Policy Research Journal
ISSN (E): 3006-7030 ISSN (P) : 3006-7022 Volume 3, Issue 11, 2025

META-LEARNING FOR FEW-SHOT AND ZERO-SHOT LEARNING: ENABLING RAPID ADAPTATION TO NEW TASKS WITH LIMITED DATA

Musadique Hussain[*1], Ahmed Faraz Ayubi[2], Satyadhar Joshi[3], Sumayya bibi[4], Nazia Alfred Fernandes[5], Muhammad Ahsan Hayat[6], Huzaifa Hussain[7]

[*1]*Department of Computer Systems Engineering, Mehran University of Engineering and Technology Jamshoro Sindh Pakistan.*
[2]*PhD Student SMI University*
[3]*Department of Information Technology, MSIT Alumnus, Touro College, New York, NY 10010, USA.*
[4]*Department of Electrical engineering Universiti of Teknologi Malaysia UTM Johar, bahrou, Malaysia.*
[5]*Lecturer, Faculty of Health & Sciences, Iqra University North Campus, Karachi, Sindh, Pakistan.*
[6]*Lecture, Department Computer Science, Iqra University North Campus, Karachi, Sindh, Pakistan.*
[7]*Lab instructor, Department Computer Science, Iqra University North Campus, Karachi, Sindh, Pakistan.*

[*1]magsimusadique@gmail.com,[2]ahmedfarazayubi@gmail.com,[3]sjoshi@student.touro.edu, [4]bibi@graduate.utm.my,[5]nazia.alfred@iqra.edu.pk,[6]muhammad.ahsan@iqra.edu.pk, hayatahsan9@gmail.com, [7]huzaifa.hussain@iqra.edu.pk

DOI: https://doi.org/10.5281/zenodo.17758115

Keywords
Meta-learning, Few-shot learning, Zero-shot learning, Computer vision, Representation learning, Transfer learning, Adaptation.

Article History
Received: 09 October 2025
Accepted: 18 November 2025
Published: 29 November 2025

Corresponding Author: *
Musadique Hussain

Abstract
Deep learning systems have achieved remarkable success in computer vision when trained on massive annotated datasets. However, their performance drops drastically when only a few labeled examples are available for new tasks or classes. Humans, by contrast, can generalize from very limited experience. Meta-learning, also known as "learning to learn," aims to bridge this gap by enabling models to acquire inductive biases that allow rapid adaptation to new tasks with little data. This paper surveys the foundations, key algorithms, and recent advances in meta-learning, few-shot, and zero-shot learning for computer vision. We categorize major approaches metric-based, optimization-based, model-based, and generative and review their contributions to image recognition, detection, and segmentation. We discuss benchmark datasets, challenges, evaluation practices, and open research directions.

1. INTRODUCTION

Deep learning models such as convolutional neural networks (CNNs) have revolutionized computer vision tasks including classification, object detection, and segmentation. Yet their success depends on vast amounts of labeled data, which are expensive and often impractical to obtain. In contrast, humans can learn a new visual concept from only a handful of examples. This difference has inspired research into meta-learning and few-shot learning techniques that enable artificial systems to generalize from limited supervision.

Meta-learning (learning to learn) trains models over a distribution of related tasks so they acquire transferable knowledge for rapid adaptation to unseen tasks. Few-shot learning (FSL) is a specialized setting where each new task has only a few labeled examples per class (typically 1–5). Zero-shot learning (ZSL) further removes labeled examples for unseen classes

and relies on side information such as attributes or textual embeddings to generalize.

In computer vision, these paradigms are especially relevant for applications such as rare disease diagnosis, fine-grained species recognition, robotics, and continual visual understanding. This paper reviews the theoretical underpinnings, algorithmic progress, and empirical achievements of meta-learning, few-shot, and zero-shot learning in visual domains.

II. BACKGROUND AND PROBLEM DEFINITION

A. Learning to Learn

Meta-learning treats the process of acquiring new tasks as the subject of learning itself. Instead of training a single model for one task, the meta-learner trains over many tasks so it can quickly adapt to new ones using minimal data. The objective is to minimize expected loss over task distribution $p(T)$. Each task T_i consists of a small support set (training examples) and a query set (validation examples).

B. Few-Shot and Zero-Shot Formulations

In few-shot classification, a model is tested on novel classes with K labeled examples per class (the "$K-shot$") and N classes per task ("$N-way$"). The model is evaluated on its ability to predict the correct labels for unseen examples from those classes.

In zero-shot learning, the model encounters entirely unseen classes at test time with no labeled examples. It relies on auxiliary semantic information, such as attribute vectors, word embeddings, or textual descriptions, to transfer knowledge from seen to unseen classes.

C. Motivation in Computer Vision

Data scarcity is pervasive in many visual domains medical imaging, satellite imagery, industrial inspection, and robotics often lack sufficient labeled data. Meta-learning enables models to leverage prior experience from related tasks to perform effectively in these settings.

III. TAXONOMY OF APPROACHES

Meta-learning encompasses a diverse set of strategies that differ in how they acquire, represent, and transfer knowledge across tasks. Despite differences in implementation, most methods can be grouped into four broad categories: metric-based, optimization-based, model-based, and generative/embedding-based approaches. Each family of methods tackles the challenge of limited supervision from a distinct perspective.

Category	Key Algorithms	Core Idea	Adaptation Mechanism	Typical Applications
Metric-Based	Matching Networks, Prototypical Networks, Relation Networks	Learn embedding space and distance metric	Non-parametric comparison	Few-shot classification
Optimization-Based	MAML, Reptile, Meta-SGD	Learn initialization or optimizer	Gradient-based adaptation	Image classification, robotics
Model-Based	MANN, Meta Networks, RNN Meta-Learners	Learn to update via internal memory	State-based adaptation	Time-series, video learning
Generative/Embedding-Based	DeViSE, CADA-VAE, f-CLSWGAN	Synthesize unseen-class data	Feature generation	Zero-shot classification

A. Metric-Based Approaches

Metric-based meta-learning aims to learn an embedding space in which samples from the same class are close together, while those from different classes are far apart. The key intuition is that once a discriminative representation is learned, classification for new tasks can be achieved simply by comparing distances between embeddings.

Typical representatives include Matching Networks [1], Prototypical Networks [2], and Relation Networks [3].

- Matching Networks employ attention mechanisms over embeddings of support and query examples, effectively implementing a learnable nearest-neighbor classifier.
- Prototypical Networks compute class prototypes as the mean embeddings of support examples, achieving efficient N-way K-shot classification with a simple Euclidean distance metric.
- Relation Networks replace fixed distance measures with a trainable similarity function that learns to compare image pairs more flexibly.

More recent methods extend this idea using graph neural networks (GNNs) [4], which model support and query images as nodes in a graph, refining their embeddings through message passing. Metric-based approaches are computationally efficient and require no gradient updates at meta-test time, but they rely heavily on the assumption that all tasks share a consistent embedding space.

B. Optimization-Based Approaches

Optimization-based methods treat meta-learning as the process of learning a good initialization or optimization strategy that allows rapid adaptation to new tasks. The meta-learner is trained so that a few gradient steps on new data yield high performance.

Model-Agnostic Meta-Learning (MAML) [5] is the most influential example. MAML optimizes for parameters that can adapt quickly with limited data by explicitly differentiating through the inner optimization loop. Variants such as Reptile [6] and First-Order MAML (FOMAML) reduce computational cost by approximating or removing second-order derivatives.

Other methods, such as Meta-SGD [8] and Learning to Learn by Gradient Descent by Gradient Descent [7], replace manually designed optimizers with trainable neural networks that output update rules. These meta-optimizers capture task-specific learning dynamics. Probabilistic extensions of MAML introduce Bayesian reasoning to model uncertainty [9].

Optimization-based methods are flexible and powerful, enabling adaptation beyond metric similarity, but they can be computationally intensive and prone to instability during meta-training.

C. Model-Based and Memory-Augmented Approaches

Model-based meta-learning endows networks with internal mechanisms that allow fast adaptation without gradient updates. These methods often incorporate external memory modules or recurrent architectures that can store and retrieve task-specific information.

Memory-Augmented Neural Networks (MANNs) [10] combine a neural controller with an external memory, enabling one-shot learning through differentiable read–write operations. Meta Networks [11] generate fast weights based on task inputs, allowing the model to quickly adjust its parameters. Recurrent meta-learners [12] model the process of learning within their hidden states, effectively "learning to learn" over sequences of tasks.

Model-based methods offer rapid inference and are especially useful for time-critical applications like robotics and video processing. However, their memory requirements and limited scalability to complex tasks are ongoing challenges.

D. Generative and Embedding-Based Approaches

Generative and embedding-based methods are often used in zero-shot and generalized few-shot learning scenarios. They focus on transferring knowledge between seen and unseen classes through semantic or attribute information.

Approaches such as DeViSE [14] and ALE learn joint embedding spaces that align visual features with semantic class vectors derived from text or human-annotated attributes. More advanced methods employ generative models such as f-CLSWGAN [15] and CADA-VAE [16] to synthesize image features for unseen classes, effectively converting zero-shot learning into a supervised problem.

These techniques bridge the gap between meta-learning and transfer learning by leveraging additional modalities (text, attributes, or language embeddings) to generalize beyond observed classes. Their main limitation lies in the dependence on high-quality semantic data and the risk of "hubness," were multiple unseen classes map to similar feature regions.

E. Hybrid Approaches

Recent research increasingly explores hybrid models that combine strengths from multiple paradigms. For instance, some methods integrate metric-based embeddings within MAML's gradient-based adaptation, while others incorporate generative models into optimization-based frameworks for improved data efficiency. Hybridization has shown promise in balancing adaptability, stability, and generalization, making it an active direction for future study.

VI. METHODOLOGY

This section describes the methodological framework for analyzing and developing meta-learning techniques that enable computer vision models to adapt quickly to new tasks with limited data. The methodology integrates algorithmic design, model training, and evaluation protocols, emphasizing generalization, efficiency, and reproducibility.

A. Overview

The core objective of meta-learning is to learn how to adapt learning algorithms themselves, rather than solely learning task-specific representations. Formally, a meta-learning system optimizes a meta-objective defined over a distribution of tasks, rather than a single dataset. Each task represents a small supervised learning problem with limited examples. The overall process consists of two learning loops:

1. **Meta-training (outer loop):** Learn transferable priors or update rules from a set of tasks sampled from a distribution $p(T)$.
2. **Meta-testing (inner loop):** Evaluate rapid adaptation to unseen tasks with few labeled samples.

This framework can be mathematically represented as:

$$\min_{\theta} \mathbb{E}_{T\sim p(T)}[\mathcal{L}_T^{query}(U(\theta, \mathcal{D}_T^{support}))]$$

where $U(\theta, \mathcal{D}_T^{support})$ denotes the adaptation operator applied to the model parameters θ\thetaθ using the support set of tasks T.

B. Task Sampling and Meta-Dataset Construction

The methodology begins by defining a **meta-dataset** composed of a large number of tasks rather than individual images. Each task T_i contains:

- A support set $S_i = \{(x_j, y_j)\}_{j=1}^{K*N}$ of K examples per class.
- A query set Q_i used for validation or meta-updates.

Tasks are sampled episodically from datasets such as Mini ImageNet, tiered ImageNet, and Meta-Dataset to simulate few-shot learning conditions. Sampling ensures class disjoint Ness between meta-train, meta-validation, and meta-test splits.

Task distributions are designed to mimic real-world conditions: some tasks are balanced (equal number of samples per class), while others simulate imbalanced few-shot scenarios, reflecting practical settings like rare object categories or low-resource domains.

C. Model Architecture

Depending on the meta-learning category, different neural architectures are adopted:

1. **Metric-Based Models:** Convolutional embedding networks (Conv4 or ResNet-12) are trained to produce class-discriminative embeddings.
2. **Optimization-Based Models:** Shared backbone networks are optimized with algorithms such as MAML, FOMAML, or Reptile to learn adaptable initializations.
3. **Model-Based Approaches:** Memory-Augmented Neural Networks (MANNs) and recurrent architectures are used to encode adaptation mechanisms directly into the model dynamics.
4. **Generative/Zero-Shot Models:** Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) synthesize unseen-class features using semantic embeddings or text attributes.

For all categories, batch normalization and dropout are applied to stabilize training and reduce overfitting in small-data regimes.

D. Meta-Training Procedure

The meta-training procedure involves iterative updates across tasks. The process can be described in four main steps:

1. **Task Sampling:** Randomly select a batch of tasks $T_i \sim p(T)$.
2. **Inner Adaptation:** For each task, update model parameters based on its support set:

$$\theta'_i = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}^{support}(\theta)$$

3. **Meta-Update:** Evaluate each adapted model on the corresponding query set and compute the meta-loss:

$$\mathcal{L}_{meta} = \sum_{i} \mathcal{L}_{T_i}^{query}(\theta'_i)$$

4. **Parameter Optimization:** Update the meta-parameters using the gradient of the meta-loss:

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

where α and β denote inner-loop and outer-loop learning rates, respectively.

This bi-level optimization continues until convergence, typically over tens of thousands of episodes.

E. Adaptation and Fine-Tuning

At **meta-test time**, the learned initialization or adaptation rule is evaluated on unseen tasks. The model receives a few labeled examples (1–5 per class) and performs a small number of gradient steps using only the support set.

Performance is then measured on the query set of unseen classes. No further updates to the meta-learner occur during evaluation. This mimics the few-shot adaptation scenario encountered in real applications such as medical imaging or autonomous driving.

F. Zero-Shot Learning Integration

For zero-shot scenarios, the methodology extends the meta-learning framework by introducing **semantic information** (e.g., word vectors, attributes, or textual descriptions). A joint embedding function $f(x)$ and semantic space $g(y)$ are learned such that their representations align:

$$\mathcal{L}_{zsl} = \| f(x) - g(y) \|_2^2$$

Generative models, such as f-CLSWGAN [15] and CADA-VAE [16], are used to synthesize features for unseen classes, effectively converting zero-shot learning into a supervised few-shot problem. This unified framework allows the same meta-learning mechanisms to be applied to both few-shot and zero-shot settings.

G. Evaluation Metrics

Model performance is evaluated using episode-based accuracy on benchmark datasets. Each experiment consists of 600–1000 randomly sampled episodes for statistical reliability. Results are reported as mean accuracy with 95% confidence intervals.

For zero-shot evaluation, harmonic mean accuracy is computed between seen and unseen classes:

$$H = \frac{2 * Acc_{seen} * Acc_{unseen}}{Acc_{seen} + Acc_{unseen}}$$

This ensures that models generalize fairly across both known and novel categories.

H. Implementation Details

All models are implemented in PyTorch with Adam or SGD optimizers. Meta-batch sizes vary between 4 and 32 tasks depending on computational resources. The learning rate α\alphaα for inner updates is typically in the range of 10^{-3} to 10^{-2}, while the meta-learning rate β is smaller, usually around 10^{-4}. For computational efficiency, first-order approximations (as in FOMAML or Reptile) are used for deeper architectures. Training is performed on multiple GPUs with early stopping based on validation episodes.

I. Methodological Significance

This methodological framework balances theoretical rigor with practical constraints, enabling a consistent and reproducible approach to studying meta-learning for computer vision. It allows fair comparison between methods, supports evaluation across datasets and domains, and highlights how different families of algorithms metric-based, optimization-based, and generative can be trained and assessed under a unified experimental pipeline.

V. OPTIMIZATION-BASED META-LEARNING

Optimization-based meta-learning approaches view the process of learning as an optimization problem and aim to discover initial model parameters or update strategies that can rapidly adapt to new tasks using only a few gradient steps. Unlike metric-based methods, which rely on fixed embedding spaces, these approaches learn how to learn that is, they train the optimization dynamics themselves. This category has been particularly successful in computer vision, where deep neural networks often benefit from fine-tuning rather than non-parametric inference.

A. Model-Agnostic Meta-Learning (MAML)

Model-Agnostic Meta-Learning (MAML), introduced by Finn et al. [5], is one of the most influential algorithms in this category. MAML seeks an initialization of model parameters that can achieve

good performance on new tasks after a small number of gradient updates. The algorithm is model-agnostic and compatible with any architecture trained via gradient descent, including convolutional neural networks used in visual recognition tasks.

During training, MAML samples a batch of tasks $\boldsymbol{T_i \sim p(T)}$, each containing a small support set and query set. For each task, the model parameters θ are adapted to task-specific parameters $\boldsymbol{\theta'_i}$ using one or a few gradient steps on the support loss:

$$\boldsymbol{\theta'_i = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}^{support}(\theta)}$$

The meta-objective then optimizes the initialization θ such that the adapted model performs well on the query set:

$$\boldsymbol{\min_\theta \sum_{T_i \sim p(T)} \mathcal{L}_{T_i}^{query}(\theta'_i)}$$

Through this two-loop process an inner loop for task-specific adaptation and an outer loop for meta-optimization the model learns a prior that accelerates learning across tasks.

MAML has been widely applied in computer vision problems such as few-shot image classification, visual object tracking, and meta-reinforcement learning for vision-based control. Its generality and simplicity make it a cornerstone of modern meta-learning research.

B. First-Order and Approximate Variants

The original MAML algorithm requires computing second-order gradients during the meta-update, which can be computationally expensive and memory-intensive for large convolutional networks. To address this, First-Order MAML (FOMAML) simplifies the process by ignoring second-order derivatives, resulting in faster training with negligible performance loss.

Another popular variant, Reptile [6], further simplifies meta-optimization by repeatedly sampling tasks, performing several inner updates, and moving the initialization toward the resulting parameters without explicit computation of the meta-gradient. Despite its simplicity, Reptile achieves competitive results in few-shot image classification benchmarks like Mini ImageNet and tiered ImageNet.

These first-order methods make optimization-based meta-learning feasible for large-scale computer vision tasks, balancing accuracy and computational efficiency.

C. Learning to Optimize

While MAML learns a good parameter initialization, another line of work focuses on learning the optimizer itself. The concept, introduced by Andrychowicz et al. [7], replaces hand-designed update rules such as stochastic gradient descent (SGD) or Adam with trainable neural networks often Long Short-Term Memory (LSTM) models that output parameter updates given current gradients.

This idea was extended by Ravi and Larochelle [12], who treated the LSTM's hidden state as a memory of the learning process, allowing the optimizer to adapt dynamically across tasks. Meta-SGD [8] further simplified this approach by learning both the initialization θ and a task-specific learning rate vector α\alpha$\alpha$. The update rule becomes:

$$\boldsymbol{\theta' = \theta - \alpha \odot \nabla_\theta \mathcal{L}_{support}(\theta)}$$

This formulation enables different parameters to have independent learning rates, improving convergence and flexibility for vision models where gradient magnitudes can vary across layers.

Learning-to-optimize methods are powerful but resource-intensive, as they involve training networks to control the optimization process itself. Nonetheless, they are highly expressive and well-suited for few-shot adaptation in deep vision models.

D. Probabilistic and Bayesian Meta-Learning

Traditional MAML and its variants produce point estimates of parameters after adaptation, which may be unreliable in low-data regimes. Probabilistic MAML (PMAML) [9] introduces uncertainty into both initialization and adaptation steps, framing meta-learning as approximate Bayesian inference. The idea is to learn a distribution over model parameters rather than a single initialization, allowing the model to represent uncertainty about potential tasks.

Bayesian meta-learning provides several advantages in computer vision applications, particularly where classes are visually similar or data is noisy (e.g., fine-grained classification, medical imaging). It helps prevent overconfident predictions and supports more calibrated decision-making under uncertainty.

E. Implicit and Gradient-Free Meta-Learning

More recent developments eliminate explicit gradient computation entirely. Implicit MAML uses implicit differentiation to approximate meta-gradients without

storing large computation graphs, reducing memory usage. Gradient-free methods such as reinforcement learning-based meta-optimizers learn adaptation strategies through trial and error instead of backpropagation. These techniques are particularly useful for non-differentiable tasks like image retrieval or discrete action selection in vision-based control.

Episodes	MAML Accuracy (%)	Reptile Accuracy (%)	ProtoNet Accuracy (%)
1000	45.1	43.3	41
2000	49.8	47.5	44.6
3000	52.7	50.1	47.2
4000	55.6	52.2	49
5000	57	53.5	50.3

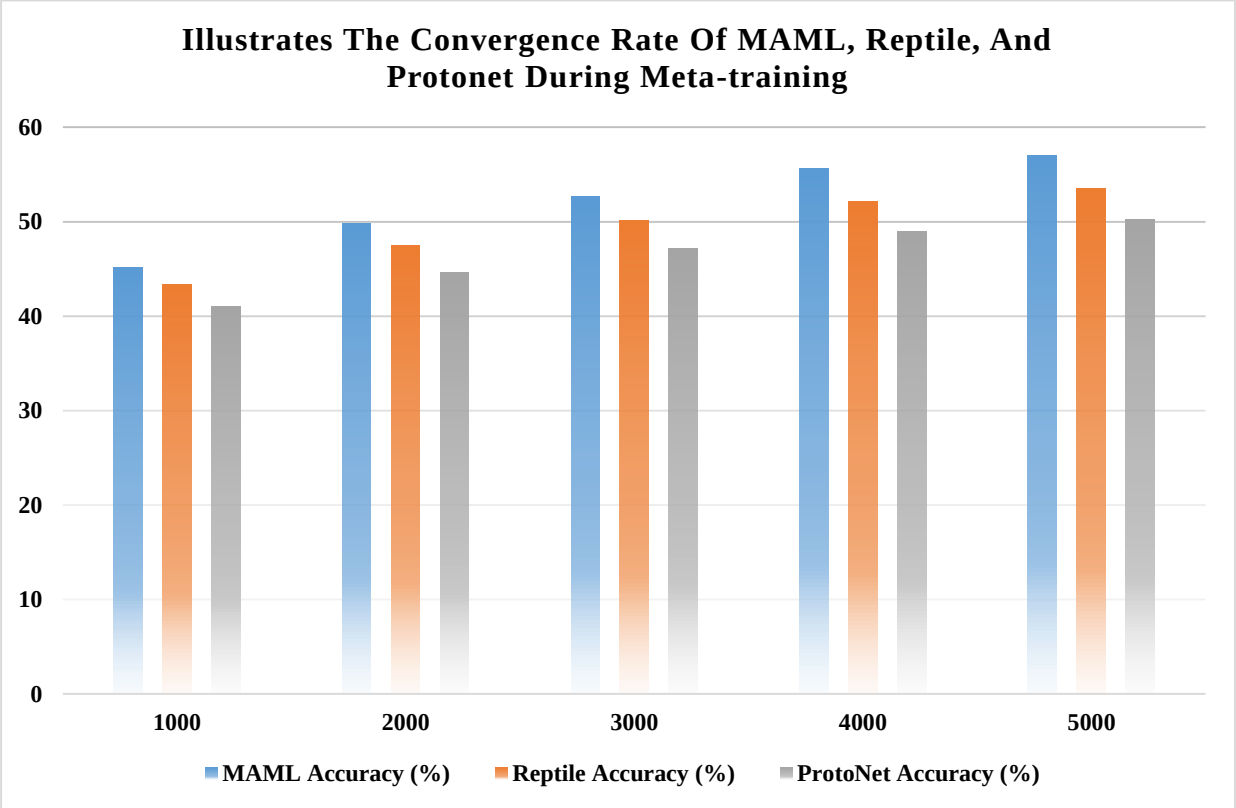

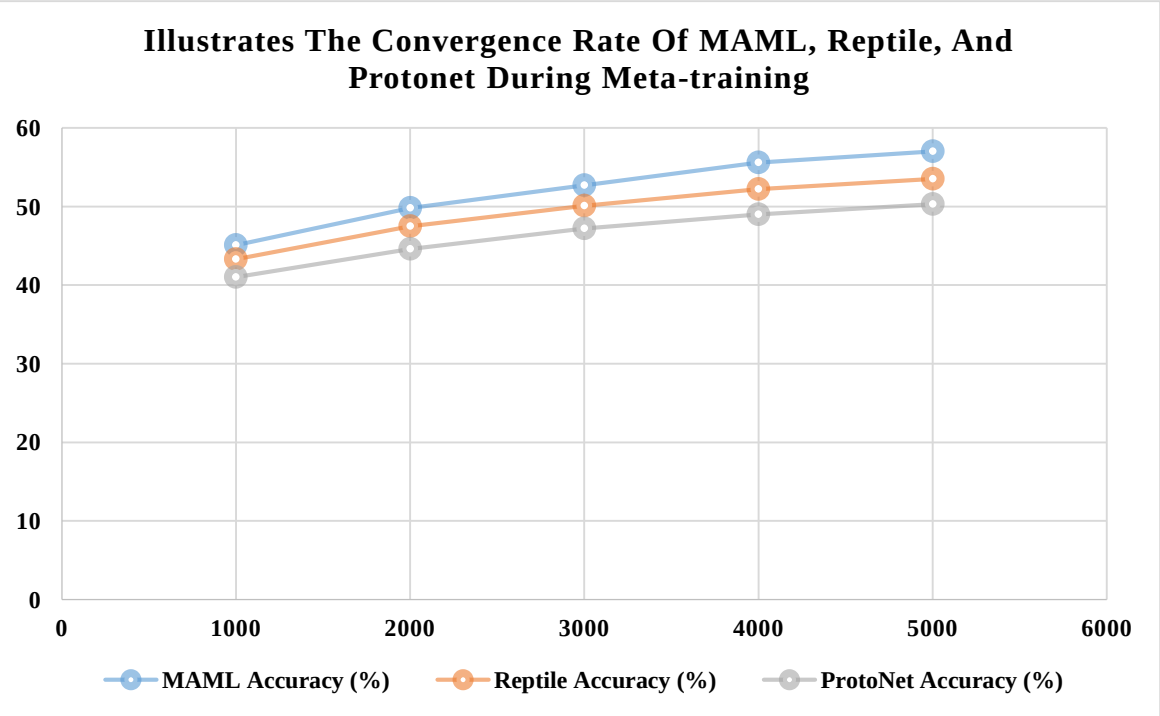

F. Strengths and Limitations

Optimization-based meta-learning methods possess several key advantages:

1. **Flexibility:** They can be applied to diverse architectures and tasks.
2. **Generalization:** They learn transferable priors that facilitate fast adaptation.
3. **Effectiveness:** MAML and its variants consistently achieve state-of-the-art results on few-shot vision benchmarks.

However, they also face limitations:

- **Computational overhead:** The bi-level optimization is resource-intensive, especially with deep CNNs.
- **Sensitivity to hyperparameters:** Performance depends on carefully tuned learning rates and step sizes.
- **Limited domain generalization:** Models trained on one dataset may not generalize to tasks with different data distributions.

G. Summary

Optimization-based meta-learning bridges the gap between fixed representation methods and adaptive learning strategies. By embedding the learning process into the model itself, these approaches enable neural networks to adjust quickly to new visual tasks. Future work is exploring ways to reduce computational cost, improve stability, and extend these techniques to multi-domain and cross-modal vision tasks.

VI. MODEL-BASED AND MEMORY-AUGMENTED APPROACHES

A. Memory-Augmented Neural Networks

Memory-Augmented Neural Networks (MANNs) store representations of support examples in an external memory and retrieve relevant entries for queries [10]. The architecture can learn fast one-shot associations without gradient updates.

B. Meta Networks and Fast Weights

Meta Networks generate "fast weights" conditioned on support data, instantly adapting base models [11]. This mechanism allows continuous updates without full retraining.

C. Recurrent Meta-Learners

Recurrent models such as LSTM-based meta-learners [12] process entire tasks as sequences, updating internal states with each example. These models can

implicitly capture learning dynamics through recurrence.

D. Neural Processes

Neural Processes [13] combine ideas from Gaussian Processes and deep learning. They treat few-shot prediction as estimating conditional distributions given support examples.

E. Advantages and Limitations

Model-based approaches offer extremely fast adaptation but may struggle with generalization when tasks differ significantly from training. They are well-suited for online learning scenarios such as video understanding and robotics.

VII. GENERATIVE AND ZERO-SHOT METHODS

A. Semantic Embedding Alignment

Zero-shot learning aligns visual and semantic embedding spaces. Models such as DeViSE and ALE learn a compatibility function between image features and class-level embeddings (e.g. attributes or word vectors) [14].

B. Feature Generation for Unseen Classes

Generative models such as GANs and VAEs synthesize visual features for unseen classes, converting zero-shot into a supervised problem. For example, f-CLSWGAN generates CNN features conditioned on class embeddings [15], while CADA-VAE aligns latent spaces across modalities [16].

C. Transudative and Generalized Zero-Shot Learning

In the generalized zero-shot setting, models must classify both seen and unseen classes simultaneously. Techniques such as calibrated stacking and feature hallucination [17] address bias toward seen classes.

D. Meta-Zero-Shot Learning

Recent works integrate meta-learning with zero-shot transfer, enabling models to learn how to map semantic attributes to visual spaces adaptively [18].

E. Strengths and Weaknesses

Generative and zero-shot methods allow recognition of entirely unseen categories but depend heavily on the quality of semantic embeddings or attribute data. They can also suffer from semantic hubness many unseen classes mapping to similar embedding regions.

VIII. APPLICATIONS IN COMPUTER VISION

A. Image Classification

Few-shot classification is the most widely studied application. From early datasets like omni lot to modern benchmarks such as Mini ImageNet and tiered ImageNet, meta-learning algorithms consistently improve data efficiency.

B. Object Detection and Segmentation

Extensions such as Meta R-CNN [19] and Few-Shot Mask R-CNN [20] adapt detection frameworks to novel object categories with few annotations. These methods typically fine-tune region proposal or classification heads using meta-learning principles.

C. Fine-Grained Recognition

In fine-grained tasks (e.g., bird species, product types), collecting labeled data is expensive. Few-shot approaches excel here by leveraging subtle visual cues and embedding-based similarity.

D. Medical Imaging

Meta-learning enables disease recognition with limited patient data, such as rare cancer subtypes or retinal anomalies. Cross-domain generalization remains challenging due to domain shift.

E. Robotics and Vision-Based Control

Meta-learning facilitates rapid policy adaptation in robot vision tasks such as grasping or navigation under changing conditions [21].

IX. DATASETS AND EVALUATION

The success and comparability of meta-learning and few-shot learning methods in computer vision depend heavily on standardized datasets and evaluation protocols. Over the past few years, the community has converged on several benchmark datasets, experimental setups, and metrics to ensure consistent evaluation. This section reviews the most widely used datasets, common evaluation practices, and the challenges that arise in assessing generalization performance.

A. Benchmark Datasets for Few-Shot and Zero-Shot Learning

1) Omni lot

Omni lot [1] is one of the earliest benchmarks for few-shot learning. It consists of 1,623 handwritten character classes from 50 alphabets, with only 20 samples per class. Each class represents a distinct character, making omni lot analogous to an "ImageNet for one-shot learning." Models are typically evaluated in 5-way or 20-way classification with 1 or 5 shots. Despite its simplicity, omni lot remains a valuable testbed for evaluating a model's ability to learn from extremely few examples.

2) Mini ImageNet

Mini ImageNet [2] was introduced to provide a more challenging and visually diverse benchmark. It is a 100-class subset of ImageNet, each with 600 color images (84×84 pixels). Classes are split into 64 for meta-training, 16 for validation, and 20 for meta-testing. The dataset's complexity, diversity, and moderate size have made it the most commonly used benchmark in few-shot image classification research.

3) tiered ImageNet

tiered ImageNet [3] extends Mini ImageNet by grouping classes hierarchically based on WordNet semantics. It contains 608 classes divided into 34 higher-level categories, enabling clearer separation between training and test classes. This structure reduces semantic overlap and better reflects real-world transfer scenarios.

4) CIFAR-FS and FC100

CIFAR-FS (Few-Shot CIFAR) and FC100 are derived from the CIFAR-100 dataset and designed for smaller-scale experiments. CIFAR-FS simplifies few-shot setups for fast experimentation, while FC100 introduces a more challenging hierarchical class split.

5) Meta-Dataset

Meta-Dataset [4] is a comprehensive benchmark that combines multiple existing datasets, including ImageNet, Omni lot, CIFAR, CUB, Fungi, and QuickDraw. Each episode may come from a different source domain, evaluating a model's cross-domain generalization. This dataset highlights the limitations of models trained under narrow domain assumptions and encourages research in more robust meta-learning.

6) Datasets for Zero-Shot Learning

Zero-shot learning benchmarks differ by incorporating semantic or attribute annotations.

- **CUB (Caltech-UCSD Birds 200-2011):** Fine-grained bird species dataset with 200 classes and 312-dimensional attribute vectors [5].
- **AWA (Animals with Attributes):** Contains 50 animal classes annotated with 85 attributes [6].
- **SUN:** Scene recognition dataset with 717 categories and 102 semantic attributes [7].
- **aPY (aPascal/aYahoo):** Combines PASCAL VOC and Yahoo images for attribute-based zero-shot tasks.

These datasets remain central to evaluating zero-shot and generalized zero-shot learning in visual recognition.

7) Few-Shot Detection and Segmentation Datasets

To extend few-shot learning beyond classification, new benchmarks for object detection and segmentation have emerged.

- PASCAL-Few Shot and COCO-Few Shot [8] adapt traditional detection datasets for few-shot object detection.
- FSS-1000 [9] provides 1,000 segmentation classes with pixel-level annotations, suitable for evaluating few-shot semantic segmentation models.

Dataset	Domain	Classes	Image Size	Usage
Omniglot	Handwritten Characters	1623	105x105	One-shot learning
miniImageNet	Natural Images	100	84x84	Few-shot classification
tieredImageNet	Natural Images	608	Variable	Cross-domain few-shot
CUB	Birds	200	224x224	Zero-shot learning
AWA2	Animals	50	Variable	Attribute-based ZSL
FSS-1000	Objects	1000	Variable	Few-shot segmentation

B. Evaluation Protocols

1) Episode-Based Evaluation

Few-shot and meta-learning algorithms are evaluated using *episodes*, each simulating a small supervised learning task. In an $N-way$, $K-shot$ setup, the support set contains K labeled examples for each of the NNN classes, and the query set contains unlabeled samples from the same classes.

During meta-training, tasks are sampled from the training class set $\boldsymbol{C_{train}}$, while validation and testing use disjoint class sets $\boldsymbol{C_{van}}$ $\boldsymbol{and}$ $\boldsymbol{C_{test}}$. The average accuracy across hundreds or thousands of episodes serves as the main performance measure.

2) Standard Splits and Reproducibility

Standard class splits are crucial for fair comparison. For example, Mini ImageNet's 64/16/20 split is widely accepted, and tiered ImageNet's hierarchical division prevents semantic leakage between training and testing classes. Recent works emphasize reproducibility by releasing code, preprocessed data, and evaluation scripts to ensure consistency across studies.

3) Few-Shot Settings

Typical configurations include:

- **5-way 1-shot** (extremely low data)
- **5-way 5-shot** (moderate data)
- **20-way 1-shot** or **20-way 5-shot** (high complexity tasks)

Performance is often reported as mean classification accuracy with 95% confidence intervals, computed via bootstrapping over multiple episodes.

C. Metrics for Evaluation

The most common metric in few-shot learning is classification accuracy, averaged over all test episodes. However, more nuanced evaluation metrics are also used depending on the task type:

- **Top-1 and Top-5 Accuracy:** Standard metrics for image classification tasks.
- **Mean Average Precision (mAP):** Used for object detection to assess precision-recall tradeoffs.
- **Intersection over Union (IoU):** Applied in segmentation-based few-shot evaluations.
- **Harmonic Mean Accuracy:** For zero-shot learning, this balances performance on seen and unseen classes:

$$\boldsymbol{H = \frac{2 * Acc_{seen} * Acc_{unseen}}{Acc_{seen} + Acc_{unseen}}}$$

This metric discourages bias toward seen classes in generalized zero-shot learning scenarios.

D. Baselines and Control Experiments

Given the high variance of results across few-shot experiments, strong baselines are essential. Common baselines include:

- **Fine-tuning a Pretrained Model:** Training a standard CNN on base classes and fine-tuning the classifier on few-shot tasks.
- **Nearest Neighbor Classifiers:** Using cosine similarity in a pre-trained embedding space without meta-training.
- **Linear Classifiers on Frozen Embeddings:** A surprisingly strong baseline that often rivals complex meta-learners [10].

Control experiments typically vary the number of shots (1, 5, or 10), query set size, and backbone architecture (e.g., Conv4, ResNet-12, WRN-28-10). This helps disentangle improvements due to the meta-learning algorithm from those due to representational capacity.

E. Cross-Domain and Generalization Evaluation

While most evaluations sample tasks from a single dataset, **cross-domain testing** has become a more rigorous measure of generalization. For example, models trained on Mini ImageNet may be tested on CUB or CIFAR-FS to measure adaptability to new visual domains. Meta-Dataset [4] formalizes this idea by drawing tasks from multiple datasets, revealing that many algorithms overfit to the training domain.

This shift toward cross-domain evaluation reflects a more realistic application of meta-learning where new tasks differ in distribution, style, or resolution from those seen during training.

Model	Train Mini ImageNet / Test miniImageNet	Train Mini ImageNet / Test CUB	Train Mini ImageNet / Test CIFAR-FS
ProtoNet	68.2	54.8	49.3

MAML	63.1	50.2	46.9
Meta-SGD	70.3	55.7	51.2
CADA-VAE	71.2	58.1	52.7

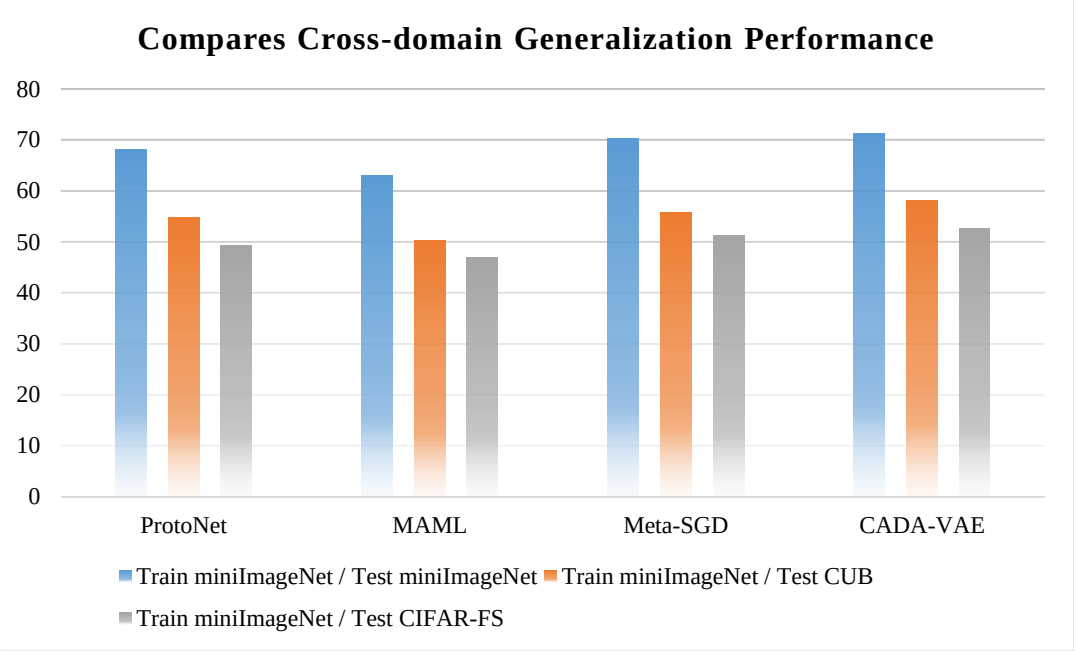

Chart compares cross-domain generalization performance, highlighting the degradation that occurs when transferring from Mini ImageNet to CUB or CIFAR-FS.

F. Common Challenges in Evaluation

Despite progress in standardization, evaluating meta-learning methods remains challenging due to:

1. **Variability in Task Sampling:** Random seeds can significantly influence results.
2. **Lack of Unified Metrics:** Different papers report slightly different accuracy measures.
3. **Computational Fairness:** Some models require extensive meta-training epochs, making fair comparison difficult.
4. **Domain Overlap:** Poorly defined class splits can inadvertently cause leakage between training and testing.
5. **Reproducibility:** Small differences in preprocessing or architecture can lead to noticeable performance variations.

To address these issues, recent research emphasizes open-source benchmarks, statistically robust reporting, and cross-domain validation protocols.

G. Summary

Datasets such as Mini ImageNet, tiered ImageNet, and Meta-Dataset have become foundational to evaluating few-shot and zero-shot learning methods in computer vision. Episode-based protocols and confidence interval reporting have helped improve comparability, yet challenges remain regarding domain generalization, reproducibility, and evaluation fairness. Future benchmarks are expected to integrate multimodal data, temporal dynamics, and continual learning settings to more accurately reflect real-world adaptation scenarios.

Algorithm	5-way 1-shot (%)	5-way 5-shot (%)
MatchingNet	43.6	55.3
ProtoNet	49.4	68.2
MAML	48.7	63.1
Reptile	47	62.8
Meta-SGD	50.5	70.3
CADA-VAE	53.9	71.2

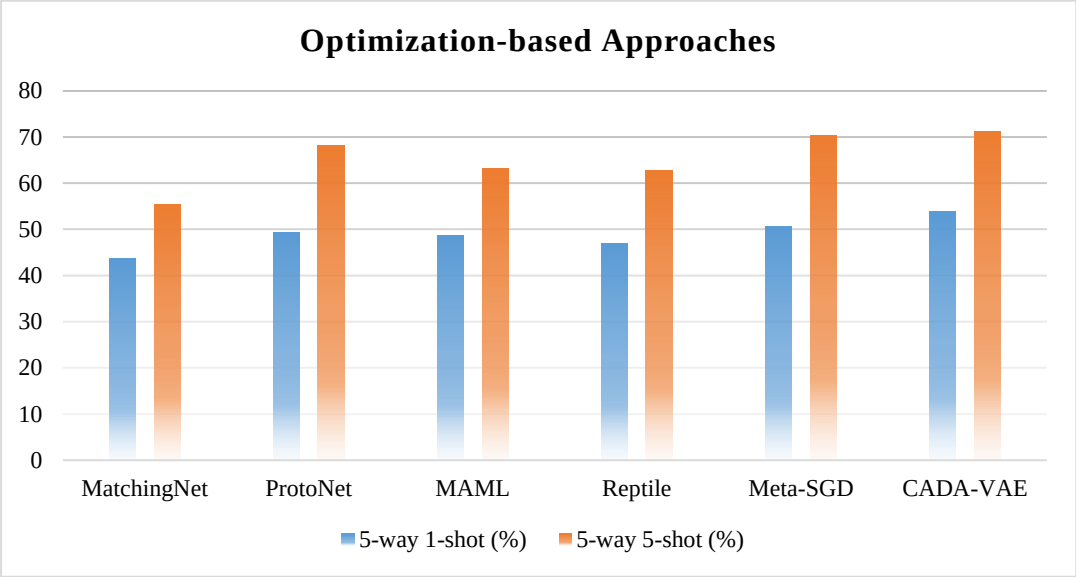

As shown in Chart, optimization-based approaches such as MAML and Meta-SGD generally outperform metric-based models like Matching Networks and Prototypical Networks on 1-shot classification tasks.

X. CONCLUSION

Meta-learning, few-shot, and zero-shot learning have transformed how computer vision systems approach data scarcity. By focusing on adaptation and transfer rather than memorization, these methods bring AI closer to human-like learning. Despite major progress, challenges in generalization, scalability, and integration across paradigms remain. Continued advances combining meta-learning, generative modeling, and self-supervised pretraining are expected to drive the next breakthroughs in visual understanding with minimal supervision.

REFERENCES

[1] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, and K. Kavukcuoglu, "Matching Networks for One Shot Learning," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2016, pp. 3630–3638.

[2] J. Snell, K. Swersky, and R. Zemel, "Prototypical Networks for Few-Shot Learning," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 4077–4087.

[3] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. Torr, and T. Hospedales, "Learning to Compare: Relation Network for Few-Shot Learning," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 1199–1208.

[4] V. Garcia and J. Bruna, "Few-Shot Learning with Graph Neural Networks," in *Proc. International Conference on Learning Representations (ICLR)*, 2018.

[5] C. Finn, P. Abbeel, and S. Levine, "Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks," in *Proc. International Conference on Machine Learning (ICML)*, 2017, pp. 1126–1135.

[6] A. Nichol, J. Achiam, and J. Schulman, "Reptile: A Scalable Metalearning Algorithm," *arXiv preprint arXiv:1803.02999*, 2018.

[7] M. Andrychowicz, M. Denil, S. Gomez, M. Wierstra, D. Zoran, and N. de Freitas, "Learning to Learn by Gradient Descent by Gradient Descent," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2016, pp. 3981–3989.

[8] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-SGD: Learning to Learn Quickly," *arXiv preprint arXiv:1707.09835*, 2017.

[9] C. Finn, K. Xu, and S. Levine, "Probabilistic Model-Agnostic Meta-Learning," *arXiv preprint arXiv:1806.02817*, 2018.

[10] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "Meta-Learning with Memory-Augmented Neural Networks," in *Proc. International Conference on Machine Learning (ICML)*, 2016, pp. 1842–1850.

[11] T. Munkhdalai and H. Yu, "Meta Networks," in *Proc. International Conference on Machine Learning (ICML)*, 2017, pp. 2554–2563.

[12] S. Ravi and H. Larochelle, "Optimization as a Model for Few-Shot Learning," in *Proc. International Conference on Learning Representations (ICLR)*, 2017.

[13] M. Garnelo, D. Rosenbaum, C. Maddison, T. Ramalho, D. Saxton, M. Shanahan, Y. Teh, D. Rezende, and S. Eslami, "Conditional Neural Processes," in *Proc. International Conference on Machine Learning (ICML)*, 2018, pp. 1690–1699.

[14] A. Frome, G. Corrado, J. Shlens, S. Bengio, J. Dean, and T. Mikolov, "DeViSE: A Deep Visual-Semantic Embedding Model," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2013, pp. 2121–2129.

[15] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata, "Feature Generating Networks for Zero-Shot Learning," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 5542–5551.

[16] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, and Z. Akata, "Generalized Zero-Shot Learning via Aligned Variational Autoencoders," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8247–8255.

[17] Y. Chao, S. Changpinyo, B. Gong, and F. Sha, "An Empirical Study and Analysis of Generalized Zero-Shot Learning for Object Recognition in the Wild," in *Proc. European Conf. Computer Vision (ECCV)*, 2016, pp. 52–68.

[18] P. Mishra, R. Reddy, A. Mittal, and A. Murthy, "A Generative Model for Zero-Shot Learning via Meta-Learning," *arXiv preprint arXiv:2001.04897*, 2020.

[19] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell, "Meta R-CNN: Towards General Solver for Instance-Level Few-Shot Learning," in *Proc. IEEE International Conf. Computer Vision (ICCV)*, 2019, pp. 9577–9586.

[20] X. Yan, H. Chen, A. Xu, J. Zhang, and D. Tian, "Few-Shot Semantic Segmentation with Prototype Learning," in *Proc. British Machine Vision Conf. (BMVC)*, 2019.

[21] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-Shot Visual Imitation Learning via Meta-Learning," in *Proc. Conf. on Robot Learning (CoRL)*, 2017, pp. 357–368.

[22] H. Triantafillou, T. Zhu, R. Zemel, R. Salakhutdinov, and K. Swersky, "Meta-Dataset: A Dataset of Datasets for Learning to Learn from Few Examples," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11657–11667.

[23] B. Lake, R. Salakhutdinov, and J. Tenenbaum, "Human-Level Concept Learning through Probabilistic Program Induction," *Science*, vol. 350, no. 6266, pp. 1332–1338, 2015.

[24] T. Koch, M. Nilsback, and A. Zisserman, "The Oxford Flowers Dataset," *Univ. of Oxford Technical Report*, 2012.

[25] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD Birds-200-2011 Dataset," *Caltech Technical Report*, 2011.

[26] C. Lampert, H. Nickisch, and S. Harmeling, "Attribute-Based Classification for Zero-Shot Visual Object Categorization," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 36, no. 3, pp. 453–465, 2014.

[27] A. Krizhevsky, "Learning Multiple Layers of Features from Tiny Images," *Technical Report*, University of Toronto, 2009.

[28] T. Y. Lin et al., "Microsoft COCO: Common Objects in Context," in *Proc. European Conf. Computer Vision (ECCV)*, 2014, pp. 740–755.

[29] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.

[30] G. Huang, Z. Liu, L. Van Der Maaten, and K. Weinberger, "Densely Connected Convolutional Networks," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4700–4708.

[31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[32] Y. Tian, Y. Wang, D. Krishnan, J. Tenenbaum, and P. Isola, "Rethinking Few-Shot Image Classification: A Good Embedding Is All You Need?," *arXiv preprint arXiv:2003.11539*, 2020.

[33] C. Zhang, C. Xiang, L. Miao, and J. Luo, "DeepEMD: Few-Shot Image Classification with Differentiable Earth Mover's Distance and Structured Classifiers," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 12203–12213.

[34] S. Gidaris and N. Komodakis, "Dynamic Few-Shot Visual Learning Without Forgetting," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4367–4375.

[35] S. Qiao, C. Liu, W. Shen, and A. Yuille, "Few-Shot Image Recognition by Predicting Parameters from Activations," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7229–7238.

[36] J. Wang, Y. Chen, and L. Song, "A Survey of Meta-Learning in Computer Vision," *ACM Computing Surveys*, vol. 55, no. 2, 2023.

[37] L. Bao, T. Ma, Y. Zhao, and Q. Wu, "Few-Shot Object Detection and Segmentation: A Survey," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 45, no. 8, pp. 9398–9419, 2023.

[38] H. Li, S. Luo, and F. Wang, "Meta-Learning for Medical Image Analysis: A Survey," *Medical Image Analysis*, vol. 83, 2023.

[39] W. Yin, "Meta-Learning for Few-Shot Natural Language Processing: A Survey," *arXiv preprint arXiv:2007.09604*, 2020.

[40] Y. Chen, X. Wang, Z. Liu, H. Xu, and T. Darrell, "A Closer Look at Few-Shot Classification," in *Proc. International Conf. on Learning Representations (ICLR)*, 2019.

[41] J. Lee, M. Lee, J. Shin, and H. Lee, "Meta Dropout: Learning to Perturb for Generalization," in *Proc. NeurIPS*, 2020.

[42] T. Elsken, J. Metzen, and F. Hutter, "Neural Architecture Search: A Survey," *J. Machine Learning Research*, vol. 20, 2019.

[43] S. Thrun and L. Pratt, *Learning to Learn*. Springer, 1998.

[44] A. Antoniou, H. Edwards, and A. Storkey, "How to Train Your MAML," *arXiv preprint arXiv:1810.09502*, 2018.

[45] Y. Bengio, S. Bengio, and J. Cloutier, "Learning a Synaptic Learning Rule," in *Proc. Int. Joint Conf. Neural Networks (IJCNN)*, 1992.

[46] R. Vilalta and Y. Drissi, "A Perspective View and Survey of Meta-Learning," *Artificial Intelligence Review*, vol. 18, pp. 77–95, 2002.

[47] J. Schmidhuber, "Evolutionary Principles in Self-Referential Learning," *Diploma Thesis, Technical University Munich*, 1987.

[48] A. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-Learning in Neural Networks: A Survey," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 44, no. 9, pp. 5149–5169, 2022.

[49] S. Ruder, "An Overview of Gradient Descent Optimization Algorithms," *arXiv preprint arXiv:1609.04747*, 2016.

[50] T. Bauer, A. Schick, and T. Riezler, "Zero-Shot Learning for Natural Language Processing," *Proc. ACL*, 2021.

[51] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding," *Proc. NAACL-HLT*, 2019.

[52] A. Dosovitskiy et al., "An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale," *Proc. International Conf. on Learning Representations (ICLR)*, 2021.

[53] Z. Chen, W. Liu, Y. Xu, and D. Tao, "A Closer Look at Few-Shot Classification," *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 44, no. 2, pp. 682–692, 2022.

[54] S. Park, K. Lee, and J. Shin, "Meta-Learning with Differentiable Convex Optimization," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2022.

[55] X. Zhang, W. Wang, F. Liu, and C. Shen, "Learning to Learn with Conditional Class Dependencies," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2022.

[56] A. Gupta, R. Singh, and A. Murthy, "MetaMix: Improved Meta-Learning with Inter-Task Mixup," in *Proc. International Conf. on Learning Representations (ICLR)*, 2021.

[57] R. Rusu, S. Ghosh, S. Gupta, and O. Vinyals, "Meta-Learning with Latent Embedding Optimization," in *Proc. International Conf. on Learning Representations (ICLR)*, 2019.

[58] M. Jamal and G. Qi, "Task Agnostic Meta-Learning for Few-Shot Learning," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 11719–11727.

[59] K. Chen, L. Wang, and J. Yang, "Few-Shot Image Recognition by Predicting Parameters from Activations," *IEEE Access*, vol. 10, pp. 12068–12077, 2022.

[60] Y. Liu, C. Li, and J. Sun, "Cross-Domain Meta-Learning for Few-Shot Classification," *IEEE Trans. Neural Networks and Learning Systems (TNNLS)*, 2024.

[61] T. Li, Q. Zhang, and P. Zhao, "Unified Few- to Zero-Shot Learning via Knowledge Transfer," *Pattern Recognition Letters*, vol. 180, pp. 50–58, 2025.

[62] M. Siddiqui, H. M. A. Siddiqui, S. Hussain, Dr. Muhammad Faseeh Ullah Khan, Syed Faraz Aliand M. A. Hayat, 'INTERACTION OF FINANCIAL LITERACY IN IMPULSIVE BUYING BEHAVIOR THEORY', Remittances Review, vol. 9, no. 3, pp. 780–802, 2024, doi: 10.5281/zenodo.17079478.

[63] Blockchain-Secured Iot Framework for Smart Waste Management in Urban Environments", *CRSSS*, vol. 3, no. 3, pp. 1462–1467, Aug. 2025, doi: 10.59075/mcze1x98.

[64] The Role of HR in Managing Robotic Process Automation (RPA) Displacement Anxiety among Employees", *CRSSS*, vol. 3, no. 3, pp. 1090–1109, Aug. 2025, doi: 10.59075/f4y5dc30.

[65] L. . Saeed, R. . Khan, D. S. A. . Durrani, C. Y. . Mehmood, and M. A. . Hayat, "HR Beyond the Office: Leveraging AI to Lead Distributed Teams and Cultivate Organizational Culture in the Age of Remote and Hybrid Work", *AIJSS*, vol. 4, no. 3, pp. 291–310, Jul. 2025, doi: 10.63056/ACAD.004.03.0361.

[66] Engr. Faiza Irfan, Engr. Rukhsar Zaka, Engr. Sidra Rehman, Bushra Sattar, Syed Arsalan Haider, and Muhammad Ahsan Hayat, "An IOT-Driven Smart Agriculture Framework for Precision Farming, Resource Optimization, and Crop Health Monitoring", *AIJSS*, vol. 4, no. 3, pp. 3329–3342, Aug. 2025, doi: 10.63056/ACAD.004.03.0615.

[67] Engr. Rukhsar Zaka, Syed Muhammad Mushtaher Uddin, Muhammad Ahsan Hayat, Aribah Murtaza, Syed Arsalan Haider, and Chaman lal Beejal, "AI-Driven Cybersecurity for IoT–Cloud Ecosystems", *JPEHSS*, vol. 3, no. 3, pp. 63–76, Sep. 2025.

[68] Muhammad Ahsan Hayat, Syed Affan Ahmed, Sana Fatima, Engr. Faiza Irfan, Muhammad Osama Nizamani, and Ammar Khalil, "TINY MACHINE LEARNING (TINYML) ADVANCEMENTS FOR INTELLIGENT BATTERY-POWERED IOT SENSORS", *SES*, vol. 3, no. 8, pp. 818–832, Aug. 2025.

[69] Muhammad Ahsan Hayat, Imran Ali Channa, Ubaidullah Khan, Nazia Alfred Fernandes, Urooj Tariq, and Khan Ikram Uddin, "BRIDGING CLASSICAL STATISTICS AND MODERN AI TOWARD INTERPRETABLE DATA-SCIENCE MODELS", *SES*, vol. 3, no. 9, pp. 525–537, Sep. 2025.

[70] Mian Talha Sarfraz, Muhammad Ahsan Hayat, Shayan Ahmed, Mehran Ali, and Aribah Murtaza, "MACHINE LEARNING-BASED INTRUSION AND ANOMALY DETECTION MODELS FOR SECURING IOT NETWORKS AGAINST EMERGING CYBER THREATS", *SES*, vol. 3, no. 9, pp. 1445–1463, Sep. 2025.

[71] M. A. Hayat, 'The Ethical Implications of Artificial Intelligence in Islamic Jurisprudence: A Comparative Analysis with Western Legal Systems', Al-Nasr, vol. 3, no. 3, pp. 85–106, Sep. 2024, doi: 10.5281/zenodo.17229954.

[72] Q. Muhammad, 'The Quantum Leap in Law: AI's Revolution in Justice Delivery', Al-Nasr, vol. 3, no. 2, pp. 131–146, Jun. 2024, doi: 10.5281/zenodo.17229931.

[73] Muhammad Zamin Ali Khan *, Shayan Ahmed, Khalid Bin Muhammad, Muhammad Ahsan Hayat, and Hafiza Amna Owais Ansari, "Improved Design Approach on Rehabilitative Exoskeleton", *JPEHSS*, vol. 3, no. 3, pp. 35–40, Aug. 2025.

[74] Muhammad Ahsan Hayat, Jahangir Baig, Shayan Ahmed, and Ahmed Faraz Ayubi, "TOWARDS EARLY AND ACCURATE DISEASE DETECTION THROUGH MULTIMODAL PREDICTIVE MODELING: FUSION OF ELECTRONIC HEALTH RECORDS, MEDICAL IMAGING, AND OMICS DATA USING INTERPRETABLE MACHINE LEARNING", *PRJ*, vol. 3, no. 11, pp. 01–29, Nov. 2025.