

Edge-AI Machine Learning System for Intelligent Customer Face Monitoring Using Raspberry Pi

Tae-Gyu LEE*

*Pyeongtaek University, 3825 Seodong-daero, Pyeongtaek-si, Gyeonggi-Do 17869, Korea

Corresponding Author: tglee@ptu.ac.kr, Fax: +82 31 659 8011, Tel: +82 31 659 8370

Abstract—This study presents a novel embedded artificial intelligence (Edge-AI) framework designed for resource-constrained environments, enabling intelligent customer interaction and real-time content adaptation through personalized digital signage. Leveraging the computational capabilities of the Raspberry Pi and lightweight convolutional neural networks (TinyML), the system performs real-time demographic inference (e.g., gender, age group) entirely on-device. By integrating a dynamic QR-code-based feedback mechanism, the platform facilitates participatory user engagement and refinement of content delivery strategies. Through comprehensive evaluations—including latency benchmarks, classification accuracy, user response time, and QR feedback participation—the proposed system demonstrates feasibility for scalable deployment in retail environments, public facilities, educational campuses, and healthcare kiosks. The study advances the design principles of privacy-preserving, responsive, and autonomous digital signage platforms for next-generation human-AI interfaces.

Keywords— Edge-AI, TinyML, Lightweight, Raspberry Pi, Face detect

I. INTRODUCTION

The increasing demand for responsive, autonomous systems in public and semi-public spaces has accelerated interest in Edge-AI architectures that minimize reliance on remote infrastructure. Traditional cloud-dependent smart kiosks, while powerful, suffer from latency bottlenecks, data privacy concerns, and high operational costs [5,16]. In contrast, Edge-AI platforms perform inference directly on embedded devices, enabling decentralized intelligence and real-time decision-making [18].

Raspberry Pi, owing to its affordability and versatility, provides an ideal hardware base for embedded vision systems [10,14]. However, constrained computational resources necessitate the use of highly optimized AI models and real-time data processing pipelines. This research addresses this challenge by presenting a self-contained Edge-AI system capable of performing end-to-end user monitoring, demographic inference, and dynamic content rendering on a single Raspberry Pi board.

The system architecture integrates real-time vision processing, a content recommendation engine, and a feedback interface via QR codes. Without reliance on cloud servers or external storage. By enabling two-way human-computer interaction through participatory feedback, the platform not

only adapts content based on the user profile but also continuously evolves based on user response data.

Moreover, in the context of campus environments, the need for context-aware navigation, safety communication, and event information has become more pronounced. A kiosk system capable of understanding user types (e.g., new students, faculty, visitors) and delivering tailored guidance not only improves user experience but also enhances operational efficiency in university administration. For example, a student at the engineering building may require directions to the registrar, while a parent visitor might need parking and event details. An intelligent system must accommodate such diversity in real time.




Figure 1. Edge-AI based Display System for Customer Monitoring using Raspberry Pi

In this study, we also explore mathematical models to quantify and optimize the latency-performance trade-off inherent in on-device AI. Through a latency-aware architecture and optimized inference pipeline, the system delivers sub-second performance while maintaining high classification accuracy. A mathematical formulation is introduced to model total latency as a function of model size, and experimental evaluations support the theoretical predictions.

Lastly, we integrate a QR-based feedback mechanism, which allows users to refine content interactively and enables the system to adapt over time as illustrated in Fig. 1. This user-centric loop transforms the system from a static

recommendation engine to a self-improving interface that learns from its environment.

The contributions of this paper are as follows:

- Design of a fully embedded low-latency inference system for customer demographic classification using Raspberry Pi and lightweight CNN models.
- Development of an adaptive display logic integrated with QR-based user feedback, forming a closed-loop recommendation cycle.
- Performance evaluation under realistic operating conditions using latency, accuracy, and power metrics, including detailed latency-by-phase analysis.
- Demonstration of field-deployable prototype in a campus navigation use case, with potential for scalability and cross-domain adaptation.
- Proposal and validation of a mathematical model linking model complexity, accuracy, and latency, enabling optimization of on-device AI pipelines.

II. RELATED WORKS

A. TinyML and Embedded Inference

The advancement of Tiny Machine Learning (TinyML) has catalysed a new generation of intelligent devices capable of executing machine learning inference on resource-constrained hardware. TinyML techniques have made it feasible to deploy neural networks on microcontrollers and embedded systems, particularly in scenarios where latency, power efficiency, and privacy are critical. Lane et al. [1] introduced DeepX, a runtime system enabling energy-efficient inference on mobile devices through dynamic resource-aware optimization. This work demonstrated how inference pipelines could be partitioned and executed with fine-grained control, optimizing both energy and latency.

TensorFlow Lite Micro, a specialized extension of TensorFlow Lite, has facilitated the deployment of sub-megabyte deep learning models on ARM Cortex-M processors and System-on-Chips (SoCs) [12-14]. These developments have enabled always-on sensing applications such as wake-word detection, gesture recognition, and simple vision tasks without requiring cloud connectivity [14]. In the context of this research, TinyML allows real-time face and demographic inference to occur entirely on embedded hardware like Raspberry Pi, enabling personalized digital signage in a decentralized and privacy-preserving manner [5,6].

B. Lightweight Neural Network Models

To operate under strict memory and computational constraints, lightweight neural network architectures have been proposed and widely adopted. The MobileNet family (MobileNetV1, V2, and V3) has emerged as a leading solution for efficient image classification and feature extraction on edge devices [3,4]. These architectures leverage depth wise separable convolutions to reduce computation and memory usage while maintaining high levels of accuracy. MobileNetV2 introduced inverted residual blocks and linear bottlenecks, further optimizing the balance between latency and representational power.

Other architectures, such as BlazeFace, were specifically designed for fast and accurate face detection on mobile and embedded hardware. BlazeFace achieves sub-200ms inference times on low-end mobile CPUs and DSPs, making it well-suited for real-time facial attribute analysis in embedded signage systems [17]. Pruning, quantization, and knowledge distillation are complementary techniques used to compress large models, further reducing latency and power consumption while maintaining acceptable accuracy [10].

In this project, MobileNetV2 is used for demographic classification (e.g., age, gender), while BlazeFace serves as the face detection frontend. Both models are deployed using TensorFlow Lite with post-training quantization to achieve sub-300ms end-to-end response times on the Raspberry Pi.

C. Edge-Based Digital Signage and Feedback Systems




Figure 2. Edge-Based Digital Signage and Feedback Systems

As shown in Fig. 2, digital signage has evolved from static content displays to intelligent, context-aware systems capable of adapting content based on environmental inputs and user profiles. In recent years, research and industry applications have focused on leveraging sensors and AI to personalize displayed information in retail, transportation, and educational settings. Kim et al. [6] demonstrated an AI-augmented signage system in public transit environments, where content was dynamically selected based on preconfigured data streams. However, this implementation relied on centralized cloud inference and lacked direct user feedback integration.

Edge-based signage platforms represent a shift toward decentralization. Hybrid edge-cloud models have been proposed to reduce latency and enhance scalability, such as those outlined in Li et al. [5,16]. While these systems offload certain computations to edge devices, they often depend on server-side components for training, feedback processing, or orchestration. In contrast, the system proposed in this study is fully independent of cloud infrastructure.

Our approach uniquely integrates on-device inference with real-time feedback through QR code interaction, forming a closed-loop system that continuously adapts its behavior based on user responses. This enables interactive, privacy-conscious

personalization that is both scalable and responsive and aligns with current trends in user-centric edge computing [8][16].

III. SYSTEM ARCHITECTURE

This section presents an in-depth overview of the hardware and software architecture of the proposed Edge-AI digital signage system, which is tailored for intelligent customer interaction in public or semi-public environments. The design emphasizes modularity, low power consumption, privacy preservation, and on-device processing capabilities. Each component, both hardware and software, is carefully selected and integrated to support real-time, personalized display operations using embedded AI models.

A. Hardware Configuration

The hardware architecture is composed of a suite of compact, cost-effective, and energy-efficient components that together enable seamless data capture, inference, interaction, and output. The system is intended for deployment in campus kiosks, public venues, and commercial spaces.

- **Raspberry Pi 4B (8GB RAM):** Serves as the central processing unit [10]. It executes the full AI pipeline, from image capture and model inference to content rendering and QR code generation. The 8GB variant is chosen to accommodate concurrent processes without memory bottlenecks.
- **Raspberry Pi Camera Module V2:** Equipped with an 8MP Sony IMX219 sensor, it captures high-resolution facial images for demographic inference. It supports a frame rate of up to 60fps at lower resolutions, ensuring responsiveness.
- **7" HDMI Touchscreen Display:** A capacitive touch interface providing visual output and interactive capabilities. The display is optimized for daylight readability and multi-point touch, offering users dynamic maps, alerts, and event information.
- **BLE (Bluetooth Low Energy) Module:** Detects user proximity through signal strength and triggers the system from standby mode to active capture mode. This not only conserves energy but also optimizes resource usage by avoiding unnecessary processing.
- **Uninterruptible Power Supply (UPS):** Maintains operation during power interruptions and allows safe shutdowns to protect system integrity. It is essential for high-reliability installations in field environments.
- **Enclosure and Mounting Frame:** Though not electronic, the weatherproof and tamper-resistant physical casing is essential for public deployments. It includes ventilation for passive cooling and support for VESA mounting.

B. Software Stack

The software components are modular, lightweight, and selected to maximize compatibility and performance on ARM-based embedded systems. The stack supports hardware acceleration, multi-threaded operations, and asynchronous event handling.

- **Operating System:** Raspbian Lite (64-bit)? stripped down to exclude unnecessary desktop elements, this distribution offers improved boot times, reduced disk usage, and optimal control over background processes.
- **Inference Framework:** TensorFlow Lite Micro? tailored for or embedded inference, this framework runs quantized CNN models with reduced computational overhead [12-14]. It is integrated with C/C++ bindings for real-time pipeline execution.
- **Pre-trained AI Models:**
 - **BlazeFace:** A lightweight and fast face detection model optimized for mobile CPUs [17]. Its architecture enables detection in under 20ms per frame on the Raspberry Pi.
 - **MobileNetV2:** Employed for age and gender classification [3,4]. The model is fine-tuned on datasets such as UTKFace and Audience, then converted into TFLite format with post-training quantization for faster inference.
- **QR Feedback Server:** A Flask-based microservice running locally on the Pi. It generates QR codes embedded with unique session tokens, hosts feedback forms via REST endpoints, and logs user interaction anonymously.
- **UI & Display Renderer:** Uses a combination of OpenCV for low-level image operations and Kivy or PyQt5 for the GUI layer. Content rendering is hardware-accelerated via OpenGL ES 2.0, ensuring fluid visual updates.
- **Content Management Module:** An internal logic engine that maps inferred demographic attributes to context-aware display content (e.g., event announcements for students, navigation for visitors).

C. System Pipeline

The digital signage system in Fig. 3 operates through a tightly orchestrated multi-stage pipeline that performs sensing, interpretation, response, and learning in real-time:




Figure 3. Digital Signage System Pipeline with Feedback Integration

- 1) **Proximity Triggering:** BLE beaconing determines the presence and distance of a user. Upon reaching a pre-set RSSI threshold, the system exits idle mode and begins video acquisition.

2) **Face Detection:** The camera stream is piped through BlazeFace, which returns bounding boxes of detected faces. Only cropped face regions proceed to the classification stage, reducing computation.

3) **Demographic Inference:** Each face is resized and normalized before being input into MobileNetV2. The model outputs predicted probabilities across predefined age brackets (e.g., 10s, 20s, etc.) and gender.

4) **Content Mapping:** A rules-based or probabilistic recommender system maps the demographic output to content. For instance, college-age prediction triggers student-focused content such as class schedules.

5) **Display Output:** The selected content is rendered to the screen with minimal delay. Dynamic overlays (e.g., weather, alerts, news tickers) are also incorporated based on kiosk location and time of day.

6) **Feedback Loop:** A unique QR code is generated and shown alongside the content. When scanned, it directs users to a form or action portal where they can rate the content or request specific information.

7) **Session Logging:** Metadata such as timestamps, inferred attributes, and feedback interactions are stored locally in a secure, anonymized format. This data informs future content updates and can be used for federated learning.

This pipeline structure ensures asynchronous, low-latency operation while enabling adaptive learning over time. Each component is optimized for embedded execution, enabling deployment in bandwidth-limited, offline, or privacy-sensitive environments. The architecture is scalable across multiple kiosks and can incorporate federated updates for model refinement, supporting long-term performance without centralized data storage.

IV. EXPERIMENTAL ANALYSIS

A. Test Conditions

To evaluate the performance and real-time applicability of the proposed lightweight Edge-AI display system, we designed a comprehensive experimental testbed. The setup simulates a realistic indoor retail or kiosk environment, applicable to smart campus navigation, information booths, and commercial settings.

- **Environment:** A controlled indoor booth with consistent ambient lighting to ensure uniform camera exposure and minimize inference variability due to shadows or brightness.
- **Participants:** A sample group of 30 volunteers was recruited (15 male, 15 female), spanning ages 10 to 70. The participant diversity enabled assessment of the demographic classifier across a wide spectrum of real-world users [9]
- **Devices:** The hardware configuration included a Raspberry Pi 4B (4 GB RAM), Raspberry Pi Camera V2 Module, and a 7-inch HDMI touch display.
- **Dataset:** The model was trained and validated on a custom dataset comprising live camera captures annotated via

cross-verification with OpenCV's DNN-based gender/age estimators. This helped simulate real-time inference under natural variability while anchoring accuracy with pre-validated ground truth.

B. Evaluation Metrics

Performance metrics were selected to assess the latency, inference accuracy, power efficiency, and user engagement under operational conditions. These are defined in Table 1 below.

TABLE 1. METRICS FOR SYSTEM PERFORMANCE EVALUATION.

Metric	Description
Inference Latency (ms)	Time elapsed from image acquisition to final classification output
Demographic Accuracy (%)	Ratio of correct gender/age predictions compared to labelled ground truth
Power Consumption (W)	Measured active power draw using inline USB power meter during inference
Display Response Time (s)	Time from model inference completion to full content update on display
QR Engagement Rate (%)	Percentage of users who voluntarily scanned and responded via QR feedback

C. Baseline Comparison

To contextualize the performance of our Edge-AI system, a baseline comparison was conducted using a cloud-based inference pipeline leveraging Google Vision API. In this setup, images were captured locally and then transmitted via REST API to the cloud service. Responses were parsed and displayed on the same kiosk interface.

- **Latency:** Frame-level timestamps were used to compute the end-to-end latency for both Edge-AI and cloud pipelines.
- **Power:** Power consumption on the Raspberry Pi was monitored using a USB inline meter across 30 inference cycles, comparing idle, inference, and display update phases.
- **Stability:** The thermal load and CPU utilization were tracked to ensure sustainable real-time operation without throttling.

The Edge-AI configuration outperformed cloud-based alternatives in terms of latency (142ms vs. >800ms) and power efficiency, with minimal compromise in inference accuracy [5,16]

D. Mathematical Modeling

To provide a theoretical framework for understanding system performance, we introduce a latency-accuracy trade-off model dependent on model parameter size P [2,10,15].

1) Latency Model

The total processing time T_{total} is composed of four stages:

$$T_{total}(P) = T_{cap} + T_{pre} + T_{inf}(P) + T_{disp}$$

Where:

- T_{cap} : Image capture time (typically ~50ms)
- T_{pre} : Preprocessing time (resize, normalization ~20ms)
- $T_{inf}(P)$: Inference time, linear in model size

- T_{disp} : Display response time (~50ms)

Assuming $T_{inf}(P) = \alpha P$, where α is a hardware-dependent constant, the total latency increases linearly with model complexity [2][10][15].

2) Accuracy Model

Model accuracy $A(P)$ is assumed to be a saturating exponential function of parameter size:

$$A(P) = A_{\infty} (1 - e^{-kP})$$

Where:

- A_{∞} : Maximum achievable accuracy
- k : Saturation rate constant

This reflects the empirical observation that larger models improve accuracy, but with diminishing returns beyond a threshold [2,11].

3) Optimization Objective

We define an objective to minimize latency while maintaining acceptable accuracy:

Alternatively, a composite cost function can be formulated:

$$\min_P T_{total}(P) \text{ subject to } A(P) \geq A_{min}$$

$$C(P) = T_{total}(P) + \lambda \cdot (A_{max} - A(P))$$

Where λ balances inference delay and classification performance.

V. RESULTS AND EVALUATION

The Edge-AI kiosk system was evaluated against a cloud-based counterpart across several performance metrics, with particular focus on responsiveness, accuracy, power consumption, and user engagement. Table 2 highlights the essential performance comparison between the two systems:

TABLE 2. COMPARISON BETWEEN EMBEDDED AND CLOUD-ASSISTED SYSTEMS.

Metric	Edge-AI System	Cloud-based System
Inference Latency	142ms	820ms (avg + network)
Accuracy	86.4%	89.2%
Power Usage	4.3W	N/A (external)
Display Response	0.42 s	1.26 s
QR Feedback Rate	31.8%	Not available

The Edge-AI system demonstrated significantly lower latency and near-equivalent accuracy compared to the cloud-based solution. The self-contained nature of the system contributed to reduced display response time, and its moderate power consumption indicates suitability for battery-powered or mobile deployments. The relatively high QR engagement rate (31.8%) indicates strong user interaction and validates the system's accessibility.

A. Model Size vs. Latency and Accuracy

To further analyse the inference performance, we evaluated several versions of the MobileNet architecture, both in quantized and non-quantized formats, with varying model sizes. The results validate the previously introduced latency? Accuracy trade-off model. Specifically, inference latency increases approximately linearly with model size, while

classification accuracy exhibits a saturating growth curve. These trends are consistent with theoretical expectations and reinforce the importance of selecting an appropriately sized model for edge deployment.

1) Latency vs. Model Size by Phase

Fig. 4 shows that the stacked bar chart depicts the latency of each processing phase—image capture, preprocessing, inference, and display rendering—across different model sizes (in megabytes). The results clearly indicate that inference latency dominates the overall processing time and increases substantially with larger model sizes [15]. In contrast, the latencies for capture, preprocessing, and display remain relatively constant, confirming that inference is the primary bottleneck in end-to-end system responsiveness.




Figure 4. Latency Breakdown by Model Size Across Processing Phases

2) Accuracy vs. Model Size

This plot illustrates a **saturating exponential relationship** between model size and accuracy. While accuracy improves with larger models, the rate of improvement diminishes beyond approximately 8-10MB. This observation suggests the existence of an **optimal trade-off point**, where additional model complexity yields minimal accuracy gains relative to the increased inference cost. Selecting a model within this range can ensure real-time responsiveness while maintaining acceptable predictive performance.




Figure 5. Accuracy Improvement with Increasing Model Size

Fig. 5 shows that as the model size increases, accuracy follows an exponential saturating curve, gradually improving and then plateauing. Around the 8-10 MB range, the performance gains begin to level off, indicating an optimal trade-off point.

B. QR Feedback Response Time

A user study involving 30 participants revealed notable variations in response times and engagement rates with the QR code-based feedback interface [6][9]. The average scan time across all participants was 2.8 seconds. The fastest response recorded was 1.2 seconds, primarily among tech-savvy university students, while the slowest response, at 7.6 seconds, was observed among older participants or visitors unfamiliar with QR scanning.

Engagement rates also varied by user category. Students exhibited the highest participation, with an engagement rate of 42.1%, followed by faculty at 28.7%, and visitors at 22.5%. These results are illustrated in Fig. 6, which compares QR feedback engagement by user type.




Figure 6. QR Code Feedback Engagement by User Category

This data indicates that demographic factors influence both response time and participation in QR code interactions. Younger and more digitally literate users were not only faster to engage but also more likely to participate in voluntary feedback tasks. In contrast, older or less tech-oriented users demonstrated hesitation or unfamiliarity, resulting in slower response times and lower engagement.

The findings underscore the need to incorporate intuitive visual cues, simplified interaction workflows, and multi-modal feedback options (e.g., touch input, voice prompts) to ensure inclusive design. Future iterations of the kiosk interface should prioritize user-centric adaptations that accommodate varying levels of digital literacy, thus enhancing overall system accessibility and feedback effectiveness.

VI. CONCLUSION AND FUTURE WORK

This study has introduced a novel Edge-AI display system, fully embedded on a Raspberry Pi platform, optimized for real-time demographic inference and dynamic content personalization. By leveraging lightweight convolutional neural networks (CNNs) and on-device inference, the system demonstrates that it is possible to achieve meaningful interactivity and privacy-preserving personalization without reliance on cloud infrastructure or continuous internet connectivity.

A key innovation lies in the integration of a QR-code-based feedback loop, which enhances user engagement while enabling adaptive system behaviour. Experimental results validate the feasibility of deploying TinyML models on resource-constrained hardware, achieving latency and accuracy levels suitable for real-world applications such as

smart kiosks. The system's modularity further supports scalability across diverse domains including retail, public transit, and higher education settings.

Looking forward, several research directions are proposed to enhance system functionality and user experience. First, incorporating facial emotion recognition would allow for contextual adjustment of content tone or urgency based on users' affective states? An approach particularly valuable in emotionally charged scenarios such as healthcare or crisis response.

Second, gesture-based interaction mechanisms may be introduced to support touchless control of kiosk interfaces, improving accessibility and hygiene, particularly in post-pandemic or high-traffic environments.

To overcome the limitations of isolated model learning, the system may adopt federated learning across distributed kiosk networks [8,11]. This approach enables collaborative model training using localized, anonymized data while preserving user privacy, thereby improving model robustness without centralized data aggregation.

Another promising extension is the application of reinforcement learning to optimize content delivery strategies based on real-time user feedback [7]. This would enable the system to dynamically adjust its decision-making policies in response to evolving user behaviour patterns and environmental contexts.

Lastly, longitudinal field deployments in a variety of operational settings will be conducted to assess system performance, usability, and user satisfaction over time. These trials will offer essential insights into the system's robustness, adaptability, and sustainability, ultimately guiding further iterations of design.

In conclusion, the convergence of embedded AI, lightweight inference, and user-driven feedback mechanisms points to a new paradigm in intelligent digital signage? One that is not only interactive and responsive, but also ethically aware and contextually intelligent. The proposed system represents an important foundational step toward that future.

REFERENCES

- [1] Lane, N. D., et al. "DeepX: A software accelerator for low-power deep learning inference on mobile devices." Proceedings of the 15th International Conference on Information Processing in Sensor Networks (IPSN), 2016.
- [2] Sze, V., Chen, Y. H., Yang, T. J., and Emer, J. S. "Efficient processing of deep neural networks: A tutorial and survey." Proceedings of the IEEE, vol. 105, no. 12, pp. 2295-2329, 2017.
- [3] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L. C. "MobileNetV2: Inverted residuals and linear bottlenecks." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
- [4] Howard, A. G., et al. "Searching for MobileNetV3." Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2019.
- [5] Li, Y., Zhang, Y., Li, Y., and Wang, X. "Edge AI: On-demand accelerating deep neural network inference via edge computing." IEEE Transactions on Wireless Communications, vol. 19, no. 1, pp. 447-457, 2020.
- [6] Kim, H., Lee, K., and Kwon, H. "Implementation of AI-based digital signage in public transportation." ICT Express, vol. 7, no. 4, pp. 407-412, 2021.
- [7] Ren, S., He, K., Girshick, R., and Sun, J. "Faster R-CNN: Towards real-time object detection with region proposal networks." IEEE

- Transactions on Pattern Analysis and Machine Intelligence (TPAMI), vol. 39, no. 6, pp. 1137-1149, 2017.
- [8] Zhu, J., Liu, Y., Zhou, Y., and Zhang, H. "Federated learning on non-IID data silos: A survey." *ACM Computing Surveys (CSUR)*, vol. 54, no. 12, pp. 1-36, 2021.
- [9] Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., and Torralba, A. "Learning deep features for discriminative localization." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [10] Han, S., Mao, H., and Dally, W. J. "Deep Compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding." *International Conference on Learning Representations (ICLR)*, 2016.
- [11] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., and Arcas, B. A. Y. "Communication-efficient learning of deep networks from decentralized data." *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.
- [12] Lane, N. D., Bhattacharya, S., Georgiev, P., Forlivesi, C., and Kawsar, F. "Squeezing deep learning into mobile and embedded devices." *IEEE Pervasive Computing*, vol. 16, no. 3, pp. 82-88, 2018.
- [13] Raschka, S., and Mirjalili, V. "Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow 2." *Packt Publishing*, 2020.
- [14] Warden, P., and Situnayake, D. "TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers." *O'Reilly Media*, 2019.
- [15] Reddi, V. J., et al. "MLPerf Inference Benchmark." *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2020.
- [16] Jiang, Y., et al. "User-centric edge computing: Vision, challenges and future directions." *IEEE Network*, vol. 32, no. 4, pp. 136-143, 2018.
- [17] Roy, A., et al. "Efficient On-Device Machine Learning: Challenges and Opportunities." *IEEE Micro*, vol. 41, no. 4, pp. 73-81, 2021.
- [18] Akbari, M., and Sarrafzadeh, M. "Hierarchical Deep Learning Models for Efficient Inference on Embedded Devices." *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 6, pp. 1183-1195, 2021.



Tae-Gyu Lee (BSc'92–MSc'96–PhD'06) He received the B.Sc. degree from Kunsan National University, Kunsan, Korea in 1992, the M.Sc. degree from Soongsil University, Seoul, Korea in 1996, and the Ph.D. degree from Korea University in 2006. He is currently a Professor in the ICT convergence Division, and the Dept. of AI Convergence in Graduate School of AI Convergence, Pyeongtaek University, Gyeonggi, Korea. He has been a Professional Researcher in Advanced Convergent Technology R&D Group, Korea Institute of Industrial Technology (KITECH), Ansan, Korea for 2009-2013. He has also been a President in the JIGUNET Corporation, Seoul, Korea, from 1999. His research interests are in distributed systems, ubiquitous computing, middleware, networks, wearable and robot computing.