

Decentralized Instruction Tuning: Conflict-Aware Splitting and Weight Merging

Minsik Choi^{1,2 *} and Geewook Kim^{1,3 *†}

* Equal Contribution


¹ **NAVER Cloud**



³ **KAIST AI**
Kim Jaechul Graduate School



MERIT removes cross-partition gradient sync from instruction tuning and **gains +2.7** on 8 multimodal benchmarks.




MERIT pipeline

STAGE 1
 θ^0 basin : Gradient extraction at the shared init ckpt ($n = 200$ calibration / task)

STAGE 2
PCA on cosine sim : Top- r PCA, then recursive sample-balanced split.

STAGE 3
 $K = 2^r$ branches : Communication-free fine-tuning per branch.

STAGE 4
 $\bar{\theta}$ merge : One-shot token-weighted avg.



No cross-branch communication. · One-shot merge.

Theoretical implications

1 Merging gain

$$\text{THM 1} \quad \sum_i w_i L(\theta_i) - L(\bar{\theta}_w) = \frac{1}{2} \sum_\ell \lambda_\ell \text{Var}_w(u_\ell^T \delta) \geq 0$$

Curvature-weighted variance reduction :
Never worse than averaging the losses.

2 PCA-aligned splitting maximizes the gain

PROP 1

$$\mathcal{G}_{var}^{PCA} \propto \lambda^3 \cdot \alpha^2$$

PCA aligned split wins :
Advantage scales cubically with curvature gap.

3 Spectral filtering + implicit norm regularization

SEC 4.3

$$\kappa_{eff} \ll \kappa, \|\bar{\theta} - \theta^{(0)}\| < \|\theta_{joint} - \theta^{(0)}\|$$

Faster optimization & tighter **PAC-Bayes** bound.

Merge-readiness

LMC	DISPLACEMENT	TRAIN ↔ HELD-OUT	PERTURBATION
0.0×10	$2.4\text{--}2.9\times$	$+0.49 \rightarrow 1.27$	$\sigma \leq 0.1$
All 10 barriers = 0.0 (3B, K = 4)	merged stays closer to θ^0 across epochs	higher train loss, better held-out → implicit reg.	robust at σ $\in \{0.01, 0.05, 0.1\}$

Qwen 2.5-3B Instruct

- Trained with 66 text-only tasks

Method	Knowledge QA		Commonsense		Text Inference		Problem Solving		Avg.
	MMLU	GPQA	HellaSwag	WinoGrande	BoolQ	XNLI	ARC-C	HumanEval	
Base model	65.5	29.5	74.5	70.1	77.3	42.9	55.6	37.8	56.7
Joint training (1 ep)	65.8	30.1	74.0	69.9	83.9	41.9	55.5	39.7	57.6
Joint training (2 ep)	66.1	30.6	76.3	69.2	82.9	41.9	56.5	42.7	58.3
Random (2 groups)	65.8	29.5	74.4	<u>70.3</u>	85.3	42.4	55.1	<u>42.4</u>	58.2
Random (4 groups)	<u>66.0</u>	30.3	<u>74.7</u>	<u>70.3</u>	85.7	<u>42.5</u>	54.9	<u>42.1</u>	<u>58.3</u>
Conflict-induced (2 groups)	65.8	28.7	74.0	69.4	85.3	42.2	53.8	42.1	57.7
Uniform soup (2)	65.9	<u>30.6</u>	74.0	70.2	83.9	41.9	55.2	40.2	57.7
Uniform soup (3)	65.7	30.8	74.0	70.2	84.2	42.1	55.7	39.0	57.7
MERIT (Proposed, 1D split, 2 groups)	66.0	<u>30.6</u>	74.4	69.9	86.3	42.1	54.2	41.2	58.1
MERIT (Proposed, 2D split, 4 groups)	66.1	30.0	<u>74.7</u>	70.8	<u>86.1</u>	42.4	55.3	41.5	58.4

Qwen 2.5-3B VL

- Trained with 136 Vision-FLAN tasks

Method	General MCQA		User Preference & Fluency		Text-Rich VQA		Image Reasoning		Avg.
	SeedBench	MMBench	LLaVA-W	MMVet	TextVQA	Ai2D	MathVista	MMMU	
Base model	66.8	79.7	53.2	34.0	61.2	63.8	29.6	41.2	53.7
Joint training (0.5 ep)	67.4	79.4	40.2	33.1	67.2	60.9	31.4	41.6	52.7
Joint training (1 ep)	69.2	80.5	41.9	36.4	68.0	62.6	34.2	41.9	54.3
Joint training (2 ep)	70.0	81.4	42.8	<u>37.6</u>	63.4	62.5	<u>36.5</u>	43.0	54.7
Random (2 groups)	69.4	80.1	44.5	34.7	70.4	62.7	34.0	41.2	54.6
Random (4 groups)	70.4	81.0	40.6	34.7	70.4	63.1	34.0	40.8	54.4
Random (8 groups)	69.5	79.9	42.2	35.0	73.7	61.7	33.5	40.5	54.5
Conflict-induced (2 groups)	70.7	80.6	42.6	35.4	70.0	62.9	34.4	42.3	54.9
Uniform soup (2)	70.2	<u>81.1</u>	45.0	35.3	68.9	<u>63.4</u>	36.8	42.2	55.4
Uniform soup (3)	70.1	<u>81.1</u>	42.3	36.3	68.8	63.1	35.8	42.5	55.0
Uniform soup (4)	70.2	<u>81.1</u>	41.8	36.3	68.4	<u>63.4</u>	35.9	42.2	54.9
MERIT (Proposed, 1D split, 2 groups)	71.0	80.0	43.1	35.0	72.4	62.1	<u>36.5</u>	41.4	55.2
MERIT (Proposed, 2D split, 4 groups)	<u>70.8</u>	78.4	47.4	36.6	<u>74.1</u>	61.5	36.0	40.7	<u>55.7</u>
MERIT (Proposed, 3D split, 8 groups)	70.5	80.1	<u>52.0</u>	37.7	75.2	62.5	35.4	<u>42.7</u>	57.0

Qwen 2.5-7B VL

- Trained with 176 Vision tasks

Model	Train Data	General MCQA		User Preference & Fluency		Text-Rich VQA		Image Reasoning		Avg.
		SeedBench	MMBench	LLaVA-W	MMVet	TextVQA	Ai2D	MathVista	MMMU	
LLaVA-7B	0.6M	37.0	38.7	57.2	25.5	—	48.3	25.4	34.1	—
LLaVA-1.5-7B	0.7M	65.9	64.3	59.6	31.1	58.2	54.8	25.6	35.3	49.4
LLaVA-1.5-13B	0.7M	68.2	67.7	66.1	36.1	61.3	59.5	27.7	33.6	52.5
+ Joint FFT (Xu et al., 2024b)	0.7M + 0.2M	—	69.8	38.5	33.4	—	—	—	34.4	—
Base VLM 7B	0.7M	64.5	67.2	57.2	28.9	57.6	46.5	26.9	32.8	47.7
+ Joint FFT	0.7M + 1.6M	70.1	76.2	58.8	32.5	73.8	58.6	32.9	36.4	54.9
+ MERIT (Proposed, 2D)	0.7M + 1.6M	70.6	75.7	59.2	35.1	<u>72.4</u>	<u>58.4</u>	35.4	36.1	55.4
Scaled base VLM 7B	3.6M	69.8	74.5	67.1	34.0	72.4	69.2	39.4	45.1	58.9
+ Joint FFT	3.6M + 1.6M	71.3	75.3	50.2	41.5	80.2	71.6	49.7	47.4	60.9
+ MERIT (Proposed, 2D)	3.6M + 1.6M	71.5	75.8	66.2	39.1	79.8	71.9	43.0	44.8	61.5

Stage	V100 System	A100 System
Basin preparation (pre-alignment)	4h 24m / 20h 13m (LLaVA recipe stage 1/2)	39 h
Gradient extraction	1h 38m (136 datasets)	1h 59m (176 datasets)
Cosine similarity matrix computation	28m	38m
PCA on cosine similarity matrix	0.7s (load) + 0.2s (compute)	0.6s (load) + 0.12s (compute)
Joint training (1 ep)	4h 22m	43 h 18m
Joint training (2 ep)	8h 40m	—
MERIT (CosSim PCA, 3D split)	5h 24m	43 h 39m

