U-MOF: Uncertainty-Guided Parameter-Efficient Multi-Objective Fine-Tuning for Long-Tailed Recognition

Di Wu[1,†], Yuan Dong[1,†], Zhe Zhao[1], Liheng Yu[1], Xiaofeng Cao[3], Pengkun Wang[1,2*] Yang Wang[1,2*]
[1]University of Science and Technology of China [2]Suzhou Institute for Advanced Research [3]Tongji University

1. Background & Motivation

The Challenge of Long-Tailed Recognition

Long- tailed recognition (LTR) requires balancing head- class preservation and rare-class recovery. Two -stage fine- tuning offers efficient adaptation by freezing the pretrained backbone and optimizing lightweight components. However, coordinating heterogeneous objectives while maintaining sufficient tail discrimination remains challenging.

Ineffective Stage-Two Adaptation

- **Root Cause 1 — Unreliable Objective Routing:** Existing methods often rely on delayed validation feedback or historical class- wise statistics. Such signals can become noisy or misleading under skewed distributions and may fail to distinguish different sources of prediction difficulty.
- **Root Cause 2 — Limited Feature Plasticity:** Classifier only adaptation restricts feature refinement, making it difficult to recover discriminative representations for rare classes.
- **Consequence:** Inappropriate objective allocation and insufficient feature adaptation can limit tail recovery and weaken overall recognition.

Limitations of Existing Approaches

- **Empirical Routing:** Relies on delayed performance statistics that may not reflect current optimization needs.
- **Classifier-Only Fine-Tuning:** Preserves efficiency but lacks feature-level flexibility for rare categories.
- **Naive Multi-Objective Fusion:** Combines heterogeneous objectives without adequately resolving gradient interference.

Our Key Insight

Diagnose Uncertainty, Route Objectives, and Preserve Feature Plasticity.
Effective stage-two adaptation should identify different prediction difficulties through online committee diagnostics, dynamically allocate objective emphasis, and coordinate conflicting gradients while retaining lightweight feature-level adaptation.

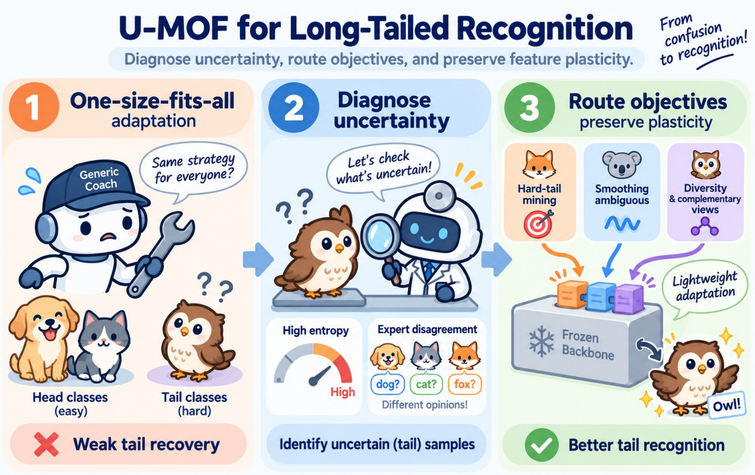

2. Method: U-MOF

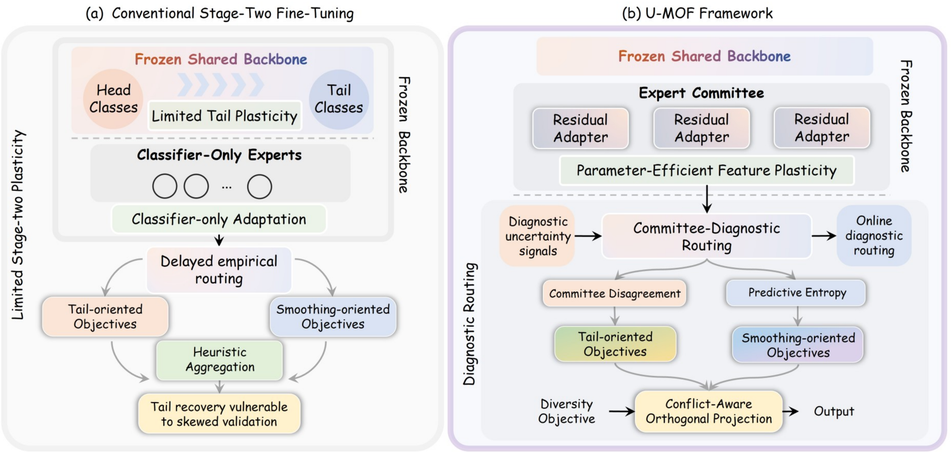

Overview: Uncertainty-Guided Multi-Objective Fine-Tuning

U-MOF introduces a **parameter-efficient framework combining residual expert adapters, committee-diagnostic objective routing, and conflict-aware gradient projection**. It coordinates heterogeneous long-tailed objectives while preserving feature-level plasticity under a shared frozen backbone.

Component 1: Parameter-Efficient Expert Adaptation

- **Residual Adapters:** Lightweight, centered residual modulation restores limited feature-level plasticity without updating the pretrained backbone. Near-identity initialization stabilizes early adaptation.
- **Expert Committee:** Multiple adapter-head experts share one frozen backbone while learning complementary feature representations and predictions.

Component 2: Committee-Diagnostic Objective Routing

U-MOF uses online committee statistics to dynamically allocate emphasis across tail-oriented, smoothing-oriented, and diversity objectives.

- **Expert Disagreement:** Higher cross-expert disagreement increases emphasis on tail-oriented mining objectives (KPS/BCL).
- **Predictive Entropy:** Higher within-expert entropy increases smoothing-oriented emphasis (LOS) to stabilize ambiguous predictions.
- **U-PLA Routing:** Converts normalized mini-batch diagnostics into adaptive objective weights without relying on delayed validation feedback.

Component 3: Conflict-Aware Gradient Coordination

- **Classification-First Projection:** Coordinates classification-related gradients before projecting conflicting diversity gradients, reducing destructive interference.
- **Diversity Preservation:** Encourages complementary expert specialization while preventing diversity objectives from directly opposing classification improvement.

Final Prediction

- **Committee Averaging:** Expert predictions are uniformly averaged after a shared backbone forward pass, without routing or projection at inference time.
- **Core Advantage:** U-MOF combines **adaptive objective coordination and lightweight feature plasticity** to improve tail recovery while preserving efficient stage-two adaptation.

3. Experiments

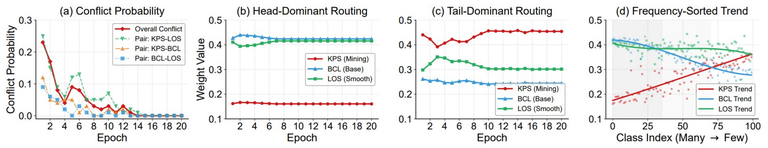

Experimental Setup

U-MOF is evaluated on CIFAR-100-LT (IR=10/50/100), ImageNet-LT, and iNaturalist 2018 using Top-1 accuracy. Results are averaged over five independent runs.

Mechanism Analysis

- **Reduced Gradient Conflict:** Classification-first projection mitigates destructive interference among heterogeneous objectives.
- **Adaptive Objective Routing:** Committee diagnostics adjust mining and smoothing emphasis across different training regimes.
- **Tail-Oriented Adaptation:** Routing dynamics show increased mining emphasis toward lower-frequency categories.

Key Findings

On CIFAR-***100*** LT (IR=100), U-MOF improves **Tail accuracy by 2.0 pp** and **overall accuracy by 0.7 pp** over TS-MOF under matched settings.
U-MOF achieves **57.6% overall / 51.5% Tail** on ImageNet-LT and **73.1% overall / 75.1% Tail** on iNaturalist 2018, with lightweight adaptation overhead.

Method	IR=10				IR=50				IR=100			
	Head	Medium	Tail	All	Head	Medium	Tail	All	Head	Medium	Tail	All
CE	63.2	40.3	–	56.5	63.9	36.2	15.2	43.8	65.6	36.2	8.2	38.1
CE-DRW [1]	62.5	48.6	–	58.2	60.6	39.0	22.9	45.0	63.4	41.2	15.7	41.4
LDAM-DRW [1]	62.7	46.1	–	57.5	63.0	41.2	25.1	47.2	62.8	42.6	21.1	43.2
Balanced Softmax [4]	61.5	50.6	–	58.1	60.3	41.3	34.3	47.9	59.6	42.3	23.7	42.8
RIDE (3 experts) [5]	66.4	49.4	–	61.1	65.7	47.7	31.8	52.2	65.7	48.6	25.0	47.5
BCL [31]	62.2	51.8	–	58.9	61.6	43.1	34.3	49.1	63.1	42.9	23.9	44.2
KPS [32]	61.7	58.7	–	59.5	51.6	49.5	52.4	50.5	41.9	39.5	48.7	42.2
SHIKE [33]	66.0	45.0	–	59.0	67.0	43.0	23.0	49.5	66.0	39.0	12.0	46.9
LOS [34]	71.9	62.3	–	69.0	72.4	51.4	40.2	58.0	70.3	52.3	36.6	53.9
TS-MOF [17]	74.7	62.0	–	70.8	75.2	50.7	47.5	60.2	79.0	49.2	39.9	56.8
U-MOF	73.2 ±0.3	**66.6** ±0.4	–	**71.1** ±0.2	74.9 ±0.2	**52.2** ±0.3	**49.9** ±0.4	**61.1** ±0.2	76.8 ±0.3	**51.4** ±0.3	**41.9** ±0.4	**57.5** ±0.2

Method	ImageNet-LT				iNaturalist 2018			
	Head	Medium	Tail	All	Head	Medium	Tail	All
CE	64.0	33.8	5.8	41.6	73.9	63.5	55.5	61.0
CE-DRW [1]	61.7	47.3	28.8	50.1	68.2	67.3	66.4	67.0
cRT [2]	58.8	44.0	26.1	47.3	69.0	66.0	63.2	65.2
LDAM-DRW [1]	60.4	46.9	30.7	49.8	–	–	–	66.1
BS [4]	60.9	48.8	32.1	51.0	65.7	67.4	67.5	67.3
KPS [32]	59.7	49.2	35.9	52.3	68.1	69.5	70.2	69.6
RIDE (3 experts) [5]	64.9	50.4	34.4	53.6	69.5	71.0	70.4	70.6
BCL [31]	65.3	53.5	36.3	55.6	69.5	70.9	71.3	70.9
LOS [34]	63.2	50.7	42.3	54.4	69.2	70.7	71.3	70.8
TS-MOF [17]	65.9	52.5	42.6	56.3	72.8	73.6	70.5	72.3
U-MOF	**67.6** ±0.2	51.3 ±0.3	**51.5** ±0.3	**57.6** ±0.1	**76.4** ±0.2	70.8 ±0.2	**75.1** ±0.3	**73.1** ±0.1

Configuration	Adapter	Div. Obj.	Projection	U-PLA	Head	Medium	Tail	Overall
Classifier-Only Multi-Obj. Baseline	×	×	×	×	66.0 ±0.4	35.0 ±0.2	6.8 ±0.6	43.1 ±0.3
+ Feature Tuning	✓	×	×	×	66.5 ±0.3	42.1 ±0.1	18.3 ±0.5	48.2 ±0.2
+ Diversity and Projection	✓	✓	✓	×	67.4 ±0.2	48.5 ±0.2	35.6 ±0.4	54.8 ±0.2
Full U-MOF	✓	✓	✓	✓	**67.6** ±0.2	**51.3** ±0.3	**51.5** ±0.3	**57.6** ±0.1