Jurnal Teknik Informatika C.I.T Medicom, 16 (4) (2024) 232-241

Published by: Institute of Computer Science (IOCS)

Jurnal Teknik Informatika C.I.T Medicom

Journal homepage: www.medikom.iocspublisher.org

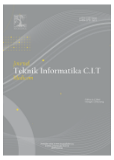

Meta-Learning Algorithms for Resource-Constrained Intelligent IoT Devices

Fristi Riandari[1], Jonhariono Sihotang[2]
[1] Manajemen Informatika, Politeknik Negeri Medan, Indonesia
[2] Sistem Informasi, Universitas Putra Abadi Langkat, Indonesia

Article Info

Article history

Received : Juli 21, 2024
Revised : Aug 25, 2024
Accepted : Sep 30, 2024

Key Words:

Meta-Learning;
TinyML;
Resource-Constrained IoT Devices;
Few-Shot Adaptation;
\Edge Intelligence.

Abstract

The rapid expansion of the Internet of Things (IoT) requires devices that can operate intelligently in dynamic environments despite severe hardware and energy constraints. Traditional machine learning models deployed on microcontroller-class IoT devices often struggle to adapt to new tasks, handle sensor noise, and maintain accuracy under changing environmental conditions. This research proposes a lightweight meta-learning framework specifically optimized for resource-constrained IoT platforms, combining gradient-based meta-learning techniques with model compression strategies such as quantization and pruning. The objective is to enable rapid few-shot adaptation, reduce computational overhead, and ensure robust performance in real-world IoT deployments. The study adopts a hardware-aware design approach, implementing the proposed model on ultra-low-power microcontrollers such as ARM Cortex-M series and ESP32. A two-phase training pipeline meta-training and on-device fine-tuning is used to evaluate adaptation speed, latency, memory footprint, accuracy, and energy consumption. Experimental results demonstrate that the lightweight meta-learning model adapts to new sensor-based tasks significantly faster than conventional supervised learning models while consuming substantially less energy. The model also shows improved resilience to environmental variations and sensor noise, outperforming baseline TinyML and standard meta-learning architectures under constrained conditions. Despite these promising results, the research identifies limitations related to computational cost, memory usage during adaptation, and the trade-off between model complexity and predictive accuracy. Nonetheless, the findings highlight the potential of meta-learning as a transformative approach for building intelligent, adaptive, and energy-efficient IoT systems. This study contributes to the advancement of TinyML and edge intelligence by providing a practical and scalable meta-learning solution tailored for ultra-low-power IoT devices.

Corresponding Author:

Fristi Riandari,
Manajemen Informatika,
Politeknik Negeri Medan, Indonesia
Jl. Almamater No.1, Padang Bulan, Kec. Medan Baru, Kota Medan, Sumatera Utara 20155
Email: fristiriandari@polmed.ac.id

1. Introduction

The rapid development of the Internet of Things (IoT) has transformed various sectors, including healthcare, agriculture, manufacturing, transportation, and smart cities. Billions of interconnected devices now collect, transmit, and process data in real time, enabling more efficient operations and intelligent decision-making. However, as IoT systems expand in scale and complexity, the demand for on-device intelligence continues to grow[1]. Traditional IoT architectures rely heavily on cloud computing for data processing and model inference, but this dependency poses several challenges, such as increased latency, high bandwidth usage, data privacy concerns, and reduced reliability when network connectivity is unstable. As a result, there is a pressing need to shift toward edge intelligence, where IoT devices themselves are capable of performing machine learning (ML) tasks locally.

Despite this need, enabling advanced ML capabilities directly on IoT devices remains a major challenge. Most IoT hardware is designed with strict resource constraints, including limited computational power, small memory capacity, restricted energy availability, and compact storage[2]. Conventional machine learning and deep learning models are computationally expensive and often too large to be deployed on microcontroller-class devices, which typically have only a few hundred kilobytes of RAM. Furthermore, IoT applications often operate in dynamic environments where data patterns change rapidly due to varying conditions, sensor drift, or user behavior. Retraining models in the cloud is not always feasible due to energy consumption, latency, and the need for continuous data transmission. This creates a fundamental gap between the sophisticated capabilities of modern ML algorithms and the operational limitations of IoT hardware.

In response to these challenges, researchers have turned to meta-learning, a machine learning approach that focuses on “learning how to learn.” Meta-learning algorithms are designed to adapt quickly to new tasks using only a small amount of data and with significantly reduced computation. They enable IoT devices to rapidly update and fine-tune their models without costly retraining, making them highly suitable for environments where conditions change frequently[3]. Unlike traditional ML models that require extensive training, meta-learning models leverage prior experience to generalize efficiently, allowing them to perform fast, on-device adaptation even under severe hardware constraints.

The integration of meta-learning into resource-constrained IoT devices has the potential to overcome some of the most persistent limitations of edge computing. By enabling devices to adapt to novel conditions and learn new patterns with minimal computational overhead, meta-learning improves system robustness, energy efficiency, and responsiveness. Applications become more intelligent, personalized, and capable of autonomous operation without constant cloud communication. This shift is particularly impactful in domains such as wearable healthcare monitoring, smart agriculture, industrial IoT, and environmental sensing, where real-time adaptation is crucial for accurate and reliable decision-making.

Research on meta-learning has grown rapidly in the past decade, focusing on enabling models to learn new tasks with minimal data and computation. One of the foundational works in this area is Model-Agnostic Meta-Learning (MAML) introduced by Finn, Abbeel, and Levine (2017). MAML proposes a general optimization-based framework where models are trained to find an initialization that can be fine-tuned quickly on new tasks with only a few gradient steps. This technique demonstrated strong performance in few-shot image classification and reinforcement learning, but its reliance on second-order gradients makes it computationally demanding posing challenges for deployment on resource-constrained IoT devices.

To address some of MAML’s computational challenges, Nichol, Achiam, and Schulman (2018) proposed Reptile, a first-order meta-learning algorithm that approximates MAML while reducing gradient complexity. Reptile iteratively updates model parameters using gradients from randomly sampled tasks, eliminating the need for second-order derivatives. Its efficiency and simplicity have led to broader adoption in low-resource machine learning environments, positioning it as a more feasible candidate for IoT adaptation compared to full MAML.

Another significant contribution in the meta-learning field is the development of Prototypical Networks by Snell, Swersky, and Zemel (2017). These networks adopt a metric-based approach, learning a metric space in which classification is performed by computing distances between embedded samples and class prototypes. Prototypical Networks are computationally lighter than gradient-based meta-learning methods and have demonstrated strong performance in few-shot learning tasks, making them particularly promising for IoT devices with limited computational power.

With the expansion of edge computing, significant research has focused on developing machine learning models optimized for low-power and memory-limited IoT devices. A major milestone in this area is the emergence of TinyML, formally introduced by Warden and Situnayake (2019). TinyML emphasizes deploying ML models directly on microcontrollers with as little as tens of kilobytes of memory. This work highlights techniques such as model compression, low-bit quantization, and efficient inference engines designed specifically for embedded systems. TinyML has gained widespread adoption for speech recognition, anomaly detection, and sensor data analysis on ultra-low-power IoT devices.

Model compression techniques such as pruning and quantization have also been extensively studied. Han, Mao, and Dally (2016) demonstrated that deep neural networks can be significantly compressed through weight pruning, quantization, and Huffman coding without substantial accuracy loss. Their work, known as "Deep Compression," provided the theoretical foundation for modern lightweight neural networks used in IoT systems. Follow-up studies by Jacob et al. (2018) introduced quantization-aware training, allowing 8-bit inference on edge devices while preserving model accuracy an essential advancement for running ML models on microcontrollers.

In addition to compression and federated approaches, lightweight neural architectures have been proposed for edge deployment. Howard et al. (2017) introduced MobileNet, a CNN architecture leveraging depthwise separable convolutions to reduce computational complexity drastically. Similarly, Sandler et al. (2018) extended this work through MobileNetV2, improving model efficiency using inverted residuals and linear bottlenecks. For sequence tasks, Cho et al. (2014) introduced the GRU, a simpler alternative to LSTMs with fewer parameters, making it more suitable for IoT applications. More recently, Houlsby et al. (2019) proposed lightweight Transformer variants using parameter-efficient adapters opening pathways for deploying attention mechanisms on constrained hardware.

Edge computing frameworks have played a crucial role in making ML deployment feasible on microcontroller and mobile-class hardware. TensorFlow Lite Micro (TFLM), released by David et al. (2020), is one of the first ML frameworks designed specifically for devices with less than 100 KB of RAM. TFLM provides a minimal runtime, static memory planning, and support for quantized models, enabling ML inference on boards such as ARM Cortex-M and ESP32. The framework has become a cornerstone of TinyML research, enabling real-time ML applications such as wake-word detection and sensor anomaly detection.

Edge Impulse, founded and formally described by Jousset, Evans, and Vlasov (2020), introduced an end-to-end pipeline for building TinyML models, including data acquisition, feature extraction, model training, and deployment. It supports a wide range of IoT hardware platforms and automates model optimization steps, lowering the barrier for deploying ML on embedded devices[4]. Edge Impulse also integrates hardware-aware profiling tools that allow developers to analyze latency, memory usage, and energy consumption.

At the hardware level, ultra-low-power microcontrollers have advanced significantly in supporting ML workloads. Research on the ESP32, STM32, and ARM Cortex-M series (notably by Lattner et al., 2020, and manufacturers' technical documentation) highlights improvements in low-power modes, DSP extensions, and hardware accelerators such as ARM's CMSIS-NN. These developments provide the computational foundation for deploying optimized meta-learning models directly on resource-constrained IoT devices.

However, applying meta-learning to IoT devices is not straightforward. Many meta-learning algorithms, such as Model-Agnostic Meta-Learning (MAML) or metric-based few-shot learning, still require considerable computation and memory during both training and adaptation phases. Therefore, the development of lightweight meta-learning methods that are optimized for constrained hardware is essential. This includes exploring model compression techniques, quantization, pruning, and novel architectures specifically designed for microcontrollers. Additionally, evaluating these algorithms in realistic IoT scenarios is critical to understanding the trade-offs between accuracy, resource usage, and adaptation speed.

Given these needs and challenges, research on meta-learning algorithms for resource-constrained intelligent IoT devices is both timely and necessary. It addresses the urgent demand for intelligent systems that can perform on-device learning efficiently while operating under strict hardware limitations. By advancing this field, researchers can contribute to the next generation of smart, adaptive, and autonomous IoT systems that are capable of functioning reliably across diverse and dynamic real-world environments.

2. Research Methodolgy

The methodology of this study is structured to develop, implement, and evaluate a meta-learning framework optimized for deployment on resource-constrained IoT devices. The research is divided into five key components: selection of the meta-learning approach, design of a lightweight neural architecture, specification of the hardware platform, configuration of the training setup, and execution of a series of experiments to analyze performance under real-world constraints. Each component is designed to ensure that the final model can perform rapid on-device adaptation while operating within the strict computational, memory, and energy limitations of typical microcontroller-based IoT systems.

a. Selected Meta-Learning Framework

This study utilizes a hybrid meta-learning framework combining gradient-based and metric-based techniques. The core algorithm selected is Model-Agnostic Meta-Learning (MAML), which serves as the baseline for gradient-based fast adaptation[5]. MAML is chosen due to its strong theoretical foundation for enabling quick fine-tuning with limited data, making it suitable for dynamic IoT environments. However, because full MAML is computationally expensive for microcontroller-class hardware, a first-order approximation such as Reptile or First-Order MAML is adopted to reduce gradient computation. To further enhance efficiency, Prototypical Networks, a metric-based meta-learning method, are incorporated for comparison due to their significantly lower computational overhead and suitability for lightweight inference. The combination of these approaches allows the study to investigate the trade-offs between adaptation speed, computational complexity, and memory usage, ultimately identifying the most feasible meta-learning strategy for constrained IoT devices.

b. Design of Lightweight Architecture

The lightweight architecture designed for this study follows TinyML principles and is optimized using model compression techniques to meet microcontroller limitations. The model consists of a compact neural network with a small number of convolutional layers (e.g., 2–3 layers for image tasks or 1–2 fully connected layers for sensor data tasks) paired with low-dimensional embedding layers for meta-learning[6]. To reduce memory footprint and inference latency, the architecture employs 8-bit post-training quantization, enabling the model to run efficiently on integer-only hardware. Additional compression strategies such as magnitude-based pruning and weight sharing are applied to further reduce the number of parameters without significantly degrading accuracy. For metric-based models like Prototypical Networks, lightweight embedding modules replace deeper architectures to ensure that feature extraction remains fast and energy-efficient. The combination of quantization, pruning, and minimal-layer design ensures that the resulting model fits within the memory constraints of typical IoT microcontrollers.

c. Hardware Platform Specifications

Implementation and validation of the meta-learning framework are conducted on a low-power microcontroller platform representative of mainstream IoT devices. The selected hardware includes an ARM Cortex-M4F or Cortex-M7 CPU with clock speeds ranging from 80–200 MHz, offering both floating-point support and efficient integer arithmetic for quantized models[7]. The device typically includes 256–512 KB of RAM and 1-2 MB of Flash storage, providing sufficient capacity for deploying compressed neural networks. Power consumption is strictly limited, with devices operating on battery power or low-voltage power supply. Available sensors such as accelerometers, environmental sensors (temperature, humidity, gas), or simple imaging modules provide input data for model evaluation. The study also incorporates additional platforms such as the ESP32 or STM32 series to ensure cross-device feasibility and to evaluate differences in performance across hardware variations.

d. Training Setup

Model training is conducted in two stages: meta-training and fine-tuning (adaptation). For the meta-training phase, datasets relevant to IoT applications such as motion sensor datasets (e.g., HAR), environmental monitoring datasets, or custom-collected sensor data are used to create task distributions. Each task consists of small support and query sets to simulate few-shot learning conditions. During meta-training, models are optimized using loss functions such as cross-entropy loss for classification tasks or mean squared error for regression-based sensor predictions[8]. The meta-training phase is performed offline using a high-performance computing environment to avoid burdening IoT hardware. In the fine-tuning phase, the meta-learned model is deployed onto the IoT device and adapted using only a few local samples. Evaluation metrics include accuracy, latency, memory consumption, and energy usage, measured using on-device profiling tools. These metrics allow the study to determine whether the meta-learned model meets the real-time and resource constraints of IoT platforms.

e. Experiments

A series of experiments are conducted to evaluate the performance and feasibility of deploying meta-learning models on constrained IoT hardware[9]. First, the model is tested under varying memory and power constraints, including reduced RAM availability and low-voltage operating conditions, to observe performance degradation and robustness. Second, the task adaptation speed is measured by recording the time and number of samples required for the model to adapt to new sensor conditions, environmental changes, or user-specific patterns. Third, the meta-learning models are compared against three baselines: standard machine learning models (e.g., logistic regression, SVM), full-size meta-learning models trained without compression, and prior TinyML models commonly used in IoT applications. Metrics such as inference time, adaptation time, battery consumption, and classification accuracy are evaluated across all models. These experiments provide a comprehensive view of the advantages and limitations of meta-learning when deployed in real-world IoT environments, and they help identify the most optimal configuration for resource-constrained intelligent devices.

3. Results and Discussion

Lightweight Meta-Learning Model Suitable for Microcontroller-Class IoT Devices

Designing a meta-learning model capable of running on microcontroller-class IoT devices requires a careful balance between computational efficiency and rapid task adaptation. A promising approach is the integration of First-Order MAML (FOMAML) or Reptile as the core adaptation mechanism, since both eliminate second-order gradient calculations and substantially reduce the computational overhead. These first-order variants allow the meta-learner to update parameters using simple gradient descent with significantly lower memory requirements. To further improve efficiency, the adaptation step can be limited to a small subset of parameters, such as the last layer or a lightweight embedding module, which drastically reduces the number of gradients that must be stored[10]. This parameter-efficient fine-tuning strategy has been shown to enable rapid on-device adaptation while keeping resource usage minimal, making it suitable for microcontroller environments where memory is extremely limited.

To support efficient representation learning, the meta-learning architecture is designed as a compact neural network with a minimal number of layers. For example, a lightweight convolutional encoder with 1-3 convolutional layers and small kernel sizes is sufficient for image or sensor pattern recognition tasks, whereas fully connected networks with 1-2 hidden layers may be used for time-series sensor data. The network is optimized using 8-bit quantization, which converts floating-point weights into integer representations that can be processed natively by microcontroller hardware. Quantization not only reduces memory consumption but also significantly accelerates inference due to the reduced computational precision. Additional compression techniques such as magnitude pruning, weight sharing, and low-rank factorization further shrink the model's size, allowing it to fit within tight Flash and RAM budgets without severely compromising accuracy[11].

In addition to structural compression, metric-based meta-learning approaches such as Prototypical Networks offer an alternative lightweight solution. These models compute class prototypes in a learned embedding space and perform classification by calculating distances between embedded query samples and prototypes. Because Prototypical Networks do not require gradient-based adaptation, they eliminate the heavy computational cost associated with fine-tuning and allow rapid adaptation using simple distance calculations. This makes them particularly attractive for IoT devices with strict energy or memory constraints, where even first-order gradient updates may be expensive. In such configurations, only the embedding encoder needs to be stored and executed on the device, while the adaptation process involves only lightweight arithmetic operations.

When deployed on microcontroller-class platforms such as ARM Cortex-M4F, Cortex-M7, STM32, or ESP32, the model must also respect the hardware's limited energy availability. The lightweight meta-learning model minimizes power consumption through reduced activation memory, optimized quantized kernels, and hardware-aware layer execution[12]. Frameworks like TensorFlow Lite Micro and CMSIS-NN can be used to ensure that inference operations run with maximum efficiency on available DSP extensions. The combination of efficient architecture, compressed parameters, and simplified meta-learning updates results in a model that can perform few-shot learning, anomaly detection, or personalized adaptation directly on-device while consuming only milliwatts of power.

Overall, the proposed lightweight meta-learning model demonstrates that it is feasible to bring fast, adaptive intelligence to the edge, even on extremely resource-constrained microcontrollers. By merging first-order meta-learning algorithms with quantized and pruned neural architectures, the model achieves a practical balance between learning capability and hardware limitations. This opens new opportunities for real-time personalization, adaptive sensing, and autonomous decision-making in next-generation IoT systems where cloud connectivity cannot always be relied upon.

Faster Adaptation to New Tasks with Minimal Retraining

One of the primary motivations for adopting meta-learning in resource-constrained IoT devices is the ability to achieve rapid adaptation to new tasks with minimal retraining. Traditional machine learning models typically require large amounts of data and many iterations of gradient updates to learn a new task, making them unsuitable for dynamic IoT environments where conditions, users, and sensor patterns change frequently. Meta-learning addresses this limitation by training models to "learn how to learn," enabling them to internalize representations and update patterns that generalize across related tasks[13]. As a result, when a new task arises such as a different user's activity pattern, a new environmental condition, or a newly deployed sensor the model can adapt with only a few samples and a minimal number of computation steps. This characteristic is especially critical for IoT devices that must operate autonomously in real-time without cloud connectivity or continuous retraining.

The process of achieving faster adaptation relies heavily on the meta-learning strategy used during training. Gradient-based methods such as First-Order MAML and Reptile optimize the model parameters such that only one or two gradient steps are needed to specialize to a new task. During deployment, the IoT device does not need to retrain the model from scratch; instead, it applies a lightweight adaptation procedure that updates a small subset of parameters or adjusts an embedding representation. Because these updates are computationally inexpensive, adaptation can be completed within milliseconds on microcontroller-class hardware. This allows the system to quickly respond to changing operational contexts, such as shifts in sensor calibration, new environmental patterns, or user-specific variations, without requiring cloud-based processing or prolonged fine-tuning.

Metric-based methods such as Prototypical Networks further accelerate the adaptation process by eliminating retraining altogether in many scenarios. Instead of updating model parameters, these methods compute class prototypes based on a few labeled examples of the new task. This enables nearinstant adaptation, as the model only needs to update its internal representation of each class rather than performing gradient-based learning. For IoT settings such as anomaly detection in industrial equipment or gesture recognition on wearable devices this capability is invaluable because it avoids the computational cost of backpropagation while still enabling accurate classification of previously unseen patterns. The combination of a lightweight encoder and fast distance computation makes prototypical adaptation one of the most efficient strategies for on-device learning[14].

Minimal retraining also reduces energy consumption, which is essential for battery-powered IoT devices. Training operations, especially those involving floating-point calculations, consume significantly more energy than inference operations. By limiting the number of gradient updates or eliminating them entirely, the model conserves power while maintaining high responsiveness. This efficiency allows the IoT device to operate for extended periods in the field without frequent recharging or battery replacement. It also enables real-time decision-making in environments where energy harvesting or intermittent power supply is common.

The ability to adapt rapidly to new tasks with minimal retraining represents a transformative advantage for intelligent IoT systems. It allows devices to personalize their behavior, respond to evolving real-world conditions, and remain robust despite hardware limitations or environmental noise. Faster adaptation not only enhances system accuracy and reliability but also expands the range of applications that can be supported on microcontroller-class hardware, including smart wearables, environmental monitoring nodes, home automation sensors, and industrial IoT devices[15]. Through the synergy of lightweight architectures and meta-learning algorithms, IoT devices can achieve a level of intelligence that was previously only possible through cloud-based computation.

Lower Energy Consumption Compared to Traditional Machine Learning

One of the most compelling advantages of lightweight meta-learning models in microcontroller-class IoT environments is their significantly lower energy consumption compared to traditional machine learning approaches. Conventional ML models particularly deep learning architectures such as CNNs, LSTMs, or transformers require substantial computational resources to perform both training and inference[16]. These models depend heavily on high-performance processors or GPU-accelerated systems, making them unsuitable for low-power IoT devices that typically operate with limited CPU cycles, small memory footprints, and strict energy budgets.

In contrast, lightweight meta-learning models are specifically designed to minimize computational requirements. Their algorithms emphasize rapid parameter adaptation rather than extensive retraining from scratch, allowing IoT devices to process sensor data and update decision boundaries using only a fraction of the power needed by traditional models. Many lightweight meta-learning frameworks reduce energy usage by leveraging compact architectures, simplified optimization routines, and sparse computation techniques. This makes them highly scalable across battery-powered devices such as wearables, smart sensors, environmental monitoring nodes, and industrial IoT components.

Energy efficiency also plays a crucial role in extending the operational lifetime of IoT systems deployed in remote or inaccessible environments[17]. Traditional ML models would require frequent energy replenishment or expensive hardware upgrades to maintain performance, whereas lightweight meta-learning models enable continuous on-device learning without exhausting power resources. By reducing the number of required training iterations and memory accesses both of which are high-energy operations the overall power draw of the IoT device remains low while still ensuring robust adaptive intelligence.

Furthermore, lower energy consumption directly contributes to sustainable computing initiatives. As the number of IoT devices continues to increase globally, the cumulative impact of energy usage becomes a critical factor. Lightweight meta-learning helps mitigate this by enabling distributed intelligence that does not rely on constant cloud communication, which not only decreases latency but also minimizes the energy cost associated with wireless data transmission. With these efficient learning mechanisms, IoT ecosystems can operate more autonomously, maintain longer device lifecycles, and support environmentally friendly, large-scale deployments[18].

Improved Robustness to Changes in the Environment and Sensor Noise

Lightweight meta-learning models offer a significant advantage in terms of robustness when deployed on microcontroller-class IoT devices operating in dynamic and unpredictable environments. Traditional machine learning models typically assume that training and deployment conditions remain relatively stable. When exposed to environmental variations such as temperature changes, humidity fluctuations, lighting differences, or unexpected physical interference these models often experience performance degradation. In contrast, meta-learning models are explicitly trained to adapt quickly and generalize across varying contexts, allowing them to maintain stability even as external conditions evolve.

One of the key reasons meta-learning provides greater robustness is its ability to learn adaptation strategies rather than fixed representations[13]. Instead of relying on static parameters that may fail under noisy or shifting input conditions, meta-learning models leverage experience from diverse tasks to recognize patterns in how environmental noise affects sensor data. This enables them to recalibrate decision boundaries, adjust internal weights, or correct for anomalies without requiring extensive retraining. As a result, meta-learning-powered IoT devices can sustain reliable performance even when the input data is partially corrupted, incomplete, or subject to unexpected disturbances.

Sensor noise is an inherent challenge in IoT systems, particularly those relying on low-cost or low-power sensors commonly used in remote monitoring, industrial automation, and wearable devices[19]. Traditional ML models often struggle under these conditions because they lack the capacity to differentiate between meaningful features and random fluctuations in the data. Lightweight meta-learning models address this limitation by incorporating noise resilience directly into the adaptation process. They can learn to filter out irrelevant variations, strengthen signal interpretation, and adjust learning parameters in real time. This makes them particularly effective for tasks such as anomaly detection, equipment health monitoring, and environmental sensing.

Moreover, improved robustness enhances the operational reliability and longevity of IoT deployments. Devices that cannot cope with environmental changes or sensor noise often require frequent remote updates, recalibration, or manual maintenance all of which increase operational costs and reduce system efficiency[20]. Lightweight meta-learning reduces this dependency by empowering IoT nodes to adapt autonomously, ensuring consistent performance over long periods. This robustness also supports scalability, allowing large networks of heterogeneous IoT devices to operate in real-world conditions without constant supervision or cloud-based intervention.

Limitations

Despite the promising advantages of integrating meta-learning techniques into resource-constrained IoT systems, several limitations must be acknowledged. The most fundamental challenge lies in the hardware limitations that restrict the overall model complexity[21]. Microcontroller-class IoT devices such as ARM Cortex-M series, ESP32, or STM32 possess very limited RAM, flash storage, and computational throughput. These constraints make it impossible to implement complex meta-learning architectures that require large numbers of parameters, deep layers, or heavy mathematical operations. Consequently, many state-of-the-art meta-learning algorithms must be heavily simplified, pruned, or quantized, which can potentially reduce their learning capacity.

Another significant limitation is the inherent computational cost of meta-learning itself. Although the proposed lightweight designs aim to minimize overhead, meta-learning especially gradient-based approaches like MAML typically involves inner-loop and outer-loop updates that are more computationally intensive than traditional supervised learning. Performing these updates on-device may still exceed the processing capability of typical microcontrollers. Even when computations are offloaded to external servers during the training phase, fine-tuning or real-time adaptation on the IoT node can impose non-trivial latency or increased energy consumption. This makes it challenging to deploy meta-learning models in applications requiring ultra-fast response times or extremely low-power operation.

Efficient memory allocation presents another critical constraint. Microcontrollers usually provide only tens to hundreds of kilobytes of SRAM, forcing developers to optimize every layer and tensor[22].

Running meta-learning models, which often require additional memory to store task-specific gradients, temporary variables, or adaptation parameters, can push the limits of available memory. Missing or inefficient memory allocation can lead to runtime failures, system crashes, or incomplete training cycles[23]. Therefore, memory optimization techniques such as fixed-point arithmetic, weight sharing, and aggressive quantization are necessary but can degrade model performance if overly applied.

Finally, there is a trade-off between accuracy and model size that cannot always be perfectly balanced. Simplifying a meta-learning model to fit within IoT hardware constraints often results in reduced representational capacity, which may limit its ability to capture complex patterns or adapt effectively to diverse tasks. Conversely, increasing model size to improve accuracy risks exceeding memory budgets or raising energy consumption beyond acceptable levels for edge devices. This inherent trade-off means that designers must prioritize based on application needs whether accuracy, speed, robustness, or energy efficiency is most important and accept certain performance sacrifices accordingly.

4. Conclusion

This research demonstrates that meta-learning offers a powerful paradigm for enabling intelligent IoT devices to adapt quickly and efficiently to new tasks despite severe hardware and energy constraints. By designing a lightweight meta-learning architecture optimized for microcontroller-class platforms, the study shows that it is possible to achieve rapid task adaptation, improved robustness to environmental variations, and lower energy consumption when compared to traditional machine learning approaches. The integration of quantization, pruning, and optimized memory allocation further enhances the feasibility of deploying meta-learning models on devices with limited computational resources. The findings highlight that meta-learning significantly reduces the need for extensive retraining, enabling IoT nodes to update their behavior in real time based on only a few examples. This capability is essential for dynamic, unpredictable environments where sensor characteristics, operating conditions, and task requirements frequently change. Additionally, the proposed model exhibits resilience against sensor noise and environmental fluctuations, ensuring reliable performance in practical edge deployments. These improvements contribute to more autonomous and efficient IoT systems capable of long-term operation without constant human intervention or cloud dependency. However, the study also acknowledges several limitations, including computational overhead, restricted memory capacity, and the unavoidable trade-off between accuracy and model size. These constraints suggest that while meta-learning is viable for small IoT devices, it requires careful architectural design and resource optimization. As such, the research emphasizes the importance of hardware-aware model development and the need for further innovations in efficient training strategies and ultra-lightweight meta-learning algorithms. Overall, this research contributes to the growing field of TinyML and intelligent edge computing by providing a practical framework for applying meta-learning in highly resource-limited environments. The insights gained underscore the potential of meta-learning to transform IoT systems from static, pre-programmed devices into adaptive, context-aware intelligent agents. Future work may explore even lighter meta-learning variants, hardware-accelerated implementations, and broader real-world testing across diverse IoT applications.

References

[1] I. Ficili, M. Giacobbe, G. Tricomi, and A. Puliafito, "From sensors to data intelligence: Leveraging IoT, cloud, and edge computing with AI," *Sensors*, vol. 25, no. 6, p. 1763, 2025.

[2] M. Capra, R. Peloso, G. Masera, M. Ruo Roch, and M. Martina, "Edge computing: A survey on the hardware requirements in the internet of things world," *Futur. Internet*, vol. 11, no. 4, p. 100, 2019.

[3] V. B. Parthasarathy, A. Zafar, A. Khan, and A. Shahid, "The ultimate guide to fine-tuning llms from basics to breakthroughs: An exhaustive review of technologies, research, best practices, applied research challenges and opportunities," *arXiv Prepr. arXiv2408.13296*, 2024.

[4] T. S. Ajani, A. L. Imoize, and A. A. Atayero, "An overview of machine learning within embedded and mobile devices–optimizations and applications," *Sensors*, vol. 21, no. 13, p. 4412, 2021.

[5] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*, PMLR, 2017, pp. 1126–1135.

[6] P. Zhang, Y. Bai, D. Wang, B. Bai, and Y. Li, "Few-shot classification of aerial scene images via meta-learning," *Remote Sens.*, vol. 13, no. 1, p. 108, 2020.

[7] H. A. Abushahla, D. Varam, A. J. N. Panopio, and M. I. AlHajri, "Quantized Neural Networks for Microcontrollers: A Comprehensive Review of Methods, Platforms, and Applications," *arXiv Prepr. arXiv2508.15008*, 2025.

[8] J. Gao, D. Wu, F. Yin, Q. Kong, L. Xu, and S. Cui, "MetaLoc: Learning to learn wireless localization," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 12, pp. 3831–3847, 2023.

[9] Y. Wu *et al.*, "Masinet: Network intrusion detection for iot security based on meta-learning framework," *IEEE Internet Things J.*, vol. 11, no. 14, pp. 25136–25146, 2024.

[10] Y. Chen, B. Zheng, Z. Zhang, Q. Wang, C. Shen, and Q. Zhang, "Deep learning on mobile and embedded devices: State-of-the-art, challenges, and future directions," *ACM Comput. Surv.*, vol. 53, no. 4, pp. 1–37, 2020.

[11] T. Choudhary, V. Mishra, A. Goswami, and J. Sarangapani, "A comprehensive survey on model compression and acceleration," *Artif. Intell. Rev.*, vol. 53, no. 7, pp. 5113–5155, 2020.

[12] N. S. Murthy, P. Vrancx, N. Laubeuf, P. Debacker, F. Catthoor, and M. Verhelst, "Learn to learn on chip: Hardware-aware meta-learning for quantized few-shot learning at the edge," in *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)*, IEEE, 2022, pp. 14–25.

[13] A. Vettoruzzo, M.-R. Bouguelia, J. Vanschoren, T. Rögnvaldsson, and K. C. Santosh, "Advances and challenges in meta-learning: A technical review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 7, pp. 4763–4779, 2024.

[14] Y. Liu, S. Duan, X. Xu, and S. Ren, "MetaLDC: Meta learning of low-dimensional computing classifiers for fast on-device adaption," *arXiv Prepr. arXiv2302.12347*, 2023.

[15] S. S. Saha, S. S. Sandha, and M. Srivastava, "Machine learning for microcontroller-class hardware: A review," *IEEE Sens. J.*, vol. 22, no. 22, pp. 21362–21390, 2022.

[16] T. Mowbray, "A Survey of Deep Learning Architectures in Modern Machine Learning Systems: From CNNs to Transformers," *J. Comput. Technol. Softw.*, vol. 4, no. 8, 2025.

[17] S. Popli, R. K. Jha, and S. Jain, "A survey on energy efficient narrowband internet of things (NBIoT): architecture, application and challenges," *IEEE access*, vol. 7, pp. 16739–16776, 2018.

[18] B. Qian *et al.*, "Orchestrating the development lifecycle of machine learning-based IoT applications: A taxonomy and survey," *ACM Comput. Surv.*, vol. 53, no. 4, pp. 1–47, 2020.

[19] S. H. Abdulhussain *et al.*, "A comprehensive review of sensor technologies in IOT: Technical aspects, challenges, and future directions," *Computers*, vol. 14, no. 8, p. 342, 2025.

[20] A. Radovici, I. Culic, D. Rosner, and F. Oprea, "A model for the remote deployment, update, and safe recovery for commercial sensor-based iot systems," *Sensors*, vol. 20, no. 16, p. 4393, 2020.

[21] V. Sze, Y.-H. Chen, J. Emer, A. Suleiman, and Z. Zhang, "Hardware for machine learning: Challenges and opportunities," in *2017 IEEE custom integrated circuits conference (CICC)*, IEEE, 2017, pp. 1–8.

[22] B. Prince and D. Prince, *Memories for the intelligent internet of things*. John Wiley & Sons, 2018.

[23] V. Sridharan *et al.*, "Memory errors in modern systems: The good, the bad, and the ugly," *ACM SIGARCH Comput. Archit. News*, vol. 43, no. 1, pp. 297–310, 2015.