

PAPER

Federated Swarm Intelligence for Privacy-Preserving Mobile Collaborative Learning: An Energy-Efficient Client-Selection Framework

Arar Al Tawil^{1,2}, Siti Hazyanti
 Mohd Hashim¹ ,
 Laiali Almazaydeh³

¹Universiti Sains Malaysia,
 Penang, Malaysia

²Applied Science Private
 University, Amman, Jordan

³Abu Dhabi University, Abu
 Dhabi, United Arab Emirates

sitihazanti@usm.my

ABSTRACT

The increase in mobile learning escalates the friction between data-led personalization and safeguarding sensitive, institutionally fragmented student data. Centralized analytics necessitates the aggregation of raw records, which conflicts with privacy regulation and mobile resource limits. A novel federated swarm-intelligence framework for mobile collaborative learning to train a student-outcome prediction model using simulated devices without any raw data transmitted from them. A three-tier structure where cloud, edge and mobile interfaces work together is proposed, along with a particle swarm optimization (PSO) controller for energy-aware and latency-aware client selection. Further, we deploy differential privacy via DP-SGD with a formal accountant and employ membership-inference audit. Federated training obtained in all simulations accuracy (AUC up to 0.91) equal to that of centralized while a PSO selection strategy achieved 40–46% reduction in communication rounds and 73% less energy use compromising accuracy only marginally. Differential privacy provides a certifiable guarantee ($\epsilon \approx 1-8$) at low cost.

KEYWORDS

mobile collaborative learning, federated learning, particle swarm optimization (PSO), differential privacy, client selection, edge computing, learning analytics

1 INTRODUCTION

Learning has now mostly transitioned to mobile devices. They relate what learners are engaged with in the classroom to what they subsequently do, leaving behind a steady stream of clicks, submitted assignments, minutes on a page [1], [2]. The trail can help. Personalization can be driven through this, Students who are slipping can be flagged, and material can be adapted to what a learner actually needs. The challenge lies in the sensitivity of the data, its dispersion across various

Tawil, A. A., Hashim, S. H. M., Almazaydeh, L. (2026). Federated Swarm Intelligence for Privacy-Preserving Mobile Collaborative Learning: An Energy-Efficient Client-Selection Framework. *International Journal of Interactive Mobile Technologies (IJIM)*, 20(18), pp. 158–170. <https://doi.org/10.3991/ijim.v20i18.62774>

Article submitted 2026-06-28. Revision uploaded 2026-08-02. Final acceptance 2026-08-05.

© 2026 by the authors of this article. Published under CC-BY.

institutions, and the fact that it resides on phones with low battery life and unreliable connections. The majority of learning analytics systems achieve this by collecting all the data to a central server before a model is trained. Making that choice may violate data-protection law, give attackers a compelling target [3], and disregard the circumstances under which the data was originally generated. Federated learning takes the opposite route instead of gathering the data, it transmits the model to the location of data: every device trains at its location and only shares its updates, not the raw records [4]. In education, where confidentiality matters and institutional control matters, it seems like a natural fit, and yet still rarely used in practice [5]. Two barriers are blocking the path. There is cost first. In each round of training, energy and bandwidth are spent. That consumes time in sending out requests and the actual task of training. Therefore, it is a waste of resources to engage every device. Some are slow, some low on battery, some barely connected etc. The second one is more subtle: just because something is shared does not mean it's private. A description of updates can still leak knowledge of the underlying data behind the updates [6]. So privacy has to be measured and proved, not assumed. We deal with both problems simultaneously. As it operates on a cloud, edge, and mobile stack, this framework considers which device should be included as something worth optimizing. In every round, a controller that utilizes particle swarm optimization version 2 (PSO-V2) searches for a set of clients which provide the most model up-date at the least cost. Updates produced during the training are differentially private with formal guarantees and membership-inference attack is used to ascertain how much can still be inferred afterward. We do not say that it is more accurate than the alternatives. Our aim is to achieve the same level of accuracy as that of the centralized training while reducing the communication cost significantly and providing a privacy guarantee.

The paper consists of four main contributions. Client selection in federated mobile learning can be modeled as a multi-objective particle swarm optimization (PSO) problem. The aim is to achieve a balance between predictive usefulness, energy, and latency. The membership inference attack combines the DP-SGD and a formal accountant with an empirical audit, so privacy is both guaranteed and measured. The framework is evaluated using two public education datasets, three federated algorithms, four selection strategies, and averaging over several seeds along with a privacy-utility analysis, a heterogeneity analysis, and ablation studies. It is clear where the gains come from: the efficiency that swarm-based selection delivers is clearly defined, as is the formal guarantee of differential privacy.

The following section outlines the organization of the paper: In this section, we are going to reviews related work. Section 3 outlines the structure, types of federated architecture and PSO controller, privacy mechanisms, datasets, and experimental setup. The results are found in Section 4, followed by Section 5, which discusses the study's limitations, and lastly, Section 6 concludes.

2 RELATED WORK

In the domain of educational technology, swarm intelligence has a rich history of utilization, with PSO being the most preferred technique among researchers. Using this approach, many different optimization problems have been solved. Examples include the construction of learning paths which respect both a knowledge graph and the profile of the learner [7], [8], balancing competing teaching goals when

recommending resources [9], and only recently, making performance predictions interpretable by combining optimization with XAI [10]. According to surveys of this work, the techniques work but almost always centrally over data gathered in one place and paying little attention to the devices that produced it [11]. Our work is closest to a recent study. This study combines swarm intelligence with mobile interaction in a cloud-edge-mobile design. Its authors candidly pointed out privacy and low-end device support issues remain unsolved [12].

Federated learning has not been developed in a vacuum, but speaks directly to these gaps. The initial proposal was to train a shared model across many devices without gathering data [4] and later methods such as FedProx cope with harder cases where clients (and their data) differ a lot [13]. Considerable research has been done on moving federated learning to the network edge, with recent surveys detailing its promise and challenges, which include communicating overhead, uneven devices, and coordination [14], [15]. A different line of work leverages biology. They use distributed swarm learning to boost the robustness and communication efficiency of edge training for the Internet of Things [16], [17]. The education plays a vital role in the technology. There are edge architectures which help to minimize latency and also help to keep data local [1]. There are edge-AI systems which help to determine quality of the teaching over the wireless network [18]. The absence of any plan to optimize the client participation against the real energy and latency budgets holds a lot of importance.

Both literatures address privacy issues, but in education, where data frequently relates to minors and high-stakes outcomes, privacy is particularly sensitive. The literature on privacy-preserving federated learning surveys the attacks and defenses that matter for decentralized training [3], while differential privacy provides a rigorous standard [19] that DP-SGD advances as a practical workhorse [20]. Membership-inference attacks provide a concrete way of testing whether a model has memorized its training data [6]. According to a report, in the field of education, some studies show federated learning achieving the same accuracy as centralized learning on tasks such as dropout prediction but being better able to resist attack [21], [5], and others construct cross-institutional models which never pool student records [22]. In contrast to mobile collaborative learning, evaluation on a single axis is also conducted [2]. Our paper shows that this can be integrated: we optimize who participates, we ensure privacy formally rather than by assumption, and we audit the leakage that remains.

3 MATERIALS AND METHODS

3.1 System architecture

The framework adopts a three-tier cloud-edge-mobile architecture, illustrated in Figure 1. The cloud tier hosts the global model, the PSO client-selection controller, and the differential-privacy accountant; it never receives raw learner data. The edge tier comprises regional relay nodes that aggregate and forward model updates, reducing the communication distance between mobile clients and the cloud. The mobile tier consists of heterogeneous client devices, phones and tablets of differing compute throughput and bandwidth, each holding a private partition of student data and performing local training. Only differentially private model updates flow upward; selection decisions and the global model flow downward.




Fig. 1. Three-tier cloud–edge–mobile architecture

Notes: Raw data remains on devices; only DP-protected updates are shared, and PSO selects an efficient client subset each round.

3.2 Federated learning framework

Let K clients be assigned the learners' population, each storing a private dataset. A shared classifier predicts whether a student will withdraw from a course based on various features. The training is conducted in round synchronously. In every round, a group of clients is chosen as participants, the most recent global model is sent to each of the selected clients, each selected client performs several local epochs of stochastic gradient descent on its own data, the resulting updates are aggregated into a new global model. The framework includes support for three aggregation algorithms: federated averaging (FedAvg) [4], FedProx with a proximal term for heterogeneous clients [13], and FedAvg with server momentum (FedAvgM). A centralized model that uses pooled data provides an upper-bound reference, while federated training with full participation provides a no-selection reference.

In choosing the classifier, a small multilayer perceptron with two hidden layers and dropout regularization was used for its low computational cost on mobile devices. Data are divided in a non-independent and non-identically-distributed (non-IID) way to reflect institutional heterogeneity through a Dirichlet partition that concentrates more with a parameter to control the level of heterogeneity.

3.3 PSO-based client selection

Every round, multi-objective optimization is solved by PSO formulated for client participation. Every particle contains subset of clients which is a candidate. The expected learning utility of the objective rewards, estimated from the volume of data received from clients and the most recent loss, is parameterized by the penalization of simulated cost (i.e., combination of per-client energy and latency). The cost model for a candidate subset S estimates energy as a function of local computation (epochs, samples) and communication volume. Its latency is equal to compute time + (communication volume or bandwidth)/client, where the per-client compute throughput and bandwidth are drawn from a fixed, seeded heterogeneous distribution to model strong and weak devices. The configurable trade-off weight between utility and cost balances inefficiency against one. By varying this weight, the efficiency frontier that can actually be achieved can be traced. The swarm iterates with standard inertia, cognitive, and social coefficients and returns the best subset for the round. In comparison, PSO considers three baselines: random selection, power-of-choice selection [23], which always prefers high-loss clients, and full participation.

3.4 Differential privacy and leakage audit

Certain clients get trained locally with DP-SGD [20] which clips per-sample gradients and adds calibrated Gaussian noise to give official guarantee. We calculate the privacy budget ϵ using a formal privacy accountant instead of a hand derivation, making our reported guarantee certifiable. We sweep the budget over several values to characterize the privacy–utility trade-off. A black-box membership-inference attack [6] is utilized to empirically audit residual leakage. Specifically, the attack is trained on the prediction-confidence features to distinguish members from non-members, with the success rate averaged over multiple data splits. An attack success level that is approximately 0.5 results in behavior that is close to random guessing and, as such, inflicts not much leakage of information.

3.5 Datasets

Dataset used for the study was public and educational. The open university learning analytics dataset (OULAD) contains information on 32,593 students across seven modules. It contains demographic details, counts of interaction with the virtual-learning-environment, counts of assessment submissions and final outcomes [24]. The prediction target is binary withdrawal, with a base withdrawal rate of 31.2%. The UCI Student Performance dataset is made up of records for secondary-school students in two subjects, with demographic, social and school-related attributes [25]. It is a smaller, harder benchmark that tests generalization. For both datasets, the features were cleaned, encoded, and scaled; a held-out test partition was reserved; and the training population was split across clients using either a non-IID partition by course module or by a Dirichlet allocation.

3.6 Experimental protocol and implementation

Every setting was assessed using five random seeds, and all reported metrics are the mean and standard deviation over seeds. The accuracy, F1 score, and area under the curve measured predictive quality (AUC). The total energy (joules), communication latency (seconds), and communication volume (megabytes) of system cost together with total rounds to convergence was measured. Membership-inference attack success together with the certified ϵ characterized privacy. Statistical comparisons across methods used the Friedman omnibus test pairwise Wilcoxon signed-rank comparisons against the strongest baseline were also computed with the caveat that five paired seeds cannot attain a two-sided p-value below 0.0625 so the Friedman test is treated as the primary inferential tool. The Python implementation of the framework makes use of, firstly, PyTorch for modeling. Secondly, a federated simulation engine for orchestration. Thirdly, a library for the DP-SGD and its accountant. Fourthly, the standard metaheuristics library for PSO. Finally, the tool for adversarial-robustness for the membership-inference attack.

4 RESULTS

4.1 Predictive performance and federated–centralized parity

The main comparison on OULAD across all algorithm-selection combinations is presented in Table 1. The performance of the federated configurations matched

centralized upper bound as the best federated AUC of 0.910 slightly exceeded the reference centralized one at 0.898, while the accuracies were between 0.82 and 0.84 for the strategy selection. The difference in accuracy across selection strategies was low compared to the standard deviations. The Friedman test across methods was significant ($p = 0.015$). The weaker performance of the momentum variant was the major contributor. Further pairwise Wilcoxon comparisons against the strongest baseline were not significant. Hence, methods are comparable in predictive quality.

Table 1. Main comparison on OULAD (mean over five seeds)

Method	Acc	F1	AUC	Energy (J)	Latency (s)	Attack
Centralized	0.818	0.733	0.898	0.0	0.00	0.538
FedAvg/random	0.827	0.753	0.910	227.8	3.09	0.493
FedAvg/PSO	0.828	0.749	0.910	209.2	1.85	0.490
FedAvg/power-of-choice	0.828	0.709	0.907	226.9	2.94	0.500
FedAvg/full	0.839	0.745	0.908	764.7	4.96	0.501
FedProx/random	0.834	0.759	0.910	227.8	3.09	0.499
FedProx/PSO	0.827	0.747	0.910	209.2	1.85	0.489
FedProx/power-of-choice	0.830	0.719	0.907	226.9	3.06	0.501
FedProx/full	0.841	0.750	0.909	764.7	4.96	0.506
FedAvgM/random	0.819	0.698	0.900	227.8	3.09	0.497
FedAvgM/PSO	0.820	0.733	0.901	209.2	1.85	0.497
FedAvgM/power-of-choice	0.817	0.720	0.894	225.1	2.93	0.493
FedAvgM/full	0.816	0.678	0.890	764.7	4.96	0.504

Table 2 reports the same comparison on the smaller UCI dataset. Absolute performance was lower and more variable, as expected for a harder, smaller benchmark, and full participation tended to yield the highest accuracy. Crucially, the efficiency trends observed on OULAD persisted: PSO selection again achieved the lowest latency among selection strategies at accuracy comparable to random selection, supporting the generalization of the efficiency benefit even where absolute accuracy is modest.

Table 2. Main comparison on UCI Student Performance (mean over five seeds)

Method	Acc	F1	AUC	Energy (J)	Latency (s)	Attack
Centralized	0.733	0.373	0.676	0.0	0.00	0.575
FedAvg/random	0.726	0.442	0.696	204.0	2.27	0.489
FedAvg/PSO	0.717	0.401	0.677	201.2	1.23	0.493
FedAvg/power-of-choice	0.612	0.424	0.700	200.0	1.91	0.476
FedAvg/full	0.763	0.363	0.671	678.3	3.26	0.503
FedProx/random	0.727	0.443	0.696	204.0	2.27	0.495
FedProx/PSO	0.718	0.405	0.677	201.2	1.23	0.496
FedProx/full	0.765	0.373	0.671	678.3	3.26	0.490
FedAvgM/random	0.733	0.345	0.650	204.0	2.27	0.488
FedAvgM/PSO	0.711	0.420	0.666	201.2	1.23	0.509
FedAvgM/full	0.757	0.384	0.661	678.3	3.26	0.514

4.2 Efficiency of swarm-based client selection

The main advantage of PSO selection is its efficiency. On OULAD with FedAvg, PSO reduced communication latency from 3.09 s (random) to 1.85 s, a 40% reduction, while altering accuracy by less than 0.001. When we compare it to full participation, PSO saves 72.6% energy (764.7 J to 209.2 J), latency by 62.6% and communication volume (66.8 MB to 20.0 MB) by 70%. According to Figure 2, various strategies are summarized. In UCI too, random selection gave significantly improved latency even by 46% by PSO. According to these outcomes, PSO is classified on the accuracy-energy frontier illustrated on Figure 3 in the favorable region, with significant savings in energy and latency for negligible loss in accuracy. Due to PSO's speed in leveraging a larger network, any latency advantage will outweigh any raw energy advantage (which is the more relevant property for responsive mobile classrooms).




Fig. 2. Communication latency (a) and energy (b) by client-selection strategy on OULAD

Note: PSO attains the lowest latency and energy among practical strategies.




Fig. 3. Accuracy-energy trade-off across all algorithm-selection combinations

Note: PSO selection attains near-best accuracy at a fraction of the energy required by full participation.

4.3 Privacy–utility trade-off

Table 3 along with Figure 4 illustrates the impact of differential privacy on OULAD. The accuracy level was slightly decreased compared to not using the reformulation method. In fact, with no coupling ($\epsilon = \infty$), the accuracy of the model was 0.828 and the AUC was 0.910. When using DP ($\epsilon = 1, 4, 8$), the accuracy level went down to approximately 0.74. Coupling had a more beneficial impact on the AUC which only went down (without coupling) to about 0.87–0.88 on successive steps. In all situations membership-inference success was close to 0.5 indicating limited empirically measurable leakage even before DP. This is because the dense model trained on low-dimensional tabular data does not very strongly memorize records. In this world, differential privacy is probably best understood as providing a certifiable formal guarantee, not as yielding any large empirical reduction in the success of such attacks.

Table 3. Privacy–utility trade-off on OULAD (DP-SGD, mean over seeds)

Privacy Budget ϵ	Accuracy	AUC	Attack Success
1.0	0.746	0.873	0.501
4.0	0.741	0.878	0.498
8.0	0.741	0.879	0.501
∞ (no DP)	0.828	0.910	0.490




Fig. 4. Privacy–utility trade-off as the budget ϵ varies

Note: Stronger privacy (smaller ϵ) lowers accuracy while attack success stays near random guessing.

4.4 Robustness to data heterogeneity

According to Table 4, we see robustness to non-IID severity controlled by Dirichlet concentration α (smaller α means more heterogeneous clients). The accuracy gradually declined from 0.832 at α equal to 1.0 to 0.786 at the very heterogeneous α equal to 0.1 with AUC staying above 0.89 throughout. The framework launches

without collapsing under the realistic heterogeneity of cross-institutional educational data, a property necessary for deployment across schools with different populations. The trend is shown in Figure 5.

Table 4. Robustness to non-IID heterogeneity on OULAD (Dirichlet α)

Dirichlet α	Accuracy	AUC
0.1 (most heterogeneous)	0.786	0.898
0.5	0.827	0.910
1.0 (least heterogeneous)	0.832	0.912




Fig. 5. Accuracy and AUC versus heterogeneity

Note: Performance degrades gracefully as client distributions become more non-IID.

4.5 Ablation studies

The framework's behavior is clarified with two ablations. In the assumption of the PSO cost-weight (refer to Table 5), we confirmed that this is the control that leads the method to the efficient frontier: the weight was raised from 0.0 to 0.8 and it reduced the energy from 244.8 J to 190.7 J – a 22% saving, and the accuracy drops only from 0.837 to 0.832. When we increased the number of clients, we could see the scaling we expected. The energy increased with the size of the federation. From 10 clients we had 136.7 J to 501.4 J at 50 clients. The accuracy started at 0.815 and increased to only 0.838 which is only a small increase. This shows how we have a large training population in terms of effective training. The cost-weight ablation is visualized in Figure 6.

Table 5. Ablation of the PSO cost-weight on OULAD

Cost-Weight	Accuracy	AUC	Energy (J)
0.0 (utility only)	0.837	0.912	244.8
0.4	0.836	0.913	223.8
0.8 (efficiency-leaning)	0.832	0.904	190.7




Fig. 6. Effect of the PSO cost-weight: higher weights trade a small amount of accuracy for lower energy

5 DISCUSSION

The findings communicate a consistent and, we believe, honest story. To begin with, federated training is a legitimate substitute for centralized analytics in this environment. Through the use of two datasets and three aggregation methods, The total privacy model achieved the same predictiveness as a centralized model without mixing the raw data generating the prediction policy. This response to the primary privacy problem with data-driven education leaves institutions in control of their own records. Moreover, the swarm-intelligence element finds its application not in enhancing accuracy but in improving performance. The selection of PSO decreased latency and energy considerably while achieving the same level of accuracy as the other strategies. In a mobile classroom where battery life and response time determine the adoption of the tool, a 40 to 46 percent reduction in latency at no cost to accuracy is a significant gain. As we state directly, PSO does not enhance accuracy; the selection mechanisms are statistically tied on that metric, and the contribution is efficiency, not prediction. The privacy findings require a careful approach. Most of the reduction in membership-inference success stems from aggregation rather than differential privacy. When a model is averaged across many clients, it hardly memorizes any single record. Therefore, even without DP, the attack is already near chance. Differential privacy bestows on systems a formal, provable guarantee on top of the empirical robustness and at only a modest cost to accuracy. To say that MPPI and MPCC produced a huge drop in attack success is to misread our own data. The actual usefulness of federated learning is that the certificate average cannot provide and this is precisely the point for educational use. The mobile tier is simulated: the energies and the latencies are supplied by a documented cost model with seeded device variation rather than by measurements on real handsets, so the absolute figures are best read as comparisons rather than ground truth. In the leakage audit, we are depending on attack of single type, and a stronger adversary might bring to surface leakage we did not. The accuracy of UCI is modest, meaning the compact model and feature set may be better suited to richer datasets such as OULAD. Having only five seeds limits statistical power. The Friedman test can inform us, but having further seeds will sharpen the pairwise comparisons. These all point to an obvious

next step: measuring on physical phones, testing against a wider array of attacks, modeling clickstreams with sequential architectures, and scaling the federation.

6 CONCLUSION

This study proposes a federated swarm-intelligence framework for privacy-preserving mobile collaborative learning. This is developed based on the cloud, edge and mobile architecture with PSO-based client selection, differential privacy with formal accountant and membership inference audit. By conducting experimentation on two different public education datasets with a wide range of algorithms and selection strategies, while at the same time averaging over multiple seeds, it was found that the researchers system managed to reach centralized-level accuracy without ever centralizing data. It also managed to cut latency by around 40 to 46 percent, and cut energy by as much as 73 percent through swarm-based selection at the same accuracy. Last but not the least, the researchers system also delivers a certifiable privacy guarantee. The framework provides a transparent and deployable path to energy-aware, privacy-preserving learning analytics on heterogeneous mobile devices, while keeping the efficiency benefit of selection distinct from the formal guarantee of differential privacy. Subsequent efforts will involve examining the cost model on actual hardware, expanding the privacy examination process, and also extending the method to more complex models. These include, studying how students acquire knowledge over time.

7 REFERENCES

- [1] K. Mohiuddin *et al.*, “Mobile learning evolution and emerging computing paradigms: An edge-based cloud architecture for reduced latencies and quick response time,” *Array*, vol. 16, p. 100259, 2022. <https://doi.org/10.1016/j.array.2022.100259>
- [2] A. Peramunugamage, U. W. Usoof, and J. A. K. S. Hapuarachchi, “Systematic review on mobile collaborative learning for engineering education,” *J. Comput. Educ.*, vol. 9, pp. 535–567, 2022.
- [3] S. Saha, T. Haque, and M. Rahman, “A multifaceted survey on privacy preservation of federated learning: Progress, challenges, and opportunities,” *Artif. Intell. Rev.*, vol. 57, p. 184, 2024. <https://doi.org/10.1007/s10462-024-10766-7>
- [4] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, 2017, pp. 1273–1282.
- [5] C. Fachola, A. Tornaria, P. Bermolen, G. Capdehourat, L. Etcheverry, and M. I. Fariello, “Federated learning for data analytics in education,” *Data*, vol. 8, no. 2, p. 43, 2023. <https://doi.org/10.3390/data8020043>
- [6] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *Proc. IEEE Symp. Secur. Privacy (S&P)*, 2017, pp. 3–18. <https://doi.org/10.1109/SP.2017.41>
- [7] M. Chen and A. Tserenchimed, “Design of the English personalized learning path recommendation algorithm based on MABPSO,” in *Proc. 5th Asia-Pacific Conf. Commun. Technol. Comput. Sci. (ACCTCS)*, 2025, pp. 1–6. <https://doi.org/10.1109/ACCTCS66275.2025.00183>
- [8] C. Liu, J. Wang, and Y. Zhang, “Design of a learning path recommendation system based on a knowledge graph,” *Int. J. Inf. Commun. Technol. Educ.*, vol. 19, no. 1, pp. 1–18, 2023. <https://doi.org/10.4018/IJICTE.319962>

- [9] S. Zhong, "Research on multi-objective optimization recommendation algorithms for work-study platforms," *J. Comput. Signal Syst. Res.*, vol. 2, no. 1, pp. 1–12, 2025. <https://doi.org/10.71222/bwj2jw48>
- [10] R. Parkavi, P. Karthikeyan, and A. S. Abdullah, "Enhancing personalized learning with explainable AI: A chaotic particle swarm optimization based decision support system," *Appl. Soft Comput.*, vol. 167, p. 112255, 2024. <https://doi.org/10.1016/j.asoc.2024.111451>
- [11] A. G. Gad, "Particle swarm optimization algorithm and its applications: A systematic review," *Arch. Comput. Methods Eng.*, vol. 29, no. 5, pp. 2531–2561, 2022. <https://doi.org/10.1007/s11831-021-09694-4>
- [12] J. Sun and Y. Mei, "A swarm intelligence framework for mobile collaborative learning," *Int. J. Interact. Mobile Technol.*, vol. 20, no. 12, pp. 4–19, 2026. <https://doi.org/10.3991/ijim.v20i12.62260>
- [13] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. Mach. Learn. Syst. (MLSys)*, vol. 2, 2020, pp. 429–450.
- [14] Q. Duan, J. Huang, S. Hu, R. Deng, Z. Lu, and S. Yu, "Combining federated learning and edge computing toward ubiquitous intelligence in 6G network: Challenges, recent advances, and future directions," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 4, pp. 2892–2950, 2023. <https://doi.org/10.1109/COMST.2023.3316615>
- [15] S. Zhan, Y. Wang, and L. Chen, "A review on federated learning architectures for privacy-preserving AI: Lightweight and secure cloud-edge-end collaboration," *Electronics*, vol. 14, no. 3, p. 512, 2025. <https://doi.org/10.3390/electronics14132512>
- [16] Y. Wang, Y. Xu, Q. Shi, and T.-H. Chang, "Distributed swarm learning for edge Internet of Things," *IEEE Commun. Mag.*, vol. 62, no. 7, pp. 110–116, 2024. <https://doi.org/10.1109/MCOM.003.2300610>
- [17] X. Fan, Y. Wang, Y. Huo, and Z. Tian, "Efficient distributed swarm learning for edge computing," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2023, pp. 3627–3632. <https://doi.org/10.1109/ICC45041.2023.10279508>
- [18] F. Li, Y. Wang, and X. Liu, "Artificial intelligence and edge computing for teaching quality evaluation based on 5G-enabled wireless communication technology," *J. Cloud Comput.*, vol. 12, p. 75, 2023. <https://doi.org/10.1186/s13677-023-00418-6>
- [19] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014. <https://doi.org/10.1561/04000000042>
- [20] M. Abadi et al., "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2016, pp. 308–318. <https://doi.org/10.1145/2976749.2978318>
- [21] M. Khalil, M. Slade, and S. Prinsloo, "Towards privacy-preserving data-driven education: The potential of federated learning," in *Proc. Int. Conf. New Trends Comput. Sci. (ICTCS)*, 2025, pp. 1–8. <https://doi.org/10.1109/ICTCS65341.2025.10989403>
- [22] Z. Huang, L. Wang, and Y. Chen, "Federated learning for AI-assisted education: Privacy-preserving and cross-platform collaborative modeling in higher education," in *Proc. 2nd Int. Conf. Mach. Intell. Digital Appl.*, 2025, pp. 1–9. <https://doi.org/10.1145/3744464.3744487>
- [23] Y. J. Cho, J. Wang, and G. Joshi, "Towards understanding biased client selection in federated learning," in *Proc. 25th Int. Conf. Artif. Intell. Statist. (AISTATS)*, PMLR, vol. 151, 2022, pp. 10351–10375.
- [24] J. Kuzilek, M. Hlosta, and Z. Zdrahal, "Open university learning analytics dataset," *Sci. Data*, vol. 4, no. 1, p. 170171, 2017. <https://doi.org/10.1038/sdata.2017.171>
- [25] P. Cortez and A. Silva, "Using data mining to predict secondary school student performance," in *Proc. 5th Future Business Technol. Conf. (FUBUTEC)*, 2008, pp. 5–12.

8 AUTHORS

Arar Al Tawil received the B.Sc. in Computer Science from Al-Hussein Bin Talal University, Jordan, in 2018, and the M.Sc. from the University of Jordan in 2021. He is currently pursuing a Ph.D. at the School of Computer Sciences, Universiti Sains Malaysia (USM), and serves as a Lecturer and Developer at the Computer Sciences Department, Faculty of Information Technology, Applied Science Private University (ASU), Amman, Jordan. His research interests include virtual reality, reinforcement learning, optimization algorithms, machine learning, deep learning, and NLP. He has authored several papers in IEEE, Springer, and IGI Global venues, with projects spanning VR-based rehabilitation and adaptive gamified learning systems.

Siti Hazyanti Mohd Hashim obtained her B.Sc., M.Sc., and Ph.D. in Computer Science from Universiti Teknologi MARA (UiTM), Malaysia. She is currently a Lecturer at the School of Computer Sciences, Universiti Sains Malaysia (USM), Pulau Pinang, Malaysia. Her research interests include computer technology, multimedia, game and mobile applications, virtual reality, and software engineering. She actively contributes to academic supervision and collaborative research in immersive technologies and educational innovation (E-mail: sitihazanti@usm.my).

Laiali Almazaydeh received her doctorate degree in computer science and engineering from University of Bridgeport in USA in 2013, specializing in human computer interaction. She is currently assistant professor in College of Engineering, Abu Dhabi University in the UAE. Laiali has published more than ninety research papers in various international journals and conferences proceedings, her research interests include human computer interaction and pattern recognition. She received best paper awards in three conferences, ASEE 2012, ASEE 2013 and ICUMT 2016. Recently she has been awarded two postdoc scholarships from European Union Commission and Jordanian-American Fulbright Commission (E-mail: laiali.almazaydeh@adu.ac.ae).