

MAPLE: Modality-Aware Post-training and Learning Ecosystem

Anonymous Authors¹

Abstract

Multimodal language models now integrate text, audio, and video for unified reasoning. Yet existing RL post-training pipelines treat all input signals as equally relevant, ignoring which modalities each task actually requires. This modality-blind training inflates policy-gradient variance, slows convergence, and degrades robustness to real-world distribution shifts where signals may be missing, added, or reweighted. We introduce **MAPLE**, a complete modality-aware post-training and learning ecosystem comprising: (1) MAPLE-bench, the first benchmark explicitly annotating minimal signal combinations required per task; (2) MAPO, a modality-aware policy optimization framework that stratifies batches by modality requirement to reduce gradient variance from heterogeneous group advantages; (3) Adaptive weighting and curriculum scheduling that balances and prioritizes harder signal combinations. Systematic analysis across loss aggregation, clipping, sampling, and curriculum design establishes MAPO’s optimal training strategy. Adaptive weighting and curriculum focused learning further boost performance across signal combinations. MAPLE narrows uni/multi-modal accuracy gaps by 30.24%, converges 3.18× faster, and maintains stability across all modality combinations under realistic reduced signal access. MAPLE constitutes a complete recipe for deployment-ready multimodal RL post-training.

1. Introduction

Recent advances in Multimodal Language Models (MLMs) enable unified reasoning across text, audio, and video sensory streams (Comanici et al., 2025; Hurst et al., 2024; Xu et al., 2025b). RL post-training has become central to

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

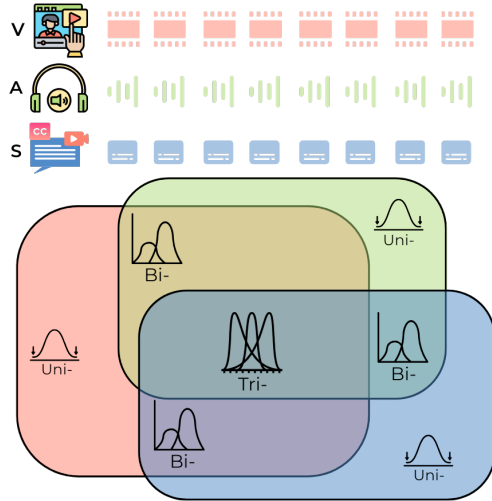


Figure 1. **Heterogeneous multimodal latent geometry needs stratified RL.** Unimodal clusters (V: video, A: audio, S: subtitle) connect through bi-modal (VA, VS, AS) and tri-modal (VAS) manifolds. Modality-blind RL mixes trajectories from these 7 distinct regions indiscriminately, inflating gradient variance across heterogeneous reward scales. MAPLE’s stratified training assigns each trajectory to its native signal combination for stable optimization.

adapting these models for complex decision-making and alignment objectives (Schulman et al., 2017; Yin et al., 2023; Rafailov et al., 2023). Effective post-training must simultaneously sustain high task accuracy, robustness to distributional shifts, and efficiency under realistic compute limits (Khatri et al., 2025; McKinzie et al., 2023b). Yet most current approaches, inherited from unimodal language settings, treat all multimodal trajectories as if drawn from a single joint distribution. This assumption neglects modality-specific cues that determine which input channels truly matter for solving a given query (Wei et al., 2025).

In real-world deployments, queries may typically depend on subset of input streams (text, audio, or video) rather than all modalities simultaneously. Fig. 1 illustrates this: distinct unimodal clusters gradually fuse into shared bi- and tri-modal manifolds in latent space. However, existing RL post-training pipelines process every rollout as if all modalities contribute equally, blending fundamentally heterogeneous interactions. Under value-model-free methods like Group Relative Policy Optimization (GRPO) (Shao et al., 2024), trajectories that leverage rich multimodal redundancy are

jointly optimized with single-channel reasoning tasks. This uniform treatment biases policy updates toward dominant modality combinations while under-training regimes with sparse or noisy signals, precisely those most common in real test conditions.

This representational heterogeneity across modality regimes translates directly into measurable optimization challenges. When reward scales and noise characteristics are conflated across modality subsets, modality-blind batching amplifies policy-gradient variance, mirroring known batch-composition effects in policy optimization (Che et al., 2023; Vasan et al., 2024) and cross-modality gradient conflicts (Anonymous, 2025). Empirically, our analysis (Sec 5.2) shows that mixed-modality batches exhibit higher gradient-norm volatility than modality-stratified batches, leading to slower convergence and instability patterns (Rha et al., 2025). More fundamentally, modality-blind optimization collapses modality-specific structure.

The consequences of this mismatch intensify under deployment-time distribution shifts. In real-world settings, models frequently encounter “missing-at-test” or partial-observation regimes that violate the full-signal assumptions made during training (McKinzie et al., 2023b). Robust policies must therefore operate reliably under absent, spurious, or reweighted input signal, conditions where failures tend to be systematic rather than random. Modality-unaware optimization exacerbates this gap by training on noisy, mixed-modality batches while evaluating under shifted signal configurations. Across benchmark evaluations, omni-modal models experience substantially greater degradation than modality-specialized counterparts (Jiang et al., 2025b; Li et al., 2025b).

We address these challenges through three concrete contributions that together constitute **MAPLE**, a *modality-aware post-training and learning ecosystem* for multimodal reinforcement learning. First, we present **MAPLE-bench**, a modality-aware benchmark where each instance is explicitly annotated with the minimal subset of signals, drawn from video, audio, and subtitles, necessary to solve the target task. This design enables principled conditioning of both data sampling and policy optimization, bridging the gap between synthetic multimodal mixtures and realistic input dependencies. MAPLE-bench spans multiple-choice question answering and open-ended omni-modal captioning across seven signal subsets, including uni-modal, bi-modal, and tri-modal configurations, thereby capturing the diversity of plausible deployment regimes.

Second, we revisit value-model-free RL in multimodal settings. While recent optimizers such as GRPO (Shao et al., 2024), Dr. GRPO (Liu et al., 2025a), and DAPO (Yu et al., 2025) demonstrate strong performance in unimodal domains, their design choices do not generalize well to hetero-

geneous signal regimes. We conduct a systematic ablation of loss aggregation, clipping behaviors, dynamic sampling, and data filtering strategies, explicitly conditioning each component on the active signal subset. This analysis yields **Modality-Aware Policy Optimization (MAPO)**, a unified optimization framework that stabilizes multimodal post-training without added critics or auxiliary regularization. We benchmark MAPO against a modality-unaware policy that always accesses all signals.

Finally, we propose **Adaptive-training strategies** that progressively trains from uni-modal reasoning tasks toward multi-signal integration. This adaptation explicitly exposes how varying signal diversity influences convergence dynamics and robustness. Across all settings, MAPLE attains stable and competitive results, while operating under reduced modality access. It also narrows accuracy gaps between uni- and multi-modal regimes by 30.24% and converges $3.18\times$ faster than full-signal training dependent on redundant modalities. Taken together, **MAPLE** constitutes a complete and practical recipe that integrates benchmark design, policy optimization, and curriculum learning for efficient, robust, and deployment-ready multimodal RL post-training.

In summary, this work:

1. Introduces **MAPLE-bench**, a modality-aware benchmark suite for QA and captioning with *Required Modality Tags* for explicit control of sampling, optimization, and evaluation.
2. Develops **MAPO** and adaptive training strategies that exploit per-sample modality structure to stabilize training and enhance performance across diverse signal configurations.
3. Demonstrates, through extensive experiments, that MAPLE yields efficient and robust multimodal RL post-training under real-world variability.

2. Related Works

Multimodal Large Language Models. Recent multimodal LLMs unify text, vision, and audio through shared token interfaces, with systems like Qwen2.5-Omni, Gemini 3, and Phi-4-Multimodal demonstrating strong omni-modal reasoning (Yin et al., 2023; Xu et al., 2025a; Abouelenin et al., 2025). These models typically assume full modality availability during both training and evaluation, treating heterogeneous signals as uniformly relevant inputs. Extended architectural comparisons appear in Appendix A.

Multimodal Benchmarks. Benchmarks such as MME, MMBench, OmniBench, and MAVERIX (Liu et al., 2024; Li et al., 2025c; Xie et al., 2025) assess cross-modal reasoning under complete signal access but rarely examine perfor-

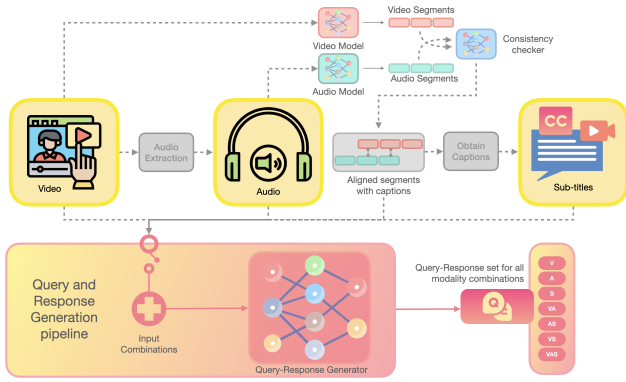


Figure 2. Overview of the MAPLE-bench curation pipeline. Seed videos are sampled from large-scale corpora and decomposed into audio, visual, and textual streams. Each modality undergoes independent annotation and cross-modal consistency alignment before generating captions and query–response pairs. A modality-isolated prompting stage then synthesizes tasks across uni-, bi-, and tri-modal configurations, producing a balanced dataset for modality-aware post-training.

mance across varying modality combinations. Even ablation studies maintain fixed ground-truth answers despite changing inputs, conflating true information deficits with fusion failures. Table 5 (Appendix A) contrasts these limitations with MAPLE-bench’s modality-conditioned evaluation.

Post-training Optimization. Value-model-free RL methods like GRPO, Dr. GRPO, and DAPO (Shao et al., 2024; Liu et al., 2025a; Yu et al., 2025) excel in text domains through group-normalized advantages and simplified reward modeling. Modality-balancing (Li et al., 2025b) and curriculum strategies (Pena et al., 2023) address signal dominance but do not condition optimization on per-query required signal subsets. MAPLE extends these paradigms with modality-aware data stratification, loss aggregation, and curriculum design. Additional post-training variants are discussed in Appendix A.

3. MAPLE-bench: A modality annotated benchmark

We introduce MAPLE-bench, a modality-aware dataset designed for evaluating post-training strategies under heterogeneous signal access. MAPLE-bench extends subsets of the Daily-Omni (Zhou et al., 2025) and VAST-Omni (Chen et al., 2023) corpora, curated to ensure strong multimodal coverage across video, audio, and text streams.

3.1. Dataset Construction

Our pipeline begins by sampling seed videos from these corpora and extracting aligned audio tracks. Each modality is independently segmented and annotated using domain-

specific models, followed by a cross-modal consistency and alignment phase driven by Gemini-2.5-Flash (Comanici et al., 2025). This stage ensures temporal coherence and semantic agreement between visual and acoustic narratives. From the aligned segments, captions and structured event descriptions are generated.

Next, we apply a *modality-isolated protocol* based query–response generation procedure, prompting a reasoning model with controlled modality combinations to synthesize datasets for two complementary tasks: (i) open-ended captioning (generative) and (ii) multiple-choice question answering (discriminative). Each example is labeled with a set of Required Modality Tags (RMTs) denoting the minimal modality subset necessary to solve the task, enabling fine-grained control over modality conditioning during training and evaluation. An overview of the full curation and query–response generation pipeline is illustrated in Fig. 2.

The curated corpus yields query–response pairs across uni-, bi-, and tri-modal regimes. To validate data integrity, we apply multiple LLM-based consistency checks and manual audits on the test split. This process establishes MAPLE-bench as a robust foundation for studying modality-aware reinforcement learning. Further synthetic dataset curation details are provided in Appendix B, and dataset statistics are summarized in the following section.

3.2. Dataset Statistics

MAPLE-bench spans all seven Required Modality Tags: {V, A, S, VA, VS, AS, VAS}, where V, A, and S denote video, audio, and subtitles respectively.

MAPLE-QA. MAPLE-QA-train provides 47,893 QA pairs across 546 videos, each tagged with minimal required RMTs and structured as discriminative 4-option multiple choice questions for verifiable reward signals (Guo et al., 2025). The human-evaluated MAPLE-QA-eval contains 5,001 samples from a different set of 68 videos, following identical generation protocols. Automatic filtering via GPT-4o, GPT-5, and Gemini-2.5-Flash (Hurst et al., 2024; Comanici et al., 2025) in data generation pipeline removes trivially solvable pairs.

MAPLE-Caption. We construct MAPLE-Caption-train with 5,120 captioning samples for post-training (no human curation). The human-curated MAPLE-Caption-eval benchmark comprises 5,348 samples, perfectly balanced at 764 samples per RMT (764 videos total). All evaluation captions are verified for factual consistency and modality coverage relative to provided signals.

A detailed comparison of MAPLE-bench against existing multimodal benchmarks appears in Appendix Table 5.

3.3. Evaluation Protocol

We evaluate both base model and trained policy checkpoints using consistent inference settings. For QA tasks, pass@1 accuracy is estimated from 5 samples per query, leveraging the discriminative 4-option structure for exact binary rewards. Captioning uses LLM-as-Judge scoring (Gu et al., 2024) with modality-conditioned reference captions to detect hallucination errors and modality-fusion across RMTs. Scores lie in $[0, 1]$. To assess modality-aware post-training beyond standard accuracy, we introduce complementary metrics aligned with our core research problem. All metrics use two-stage normalization: per-sample superset averaging, then RMT-balanced macro-average.

- **Modality Gap:** Relative performance scaling across Uni/Bi/Tri-modal RMT groups.
- **Training Efficiency:** Wall-clock time per optimization step, normalized by batch size.
- **Fusion Gain:** Fraction of samples where multi-modal signal information outperform best uni-modal response.

These metrics directly quantify the stability, robustness, and efficiency claims from our introduction. Detailed prompts, judge model specifications, and additional analysis appear in Appendix C.

4. Modality-Aware Post-Training

4.1. RL Training Pipeline

Post-training optimizes policy π_θ via GRPO (Shao et al., 2024) without KL regularization. The old policy used for generation $\pi_{\theta_{\text{old}}}^{\text{gen}}$ produces G rollouts $\{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}^{\text{gen}}(\cdot|x)$ for query $x \sim \mathcal{D}$ distribution, each scored $r_i = r(x, y_i)$. Group advantages normalize within the generator batch:

$$\hat{A}_i^G = \frac{r_i - \mu_G}{\sigma_G}, \quad \mu_G = \frac{1}{G} \sum r_i, \quad \sigma_G = \text{std}(\{r_i\}).$$

Trainer $\pi_\theta^{\text{train}}$ optimizes the token-level objective with symmetric clipping. Importance sampling ratios are:

$$\rho_{i,t}(\theta) = \frac{\pi_\theta^{\text{train}}(y_{i,t} | x, y_{i,<t})}{\pi_{\theta_{\text{old}}}^{\text{gen}}(y_{i,t} | x, y_{i,<t})}, \quad \text{clip}(\rho, 1 - \epsilon, 1 + \epsilon)$$

The surrogate loss aggregates across all tokens:

$$\mathcal{L}_\theta^{\text{GRPO}, \mathcal{D}} = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}^{\text{gen}}(\cdot|x)} \left[\frac{1}{\sum_{i=1}^G |y_i|} \sum_{i=1}^G \sum_{t=1}^{|y_i|} \min \left(\rho_{i,t}(\theta) \hat{A}_i^G, \text{clip}(\rho_{i,t}(\theta), \epsilon) \hat{A}_i^G \right) \right] \quad (1)$$

Each query x requires modality subset $M_x \subseteq \mathcal{M} = \{\text{V}, \text{A}, \text{S}\}$ (video/audio/subtitles), the minimal signals needed (e.g., $M_x = \{\text{V}, \text{A}\}$ for audio-visual, $M_x = \{\text{S}\}$ for subtitles-only). Let R_M be rewards under requirement M . Different M yield heterogeneous R_M .

Modality-Unaware baseline. Ignores M_x , feeds full (V, A, S) signals, and mixes all R_M in batches $\mathcal{B} = \{(x_j, y_j, r_j)\}_{j=1}^B$. Therefore the gradient estimate takes the standard form:

$$\hat{g}_{\text{MU}} = \frac{1}{B} \sum_{j=1}^B \nabla_\theta \log \pi_\theta(y_j | x_j) \hat{A}_j$$

so rewards from heterogeneous R_M are normalized together as if drawn from a single distribution. This implicitly treats all modality regimes as homogeneous and inflates variance through between-subset differences in means and scales.

Modality-Aware scheme. A modality-aware scheme respects the requirement tags by forming per-subset (or compatible union) mini-batches

$$\mathcal{B}_M = \{(x_j, y_j, r_j) : M_{x_j} = M\}$$

and computing per-group advantages using only rewards from the same requirement set:

$$\hat{g}_{\text{MA}} = \sum_M \frac{1}{|\mathcal{B}_M|} \sum_{(x_j, y_j, r_j) \in \mathcal{B}_M} \nabla_\theta \log \pi_\theta(y_j | x_j) \hat{A}_j^{(M)}$$

where each $\hat{A}_j^{(M)}$ is normalized within R_M . Intuitively, this stratified normalization removes between-subset reward variance and aligns gradient updates with the difficulty and noise characteristics of each modality regime.

Under standard assumptions (independent samples and bounded score function $\nabla_\theta \log \pi_\theta$), the variance of the modality-unaware estimator \hat{g}_{MU} decomposes into within-subset and between-subset components, whereas the modality-aware estimator \hat{g}_{MA} retains only the within-subset term. As long as the reward distributions R_M differ across modality requirements (non-zero between-subset variance), this yields

$$\text{Var}(\hat{g}_{\text{MA}}) \leq \text{Var}(\hat{g}_{\text{MU}})$$

suggesting faster and more stable convergence when optimization is conditioned on required modalities instead of mixing all regimes in a modality-unaware manner (McKinzie et al., 2023b; Jiang et al., 2025a).

4.2. MAPLE’s Modality-Aware Policy Optimization

Building on the modality-aware training foundation above, we establish the optimal strategy for policy optimization.

With stratified batches \mathcal{B}_M , we systematically study four design axes—loss aggregation, clipping, sampling, and modality curriculum—each conditioned on RMTs M_x and compared against the modality-unaware baseline. We measure per-RMT accuracy, robustness gaps across complexity levels, and wall-clock convergence.

Baseline MAPO uses stratified GRPO:

$$\mathcal{L}_\theta^{\text{MAPO}} = \sum_M \frac{1}{|\mathcal{B}_M|} \mathcal{L}_\theta^{\text{GRPO}, \mathcal{B}_M} \quad (2)$$

The modality-aware gradient scales with per-group advantage: $|\hat{g}_{\text{MA}}| \propto |\hat{A}_j^{(M)}|$. Harder queries produce systematically lower advantages (smaller gradients), while easier queries yield stronger signals. To detect and compensate for this skew, we measure batch difficulty via KL divergence between empirical reward distribution p_{emp} and target $p_{\text{tgt}} \sim \text{Beta}(100, 1)$ (mode 0.99):

$$D_{\text{KL}}(p_{\text{emp}} \| p_{\text{tgt}})$$

Hard queries show $p_{\text{emp}} \ll p_{\text{tgt}}$ (large D_{KL}); easy queries converge to zero. We introduce per-RMT difficulty-adaptive weights $w_M \in (0, 1)$:

$$\mathcal{L}_\theta^{\text{MAPO}_{\text{adv}}} = \sum_{M \in \mathcal{M}} \frac{w_M}{|\mathcal{B}_M|} \mathcal{L}_\theta^{\text{GRPO}, \mathcal{B}_M} \quad (3)$$

Each w_M adapts via KL history \mathcal{H}_M over window of size \approx steps-per-epoch:

$$z_M = \frac{D_{\text{KL}, M} - \mu_{\mathcal{H}_M}}{\sigma_{\mathcal{H}_M} + \varepsilon}, \quad w_M = \text{sigmoid}(z_M) \quad (4)$$

This aggregates across tagged batches per training cycle for robust difficulty estimation.

Further, during each epoch we use KL statistics to arrange stratified batches by increasing complexity, counteracting gradient skew from reward-scale differences. Harder batches produce smaller gradient updates, so when they are naively interleaved with easier batches that have larger gradient norms, their learning signals are effectively overshadowed. We use KL statistics as dynamic curriculum signal, front-loading underperforming RMTs so the model learns their weaker signals first.

Formally, for each RMT M we maintain KL history \mathcal{H}_M and compute priority:

$$s_M = \frac{1}{|W_M|} \sum_{k \in W_M} D_{\text{KL}, M}^{(k)}$$

where $|W_M| = \min(L_W, |\mathcal{H}_M|)$ and $L_W = 5$. Higher s_M schedules harder tags earlier. For the first epoch—lacking KL history—we initialize via zero-shot performance, assigning higher priority to lower-accuracy tasks. See Appendix D for complete MAPO pseudocode.

5. Experiments & Results

5.1. Experimental Setup

We conduct all experiments using the Qwen2.5-Omni-3B (Jin Xu, 2025) multimodal language model, which natively processes text, audio, and video inputs. Training uses a sequence length of 10,240 tokens (8,192 input + 2,048 response) and the veRL (Sheng et al., 2025) library for RL optimization.

Training hyperparameters: In our setup, we trained for multiple epochs over the training prompts with AdamW optimizer (Loshchilov & Hutter, 2017) and a constant learning rate 2×10^{-6} . We used global batch size 256 (mini-batch 32, yielding 8 updates per rollout). Rollouts generate $G = 8$ responses per prompt with temperature 1.0, top- $p = 1.0$, within a 8,096-token context limit. All runs employ mixed precision with the FP32 language model head fix (Chen et al., 2025) to eliminate generator-trainer probability mismatches that destabilize importance sampling ratios.

Training runs were conducted on $4 \times$ NVIDIA H100-80GB nodes. We first establish a strong modality-unaware baseline (MUPO) that always receives full (V, A, S) signals, leveraging multimodal redundancy for robust task solving. This follows standard post-training practice: token-level importance sampling and loss aggregation with symmetric clipping (± 0.2), full-dataset on-policy updates, and per-group advantage normalization. All experiments report pass@1 accuracy estimated from 5 samples per query, with other hyperparameter sweeps detailed in Appendix E. On MAPLE-QA, this baseline achieves an average task performance of 58.58%, while on MAPLE-Caption achieves 67.67%.

5.2. Optimal Training Strategy

In this section, we conduct experiments on MAPLE-QA due to its verifiable 4-option multiple-choice rewards, enabling precise algorithmic analysis without reward hacking concerns.

We first establish MAPO’s effectiveness using the same configuration as our baseline MUPO, differing only in modality-stratified batching \mathcal{B}_M . Despite receiving strictly limited signals matching each query’s RMT, MAPO achieves 58.68% pass@1, better than baseline performance while reducing policy gradient variance by 12.89%. Building on this foundation, we systematically optimize four algorithmic axes critical to multimodal RL stability. To validate that these choices remain optimal when combined, we conduct leave-one-out experiments. Each variant improves accuracy, robustness (uni/multi-modal gaps), or training efficiency, with results shown in Table 1 and all training characteristic curves are mentioned in Appendix E.

Loss aggregation. Loss aggregation controls credit assign-

Table 1. Results of MAPLE’s four design axes with MAPLE-QA on Qwen2.5-Omni-3B model, studying loss aggregation, clipping, dynamic sampling, and curriculum learning. Each row optimizes one axis atop modality-stratified MAPO while others follow standard settings. MAPLE’s optimal recipe achieves average accuracy of 58.72%, lowest modality gap (1.74%), and fastest training (164.72 s/step).

Method	Variant	Modality Accuracy (Pass@1 %)								Modality Gap (%)				Efficiency
		V	A	S	VA	VS	AS	VAS	Avg	U-B	U-T	B-T	Avg	Time (s/step)
	<i># Samples</i>	552	75	76	1145	1784	171	1198	5001					
Zero-shot		34.54	36.89	43.42	37.79	40.83	42.30	40.46	39.38	11.30	13.17	1.68	8.72	–
MUPO		55.08	65.34	63.82	57.47	60.15	58.77	58.14	58.58	3.43	1.79	1.59	2.27	523.28
MAPO														
Loss aggregation	Token-level	55.79	65.78	61.40	57.85	60.26	59.06	57.76	58.68	3.20	0.52	2.60	2.11	389.11
	Sample-level	55.74	65.78	62.28	58.69	60.42	58.87	57.48	58.86	3.79	0.07	3.72	2.53	395.03
Clipping	Asymmetric	55.44	68.89	57.89	58.02	60.65	57.70	57.18	58.62	4.17	0.07	3.93	2.72	379.40
Dynamic Sampling	Early filtering	55.50	64.89	63.16	64.63	59.70	58.48	57.12	58.00	2.04	0.37	2.36	1.59	197.02
	Late filtering	55.32	69.33	61.84	57.29	60.18	59.65	57.10	58.39	2.71	0.73	3.35	2.26	265.13
Curriculum learning	Modality-based	56.40	65.78	60.53	58.05	60.83	59.65	57.96	59.05	3.27	0.19	2.98	2.15	424.03
	Full-recipe	56.22	67.55	63.16	57.50	60.58	56.92	57.71	58.72	1.82	0.81	2.58	1.74	164.72




Figure 3. MAPLE algorithmic variations across four design axes: (a) Sample-level loss aggregation outperforms others, (b) Asymmetric clipping maintains low clip fractions, (c) Dynamic filtering accelerates reward gains and (d) Curriculum improves accuracy by modality complexity.

ment across heterogeneous signal regimes. We evaluate three strategies on MAPLE-QA: (1) *token-level* (average across all tokens in batch); (2) *sample-level* (per-sequence mean, then batch-average); (3) *prompt-level* (fixed sum per sequence). Sample-level aggregation achieves 58.86% pass@1, outperforming token-level (58.68%). Token granularity dilutes per-query signal structure, while prompt-level causes divergence (Appendix E). Sample-level preserves query-specific credit assignment while remaining computationally tractable.

Asymmetric clipping. Symmetric clipping ($\epsilon = 0.2$) is a common choice for post-training with RL pipeline but it destabilizes multimodal training due to varying information content across signal combinations. We adopt asymmetric thresholds $\epsilon^+ = 0.3$, $\epsilon^- = 0.2$, expanding the upper trust region for positive advantages while maintaining conservative negative updates. This reduces clip fractions by 71.65% as shown in Fig. 3b, yielding healthier entropy and smoother policy losses compared to symmetric clipping.

Dynamic sampling. Some prompts produce identical rewards across all $G = 8$ generations (zero-variance samples) that contribute nothing to policy gradients ($\nabla_{\theta} \log \pi_{\theta} \hat{A}_i^G = 0$). Baseline training includes them, artificially inflating effective batch size while diluting signal. Early filtering at batch construction accelerates wall-clock training by $1.98\times$

(58.00% accuracy) due to reduced computation per step. Mid-training filtering (training on full dataset for 1 epoch, then filtering) recovers accuracy to 58.39% but slows training by $0.35\times$ relative to early filtering. Applying any filtering has nearly no change in entropy reduction as shown in Fig. 3c.

Curriculum learning. We sequence training dataset by signal complexity: uni-modal \rightarrow bi-modal \rightarrow tri-modal queries. Curriculum training produces cleaner gradient magnitudes and monotonically increasing rollout rewards compared to mixed stratified training (Fig. 3d). Observing first epoch, uni-modal batches show gradual reward gains under curriculum ordering, while bi- and tri-modal batches accelerate sharply, outpacing non-curriculum settings. This pattern persists into epoch 2, confirming effective policy adaptation to increasing signal complexity. Final accuracy reaches 59.05%, validating that structured progression prevents overfitting to simple signal patterns.

To validate our optimizations, leave-one-out experiments revert each axis individually to MUPO baseline while keeping others active, showing every component contributes positively with no destructive interactions. From these analyses, we consolidate the optimal configuration, our complete **Full-recipe**, using sample-level loss aggregation, asymmetric clipping with a wider upper trust region, early zero-variance

Table 2. MAPLE-QA results on the Qwen2.5-Omni-3B model. The table reports modality-wise accuracy for MAPO adaptive strategies. The configuration with reweighting and adaptive curriculum ($adp_w + adp_{cur}$) achieves the highest average accuracy (59.82%).

Method	Modality Accuracy (Pass@1 %)							Avg
	V	A	S	VA	VS	AS	VAS	
Zero-shot	34.03	33.60	42.16	38.58	41.40	43.45	40.86	39.78
MUPO	55.08	65.34	63.82	57.47	60.15	58.77	58.14	58.58
MAPO	55.79	65.78	61.40	57.85	60.26	59.06	57.76	58.68
+ adp_w	55.34	67.47	62.16	58.28	60.77	58.57	57.96	58.98
+ cur	56.40	65.78	60.53	58.05	60.83	59.65	57.96	59.05
+ adp_{cur}	56.87	65.33	61.89	58.62	61.00	59.52	58.27	59.38
+ $adp_w + cur$	57.24	65.87	59.46	59.37	61.63	58.81	58.13	59.73
+ $adp_w + adp_{cur}$	57.13	66.40	61.35	59.16	61.64	58.33	58.67	59.82

Table 3. MAPLE-Caption results on the Qwen2.5-Omni-3B model. The table reports modality-wise accuracy for MAPO adaptive strategies. The configuration with reweighting and adaptive curriculum ($adp_w + adp_{cur}$) achieves the highest average accuracy (74.00%).

Method	Modality Accuracy (LLM-as-Judge score)							Avg
	V	A	S	VA	VS	AS	VAS	
Zero-shot	56.41	63.25	61.98	59.51	59.10	63.63	58.38	60.32
MUPO	63.78	73.69	68.20	64.46	65.69	71.71	65.18	67.67
MAPO	66.78	83.46	87.50	66.44	67.99	78.59	66.40	73.88
+ adp_w	65.65	82.39	85.82	65.89	67.72	77.57	65.18	72.89
+ cur	65.03	81.65	86.39	64.71	66.21	76.83	63.67	72.07
+ adp_{cur}	65.11	79.89	86.14	65.12	66.94	76.26	64.17	71.95
+ $adp_w + cur$	66.43	82.91	85.57	66.30	68.48	78.23	66.38	73.47
+ $adp_w + adp_{cur}$	66.20	81.42	84.97	66.98	69.10	81.71	67.58	74.00

filtering at batch construction to prioritize speed, and curriculum learning that sequences training by increasing signal complexity. Compared to the modality-unaware baseline, this recipe achieves 58.72% accuracy, the lowest modality gap of 1.74%, and most efficient gains (164.72 secs per training step) that is 3.18x faster than MUPO. Training characteristic curves for each setting appear in Appendix E.

5.3. Adaptive Training Strategy

Having established the optimal static training configuration, we next explore adaptive strategies to further boost accuracy by addressing gradient skewness observed in multimodal RL. As detailed in Sec. 4.2, we evaluate two adaptation techniques: (1) adp_w : batch weighting by historical KL divergence between empirical and target reward distributions; (2) adp_{cur} : KL-based dynamic curriculum adaptation. Since our static curriculum already delivers the strongest gains on MAPLE-QA [Table 1], we include it (cur) for fair comparison. We also test combinations: $adp_w + cur$ (static curriculum + weighting) and $adp_w + adp_{cur}$ (dynamic curriculum + weighting). Results across both benchmarks from Sec. 3 appear in Tables 2 and 3.

For QA with discrete rewards, MAPO baseline achieves 58.68% average accuracy, improving to 59.82% with adaptive weighting and dynamic curriculum ($adp_w + adp_{cur}$). Hard modalities gain most (V: 55.79% \rightarrow 57.13%; VA:




Figure 4. Tag-wise disparity heatmaps across MAPLE-QA-eval. Left: Basic MAPO reduces cross-tag variance. Right: Full adaptive MAPO ($adp_w + adp_{cur}$) achieves minimal disparity (smaller values better).




Figure 5. Tag-wise disparity heatmaps across MAPLE-Caption-eval. Left: Basic MAPO reduces cross-tag variance. Right: Full adaptive MAPO ($adp_w + adp_{cur}$) achieves minimal disparity (smaller values better).

57.85% \rightarrow 59.16%) while easy text tags remain comparable, showing reweighting corrects scale imbalances and adaptive curriculum controls update timing. Captioning also shows better performance with continuous rewards: from MAPO’s 73.88% baseline, it reaches 74.00% while lifting heavy tags (VS: 67.99% \rightarrow 69.10%; AS: 78.59% \rightarrow 81.71%; VAS: 66.40% \rightarrow 67.58%), confirming both scale (reweighting) and order (adaptive curriculum) corrections counteract easy-hard skew.

Our full adaptive MAPO strategy consistently outperforms all variants per-sample, achieving the highest average accuracy across both benchmarks. This combination of reweighting and adaptive curriculum secures top performance across modality subsets and demonstrates superior stability over basic MAPO. This dominance beyond aggregate averages, establishes the full adaptive recipe for heterogeneous multimodal training. MAPLE’s full adaptive strategy substantially enhances multimodal fusion ability, with 30.24% of MAPLE-Caption samples showing multi-modal captions outperforming the best uni-modal baseline, the 18.19% achieved by basic MAPO. Heatmap analysis further confirms reduced cross-modality disparities: adaptive strategies match or exceed base MAPO’s variance reduction on QA

Table 4. MAPLE-QA+ ablation: MAPO trained across modality-exact, -superset, and -deficit configurations. Model generalizes beyond fixed RMT patterns, correctly abstaining via “None” option when evidence is insufficient.

Training dataset	Modality Accuracy (Pass@1 %)							Avg
	V	A	S	VA	VS	AS	VAS	
Zero-shot	56.26	52.19	52.12	56.05	51.46	50.34	47.94	52.19
MAPLE-QA	62.58	61.52	57.72	54.93	53.99	44.90	40.91	51.90
MAPLE-QA+	72.68	78.67	72.20	79.72	78.31	74.23	74.98	76.99

(Fig. 4), while captioning also exhibits the same pattern. Together, these metrics demonstrate that adaptive MAPO optimally balances overall accuracy, tag-wise robustness, per-sample consistency, and genuine multimodal integration by addressing both *how much* (reweighting) and *when* (adaptive curriculum) to update under heterogeneous signal distributions.

5.4. Ablations

5.4.1. DATA AUGMENTATION

To investigate whether MAPO simply overfits to the exact modality patterns seen during training or truly learns modality-aware reasoning, we construct an augmented QA dataset that perturbs the mapping between queries and their available modalities. Instead of always supplying only the modalities specified as the required modalities (RMT), we generate three strata for each query: 1) modality-exact (information exactly matches the RMT specification), 2) modality-superset (additional modalities are provided beyond those indicated by the RMT), and 3) modality-deficit (one or more RMT-specified modalities are removed). We refer to this augmented dataset as MAPLE-QA+, which contains 137,313 QA pairs with 103,265 and 34,048 samples as a train–test split, partitioned across the three strata. To further discourage superficial pattern matching, we also introduce a “None” option. For 25% of modality-exact and modality-superset examples, the original correct answer is replaced by this new option. Also for all modality-deficit cases the “None” option is treated as correct.

Table 4 shows MAPO on MAPLE-QA+ achieving 77.0% average accuracy across all modality combinations vs. modality-exact MAPLE-QA with largest gains on mixed/incomplete settings. QA’s discrete rewards isolate true modality-awareness from scoring noise. This confirms MAPO learns robust reasoning under missing/redundant signals, correctly abstaining via “None” when evidence is insufficient.

5.4.2. CONTRASTIVE TRAINING

Despite MAPO’s gains, captioning policies still suffer from *mode collapse*: high-reward captions remain nearly identical




Figure 6. Intra-group dispersion histogram showing a rightward shift of the distribution under CRW.

regardless of input modality (VAS vs. VA vs. V-only), failing to reflect available signals. We address this through Contrastive Reward Weighting (CRW), which upweights high-reward “positives” from complete modalities only when they maintain embedding separation from deliberately-generated “negatives” under modality-deficit conditions for the same prompt (Appendix E.3). CRW is applied exclusively to captioning, where continuous $[0, 1]$ rewards enable rich contrastive gradients unlike QA’s binary signals. CRW improves fusion ability from 18.19% to 18.46% (multimodal > best uni-modal captions), intra-group dispersion from 0.35 to 0.36 (greater caption diversity within each RMT), and inter-group separation from 0.41 to 0.46 (clearer modality-conditioned manifolds). Pairwise distance histograms shift rightward and t-SNE projections show more distinct modality clusters (Figs. 6, 16), confirming CRW simultaneously boosts within-condition diversity and cross-condition separation for truly modality-aware generation.

Conclusion

MAPLE addresses a fundamental oversight in multimodal RL post-training: treating heterogeneous signal combinations as homogeneous training data. Through MAPLE-bench, MAPO, and adaptive strategies, we achieve highest performance under realistic reduced-signal conditions, while converging faster and reducing uni- to multi-modal gaps. This work reveals that modality-stratified optimization eliminates between-group variance, a pervasive yet unrecognized source of instability in omni-modal training. MAPLE constitutes a complete, production-ready recipe that will enable robust multimodal RL deployment across diverse real-world signal configurations, establishing modality-aware post-training as essential infrastructure for next-generation multimodal reasoning systems.

Impact Statement

This paper advances the field of machine learning through MAPLE, a modality-aware RL post-training ecosystem that improves efficiency, stability, and robustness for multimodal language models. MAPLE enables more reliable deployment of multimodal reasoning systems under realistic partial-signal conditions commonly encountered in production environments. While our technical contributions focus on benchmark dataset creation and algorithmic improvements, we note potential downstream applications in accessibility (e.g., robust audio-visual processing for assistive technologies) and content analysis across diverse signal availability. We do not foresee specific ethical concerns beyond standard responsible AI practices for improved model performance, and no negative societal consequences are anticipated from these optimization techniques.

References

- Abouelenin, A., Ashfaq, A., Atkinson, A., Awadalla, H., Bach, N., Bao, J., Benhaim, A., Cai, M., Chaudhary, V., Chen, C., Chen, D., Chen, D., Chen, J., Chen, W., Chen, Y.-C., ling Chen, Y., Dai, Q., Dai, X., Fan, R., Gao, M., Gao, M., Garg, A., Goswami, A., Hao, J., Hendy, A., Hu, Y., Jin, X., Khademi, M., Kim, D., Kim, Y. J., Lee, G., Li, J., Li, Y., Liang, C., Lin, X., Lin, Z., Liu, M., Liu, Y., Lopez, G., Luo, C., Madan, P., Mazalov, V., Mitra, A., Mousavi, A., Nguyen, A., Pan, J., Perez-Becker, D., Platin, J., Portet, T., Qiu, K., Ren, B., Ren, L., Roy, S., Shang, N., Shen, Y., Singhal, S., Som, S., Song, X., Sych, T., Vaddamanu, P., Wang, S., Wang, Y., Wang, Z., Wu, H., Xu, H., Xu, W., Yang, Y., Yang, Z., Yu, D., Zabir, I., Zhang, J., Zhang, L. L., Zhang, Y., and Zhou, X. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*, 2025. URL <https://arxiv.org/abs/2503.01743>.
- Anonymous. GOAL: Balance multimodal learning with gradient orthogonalization and adaptive leveraging. In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=I3uFqoUZ2Y>. under review.
- Che, F., Vasan, G., and Mahmood, A. R. Correcting discount-factor mismatch in on-policy policy gradient methods. In *International Conference on Machine Learning*, pp. 4218–4240. PMLR, 2023.
- Chen, A., Li, A., Gong, B., Jiang, B., Fei, B., Yang, B., Shan, B., Yu, C., Wang, C., Zhu, C., et al. Minimax-m1: Scaling test-time compute efficiently with lightning attention. *arXiv preprint arXiv:2506.13585*, 2025.
- Chen, S., Li, H., Wang, Q., Zhao, Z., Sun, M., Zhu, X., and Liu, J. Vast: A vision-audio-subtitle-text omni-modality foundation model and dataset. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL <https://arxiv.org/abs/2305.18500>.
- Cheng, Z., Leng, S., Zhang, H., Xin, Y., Li, X., Chen, G., Zhu, Y., Zhang, W., Luo, Z., Zhao, D., and Bing, L. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. URL <https://arxiv.org/abs/2406.07476>.
- Comanici, G., Bieber, E., Schaeckermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Gu, J., Jiang, X., Shi, Z., Tan, H., Zhai, X., Xu, C., Li, W., Shen, Y., Ma, S., Liu, H., et al. A survey on llm-as-a-judge. *The Innovation*, 2024.
- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Jiang, Q., Huang, L., and Yang, Y. Rethinking multimodal learning from the perspective of mitigating classification ability disproportion. *arXiv preprint arXiv:2502.20120*, 2025a.
- Jiang, S., Liang, J., Wang, J., Dong, X., Chang, H., Yu, W., Du, J., Liu, M., and Qin, B. From specific-mlms to omni-mlms: a survey on mlms aligned with multi-modalities. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 8617–8652, 2025b.
- Jin Xu, Zhifang Guo, J. H. H. H. T. H. S. B. K. C. J. W. Y. F. K. D. B. Z. X. W. Y. C. J. L. Qwen2.5-omni technical report. *arXiv preprint arXiv:2503.20215*, 2025.
- Khatri, D., Madaan, L., Tiwari, R., Bansal, R., Duvvuri, S. S., Zaheer, M., Dhillon, I. S., Brandfonbrener, D., and Agarwal, R. The art of scaling reinforcement learning compute for llms. *arXiv preprint arXiv:2510.13786*, 2025.
- Li, C., Chen, Y., Ji, Y., Xu, J., Cui, Z., Li, S., Zhang, Y., Tang, J., Song, Z., Zhang, D., He, Y., Liu, H., Wang, Y., Wang, Q., Wu, Z., Luo, J., Pan, Z., Xie, W., Zhang, C.,

- 495 Wang, Z., Tian, J., Wang, Y., Cao, Z., Dai, M., Wang, K.,
496 Wen, R., Ma, Y., Pan, Y., Chang, S., Taheri, T., Xia, H.,
497 Plachouras, C., Benetos, E., Li, Y., Zhang, G., Yang, J.,
498 Peng, T., Wang, Z., Liu, M., Peng, J., Zhang, Z., and Liu,
499 J. Omnivideobench: Towards audio-visual understanding
500 evaluation for omni mlms, 2025a. URL [https://](https://arxiv.org/abs/2510.10689)
501 arxiv.org/abs/2510.10689.
- 502 Li, S., Hu, Y., Ning, X., Liu, X., Hong, K., Jia, X., Li, X.,
503 Yan, Y., Ran, P., Dai, G., et al. Mbq: Modality-balanced
504 quantization for large vision-language models. In *Pro-*
505 *ceedings of the Computer Vision and Pattern Recognition*
506 *Conference*, pp. 4167–4177, 2025b.
- 507 Li, Y., Ma, Y., Zhang, G., Yuan, R., Zhu, K., Guo, H.,
508 Liang, Y., Liu, J., Wang, Z., Yang, J., Wu, S., Qu,
509 X., Shi, J., Zhang, X., Yang, Z., Wen, Y., Wang, Y.,
510 Li, S., Zhang, Z., Liu, Z., Benetos, E., Huang, W.,
511 and Lin, C. Omnibench: Towards the future of uni-
512 versal omni-language models, 2025c. URL [https:](https://arxiv.org/abs/2409.15272)
513 [//arxiv.org/abs/2409.15272](https://arxiv.org/abs/2409.15272).
- 514 Liu, S.-Y., Dong, X., Lu, X., Diao, S., Belcak, P., Liu,
515 M., Chen, M.-H., Yin, H., Wang, Y.-C. F., Cheng, K.-T.,
516 et al. Gdpo: Group reward-decoupled normalization pol-
517 icy optimization for multi-reward rl optimization. *arXiv*
518 *preprint arXiv:2601.05242*, 2026.
- 519 Liu, Y., Duan, H., Zhang, Y., Li, B., Zhang, S., Zhao,
520 W., Yuan, Y., Wang, J., He, C., Liu, Z., Chen, K., and
521 Lin, D. Mmbench: Is your multi-modal model an all-
522 around player?, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2307.06281)
523 [abs/2307.06281](https://arxiv.org/abs/2307.06281).
- 524 Liu, Z., Chen, C., Li, W., Qi, P., Pang, T., Du, C., Lee,
525 W. S., and Lin, M. Understanding r1-zero-like training:
526 A critical perspective. *arXiv preprint arXiv:2503.20783*,
527 2025a.
- 528 Liu, Z., Liu, J., He, Y., Wang, W., Liu, J., Pan, L., Hu,
529 X., Xiong, S., Huang, J., Hu, J., et al. Part i: Tricks or
530 traps? a deep dive into rl for llm reasoning. *arXiv preprint*
531 *arXiv:2508.08221*, 2025b.
- 532 Loshchilov, I. and Hutter, F. Decoupled weight decay regu-
533 larization. *arXiv preprint arXiv:1711.05101*, 2017.
- 534 McKinzie, B., Cheng, J., Shankar, V., Yang, Y., Shlens, J.,
535 and Toshev, A. Robustness in multimodal learning under
536 train-test modality mismatch. In *ICML*, 2023a. URL
537 <https://arxiv.org/abs/2304.04385>.
- 538 McKinzie, B., Shankar, V., Cheng, J. Y., Yang, Y., Shlens,
539 J., and Toshev, A. T. Robustness in multimodal learning
540 under train-test modality mismatch. In *International Con-*
541 *ference on Machine Learning*, pp. 24291–24303. PMLR,
542 2023b.
- 543 Park, J. R., Kim, J., Kim, G., Jo, J., Choi, S., Cho, J., and
544 Ryu, E. K. Clip-low increases entropy and clip-high
545 decreases entropy in reinforcement learning of large lan-
546 guage models. *arXiv preprint arXiv:2509.26114*, 2025.
- 547 Pena, N. et al. Intracurriculum learning for multimodal
548 tasks. *ICLR*, 2023. Focuses on sequencing multimodal
549 task complexity.
- 550 Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D.,
551 Ermon, S., and Finn, C. Direct preference optimiza-
552 tion: Your language model is secretly a reward model.
553 *Advances in neural information processing systems*, 36:
554 53728–53741, 2023.
- 555 Rha, H., Yeo, J. H., Won, J., Park, S. J., and Ro, Y. M. Learn-
556 ing what to attend first: Modality-importance-guided rea-
557 soning for reliable multimodal emotion understanding.
558 *arXiv preprint arXiv:2512.02699*, 2025.
- 559 Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and
560 Klimov, O. Proximal policy optimization algorithms.
561 *arXiv preprint arXiv:1707.06347*, 2017.
- 562 Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang,
563 H., Zhang, M., Li, Y., Wu, Y., et al. Deepseekmath: Push-
564 ing the limits of mathematical reasoning in open language
565 models. *arXiv preprint arXiv:2402.03300*, 2024.
- 566 Sheng, G., Zhang, C., Ye, Z., Wu, X., Zhang, W., Zhang, R.,
567 Peng, Y., Lin, H., and Wu, C. Hybridflow: A flexible and
568 efficient rlhf framework. In *Proceedings of the Twentieth*
569 *European Conference on Computer Systems*, pp. 1279–
570 1297, 2025.
- 571 Sun, G., Yu, W., Tang, C., Chen, X., Tan, T., Li, W., Lu,
572 L., Ma, Z., Wang, Y., and Zhang, C. video-salmonn:
573 Speech-enhanced audio-visual large language models.
574 *arXiv preprint arXiv:2406.15704*, 2024. URL [https:](https://arxiv.org/abs/2406.15704)
575 [//arxiv.org/abs/2406.15704](https://arxiv.org/abs/2406.15704).
- 576 Vasan, G., Elsayed, M., Azimi, S. A., He, J., Shahriar, F.,
577 Bellinger, C., White, M., and Mahmood, R. Deep policy
578 gradient methods without batch updates, target networks,
579 or replay buffers. *Advances in Neural Information Pro-*
580 *cessing Systems*, 37:845–891, 2024.
- 581 Wei, L., Li, Y., Wang, C., Wang, Y., Kong, L., Huang, W.,
582 and Sun, L. Unsupervised post-training for multi-modal
583 llm reasoning via grpo. *arXiv preprint arXiv:2505.22453*,
584 2025.
- 585 Xiang, B., Zhao, S., Guo, T., and Zou, W. Understanding
586 the modality gap: An empirical study on the speech-
587 text alignment mechanism of large speech language mod-
588 els. In *Proceedings of the 2025 Conference on Empiri-*
589 *cal Methods in Natural Language Processing*, pp. 5187–
590 5202, 2025.

- 550 Xie, L., Kuthiala, A., Wei, G. Z., Zheng, C., Bal, A.,
551 Dabhi, M., Wen, L., Rustagi, T., Lai, E., Khyalia, S.,
552 Choudhury, R., Ziyadi, M., Zhang, X., Yang, H., and
553 Jeni, L. A. Maverix: Multimodal audio-visual eval-
554 uation and recognition index, 2025. URL <https://arxiv.org/abs/2503.21699>.
555
- 556 Xie, Z. and Wu, C. Mini-omni2: Towards open-source gpt-
557 4o with vision, speech and duplex capabilities. *arXiv*
558 *preprint arXiv:2410.11190*, 2024. URL <https://arxiv.org/abs/2410.11190>.
559
- 560 Xu, J., Guo, Z., He, J., Hu, H., He, T., Bai, S., Chen, K.,
561 Wang, J., Fan, Y., Dang, K., Zhang, B., Wang, X., Chu,
562 Y., and Lin, J. Qwen2.5-omni technical report. *arXiv*
563 *preprint arXiv:2503.20215*, 2025a. URL <https://arxiv.org/abs/2503.20215>.
564
- 565 Xu, J., Guo, Z., He, J., Hu, H., He, T., Bai, S., Chen, K.,
566 Wang, J., Fan, Y., Dang, K., et al. Qwen2. 5-omni techni-
567 cal report. *arXiv preprint arXiv:2503.20215*, 2025b.
568
- 569 Yamaguchi, S., Feng, D., Kanai, S., Adachi, K., and Chijiwa,
570 D. Post-pre-training for modality alignment in vision-
571 language foundation models. In *Proceedings of the Com-*
572 *puter Vision and Pattern Recognition Conference*, pp.
573 4256–4266, 2025.
574
- 575 Ye, Q., Xu, H., Ye, J., Yan, M., Hu, A., Liu, H., Qian,
576 Q., Zhang, J., and Huang, F. mplug-owl2: Revolution-
577 izing multi-modal large language model with modality
578 collaboration. In *Proceedings of the ieee/cvf conference*
579 *on computer vision and pattern recognition*, pp. 13040–
580 13051, 2024.
581
- 582 Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., and
583 Chen, E. A survey on multimodal large language models.
584 *arXiv preprint arXiv:2306.13549*, 2023. URL <https://arxiv.org/abs/2306.13549>. v4, Nov. 2024.
585
- 586 Yu, Q., Zhang, Z., Zhu, R., Yuan, Y., Zuo, X., Yue, Y., Dai,
587 W., Fan, T., Liu, G., Liu, L., et al. Dapo: An open-source
588 llm reinforcement learning system at scale. *arXiv preprint*
589 *arXiv:2503.14476*, 2025.
590
- 591 Zhang, Y., Luo, Z., Yan, Q., He, W., Jiang, B., Chen, X., and
592 Han, K. Omnieval: A benchmark for evaluating omni-
593 modal models with visual, auditory, and textual inputs.
594 *arXiv preprint arXiv:2506.20960*, 2025.
595
- 596 Zhao, T., Zhang, L., Ma, Y., and Cheng, L. A survey on safe
597 multi-modal learning systems. In *Proceedings of the 30th*
598 *ACM SIGKDD Conference on Knowledge Discovery and*
599 *Data Mining*, pp. 6655–6665, 2024.
600
- 601 Zhou, Z., Wang, R., and Wu, Z. Daily-omni: Towards
602 audio-visual reasoning with temporal alignment across
603 modalities, 2025. URL <https://arxiv.org/abs/2505.17862>.
604

A. Extended Discussion of Related Literature

A.1. Multimodal LLM Architectures

Recent MLLMs follow two dominant architectural paradigms: *deep fusion* via cross-attention within LLM layers (Type-A/B) versus *early fusion* at the input stage (Type-C/D).

Deep fusion approaches (Type-A) employ standard cross-attention, as in early VideoLLaMA 2 (Cheng et al., 2024), while Type-B uses custom layers like mPLUG-Owl2’s modality-adaptive attention (Ye et al., 2024). **Early fusion** dominates recent work: Type-C uses non-tokenized modality encoders (e.g., CLIP-ViT + Whisper feeding directly into LLaMA (Sun et al., 2024)), while Type-D applies discrete tokenization across modalities (e.g., Qwen2.5-Omni’s TMRoPE for synchronized audio-visual streaming (Xu et al., 2025a)).

Industry systems emphasize end-to-end omni-modality: Gemini-3 offers native multimodal I/O with developer controls for reasoning depth (Comanici et al., 2025); Mini-Omni2 implements staged vision-to-speech alignment (Xie & Wu, 2024); and Phi-4-Multimodal uses Mixture-of-LoRAs for parameter-efficient vision-speech augmentation (Abouelenin et al., 2025). Despite architectural sophistication, all assume static, fully-available modality combinations during post-training (Yin et al., 2023).

A.2. Detailed Benchmark Comparison

Table 5. Comparison of multimodal benchmarks. I: Image, V: Video, A: Audio, T: Text. **Cond. Output** indicates whether valid QA answers or captions vary given identical VAT instances under different modality combinations.

Benchmark	Modality	Task	Size (#Q)	Cond. Output
MME	I, T	QA (Y/N)	–	No
MMBench	I, T	QA	2.9k	No
MM-Vet (v2)	I, T	Free-form	0.5k	No
SEED-Bench-2	I / V, T	QA	24k	No
OmniBench	I, A, T	QA	1.14k	No
MuAViC	A, V, T	Free-form	1.1M	No
AVHBench	A, V, T	QA, Free-form	5.3k, 1.1k	No
MAVERIX	A, V, T	QA	2.5k	No
OmniVideoBench	A, V, T	QA	1k	No
JointAVBench	A, V, T	QA	2.8k	No
MAPLE-bench (ours)	A, V, T	QA, Free-form	≈ 5k, ≈ 5k	Yes

Existing benchmarks reveal two critical limitations: (i) *fixed ground truth* despite varying inputs, and (ii) lack of explicit *modality-deficit awareness*. For example, OmniBench (Li et al., 2025c) varies I/A/T inputs but maintains identical QA options/answers; OmniVideoBench (Li et al., 2025a) offers A/V-only ablations under static problem definitions. Neither distinguishes tasks unsolvable due to true information gaps from fusion failures, nor adapts caption references to available signals.

MAPLE-bench addresses this through RMTs and combination-conditioned outputs, enabling fine-grained diagnosis of modality dependence.

A.3. Post-training Method Taxonomy

Post-training has evolved from PPO-based RLHF (Schulman et al., 2017) toward value-model-free methods:

- **GRPO family:** Group Relative Policy Optimization (Shao et al., 2024) eliminates critics via multi-sample group normalization; Dr. GRPO (Liu et al., 2025a) and DAPO (Yu et al., 2025) refine clipping and aggregation.
- **Scale-focused:** ScaleRL (Khatri et al., 2025) drops KL regularization for 100k+ GPU scaling; GDPO (Liu et al., 2026) decouples the normalization of individual rewards, more faithfully preserving their relative differences
- **Unimodal ablations:** Recent analysis (Park et al., 2025; Liu et al., 2025b) dissects loss aggregation and asymmetric clipping on reasoning benchmarks but excludes multimodal data.

Modality-specific strategies exist but remain disconnected from core RL optimization: MBQ (Li et al., 2025b) balances signal contributions pre-RLHF; IntraCurriculum (Pena et al., 2023) sequences task complexity without per-query signal tagging. No prior work conditions batch composition, reward normalization, or clipping on minimal required signals—the core innovation of MAPLE’s MAPO.

A.4. Robustness and Modality Gap Evidence

Recent analyses reveal three fundamental limitations of modality-agnostic training that MAPLE directly addresses:

1. Train-Test Modality Mismatch Omni-modal training often underperforms uni-modal specialists under deployment shifts. McKinzie et al. (2023b) documents large gaps between average and worst-case performance across modality configurations, with follow-ups (McKinzie et al., 2023a) confirming additional modalities can be under-utilized or even harmful without targeted interventions. Jiang et al. (2025b) similarly shows omni-MLLMs trail specialists on per-modality benchmarks due to mixed-alignment interference.

2. Modality Imbalance and Dominance Strong modalities systematically dominate weaker ones during training. Jiang et al. (2025a) demonstrates this causes lagging performance even when all signals are present, while explicit metrics reveal average vs. worst-case gaps over train-test setups. MBQ (Li et al., 2025b) confirms symmetric treatment (quantization, optimization) degrades performance—modality-balanced allocation is essential. CLIP-Refine (Yamaguchi et al., 2025) shows persistent vision-language “modality gaps” post-training under naive omni approaches.

3. Optimization Instability Across Configurations Even advanced post-training exhibits uneven gains. Safe multimodal surveys (Zhao et al., 2024) highlight vulnerability to imbalance; unsupervised GRPO-style methods (Wei et al., 2025) yield inconsistent improvements despite significant compute. OmniEval (Zhang et al., 2025) confirms wide performance variance across modality combinations despite strong aggregate scores.

MAPLE resolves these through per-query Required Modality Tags: MAPLE-bench provides tagged supervision, MAPO conditions batching/normalization/clipping on minimal signal subsets, and curriculum learning sequences by signal sparsity.

B. Dataset generation details

B.1. MAPLE-QA task

We performed several refinement and validation steps to ensure high-quality, modality-dependent supervision. Within the generation pipeline outlined in Fig 2, all captions and annotations undergo cross-modal consistency checks. Respective captions are then temporally aligned across audio and visual streams to maintain coherent event ordering and sub-title generation. Our dataset generation pipeline follows the automated annotation and optimization framework of Daily-Omni (Zhou et al., 2025). Specifically, we revised the Advanced QA system prompt as described below B.3, while following the remaining configuration from Daily-Omni (Zhou et al., 2025).

Using the aligned subtitles, we generate QA pairs with multiple tag label variants, subsequently refined with state-of-the-art models (e.g., GPT-4o, GPT-5 (Hurst et al., 2024)) to improve linguistic clarity and logical consistency. This process ensures that each query–response instance accurately reflects its corresponding modality cues. To guarantee faithful modality attribution, we apply an automatic filtering stage in which only QA items solvable exclusively under their assigned Required Modality Tags ((RMTs)) are retained. Multi-model voting is used for this filtering—questions that can be answered without the required modalities are discarded. Despite leveraging multiple high-performance models, automatic TAG labeling remains imperfect.

For both the MAPLE-bench subsets, human reviewers inspect and revise the test portion, with approximately 74.2% (for MAPLE-QA) of TAG assignments updated after manual validation. This highlights the current limitations of state-of-the-art models in accurately understanding and attributing modality-specific reasoning. These additional refinement steps close several observed “generation gaps,” addressing inconsistencies in modality correlation, annotation drift, and task solvability. They collectively ensure that MAPLE-bench provides rigorous, modality-grounded supervision for evaluating multimodal post-training.

B.2. MAPLE-QA+ task

To further assess whether MAPO genuinely acquires modality-aware reasoning rather than memorizing modality–query correspondences, we extend the MAPLE-QA dataset into an augmented version termed **MAPLE-QA+**. This extension deliberately perturbs the alignment between each question and its accessible modalities, thereby testing the model’s robustness under incomplete or redundant sensory conditions.

For every QA instance, we construct three distinct strata:

1. **Modality-Exact:** The available modalities precisely match the Required Modality Tags (RMTs).
2. **Modality-Superset:** Additional modalities beyond those specified by the RMTs are included.
3. **Modality-Deficit:** One or more RMT-specified modalities are intentionally omitted.

The resulting MAPLE-QA+ dataset comprises a total of **137,313** QA pairs, partitioned into **103,265** for training and **34,048** for testing, evenly distributed across the three modality strata. To discourage superficial pattern recognition and encourage genuine modality reasoning, we introduce a set of “*None-answer*” options that represent cases where the provided modalities are insufficient to yield a valid response.

Specifically, for all **modality-deficit** samples, the correct answer is replaced with one of the predefined *none-answer* statements, compelling the model to recognize when information is incomplete. In addition, for **25%** of the **modality-exact** and **modality-superset** examples, the original incorrect answer is substituted with a randomly selected *none-answer* to further regularize the model’s confidence calibration. The training set includes **100** distinct *none-answer* variants, each expressing the notion of insufficient multimodal evidence in slightly different linguistic forms.

To construct the MAPLE-QA+ dataset, we first expanded the human-verified MAPLE-QA test set of **4,864** samples across all modality combinations, resulting in a total of **34,048** test samples. For the training set, to balance training efficiency and the ratio between *None-answer* and valid answer choices, we did not expand all seven modality combinations. Instead, we included the all original MAPLE-QA training set of **47,893** samples as the **Modality-Exact** subset, and added **31,635** **Modality-Superset** and **23,737** **Modality-Deficit** samples. The 25% substitution ratio for *None-answer* replacement in the Modality-Exact and Modality-Superset strata is summarized in Table 6.

Table 6. Number of *None-answer* substitutions in **MAPLE-QA+** Train set by strata.

Strata	None-answer Count
Modality-Exact	10,969
Modality-Superset	7,316
Modality-Deficit	23,614

Below is the complete list of *none-answer* candidates used in MAPLE-QA+:

Tag-specific QA generation Prompt

- Cannot determine the answer from the given information.
- The provided data is insufficient to infer a valid response.
- Unable to answer with the current set of modalities.
- The available inputs do not contain enough evidence.
- Not enough multimodal context is provided.
- The input lacks sufficient cross-modal information.
- Missing modality prevents accurate reasoning.
- Unable to infer due to incomplete sensory input.
- The visual, audio, or speech data is inadequate.
- Cannot make a reliable decision with missing modalities.
- The input lacks necessary multimodal cues.
- Insufficient perceptual evidence to determine the answer.

770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824

- The question cannot be resolved with the current data.
- Incomplete sensory input prevents a valid conclusion.
- The available modalities do not support a confident response.
- Unable to conclude due to missing sensory information.
- The input data is too limited to determine an answer.
- Cannot infer without complete multimodal evidence.
- The response cannot be generated with partial modalities.
- Missing visual or speech information restricts reasoning.
- The given input lacks essential modality signals.
- Unable to decide due to insufficient multimodal evidence.
- The data provided is incomplete for this task.
- Cannot produce an answer without all required modalities.
- The current input does not provide enough context.
- Missing modality information leads to uncertainty.
- The answer cannot be derived from partial inputs.
- Insufficient multimodal alignment to determine output.
- The available data does not support a clear conclusion.
- Unable to respond accurately with missing modality data.
- The input lacks critical sensory components.
- Cannot determine the answer without full modality coverage.
- Incomplete multimodal input prevents reasoning.
- The provided information is too sparse to answer.
- Missing vision or audio data limits understanding.
- The question requires more complete sensory input.
- Unable to infer due to lack of multimodal coherence.
- The current modalities do not provide enough evidence.
- Cannot answer confidently with incomplete data.
- The input is missing key perceptual information.
- Insufficient modality coverage for reliable inference.
- The available inputs are inadequate for this question.
- Missing sensory data prevents accurate interpretation.
- Cannot determine outcome with partial modality input.
- The multimodal data is incomplete for reasoning.
- Unable to conclude due to missing perceptual signals.
- The input lacks sufficient cross-modal information.
- Not enough multimodal evidence to support an answer.
- The provided modalities are insufficient for inference.
- Cannot generate a valid response with missing data.
- Incomplete modality input leads to uncertainty.
- The question cannot be answered with current inputs.
- Missing modality context prevents understanding.
- The available sensory data is too limited.
- Unable to reason effectively with incomplete modalities.
- The input does not contain enough multimodal cues.
- Cannot infer due to lack of modality completeness.
- The data lacks necessary cross-modal features.
- Insufficient perceptual input to determine the answer.
- The current modalities do not provide enough clarity.
- Missing modality signals hinder reasoning.
- The input is incomplete for multimodal inference.
- Cannot produce a confident answer with missing inputs.
- The available data lacks multimodal consistency.
- Incomplete sensory evidence prevents decision-making.
- The question requires additional modality input.
- Unable to determine due to missing visual or audio data.
- The input lacks sufficient multimodal representation.
- Cannot answer without complete sensory information.
- The provided modalities are incomplete for reasoning.
- Insufficient multimodal context to generate an answer.
- The data is too limited to support a conclusion.
- Missing modality features prevent accurate inference.
- The input does not provide enough perceptual evidence.
- Unable to respond due to incomplete modality coverage.

- The available inputs are insufficient for reasoning.
- Cannot determine the answer with partial sensory data.
- The multimodal input is incomplete for this task.
- Missing modality information limits interpretability.
- The question cannot be resolved with current modalities.
- Insufficient sensory evidence to produce an answer.
- The input lacks necessary multimodal alignment.
- Unable to infer due to missing perceptual context.
- The data provided is not enough for reliable reasoning.
- Cannot generate an answer with incomplete modalities.
- The available sensory inputs are too limited.
- Missing modality data prevents accurate response.
- The input lacks full multimodal representation.
- Insufficient modality information to conclude.
- The question cannot be answered with partial data.
- Unable to reason due to incomplete sensory input.
- The provided modalities do not cover all required aspects.
- Cannot determine outcome with missing multimodal evidence.
- The input data is incomplete for this inference.
- Missing modality signals reduce confidence in response.
- The available information is insufficient for reasoning.
- Unable to produce an answer due to missing modalities.
- The multimodal input lacks completeness for inference.
- Cannot answer accurately with partial modality data.
- The input does not contain enough multimodal evidence.

By exposing the model to all modality configurations and explicitly teaching it to abstain when evidence is lacking, MAPLE-QA+ acts as a strong regularizer. It encourages the model to calibrate its confidence based on information sufficiency, thereby reducing hallucinations and improving factual grounding under realistic multimodal conditions.

B.3. Tag-specific QA generation Prompt

Tag-specific QA generation Prompt

Task:

Given a detailed description that summarizes the visual and audio(sound event sequence and speech) content of a video and a series of audio and visual events that occurs at the same time, generate question-answer pairs that based on the description to help human better understand the video.

Guidelines For Question-Answer Pairs Generation:

- The QAs you generate should be answered with and only with BOTH audio(human speech and object sound) and visual information.
- When generating choices, make sure the choices are equally long so that there won't be data bias.
- To increase the difficulty, generate answer choices with some ambiguity or confusion, ensuring that it requires careful attention to both visual and audio elements to answer correctly.
- Make the answer choices more deceptive so that multiple options appear plausible, rather than having only one or two clearly reasonable choices. This will increase the question's difficulty.
- For all types of questions, don't explicitly mention too much information in the question otherwise the question can be answered correctly without using information from the video.
- For every type of question, you need to write a QA for each Tag Types (Modality-specific QA generation). For example, you need to create a QA for each type of V/A/S/VS/VA/AS/VAS for the reasoning type, and for the rest of the question types, make a question for all combinations as well.
- Please tag each modality combination carefully. For example, if you can answer with just vision information without needing any other modality information, you should tag it with "V", if you can answer with sound information only, "A", and "S" if you answer with speech information only. In the same way, the rest of the combinations "VA", "VS", "AS", and "VAS" can be answered with just that modality. Please verify and tag whether you need that modality.

Question Types:

You should generate the following types of questions:

- AV Event Alignment: Formulate questions to determine which audio and visual events occurred simultaneously with each other.
- Event Sequence: Formulate questions to determine the temporal sequence of visual and audio events in the video.
- Reasoning: Formulate questions to explain the cause or reason behind the occurrence of a visual or audio event in the video.
- Inference: Formulate questions to speculate on information not explicitly presented in the video.
- Comparative: Formulate questions to compare the similarity or difference between the audio and visual information of two or more events in the video.
- Context understanding: Formulate questions to determine the contextual information surrounding a specific event in the video.

Tag Types (Modality-specific QA generation):

When generating each Question-Answer pair, assign one of the following tags to indicate which modalities are considered:

- 'V': Use only the visual information from the video. Do not use speech or other audio information.
- 'A': Use only the non-speech audio information from the video. Do not use speech or visual information.
- 'S': Use only the speech information from the video. Do not use other audio or visual information.
- 'VA': Use both visual and non-speech audio information. Do not use speech information.
- 'VS': Use both visual and speech information. Do not use non-speech audio information.
- 'AS': Use both non-speech audio and speech information. Do not use visual information.
- 'VAS': Use all available information from the video (visual, non-speech audio, and speech).

Output Format:

The questions should be in the form of multiple choice. The choices should look like: ["A. Choice 1", "B. Choice 2", "C. Choice 3", "D. Choice 4"]

The answer should be a single capital letter A, B, C, or D.

Your output should be formed in a JSON file.


Your response should look like:

```
```json
[{"Type": <type-1>, "Question": <question-1>, "Choice": <choice-1>, "Answer": <answer-1>, "Explanation":
<explanation-1>, "tag": <tag-1>}, {"Type": <type-2>, "Question": <question-2>, "Choice": <choice-2>,
"Answer": <answer-2>, "Explanation": <explanation-2>, "tag": <tag-2>}, ...]
```

**B.4. Benchmark Example:MAPLE-QA+**

MAPLE-QA+ VA Question Example Instance

Seed image (from video -spqgACI6UQ.mp4).



**Subtitle: 0-10s:** In this section, a man is speaking with an instructional tone. He says: “And as it’s coming up around here, then you push down with your feet to come back down again.” **10-20s:** In this section, there is the sound of a skateboard hitting the ground. Then, there are generic sound effects. Finally, there are more sound effects. **20-30s:** In this section, a male speaker is instructing someone in a casual tone. Then, the male speaker says, “Now we’re going to drop in. Okay. So, you leave your foot on the tail, leave it up.”

{Question}

{Choices}

The final option MUST BE put in [].

**Question:**

What sequence of non-speech audio events closely corresponds to the second skateboarder’s aerial trick and subsequent landing on his board?

**Modality-Exact & Expansion Choices VA / VAS:**

- A. A bang followed by a metallic clank, then rolling sounds continuing for several seconds.
- B. A thud, a rapid whoosh, a sharp thwack, a second whoosh, and a final heavy thud.**
- C. A prolonged scraping sound followed by rising impact noises and a clicking sound.
- D. Continuous grinding sounds with intermittent clunking and faint wind-like noises.

**Modality-Deficit Choices V / A / S / VS / AS:**

- A. A bang followed by a metallic clank, then rolling sounds continuing for several seconds.
- B. Cannot answer based on the provided information.**
- C. A prolonged scraping sound followed by rising impact noises and a clicking sound.
- D. Continuous grinding sounds with intermittent clunking and faint wind-like noises.

**B.5. MAPLE-Caption task**

We generate modality-tagged caption instances from seed videos sampled from the VAST-Omni corpora (Chen et al., 2023), using Gemini-2.5-Flash (Comanici et al., 2025) under a modality-isolated protocol that controls which modalities are exposed to the generator. We provide the exact prompts used to generate caption data for each modality tag. Our prompt consists of (i) a shared base instruction, (ii) a tag-specific question template, and (iii) a common ending instruction that enforces a unified, modality-agnostic sentence.

**Shared Base Prompt**

Base Prompt

Given the video, there are visual, audio, and subtitle information. The audio can be divided into sound and speech information, and the audio information should cover both sound and speech aspect.

**Tag-specific Question Templates** The following templates are used depending on the available modality tag: V (visual only), A (audio only), S (subtitle only), VA, VS, AS, and VAS.

Question Templates by Modality Tag

[V] Following a video without an audio, make a context understanding video description in one sentence. The answer format should be follow ' Video Description' format.

[A] Following an audio, make a context understanding video description in one sentence. The audio is a part of video. The description should include both speech and background sound. The answer format should be follow ' Video Description' format.

[S] For given subtitles, make a context understanding video description in one sentence. Only refer to subtitle. The answer format should be follow ' Video Description' format. subtitle: {subtitle}

[VA] Following a video, make an unified context understanding video description in one sentence. The description should refer to both visual and audio evidences (speech and background sound). The answer format should be follow ' Video Description' format.

[VS] Following a video, make an unified context understanding video description in one sentence, referring to both visual evidences and subtitle. Do not guess anything about sound. The answer format should be follow ' Video Description' format. subtitle: {subtitle}

[AS] Following an audio, make an unified context understanding video description in one sentence. The audio is a part of video. The description should refer to both audio evidences and subtitle. The answer format should be follow ' Video Description' format. subtitle: {subtitle}

[VAS] Following a video, make an unified context understanding video description in one sentence, referring to all of the visual, audio evidences (speech and background sound), and subtitle. The answer format should be follow ' Video Description' format. subtitle: {subtitle}

**Common Ending Instruction**

End Prompt


Please combine the given information to create a natural sentence without mentioning modality, ensuring that the sentence should be well-aligned and unified given modalities information.

**Final Prompt Composition.** For each instance, we concatenate Base Prompt + (tag-specific question template) + End Prompt. For tags that include subtitles (S, VS, AS, VAS), we replace {subtitle} with the corresponding subtitle text.

**B.6. Benchmark Example:MAPLE-Caption**

MAPLE-Caption Example Instance

Seed image (from video `A_4TRaRiLRw.10.mp4`).



Subtitle: So yes, you like this yeah. Alright, you like it.

**Modality-tagged ground-truth captions**

- V:** A woman and young boy with a colorful clapper toy, before she opens a green bag to reveal an item inside.
- A:** A female adult asking a child if he like something, as the child enthusiastically confirms he likes it amidst the rustling sounds of objects being handled, with low-volume background music.
- S:** Someone confirms with another person that they like something.
- VA:** A woman and a young boy are seen at a table, discussing a colorful clapper toy as the boy enthusiastically confirms he likes it, before the woman pulls an item from a green bag, accompanied by low-volume background music.
- VS:** A woman shows a colorful hand clapper toy to a young boy, confirming his liking for it as she puts the toy away and then retrieves a paper from a green bag.
- AS:** An adult asks a child if he like something and the child gives an enthusiastic “Yep, yep” in response, amidst the rustling sounds of objects being handled, with low-volume background music.
- VAS:** A woman holding a clapper toy that makes clapping sounds asks a young boy sitting next to her if he likes it, and after he enthusiastically agrees, she reaches into a green bag to pull out another item, accompanied by low-volume background music.

**C. Detailed Evaluation Protocol**

**C.1. Modality Gap Computation**

For a model  $M$  evaluated across all RMTs, we partition queries by complexity:

$$\begin{aligned}
 \mathcal{U} &= \{V, A, S\} & \bar{A}_{\text{Uni}} &= \frac{1}{|\mathcal{U}|} \sum_{r \in \mathcal{U}} \text{Acc}(M, r) \\
 \mathcal{B} &= \{VA, VS, AS\} & \bar{A}_{\text{Bi}} &= \frac{1}{|\mathcal{B}|} \sum_{r \in \mathcal{B}} \text{Acc}(M, r) \\
 \mathcal{T} &= \{VAS\} & \bar{A}_{\text{Tri}} &= \text{Acc}(M, VAS)
 \end{aligned}$$

The Modality Gap (Xiang et al., 2025) captures average relative improvement:

$$\Delta_{\text{gap}} = \frac{1}{3} \left( \left| \frac{\bar{A}_{\text{Bi}} - \bar{A}_{\text{Uni}}}{\bar{A}_{\text{Uni}}} \right| + \left| \frac{\bar{A}_{\text{Tri}} - \bar{A}_{\text{Uni}}}{\bar{A}_{\text{Uni}}} \right| + \left| \frac{\bar{A}_{\text{Tri}} - \bar{A}_{\text{Bi}}}{\bar{A}_{\text{Bi}}} \right| \right)$$

**C.2. LLM-as-Judge Scoring (J)**

We use an LLM-as-Judge with modality-conditioned *textual references* to detect hallucination errors and modality-fusion in the captioning task:

$$J \in [0, 1] \quad (\text{per RMT configuration}).$$

The Judge assigns three axis scores per sample: *modality missing* ( $s_{\text{miss}}$ ), *hallucination* ( $s_{\text{hall}}$ ), and *fusion quality* ( $s_{\text{fus}}$ ), each taking a value in  $[0, 1]$ . The final captioning reward is computed as their average:

$$R_{\text{cap}} = \frac{1}{3} (s_{\text{miss}} + s_{\text{hall}} + s_{\text{fus}}) \in [0, 1].$$

For precise multimodal assessment, we adopt a *ref-combine* protocol. Specifically, the Judge is provided with a reference bundle  $\mathcal{R}_{\text{ref}} = \{\text{GT}_{\text{tag}}, \text{Ref}_V, \text{Ref}_A, \text{Ref}_S\}$ , where all elements are given as *textual descriptions*: the ground-truth caption corresponding to the current RMT, together with uni-modal reference captions for visual, audio, and subtitle streams. The Judge evaluates the generated caption  $\hat{y}$  against these references using an explicit scoring rubric. We used GPT-4o (Hurst et al., 2024) as the Judge model.

**Judge Prompt Template.** We use the following text-only prompt template for GPT-4o as the captioning Judge. All references (uni-modal captions) are provided in *text* form, and only those references whose modalities are contained in the current modality tag are included in the reference bundle.

Judge Prompt Template

You are given:

Ground Truth (GT) caption describing the video accurately.

Predicted caption generated by fusing information from one or more modalities (vision, audio, subtitle).

Separate captions for each modality (vision-only caption, audio-only caption, subtitle-only caption).

\*\* If there is only one modality, the GT caption and the separate captions for each modality are the same, so only one is provided.

Your task: Evaluate the predicted video caption by comparing it directly to the GT caption, while also using the modality captions to strengthen detection of hallucinations and assess fusion quality.

Scoring criteria:

[1. Missing Information] (Compare Predicted vs GT)

1.0 → No missing information at all.

0.8 → Minor omission, meaning unaffected.

0.6 → Moderate omission, some important details missing.

0.4 → Significant omission, core meaning partially lost.

0.2 → Severe omission, most important details missing.

0.0 → Almost all information missing.

[2. Errors / Hallucinations] (Check Predicted against GT and modality captions)

1.0 → No errors or hallucinations.

0.8 → Minor wording inaccuracy, meaning preserved.

0.6 → Moderate factual error, possible meaning distortion.

0.4 → Significant error, meaning largely affected.

0.2 → Severe error, multiple incorrect details.

0.0 → Completely incorrect description.

[3. Modality Fusion Accuracy] (Check if Predicted integrates all modality info into one coherent sentence)

1.0 → Perfect integration of all modalities into one coherent sentence.

0.8 → Mostly coherent fusion, very small detail underrepresented.

0.6 → Fusion partially successful; some modality info feels separate.

0.4 → Significant fusion issues; modalities mentioned in isolation.

1155 0.2 → Severe fusion failure; modalities not meaningfully combined.  
 1156 0.0 → No fusion at all.

1157  
 1158 Output format (MUST follow exactly, and DO NOT provide any explanation or reasoning):  
 1159

1160 Missing Info score:  $\alpha score \alpha$   
 1161 Error score:  $\beta score \beta$   
 1162 Fusion score:  $\gamma score \gamma$   
 1163

1164 Given Information as follows:  
 1165 Generated caption: {pred}  
 1166 GT caption: {gt<sub>fused</sub>}  
 1167 Vision caption: {Ref<sub>V</sub>}  
 1168 Audio caption: {Ref<sub>A</sub>}  
 1169 Subtitle caption: {Ref<sub>S</sub>}  
 1170

### 1171 C.3. Fusion Gain

1172 Out of total  $N$  samples, count the fraction of samples where multi-modal captions outperform best uni-modal caption.

$$1173 \text{ uni-modal} = \max\{J(V), J(A), J(S)\}$$

$$1174 \text{ multi-modal} = \max\{J(VA), J(VS), J(AS), J(VAS)\}$$

$$1175 \text{ Fusion Gain} = \frac{|\{\text{samples: multi-modal} > \text{uni-modal}\}|}{N}$$

## 1176 D. RL formalization: Distributions and Training Schemes for MLM

### 1177 D.1. General RL Baseline

1178 In this section, we explain the reinforcement learning from trajectories where prompts  $x \sim \mathcal{D}$  distribution trigger completions  $y \sim \pi_\theta(\cdot | x)$ . The distributed setup splits GPUs between *generators* (high-throughput rollouts via optimized kernels) and *trainers* (FSDP parameter updates), with  $\pi_{\text{gen}}^\theta$  and  $\pi_{\text{train}}^\theta$  denoting the policy on each backend.

1179 Our basic training pipeline adopt GRPO (Shao et al., 2024) without KL regularization, augmented with symmetric clipping (Yu et al., 2025) and token-level loss-aggregation because it provides a finer, more detailed optimization signal for reasoning, preventing issues like length bias and enabling better learning from complex signals (Sheng et al., 2025). For prompt  $x$ , the old policy  $\pi_{\text{gen}}^{\theta_{\text{old}}}$  generates  $G$  completions  $\{y_i\}_{i=1}^G$ , each rewarded  $r_i = r(x, y_i)$ . Group-normalized advantages follow:

$$1180 \hat{A}_i = r_i - \frac{1}{G} \sum_{j=1}^G r_j, \quad \hat{A}_i^G = \frac{\hat{A}_i}{\text{std}(\{r_j\}_{j=1}^G)}$$

1181 Token-level importance sampling ratios use symmetric clipping ( $\epsilon^- = \epsilon^+ = \epsilon$ ):

$$1182 \rho_{i,t}(\theta) = \frac{\pi_{\text{train}}^\theta(y_{i,t} | x, y_{i,<t})}{\pi_{\text{gen}}^{\theta_{\text{old}}}(y_{i,t} | x, y_{i,<t})}, \quad \text{clip}_{\text{sym}}(\rho, \epsilon^-, \epsilon^+) = \text{clip}(\rho, 1 - \epsilon^-, 1 + \epsilon^+)$$

1183 The token-level surrogate objective becomes:

$$1184 \mathcal{L}_\theta^{\text{GRPO}, \mathcal{D}} = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}^{\text{gen}}(\cdot | x)} \left[ \frac{1}{\sum_{i=1}^G |y_i|} \sum_{i=1}^G \sum_{t=1}^{|y_i|} \min(\rho_{i,t}(\theta) \hat{A}_i^G, \text{clip}(\rho_{i,t}(\theta), \epsilon) \hat{A}_i^G) \right] \quad (5)$$

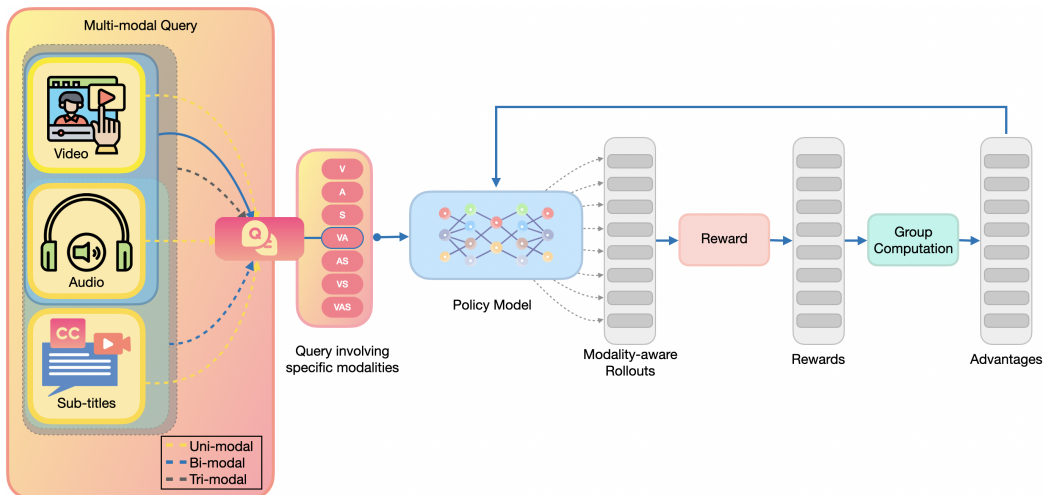


Figure 7. **Modality-Aware Policy Optimization using modality-stratified GRPO.** Algorithm flow showing RMT-driven stratification. (1) Sample  $x$  and extract required modalities  $M_x$ ; (2) Generate  $G$  rollouts  $\{y_i\}_{i=1}^G$  and compute rewards  $r_i = r(x, y_i, M_x)$  and finally group advantages  $\hat{A}_i^G$ ; (3) Form per-RMT batches  $\mathcal{B}_M$ ; (4) Modality-unaware baseline  $\hat{g}_{\text{MU}}$  mixes all  $R_M$  (high between-group variance); (5) MAPO base  $\hat{g}_{\text{MA}}$  normalizes within each  $\mathcal{B}_M$  using  $\hat{A}_j^{(M)}$  (variance  $\text{Var}(\hat{g}_{\text{MA}}) \leq \text{Var}(\hat{g}_{\text{MU}})$ ). Stratification removes cross-modal reward scale differences, enabling stable multimodal post-training.

## D.2. Modality based post-training

Consider a finite modality set  $\mathcal{M} = \{V, A, S\}$  for video, audio, and text/subtitles. Each query  $x_i \in \mathcal{D}$  is tagged with required-modality subset  $M_i \subseteq \mathcal{M}$ —the minimal signals needed to solve it reliably (e.g.,  $M_i = \{V, A\}$  for audio-visual reasoning,  $M_i = \{S\}$  for subtitles-only). The full MAPO algorithm is mentioned in Alg. 1 and pictorially described for modality based required tags in Fig. 7. The adaptive KL and dynamic curriculum based changes to MAPO are further described in Alg. 2.

---

### Algorithm 1 Modality-Aware Policy Optimization (MAPO) $\approx$ Modality-Stratified GRPO

---

**Require:** Dataset  $\mathcal{D}$  with RMTs  $M_x \subseteq \{V, A, S\}$ , group size  $G$ , clip  $\epsilon$ , learning rate  $\alpha$

**Ensure:** Policy  $\pi_\theta$

- 1: **while** not converged **do**
  - 2:   Sample  $x \sim \mathcal{D}$ ; extract modalities  $M_x$
  - 3:   Generate  $\{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}^{\text{gen}}(\cdot|x, M_x)$ ; score  $r_i = r(x, y_i, M_x)$
  - 4:   Compute group advantages:  $\hat{A}_i^G = \frac{r_i - \mu_G}{\sigma_G}$
  - 5:   **// Stratify by RMT**
  - 6:   Form  $\mathcal{B}_M = \{(x_j, y_j, r_j) : M_{x_j} = M\}$  for each  $M$
  - 7:   **Modality-Unaware (baseline):**
  - 8:    $\hat{g}_{\text{MU}} \leftarrow \frac{1}{B} \sum_{j=1}^B \nabla_{\theta} \log \pi_{\theta}(y_j|x_j) \hat{A}_j$
  - 9:   **Modality-Aware (MAPO base):**
  - 10:    $\hat{g}_{\text{MA}} \leftarrow \sum_M \frac{1}{|\mathcal{B}_M|} \sum_{(x_j, y_j, r_j) \in \mathcal{B}_M} \nabla_{\theta} \log \pi_{\theta}(y_j|x_j, M_x) \hat{A}_j^{(M)}$
  - 11:    $\theta \leftarrow \theta - \alpha * \hat{g}_{\text{MA}}$   $\triangleright$  Variance  $\text{Var}(\hat{g}_{\text{MA}}) \leq \text{Var}(\hat{g}_{\text{MU}})$
  - 12: **end while**
- 

## E. Implementation details and observations

All inference (base model and checkpoints) uses consistent per-task settings: MAPLE-QA employs temperature 0.6, nucleus sampling ( $p = 0.95$ ); MAPLE-Caption uses temperature 0.7,  $p = 0.90$ .

---

**Algorithm 2** MAPO with Adaptive KL Weighting and/or Curriculum

---

**Require:** RMTs  $\mathcal{T}$ , Beta(100, 1),  $L_W = 5$

**Ensure:**  $\pi_\theta$ , KL histories  $\{\mathcal{H}_M\}_{M \in \mathcal{T}}$

```

1265 1: for all $M \in \mathcal{T}$ do
1266 2: $\mathcal{H}_M \leftarrow []$
1267 3: end for
1268 4: // Epoch 1: Sort \mathcal{T} by zero-shot accuracy (ascending)
1269 5: for $e = 1$ to E do
1270 6: if using KL-based-curriculum then
1271 7: $s_M = \frac{1}{|W_M|} \sum_{k \in W_M} D_{\text{KL},M}^{(k)}$ $\triangleright |W_M| = \min(L_W, |\mathcal{H}_M|)$
1272 8: Reorder \mathcal{T} by s_M (descending)
1273 9: end if
1274 10: for $M \in \mathcal{T}$ do
1275 11: Sample $\mathcal{B}_M = \{(x_j, y_j, r_j) : M_{x_j} = M\}$
1276 12: Compute $\hat{A}_j^{(M)}$ within \mathcal{B}_M
1277 13: $p_{\text{emp}} \leftarrow \text{histogram}(\{r_j\}_{j \in \mathcal{B}_M})$
1278 14: $D_{\text{KL},M} \leftarrow D_{\text{KL}}(p_{\text{emp}} \| \text{Beta}(100, 1))$
1279 15: $\mathcal{H}_M.\text{append}(D_{\text{KL},M})$
1280 16: $z_M = \frac{D_{\text{KL},M} - \mu_{\mathcal{H}_M}}{\sigma_{\mathcal{H}_M} + \epsilon}$
1281 17: $w_M = \text{sigmoid}(z_M)$
1282 18: $\mathcal{L}_\theta^{\text{MAPO}_w} = \frac{w_M}{|\mathcal{B}_M|} \mathcal{L}_\theta^{\text{GRPO}, \mathcal{B}_M}$
1283 19: $\theta \leftarrow \text{Optimizer step on } \mathcal{L}_\theta^{\text{MAPO}_w}$
1284 20: end for
1285 21: end for

```

---

### E.1. Optimal Training Strategy - Training characteristic curves

This appendix presents training curves for algorithmic variants from Sec. 5.2 on MAPLE-QA (2 epochs, 380 steps). Training stabilizes beyond 2 epochs, justifying this duration. Fig. 8 reveals MAPO achieving smaller, lower-variance policy gradient norms vs. MUPO with similar training configurations. Figure 9b shows MAPO achieving lower, more stable entropy and reduced step-time vs. MUPO by processing only required modality signals. Fig. 10 shows sample-level aggregation achieving superior entropy reduction despite modestly higher step-time vs. token-level aggregation. Fig. 11 shows asymmetric clipping reducing clip fractions 71.65% vs. symmetric, enabling more stable policy gradient loss through increased exploration. Fig. 12 shows early/mid-training filtering maintaining equivalent entropy and rollout scores vs. no filtering, but substantially reducing training time per step. Fig. 13 shows curriculum learning (uni→bi→tri-modal) achieving higher rollout scores and more stable policy gradient loss vs. mixed training. Fig. 14 shows Full-recipe outperforming MUPO across training time, eval scores, and entropy—delivering superior performance, robustness, and efficiency.

Additional observations: Fig. 15 shows prompt-level aggregation failing to converge entropy in MAPO even with more epochs of training.




Figure 8. MAPO achieving smaller, lower-variance policy gradient norms vs. MUPO with similar training configurations.




Figure 9. Entropy and step-time comparison for MUPO vs. MAPO under identical training settings. Left: MAPO exhibits lower and more stable policy entropy than the modality-unaware MUPO baseline, indicating better-regularized optimization. Right: MAPO achieves reduced wall-clock time per training step by processing only the required modality signals, improving efficiency without changing the overall configuration.

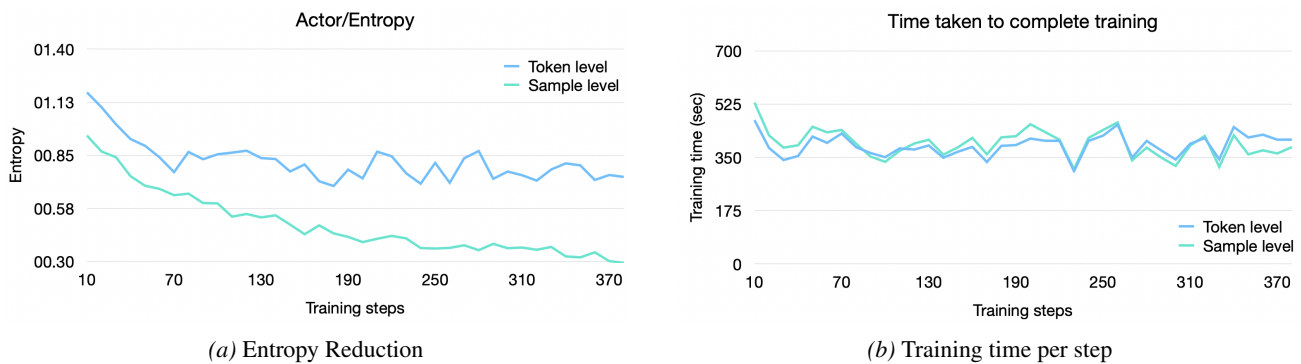


Figure 10. Entropy and step-time comparison across loss aggregation strategies for MAPO. Left: Sample-level aggregation achieves substantially lower and more stable policy entropy than token-level aggregation. Right: Sample-level incurs slightly higher wall-clock time per step but yields superior optimization stability.




Figure 11. **Clipping behavior and policy loss stability: symmetric vs. asymmetric clipping in MAPO.** Left: Asymmetric clipping ( $\epsilon^+ = 0.3, \epsilon^- = 0.2$ ) reduces clip fractions by 71.65% vs. symmetric clipping ( $\epsilon = 0.2$ ). Right: Lower clipping enables more exploration, yielding more stable final policy gradient loss.




Figure 12. **Entropy and rollout scores across sampling strategies in MAPO.** Left: All strategies (no sampling, early filtering, mid-training filtering) achieve equivalent entropy reduction and rollout scores (solid: effective scores; dotted: including filtered samples). Right: Early filtering substantially reduces training time per step while maintaining optimization quality.




Figure 13. **Training dynamics across curriculum strategies in MAPO.** Left: Prompt length progression (uni→bi→tri-modal) for curriculum vs. mixed training. Middle/Right: Curriculum yields higher rollout scores and more stable policy gradient loss throughout training.

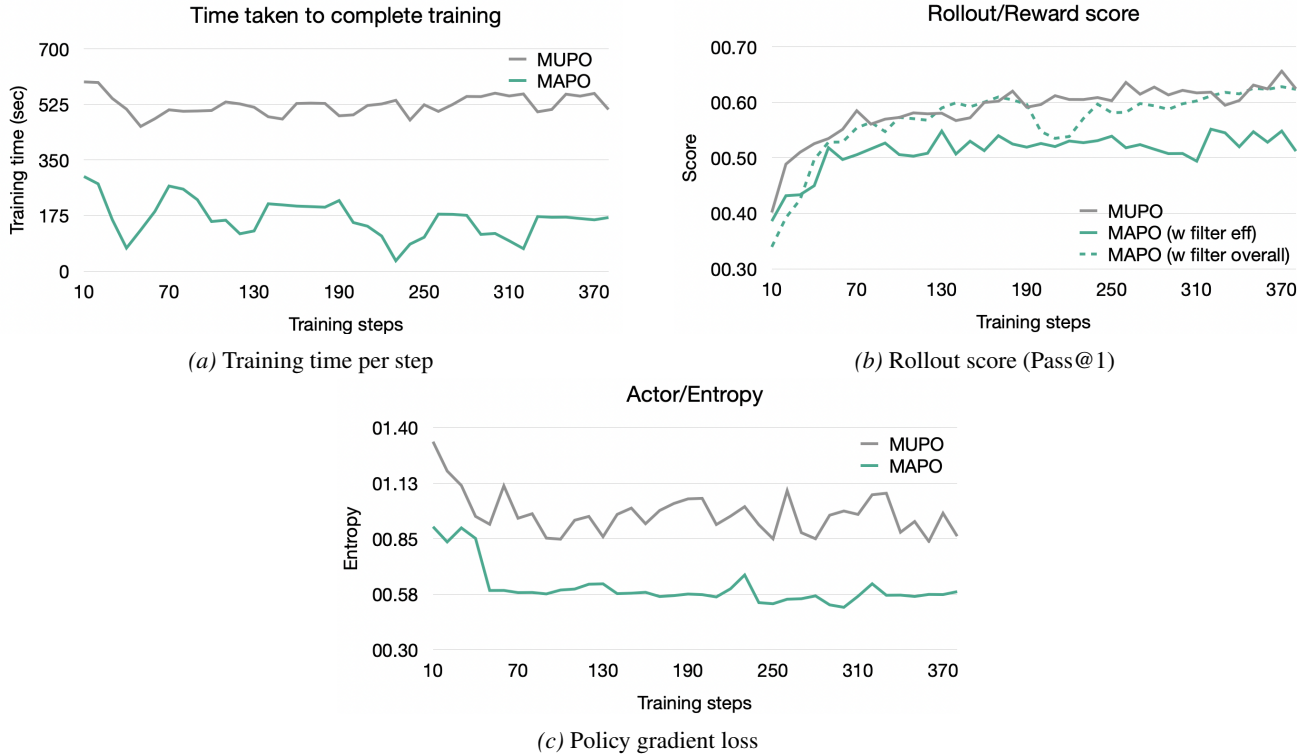


Figure 14. **MUPO vs. final Full-recipe: comprehensive training dynamics.** a): Full-recipe reduces training time per step. b): Equivalent rollout scores throughout training. c): Lower, more stable entropy. Full-recipe improves output performance, robustness, and cost-efficiency across all metrics.




Figure 15. **Entropy curve for loss aggregation at prompt level.** We noticed no convergence of entropy by this mode.

1485 **E.2. Adaptive Training Strategy**

1486 For adaptive training strategies, we train all variants for 5 epochs on both datasets to ensure full convergence and accurately  
 1487 measure performance gains from each component. In modality-aware post-training, data are partitioned into modality tags  
 1488  $M \in \mathcal{M}$  (e.g., V, A, S, VA, VS, AS, VAS), and GRPO-style updates scale with observed advantages. This creates tag-wise  
 1489 imbalance where harder tags yield smaller gradient signals and can be overshadowed by easier tags. Accordingly, our  
 1490 implementation consists of: (i) a KL-based difficulty metric, (ii) difficulty-adaptive reweighting (controlling *how much* to  
 1491 update per tag), and (iii) KL-driven curriculum scheduling (controlling *when* to update each tag).

1493 E.2.1. KL-BASED DIFFICULTY METRIC

1495 **Target distribution.** We quantify tag difficulty using the KL divergence between an empirical reward distribution  $p_{\text{emp}}$   
 1496 (measured on a tag-batch) and a fixed near-optimal target  $p_{\text{tgt}}$ . For captioning, we use  $p_{\text{tgt}} \sim \text{Beta}(100, 1)$ , whose mode is  
 1497 0.99.

1498 **Captioning (continuous rewards).** Let  $R \in [0, 1]$  denote per-sample captioning rewards. We approximate  $p_{\text{emp}}$  on  $[0, 1]$   
 1499 via histogram discretization: using  $K$  uniform bins (width  $\Delta x$ ), bin centers  $\{x_i\}_{i=1}^K$ , and a normalized histogram with  
 1500 `density=True` to obtain an empirical *density*  $p(x_i)$ . With  $q(x_i) = \text{BetaPDF}(x_i; 100, 1)$ , we compute

$$1502 D_{\text{KL},M}^{\text{Cap}} \approx \sum_{i=1}^K p(x_i) \log \frac{p(x_i) + \varepsilon}{q(x_i) + \varepsilon} \Delta x, \tag{6}$$

1505 where we use  $K = 50$  and a small  $\varepsilon$  for numerical stability.

1507 **QA (discrete rewards).** For QA, rewards are binary  $R \in \{0, 1\}$ . To match the captioning target, we set the Bernoulli  
 1508 target success rate to the mean of  $\text{Beta}(100, 1)$ :

$$1509 p_{\text{tgt}} = \frac{100}{101} \approx 0.9901. \tag{7}$$

1512 For a tag-batch of size  $B_M$ , the empirical accuracy is

$$1514 p_{\text{emp}} = \frac{1}{B_M} \sum_{j=1}^{B_M} 1\{R_j = 1\}.$$

1517 and we compute the Bernoulli–Bernoulli KL:

$$1519 D_{\text{KL},M}^{\text{QA}} = p_{\text{emp}} \log \frac{p_{\text{emp}}}{p_{\text{tgt}}} + (1 - p_{\text{emp}}) \log \frac{1 - p_{\text{emp}}}{1 - p_{\text{tgt}}}, \tag{8}$$

1521 with probabilities clipped to  $[\varepsilon, 1 - \varepsilon]$  if needed.

1523 **Interpretation and usage.** Hard tags typically exhibit  $p_{\text{emp}} \ll p_{\text{tgt}}$  (large KL), while easy tags approach the target (KL  
 1524 near zero). We use KL as a unified difficulty signal for both reweighting and curriculum scheduling.

1526 E.2.2. DIFFICULTY-ADAPTIVE REWEIGHTING: PRACTICAL DETAILS

1527 **What is adaptive about the weights.** The difficulty-adaptive reweighting (defined in the main paper) is intended to correct  
 1528 *relative* imbalance across tags, rather than to rescale absolute reward magnitudes. We therefore base the weight update on  
 1529 each tag’s KL statistics and normalize them using a short KL history to make the weights robust to reward drift over training.

1531 **History window size  $L_{\mathcal{H}}$ .** We set the history length  $L_{\mathcal{H}}$  to be approximately the number of steps in one epoch (i.e.,  
 1532 steps-per-epoch). This choice ensures that the history aggregates KL values across (almost) all modality tags within a  
 1533 training cycle, which better approximates the global reward distribution and stabilizes *relative* difficulty estimation. In  
 1534 contrast, overly short windows yield noisy, batch-specific weights, while overly long windows make the weights slow to  
 1535 react to changes in learning dynamics.

1536 **Normalization and squashing.** We compute a z-score from the running mean and standard deviation over the KL history,  
 1537 and squash it with a sigmoid to map to  $(0, 1)$ . In practice, most z-scores fall within a moderate range (e.g.,  $[-3, 3]$ ), which  
 1538 avoids extreme scaling and reduces oscillatory updates caused by noisy per-batch rewards.

**Why this is complementary to curriculum.** Reweighting controls *how much* each tag updates, but when difficulty gaps are extreme, hard-tag gradients can still be overshadowed by easy-tag batches if trained in an unfavorable order. We therefore additionally control *when* each tag is optimized via curriculum scheduling.

### E.2.3. KL-DRIVEN CURRICULUM: PRACTICAL DETAILS

**Window size ( $L_W$ ).** We compute curriculum priorities from a short recent window of KL values and set  $L_W = 5$ . This value was chosen as a practical heuristic rather than through explicit statistical tuning: we use the *last five* tag-batches because their performance signals are typically the most *mature* (i.e., after the initial transient), providing a stable yet responsive estimate of current tag difficulty.

**First-epoch initialization.** In the first epoch, KL histories are unavailable. We therefore initialize the curriculum order using zero-shot performance, assigning higher priority to tags with lower zero-shot accuracy. After the first epoch, once sufficient KL history has accumulated, we switch to KL-based scheduling.

**Why ordering matters.** Curriculum scheduling complements reweighting by controlling *when* each tag is optimized. By front-loading underperforming (high-KL) tags, the model learns from their weaker gradient signals before the optimization trajectory is dominated by easier tags with stronger gradients.

### E.2.4. REWARD EVALUATION NOTE (POINTER TO D.2)

**Ref-combine protocol.** We adopt a ref-combine evaluation setup for captioning rewards to minimize ambiguity in assessing *hallucination* and *modality-fusion* quality in multimodal captions. The judge is provided with the tag-specific ground-truth caption along with uni-modal reference captions (denoted  $\{GT_M, Ref_V, Ref_A, Ref_S\}$ ), and scores the generated fusion caption along three axes (missing, hallucination, fusion), producing a continuous reward in  $[0, 1]$ . See Appendix C.2 for the full judge prompt and rubric.

**Modality-unaware Policy Optimization for Captioning.** While QA tasks naturally accommodate full-modality inputs (V, A, S) because answers are derived from content rather than fixed reference captions, captioning tasks differ fundamentally: each video–modality pair has a *fixed ground-truth caption* determined by its specific tag combination. This creates a structural mismatch for MUPO training, as modality dropout or random subsampling cannot be applied without invalidating the supervision signal. Consequently, similar to MAPO, we restrict caption-level MUPO training to the *VAS* condition on the same underlying videos, ensuring consistency between the provided inputs and the available ground-truth captions. This necessity leads to a significantly reduced training scale (*all-combi*: 5140 captions; *VAS-only*: 730 captions), but preserves the correctness of supervision while enabling modality-agnostic optimization in the captioning setting.

## E.3. Contrastive Reward Weighting

We apply a modality-aware reward weighting mechanism that adjusts the reward of each positive rollout according to its semantic similarity to a negative response. The final reward used for optimization is defined as:

$$R_{\text{weighted}} = \text{clip}_{[0,1]}(R_{\text{pos}} \cdot [1 + \alpha w(s; \tau)]), \tag{9}$$

where  $R_{\text{pos}}$  is the base positive reward,  $\alpha$  is a scaling factor,  $s$  is the similarity score between the positive and negative responses, and  $\tau$  is a similarity threshold.

**Positive and Negative Responses.** For each sample  $x$ , we denote by  $y^+$  a positive rollout generated under the complete modality set  $M$ , and by  $y^-$  a negative response generated from the same sample under a deficit condition, where one or more modalities are removed or degraded. Thus,  $M$  represents the full available modalities, while “deficit” indicates missing-modality or degraded-modality inputs.

**Similarity Function.** The similarity score  $s$  is computed as:

$$s = \text{Sim}(y^+ | x, y^- | x), \tag{10}$$

where  $\text{Sim}(\cdot)$  measures semantic closeness between the embeddings of  $y^+$  and  $y^-$ . In our experiments, we instantiate  $\text{Sim}(\cdot)$  using BERTSCORE, which produces a similarity value in the range  $[0, 1]$ .

**Gating Function.** We use discrete gating  $w(s; \tau) = 1\{s < \tau\}$ , where  $s$  is semantic distance and  $\tau$  the threshold.

**Interpretation.** The formulation in Eq. 9 increases the effective reward when the positive rollout remains semantically distant from deficit-induced responses, while the clipping operator  $\text{clip}_{[0,1]}(\cdot)$  ensures numerical stability. This encourages the model to produce modality-consistent and semantically distinct outputs during training. In the single-modality setting, CRW has no effect because no modality variation exists to form deficit-induced negative responses.

CRW - Qualitative Example

Seed image (from video A4H8CsKWDk.14.mp4).



Subtitle: I am going to be sensibly splitting the money between buying a daily and buying a supercar.

MAPO w/o CRW :

- A:** Someone discusses their plans to sensibly split money between buying a daily vehicle and a supercar while driving in a car.
- AS:** Someone describes their plan to sensibly allocate money between purchasing a daily vehicle and a supercar.
- S:** Someone explains their plan to responsibly divide their budget between buying a daily vehicle and a supercar.
- V:** The video shows a man sitting in the driver’s seat of a car, wearing a white T-shirt, a green cap, and a seatbelt. He gestures with his hands as he speaks, with the car’s interior and outdoor scenery visible through the windows.
- VA:** A person in a moving car discusses their plan to sensibly split spending between a daily driver and a supercar.
- VAS:** A person inside a car explains their plan to responsibly allocate money between purchasing a daily driver and a supercar.
- VS:** The video shows a man driving a car, wearing a white shirt and a cap, gesturing as he talks about how he plans to sensibly divide his budget between buying a daily car and a supercar.

MAPO w/ CRW:

- A:** The audio describes someone discussing their plan to sensibly split money between buying a daily car and a supercar while in a vehicle.
- AS:** The audio describes someone planning to sensibly split their money between purchasing a daily car and a supercar, while the subtitle reiterates this decision regarding financial allocation for these two types of vehicles.
- S:** The context is about someone sensibly planning to split their money between purchasing a daily use item and a supercar.
- V:** A man in a car, wearing a cap and white shirt, is gesturing with his hands while driving, with the car’s interior and some outdoor scenery visible through the windows.
- VA:** The video shows a man inside a car, wearing a cap and white shirt, gesturing with his hands, while speaking about sensibly splitting money between buying a daily driver and a supercar.
- VAS:** The video shows a man inside a car, wearing a cap and white shirt, talking with hand gestures about sensibly splitting money between buying a daily car and a supercar.
- VS:** The video shows a man in a car gesturing with his hands while talking about sensibly splitting money between buying a daily car and a supercar.

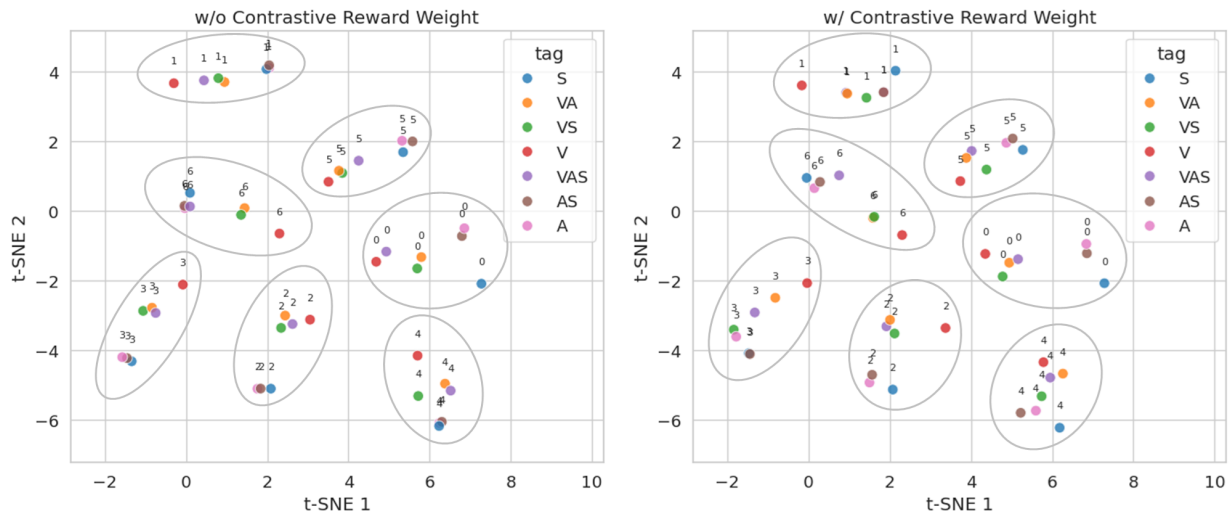


Figure 16. : Joint t-SNE visualization of response representations before and after applying CRW. Points that were densely overlapped within a sample group are more widely dispersed under CRW, indicating reduced representation collapse. The number annotated above each point denotes the corresponding video ID.