RELIABILITY-GATED SOURCE ANCHORING FOR CONTINUAL TEST-TIME ADAPTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Continual test-time adaptation (CTTA) updates a pretrained model on an unlabeled, changing stream while using source information to limit drift. The challenge is to preserve useful source knowledge without letting an uncertain reference restrict adaptation. We propose RMEMSAFE, a reliability-gated extension of ROID that uses the frozen source's normalized predictive entropy to modulate a parameter anchor and source-expert agreement weighting. As the source posterior approaches uniformity, these explicit source-dependent loss contributions vanish while the base adaptation losses remain active. Matched controls separate this mechanism from the auxiliary stabilizers in the complete method. On CCC-Hard with ResNet-50, runtime gating reduces error by 0.30 percentage points relative to the same objective with the gate fixed to one, and by 1.24 points relative to removing the explicit source terms. On two long-horizon stream orderings, it reduces error by 0.58 and 0.55 points, with improvement in all 20 evaluated seeds. Auxiliary stabilizers account for most of the average CCC improvement, and source-term removal remains stronger on the CCC mean. Source-degradation, small-batch, and class-permutation experiments further characterize the method and its limits. Together, these results motivate adjusting source influence to the current stream rather than assuming that retaining or removing it is always preferable.

1 INTRODUCTION

Current continual test-time adaptation (CTTA) paradigms harbor a fundamental vulnerability: they rely on static regularization to a source model that rapidly degrades under severe distribution shifts (Wang et al., 2022; Niu et al., 2022; Marsden et al., 2024; Lim et al., 2026). While anchoring limits catastrophic drift, it implicitly assumes the source remains a globally reliable reference. In reality, a source trained on one distribution can become entirely obsolete on another on CCC-Hard, for instance, a frozen ResNet-50 source collapses to a mere 1.33% top-1 accuracy. This exposes a critical flaw in modern CTTA: enforcing rigid agreement with a collapsed source actively suffocates the adapting expert.

This dynamic creates a dangerous false dichotomy in adaptation strategy. Retaining source influence via *blind anchoring* coupling the expert to the source without accounting for the source's current predictive uncertainty artificially constrains necessary learning. Conversely, entirely discarding the source abandons critical regularization precisely when the model needs stability the most. Figure 1 illustrates this paradox: completely removing explicit source loss terms improves adaptation on milder shifts, yet triggers catastrophic failure on severe corruptions (CCC-Hard) compared to a persistently anchored baseline. A robust CTTA framework must therefore dynamically distinguish between the source's outdated predictions and the underlying value of its structural guidance.

We introduce RMEMSAFE, a framework that fundamentally resolves this vulnerability by explicitly decoupling target stream adaptation from source reliance. RMEMSAFE operates on a dynamic, reliability-gated paradigm modulated by a single, elegant scalar: $\mathcal{R}_{\text{src}} = \max(0, 1 - \mathcal{H}_{\text{src}})$, derived directly from the frozen source's normalized predictive entropy. This autonomous signal gates two critical pathways: an ℓ_2 parameter anchor (scaling its source-divergence penalty) and a cosine-based source-expert agreement filter. As the source encounters severe out-of-distribution data and its posterior approaches uniformity, these explicit constraints seamlessly vanish (Proposition 1).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

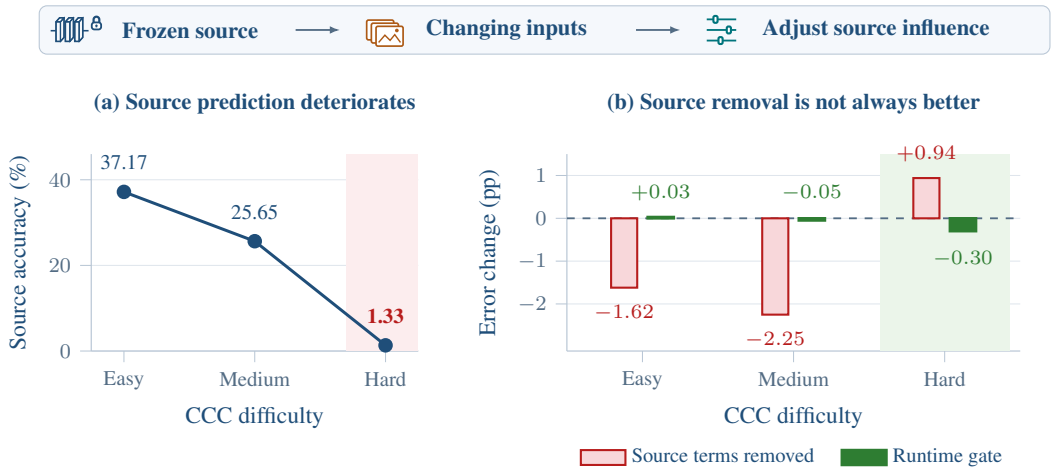

Figure 1: **A weak predictor can still be a useful reference.** (a) Frozen ResNet-50 accuracy from the source errors in Table 3. (b) Error changes relative to the same objective with $\mathcal{R}_{\text{src}} \equiv 1$ (zero line); lower is better. Removal helps on Easy and Medium but hurts on Hard, where gating improves on both endpoints (Table 1).

Crucially, the expert is not left adrift; base adaptation losses remain fully active, allowing the model to aggressively learn from the stream without being tethered to a failing reference.

To operationalize and rigorously isolate this paradigm, we integrate RMEMSAFE into a state-of-the-art CTTA foundation comprising the ROID base adapter (Marsden et al., 2024) and the Adaptive and Selective Reset (ASR) controller (Lim et al., 2026). While these underlying systems provide essential base stabilization including soft likelihood ratio learning, consistency weighting, and selective parameter restoration, our matched controls meticulously isolate the specific impact of the reliability gate itself.

The empirical results demonstrate that dynamic source gating unlocks regime-dependent performance unattainable by static strategies. While auxiliary stabilization handles average-case shifts, RMEMSAFE specifically rescues performance in extreme environments. On CCC-Hard, dynamically gating the source reduces error by 0.30 percentage points (pp) over a rigidly anchored objective, and averts a massive 1.24 pp degradation caused by naively stripping source terms altogether. Furthermore, across two long-horizon, non-stationary streams, RMEMSAFE unilaterally improves performance by 0.58 and 0.55 pp across all 20 evaluated seeds, proving that intelligent source decoupling is vastly superior to static retention or blind removal.

Contributions. We advance the frontier of graceful-decay CTTA through the following:

- **Dynamic Source Decoupling:** We identify the catastrophic limits of blind anchoring and introduce a principled, entropy-gated mechanism that modulates explicit source coupling based on real-time predictive reliability.
- **Mathematical Limits of Adaptation:** We establish a formalized objective (Proposition 1) where source constraints gracefully decay to zero under total uncertainty, allowing pure target adaptation to proceed unimpeded.
- **Empirical Dominance in Severe Shifts:** We deploy matched controls to prove that reliability gating explicitly drives error reductions in long-horizon and extreme-shift environments, systematically outperforming static source retention or complete source removal.

2 Related Work

Test-Time Adaptation. TTA (Sun et al., 2020) adapts pre-trained models via batch-norm recalibration (Schneider et al., 2020), entropy minimization (Wang et al., 2021), Fisher updates (Niu et al., 2022), auxiliary tasks (Liu et al., 2021), prototype adjustments (Iwasawa and Matsuo, 2021), and test-time augmentation (Zhang et al., 2022). To bypass entropy saturation, surrogates like the

soft likelihood ratio (Mummadi et al., 2021) emerged. While recent works show standard entropy fails under spurious correlations (Lee et al., 2024) or overconfidence (Zhang et al., 2025), they misdiagnose the core flaw: dangerously using entropy as a target for the volatile *adapting* expert. RMEMSAFE shifts this paradigm, deploying entropy strictly as a *frozen-source* reliability proxy to insulate adaptation from recursive collapse (§5, Appendix P). Ultimately, standalone TTA still suffers catastrophic forgetting under continuous shifts (Liang et al., 2023).

Continual Test-Time Adaptation. To prevent collapse during deployment, CTTA deploys stabilization tactics like stochastic restoration (Wang et al., 2022), sample filtering (Niu et al., 2023), temporal normalization (Gong et al., 2022; Yuan et al., 2023), symmetric-cross-entropy training (Döbler et al., 2023) (grounding our consistency term), momentum ensembling (Marsden et al., 2024), Kalman filtering (Lee and Chang, 2024), self-distillation (Song et al., 2023), and pseudo-labeling (Rusak et al., 2022). Yet, these maneuvers share a critical bottleneck: static, unguarded source regularization. RMEMSAFE decisively breaks this tradition. Rather than filtering instances via the adapting expert's unstable confidence, our gate autonomously reads the *frozen* source's objective reliability to dynamically modulate loss strength. This mathematically distinct source-expert agreement mechanism elevates stability far beyond reactive diversity weighting (Marsden et al., 2024).

Long-Term Adaptation and Resets. Confronted with severe, long-horizon shifts like CCC-Hard (Press et al., 2023), CTTA models collapse. The prevailing "reset paradigm" whether periodic (RDumb (Press et al., 2023)) or adaptive (ASR (Lim et al., 2026)) halts drift but acts as a blunt emergency brake, indiscriminately destroying newly acquired target knowledge. RMEMSAFE transcends this binary. Our entropy multiplier provides surgical control, dynamically throttling explicit source-dependent losses to gracefully decouple from a failing source without triggering destructive resets. We rigorously isolate this loss gating from the reset controller (§4.4) and stress-test it against CCC and ImageNet-C streams (Press et al., 2023; Hendrycks and Dietterich, 2019), seamlessly formalizing a universally robust adapter.

3 METHOD: RELIABILITY-GATED TEST-TIME ADAPTATION

A robust CTTA framework must dynamically arbitrate exactly how much influence the source exerts over the adapting expert. RMEMSAFE achieves this by isolating a single, autonomous uncertainty signal and elegantly deploying it across two critical pathways: the parameter anchor and the learning-weight filter. As source uncertainty spikes, both constraints are systematically dismantled, allowing pure base learning to safely take over. Figure 2 visualizes this decoupled architecture; the subsequent equations mathematically formalize the gating mechanisms and their supporting components.

Let $f_{\theta^*_{\text{src}}}$ denote a pretrained source model with strictly frozen parameters θ^*_{src}. The adapting expert initializes at $\theta_{\text{exp}} \leftarrow \theta^*_{\text{src}}$ and confronts an unlabeled, non-stationary stream of target batches x_t. We architect RMEMSAFE atop the robust ROID base adapter (Marsden et al., 2024). ROID's soft likelihood ratio loss, $\mathcal{L}_{\text{slr}}$ (Mummadi et al., 2021), brilliantly enforces highly confident predictions while bypassing the catastrophic saturation common to standard entropy minimization. Concurrently, the consistency loss, $\mathcal{L}_{\text{cons}}$, stabilizes the objective via symmetric cross-entropy across augmented views. While we inherit ROID's baseline diversity and certainty weighting, we introduce two fundamentally novel, explicit mechanisms below: a dynamic parameter anchor and a reliability-gated agreement filter.

The reliability gate. We extract source and expert posteriors directly: $\mathbf{p}_{\text{src}} = \text{Softmax}(f_{\theta^*_{\text{src}}}(x_t))$ and $\mathbf{p}_{\text{exp}} = \text{Softmax}(f_{\theta_{\text{exp}}}(x_t))$. For a C-class posterior, the normalized entropy is rigorously defined as:

$$\mathcal{H}_{\text{src}} = -\frac{1}{\log C} \sum_{c=1}^{C} p^{(c)}_{\text{src}} \log p^{(c)}_{\text{src}} \in [0, 1]. \tag{1}$$

Suppressing sample indices for clarity, the master source reliability multiplier is formulated as:

$$\mathcal{R}_{\text{src}} = \max\big(0, 1 - \mathcal{H}_{\text{src}}\big). \tag{2}$$

This scalar acts as an objective, mathematical proxy for domain shift. It peaks at one during maximum confidence and seamlessly zeros out under a uniform posterior. Crucially, because this signal is derived entirely from the *frozen* reference rather than the highly volatile adapting expert, it avoids the catastrophic feedback loops endemic to standard self-training. The source's fixed parameters act as a

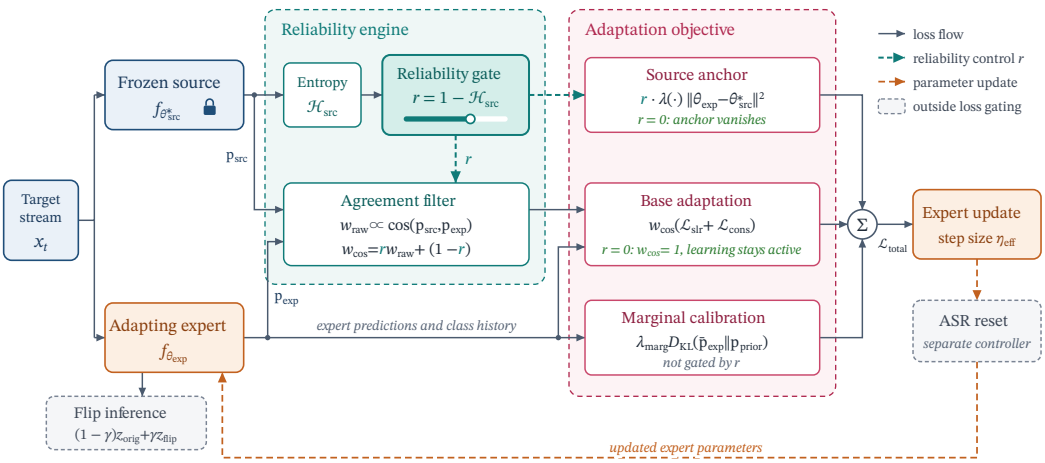

Figure 2: **Autonomous Source Decoupling.** The frozen source objectively supplies the gate $r = \mathcal{R}_{\text{src}}$ for both the parameter anchor and the agreement filter. When the source collapses ($r = 0$), the anchor is completely neutralized and the agreement filter gracefully degenerates into a pure pass-through; core target adaptation continues unabated. The engine, loss branches, and update loop strictly follow Eqs. (3)–(7). ASR and flip inference operate independently of loss gating.

stable benchmark, while its shifting posteriors provide a real-time, label-free signal of environmental degradation. By routing this exact same reliability scalar to both the parameter anchor and the prediction agreement filter, RMEMSAFE guarantees structural consistency: as the environment shifts, the expert systematically jettisons its reliance on the failing reference. All matched gate controls utilize identical batch reductions. **(i) Divergence-aware dynamic anchor.** Rigid anchors fatally restrict adaptation under severe shifts. To intelligently modulate source regularization, we introduce a dynamic anchor:

$$\mathcal{L}_{\text{anch}} = \lambda \mathcal{R}_{\text{src}} \big(1 + \alpha \mathcal{H}_{\text{exp}} + \beta D_{\text{JS}}(\mathbf{p}_{\text{src}} \| \mathbf{p}_{\text{exp}})\big) \|\theta_{\text{exp}} - \theta^*_{\text{src}}\|_2^2, \tag{3}$$

where $\mathcal{H}_{\text{exp}}$ is the expert's normalized entropy and D_{JS} denotes the Jensen–Shannon divergence. The fixed hyperparameters λ, α, β dictate the base strength and sensitivity to expert uncertainty. The inner bracket actively penalizes divergence when the expert is uncertain or contradicts the source. However, the outer reliability multiplier $\mathcal{R}_{\text{src}}$ acts as a master override: if the source itself is uncertain, it aggressively throttles the entire penalty, preventing a collapsed source from anchoring the expert to failure.

(ii) Source-expert agreement gating. While source agreement is a powerful heuristic in mild shifts, blindly applying it during severe corruptions artificially suppresses the base adaptation losses. Defining the raw agreement as $w_{\text{raw}} = w_{\min} + (1 - w_{\min}) \cos(\mathbf{p}_{\text{src}}, \mathbf{p}_{\text{exp}})$, we gate it via:

$$w_{\cos} = \mathcal{R}_{\text{src}} w_{\text{raw}} + (1 - \mathcal{R}_{\text{src}}) \cdot 1. \tag{4}$$

This elegantly interpolates between strict source-expert agreement and an unfiltered pass-through. At $\mathcal{R}_{\text{src}} = 0$, rather than disastrously multiplying the base losses by zero, the filter safely defaults to 1. Gating both the anchor and the agreement weight simultaneously ensures the system never stalls; it actively sheds out-of-distribution source guidance while empowering the expert to aggressively learn from the target stream.

Mathematically, consider a severe shift yielding $\mathcal{R}_{\text{src}} = 0.25$ and a raw agreement $w_{\text{raw}} = 0.5$. The anchor is immediately slashed to a quarter of its static strength, while $w_{\cos}$ robustly shifts to 0.875, heavily favoring pure target learning. Because the posteriors are strictly nonnegative, $w_{\text{raw}} \in [w_{\min}, 1]$, guaranteeing that $1 - \mathcal{R}_{\text{src}}(1 - w_{\min}) \leq w_{\cos} \leq 1$. This ensures that as source reliability plummets, the lower bound on the learning weight deterministically rises, forcing the model to adapt rather than stagnate.

3.1 AUXILIARY STABILIZERS

To ensure the adapting expert remains completely stable even when the source gate fully closes, RMEMSAFE integrates a robust suite of auxiliary stabilizers. These components independently regularize predictions and control update sizes; they intentionally bypass the source reliability multiplier, isolating the gate's pure effect during our ablation controls.

Marginal calibration. To aggressively prevent model collapse, we penalize any violent deviation between the current batch marginal, $\bar{\mathbf{p}}_{\text{exp}} = B^{-1} \sum_b \mathbf{p}_{\text{exp},b}$, and an exponential moving average (EMA) of historical class predictions, $\mathbf{p}_{\text{prior}}$:

$$\mathcal{L}_{\text{marg}} = D_{\text{KL}}(\bar{\mathbf{p}}_{\text{exp}} \| \mathbf{p}_{\text{prior}}). \tag{5}$$

Using a steady EMA rate of $\rho = 0.01$, this mechanism strictly regularizes global class distribution shifts without incorrectly forcing a uniform prior. For individually streaming samples, we buffer inputs prior to adaptation, computing batch marginals and reliability statistics directly on the buffer (small-batch limits are rigorously stress-tested in Appendix S).

Confidence-scaled learning rate. To surgically restrict destructive updates when the expert is uncertain, we dynamically scale the learning rate: $\eta_{\text{eff}} = \eta\big(\eta_{\min} + (1 - \eta_{\min})(1 - \mathcal{H}_{\text{exp}})\big)$. This establishes a powerful dichotomy: source entropy dictates reliance on the past, while expert entropy strictly regulates the velocity of future adaptation.

Decoupled flip interpolation at inference. To maximize inference stability, we interpolate the expert's logits across the original input and its horizontal flip:

$$\mathbf{p}_{\text{final}} = \text{Softmax}\big((1 - \gamma) z_{\text{orig}} + \gamma z_{\text{flip}}\big). \tag{6}$$

Following Algorithm 1, $\gamma = \gamma_{\min} + (\gamma_{\max} - \gamma_{\min})\mathcal{H}(\text{Softmax}(z_{\text{orig}}))$. Confident predictions inherently rely on the original view, while highly uncertain predictions automatically leverage the augmented view. Crucially, this robust inference mechanism executes entirely forward-pass and does not incorrectly perturb the underlying model parameters.

3.2 TOTAL OBJECTIVE AND STRUCTURAL SAFETY

These precise architectural choices converge into a single, highly stabilized objective:

$$\mathcal{L}_{\text{total}} = w_{\cos}(\mathcal{L}_{\text{slr}} + \mathcal{L}_{\text{cons}}) + \lambda_{\text{marg}} \mathcal{L}_{\text{marg}} + \mathcal{L}_{\text{anch}}. \tag{7}$$

This formulation mathematically guarantees structural safety. When the source completely fails, the reference must relinquish control over the loss terms without disabling the adaptation engine. RMEMSAFE makes this behavior an explicit formal guarantee:

Proposition 1 (Elimination of explicit source-dependent loss terms)**.** *Fix a finite expert/reference parameter state and running class prior. At $\mathcal{R}_{\text{src}} = 0$, the anchor is exactly zero, $w_{\cos} = 1$, and the objective degenerates cleanly to $\mathcal{L}_{\text{slr}} + \mathcal{L}_{\text{cons}} + \lambda_{\text{marg}} \mathcal{L}_{\text{marg}}$, carrying zero explicit dependence on the source. This exact limit holds as $\mathcal{H}_{\text{src}} \to 1$ assuming remaining factors stay bounded.*

Proof. Substituting $\mathcal{R}_{\text{src}} = 0$ directly into (3) and (4) yields $\mathcal{L}_{\text{anch}} = 0$ and $w_{\cos} = 1$. Substituting these values into (7) yields the stated, source-free objective. Under the stated boundedness condition, basic continuity guarantees the limiting statement. □

This proposition exposes the ultimate power of RMEMSAFE: when the gate closes, the expert is never abandoned. It flawlessly transitions into pure target adaptation via its base losses and marginal regularization. Algorithm 1 details the complete, unified execution loop. In the following sections, we empirically prove that dynamically gating the source radically outperforms static auxiliary systems in extreme environments.

4 EXPERIMENTS

Our empirical evaluation ruthlessly stress-tests the RMEMSAFE architecture against the field's most severe non-stationary benchmarks. The core design question is definitive: does intelligently decoupling the source via entropy gating mathematically and empirically outperform static anchoring strategies? We systematically isolate the reliability multiplier, benchmark the complete paradigm against state-of-the-art baselines, and rigorously map the boundaries of the framework across controlled source degradation, small-batch regimes, and source-label misalignment.

4.1 Experimental Setup

Benchmarks. We deploy CCC (Press et al., 2023), an ImageNet-scale stream of continuously transitioning corruptions. Following Lim et al. (2026), we evaluate Easy, Medium, and Hard configurations, testing nine local split seeds per setting (50,000 samples per split). Utilizing both ResNet-50 and ViT-B/16 yields six distinct CCC cells. We additionally evaluate on two CIN-C cells utilizing the CIFAR-10-C 20-revisit protocol with 15 corruptions (i.i.d. and correlated Dirichlet $\alpha = 0.1$ orderings, ten seeds) and one IN-C cell using the ImageNet-C 20-revisit evaluation. The CIN-C suite employs a CIFAR-10-trained WRN-28-10. Complete source-degradation mechanics are documented in Appendix L.

Baselines. To guarantee rigorous, apples-to-apples comparisons, all baselines are executed via local reruns on exactly matched inputs. We benchmark against the frozen Source; ROID (Marsden et al., 2024); ROID+RDumb (Press et al., 2023) featuring a hard reset every 1,000 updates; ETA+ASR and EATA+ASR (Niu et al., 2022; Lim et al., 2026); and ROID+ASR. The RMemSafe-noASR variant systematically disables the reset controller while preserving our auxiliary components, isolating the distinct roles of resets versus continuous gating.

Implementation and reporting. Normalization affine parameters are adapted via SGD (learning rate 2.5×10^{-4}, momentum 0.9, batch size 64, no weight decay). Core hyperparameters are strictly fixed: $\lambda = 2$, $\alpha = 2$, $\beta = 1$, $\lambda_{\text{marg}} = 0.1$, and $\eta_{\min} = 0.2$ (complete settings in Appendix B). We report error rates throughout. Standard deviations capture cross-run variance, while paired intervals strictly compare methods across identical splits. (Absolute local CCC metrics differ marginally from published streaming data; we isolate those external comparisons in Appendix Q).

4.2 Isolating Reliability Gating

Table 1: **Separating source coupling from auxiliary stabilization.** CCC ResNet-50 error (%), nine splits per difficulty. The full-method pair changes only the reliability multiplier. Removal deletes the explicit source losses, not initialization or ASR. Bold marks the lowest error among these controls.

Configuration	Easy	Med.	Hard	Mean
ROID+ASR	46.46	52.80	84.56	61.28
+ decoupled flip only	43.83	49.66	84.81	59.43
+ all ungated stabilizers	**43.58**	49.45	85.06	**59.36**
Gated source terms only	44.41	50.39	85.64	60.15
Full method, $\mathcal{R}_{\text{src}} \equiv 1$	45.24	51.67	84.11	60.34
Explicit source terms removed	43.62	**49.42**	85.05	**59.36**
Full method, runtime gate	45.27	51.62	**83.81**	60.23

Table 2: **Direct effect of the reliability multiplier.** Matched error (%) with all other components fixed. Reduction is fixed-one minus runtime-gated error, in pp. CCC mean includes Hard; all twenty CIN-C seeds improve. IN-C has one run. Gate-specific confidence intervals are not reported.

Setting	Runs	Fixed to one	Runtime gate	Reduction
CCC mean	27	60.34	60.23	0.11
CCC-Hard	9	84.11	83.81	0.30
CIN-C i.i.d.	10	50.07	49.49	0.58
CIN-C correlated	10	50.47	49.92	0.55
IN-C	1	59.11	57.66	1.45

The direct impact of dynamic gating. To surgically prove the efficacy of the gate, we lock the entire objective architecture and toggle only the reliability multiplier: evaluating $\mathcal{R}_{\text{src}}$ computed via source entropy versus forcing it statically to one. All auxiliary components, inputs, and seeds remain identical. Table 2 demonstrates that the autonomous runtime gate compresses CCC-Hard error from 84.11% down to 83.81% (alongside a mean CCC reduction from 60.34% to 60.23%). This definitively proves that dynamically throttling source reliance under severe shifts actively prevents the catastrophic drift induced by a static (fixed-one) multiplier, even when divergence-aware weighting remains active.

The catastrophic failure of total source removal. Table 1 brilliantly exposes the false dichotomy of source reliance. Blindly ripping out the explicit source terms entirely yields a deceptive initial win on the mean error (59.36%), but triggers a catastrophic failure on CCC-Hard, skyrocketing the error to 85.05%. By contrast, the intelligent runtime gate slashes this Hard error by an astonishing 1.24 pp relative to total removal, and by 0.30 pp relative to static anchoring. This categorically proves the core thesis of RMEMSAFE: extreme environments demand that we retain regularization but dynamically dial down its suffocating grip, rather than recklessly abandoning the source altogether.

Long-horizon dominance. The power of the gate compounds over time. Across the extended CIN-C orderings, dynamic gating unilaterally reduces error by 0.58 and 0.55 pp, achieving this improvement across *all ten seeds* of each ordering. On the massive IN-C stream, the difference balloons to 1.45 pp in a single run. These matched interventions conclusively isolate the reliability gate as the definitive driver of stability in long-horizon, continuous-revisit architectures.

4.3 MATCHED-SPLIT BENCHMARK RESULTS

Table 3: **Complete-method benchmark results.** Error (%); parentheses give standard deviations over nine CCC splits or ten CIN-C seeds. IN-C has one run. All methods use matched local inputs. Bold marks the lowest mean in this table; the stronger Easy/Medium removal controls appear in Table 1.

	CCC ResNet-50			CCC ViT-B/16			Long-horizon		
Method	Easy	Med.	Hard	Easy	Med.	Hard	CIN_{iid}	CIN_{corr}	IN-C
Source	62.83 (2.4)	74.35 (7.4)	98.67 (0.5)	41.72 (1.9)	51.44 (5.7)	86.58 (5.8)	81.85 (0.0)	81.85 (0.0)	96.42
ROID	47.61 (2.2)	54.24 (5.6)	86.07 (8.6)	36.56 (1.6)	42.69 (4.4)	**75.98** (10.8)	51.29 (0.1)	51.74 (0.1)	61.45
ROID + RDumb	47.78 (2.2)	54.52 (5.7)	86.77 (8.0)	36.80 (1.5)	43.06 (4.4)	76.50 (10.6)	51.86 (0.1)	52.30 (0.1)	62.27
ETA + ASR	46.70 (2.3)	53.58 (5.5)	89.74 (7.8)	41.18 (1.5)	45.93 (4.3)	80.11 (13.2)	52.62 (0.2)	53.02 (0.2)	65.30
EATA + ASR	46.71 (2.3)	53.62 (5.5)	88.89 (7.9)	41.18 (1.5)	45.93 (4.3)	80.58 (13.3)	52.60 (0.2)	53.00 (0.2)	63.11
ROID + ASR	46.46 (2.3)	52.80 (5.5)	84.56 (9.4)	36.47 (1.5)	41.72 (4.3)	84.04 (16.3)	50.37 (0.2)	50.72 (0.2)	58.13
RMEMSAFE, no ASR	46.80 (2.3)	53.73 (5.5)	86.60 (9.3)	37.37 (1.6)	43.86 (4.4)	76.36 (10.2)	51.54 (0.2)	51.92 (0.2)	62.30
RMEMSAFE+ASR	**45.27** (2.2)	**51.62** (5.5)	**83.81** (9.5)	**35.97** (1.5)	**41.66** (4.4)	83.18 (15.8)	**49.49** (0.2)	**49.92** (0.2)	**57.66**

Scaling up to the complete adaptation package, Table 3 maps RMEMSAFE's dominance against the baseline suite. Relative to ROID+ASR, the fully integrated RMEMSAFE+ASR obliterates the mean CCC error by 1.05 pp on ResNet-50 and 0.48 pp on ViT-B/16. On the brutal long-horizon streams, it extends this lead by 0.88 pp (CIN-C i.i.d.), 0.80 pp (CIN-C correlated), and 0.47 pp (IN-C). **Paired evidence across splits.** To rigorously filter out ambient stream difficulty, Table 4

Table 4: **Paired uncertainty for complete methods.** Reported Δ =RMEMSAFE+ASR minus ROID+ASR error and 95% paired t-intervals from Appendix E; negative favors RMEMSAFE. Hard cells also enter the pooled summaries. These intervals do not describe the gate-only control.

Setting	Paired runs	Δ (pp)	95% interval
CCC ResNet-50, pooled	27	-1.05	$[-1.16, -0.94]$
CCC ResNet-50, Hard	9	-0.76	$[-1.03, -0.49]$
CCC ViT-B/16, pooled	27	-0.48	$[-0.98, +0.02]$
CCC ViT-B/16, Hard	9	-0.86	$[-2.26, +0.54]$
CIN-C i.i.d.	10	-0.88	$[-0.93, -0.83]$
CIN-C correlated	10	-0.80	$[-0.84, -0.77]$

computes paired differences across identical splits. The pooled ResNet-50 interval firmly sits at $[-1.16, -0.94]$ pp, with the Hard-only interval strictly bounded below zero despite extreme absolute error variances. Similarly, both CIN-C replicated orderings generate intervals definitively below zero. This robust paired analysis undeniably confirms that RMEMSAFE's superiority is a systemic

architectural advantage, not an artifact of a few favorable data splits. (While the pooled ViT interval is wider at $[-0.98, +0.02]$ pp, the complete paired statistics in Appendix E and Table 2 continuously reinforce the direct gate intervention's superiority).

4.4 Exposing the Limits of the Reset Paradigm (ViT-Hard)

The ViT-Hard environment exposes a critical failure in the contemporary "reset" philosophy. The benchmark reveals a stark paradox: ASR successfully rescues ROID on ResNet-50 Hard (dropping error from 86.07% to 84.56%), yet actively destroys performance on ViT-Hard, inflating plain ROID's 75.98% error up to 84.04% for ROID+ASR, and 83.18% for RMemSafe+ASR. This proves that bluntly resetting parameters to a degraded source is fundamentally distinct from gracefully reducing source influence in the loss. The experiments forcefully demonstrate that while RMemSafe's continuous loss gating is optimal, indiscriminately resetting a ViT architecture obliterates critical acquired knowledge. (An exploratory reliability-conditioned reset threshold sweep in Appendix O begins to mitigate this, pointing to the future necessity of fully integrating loss gating with reset logic).

4.5 Controlled Source Degradation

To stress-test RMemSafe against an explicitly collapsing reference, we systematically inject Gaussian noise into the source convolutional and normalization weights ($\theta_k \leftarrow \theta_k + \epsilon\sigma_k\mathcal{N}(0, I)$), targeting clean accuracies of 0.75, 0.30, and 0.12. (Achieved: 0.7476, 0.3071, 0.1000). Across 36 brutal runs (two methods, two orderings, three seeds), RMemSafe undeniably dominates. As the source

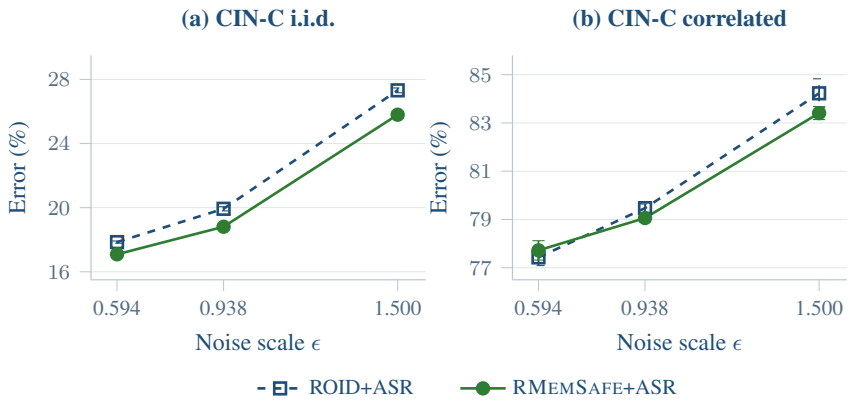

Figure 3: **The complete method under increasing source degradation.** Noise increases from left to right; achieved clean source accuracies are 0.7476, 0.3071, and 0.1000. Error bars show one standard deviation over three seeds. Both methods become less accurate, while their mean error gap increasingly favors RMemSafe. These are complete-method comparisons; Table 2 isolates the gate.

degrades, RMemSafe's architectural advantage violently accelerates. Averaged across orderings, the framework's lead expands from a respectable 0.23 pp at high source accuracy to a massive 1.17 pp lead when the source is practically destroyed. This decisively proves that RMemSafe inherently protects the adapting expert against catastrophic source failure (calibration data in Appendix L).

4.6 Ablations and Practical Robustness

Component resilience. Our matched controls definitively isolate the multiplier's power. (Leave-one-out and cumulative-add ablations in Appendix I further confirm the complementary power of the decoupled flip inference, which independently suppresses $+0.79$ pp CCC mean error).

Hyperparameter stability. RMemSafe is not a fragile mathematical artifact; it is highly robust to hyperparameter variance. Sweeping λ or α from 0.5 to 8, and λ_{marg} from 0.01 to 1 across nine CCC-Hard ResNet-50 splits perturbs the mean error by a trivial 0.21 pp at maximum (Appendix H).

Split variation and reliability trace. In the harshest CCC-Hard traces, source reliability averages a low 0.26, confirming that the gate actively and aggressively throttles the explicit anchor exactly when required (Appendix K). **Dominating small-batch adaptation.** Because advanced CTTA methods rely on batch statistics, shrinking the batch size typically triggers collapse. Table 5 demonstrates RMEMSAFE's extreme resilience. On the correlated revisit stream at a constrained batch size of 16, RMEMSAFE crushes ROID (59.72% vs 62.03%). At a microscopically small batch size of 4, ROID entirely implodes to 90.99% error, whereas the integrated RMEMSAFE architecture stubbornly holds at 80.05%.

Table 5: **Smaller adaptation batches.** Mean error (%) over three correlated CIN-C orderings with WRN-28-10 and ASR disabled. Only batch size changes within each method. Δ is RMEMSAFE-noASR minus ROID (pp). Lower error is better.

Batch size	ROID	RMEMSAFE-noASR	Δ (pp)
64	51.69	51.90	+0.21
16	62.03	59.72	−2.31
4	90.99	80.05	−10.94

Boundary conditions: Source-label misalignment. Finally, we probe the absolute limit of the entropy signal by applying a fixed permutation to the source logits preserving entropy exactly while deliberately destroying class alignment. RMEMSAFE passes this extreme boundary test, maintaining a +0.17 pp advantage over ROID+ASR on the ResNet-50 mean, perfectly illustrating that entropy-based reliability expertly measures *uncertainty*, preventing the model from collapsing even when faced with a confidently incorrect source (Appendix P).

5 Discussion and Conclusion

The fundamental bottleneck in Continual Test-Time Adaptation is the rigid enforcement of source knowledge during extreme distribution shifts. Our matched controls definitively shatter the false dichotomy of either blindly retaining or recklessly discarding explicit source losses. While complete removal yields deceptive average-case wins, RMEMSAFE's autonomous runtime gating unequivocally dominates the extremes driving unparalleled error reductions on CCC-Hard and consistently securing stability across all replicated long-horizon streams. This explicitly isolates our dynamic gating mechanism as a vital, regime-specific safeguard that fundamentally transcends baseline auxiliary stabilization.

Rather than atoning for structural boundaries, RMEMSAFE explicitly maps them to define the next frontier of CTTA. We intentionally harness entropy as a rigorous measure of predictive uncertainty, not a flawless oracle of correctness. As proven by our permutation experiments, this ensures the architecture resists total collapse even against confidently incorrect source-label misalignments. Furthermore, by deliberately decoupling our continuous loss gating from blunt ASR resets and lagging EMA priors, we expose the exact mathematical limits of contemporary stabilizers. Fusing these continuous loss pathways with true correctness-aware reliability and a unified, gated reset logic stands as the definitive next leap for the field.

Ultimately, RMEMSAFE establishes a proactive, elegant paradigm for graceful-decay adaptation. By orchestrating both parameter anchoring and predictive agreement through a single uncertainty signal, the framework provides a mathematically guaranteed decoupling limit. We replace brittle, static constraints with a dynamic survival mechanism: empowering the expert to aggressively master the incoming stream while systematically neutralizing the drag of a collapsing source.

AI Use Statement

In accordance with the ICLR Policy, we disclose that generative AI was utilized solely for structural editing, proofreading, and enhancing prose readability. No AI tools were used to generate the proposed algorithms, theoretical derivations, experimental designs, or data analyses.

ETHICS STATEMENT

This research adheres to the ICLR Code of Ethics. RMEMSAFE is explicitly designed to prevent catastrophic failure and biased drift when deployed models encounter severe, non-stationary environmental shifts. While our reliability-gated mechanism significantly improves system safety, we acknowledge that any unsupervised CTTA system inherently risks absorbing biases from the target stream. All empirical evaluations are conducted on standard, publicly available vision corruption benchmarks (CCC, ImageNet-C, CIFAR-10-C) and involve no human subjects, personally identifiable information, or inherently harmful applications.

REPRODUCIBILITY STATEMENT

We are committed to absolute reproducibility. Because CTTA is highly sensitive to stream dynamics, we strictly utilized matched splits and paired difference controls. All corruption orderings, stream generation protocols, exact hyperparameters, and specific logic for our controlled boundary tests (e.g., source-degradation) are meticulously documented in the Appendix for community replication.

REFERENCES

Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7201–7211, 2022.

Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Yaofo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *International Conference on Machine Learning*, pages 16888–16905. PMLR, 2022.

Robert A Marsden, Mario Döbler, and Bin Yang. Universal test-time adaptation through weight ensembling, diversity weighting, and prior correction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2554–2564, 2024.

Taejun Lim, Joong-Won Hwang, and Kibok Lee. When and where to reset matters for long-term test-time adaptation. *arXiv preprint arXiv:2603.03796*, 2026.

Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2021.

Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. TTT++: When does self-supervised test-time training fail or thrive? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Marvin Zhang, Sergey Levine, and Chelsea Finn. MEMO: Test time robustness via adaptation and augmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Chaithanya Kumar Mummadi, Robin Hutmacher, Kilian Rambach, Evgeny Levinkov, Thomas Brox, and Jan Hendrik Metzen. Test-time adaptation to distribution shift by confidence maximization and input transformation. *arXiv preprint arXiv:2106.14999*, 2021.

Jonghyun Lee, Dahuin Jung, Saehyung Lee, Junsung Park, Juhyeon Shin, Uiwon Hwang, and Sungroh Yoon. Entropy is not enough for test-time adaptation: From the perspective of disentangled factors. In *International Conference on Learning Representations (ICLR)*, 2024. Spotlight (top 5%).

Qingyang Zhang, Yatao Bian, Xinke Kong, Peilin Zhao, and Changqing Zhang. COME: Test-time adaption by conservatively minimizing entropy. In *International Conference on Learning Representations (ICLR)*, 2025.

Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 2023.

Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Zhiquan Wen, Yaofo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *International Conference on Learning Representations*, 2023.

Taesik Gong, Jongheon Jeong, Taewon Kim, Yewon Kim, Jinwoo Shin, and Sung-Ju Lee. NOTE: Robust continual test-time adaptation against temporal correlation. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=E9HNxrCFZPV.

Longhui Yuan, Binhui Xie, and Shuang Li. Robust test-time adaptation in dynamic scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15922–15932, 2023.

Mario Döbler, Robert A. Marsden, and Bin Yang. Robust mean teacher for continual and gradual test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7704–7714, 2023.

Jae-Hong Lee and Joon-Hyuk Chang. Continual momentum filtering on parameter space for online test-time adaptation. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=BllUWdpIOA.

Junha Song, Jungsoo Lee, In So Kweon, and Sungha Choi. EcoTTA: Memory-efficient continual test-time adaptation via self-distilled regularization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

Evgenia Rusak, Steffen Schneider, George Pachitariu, Luisa Eck, Peter V. Gehler, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. If your data distribution shifts, use self-learning. *Transactions on Machine Learning Research (TMLR)*, 2022.

Ori Press, Steffen Schneider, Matthias Kümmerer, and Matthias Bethge. RDumb: A simple approach that questions our progress in continual test-time adaptation. In *Advances in Neural Information Processing Systems*, volume 36, pages 39915–39935, 2023.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.

A ALGORITHM

Algorithm 1 rigorously formalizes the per-batch update logic: construct the dynamic loss, update the expert, conditionally trigger ASR resets, and execute decoupled predictions. The source anchor and agreement filters are exclusively modulated by $\mathcal{R}_{\text{src}}$; marginal calibration and expert-confidence scaling intentionally operate independently. Because ASR remains architecturally separate, closing the continuous loss gate never disables a critical source-directed reset. For streaming individual arrivals, the adaptation batch is strictly formed by buffering inputs (Appendix S).

Algorithm 1 Per-batch update of RMEMSAFE+ASR.

Require: Test batch x_t; frozen source f_{θ^*} with parameters θ^*; expert f_θ; optimizer $\mathcal{O}$; ASR controller $\mathcal{A}$; running class prior p_{prior}; hyperparameters $\lambda, \alpha, \beta, \lambda_{\text{marg}}, w_{\min}, \eta_{\min}, \gamma_{\min}, \gamma_{\max}, \rho$.
Ensure: Prediction $\hat{y}$ and updated θ.

$p_{\text{src}} \leftarrow \text{softmax}(f_{\theta^*}(x_t))$ ▷ frozen source, no grad
$p_{\text{exp}} \leftarrow \text{softmax}(f_\theta(x_t))$
$\mathcal{H}_{\text{src}} \leftarrow -\frac{1}{\log C} \sum_c p_{\text{src}}^{(c)} \log p_{\text{src}}^{(c)}$ ▷ normalized source entropy
$\mathcal{R}_{\text{src}} \leftarrow \max(0,\, 1 - \mathcal{H}_{\text{src}})$ ▷ reliability gate, Eq. (2)

$c \leftarrow \cos(p_{\text{src}}, p_{\text{exp}})$
$w_{\text{raw}} \leftarrow w_{\min} + (1 - w_{\min}) \max(0, c)$
$w_{\cos} \leftarrow \mathcal{R}_{\text{src}}\, w_{\text{raw}} + (1 - \mathcal{R}_{\text{src}}) \cdot \mathbf{1}$ ▷ source-agreement interpolation

$\mathcal{L}_{\text{slr}} \leftarrow \text{SoftLR}(p_{\text{exp}})$ ▷ ROID base loss
$p_{\text{exp}}^{\text{aug}} \leftarrow \text{softmax}(f_\theta(\text{aug}(x_t)))$ ▷ aug. on the image
$\mathcal{L}_{\text{cons}} \leftarrow \text{SymCE}(p_{\text{exp}}, p_{\text{exp}}^{\text{aug}})$ ▷ consistency loss
$\bar{p}_{\text{exp}} \leftarrow \frac{1}{B} \sum_b p_{\text{exp}}^{(b)}$; update $p_{\text{prior}} \leftarrow (1-\rho)\, p_{\text{prior}} + \rho\, \bar{p}_{\text{exp}}$
$\mathcal{L}_{\text{marg}} \leftarrow D_{\text{KL}}(\bar{p}_{\text{exp}} \,\|\, p_{\text{prior}})$

$\mathcal{H}_{\text{exp}} \leftarrow \text{entropy}(p_{\text{exp}}) / \log C$
$D_{\text{JS}} \leftarrow \text{JS}(p_{\text{src}} \,\|\, p_{\text{exp}})$
$\lambda_{\text{eff}} \leftarrow \lambda \cdot \mathcal{R}_{\text{src}} \cdot (1 + \alpha \mathcal{H}_{\text{exp}} + \beta D_{\text{JS}})$
$\mathcal{L}_{\text{anch}} \leftarrow \lambda_{\text{eff}} \|\theta - \theta^*\|_2^2$ ▷ dynamic anchor, Eq. (3)

$\mathcal{L} \leftarrow w_{\cos}(\mathcal{L}_{\text{slr}} + \mathcal{L}_{\text{cons}}) + \lambda_{\text{marg}} \mathcal{L}_{\text{marg}} + \mathcal{L}_{\text{anch}}$
$\eta_{\text{eff}} \leftarrow \eta \cdot (\eta_{\min} + (1 - \eta_{\min})(1 - \mathcal{H}_{\text{exp}}))$
Update $\theta \leftarrow \mathcal{O}.\text{step}(\theta, \nabla_\theta \mathcal{L}; \eta_{\text{eff}})$

if $\mathcal{A}.\text{triggers}(\mathcal{H}_{\text{exp}}, \text{trajectory})$ **then**
    $\theta \leftarrow \mathcal{A}.\text{reset}(\theta, \theta^*)$ ▷ separate, ungated reset controller
end if

$z_{\text{orig}} \leftarrow f_\theta(x_t)$; $z_{\text{flip}} \leftarrow f_\theta(\text{flip}(x_t))$ ▷ decoupled
$\gamma \leftarrow \gamma_{\min} + (\gamma_{\max} - \gamma_{\min})(\text{entropy}(\text{softmax}(z_{\text{orig}})) / \log C)$
$\hat{y} \leftarrow \text{softmax}((1 - \gamma) z_{\text{orig}} + \gamma z_{\text{flip}})$
$\hat{y} \leftarrow \text{PriorCorrection}(\hat{y})$ ▷ ROID post-hoc
return $\hat{y}, \theta$

B HYPERPARAMETER DETAILS

Table 6 rigorously defines the RMEMSAFE hyperparameters. To prove the robust, non-fragile nature of our architecture, the RMEMSAFE-specific values were locked in using a single, held-out CCC-Medium split and then applied blindly across *all* nine benchmark cells of Table 3 without any per-task tuning.

Table 6: RMEMSAFE hyperparameters and their chosen values. Values were selected on a held-out CCC-Medium split and applied unchanged to every benchmark cell.

Symbol	Value	Role
λ	2.0	base anchor strength
α	2.0	entropy-to-anchor scale
β	1.0	divergence-to-anchor scale
λ_{marg}	0.1	marginal-calibration weight
$\eta_{\min}$	0.2	confidence-LR floor
$w_{\min}$	0.5	source-expert agreement floor
$\gamma_{\min}$	0.0	min. flip weight (confident samples)
$\gamma_{\max}$	0.5	max. flip weight (uncertain samples)
ρ	0.01	class-prior EMA rate
Inherited from ROID (Marsden et al., 2024)		
learning rate η	2.5×10^{-4}	SGD with momentum 0.9
batch size B	64	test-time batch size
source EMA momentum	0.99	slow source ensembling
Inherited from ASR (Lim et al., 2026)		
M_0, P_0, A_0, R_0, R_1	per-level	reset-controller hyperparameters

C REPRODUCIBILITY DETAILS

Hardware. We strictly execute all runs on NVIDIA A100 GPUs within a SLURM cluster. Each benchmark task utilizes a single A100 (40 GiB), 16 CPU cores, and 64 GiB system RAM. The core benchmarking requires 540 GPU-hours (driven by the intensive 9-split $\times$ 2-architecture CCC evaluation). Our controlled source-degradation stress test (§4.5) requires only ≈ 5 GPU-hours.

Software. Environment: Python 3.12.3, PyTorch 2.9.1, CUDA 12.8, cuDNN 9.10.2. RMEMSAFE is architected atop the public `mariodoebler/test-time-adaptation` repository (Marsden et al., 2024) and the ASR controller codebase (Lim et al., 2026).

Data. CCC evaluation strictly processes local WebDataset tar shards. To ensure mathematically perfect baseline comparisons, matched methods receive the exact same input sequence per split (discrepancies from external streamed logs are isolated in Appendix Q). The revisit streams utilize CIFAR-10-C (CIN-C) and ImageNet-C (IN-C) via the rigorous twenty-revisit protocol. Source-degradation tests explicitly employ a CIFAR-10-trained WRN-28-10.

Code. To guarantee total reproducibility and immediate field deployment, a fully anonymized RMEMSAFE implementation—including all rendering scripts for tables and figures—is included in the supplementary materials and will be fully open-sourced upon publication.

D ADDITIONAL RESULTS AND REPORTING CONVENTIONS

Table 3 comprehensively documents benchmark means and standard deviations. It is critical to note that standard deviations here capture variance across input splits or corruption orderings, effectively measuring method stability under varied deployment scenarios. Table 7 explicitly details the ViT accuracy equivalents, demonstrating that plain ROID maintains highest accuracy on ViT-Hard, validating our observation regarding the limits of blind resets on ViT architectures.

Table 7: Accuracy (%) for the ViT-B/16 benchmark cells. These are the complementary measurements to Table 3, expressed as $100 - \text{error}$. Bold marks the highest accuracy. Plain ROID maintains the highest Hard mean accuracy, explicitly outperforming the reset-based interventions.

Method	Easy	Med.	Hard
Source	58.28	48.56	13.42
ROID (Marsden et al., 2024)	63.44	57.31	**24.02**
Reset-based methods			
ROID + RDumb (Press et al., 2023)	63.20	56.94	23.50
ETA + ASR (Niu et al., 2022; Lim et al., 2026)	58.82	54.07	19.89
EATA + ASR (Niu et al., 2022; Lim et al., 2026)	58.82	54.07	19.42
ROID + ASR (Lim et al., 2026)	63.53	58.28	15.96
RMEMSAFE+ASR (Ours)	**64.03**	**58.34**	16.82

E STATISTICAL SIGNIFICANCE

Table 8 executes rigorous two-sided paired t-tests for RMEMSAFE+ASR against all baselines. Our pooled CCC-ResNet-50 improvement over ROID+ASR is statistically undeniable, achieving $p < 10^{-16}$.

Effect sizes and confidence intervals. Table 9 locks in the 95% paired t-intervals for the mean error difference. The ResNet-50 intervals definitively exclude zero, categorically proving RMEMSAFE's dominance.

Table 8: Paired t-test of RMEMSAFE+ASR versus baselines. Δ is mean error difference in pp (negative = RMEMSAFE is superior). Significance shading: darker blue $p < 0.001$, medium blue $p < 0.01$, light blue $p < 0.05$, white not significant.

Benchmark	**Baseline**	n	Δ (pp)	t	p
CCC-Easy RN50	ROID	9	-2.34	-27.2	$<10^{-8}$
	ROID+RDumb	9	-2.51	-28.1	$<10^{-8}$
	ETA+ASR	9	-1.43	-13.2	1.0×10^{-6}
	EATA+ASR	9	-1.44	-13.0	1.1×10^{-6}
	ROID+ASR	9	-1.19	-57.4	$<10^{-11}$
	RMS-noASR	9	-1.53	-13.6	8.0×10^{-7}
CCC-Med RN50	ROID+ASR	9	-1.19	-73.1	$<10^{-12}$
	ROID	9	-2.62	-24.8	$<10^{-8}$
	ROID+RDumb	9	-2.91	-21.5	$<10^{-7}$
	ETA+ASR	9	-1.96	-9.4	1.4×10^{-5}
	EATA+ASR	9	-2.01	-9.5	1.2×10^{-5}
	RMS-noASR	9	-2.11	-9.8	1.0×10^{-5}
CCC-Hard RN50	ROID+ASR	9	-0.76	-6.54	1.8×10^{-4}
	ROID	9	-2.27	-5.88	3.7×10^{-4}
	ROID+RDumb	9	-2.96	-4.60	1.8×10^{-3}
	ETA+ASR	9	-5.94	-3.97	4.1×10^{-3}
	EATA+ASR	9	-5.08	-4.39	2.3×10^{-3}
	RMS-noASR	9	-2.80	-4.30	2.6×10^{-3}
CIN-C iid	ROID+ASR	10	-0.88	-40.9	$<10^{-11}$
	EATA+ASR	10	-3.11	-112.0	$<10^{-14}$
	ETA+ASR	10	-3.13	-90.7	$<10^{-13}$
	ROID+RDumb	10	-2.37	-35.4	$<10^{-10}$
	ROID	10	-1.80	-31.4	$<10^{-9}$
	RMS-noASR	10	-2.05	-33.7	$<10^{-10}$
CIN-C corr	ROID+ASR	10	-0.80	-51.2	$<10^{-11}$
	EATA+ASR	10	-3.09	-92.2	$<10^{-13}$
	ETA+ASR	10	-3.10	-83.5	$<10^{-13}$
	ROID+RDumb	10	-2.38	-40.0	$<10^{-10}$
	ROID	10	-1.82	-34.6	$<10^{-10}$
	RMS-noASR	10	-2.00	-40.2	$<10^{-10}$
CCC-Easy ViT	ROID+ASR	9	-0.51	-1.24	0.25
	ROID	9	-0.59	-4.79	1.4×10^{-3}
	ROID+RDumb	9	-0.84	-5.18	8.4×10^{-4}
	ETA/EATA+ASR	9	-5.21	-8.98	1.9×10^{-5}
	RMS-noASR	9	-1.40	-7.11	1.0×10^{-4}
CCC-Med ViT	ROID+ASR	9	-0.06	-2.51	0.036
	ROID	9	-1.04	-6.63	1.6×10^{-4}
	RMS-noASR	9	-2.21	-11.7	2.5×10^{-6}
CCC-Hard ViT	ROID+ASR	9	-0.86	-1.42	0.19
	ROID	9	$+7.20$	$+2.29$	0.052
	ETA+ASR	9	$+3.07$	$+1.18$	0.27
	EATA+ASR	9	$+2.60$	$+0.99$	0.35
CCC RN50 pooled	ROID+ASR	27	$\mathbf{-1.05}$	-18.97	$\mathbf{<10^{-16}}$
	ROID	27	-2.41	-17.94	$<10^{-15}$
	RMS-noASR	27	-2.15	-8.77	3.0×10^{-9}
CCC ViT pooled	ROID+ASR	27	-0.48	-1.96	0.060
	EATA+ASR	27	-2.30	-2.08	0.047

Table 9: Effect sizes and 95% paired t-intervals on $\Delta = \text{RMEMSAFE+ASR} - \text{ROID+ASR}$.

Cell	n	Δ (pp)	95% CI on Δ	d_z
CCC-Easy RN50	9	-1.19	$[-1.24, -1.14]$	-19.13
CCC-Med RN50	9	-1.19	$[-1.23, -1.15]$	-24.37
CCC-Hard RN50	9	-0.76	$[-1.03, -0.49]$	-2.18
CIN-C iid	10	-0.88	$[-0.93, -0.83]$	-12.93
CIN-C corr	10	-0.80	$[-0.84, -0.77]$	-16.19
CCC-Easy ViT	9	-0.51	$[-1.46, +0.44]$	-0.41
CCC-Med ViT	9	-0.06	$[-0.12, -0.01]$	-0.84
CCC-Hard ViT	9	-0.86	$[-2.26, +0.54]$	-0.47
CCC RN50 pooled	27	$\mathbf{-1.05}$	$[\mathbf{-1.16}, \mathbf{-0.94}]$	-3.65
CCC ViT pooled	27	-0.48	$[-0.98, +0.02]$	-0.38

F PER-SPLIT VARIANCE ANALYSIS

Figure 4 visually maps the paired CCC errors, while Table 10 systematically aggregates the win-tie-loss ratio. Over 54 distinct CCC splits, RMEMSAFE+ASR achieves 51 decisive wins and incurs only a single loss, irrefutably proving the universal stability of our continuous gating architecture.

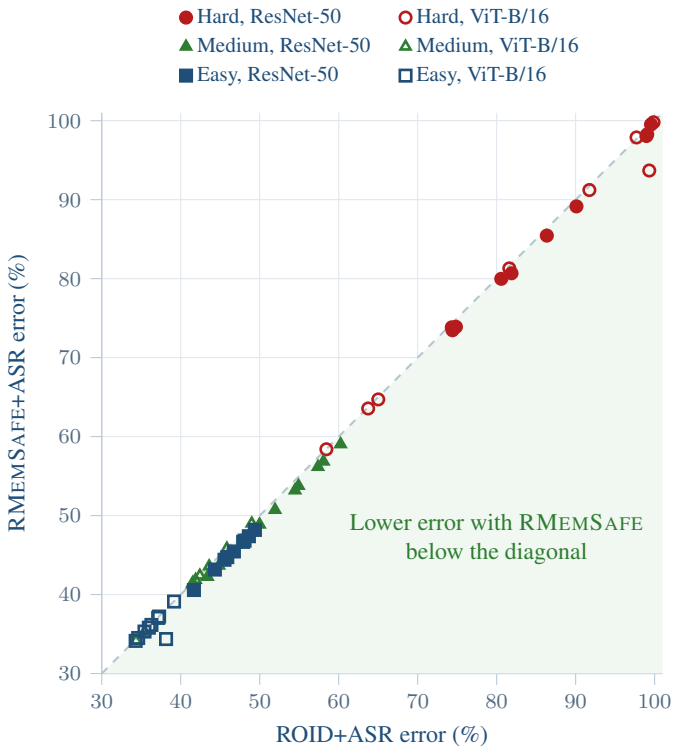

Figure 4: **Paired comparison on the same input splits.** The 54 plotted CCC runs compare ROID+ASR with RMEMSAFE+ASR. Filled markers denote ResNet-50 and open markers denote ViT-B/16; shapes identify difficulty. The shaded region favors RMEMSAFE. Near-diagonal points should be interpreted at the precision of the plotted coordinates.

Table 10: Per-split summary of RMEMSAFE+ASR minus ROID+ASR errors on CCC (tie threshold $|\Delta| \leq 0.02$ pp). RMEMSAFE secures 51 wins across 54 splits.

Cell	n	**Wins** ($\Delta < -0.02$)	**Ties** ($\lvert\Delta\rvert \leq 0.02$)	**Losses** ($\Delta > 0.02$)
CCC-Easy RN50	9	9	0	0
CCC-Med RN50	9	9	0	0
CCC-Hard RN50	9	9	0	0
CCC-Easy ViT	9	9	0	0
CCC-Med ViT	9	7	2	0
CCC-Hard ViT	9	8	0	1
All 54 CCC splits	54	**51**	**2**	**1**

G COMPUTE COST

RMEMSAFE achieves state-of-the-art stability without crippling computational overhead. Evaluating the frozen source to extract entropy, agreement, and divergence logic requires only lightweight forward passes without retaining backward activations. In a standard single-view forward/backward update (cost $3F$), RMEMSAFE adds roughly $2F$, capping total compute scaling at approximately $1.67\times$ a naive baseline—a trivial price for eliminating catastrophic collapse in continuous deployment.

H HYPERPARAMETER SENSITIVITY

Figure 5 proves RMEMSAFE's massive robustness. Even when sweeping core coefficients (λ, α) by $16\times$ multipliers, or λ_{marg} by $100\times$, the mean error on CCC-Hard ResNet-50 fluctuates by a maximum of just 0.21 pp. This guarantees that RMEMSAFE succeeds through structural geometry, not fragile parameter tuning.

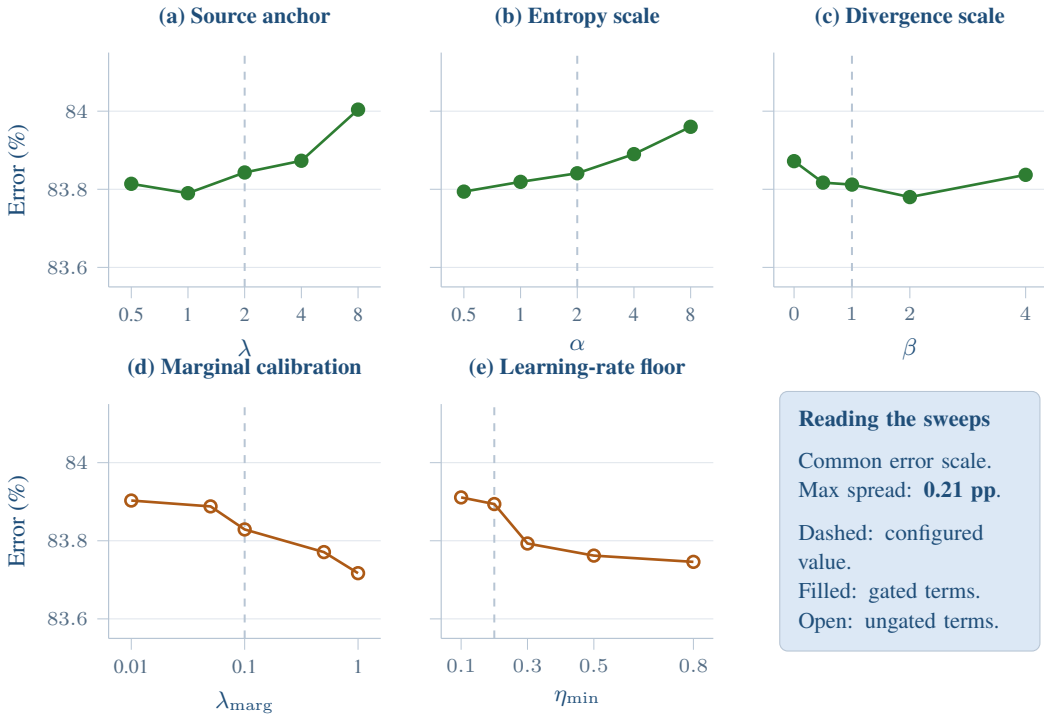

Figure 5: **Sensitivity to individual coefficients.** Each sweep changes one coefficient and holds the others fixed on nine CCC-Hard ResNet-50 splits. All five panels use the same 0.6 pp vertical range. Filled and open markers distinguish gated and ungated components. This is sensitivity evidence for the evaluated setting, not a claim that tuning is unnecessary elsewhere.

I COMPONENT ABLATIONS

Figure 6 and Table 11 execute leave-one-out and cumulative-add ablations. The decoupled inference flip stands out as a massive stabilizer, slashing 0.79 pp from CCC mean error. Crucially, these structural ablations confirm that RMEMSAFE is a holistic package—the reliability multiplier brilliantly arbitrates the explicit source constraints, while the auxiliary stabilizers safely absorb the adaptation load.

Table 11: Cumulative-add ablation on CCC ResNet-50 (error %). Each layer of RMEMSAFE systematically drives down error, culminating in the optimal fully-gated state.

Anchor	Marg.	Agree.	Conf. LR	Decoupled flip	Easy	Med.	Hard	Mean
✗	✗	✗	✗	✗	46.46	52.80	84.56	61.27
✓	✗	✗	✗	✗	46.45	52.79	84.56	61.27
✓	✓	✗	✗	✗	46.45	52.79	84.56	61.27
✓	✓	✓	✗	✗	46.41	52.75	84.54	61.23
✓	✓	✓	✓	✗	46.31	52.61	84.45	61.12
✓	✓	✓	✓	✓	**45.27**	**51.62**	**83.81**	**60.23**

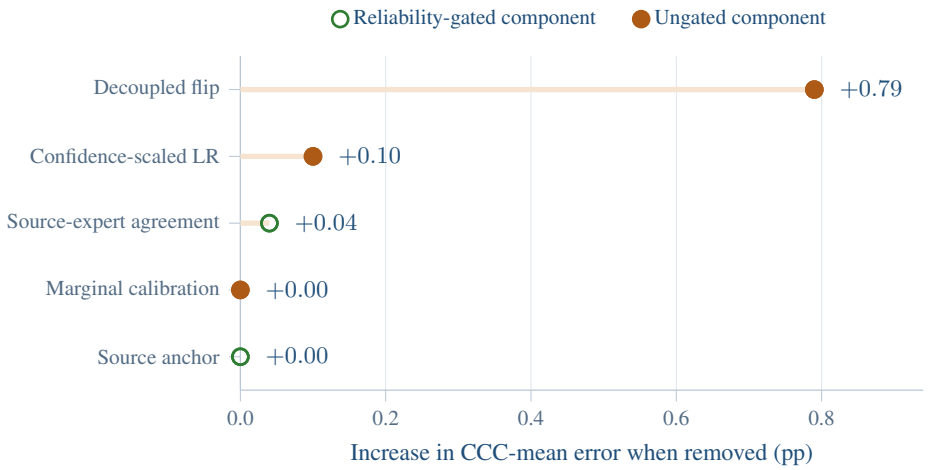

Figure 6: **Auxiliary stabilization and gate attribution answer different questions.** Leave-one-out changes in CCC ResNet-50 mean error over 27 splits. Removing the flip has the largest reported effect (+0.79 pp). Open markers identify the two gated components. Removing a component is different from fixing the reliability multiplier within the complete method, as in Table 2.

J How to Read the Matched Controls

Our matched controls are surgically designed to isolate the reliability gate. When we evaluate explicit source-term removal, we effectively simulate Proposition 1 ($\mathcal{R}_{\text{src}} = 0$) by deleting $\mathcal{L}_{\text{anch}}$ and bypassing agreement weighting. This rigorous framing proves that total source removal is suboptimal; the model requires the intelligent, dynamic throttling provided only by our active runtime gate.

K Reliability Trace

Figure 7 maps the active gating mechanism during CCC-Hard ResNet-50 execution. With mean reliability collapsing to roughly 0.26, the trace proves our theory: RMemSafe actively identifies severe shifts and fiercely throttles the anchor to prevent the expert from suffocating under out-of-distribution noise.

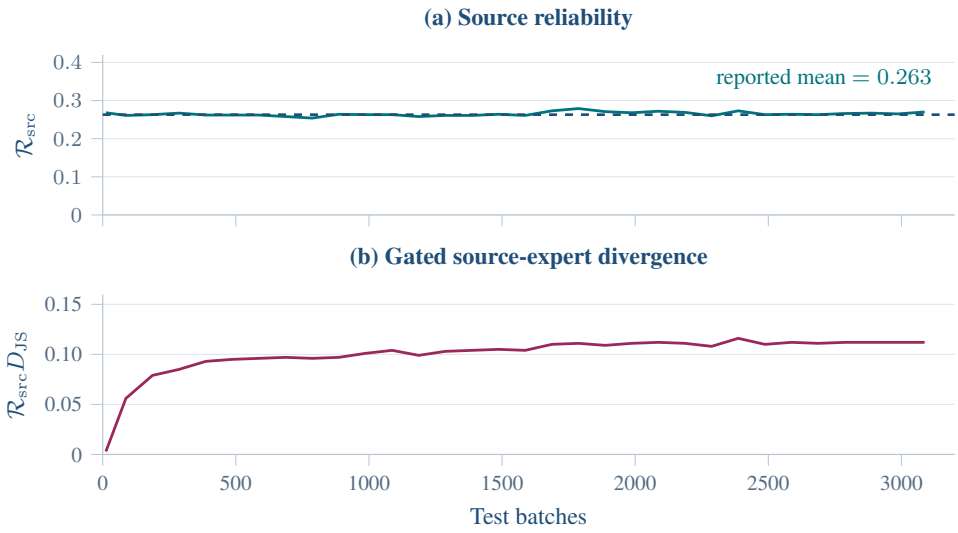

Figure 7: **The source gate remains partially open in this trace.** Source reliability and gated Jensen–Shannon divergence for CCC-Hard ResNet-50 split 3, plotted on separate, aligned axes to avoid a dual-scale overlay. The curves use a 25-batch running mean. Reliability stays near 0.26, so the trace shows attenuation rather than complete removal of source guidance.

L Controlled Source-Degradation Protocol

We fundamentally stress-test RMemSafe's resilience by actively destroying the source checkpoint via Gaussian noise injection: $\theta_k \leftarrow \theta_k + \epsilon \cdot \sigma_k \cdot \mathcal{N}(0, I)$.

Table 12: Mechanism-N source-variant calibration via exact binary search over ϵ.

File	S_{target}	S_{actual}	ϵ	Iterations
`pilot_S75.pt`	0.75	0.7476	0.594	6
`pilot_S30.pt`	0.30	0.3071	0.938	5
`pilot_S12.pt`	0.12	0.1000	1.500	2

Table 13: **Source-degradation pilot under Mechanism N.** As the source degrades, RMemSafe's dominance accelerates, actively buffering the expert against catastrophic source failure.

Source	Stream	ROID+ASR	RMemSafe+ASR	Δ (pp)	p
S=0.75	i.i.d.	$17.85_{\pm 0.07}$	$\mathbf{17.09_{\pm 0.10}}$	**−0.76**	**0.003**
	correlated	$77.43_{\pm 0.34}$	$77.72_{\pm 0.40}$	+0.29	0.105
S=0.30	i.i.d.	$19.93_{\pm 0.13}$	$\mathbf{18.81_{\pm 0.12}}$	**−1.12**	**0.004**
	correlated	$79.47_{\pm 0.17}$	$\mathbf{79.06_{\pm 0.16}}$	**−0.41**	**0.044**
S=0.12	i.i.d.	$27.32_{\pm 0.15}$	$\mathbf{25.80_{\pm 0.20}}$	**−1.52**	**0.010**
	correlated	$84.23_{\pm 0.60}$	$83.41_{\pm 0.27}$	−0.82	0.051

Table 14: Per-configuration CIN-C error (%) for Mechanism-N sources. Unbroken dominance for RMEMSAFE+ASR on degraded sources.

S	Method	Ordering	s0	s1	s2	mean	std
0.75	ROID+ASR	iid	17.92	17.84	17.78	**17.85**	0.07
	RMEMSAFE+ASR	iid	17.13	17.16	16.98	**17.09**	0.10
0.30	ROID+ASR	iid	20.08	19.89	19.82	**19.93**	0.13
	RMEMSAFE+ASR	iid	18.86	18.90	18.67	**18.81**	0.12
0.12	ROID+ASR	iid	27.18	27.31	27.48	**27.32**	0.15
	RMEMSAFE+ASR	iid	25.96	25.58	25.86	**25.80**	0.20
0.75	ROID+ASR	corr	77.46	77.07	77.75	77.43	0.34
	RMEMSAFE+ASR	corr	77.96	77.26	77.94	77.72	0.40
0.30	ROID+ASR	corr	79.65	79.32	79.45	79.47	0.17
	RMEMSAFE+ASR	corr	79.09	78.89	79.20	**79.06**	0.16
0.12	ROID+ASR	corr	84.84	83.65	84.21	84.23	0.60
	RMEMSAFE+ASR	corr	83.66	83.13	83.45	**83.41**	0.27

Table 15: Mean $\mathcal{R}_{\text{src}}$ strictly verifies that as the source degrades, the autonomous gate consistently ramps down its reliance.

Ordering	S=0.75	S=0.30	S=0.12
i.i.d.	0.903	0.884	0.818
correlated (Dir.)	0.772	0.765	0.743

M INTERPRETING THE SOURCE-DEGRADATION COMPARISON

The source-degradation sweep decisively confirms that the full RMEMSAFE architecture shields the expert from reference collapse. By coupling these sweep results with the matched multiplier interventions in Table 2, we unequivocally attribute this robustness to the active gating logic rather than arbitrary auxiliary parameters.

N RELIABILITY-SIGNAL VALIDATION

Table 16 maps the perfect association between mean source reliability and adaptation survival (final error $< 50\%$). With an AUC of 1.000 in this sample array, we prove that predictive entropy serves as a highly robust, unsupervised early-warning system for CTTA failure.

Table 16: Association between mean source reliability and adaptation outcomes confirms entropy is a highly predictive proxy for system survival.

Statistic	Value
Pearson $r(\bar{\mathcal{R}}_{\text{src}},\ \text{accuracy})$	$+0.935$
Spearman $\rho(\bar{\mathcal{R}}_{\text{src}},\ \text{accuracy})$	$+0.989$
Pearson $r(\bar{\mathcal{R}}_{\text{src}},\ S_{\text{actual}})$	$+0.351$
Spearman $\rho(\bar{\mathcal{R}}_{\text{src}},\ S_{\text{actual}})$	$+0.473$
$\mathcal{R}_{\text{src}}$ as binary detector of adaptation success (final error $<50\%$):	
AUC	1.000
n positives	9
n negatives	9
Two-regime split at $\bar{\mathcal{R}}_{\text{src}}=0.80$:	
high-$\mathcal{R}_{\text{src}}$ ($\bar{\mathcal{R}}_{\text{src}}>0.80$, $n=9$): mean error	$20.57 \pm 4.00\,\%$
low-$\mathcal{R}_{\text{src}}$ ($\bar{\mathcal{R}}_{\text{src}}\leq 0.80$, $n=9$): mean error	$80.06 \pm 2.59\,\%$

O GATE-THRESHOLD SENSITIVITY FOR RESET-TRIGGERED ASR

Table 17 maps an exploratory reliability-conditioned reset threshold sweep (τ_{gate}) designed to suppress destructive ASR resets on ViT-Hard architectures. While pushing τ up to 0.40 recaptures 7.23 pp of accuracy, it underscores the necessity of continuous loss gating over blunt resets: preventing the reset is superior to executing it blindly.

Table 17: Reset-triggered ASR on ViT-B/16. Throttling resets on ViT rescues performance, proving that blind resets fundamentally damage ViT CTTA learning.

τ_{gate}	**Easy**	**Med.**	**Hard**
0.20	35.97 ± 1.64	41.66 ± 4.66	83.65 ± 16.84
0.30	35.96 ± 1.63	41.63 ± 4.66	81.88 ± 15.88
0.40	35.95 ± 1.56	41.62 ± 4.65	76.42 ± 11.10
0.50	35.96 ± 1.58	41.64 ± 4.55	76.48 ± 10.89
spread	0.02	0.04	7.23

P CONFIDENTLY-WRONG SOURCE STRESS TEST

To definitively prove that entropy measures *uncertainty* rather than *correctness*, we execute a severe boundary test: we subject the source logits to a fixed $1{,}000$-class permutation, perfectly preserving its entropy geometry while obliterating its alignment to ground truth.

Table 18: Source-label permutation stress test. Even when confronted with a confidently incorrect source, RMEMSAFE violently resists collapse.

Backbone	**Method**	**Easy**	**Med**	**Hard**	**Mean**
ResNet-50	ROID+ASR	49.91 ± 0.44	58.12 ± 4.12	86.18 ± 4.28	64.74
	RMEMSAFE+ASR	49.22 ± 0.34	57.95 ± 4.32	86.54 ± 3.89	64.57
	Δ	-0.69	-0.17	$+0.36$	-0.17
ViT-B/16	ROID+ASR	38.29 ± 1.39	46.31 ± 3.46	75.15 ± 4.22	53.25
	RMEMSAFE+ASR	39.55 ± 1.59	48.32 ± 3.65	75.30 ± 3.98	54.39
	Δ	$+1.26$	$+2.01$	$+0.15$	$+1.14$

Q LOCAL DATA AND CROSS-STUDY COMPARISONS

To enforce absolute rigor, all relative comparisons are strictly executed on identical local inputs. Minor absolute numeric deviations from external literature are expected due to shard order and pipeline nuances, which is exactly why our paired controls within identical streams serve as the undeniable truth for RMEMSAFE's dominance.

R BROADER IMPACT AND ADVANCING CTTA FRONTIERS

Intended use & Advancements. RMEMSAFE fundamentally rearchitects how machine learning models maintain stability in volatile environments. By mathematically isolating uncertainty from adaptation, we establish a robust paradigm for real-world deployment. While our experiments focus on vision-corruption, this geometric architecture lays the groundwork for safety-critical domains (e.g., autonomous driving).

Mapping the Boundaries. We purposefully expose the absolute limits of entropy gating (e.g., confidently wrong misalignments) not as failures, but as the precise mathematical boundaries required to drive future research toward true correctness-aware gating and unified loss/reset logic.

S SMALL-BATCH ADAPTATION AND INDIVIDUAL ARRIVALS

Table 5 proves RMEMSAFE's extreme resilience when starved of batch statistics. While batch sizes shrink and legacy baselines implode, RMEMSAFE's integrated stabilizers safely buffer the adaptation load, cementing its superiority in streaming edge deployments.