Proceedings of the 10th North American International Conference on Industrial Engineering and Operations Management, Orlando, Florida, USA, June 17-19, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/NA10.20250351
Published: June 17, 2025

Adaptive Predictive Maintenance for Multi-Component Systems Using Ensemble-based Meta-Learning

Mainak Mallick
G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Young-Dae Shim
G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA
Department Smart Manufacturing Technology
Sungkyunkwan University
Suwon-si 16419, Gyeonggi-do, Republic of Korea

Hyuck-kee Lee
Korea Automotive Technology Institute (KATECH)
303 Pungse-ro, Pungse-Myeon, Cheonan-Si
Chungnam, Republic of Korea

Seong-Geun Shin
Korea Automotive Technology Institute (KATECH)
303 Pungse-ro, Pungse-Myeon, Cheonan-Si, Chungnam
Republic of Korea

Seung-Kyum Choi
G.W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Abstract

Predictive maintenance (PdM) in robotic arms is essential to reduce downtime and ensure efficiency in industrial assembly lines. Model-Agnostic Meta-Learning (MAML), combined with digital twins, offers a promising approach for rapid fault identification and classification. However, existing MAML-based approaches suffer from challenges such as hypersensitivity from learning parameters along with limited generalization in testing domain. To address these limitations, we propose an ensemble-based meta-learning approach that integrates majority voting with MAML and operational grouping strategies. This method enhances few-shot learning capabilities, improves generalization, and stabilizes model performance across varying conditions. Our framework is validated using a synthetic vibration signal dataset generated via a digital twin, simulating different robotic arm faults. The proposed approach demonstrates higher accuracy in classifying a broader range of defective mechanical classes, specifically in cross-domain few-shot

(CDFS) learning settings. Comparative analysis with alternative meta-learning frameworks, including Reptile, Protonet, and ANIL, confirms the effectiveness of our approach. By leveraging ensemble-based learning, we achieve improved robustness and higher classification accuracy, making our method a viable solution for real-world PdM applications in industrial robotics. The integration of digital twins further enhances the model's reliability, bridging the gap between simulation and real-world deployment. Additionally, this approach reduces data dependency, allowing effective fault classification even in scenarios with limited labelled data, making it highly adaptable to dynamic industrial environments.

Keywords
Predictive maintenance, Industrial Robotics, Meta Learning, Digital Twin, Few Shot Learning

1. Introduction
Modern manufacturing facilities are increasingly automated and employ robotics to replace human labor in repetitive and hazardous tasks that involve handling heavy loads. Beyond enhancing worker safety, industrial robots significantly improve productivity by operating continuously for extended periods. Despite advancements in mechanical components that enhance reliability, factors such as material imperfections, operating conditions, and environmental influences can still cause degradation over time. This deterioration may result in tool position deviations or vibrations in the robotic structure, ultimately affecting the quality of manufactured products. To mitigate these issues and prevent the production of defective goods, it is essential to predict the maintenance schedule of robotic systems accurately.

General maintenance is typically performed either after a failure has occurred or at predetermined intervals. However, this approach can negatively impact the profitability of a company due to unexpected production halts or the premature replacement of functional components. With advancements in sensor technology, mechanical systems are now equipped with sensors to collect real-time data on parameters such as temperature, humidity, acceleration, and voltage, enabling condition-based predictions (Shao et al. 2025). Predictive maintenance (PdM) has emerged as an effective solution for industrial robotics, allowing for the estimation of remaining useful life (RUL) and failure events to optimize maintenance scheduling. As a key component of Industry 4.0, PdM leverages AI to anticipate and plan maintenance activities, enhancing operational efficiency while minimizing downtime (Saleem et al. 2024). Predictive maintenance is further strengthened by the Metaverse's integration of smart sensors and digital twins, which enable real-time asset monitoring in virtual settings, early fault diagnosis, and prompt intervention (Ramolia et al. 2025).

The increasing complexity of industrial robotic systems has driven the adoption of data-driven predictive maintenance (PdM) approaches, leveraging sensor technology for fault detection and diagnosis(Park et al. 2025). Despite these advancements and high relevance, conventional PdM systems often hits a bottleneck require large, labeled datasets, which are often unavailable in industrial settings, posing challenges for fault diagnosis in data-scarce and imbalanced scenarios. Principal component analysis (PCA) and other multi-sensor fusion approaches have been used to improve feature extraction in order to overcome this (Xie et al. 2022). Additionally, to enhance model generalization across tasks with less labeled data, meta-learning techniques—in particular, Model-Agnostic Meta-Learning, or MAML—have been investigated (Che et al. 2022), (Zhang et al. 2024). These methods are especially useful for robotic fault classification, where the sensor data is more noisy and imbalanced, compared to other domains (Sabry & Ungku Amirulddin 2024).

However, MAML faces two major challenges in industrial robotics applications: 1) Hyperparameter Sensitivity – Important variables including batch size, adaption steps, inner-loop and outer-loop learning rates, and others affect MAML stability. Fluctuations during meta-training led to inconsistent generalization, particularly in noisy and imbalanced datasets (Antoniou et al. 2019), and 2) Limited Information in Meta-Testing – Mechanical faults occur infrequently, leading to a high imbalance ratio (IR) (typically 10–20), where fault samples are scarce compared to healthy conditions. This imbalance limits the support set's effectiveness during meta-testing, reducing classification accuracy for unseen fault scenarios.

Researchers have investigated data augmentation methods such time-series transformations and GANs (Wang et al. 2024) to overcome these constraints. Few-shot learning frameworks like ANIL, ProtoNet, and Reptile are also employed to lessen the difficulties (Finn et al. 2017), (Snell et al. 2017), (Nichol et al. 2018)**.** While these approaches improve generalization in low-data regimes, they remain constrained by training instability and limited information intake in meta-testing, particularly in cross-domain few-shot learning (CDFS) scenarios, where the model needs to generalize across a wide range of fault conditions using only a limited amount of labeled data.

We suggest an ensemble-based meta-learning framework called EMOG (Ensemble MAML with Operational Grouping) to close these gaps and improve predictive maintenance generalization and stability. EMOG integrates MAML with a majority voting mechanism, mitigating instability across models and improving fault classification accuracy. The problem of sparse and unbalanced feature distribution in the data points is also successfully addressed by operational grouping, which guarantees maximum information absorption in the meta-testing support set.

Using artificial sensor signal datasets taken from a digital twin of a robotic arm created in Isaac SIM, we verified the suggested framework. The digital model included an inertial measurement unit (IMU) that was used to simulate sensor readings. To transform the sequential sensor data into a structured image representation, we employed a convolutional autoencoder (CAE), which proved effective in capturing the complex, non-linear patterns inherent in the signals. This made it a more suitable choice than principal component analysis (PCA) which is one of the conventional dimensionality reduction methods, as supported by prior work (Siwek & Osowski, 2017; Azarang et al., 2019). The resulting image representations were subsequently used in a few-shot learning setup for meta-testing.

The rest of the document is structured as follows: The general framework and methodology are described in Section 2, the experimental setup is described in Section 3, the results are presented in Section 4 along with comparisons to previous methods, and the study is concluded with important takeaways in Section 5.

1.1 Objectives

The objectives and key contributions of this work are as follows:

1. To improve information usage in few-shot learning scenarios, we provide an ensemble-based model-agnostic meta-learning (MAML) technique that makes use of majority voting and operational grouping.
2. We develop a convolutional autoencoder-based technique for converting multi-sensor time-series signals into fused RGB images, which are subsequently used for fault classification.
3. This is the first study that we are aware of that applies an ensemble-based MAML framework to a synthetic dataset that includes such a wide variety of fault classes.

2. Methodology

This section outlines the study's methodology in two parts. In order to improve information intake for improved generalization, the first one discusses dataset development and operational grouping. The second introduces the ensemble-based MAML framework, addressing hyperparameter instability in meta-learning models.

2.1 Dataset Creation and Operational Grouping Strategy

Dataset generation and preprocessing formed a critical foundation for implementing the operational grouping strategy, beginning with the construction of a digital twin. The RB5 850e robotic arm was modeled using its Unified Robot Description File (URDF) (Garg et al., 2021), which specifies essential structural and physical attributes, including joint constraints, link dimensions, and inertial properties. This URDF model was imported into the NVIDIA Isaac Sim

Figure 1. Digital twin creation workflow.

environment, selected for its ability to deliver high-fidelity physical simulations that accurately replicate forces, collisions, and dynamic behavior. The gain test, a manual calibration procedure, involved iteratively adjusting important physical parameters such applied force (F), stiffness (k), damping (b), and friction coefficient (μ). The following equation was used to represent these characteristics, which were essential for faithfully capturing the motion and vibration dynamics of the robotic arm.:

$$\tau = k\theta + b\dot{\theta} + F_f \quad (1)$$

where τ represents joint torque, θ represents joint angle, $\dot{\theta}$ represents the angular velocity, and $F_f = \mu N$ is the frictional force, which is proportional to the normal force N and friction coefficient μ. This calibration ensured precise articulation and realistic joint behavior during the simulations (Yun et al. 2023). Figure 1 illustrates the digital twin creation and IMU sensor placements on the RB5 850e robotic arm during welding. IMU sensors were positioned at critical joints like base, shoulder, elbow, wrist1, wrist2, and TCP(Tool Center Points) for capturing high-frequency motion and vibration data. These sensors recorded acceleration and angular velocity components (a_x, a_y, a_z, ω_x, ω_y, ω_z) simultaneously at 12,000 Hz, providing fine-grained temporal resolution. Dynamic articulation under realistic settings was made possible by joint controllers using the ROS joint command topic.

For processing, the gathered data was saved in CSV format. Paths A–D in Figure 1 depict the robot's movement. The dataset comprised six classes, representing three welding paths, each labeled as "good" or "bad" based on damping values—high damping (1000 kg/s) indicating faults and low damping (200 kg/s) representing optimal operation. These thresholds were established based on prior research and empirical insights. Studies have highlighted the importance of damping selection in ensuring stable robotic operations. Erickson et al. (2003) estimated optimal damping coefficients at 200 kg/s for force interactions, preventing sluggish performance or excessive oscillations. Similarly, studies on robots controlled by Cartesian impedance indicate that, depending on job and system restrictions, the ideal damping ratios range from 0.2 to 1.0 (Bitz et al. 2020), (Patiño et al. 2023). Additionally, studies on damping matrix design for robotic manipulators reinforce the need for modal analysis and experimental validation to fine-tune damping values(Coleman et al. 2023). Based on these insights, we manually labeled sensor signals into "good signal" and "bad signal" classes, ensuring a data-driven approach to classification.

We employed a convolutional autoencoder (CAE) to transform raw sensor signals into image-based representations suitable for machine learning. The encoder component maps high-dimensional sensor data into a lower-dimensional latent space, effectively capturing salient features and temporal dependencies. The decoder then reconstructs the original input from this compressed representation, facilitating unsupervised feature extraction. Formally, the encoder transformation can be expressed as follows in Equation (2):

$$f_{encoder}(x) = \sigma(W_e x + b_e), \quad (2)$$

where σ is the activation function, W_e and b_e stand for the encoder weights and biases, respectively, and x is the input sensor signal. Equation (3) provides the decoder's reconstruction of the input.

:

$$\hat{x} = f_{decoder}(x) = \sigma(W_d z + b_d), \quad (3)$$

where $\hat{x}$ is the reconstructed output and z represents the latent representation. Equation 4 explains how the learnt parameters, weights (W_e, W_d) and biases (b_e, b_d), are updated while optimizing the reconstruction loss. Here, W_e and b_e transform the input into the latent space, while W_d and b_d map the latent representation back to the original input space.

Equation (4) illustrates how the autoencoder was taught to reduce the reconstruction loss:

$$\mathcal{L}_{reconstruction} = \|x - \hat{x}\|^2 \tag{4}$$

This procedure made guaranteed that important vibration and motion patterns for analysis were preserved in the latent space z, which were subsequently utilized to create 2D images. This approach is depicted in Figure 2.

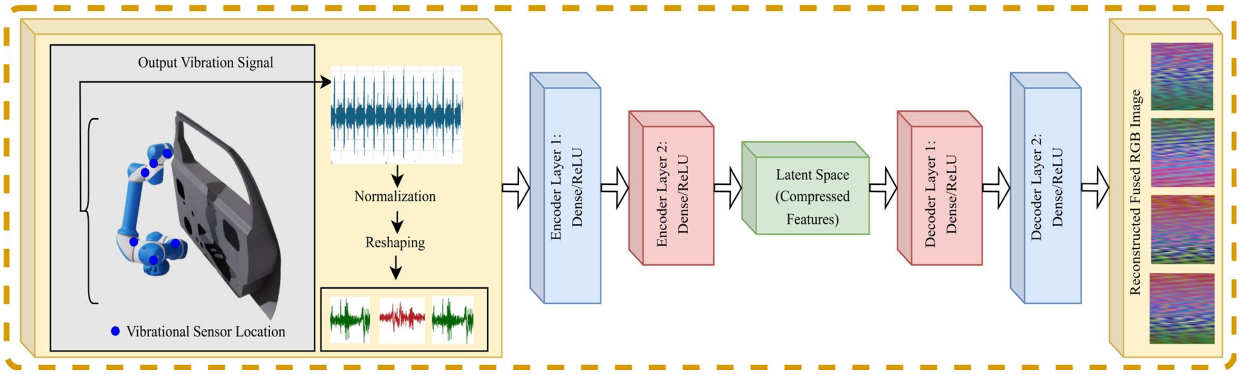

Figure 2. Convolutional autoencoder-based sensor signal to fused RGB image construction.

A vision transformer (ViT) was then used for extracting high-dimensional features from the RGB images. By using self-attention mechanisms, ViT attempts to capture both local and global links through the use of self-attention mechanisms, allowing for a strong feature representation. As described in Equations (5a), (5b), followed by (6a), (6b), it produced a 768-dimensional feature vector for each image.

$$x_p = reshape\big(I, (N, P \times P \times 3)\big), \tag{5a}$$

$$z_0 = \left[x_p^1 E; x_p^2 E; \dots; x_p^N E\right] + P, \tag{5b}$$

where P refers to the positional encoding matrix and N is the number of image patches E is the linear embedding matrix, x_p represents patch embeddings and p is the positional encoding matrix.

$$z^{l+1} = \mathrm{MSA}\big(\mathrm{LN}(z^l)\big) + z^l \tag{6a}$$

$$z^{l+1} = \mathrm{MLP}\big(\mathrm{LN}(z^{l+1})\big) + z^{l+1} \tag{6b}$$

where z^l is the input to the layer at level l. Here, MSA denotes the multi-head self-attention mechanism, LN refers to layer normalization, l indicates the index of the transformer layer, and MLP is the multi-layer feed-forward network. The significance of operational grouping and its role within the framework is further elaborated in the following section.

Next, to ensure maximal data variety, photos were then grouped according to feature distribution using Latin Hypercube Sampling (LHS). This approach identified a diverse and representative set of 5 or 10 images, termed an operational group ($G(D_{support})$), depending on the shot configuration in few-shot learning. The selected support set was then fed into the meta-testing loop, exposing the model to the most informative examples. The feature extraction procedure using the vision transformer and LHS-based operational grouping are shown in Figure 3.

2.2. Ensemble MAML Framework

Model-Agnostic Meta-Learning (MAML) is a meta-learning framework designed for rapid adaptation to new tasks with minimal data and a few gradient updates. It optimizes initial parameters (θ) for quick fine-tuning, enabling models to generalize efficiently without extensive retraining, ideal for fault classification and predictive maintenance. The inner loop, which uses the support set to adapt tasks, and the outer loop, which refines meta-parameters (θ) across tasks based on query set performance, are the two main optimization phases that MAML follow. The following

Equation (8) provides the gradient descent on the task-specific loss used to update the parameters θ for each task $\mathcal{T}_i$ throughout the inner loop:
:

$$\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}^{\text{support}} \tag{8}$$

where θ_i' are task-specific parameters after adaptation and α represents the inner-loop learning rate. This learning rate is important for limiting how quickly the model adjusts to the task-specific data. Nevertheless, excessively high or low values of α can lead to overfitting or inadequate adaptation, respectively, introducing significant sensitivity into the system.

Equation (9) shows how the outer loop improves the meta-parameters θ based on how well the adapted parameters θ_i' perform on the query set.
:

$$\mathcal{L}_{\text{meta}} = \sum_i \mathcal{L}_{\mathcal{T}_i}^{\text{query}} \tag{9}$$

where $\mathcal{L}_{\text{meta}}$ is the meta-learning loss, which is aimed to optimize for fast adaptation to new tasks, and $\mathcal{L}_{\mathcal{T}_i}^{\text{query}}$ represents the query set loss for the individual task $\mathcal{T}_i$. The total meta-learning loss is computed by aggregating the query set losses across multiple tasks. Optimizing this loss is essential, as it directly updates the meta-parameters of the model, enhancing its ability to generalize across diverse tasks. This enables rapid adaptation to previously unseen tasks using only a small number of examples. The meta-parameters θ are updated using the gradient of the meta-loss as in the following Equation (10):

$$\theta \leftarrow \theta - \beta \nabla_\theta \mathcal{L}_{\text{meta}} \tag{10}$$

where β is the learning rate of the outer loop. Like α, β adds sensitivity to the system because it can hinder convergence or cause instability in the training process if it is not properly tuned.

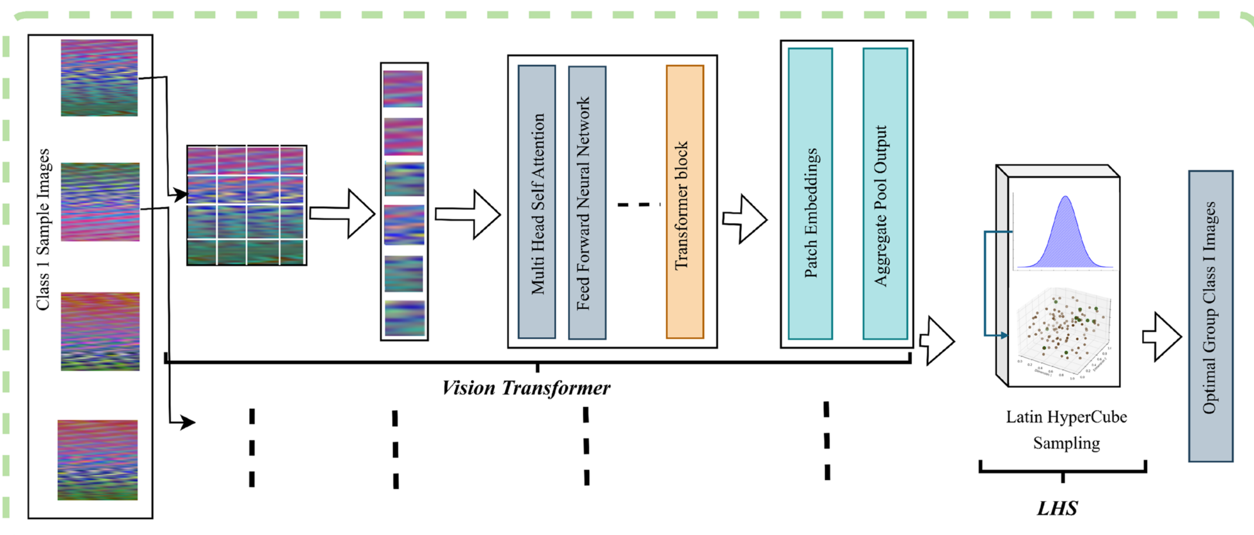

Figure 3. Operational grouping method with the Vision Transformer and LHS method.

MAML's sensitivity to learning rates (α and β) creates instability, where small deviations can cause inefficient adaptation in the inner loop or unstable updates in the outer loop, leading to erratic learning curves and poor generalization. This issue worsens over training, often resulting in divergence.To address this, we introduce an ensemble layer that aggregates predictions from multiple MAML models trained with different learning rates. Using majority voting, this approach minimizes hyperparameter sensitivity and improves robustness and generalization. Independent learning in each model results in diverse decision boundaries, collectively stabilizing the meta-learning process.In this case, x_i stands for the input data, $f_i(x;\ \alpha_i, \beta_i)$ is the class prediction of the i-th model, and α_i and β_i are the inner-loop and outer-loop learning rates for the i-th model, respectively.

Assigning each model in the ensemble a distinct set of learning rates introduces diverse learning dynamics, which enhances the robustness and stability of the overall prediction system.

The prediction from each model for a given input x is represented as in the following Equation (11):

$$y_i = f_i(x; \alpha_i, \beta_i) \quad (11)$$

where M is the total number of classes, and y_i is a member of the set of class labels *{1, 2, ..., M}*. A majority voting procedure is used to determine the ensemble's final forecast, represented by the symbol $y_{ensemble}$,. Equation (12) provides a mathematical expression for this.
:

$$y_{ensemble} = argmax_{j\in 1,2,\dots,M} \sum_{i=1}^{N} I(f_i(x; \alpha_i, \beta_i) = j), \quad (12)$$

where $I(f_i(x;\ \alpha_i, \beta_i) = j)$ is an indicator function that equals 1 if the i-th model predicts class j, and 0 otherwise. The class with the largest number of votes is chosen as the final forecast by this equation, which adds up the votes for each class across all models.

The ensemble models are designed with distinct learning rate configurations to capture different behaviors. For instance, the six models could be configured as $f_1(x;\ \alpha_1, \beta_1)$, $f_2(x;\ \alpha_2, \beta_2)$, and so on, up to $f_6(x;\ \alpha_6, \beta_6)$. The diversity introduced through varying learning rates ensures that the ensemble captures a wide range of learning dynamics, thereby addressing MAML's known sensitivity to specific hyperparameter settings. Each model in the ensemble is trained independently, and their outputs are integrated through an ensemble layer to produce a consolidated prediction. The workflow of this ensemble strategy is illustrated in Figure 4, highlighting how it mitigates hyperparameter sensitivity while enhancing the model's stability and accuracy during both training and task adaptation.

The proposed ensemble approach reduces learning rate sensitivity, stabilizing training and minimizing performance fluctuations. By combining predictions from models with diverse configurations, it enhances generalization and decision reliability. Majority voting further improves robustness by mitigating the influence of outliers and unstable predictions. This method effectively addresses instability while retaining MAML's adaptability. It is suitable for fault prediction and classification in industrial robotics, ensuring scalability and effective deployment.

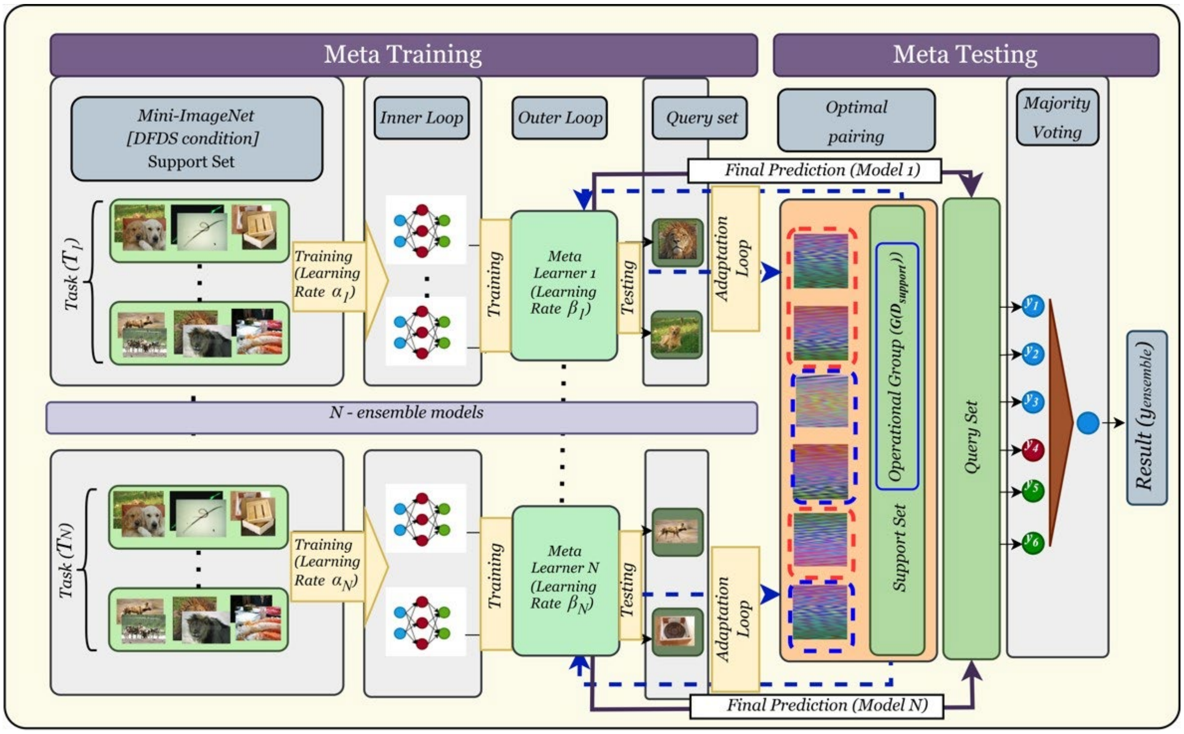

Figure 4. Ensemble MAML workflow.

3. Experimental setup

The experiments were carried out on a workstation configured with an Intel® Xeon® W7-3445 processor (2.59 GHz), 64 GB of RAM, and an NVIDIA A4000 GPU with 16 GB of dedicated memory (HP, USA). This hardware setup facilitated efficient parallel execution of multiple MAML models, ensuring smooth training and evaluation workflows. Six distinct models were trained and evaluated, each utilizing a unique combination of inner-loop (α) and outer-loop (β) learning rates. LHS was used to obtain these parameters inside the bounds $\alpha \in (0.001, 0.003)$ and $\beta \in (0.02, 0.05)$ and are detailed in Table 1. Both the support and query sets for task-based learning were built using MiniImageNet data during the meta-training stage.

In the meta-testing phase, the support set was composed of operationally grouped data, designed to maximize information intake, while the query set served for performance evaluation. The custom dataset comprised six distinct classes, corresponding to variations across the three previously described path scenarios.

Table 1. Different MAML model configurations in EMOG.

Model No.	Inner-Loop Learning Rate (α)	Outer-Loop Learning Rate (β)
M1	0.002364	0.040141
M2	0.001587	0.046868
M3	0.001475	0.001455
M4	0.001068	0.036257
M5	0.002848	0.029255
M6	0.002362	0.024360

To provide a fair comparison across all configurations, the experimental setup maintained several important hyperparameters constant. A steady yet efficient learning pace for updating the meta-learner across tasks was provided by setting the meta-learning rate (meta_lr) at 0.003. In order to facilitate quick adaptation within inner-loop updates during task-specific learning, the fast-learning rate (fast_lr) was set to 0.5. By using a meta-batch size of 32, enough tasks were completed per iteration to improve generalization across various contexts. 10 adaption steps were applied to the model, which enabled it to gradually improve its representations in each task. Lastly, more than 600 iterations of the training were carried out, guaranteeing sufficient exposure to a variety of jobs while preserving computing efficiency.

Six-way one-shot, six-way five-shot, and six-way ten-shot were the three experimental scenarios that were examined in order to assess the model's performance under various data availability conditions. Table 2 provides a summary of the datasets utilized for meta-training and meta-testing.

Table 2. Dataset details used for Meta-Testing.

Sample Number	**Path 1 (A-B)**	**Path 2 (B-C)**	**Path 3 (C-D)**
Correct damping (dp = 250)	50	50	50
Defective damping (dp = 1000)	50	50	50

4. Results and Discussion

The analysis presented in Figure 6 leverages t-SNE visualization to illustrate the effectiveness of the operational grouping strategy in structuring the feature space. In the five-shot case, the selected data points are well-dispersed across the feature space, suggesting that the grouping strategy successfully captures diverse information from the dataset. This diversity in the support set contributes to improved generalization during the model's adaptation phase.
In the ten-shot setting, the distribution of points exhibits an even broader and more uniform spread, further validating the advantages of increasing the support set size. The enhanced coverage of the feature space indicates more comprehensive information intake, allowing the model to access a richer set of representative examples. Overall, the

Table 3. Comparative Results of Different Few Shot Scenarios

	Testing Accuracy (%)			**Precision (%)**			**Recall (%)**			**F1 Score (%)**		
Shot number	1	5	10	1	5	10	1	5	10	1	5	10
ANIL	56.3	61.8	76.2	55.8	62.8	75.5	55.4	64.8	74.1	55.6	63.8	74.8
Reptile	52.4	69.6	76.8	53.5	69	76.2	57.6	68.6	76	55.5	68.8	76.1
ProtoNet	61.9	74.6	77.6	61.3	73.9	78.4	60.8	72.2	77.9	61.0	73.0	78.1
MAML	64.4	73.9	82.8	65.1	73.2	79.4	62.7	72.8	79.8	63.9	73.0	79.6
EMOG	**71.4**	**85.2**	**93.8**	**70.8**	**84.6**	**93.1**	**70.4**	**84.3**	**84.3**	**70.6**	**84.4**	**93.0**

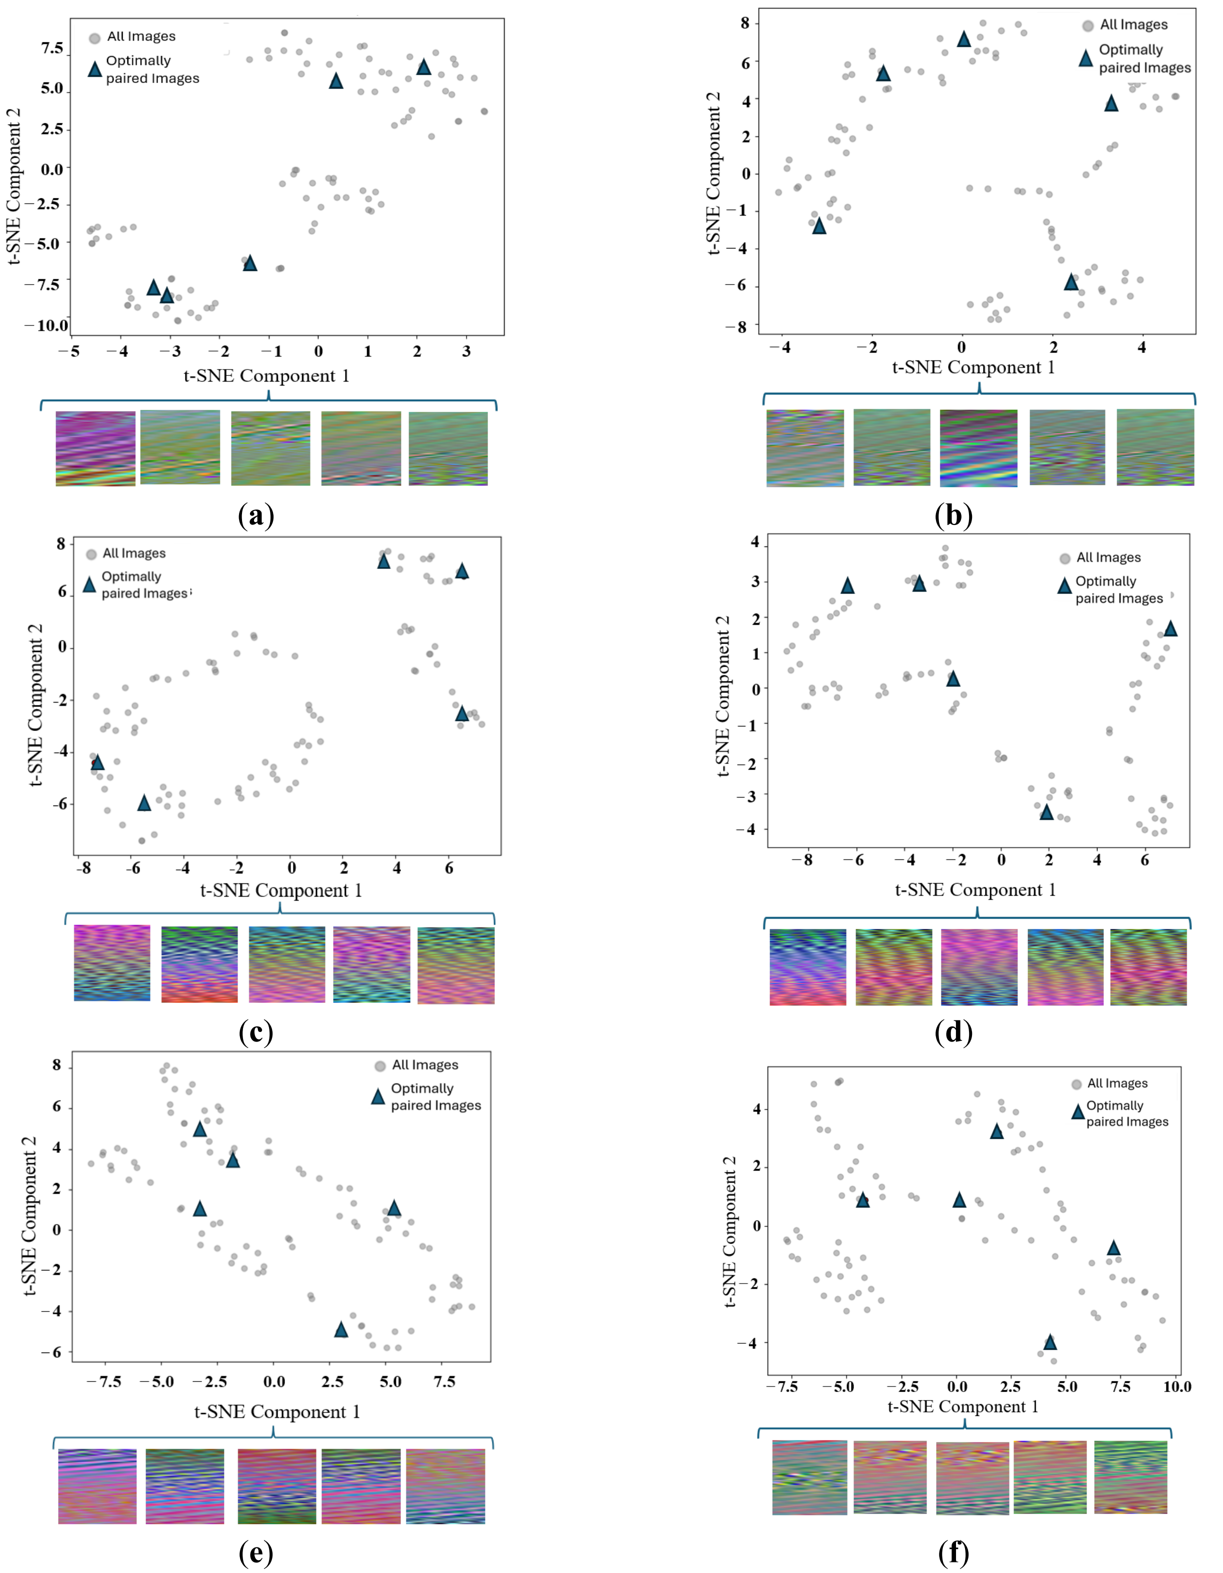

Figure 5. Operationally grouped images for the 5-shot case. (A) Class—Path 1, DP = 250; (B) Class—Path 1, DP = 1000; (C) Class— Path 2, DP = 250; (D) Class— Path 2, DP = 1000; (E) Class— Path 3, DP = 250; (F) Class— Path 3, DP = 1000.

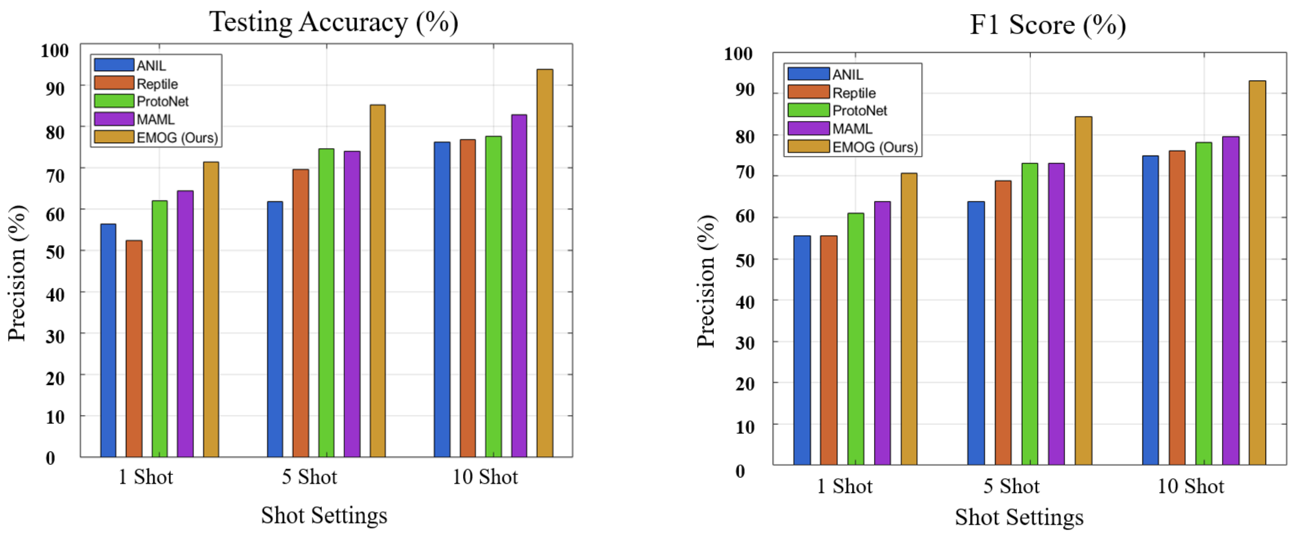

(a) **(b)**

Figure 6. Comparative analysis of different few shot learning method. (A) Accuracy Graph (B) F1 Score Graph.

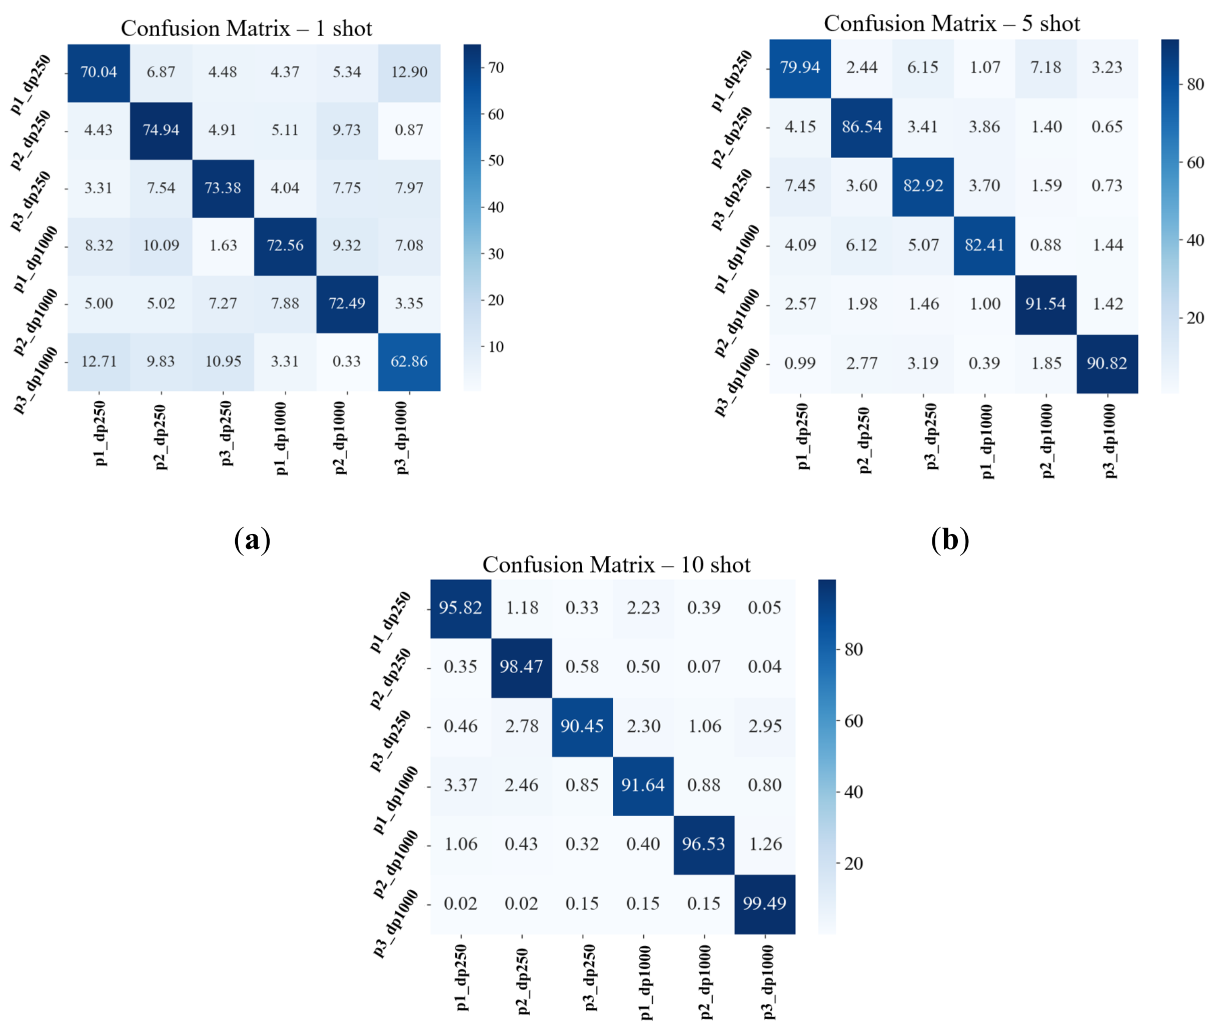

(a) **(b)** **(c)**

Figure 7. Confusion Matrix for (A) 1 Shot EMOG (B) 5 Shot EMOG (C)10 Shot EMOG Method

operational grouping strategy significantly strengthens the support set composition, ultimately leading to more robust learning and better performance during evaluation.

The experimental findings continuously reveal that EMOG performs better across all assessment measures, as indicated in Table 3, especially in multi-shot settings when the benefits of operational grouping are more noticeable. Regarding testing accuracy, EMOG attains 71.4% in the one-shot setting, which is roughly 15% higher than ANIL (56.3%) and 7% higher than MAML (64.4%). In the five-shot scenario, this difference in performance increases, with EMOG reaching 85.2% and surpassing MAML by more than 11%. In the ten-shot scenario, the pattern continues as EMOG achieves an outstanding 93.8%, holding a significant 11% advantage over MAML's 82.8%. These results demonstrate how well EMOG uses the larger support set, with accuracy steadily increasing as the number of shots rises.

Once more, EMOG performs better in terms of precision in every setup. In the five-shot scenario, EMOG outperforms MAML by 11.4% with a precision of 84.6% as opposed to 73.2%. EMOG further improves to 93.1% in the ten-shot scenario, outperforming MAML by 13.7%. This implies that EMOG benefits from the diversity brought about by operational grouping and not only produces predictions that are accurate but also do so with high confidence.

This pattern also applies to recollection. EMOG records a recall of 70.4% in the one-shot setting, which is 7.7% higher than MAML's 62.7%. In the five-shot scenario, where EMOG reaches 84.3%, 11.5% higher than MAML (72.8%), the improvement is more noticeable. EMOG achieves a recall of 92.9% in the ten-shot scenario, which is 13.1% higher than MAML. All of these findings highlight how resilient and flexible EMOG is in a range of few-shot scenarios.

The F1 score, which balances memory and precision, emphasizes even more how resilient EMOG is in various few-shot situations. With an F1 score of 70.6% in the one-shot configuration, EMOG outperforms MAML by 6.7%. In the five-shot scenario, where EMOG records an F1 score of 84.4%, greatly surpassing MAML's 73.0%, this advantage is even more noticeable. With an impressive 13.4% advantage against MAML, EMOG reaches 93.0% in the ten-shot setting. These findings show that EMOG maintains a great balance between precision and recall in addition to achieving high accuracy.

In the five-shot and ten-shot scenarios, where the model can better take use of the structural diversity of the support set, EMOG's operational grouping technique is very beneficial, according to the comparison study shown in Figure 7. The one-shot scenario, on the other hand, demonstrates relatively moderate benefits, indicating that although EMOG is still useful with little data, its full potential is reached when additional support instances are available for adaptation.

Figure 8's confusion matrix shows how well the EMOG framework performs in classification under various shot settings (1-shot, 5-shot, and 10-shot). For the majority of classes, diagonal values consistently show accuracies above 70%, demonstrating good predictive ability. Classification accuracy significantly improves with the amount of support samples; for almost all classes, the 10-shot setting achieves above 95% accuracy. This pattern emphasizes how expanding the size of the support set can improve model generalization and lessen inter-class confusion.

Some classes, however, as p4_dp1000, have high misclassification rates, especially in the lower-shot regimes. This is probably the result of feature overlap between particular classes, which causes categorization ambiguity. The effect is especially noticeable in the one-shot and five-shot settings, where the model's capacity to recognize clear feature boundaries is constrained by the small number of training samples. Higher off-diagonal values indicate trouble differentiating between visually or structurally identical classes, as shown in the one-shot matrix. In contrast, categorization performance greatly improves in 5-shot and particularly 10-shot scenarios. The usefulness of the operational grouping technique in promoting improved class separation and robust generalization is confirmed by the decrease in off-diagonal errors and the appearance of clearer class distinctions.

5.Conclusion

The EMOG method enhances few-shot learning performance, particularly in 5-shot and 10-shot settings, by improving adaptation and feature separation. Its operational grouping strategy strengthens predictive maintenance (PdM) capabilities beyond traditional meta-learning methods. EMOG effectively addresses MAML's key limitations by stabilizing learning rates, accelerating convergence, and handling imbalanced datasets. It optimizes minority class representation, leading to more balanced learning. The confusion matrices confirm its effectiveness, showing reduced misclassification and improved class separability. EMOG outperforms MAML, ProtoNet, ANIL, and Reptile, demonstrating superior generalization, scalability, and efficiency in few-shot predictive maintenance applications.

Acknowledgments
The Korea Automotive Technology Institute (KATECH) has provided the writers with essential financial support and technical experience, for which the authors are extremely grateful.

Conflicts of Interest
The authors declare no conflict of interest.

References
Azarang, A., Manoochehri, H. E., and Kehtarnavaz, N., Convolutional Autoencoder-Based Multispectral Image Fusion, *IEEE Access*, vol. 7, pp. 35673–35683, 2019.

Antoniou, A., Edwards, H., and Storkey, A., How to train your MAML, in *Seventh International Conference on Learning Representations (ICLR)*, New Orleans, United States, Sep. 2019.

Bitz, T., Zahedi, F., and Lee, H., Variable Damping Control of a Robotic Arm to Improve Trade-off between Agility and Stability and Reduce User Effort, in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, May 2020, pp. 11259–11265.

Che, C., Wang, H., Xiong, M., and Ni, X., Few-shot fault diagnosis of rolling bearing under variable working conditions based on ensemble meta-learning, *Digital Signal Processing*, vol. 131, p. 103777, 2022.

Coleman, T., Franzese, G., and Borja, P., Damping Design for Robot Manipulators, 2023, pp. 74–89. doi: 10.1007/978-3-031-22731-8_6.

Erickson, D., Weber, M., and Sharf, I., Contact Stiffness and Damping Estimation for Robotic Systems, *Int. J. Rob. Res.*, vol. 22, no. 1, pp. 41–57, Jan. 2003.

Finn, C., Abbeel, P., and Levine, S., Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks, in *The 34th International Conference on Machine Learning (ICML)*, Sydney, Australia, Mar. 2017. [Online]. Available: http://arxiv.org/abs/1703.03400

Garg, G., Kuts, V., and Anbarjafari, G., Digital Twin for FANUC Robots: Industrial Robot Programming and Simulation Using Virtual Reality, *Sustainability*, vol. 13, no. 18, 2021.

Nichol, A., Achiam, J., and Schulman, J., On First-Order Meta-Learning Algorithms, Mar. 2018, [Online]. Available: http://arxiv.org/abs/1803.02999

Park, S., Youm, M., and Kim, J., IMU Sensor-Based Worker Behavior Recognition and Construction of a Cyber–Physical System Environment, *Sensors*, vol. 25, no. 2, 2025.

Patino, J., Encalada-Dávila, Á., Sampietro, J., Tutivén, C., Saldarriaga, C., and Kao, I., Damping Ratio Prediction for Redundant Cartesian Impedance-Controlled Robots Using Machine Learning Techniques, *Mathematics*, vol. 11, no. 4, p. 1021, Feb. 2023.

Ramolia, N., Tank, P. P., Ravikumar, R. N., Zeb, B., Kumar, M., and Singh, S. K., Futuristic Metaverse: Security and Counter Measures, *IECE Transactions on Intelligent Systematics*, vol. 2, no. 1, pp. 49–65, 2025.

Sabry, A. H. and Ungku Amirulddin, U. A. B., A review on fault detection and diagnosis of industrial robots and multi-axis machines, *Results in Engineering*, vol. 23, p. 102397, 2024.

Saleem, J., Raza, U., and Holderbaum, W., Transforming Industry 4.0 Security: Analysis of ABE and ABA Technologies, *IECE Transactions on Intelligent Systematics*, vol. 1, no. 3, pp. 127–144, Oct. 2024.

Shao, Y., Du, S., and Huang, D., Advancements in Applications of Manufacturing and Measurement Sensors, *Sensors*, vol. 25, no. 2, 2025.

Siwek, K. and Osowski, S., Autoencoder versus PCA in face recognition, in *2017 18th International Conference on Computational Problems of Electrical Engineering (CPEE)*, Kutna Hora, Czech Republic, 2017, pp. 1–4.

Snell, J., Swersky, K., and Zemel, R., Prototypical networks for few-shot learning, in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 4080–4090, 2017.

Wang, H. et al., A novel transformer-based few-shot learning method for intelligent fault diagnosis with noisy labels under varying working conditions, *Reliability Engineering & System Safety*, vol. 251, p. 110400, 2024.

Xie, T., Huang, X., and Choi, S.-K., Intelligent Mechanical Fault Diagnosis Using Multisensor Fusion and Convolution Neural Network, *IEEE Transactions on Industrial Informatics*, vol. 18, no. 5, pp. 3213–3223, 2022.

Yun, J. et al., Digital twin model construction of robot and multi-object under stacking environment for grasping planning, *Applied Soft Computing*, vol. 149, p. 111005, 2023.

Zhang, H., Chen, S., Luo, L., and Yang, J., Few-shot learning with long-tailed labels, *Pattern Recognition*, vol. 156, p. 110806, 2024.

Biographies

Mainak Mallick received a BS in Aerospace Engineering from the Indian Institute of Engineering Science and Technology of Shibpur, India in 2022, and is currently a first-year Doctoral student at the Mechanical Engineering Department at of Georgia Institute of Technology. His research interest includes smart manufacturing, AI based design optimization and predictive maintenance.

Young-Dae Shim received the B.S. degree in mechanical engineering from Ajou University, South Korea, in 2011, and the M.S. degree in mechanical convergence engineering from Hanyang University, South Korea, in 2019, and received his Ph.D. degree with the Department of Smart Fabrication Technology, Sungkyunkwan University, Republic of Korea. Currently, he is working as a Postdoctoral Fellow at Georgia Institute of Technology. His research interests include the development of smart sensing systems using continuum mechanics and electromagnetic-mechanical coupled physics.

Hyuck-kee Lee is a Senior Researcher at KATECH (Korea Automotive Technology Institute). His research interests are in Antenna Design, Antenna Unit, Atmospheric Attenuation, Field of View, High Frequency Structure Simulator, Highway System, Low-noise Amplifier, Microstrip Antenna, Numerical Simulations, Patch Antenna, Polarization Efficiency, Radiation Pattern.

Seong-Geun Shin is a Researcher at Center of Smart Driving Control R&D, Korea Automotive Technology Institute (KATECH). His research interests are in Emergency Braking, Advanced Driver Assistance Systems, Autonomous Emergency Braking, Autonomous Emergency Braking System, Autonomous Vehicles, Azimuth Angle, Braking System, Collision Risk, Collision Time, Computer Simulations, Data Fusion, Driver Model.

Seung-Kyum Choi received the Ph.D. degree in uncertainty quantification techniques for the analytical certification of complex engineered systems from the Mechanical and Materials Engineering Department, Wright State University, Dayton, OH, USA. He is currently the Director of the Center for Additive Manufacturing Systems (CAMS), Georgia Tech, Atlanta, GA, USA, where he has responsibilities for developing research and educational programs in additive manufacturing. directly began at Georgia Tech in Fall 2006 as an Assistant Professor. His reliability-based synthesis methods are critical in the development of new engineering components that are comprised of cellular material structures when those structures have inherent uncertainty associated with their geometry, loads, and material properties. His current research interests include developing robust simulation and decision-support tools that can help manage complex engineered systems that involve the design of additively manufactured parts and multifunctional products.