CIT. Journal of Computing and Information Technology, Vol. 34, No. 3, September 2026, i–ii
doi: 10.20532/cit.2026.1007021

Editorial

The September 2026 (Vol. 34, No. 3) issue of *CIT. Journal of Computing and Information Technology* brings three papers from the areas of learner modeling, smart grids, and speech recognition.

The first paper in this issue, titled *Mahalanobis-enhanced K-Means++ Clustering for Multidimensional Learner Personas: Evidence from the Canvas MOOC Dataset*, deals with the topic of artificial intelligence applications in education. Namely, the authors of the paper, Yafeng Liu, Rui Yan, Yong Xiao, Ailing Chen, and Di Zhang introduce a multidimensional learner persona modeling framework to support personalized instruction in smart learning environments, using behavioral, cognitive, and motivation-related indicators derived from the Canvas open dataset. It compares K-Means++ with Euclidean and Mahalanobis distances, as well as density-based and probabilistic clustering methods. The results show that Mahalanobis-enhanced K-Means++ achieves better clustering performance than the Euclidean variant, with a higher Silhouette score (0.4277 vs. 0.3281) and improved Davies–Bouldin and Calinski–Harabasz indices. Based on the clustering results, the authors identify four learner personas and propose differentiated instructional interventions tailored to each group. The feasibility of these interventions is assessed through expert review, while their effectiveness in real classroom settings remains to be empirically validated.

In the second paper, titled *Local Differentially Private and Robust Edge–Fog Aggregation for Smart Grid Data*, the authors Song-Yao Feng, Zheng-Yan Huang, Zhou Ai, and Jun-Hao Song address the challenge of preserving user privacy in high-frequency smart-grid data aggregation while maintaining communication efficiency and robustness against anomalous or malicious reports. It proposes an edge–fog aggregation method that combines local differential privacy at the user or edge side with robust, anomaly-aware fusion at the fog layer, using clipping, random projection, Laplacian noise, coordinate-wise medians, and adaptive weighting. The method was evaluated on the London SmartMeter Energy Consumption dataset and the Pecan Street Dataport, achieving MAPE values of 4.89% and 5.02% for regional load-related tasks. Compared with representative privacy-preserving baselines, the proposed approach substantially reduced communication overhead and demonstrated particularly strong resilience to data-tampering attacks: the aggregation error increased by only 8.92%, compared with 12.46% for a related method (DP-EDA). An important contribution is the explicit formulation of the privacy–utility–robustness trade-off, particularly the observation that stronger differential-privacy noise may also make legitimate reports appear anomalous during robust aggregation. The authors note that broader validation on larger user populations, stronger formal privacy analysis, and more realistic adversarial scenarios are still needed before the method can be considered ready for large-scale deployment.

In the last paper of the issue, titled *Optimization and Development of Edge-Embedded Speech Recognition Module of Improved DFCNN-CTC*, the authors Fang Nie and Lijuan Zheng propose a lightweight edge-embedded speech recognition model for smart toys, designed to address limited computing resources, environmental noise, and the acoustic variability of children's speech. The model improves a Deep Fully Convolutional Neural Network-Connectionist Temporal Classification (DFCNN-CTC) architecture by using original spectrograms, a Transformer-based language model, and a mutual-learning knowledge-distillation framework with dual student networks. Experiments on two Chinese children's speech datasets showed that the proposed model achieved a word error rate of 8.2%, outperforming Conformer, Branchformer, and E-Branchformer models. Model com-

pression reduced the size from 125 MB to 8.1 MB while lowering inference latency to 155 ms and improving energy efficiency by 4.4 times, which is particularly important for real-time on-device speech recognition on resource-constrained smart toys with limited memory, processing capacity, and power availability. In real-world home tests, the system achieved a 96.5% wake-up success rate, 93.2% command recognition accuracy, and a false wake-up rate of 0.8 times per hour. The results demonstrate the model's suitability for accurate, low-power, real-time speech recognition on resource-constrained smart toys, although performance still decreases in noisy Chinese-English mixed-language scenarios.

Alan Jović
Editor-in-Chief