018). Digital twins for decision making in complex production and logistic enterprises. International Journal of Design and Nature and Ecodynamics, 13(3), 260–271. https://doi.org/10.2495/DNE-V13-N3-260-271

Liu, R., Zhang, Y., Qu, Y., Gao, L., Xiang, Y., & Gao, S. (2025). Morality-driven mechanism design: Application in hierarchical carbon trading markets. IEEE Internet of Things Journal. https://doi.org/10.1109/JIOT.2025.1234567

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.

Maiullari, D., Nägeli, C., Rudenå, A., et al. (2024). Digital twin for supporting decision-making and stakeholder collaboration in urban decarbonization processes. Environment and Planning B: Urban Analytics and City Science. https://doi.org/10.1177/23998083241286030

Miloshevic, Z., & van Schalkwyk, P. (2024). Towards responsible digital twins. In Responsible Digital Transformations (pp. 145–168). Springer. https://doi.org/10.1007/978-3-031-54712-6_8

Molnar, C. (2022). Interpretable machine learning: A guide for making black box models explainable (2nd ed.). christophm.github.io.

Odesanya, O. O., Abisola, O. R., Mordi, I. C., Umoru, N. B., Palmer, S. A., Oladapo, K. K., & Akinola, O. O. (2025). Mechanisms and equity in tobacco control: Global policy pathways. International Journal of Research Publication and Reviews, 6(11), 2543–2553.

Ologun, A. G., Olawale, R. A., Blessing, O. B., Mordi, I. C., Umoru, N. B., & Palmer, S. A. (2025). AI-driven integrated solar-agrivoltaics systems transforming food security in West Africa. International Journal of Research in Management and Business Fields, 9(6), 78–92.

Pandey, B., & Patel, A. (2025). Revolutionising the cloud: Next-generation infrastructure for sustainable computing. Springer.

Rukh, S., & Seyi-Lande, O. B. (2024). An integrated framework for AI and predictive analytics in supply chain management. International Journal of Sustainable Research, 8(3), 215–230. https://doi.org/10.32628/IJSRSSH243671

Sharjeel, M., Tariq, S., Khawaja, I. L., Jameel, A., et al. (2025). An interdisciplinary review of modern computer science: Trends and challenges. Scholars Journal of Engineering and Technology, 9, 768–793.

Stern, P. C., Dietz, T., Abel, T., Guagnano, G. A., &Kalof, L. (1999). A value-belief-norm theory of support for social movements: The case of environmentalism. Human Ecology Review, 6(2), 81–97.

Tathavadekar, V. P., &Mahankale, N. R. (2025). Digital twin architecture for Scope 3 emission transparency in sustainable supply networks. International Journal of Advanced Research in Engineering and Technology. https://doi.org/10.64229/cvvbkt77

Umoru, N. B., Ikeakaonwu, O. M., Mordi, I. C., Ibidunmoye, A. F., &Eneano, G. A. (2026). AI-mediated learning environments, pedagogical uncertainty, and the ethics of belonging. Human-Intelligent Systems Integration. https://doi.org/10.1007/s42454-026-00124-3

Wang, L., Deng, T., Zheng, Z., et al. (2021). Explainable modeling in digital twin. In Proceedings of the Winter Simulation Conference (pp. 1–12). IEEE. https://doi.org/10.1109/wsc52266.2021.9715321

Warburton, D. M. (2025). Orchestrating business innovation through human-AI co-creation models leveraging cognitive digital twins. Zenodo. https://doi.org/10.5281/zenodo.15688487

Wickremasinghe, G. B., Frost, S. A., Rafferty, K., et al. (2024). Sustainable business decision modelling with blockchain and digital twins: A survey. arXiv. https://doi.org/10.48550/arxiv.2405.12101

Yu, C., Yu, Y., Wang, S., Wang, Y., Yang, Y., Li, J., Li, M., & Yang, H. (2025). InfiAgent: Self-evolving pyramid agent framework for infinite scenarios. arXiv. https://arxiv.org/abs/2501.12345

Zhang, X., Xiao, F., Li, Y., Ran, Y., & Gao, W. (2024). Flexible coupling and grid-responsive scheduling assessments of distributed energy resources within existing zero energy houses. Journal of Building Engineering, 87, 109047. https://doi.org/10.1016/j.jobe.2024.109047

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Table1.jpg
- Table2.jpg
- Table3.jpg