



Von Mises-Fisher Mixture Model with Dynamic Shrinkage for Realistic Test-Time Transduction

Jiazhen Huang¹ · Zhiming Liu¹ · Changhu Wang² · Wei Ju^{3, ✉} · Ziyue Qiao^{4, ✉} · Xiao Luo⁵


¹Tsinghua University · ²Fred Hutchinson Cancer Center · ³Peking University · ⁴Great Bay University · ⁵University of Wisconsin–Madison




Overview & Background

Realistic Class Imbalance

Batch adaptation




Online adaptation



Transduction Paradigm

Zero-shot predictions

Adapted predictions



- Perform **joint prediction** by exploiting all available samples
- Being **training-free**, **model-agnostic** and **fast**
- Novel paradigm for Test-Time Adaptation (TTA)

PLE View: Shrinkage Toward a KL Anchor

For any regular minimal exponential family, a **KL-anchored PLE objective** yields an **adaptive shrinkage behavior** for transduction.

$$p(x | \eta) = h(x) \exp(\eta^\top T(x) - A(\eta)).$$

$$\min_{\{\eta_k\}} - \sum_{i=1}^N \sum_{k=1}^K z_{i,k} \log p(x_i | \eta_k) + \alpha \sum_{k=1}^K \text{KL}(q_k \| p(\cdot | \eta_k)),$$


PLE objective

$$\nabla A(\eta_k^*) = \frac{S_k + \alpha \mu'_k}{n_k + \alpha} = \beta_k \hat{\mu}_k + (1 - \beta_k) \mu'_k, \quad \beta_k = \frac{n_k}{n_k + \alpha}.$$

Data-driven convex combination in mean-parameter space.

text prior

noisy MLE



Limitation 1: no anchor

With $\alpha = 0$, PLE reduces to MLE and can overfit locally biased statistics.


Limitation 2: bounded but static shrinkage

A fixed anchor weight α assumes equal reliability across instances and classes.

Method

MOON: Training-free Test-Time Transduction

vMF Mixture Model



$$p_{i,k}^{\text{vMF}} \propto \mathcal{C}_d(\kappa_k) \exp(\kappa_k \mu_k^\top \mathbf{f}_i)$$

MOON Objective

$$\mathcal{L}_{\text{PLE}}(\mathbf{z}; \mu, \kappa) = \underbrace{\gamma \left[- \sum_i \mathbf{z}_i^\top \log p_i^{\text{vMF}} + \mathcal{R}(\mathbf{z}) \right]}_{\text{Log-Likelihood}} + \underbrace{\alpha \sum_k \text{KL}(\mathcal{V}'_k \| \mathcal{V}_k)}_{\text{KL Statistical anchor}},$$

$$\mathcal{R}(\mathbf{z}) = - \underbrace{\sum_{i,j} \omega_{ij} \mathbf{z}_i^\top \mathbf{z}_j}_{\text{Laplacian reg.}} + \underbrace{\sum_i \text{KL}(\mathbf{z}_i \| \hat{\mathbf{y}}_i)}_{\text{Text supervision}}.$$

Dynamic weights from zero-shot priors!

MOON models VLM embeddings on the unit hypersphere, and perform a **KL-anchored PLE objective** with **dynamic shrinkage** for adaptation.

TL; DR

MOON is a training-free, model-agnostic, and efficient method for test-time adaptation (TTA) of VLMs (e.g., CLIP), designed for realistic class imbalanced scenarios.

Dynamic Shrinkage

Instance-level

$$\gamma_i = 1 - \frac{H(\hat{\mathbf{y}}_i)}{\log K},$$

entropy-guided

$$H(\hat{\mathbf{y}}_i) = - \sum_{k=1}^K \hat{y}_{i,k} \log \hat{y}_{i,k}$$

High-entropy samples contribute less to parameter estimation.

Class-level

$$\alpha_k = \frac{1}{\lambda_k},$$

confidence (max prob)-guided

$$\lambda_k = \sqrt{\frac{1}{N} \sum_{i=1}^N \hat{\mathbf{y}}_{i,k} \odot \max_i \hat{\mathbf{y}}_{i,k}}.$$

Low-confidence outlier classes shrink more towards the anchor.

Closed-Form Parameter Update


$$n_k = \sum_i \gamma_i z_{i,k}, \quad \beta_k = \frac{n_k}{n_k + \alpha_k},$$

$$\mu_k = \frac{\beta_k \mathbf{v}_k + (1 - \beta_k) \mu'_k}{\|\beta_k \mathbf{v}_k + (1 - \beta_k) \mu'_k\|}, \quad \mathcal{A}_d(\kappa_k) = \|\beta_k \mathbf{v}_k + (1 - \beta_k) \mu'_k\|,$$

BSUM iteration, until convergence

Experiment

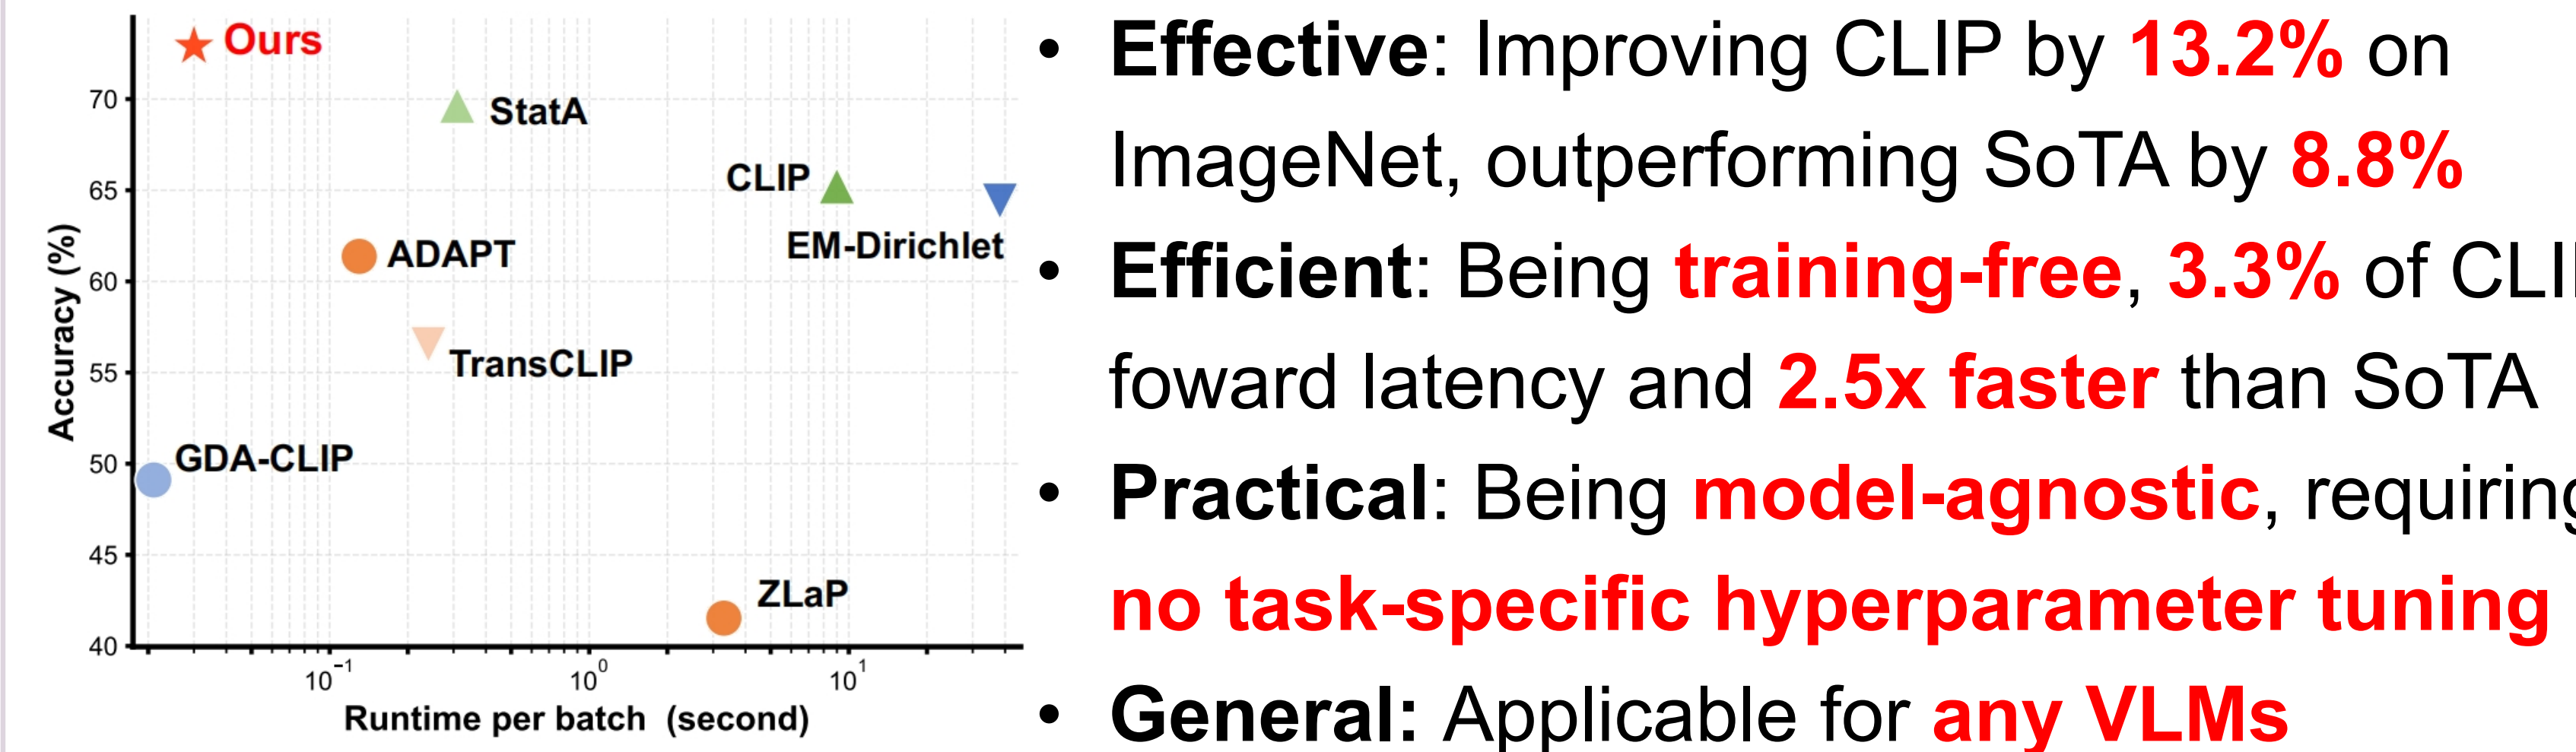

Main Results under Realistic Imbalance



K_{eff}	Method	ImageNet	SUN397	Aircraft	Flowers102	Pets	Food101	StanfordCars	UCF101	Avg.
Very Low (≤ 0.1)	CLIP	66.6	62.5	24.7	48.3	65.6	85.9	89.1	93.2	65.2
	TransCLIP	69.3 _{+1.7}	64.8 _{+1.3}	27.4 _{+1.7}	46.9 _{+1.1}	68.0 _{+1.1}	87.2 _{+1.9}	89.4 _{+1.2}	92.1 _{+1.0}	66.6 _{+1.2}
	ADAPT	70.2 _{+0.8}	75.7 _{+1.0}	28.5 _{+1.1}	47.2 _{+1.1}	68.2 _{+1.0}	88.1 _{+0.9}	87.5 _{+1.0}	91.2 _{+0.8}	68.5 _{+0.9}
	StatA	145 _{+1.1}	13.0 _{+0.8}	8.4 _{+0.8}	36.6 _{+1.7}	23.7 _{+0.9}	31.9 _{+0.8}	57.0 _{+0.8}	22.4 _{+0.8}	13.0 _{+0.8}
	MOON	72.9 _{+0.3}	60.0 _{+0.1}	29.3 _{+0.4}	58.8 _{+0.5}	70.2 _{+0.9}	90.3 _{+0.4}	95.5 _{+0.1}	93.0 _{+0.2}	72.9 _{+0.3}
	CLIP	216 _{+0.8}	21.1 _{+0.4}	11.6 _{+0.4}	45.1 _{+1.1}	34.7 _{+0.9}	59.2 _{+0.8}	64.4 _{+0.7}	62.3 _{+0.9}	37.7 _{+0.8}
	TransCLIP	60.8 _{+0.4}	56.0 _{+0.1}	21.4 _{+0.3}	45.9 _{+0.4}	59.7 _{+0.5}	81.4 _{+0.3}	84.7 _{+0.4}	68.8 _{+0.3}	60.7 _{+0.3}
	ADAPT	72.9 _{+0.3}	60.0 _{+0.1}	29.3 _{+0.4}	58.8 _{+0.5}	70.2 _{+0.9}	90.3 _{+0.4}	95.5 _{+0.1}	93.0 _{+0.2}	72.9 _{+0.3}
	StatA	62.2 _{+0.3}	63.6 _{+0.1}	24.3 _{+0.3}	55.1 _{+0.3}	67.4 _{+0.1}	80.9 _{+0.1}	92.5 _{+0.1}	92.1 _{+0.1}	68.8 _{+0.1}
	MOON	67.1 _{+0.3}	64.3 _{+0.1}	24.6 _{+0.1}	47.9 _{+0.3}	68.0 _{+0.1}	85.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	67.1 _{+0.3}
Low (2-10)	CLIP	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	TransCLIP	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	ADAPT	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	StatA	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	MOON	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	CLIP	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	TransCLIP	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	ADAPT	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	StatA	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
	MOON	68.1 _{+0.1}	64.8 _{+0.1}	24.6 _{+0.1}	56.2 _{+0.1}	68.1 _{+0.1}	81.9 _{+0.1}	89.4 _{+0.1}	93.0 _{+0.1}	68.1 _{+0.1}
Medium (5-25)	CLIP	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	TransCLIP	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	ADAPT	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	StatA	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	MOON	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	CLIP	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	TransCLIP	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	ADAPT	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	StatA	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}
	MOON	77.1 _{+1.1}	72.9 _{+0.8}	26.1 _{+1.1}	42.9 _{+1.1}	71.6 _{+0.7}	90.8 _{+0.9}	93.7 _{+0.5}	42.9 _{+0.8}	72.9 _{+0.8}

Further Discussion

α	γ	Update μ	Update κ	Bs=64	Bs=1,000
✓	✓	✓	✓	69.0	68.9
✓	✓	✓	✓	72.3	70.4
✓	✓	✗	✓	66.8	66.5
✓	✓	✓	✗	70.8	68.9
✓	✓	✓	✓	72.4	71.1



- Effective:** Improving CLIP by **13.2%** on ImageNet, outperforming SoTA by **8.8%**
- Efficient:** Being **training-free**, **3.3%** of CLIP forward latency and **2.5x faster** than SoTA
- Practical:** Being **model-agnostic**, requiring **no task-specific hyperparameter tuning**
- General:** Applicable for **any VLMs**