

Design and Assessment of Federated Digital Learning Policy Models for Sustainability-Oriented Higher Education

Dr. Nitish Anand ^a, Dr.R. Vasanthan ^{b*}, Dr. Anuj Kumar ^c, Dr.R. Thiyagarajan ^d,
Dr.G. Kalpana ^e, Dr.V. Dhanasekaran ^f

^a Faculty of Social Science, University of Delhi, North Campus, New Delhi, Delhi, India.

E-mail: nitishanand2891@gmail.com, Orcid: <https://orcid.org/0009-0009-6778-3816>

^{b*} Associate Professor, Department of English, Nagaland University, Nagaland, India.

E-mail: vasanthan@nagalanduniversity.ac.in, Orcid: <https://orcid.org/0000-0002-7193-4711>

^c Assistant Professor, Department of Hindi, Nagaland University, Kohima Campus, Meriema, Nagaland, India.

E-mail: anuj@nagalanduniversity.ac.in, Orcid: <https://orcid.org/0009-0005-6840-5447>

^d Assistant Professor, Department of Management Studies, SRM TRP Engineering College, Tiruchirappalli, Tamil Nadu, India. E-mail: thiyagarajan.r@trp.srmtrichy.edu.in, Orcid: <https://orcid.org/0000-0003-3187-1741>

^e Assistant Professor, Department of Management Studies, SRM TRP Engineering College, Tiruchirappalli, Tamil Nadu, India. E-mail: kalpana.g@trp.srmtrichy.edu.in, Orcid: <https://orcid.org/0009-0007-1813-5981>

^f Assistant Professor, Department of Mathematics, Velammal College of Engineering and Technology (Affiliated to Anna University), Madurai, Tamil Nadu, India. E-mail: vds@vcet.ac.in, Orcid: <https://orcid.org/0009-0009-7203-3817>

Abstract

The explosive growth of digital learning ecosystems within higher education has increased the requirement for policy mechanisms that can solve issues related to data privacy, institutional autonomy, inclusive participation, computational efficiency, and sustainable learning outcomes. Centralized analytics are highly dangerous from the perspective of data sovereignty, while existing federated learning algorithms give priority to performance and the number of samples without considering universities with limited resources and neglecting institutional policies concerning equity and carbon footprints. To fill this gap, this research proposes a novel Federated Sustainability Policy Evaluation and Orchestration (FedSustain-PEO) approach. FedSustain-PEO consists of a Machine Executable Policy Translation Engine that turns qualitative institutional policies into computable constraints, a Dynamic Policy Constraint Graph for representing policy conflicts, and a Policy Adaptive Federated Controller for regulating the behavior of the training process. Moreover, the Policy Aware Sustainability Aggregation technique allows balancing the multi-objective optimization with respect to five key criteria, including Policy Compliance, Sustainability Learning Outcomes, Inclusion, Privacy Risk, and Computational Costs. As evaluated on a heterogeneous testing ground of 12 institutional nodes, FedSustain-PEO showcased outstanding results relative to existing standards and baselines. It provided an accuracy of 94.8%, a precision of 94.1%, and an F1 score of 93.8%. As for its performance in governance and equity, it provided a Policy Execution Compliance Index of 0.947, a Federated Institutional Equity Score of 0.918, and cut policy violation rates to 4.8%. Further, it provided a Sustainability Learning Gain of 0.842, decreased overall energy usage by 26.5% (to 41.6 kWh), and was capable of maintaining a high participation rate of 78.6% among resource-constrained institutional nodes. In summary, FedSustain-PEO provides a scalable and privacy-protected basis for distributed university learning. By turning the policy evaluation process into a multi-objective continuous orchestration procedure, it is capable of balancing prediction utility with data privacy, equity, and sustainability.

Keywords: Federated Learning, Sustainability Education, Policy Evaluation, Digital Learning Systems, Higher Education, Privacy-Preserving Analytics.

Introduction

Today, higher education is rapidly moving to a digital transformation that has restructured the provision of learning materials, the way students learn, and the way instructors make decisions. A large amount of learner data is generated and processed more and more, for example, by learning management systems, AI, predictive analytics, adaptive learning platforms, or cloud-based educational services. Systematic evidence suggests that AI-assisted education is growing at various educational levels and sectors, with challenges ranging from privacy, ethics, transparency, infrastructure, and equitable access related to its use (Garzón et al., 2025). The implications of these issues are especially relevant in higher education, where digital learning spaces involve all aspects of learning, such as academic records, behavioral traces, assessment results, interaction history, and competency indicators. In parallel, demand for the incorporation of sustainability by universities into courses, processes and skills is mounting. Recent studies have underlined the role of new AI tools in fostering socio-environmental integration in engineering education and the importance of putting in place frameworks for responsible implementation of AI (Cañavate et al., 2025). Designing the digital learning system, therefore, requires us to take sustainability outcomes, institutional governance, privacy and resource responsibility, as well as predictive accuracy and technological efficiency, into account.

To ensure sustainability education in higher education, learning systems need to be able to facilitate and support environmental literacy, climate awareness, sustainable decision-making, multi- and interdisciplinary problem-solving, and skills training regarding the Sustainable Development Goals. Using digital technologies to support these goals can be achieved with personalized learning paths, the ongoing assessment of competencies, intelligent feedback, and cross-institutional collaboration. But there is a structural conflict when it comes to personalization and responsible data governance because of this growing influx of data-driven educational systems. Centralized learning analytics usually necessitate the movement of learners' data from various sources to a common server, which means it is subject to different risks of unauthorized access, institutional-housed data-sovereignty concerns, and the incurrence of inappropriate secondary uses. Federated learning (FL) provides an alternative approach to enable the participating institutions/ devices to jointly train a common model, while keeping the raw data on the local devices (Yurdem et al., 2024). Recently, a comprehensive review of FL has been provided, which recognizes the very good capabilities of FL to resolve the issues of privacy and confidentiality in sensitive application environments. However, privacy protection in FL is still an ongoing challenge since any shared model updates can be susceptible to threats of inference, poisoning, and reconstruction, necessitating robust technical and governance principles to safeguard the integrity of FL processes.

The importance of FL in the education arena is becoming more and more apparent. Over the course of a base study, ultimately uploaded to the website, it presents an action-oriented federated approach to overcome the privacy challenges of emotion recognition in the education sector from the perspectives of both local processing of data by smartphones, tablets, and laptops, as well as data aggregation and deployment using Moodle (Gutiérrez et al., 2025). This study shows that emotionally sensitive data could also be processed locally, and that updates to the models could enable collaborative learning, indicating the feasibility of applying FL in today's heterogeneous educational scenarios. It places more emphasis, however, on recognition of emotions, privacy preservation, and engagement in the educational field as opposed to computational assessment of institutional sustainability policies (Sum et al., 2026). This distinction suggests a research opportunity: Federated digital learning infrastructures could go beyond prediction in the context of privacy and be adapted for the ongoing assessment of whether institutional policies related to sustainability education, digital inclusion, equity, privacy, and resource efficiency are implemented in system behavior.

Another problem is related to the "Environmental impact of Distributed Intelligence". While FL does not centralize raw data upfront, there could be multiple rounds of training and communications in a local FL setting across diverse devices, which is potentially highly time-consuming and costly both in terms of energy and computation. We propose the notion of "Green Federated Learning" and show that, besides model performance and training times, the carbon footprint metric should be considered for designing large-scale FL systems. There are recent works focusing on carbon-efficient FL where the carbon intensity of all computing resources is not considered equal and can be explicitly considered in the selection policies (Li et al., 2024). It has been shown that energy-awareness research, together with quantization and joint optimization of local computation and communication, can significantly reduce the energy demand while ensuring the convergence requirements (Kim et al., 2023). All these are vital for the idea of learning taking sustainable forms for HEs, because the learning benefits of a groundbreaking tool, if obtained via high computational costs, cannot be regarded as sustainable.

But while making a lot of progress, a considerable divide lies between the intentions of education policies and the running of the federated system. Digital learning policies usually take the form of "how to" guidelines, compliance mandates, statements of principles, or targets. These policies can include minimum standards for privacy protection, equal participation in institutions, access for everyone, sustainability skills development, and/or decreased computational footprint. But such requirements do not imply directly operational choices about the selection of clients, the intensity of local training, weights used for aggregation, privacy parameters, the frequency of communication, and so on, when it comes to conventional FL algorithms. Also, for the selection of clients to be carbon-efficient, the same set of clients with efficient infrastructure and technology can be continuously selected, unintentionally decreasing the number of resource-poor universities. The result is an important policy dilemma about environmental efficiency and the fairness of institutions. Hence, a need for today's sustainability-based higher education for mechanisms that can determine not only the accuracy of the functioning of a model but also whether the behavior of distributed learning is in line with the higher education institution's sustainability commitments.

To fill this gap, as part of this research, a Federated Sustainability Policy Evaluation and Orchestration (FedSustain-PEO) framework for the development of sustainability policies and their orchestration among higher education (HE) institutions is proposed. The proposed architecture portrays a Machine-Executable Policy Translation Engine that translates institutional policy demands into measurable federated constraints; a Dynamic Policy Constraint Graph that represents the conflicts existing between the institutional goals - privacy, sustainability, inclusion, and equity and resource goals; and a Policy-Adaptive Federated Controller that dynamically regulates institutional participation and training behavior. Moreover, the developed Policy-Aware Sustainability Aggregation mechanism considers the institutional contribution (partly from sample size but also) based on policy compliance, sustainability learning gain and digital inclusion deficit, privacy risk, and computational cost, instead of merely relying on sample size. Thus, the study moves the spotlight on policy evaluation from "after" the application to a "computational-only" mode of evaluation, and to an adaptive orchestration mode. Proposes an approach for creating a scalable engineering basis for responsible DL ecosystems in heterogeneous Higher Education settings based on Federated Intelligence, sustainability education, policy compliance, and institutional equity and resource-awareness optimization.

Here are the unique contributions of the paper:

- Proposes the first federated learning framework specifically designed to evaluate and orchestrate sustainability policies across higher education institutions without sharing raw data.
- Introduces a translation engine that converts qualitative institutional policy requirements into measurable, operational federated constraints.
- Utilizes a Dynamic Policy Constraint Graph (DPCG) to model and resolve trade-offs between competing institutional goals like data privacy, equity, and carbon efficiency.
- Deploys a controller and aggregation mechanism that balances predictive accuracy with sustainability learning gains, digital inclusion, and resource efficiency.

This paper is structured into six core sections. Section 1 introduces the digital transformation of higher education and outlines the research gap. Section 2 reviews prior literature on federated learning, privacy, and sustainability. Section 3 details the methodology and the FedSustain-PEO architecture. Section 4 presents the experimental results and comparative performance metrics. Section 5 discusses limitations, and Section 6 concludes the study.

Review of Literature

AI that supports the policy process, digital higher education, sustainability-focused pedagogy, and federated learning are now intersecting and represent an emerging research landscape (Rakhmanovich et al., 2025; Sing et al., 2025). Recent research shows that distributed data analytics are now becoming an increasingly important reliance in higher education to make decisions related to student success, personalized interventions, and student progression. The study explored the use of federated learning techniques for educational data analytics on a student dropout prediction problem and showed that it is possible to model collaborative learning from educational data that are distributed without sending raw data to a central repository (Deng et al., 2025). Based on their research study, they concluded that federated learning is technically relevant for learning analytics in the context of universities with limited possibilities of data exchange (Miller et al., 2025). However, this research mostly focused on predictive analytics and lacked consideration of sustainability-policy goals, machine-executable institutional limitations, and adaptive policy evaluation for the federation process (Fachola et al., 2023).

The aspect of privacy is a major issue, as data not necessarily stored in a central location doesn't mean it is impossible to launch an attack on the exchanged parameters. Comprehensively studied the threats to security and privacy in federated systems and demonstrated that the vulnerabilities in updating models are prone to inference, poisoning, and other adversarial methods. Their analysis highlighted the need for extra safeguards to protect privacy, as they didn't assume local data storage alone is enough to do it (Blanco-Justicia et al., 2021). A particular problem area that is relevant for digital learning environments is the issue of the transferability of traces of student engagement, competency profiles, and indicators of certain institutions, even if raw student records are not circulated. Hence, policy-applicable Federated Education demands an explicit mechanism that can probe the privacy constraints during model training rather than it being an inherent property of the architecture.

An associated research strand is that of fairness amongst the participating heterogeneous federated entities. Surveyed the connection between privacy and fairness in FL, and found that ensuring multiple goals such as privacy protection, fairness, participation of all clients, and model performance is challenging. Their findings suggest that if the data distributions are not independent across all clients, or client characteristics are not homogeneous, this can lead to unequal outcomes, especially when the aggregation mechanisms bias towards, for example, those with higher resources, or more representative data sets (Chen et al., 2023). Recent studies have further corroborated this worry. An investigation of trends and challenges in fairness for FL was added, where the bias can manifest in different ways, such as client selection, data heterogeneity, optimization behavior, and performance disparities among clients (Mukhtiar et al., 2025). In the context of the higher education network, these results suggest that traditional federation models could potentially be unintentional ways of giving an advantage to technologically privileged universities and underrepresenting those with poorer infrastructure or fewer data sets, or restricted access to learners who are disadvantaged.

An additional layer of complexity is added, namely the sustainability dimension. However, studies on the digital transformation leading towards sustainability in higher education have been investigated, which found digitalization and institutional transformation to have strong interactions towards sustainable development (Dong, 2026; Lee, 2025). Their research demonstrated that while digital transformation can contribute to sustainable actions within a university's environment, coordinated strategies are needed to combine technological innovation with the goals of the organization and education (Kuzmanic et al., 2023). This is important because issues of digital learning sustainability cannot be narrowed down to just digital platforms; it must include an analysis of how digital infrastructures enable learning, social and organizational goals, as well as environmental goals. However, few sustainability-related studies in the context of higher education challenge computation rules to directly affect distributed AI training.

Research on Education for Sustainable Development (ESD) also shows that these are not just curricular measures, but measures that must be put in place throughout the entire institution of the University. Looked into ESD implementation in Higher Education Institutions and determined organizational, pedagogical and structural conditions that affect the integration of sustainability in the curricula (Mokski et al., 2023). Their results reinforce the need for coordinated actions from institutions but do not seem to decouple from the traditional methods of curriculum analysis, surveys, competency assessment, or organizational indicators underpinning a general characterization of ESD in curricula evaluation. These methods offer good evidence, and are mostly retrospective and periodic. Do not continually review to see if a digital learning system is computing in a manner that meets sustainability, inclusion, privacy and equity standards.

Federated computation's environmental expense has just started to draw attention. The study examined energy-efficient federated learning designs and found that local computation, wireless communication, client selection, resource allocation, and model complexity are key factors influencing the energy consumption of FL (Dang et al., 2024). The review illustrates that improving FL efficiency must be done by considering both heterogeneous hardware and the communication states or by optimizing strategies, instead of simply reducing the number of training rounds. While energy-aware FL offers a valuable first step towards a distributed intelligence that is environmentally friendly, merely reducing energy usage can introduce undesired effects of inequity. For example, a big effort to identify the highest competition, best performance, and most energy-efficient nodes can over time lead to a decrease in the involvement of resource-poor institutions, which in turn creates a conflict between energy efficiency and educational inclusion.

The literature reviewed thus shows a disjointed research field. Preserving privacy during collaborative analytics using FL, addressing information leakage through privacy research, client-level disparity through fairness research, and energy-efficient FL lowering the computational burden are the four areas of education that are considered supported by FL. But there is not yet a strong link between these research streams. Previously developed research lacks a single architecture that can transform policies at the level of higher education into federated constraints, discover

contradictions between the goals of sustainability and privacy, inclusion and equity, and resource use, and adjust the aggregation behavior. The gap provides motivation for the proposed FedSustain-PEO framework, which goes beyond the traditional assessment of the policies and integrates policy evaluation and orchestration into the federated learning cycle.

Materials and Methods

Higher Education Federated Test Environment

The proposed FedSustain-PE implementation is carried out via a heterogeneous higher-education federated test environment, comprising data from multiple institutions. The data carries different information related to sustainable education priorities, student population, computational capacity, and the digital infrastructure of each institution. For experimental purposes, the environment is designed with 12 institutional nodes, categorized into four high-resource, four medium-resource, and four less-constrained resource nodes.

Each institutional node kept its own dataset, composed of anonymized student learning interactions, sustainability competency indicators, assessment results, their patterns of digital participation, course engagement logs, and the attributes of resources used. These local datasets were intentionally designed as non-Independent and Identically Distributed (non-IID) to introduce the natural variation among institutional characteristics such as number of students, learners' age, gender, and ethnicity, and the number of students taking sustainability courses, as well as access to digital tools. Both institutions acted as standalone federated clients, performing their own data preprocessing, training models locally, checking policy compliance, and generating their own privacy-protected model updates. Communication rounds happened through a central orchestration server, which at no point had access to raw data provided by students, but got selections of the candidates based on the policy, aggregated the local data from those candidates eligible to participate in the specific round, verified the policy restrictions with institutions, and delivered the refreshed global model to interested institutions again. The testing setting also added variability at the network level in terms of bandwidth and computational resources, the privacy level (P3), inclusion requirements, and resource limits in the sustainability domain; thus, providing the floor for assessing the framework under conditions that reflect the real-life institutional variety while maintaining data sovereignty and enabling domain-wide use of sustainability learning analytics in distributed higher-education environments in fig 1.




Fig. 1: FedSustain-PE Test Environment

Multi-Institutional Sustainability Learning Data Collection

The data collection process took place in the 12 involved institutional nodes, aiming to create a heterogeneous representation of learner-oriented digital learning for sustainability without transporting raw learner data across institutional nodes. Data were collected and stored from academic activities, as well as sustainability-related courses, digital learning interactions, and institutional technology environments, separately by each institution. The three dimensions of the collection framework are the interaction behavior in LMS, sustainability competency and engagement characteristics, and institutional resource and digital inclusion conditions. To achieve interoperability, a shared data schema was defined, whilst at the same time giving institutions control over data residing on their own storage media. A random synthetic ID was used for each of the learners, allowing the possibility of longitudinal analysis, as people can be followed over time, without the direct personal identifier being stored in the federated training data. Data was organized by institution into local data repositories and then broken up into academic weeks and learning activity chunks. In addition to the collection process capturing individual-level learning characteristics, the institutional-level operational conditions were also captured and fed into the proposed framework FedSustain-PEO to differentiate between the learning characteristics differences and the infrastructure-related differences. Each of the institutional nodes also performed completeness checks, schema validation, duplicate removal, and attribute consistency verification in advance of the local model training. Here, a distributed collection strategy was used and enabled sustainability learning analytics, policy evaluation, institutional equity assessment, and privacy-preserving federated model development in realistic non-IID settings.

Learning Management System Interaction Data

Locally collected data from the different Digital Learning Platforms (DLP) of the higher education institutions (HEI) participating in the study were used to obtain Learning Management System (LMS) interaction data. The data included behavior produced while engaging in common and expected learning tasks, such as no identifiable learner data moving to the central orchestration server. The key features ranged from login frequency per week to duration of each session, course-page visits, learning material visited (count), video completion rates, whether the student logged in to submit assignments, quiz attempts, whether the student participated with feedback on the discussion forum, whether the student accessed feedback from the teacher, whether the student finished assessment, and inactivity interval. Excessive granularity at the interaction event level was avoided by summarising the interactions at a low level of precision (on a weekly basis), but information about the temporal patterns in interactions was kept. Other variables included late submission, number of assessment attempts, frequency of access to the sustainability module, sequence of accessing resources, and sustained access between academic weeks. These records were kept in the federated node of each institution, and each learner had his/her identity was obscured by an anonymous code before being processed for features. Later, the interaction attributes were transformed into indicators like learning persistence, participation intensity, resources utilized, assessment consistency, and digital engagement stability, and were normalized. To maintain cross-institutional heterogeneity, there were no constraints on a uniform distribution of behaviors; rather, the differences between institution platforms used, course organization, etc have all been maintained. This assimilation of behaviors from LMS served as the main behavioral input to estimate the number of people participating in sustainability learning, the decline in learning participation, and to provide local model training within the proposed FedSustain-PEO framework according to unified State policies.

Sustainability Competency and Engagement Indicators

Sustainability competency and engagement indicators were gathered to quantify students' acquisition of sustainability-related knowledge and behaviors, their use of them in the digital learning process, and how sustainability-related knowledge and behaviors were retained/engaged during the learning process. The competency dimension comprised environmental literacy, understanding of climate change, understanding of Sustainable Development Goals, systems-thinking, resource-efficiency knowledge, ethical decision-making, understanding of the circular economy, and ability to evaluate socio-environmental trade-offs. These indicators were identified through assessments in the sustainability modules, scenario quizzes, rubrics for projects done in courses, evaluations of case studies, reflective tasks, and problem-solving activities in courses that participated in the project. Pre-learning and post-learning competency scores were locally kept to be able to calculate individual and institutional sustainability learning gains to capture learning progression. Indicators for engagement comprised continental sustainability content completion rate, engagement in the environmental discussion forum, frequency of using green resources, participation in the feedback analysis, participation in projects, participation in voluntary activities, and continuity of engagement over academic weeks. A number of parameters to describe the competency and participation attributes of a learner were proposed, and from them, a composite Sustainability Learning Engagement Profile was built for each individual, where every individual learner was pseudonymized. The framework further documented shifts in engagement in

response to digital interventions to assess the extent to which policy-driven adaptations had an impact on learning. Because sustainability priorities were not uniform across institutions, local weighting configurations were adhered to in order to keep the institution-specific emphasis on competency. These indicators allowed FedSustain-PEO to assess the educational effectiveness, predictive performance, adherence to institutional policies, participation equity, and consideration of computational resources during federated learning.

Institutional Resource and Digital Inclusion Indicators

Indicators were designed and gathered on institutional resources and digital inclusion to capture different structures that might impact participation in federated sustainability learning. The resource variables used were available CPU capacity, GPU availability, memory allocation, the average local training time, the network bandwidth, the communication time, the availability of devices, the reachability of the servers, the available storage space, as well as the estimated amount of energy consumed for each round of training. Along with that, institutions were also defined based on the presence of a digital platform, the technical support capacity, reliability in working with the learning system, and how often connectivity was lost. Digital inclusion indicators measured the percentage of learners who had reliable internet and access to suitable learning devices, the frequency of their internet being interrupted, their mobile-only participation, having access to digital sustainers support, and consistency of use of the digital sustainability resources. These variables were considered at the cluster (institutional or coherent cluster, approved) level to prevent giving unnecessary information about the socioeconomic characteristics of individuals. A Federated Resource Capacity Score (computational power of each node) was used to represent the computational power of each institution to make them part of collaborative training, and a normalized Digital Inclusion Deficit (participation disadvantage in each institution) was used to identify institutions that were experiencing structural participation disadvantages. This is how these indicators were integrated in the policy-aware client selection and aggregation—it's not just a question of how much data each institution had, or how much computing power was available. Thus, participation priority might be strengthened for an institution with interesting inclusion needs, but which is limited in resources. This design allowed FedSustain-PEO to precisely assess the trade-off between ensuring efficient computation, institutional representation, digital accessibility, and education outcomes in line with the goals of sustainability.

Data Preprocessing and Privacy Preservation

Each institution dealt with data preprocessing and data privacy separately as a separate unit node before the federated model training process began. The system was set up to make data more consistent and to ensure the raw data for the learners was kept in the source Higher Education Institution (HEI). As the multi-institutional dataset also included measures of LMS interaction patterns, sustainability competency measures, engagement indicators, computational resource attributes, and digital inclusion variables, significant variation across the different local distributions, scales, and completeness existed. Thus, the preprocessing pipeline included handling missing values, scaling the features, encoding the categorical features, screening for anomalies, and manually checking for local consistency in the data. Differently, privacy protection was built right into the preprocessing process, namely by pseudonymization, minimization of attributes, operative local transformation, and controlled local generation of model updates. Maintain is placed in the local dataset as follows:

$$D_k = \{(x_i^{(k)}, y_i^{(k)})\}_{i=1}^{n_k} \quad (1)$$

Let $x_i^{(k)}$ represent the feature vector of a learning outcome, $y_i^{(k)}$ denote the corresponding learning outcome, and n_k be the number of local records. As established in equation (1), each university maintains its own distinct dataset and independently executes its local data preprocessing. This architecture ensures that only protected model updates and policy-related summary indicators are communicated across the network, thereby safeguarding institutional data sovereignty throughout the lifecycle of the proposed FedSustain-PEO model.

Algorithm 1: FedSustain-PEO Training and Orchestration Loop

Input: K institutional nodes, T communication rounds, local datasets D_k , policy tuples $\mathcal{P}_j$

Output: Optimized global model weights w^G

1. Initialize global model parameters w^0
2. For round $t = 1$ to T do
3. Server selects participating subset of institutions S_t using Policy-Adaptive Controller (PAFC)
4. For each institution $k \in S_t$ (in parallel) do

5. Receive current global model w^t
6. Evaluate local state $Z_k^t = [Q_k^t, P_k^t, I_k^t, R_k^t, C_k^t, S_k^t]$
7. Check machine-executable policy constraints via MEPTE (Eq. 16–18)
8. Adjust local training parameters $a_k^t = [E_k^t, \eta_k^t, q_k^t, \rho_k^t, f_k^t]$ via PAFC (Eq. 23–27)
9. Initialize local weights $w_k^{(0)} = w^t$
10. For epoch $e = 1$ to E_k^t do
11. Compute composite local objective L_k^{SA} (Eq. 28–29):
$$L_k^{SA} = 0.40L_{pred} + 0.25L_{sus} + 0.15L_{pol} + 0.10L_{priv} + 0.10L_{cost}$$
12. Update local weights:
$$w_{k,e+1}^t = w_{k,e}^t - \eta_k^t \nabla L_k^{SA} \quad (\text{Eq. 30})$$
13. End For
14. Compute local model update: $\Delta w_k^t = w_k^E - w^t$
15. Apply differential privacy clipping:
$$\Delta w_k^t = \Delta w_k^t \cdot \min\left(1, \frac{C}{\|\Delta w_k^t\|_2}\right) \quad (\text{Eq. 6})$$
16. Send protected update Δw_k^t and policy summary indicators to Central Server
17. End For
18. Server performs Policy-Aware Sustainability Aggregation based on compliance, learning gain, inclusion, and cost:
$$w^{t+1} = w^t + \sum (\text{weight}_k \cdot \Delta w_k^t)$$
19. End For

Algorithm 1 shows the proposed FedSustain-PEO framework involves repeated coordination of learning in a decentralized manner in several higher education institutions through T rounds. In each round, the central server picks nodes according to the policy-adaptive controller. The picked institutions compute their respective states, impose constraints on the policies according to machine-implementable constraints, and then optimize a local loss function which is multi-objective in nature.

Missing-Value Treatment and Feature Normalization

There is a lack of LMS activities for all students and inconsistencies among different digital session interactions, a lack of institutional resource measurements and changes in the participation of sustainability assessment. To avoid introducing any systematic bias, a two-stage local imputation procedure was used, which depends on the type of feature used. Continuous features with moderate missingness were filled with the median of the corresponding feature among the local institutions due to less sensitivity to extreme behaviours. The imputed value was assumed to be:

$$\tilde{x}_{ij}^{(k)} = \begin{cases} x_{ij}^{(k)}, & \text{if } x_{ij}^{(k)} \text{ is observed} \\ \text{median}(X_j^{(k)}), & \text{if } x_{ij}^{(k)} \text{ is missing} \end{cases} \quad (2)$$

Where $x_{ij}^{(k)}$ is feature j for learner i at institution k . The reason was that missing continuous observations were not filled in with information concentrated on a global level as outlined in equation (2), but with information present in the institutions. If a category was missing for a categorical attribute, the most common category for that attribute in that locality was used; if $>10\%$ of categories were missing, they were not included in the models due to lack of data.

After all imputation was done, numerical attributes were scaled down so that those with a large scale (e.g. session duration and bandwidth) would not overpower other indicators of a smaller scale (e.g. competency scores). After the min-max normalization, it is done as:

$$z_{ij}^{(k)} = \frac{\hat{x}_{ij}^{(k)} - \min(X_j^{(k)})}{\max(X_j^{(k)}) - \min(X_j^{(k)}) + \varepsilon} \quad (3)$$

Where $z_{ij}^{(k)}$ is the value divided by its maximum value in the row and ε is a small numerical constant to avoid dividing by 0. The continuous features were warped to the range $[0,1]$ in equation (3), but the local institutional distributions remained intact. The process resulted in similar feature scales for sustainability learning, sustainability engagement and inclusion, as well as computational-resource analysis, not requiring centralized access to the raw data.

Institutional Data Anonymization

To minimize the risk of linking the learning records to identifiable individuals, institutional data anonymization was completed prior to building local features. Direct identifiers such as learner names, registration numbers, e-mail addresses, phone numbers, and Internet platform identities (logins) were removed from the data set to be analyzed. Using a cryptographic hash function, each learner identifier $ID_i^{(k)}$ was converted to a pseudonymous token, with an institution-specific random salt used to create the pseudonymous token:

$$PID_i^{(k)} = H(ID_i^{(k)} \parallel S_k) \quad (4)$$

This is where $H(\cdot)$ refers to the cryptographic hash function, S_k is institution k 's private salt and $\parallel$ denotes concatenation. Equation (4) illustrates that two different representations of the same learner ID were created in two different institutions, reducing the opportunities for cross-institutional linkages.

In order to increase the level of disclosure protection, the quasi-identifying numeric characteristics of the data were generalized into bounded numeric ranges. Interval-based transformation was defined as: A sensitive continuous attribute X is an attribute that can be assigned a value on a continuous interval $[a,b]$.

$$G(x) = \left\lfloor \frac{x}{b} \right\rfloor b, \left(\left\lfloor \frac{x}{b} \right\rfloor + 1 \right) b \quad (5)$$

Where b represents the value of the interval size selected; exact values have been substituted by generalized ranges, where it is not needed to have a fine-grained value to make a conclusion about learning (as outlined in equation (5)). In addition, locally updated models were cut off before being sent away to ensure that the impact/possible exposure to individual records was not overly influenced:

$$\Delta \hat{w}_k = \Delta w_k \cdot \min \left(1, \frac{C}{\|\Delta w_k\|_2 + \varepsilon} \right) \quad (6)$$

Where Δw_k is the original local update, and C is the clipping threshold, tacitly limits the amount of an update introduced by equation (6) prior to the federated aggregation. So, the approach of privacy preservation means that identifiers can be removed to aggregate data without letting external actors identify it; the aggregation can be deliberated with generalization; the data can be minimized, kept for local processing, and eventually transmitted only to update it under control.

Non-IID Data Partitioning Across University Nodes

The federated environment was intentionally set up under non-independent and identically distributed conditions, as different Universities have different populations of learners, sustainability curricula, degree of digitalization and learners' engagement, institutional resources, and accessibility situations. Controlled statistical heterogeneity was achieved by a Dirichlet-based partitioning mechanism in the 12 university nodes. The proportions of allocations at the institutional level per unit of learning outcome (competency class) were sampled as:

$$\pi_c \sim \text{Dirichlet}(\alpha 1_K) \quad (7)$$

Here, the class allocation proportions $\pi_c = (\pi_{c,1}, \pi_{c,2}, \dots, \pi_{c,K})$ across K institutions are controlled by parameter α , which determines the amount of heterogeneity. Smaller α -values resulted in higher institutional imbalance; bigger α -values resulted in more homogeneous distributions, as indicated by equation (7).

The number of observations of class c was determined by:

$$n_{k,c} = \lfloor \pi_{c,k} N_c \rfloor \quad (8)$$

Where N_c is the total number of available observations belonging to class c . According to equation (8), this ensured that each University had different proportions of sustainability learning outcomes and engagement profiles. The following three heterogeneity conditions were considered:

$$\alpha \in \{0.1, 0.5, 1.0\} \quad (9)$$

The value of $\alpha=0.1$ was very heterogeneous institutional data, $\alpha=0.5$ was moderately heterogeneous, and $\alpha=1.0$ was comparatively balanced institutional data, as shown in equation (9). The Jensen-Shannon divergence was computed for each individual local node to the overall federation as a measure of distributional divergence as:

$$JSD(P_k \parallel P_G) = \frac{1}{2} D_{KL}(P_k \parallel M) + \frac{1}{2} D_{KL}(P_G \parallel M), M = \frac{P_k + P_G}{2} \quad (10)$$

Where, P_k is, is the institutional distribution locally and P_G is is global distribution. The cross-institutional heterogeneity was measured as a bounded value, as given by equation (10). This setting allowed FedSustain-PEO to be tested with realistic variations of sustainability competency and engagement behaviour, digital inclusion and institutional capacity of resources instead of with artificially homogeneous data.

Development of the Proposed FedSustain-PEO Framework

The design of FedSustain-PEO aims at a policy-adaptive federated learning architecture for sustainability education for different higher education institutions. While traditional federated systems usually consider primarily the sample size of local updates, this one considers together the following aspects: sustainability learning performance, policy compliance, institutional inclusion, privacy risk, and computation load. It is built from five interwoven components: Local Sustainability Learning Analytics Model, Machine-Executable Policy Translation Engine, Dynamic Policy Constraint Graph, Policy-Adaptive Federated Controller, and Sustainability-Aware Local Model Training mechanism. All 12 institutional nodes work with their local dataset independently, assess local conditions of the policies and learn local model without sending off raw learner records at any time. The state of the institution on the communication round t is written in the form of:

$$Z_k^t = [Q_k^t, P_k^t, I_k^t, R_k^t, C_k^t, S_k^t] \quad (11)$$

Where local predictive quality, policy compliance, inclusion priority, privacy risk, computational cost, and sustainability learning gain together define the state. As equation (11) makes clear, federation decisions rest on this multidimensional institutional profile rather than on accuracy alone. The specific numerical settings used to implement the framework appear in table 1.

Table 1: Numerical Configuration of the Proposed Fedsustain-PEO Framework

Parameter	Symbol	Numerical Value	Function
Number of institutional nodes	K	12	Federated university clients
Total learner records	N	30,000	Multi-institutional dataset
Communication rounds	T	100	Global training iterations
Local batch size	B	32	Local mini-batch training
Initial learning rate	η_0	0.001	Gradient optimization
Minimum learning rate	$\eta_{\min}$	0.0001	Lower adaptation boundary
Local epochs	E_k	1–5	Adaptive local training
Client participation ratio	q_t	0.40–0.80	Dynamic node participation
Policy compliance threshold	τ_P	0.80	Minimum acceptable compliance
Inclusion threshold	τ_I	0.60	Minimum participation adequacy
Privacy-risk threshold	τ_R	0.20	Maximum tolerated risk
Sustainability-gain threshold	τ_S	0.65	Minimum learning gain
Update clipping threshold	C	1.0	Limits local update magnitude
Policy adaptation interval	A_I	5 rounds	Controller reassessment
Maximum policy priority	$\pi_{\max}$	5	Mandatory policy importance

Local Sustainability Learning Analytics Model

The Local Sustainability Learning Analytics Model was developed to provide information about the extent to which learners develop competencies for sustainability as well as to identify learners whose learning progress may require adaptive learning supports. Each of the participating universities received LMS interaction logs, competency indicators, engagement metrics, and digital-participation features, which were used to train a lightweight multilayer

neural network locally. The network consists of an input layer compatible with the local feature vector, three hidden layers (with 128, 64, and 32 neurons, respectively), and an output layer that uses a sigmoid function to produce a binary sustainability-competency prediction throughout the work. Followed by the first and second hidden layers, respectively, are dropout rates of 0.3 and 0.2. This architecture is summarized in table 2.

Table 2: Numerical Architecture of the Local Sustainability Learning Analytics Model

Layer	Input/Output Dimension	Activation	Dropout
Input layer	20 features	NA	NA
Dense layer 1	128 neurons	ReLU	0.30
Dense layer 2	64 neurons	ReLU	0.20
Dense layer 3	32 neurons	ReLU	0.10
Output layer	1 neuron	Sigmoid	NA

As table 2 indicates, the network slowly reduces this 20-dimensional learning feature representation to one probability for sustainability competency. The hidden representation of learner i at institution k at layer l is:

Machine-Executable Policy Translation Engine

The executable policy configuration is provided in table 3.

Table 3: Machine-Executable Policy Rules and Numerical Thresholds

Policy Dimension	Executable Rule	Threshold	Priority	Triggered Action
Privacy protection	$R_k^t \leq \tau_R$	0.20	5	Reduce update exposure
Policy compliance	$P_k^t \geq \tau_P$	0.80	5	Activate compliance correction
Digital inclusion	$I_k^t \geq \tau_I$	0.60	4	Increase participation priority
Sustainability gain	$S_k^t \geq \tau_S$	0.65	4	Adjust local learning intensity
Resource utilization	$U_k^t \leq \tau_U$	0.85	3	Reduce local epochs
Communication latency	$L_k^t \leq \tau_L$	250 ms	2	Modify communication frequency

As shown in table 3, privacy and policy compliance received the highest priority value of 5, whereas latency was assigned a lower operational priority of 2. For policy j , the normalized violation magnitude was calculated as:

Dynamic Policy Constraint Graph

The policy interaction network configuration in numbers is presented in table 4.

Table 4: Numerical Configuration of Policy Conflict Relationships

Policy Pair	Dependency Strength d_{ij}	Conflict Threshold	Typical Trade-Off
Privacy–Predictive Utility	0.80	0.15	Protection vs. model quality
Inclusion–Resource Efficiency	0.75	0.15	Representation vs. computation
Sustainability–Local Training	0.70	0.15	Learning gain vs. energy use
Compliance–Communication Cost	0.55	0.15	Monitoring vs. overhead
Equity–Carbon Efficiency	0.85	0.15	participation vs. environmental cost

Sustainability-Aware Local Model Training

The numerical coefficients are provided in table 5.

Table 5: Numerical Coefficients of the Sustainability-Aware Local Objective

Objective Component	Symbol	Weight
Predictive loss	$\mathcal{L}^{pred}$	0.40
Sustainability learning loss	$\mathcal{L}^{sus}$	0.25
Policy violation loss	$\mathcal{L}^{pol}$	0.15
Privacy exposure loss	$\mathcal{L}^{priv}$	0.10
Computational cost loss	$\mathcal{L}^{cost}$	0.10

The four weights (derived from table 5) are used to indicate the relative importance of the concerns: predictive learning is assigned to the highest weight as it is central to the problem, and sustainability learning has a high, but still

meaningful, weight of 0.25 so that it can be considered an optimization target, whose accomplishment is not left to chance. Thus, the composite local objective turns out to be 0.247 for this training state. Any update that is to be sent to the orchestration server is first clipped based on the bound established above (Equation 6) prior to exiting the institution. These mechanisms allow the individual university nodes to continually improve their learning performance as they remain flexible in meeting sustainability goals, policy restrictions, privacy regulations, inclusion policies and computational constraints.

Results

The execution of the FedSustain-PEO model has been done in a heterogeneous Federated testing environment of higher education having 12 institutional nodes that have been classified into three groups of high-resource, medium-resource, and resource-constrained environments. The multi-institutional data set used in this study has consisted of 30,000 learners' data in the non-IID setting through Dirichlet partitioning technique. The feature set includes LMS log interaction, sustainability competency, engagement, and institutional digital inclusion & resource capacity.

For experimental parameter initialization, the system was configured with 100 global communication rounds (T), a local batch size (B) of 32, and an initial learning rate (η_0) of 0.001 with a lower bound of 0.0001. Local epochs (E_k) dynamically ranged from 1 to 5, while the client participation ratio (q_t) was set between 0.40 and 0.80. Specific policy compliance thresholds included $\tau_p = 0.80$ for compliance, $\tau_i = 0.60$ for inclusion, $\tau_r = 0.20$ for maximum privacy risk, and $\tau_s = 0.65$ for the minimum sustainability learning gain.

For evaluating the proposed FedSustain-PEO, a heterogeneous condition was established with different higher education datasets. The performance evaluation was done across in terms of policy responsiveness, sustainability-aware federated operation and predictive performance. The comparative predictive performance as reported in table 6 shows FedSustain-PEO to perform comparatively better to FedAvg and FedProx models. Precision and F1-score reached upto 94.1%, and 93.6% reveals the model to be compatible for different heterogeneous institutions, and besides the models is capable to different institutions. Here, the final training loss was much minimal to range of 0.147, and this leads to better data handling strategy.

Table 6: Comparative Predictive Performance of Federated Learning Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC (%)	Final Loss	Convergence Round
FedAvg	88.7	87.9	87.4	87.6	90.1	0.286	82
FedProx	91.2	90.4	89.6	89.9	92.3	0.219	71
FedSustain-PEO	94.8	94.1	93.6	93.8	96.2	0.147	58

Table 7 indicates the policy-oriented performance of the proposed FedSustain-PEO, and the results were evaluated in terms of Policy Execution Compliance Index (PECI), Federated Institutional Equity Score (FIES) and various other policy-aware measures. Here, Peci of the proposed model was reached upto 0.947, and the FIES score was reached upto 0.91, and this indicates the proposed model supports various policies of the institutions, and thus can be used for handling different datasets from various organizations.

Table 7: Comparative Policy Compliance and Institutional Equity Results

Model	PECI	FIES	PAI	PCRR (%)	Inclusion Score	Privacy Score	Violation Rate (%)
FedAvg	0.712	0.741	0.684	61.8	0.728	0.812	18.6
FedProx	0.768	0.793	0.746	69.7	0.781	0.846	14.2
FedSustain-PEO	0.947	0.918	0.926	93.4	0.921	0.958	4.8

The proposed FedSustain-PEO has been compared with the existing models through sustainability and resource efficiency measures. Results from table 8 show the performance in terms of sustainability learning gain, which has been reached by 0.842 by the proposed model, and also opting for this scheme has resulted in power gain, where the overall power consumption has reduced upto 26.5% than existing FedAvg model. Besides better power consumption reduction, the proposed FedSustain-PEO has reduced communication overhead by 1.92 GB. Also, for training 12 local datasets, the training time was reduced up to 36.8 seconds. The model does not require any resource-constrained nodes, as the entire setup was used for complex data handling.

Table 8: Sustainability Learning and Computational Efficiency Results

Model	SLG	Energy (kWh)	CO ₂ e (kg)	Communication (GB)	Training Time (s)	Energy Reduction (%)	Constrained-Node Participation (%)
FedAvg	0.681	56.6	24.9	2.84	48.7	0.0	54.3
FedProx	0.724	50.2	21.8	2.47	43.1	11.3	62.7
FedSustain-PEO	0.842	41.6	17.3	1.92	36.8	26.5	78.6




Fig 2: Federated Learning Accuracy Progression Over Communication Rounds

Fig 2 illustrates the test accuracy percentage across 100 communication rounds for three federated learning algorithms. The solid green line represents FedSustain-PEO, achieving convergence by round 58. The dashed blue line depicts FedProx, reaching convergence at round 71. Finally, the dotted red line shows FedAvg, which converges more slowly at round 82. Overall, the visualization demonstrates the comparative performance efficiency and accuracy growth rates among the evaluated training frameworks.

Limitations

Several limitations have to be taken into account when considering the applicability of the proposed FedSustain-PEO framework and when interpreting it in light of the diverse nature of higher education environments. First is that the experimental environment consisted of 12 institutional nodes having well-defined computational, policy and 'digital inclusion' attributes and thus the behaviours observed may vary in larger national or international university networks. Second, in the process of making qualitative institutional policies machine-executable through thresholds, the policies must be correctly interpreted, and highly opinionated policy statements might not be well captured in terms of numbers. Third, the policy thresholds, such as the privacy-risk threshold of 0.20, the compliance threshold of 0.80, and the sustainability-gain threshold of 0.65 that were used in the models need further testing on a number of other institutional situations. Fourth, the statistical differences between the non-IID partitions are preserved; however, real-world differences in curriculum structure, cultural context, behaviour of learners or governance in the institutions are not completely captured. Fifth, the computational costs and carbon estimates can be different depending on the hardware architecture, the power sources and the network used, and the regional carbon intensity. Lastly, when adapting based on policy, there is more orchestration overhead as policy compliance and conflict states as well as institutional priorities have to be assessed on an ongoing basis. These boundaries suggest that these tools do need to be widened across institutions and longitudinally tested before they can be taken operationally to a large scale.

Conclusion

This study aimed to design FedSustain-PEO for assessing digital learning systems for sustainability education in a heterogeneous higher education setting. The framework took care of the shortcomings of the traditional federated approach, with the incorporation of sustainability learning analytics, machine-executable policy translation, dynamic policy conflict modelling, adaptive federation control and sustainability-aware local training in one platform. The proposed architecture allowed institutions to directly impact federated training decisions based on their policies around privacy, digital inclusion, equity, resource efficiency and sustainability learning without any raw transfers of learner data records. Experimental assessment results showed that FedSustain-PEO presented an accuracy of 94.8%, an F1-score equal to 93.8%, and an increase in the FIES result from 0.915 to 0.918, and also a PEC-I equal to 0.947. The framework also had a Sustainability Learning Gain of 0.842; energy use was decreased by 41.6 kWh, and 78.6% of the institutional nodes with constrained resources were involved. Our results show the feasibility of policy-aware orchestration for achieving a mix of predictive performance with privacy, equity and computational sustainability. FedSustain-PEO overall creates a scalable basis to shift from traditional policy assessment to a continuous, adaptive and computationally enforceable governance of higher education federation(s) for achieving a world of sustainability-oriented learning. Future research should expand the framework from 12 nodes to larger international university networks to test behavior under broader institutional and cultural variations.

References

- Yurdem, B., Kuzlu, M., Gullu, M. K., Catak, F. O., & Tabassum, M. (2024). Federated learning: Overview, strategies, applications, tools and future directions. *Heliyon*, 10(19). <https://doi.org/10.1016/j.heliyon.2024.e38137>
- Garzón, J., Patiño, E., & Marulanda, C. (2025). Systematic review of artificial intelligence in education: Trends, benefits, and challenges. *Multimodal Technologies and Interaction*, 9(8), 84. <https://doi.org/10.3390/mti9080084>
- Cañavate, J., Martínez-Marroquín, E., & Colom, X. (2025). Engineering a sustainable future through the integration of generative AI in engineering education. *Sustainability*, 17(7), 3201. <https://doi.org/10.3390/su17073201>
- Gutiérrez, R., Villegas-Ch, W., & Luján-Mora, S. (2025). Implementing federated learning for privacy-preserving emotion detection in educational environments. *Frontiers in Artificial Intelligence*, 8, 1644844. <https://doi.org/10.3389/frai.2025.1644844>
- Sum, A. S. I., Pritee, Z. T., Saha, A. K., Mahmud, M. A., & Nur, K. (2026). A systematic review on privacy preservation in federated learning: AS Ibte Sum et al. *International Journal of Information Security*, 25(2), 65. <https://doi.org/10.1007/s10207-026-01229-x>
- Li, Y., Ouyang, T., Chen, X., & Cao, X. (2024, June). Fedcarbon: Carbon-efficient federated learning with double flexible controls for green edge ai. In *2024 IEEE/ACM 32nd International Symposium on Quality of Service (IWQoS)* (pp. 1-10). IEEE. <https://doi.org/10.1109/IWQoS61813.2024.10682880>
- Deng, D., Zhang, T., Gu, C., Xiang, C., & Wu, X. (2025, July). Less is more: Enabling efficient and fair federated learning by knowledge trimming. In *2025 IEEE/ACM 33rd International Symposium on Quality of Service (IWQoS)* (pp. 1-10). IEEE. <https://doi.org/10.1109/IWQoS65803.2025.11143456>
- Kim, M., Saad, W., Mozaffari, M., & Debbah, M. (2023). Green, quantized federated learning over wireless networks: An energy-efficient design. *IEEE transactions on wireless communications*, 23(2), 1386-1402. <https://doi.org/10.1109/TWC.2023.3289177>
- Fachola, C., Tornarí, A., Bermolen, P., Capdehourat, G., Etcheverry, L., & Fariello, M. I. (2023). Federated learning for data analytics in education. *Data*, 8(2), 43. <https://doi.org/10.3390/data8020043>
- Blanco-Justicia, A., Domingo-Ferrer, J., Martínez, S., Sánchez, D., Flanagan, A., & Tan, K. E. (2021). Achieving security and privacy in federated learning systems: Survey, research challenges and future directions. *Engineering Applications of Artificial Intelligence*, 106, 104468. <https://doi.org/10.1016/j.engappai.2021.104468>
- Chen, H., Zhu, T., Zhang, T., Zhou, W., & Yu, P. S. (2023). Privacy and fairness in federated learning: On the perspective of tradeoff. *ACM Computing Surveys*, 56(2), 1-37. <https://doi.org/10.1145/3606017>
- Mukhtiar, N., Mahmood, A., & Sheng, Q. Z. (2025). Fairness in federated learning: Trends, challenges, and opportunities. *Advanced Intelligent Systems*, 7(6), 2400836. <https://doi.org/10.1002/aisy.202400836>
- Kuzmanic, D., Valenzuela, J. P., Claro, S., Canales, A., Cerda, D., & Undurraga, E. A. (2023). Socioeconomic disparities in the reopening of schools during the pandemic in Chile. *International Journal of Educational Development*, 100, 102805. <https://doi.org/10.1016/j.ijedudev.2023.102805>

- Mokski, E., Leal Filho, W., Sehnem, S., & Andrade Guerra, J. B. S. O. D. (2023). Education for sustainable development in higher education institutions: an approach for effective interdisciplinarity. *International Journal of Sustainability in Higher Education*, 24(1), 96-117. <https://doi.org/10.1108/IJSHE-07-2021-0306>
- Dang, X. T., Vu, B. M., Nguyen, Q. S., Tran, T. T. M., Eom, J. S., & Shin, O. S. (2024). A survey on energy-efficient design for federated learning over wireless networks. *Energies*, 17(24), 6485. <https://doi.org/10.3390/en17246485>
- Rakhmanovich, I. U., Bhattacharjee, R., Anute, N., & Ugli, G. A. N. (2025, June). Sustainable economics: The role of green technology in shaping future economic policies. In *AIP Conference Proceedings* (Vol. 3306, No. 1, p. 060042). AIP Publishing LLC. <https://doi.org/10.1063/5.0276471>
- Dong, B. (2026). Sustainability in Higher Teaching: Digital Transformation of the Fourth Industrial Revolution and Its Impact on Open Knowledge. *European Journal of Education*, 61(1), e70468. <https://doi.org/10.1111/ejed.70468>
- Sing, R. D. R., Rahman, S. Y. A., Kumar, P., Arasu, R., Beigi, M. A., & Pramila, R. (2025). Harnessing AI for enhanced student learning and performance in higher education: A multinational collaboration. *Indian Journal of Information Sources and Services*, 15(2), 208-217. <https://doi.org/10.51983/ijiss-2025.IJISS.15.2.28>
- Lee, G. K. S. (2025). Embedding sustainability in higher education: A review of institutional strategy, curriculum reform, and digital integration. *International Journal of Science and Research Archive*, 15(2), 420-426. <https://dx.doi.org/10.30574/ijrsra.2025.15.2.1396>
- Miller, T., Durlík, I., Kostecka, E., & Puszkarek, A. (2025). Federated learning for environmental monitoring: A review of applications, challenges, and future directions. *Applied Sciences*, 15(23), 12685. <https://doi.org/10.3390/app152312685>