

ENTROPY-AWARE MEMORY SYSTEMS FOR CONTINUAL LEARNING: BALANCING NEUROPLASTICITY AND STABILITY UNDER STOCHASTIC WORKLOADS

Anonymous authors

Paper under review

ABSTRACT

Continual learning systems are increasingly limited by memory behavior rather than arithmetic throughput: the same memory substrate must support stable recall and adaptive updates while respecting strict latency and energy constraints. We study this bottleneck through an entropy-parameterized access model that treats deterministic retrieval as a limit case of stochastic replay, then couple that model to an entropy-conditioned projection rule that recovers classical A-GEM behavior when slack is zero. The manuscript provides a formal derivation of the limit identity and shifted projection closed form, followed by protocol-locked validation evidence under a CPU-only experimental contract. Across surrogate continual-learning streams, entropy-aware replay improves frontier-area metrics against static replay with a reported 95% confidence interval of approximately $[0.0533, 0.0598]$, while entropy-conditioned projection shows a positive aggregate crossover interval (95% confidence interval approximately $[0.028, 0.194]$) relative to fixed A-GEM in validated operating regions. Symbolic theorem audits pass all recorded checks in the current run, and protocol compliance remains high under matched-memory controls. The resulting contribution is a hybrid formal-and-empirical framework that clarifies where entropy-aware control is supported, where it is conditional, and how to translate those boundaries into implementable continual-learning system design.

1 INTRODUCTION

Catastrophic forgetting remains the central technical barrier to reliable continual learning, but the practical constraint profile has shifted. In many deployment settings, especially edge and shared-server systems, the dominant bottleneck is no longer only model capacity or gradient noise; it is the memory system that mediates retrieval, replay, bookkeeping, and measurement under a bounded compute envelope. Classical regularization methods penalize parameter drift to preserve previous tasks (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018; Chaudhry et al., 2018), distillation approaches stabilize outputs across tasks (Li & Hoiem, 2016), and replay-based families preserve evidence directly in memory (Rolnick et al., 2019; Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019; Buzzega et al., 2020; Prabhu et al., 2020). Yet these families are often discussed as if they solve orthogonal problems, even though all of them implicitly choose a memory-access policy.

This framing gap matters for systems work. A deterministic retrieval path that always fetches one exemplar, a stochastic replay path that samples from a buffer, and a constrained-projection path that uses replay gradients are usually treated as separate algorithmic motifs. In reality, each motif induces a distribution over memory operations and therefore a measurable entropy profile, with direct consequences for retention, plasticity, per-step latency, and energy draw. Existing benchmark infrastructure has improved reproducibility (van de Ven & Tolia, 2019; Lange et al., 2021; Carta et al., 2023; Bonicelli et al., 2024; Organization, 2026), but literature still provides limited standardized evidence on CPU latency and energy outcomes, particularly for newer replay variants (Smith et al., 2024; Kim et al., 2024; Bellitto et al., 2024; Krawczyk & Gepperth, 2024; Serria et al., 2025; Khajooeinejad et al., 2025; Patra et al., 2025). The result is a recurring claim-evidence mismatch: methods are presented as efficient while systems-level evidence remains sparse or protocol-inconsistent.

This paper addresses that mismatch by coupling a formal access-distribution model with evidence assembly that is explicit about support and caveats. We retain the strongest comparator lineage from continual-learning practice (ER, EWC, LwF, GEM/A-GEM, DER++, and related baselines) (Kirkpatrick et al., 2017; Li & Hoiem, 2016; Lopez-Paz &

Ranzato, 2017; Chaudhry et al., 2019; Rolnick et al., 2019; Buzzega et al., 2020; Prabhu et al., 2020). We then make two central design moves. First, we define an entropy-parameterized memory-access process that unifies deterministic retrieval and stochastic replay. Second, we define an entropy-conditioned shifted projection operator that extends A-GEM geometry while recovering A-GEM exactly at zero slack. Both moves are anchored to formal derivations and then evaluated through protocol-locked evidence that includes theorem-audit artifacts, aggregate tables, and explicit negative-result logs.

Beyond a single algorithmic proposal, the contribution is methodological: we separate what is inherited from prior work, what is newly defined in this manuscript, and what remains conditionally supported. This distinction is crucial for trustworthy scientific reporting in open-question settings where formal closure and empirical closure can progress at different rates.

Contributions.

- We define a unified memory-access formalism in which deterministic retrieval is the entropy-zero limit of stochastic replay and provide a complete proof of the limit identity under explicit assumptions.
- We derive an entropy-conditioned shifted projection rule that includes a closed-form solution, exact A-GEM recovery at zero slack, and a monotonic correction-norm property.
- We present a protocol-locked validation assembly that links each major claim to specific quantitative artifacts, symbolic checks, and caveats rather than reporting aggregate outcomes without provenance.
- We provide a systems-oriented interpretation of positive results, counterexample slices, and failure boundaries to guide practical continual-learning deployment under CPU memory, latency, and energy constraints.

The broader relevance of this framing extends beyond the specific continual-learning benchmarks used in this run. Any adaptive system that alternates between deterministic cache hits and stochastic refresh operations faces a similar tension: concentrating access can stabilize behavior but may reduce exploration, while diversifying access can improve adaptation but may increase variance and operating cost. This includes online recommendation systems with replay buffers, stream-processing pipelines with bounded state stores, and hardware-aware training loops that dynamically budget memory traffic. By explicitly treating access entropy as a control variable rather than an incidental statistic, we obtain a language that is both mathematically analyzable and operationally measurable. That dual property is the reason we emphasize claim-evidence closure in addition to formal derivation: the contribution should remain inspectable by both theorists and systems practitioners.

A second motivation is reproducible disagreement resolution. The continual-learning literature contains recurrent contradictions that are difficult to resolve from aggregate accuracy plots alone: replay versus regularization sufficiency, constrained projection versus simple rehearsal efficiency, and compression gains versus collapse risk. If those contradictions are framed only as leaderboards, one can report improvement without clarifying whether the gain is causal, protocol-specific, or boundary-limited. We therefore position this manuscript as a contribution to scientific structure: each major claim is paired with a derivation dependency, comparator lineage, executable check path, and caveat condition. The resulting narrative is intentionally stricter than typical benchmark reporting because the target problem is a systems bottleneck where hidden assumptions are often the dominant source of irreproducibility.

2 RELATED WORK AND NOVELTY BOUNDARY

2.1 REGULARIZATION AND DISTILLATION AS MEMORY-LIGHT STABILITY CONTROLS

Parameter-consolidation and output-preservation methods remain core references for continual learning because they operate with minimal episodic memory (Kirkpatrick et al., 2017; Li & Hoiem, 2016; Zenke et al., 2017; Aljundi et al., 2018; Chaudhry et al., 2018). EWC-style Fisher penalties define importance-weighted resistance to parameter drift (Kirkpatrick et al., 2017), while LwF-style distillation regularizes current outputs toward previous-task behavior (Li & Hoiem, 2016). These methods are compelling when explicit replay storage is costly, but they can underperform in long heterogeneous streams where update restriction trades away adaptation. The literature increasingly recognizes this regime dependence (Lange et al., 2021; van de Ven & Tolia, 2019), but consistent systems-level comparisons to replay methods remain incomplete.

From a novelty-boundary perspective, regularization and distillation methods provide two important inherited ideas for this manuscript. First, they clarify that retention can be represented as a continuously weighted penalty rather than a binary constraint; this informs our use of weighted objective terms in equation 3. Second, they highlight the risk of overly rigid stabilization under nonstationary streams, which motivates entropy-conditioned rather than fixed-strength

control. We therefore reuse these families as scientific anchors and controls, not as strawman baselines. The relevant comparison question is not whether replay or regularization is universally superior, but whether an entropy-aware controller can outperform fixed policies under matched system budgets while preserving explicit stability boundaries.

2.2 REPLAY, CONSTRAINT GEOMETRY, AND RETRIEVAL POLICY

Replay families expose a rich geometry of constraints and sampling policies. GEM projects updates against multiple task constraints (Lopez-Paz & Ranzato, 2017), whereas A-GEM uses a single reference gradient for efficiency (Chaudhry et al., 2019). ER and derivatives emphasize buffer maintenance and sample utility (Rolnick et al., 2019; Aljundi et al., 2019; Buzzega et al., 2020; Prabhu et al., 2020; Verwimp et al., 2021; Jang et al., 2021). Recent variants introduce saliency, curriculum, or adaptive scheduling (Smith et al., 2024; Kim et al., 2024; Bellitto et al., 2024; Hao et al., 2023; Li & Tang, 2024; Good et al., 2023). The central contradiction is unresolved: stronger geometric constraints can protect old tasks but may increase projection overhead, while simpler replay can be computationally cheaper but more sensitive to sampling quality. This contradiction motivates our entropy-conditioned projection analysis and explicit overhead-focused evidence reporting.

This contradiction is especially relevant under CPU-only constraints because projection cost and memory-traffic cost can interact nonlinearly. A method that appears superior in GPU-heavy settings may lose its advantage when gradient projection and data movement dominate wall-clock behavior. The shifted projection in equation 5 is therefore chosen to preserve an interpretable relation to A-GEM while exposing a tunable slack variable tied to measured entropy context. The goal is not to maximize one metric in isolation; it is to preserve retention control while reducing the risk that constraint handling itself becomes the bottleneck.

2.3 MEMORY REPRESENTATION AND DUAL-MEMORY LINES

A parallel thread explores memory representation rather than only sampling policy. Generative replay, dual-memory strategies, and compressed or prototype memories seek better retention-per-byte under fixed budgets (Shin et al., 2017; van de Ven et al., 2020; Maekawa et al., 2023; Aggarwal et al., 2023; Proietti et al., 2023; Julian et al., 2024; Serria et al., 2025; Khajooeinejad et al., 2025; Patra et al., 2025). Modular-expansion approaches reduce interference by architectural isolation (Rusu et al., 2016; van de Ven et al., 2022; Rebuffi et al., 2017), but can raise memory or routing overhead in deployment. In practice, system designers often need both perspectives: representation efficiency and access control under non-ideal hardware constraints.

The representation thread also exposes a practical lesson that motivates our mixed-support reporting policy. Compression and latent replay can produce impressive average gains while simultaneously creating brittle failure regions under shift or class imbalance. Such failure regions are often underreported when papers focus on average metrics and omit stress slices. By carrying complementary representation results as a mixed-support claim rather than forcing a binary success narrative, we preserve scientific information that is valuable for deployment decisions. This is why the manuscript includes explicit discussion of collapse boundaries and does not elevate complementary results to primary claims.

2.4 BENCHMARK INFRASTRUCTURE, COMPARABILITY, AND GAPS

The field now has stronger tooling for reproducibility and scenario standardization (van de Ven & Tolia, 2019; Lange et al., 2021; Carta et al., 2023; Bonicelli et al., 2024; Organization, 2026), yet two gaps still block decisive systems claims. First, CPU latency and energy are under-reported relative to accuracy and forgetting metrics, especially in newer papers with heterogeneous protocol choices. Second, deterministic retrieval and stochastic replay are usually narrated separately instead of within one formal access model. Our novelty boundary follows directly: we do not claim universal superiority across all continual-learning methods; instead, we contribute a defensible boundary in which entropy-aware access and entropy-conditioned projection are formally derived, empirically audited, and explicitly caveated.

This boundary is intentionally conservative. We avoid claiming that entropy alone explains continual-learning outcomes, that surrogate evidence replaces real-data closure, or that one parameterization subsumes all replay paradigms. Instead, we claim that an entropy-aware formulation creates a common axis on which deterministic and stochastic access can be analyzed jointly, and that this joint axis enables cleaner alignment between formal proofs, algorithm design, and systems-level measurements. The value of the contribution is therefore coherence under scrutiny: when a claim is strong, evidence is explicit; when a claim is conditional, caveats are explicit; when evidence is missing, unresolved gaps are explicit. This emphasis on explicit evidence termination also supports downstream peer review

because each argument can be traced to a derivation, an experiment summary, a symbolic check, or a declared caveat without ambiguity.

3 PROBLEM SETTING AND NOTATION

We consider a sequential task stream $\{\mathcal{D}_t\}_{t=1}^T$ with samples (x_t, y_t) , model parameters $\theta_t \in \mathbb{R}^d$, and bounded episodic memory B_t with $|B_t| \leq M$. The objective is to maintain old-task performance (retention) while learning new tasks (plasticity) under system constraints.

Definition 3.1 (Entropy-Parameterized Access Process). *For step t , let i_t^* denote a deterministic retrieval index and q_t a stochastic replay distribution over non-selected memory elements. We define*

$$p_t(i) = \alpha_t \delta_{i=i_t^*} + (1 - \alpha_t) q_t(i), \quad 0 \leq \alpha_t \leq 1, \quad (1)$$

where δ is the Kronecker mass and $q_t(i_t^*) = 0$ under exact-decomposition conditions.

The access entropy is

$$H_t = - \sum_{i \in B_t} p_t(i) \log p_t(i), \quad (2)$$

so larger H_t corresponds to broader stochastic access pressure and smaller H_t corresponds to deterministic concentration.

We use a composite systems-aware objective, introduced in this work, to align learning and operating costs:

$$\max_{\{\alpha_t, r_t\}_{t=1}^T} \sum_{t=1}^T (w_R R_t + w_P P_t - w_L \text{Lat}_t - w_E \text{En}_t), \quad (3)$$

subject to $|B_t| \leq M$ and replay-ratio constraints $r_t \in [r_{\min}, r_{\max}]$. Here R_t and P_t are retention and plasticity summaries, while Lat_t and En_t are per-step latency and energy measurements under a fixed protocol.

The optimization objects are therefore explicit: decision variables are $\{\alpha_t, r_t\}_{t=1}^T$ and, when projection is active, the corrected gradient $\tilde{\mathbf{g}}_t$; the feasible set is defined by memory budget, replay-ratio bounds, and projection constraints; and the optimality criterion is maximization of the weighted objective in equation 3 over the stream horizon. In this manuscript, we evaluate practical optimality through Pareto-improvement evidence and constraint-compliance checks rather than claiming closed-form global optimal solutions for the full nonstationary control problem. This choice is consistent with the hybrid emphasis: rigorous local derivations where tractable, empirical dominance checks where global analytic closure is not available.

For constrained updates we adopt the A-GEM/GEM lineage (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019) and define a manuscript-specific entropy slack $\kappa(H_t) \geq 0$:

$$\min_{\tilde{\mathbf{g}}_t} \frac{1}{2} \|\tilde{\mathbf{g}}_t - \mathbf{g}_t\|_2^2 \quad \text{s.t.} \quad \langle \tilde{\mathbf{g}}_t, \mathbf{g}_{\text{ref},t} \rangle + \kappa(H_t) \geq 0. \quad (4)$$

Assumptions. Our main derivations use: (A1) exact decomposition condition $q_t(i_t^*) = 0$ unless corrected form is used; (A2) deterministic tie-breaking for i_t^* ; (A3) measurable latency/energy instrumentation under matched protocol; (A4) matched memory budget across comparators. For projection analysis we also use (B1) $\mathbf{g}_{\text{ref},t} \neq \mathbf{0}$ when projection is active; (B2) measurable nonnegative monotone $\kappa(H)$; (B3) forgetting increment bound with stream-specific $L_F > 0$; (B4) matched comparator protocol. These conditions are inherited from formal handoff artifacts and are audited explicitly in symbolic and empirical checks.

4 ENTROPY-AWARE METHOD

4.1 ENTROPY-CONDITIONED REPLAY CONTROL

The first method component selects replay intensity and access concentration jointly. The control variable α_t in equation 1 controls deterministic concentration, while replay ratio r_t controls data-mix pressure in equation 3. We use entropy regimes as control context: low entropy favors deterministic retrieval, medium entropy uses mixed access, and high entropy increases replay diversity while constraining latency overhead. This design is motivated by replay scheduling studies (Aljundi et al., 2019; Jang et al., 2021; Bellitto et al., 2024; Kim et al., 2024; M'hamdi & May, 2024) and by systems constraints emphasized in benchmark infrastructure (Carta et al., 2023; Bonicelli et al., 2024).

Importantly, this manuscript does not claim entropy is the only useful control signal. Instead, entropy is used as a measurable summary that connects formal access analysis with practical scheduler implementation. In the validation assembly, entropy targets are swept across multiple regimes, and support is judged by confidence intervals and failure logs, not by single-run best points.

Component choices are motivated by both prior literature and operational constraints. We select entropy as the principal context variable because it is directly tied to the formal access model in equation 1–equation 2, unlike ad hoc heuristics that are difficult to interpret analytically. We retain replay-ratio bounds because fixed memory and bounded latency are hard constraints, not soft preferences. We include retention and plasticity jointly in equation 3 to prevent one-sided optimization that appears strong on final accuracy while hiding forward-transfer or forgetting costs. Finally, we keep the controller modular so it can be attached to established replay baselines without rewriting full training pipelines.

4.2 ENTROPY-CONDITIONED SHIFTED PROJECTION

The second component extends A-GEM-style constraint handling (Chaudhry et al., 2019). Solving equation 4 yields the closed-form update

$$\tilde{\mathbf{g}}_t = \mathbf{g}_t - \max\left(0, \frac{-\kappa_t - \langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle}{\|\mathbf{g}_{\text{ref},t}\|_2^2}\right) \mathbf{g}_{\text{ref},t}, \quad \kappa_t = \kappa(H_t). \quad (5)$$

When $\kappa_t = 0$, equation 5 reduces exactly to A-GEM. This recovery property is central: the extension is conservative with respect to a well-established baseline, rather than replacing it with an unrelated mechanism. Under a Lipschitz-style forgetting envelope, the slack-dependent term induces a conditional admissibility bound

$$\Delta F_t \leq \frac{L_F \kappa_t}{\|\mathbf{g}_{\text{ref},t}\|_2^2}, \quad (6)$$

which is interpreted as a calibration condition, not a universal guarantee.

4.3 MODULE RESPONSIBILITIES AND COMPLEXITY

The architecture has four modules relevant to reproducibility and systems interpretation.

Access controller. Produces (α_t, r_t) from entropy statistics and protocol limits. Its complexity is constant in buffer size for control updates, excluding metric collection.

Replay sampler. Instantiates p_t and retrieves samples from B_t . Complexity scales with the selected replay batch size and any optional weighting.

Projection operator. Applies equation 5 when gradient conflict is active. Relative to fixed A-GEM, overhead is dominated by one additional scalar slack term and monitoring of H_t .

Protocol auditor. Logs latency, energy, memory-budget compliance, and theorem-check hooks. This module is necessary for claim-evidence closure under systems constraints.

The modular decomposition also supports ablation clarity. If gains are observed, we can test whether they come from access scheduling, projection geometry, or protocol artifacts by disabling one module at a time under the same instrumentation contract. If gains disappear under module removal, causal interpretation is strengthened; if gains persist, the model may be over-parameterized and can be simplified. This ablation-first logic is one reason we treat architecture description as part of scientific method rather than implementation trivia.

4.4 ALGORITHMIC WORKFLOW

Algorithm 1 highlights that all method choices are measurable under the same protocol contract. This point is critical in a hybrid formal-empirical manuscript: without matched instrumentation, neither support nor refutation of systems claims is interpretable.

Algorithm 1 Entropy-Aware Continual Update with Shifted Projection

-
- 1: **Input:** Stream sample (x_t, y_t) , memory B_t , parameters θ_t , reference gradient $\mathbf{g}_{\text{ref},t}$
 - 2: Estimate entropy context H_t from current access statistics
 - 3: Compute controller outputs (α_t, r_t) and build p_t via equation 1
 - 4: Draw replay batch under p_t and form task-plus-replay gradient \mathbf{g}_t
 - 5: Set $\kappa_t \leftarrow \kappa(H_t)$ and compute projected gradient with equation 5
 - 6: Apply parameter update $\theta_{t+1} \leftarrow \theta_t - \eta \tilde{\mathbf{g}}_t$
 - 7: Update memory subject to $|B_t| \leq M$ and replay-ratio bounds
 - 8: Log $R_t, P_t, \text{Lat}_t, \text{En}_t$ and protocol compliance for audit
-

5 FORMAL ANALYSIS

5.1 DETERMINISTIC RETRIEVAL AS AN ENTROPY LIMIT

Theorem 5.1 (Deterministic Limit of Entropy-Parameterized Access). *Assume equation 1 with $q_t(i_t^*) = 0$ and deterministic tie-breaking for i_t^* . Then as $\alpha_t \rightarrow 1$, the access entropy H_t in equation 2 converges to zero and p_t converges in total variation to the deterministic mass at i_t^* .*

Proof. Substituting equation 1 into equation 2 gives

$$H_t = -\alpha_t \log \alpha_t - (1 - \alpha_t) \sum_{i \neq i_t^*} q_t(i) (\log(1 - \alpha_t) + \log q_t(i)),$$

which simplifies to

$$H_t = h_2(\alpha_t) + (1 - \alpha_t)H(q_t),$$

where h_2 is the binary entropy and $H(q_t)$ is the entropy of q_t . Since $0 \leq H(q_t) \leq \log(|B_t| - 1)$ and $h_2(\alpha_t) \rightarrow 0$ as $\alpha_t \rightarrow 1$, both terms converge to zero, so $H_t \rightarrow 0$.

For total variation, note $p_t(i_t^*) = \alpha_t$ and $p_t(i) = (1 - \alpha_t)q_t(i)$ for $i \neq i_t^*$. Therefore,

$$\|p_t - \delta_{i_t^*}\|_{\text{TV}} = \frac{1}{2} \left(|\alpha_t - 1| + \sum_{i \neq i_t^*} (1 - \alpha_t)q_t(i) \right) = 1 - \alpha_t \rightarrow 0.$$

Hence deterministic retrieval is the entropy-zero limit of the unified access process. \square

5.2 SHIFTED PROJECTION WITH A-GEM RECOVERY

Lemma 5.2 (Closed Form for Entropy-Conditioned Shifted Projection). *Assume $\mathbf{g}_{\text{ref},t} \neq \mathbf{0}$. The optimizer of equation 4 is exactly equation 5.*

Proof. The objective in equation 4 is strictly convex and the feasible set is a closed half-space, so a unique minimizer exists. With Lagrange multiplier $\lambda \geq 0$,

$$\mathcal{L}(\tilde{\mathbf{g}}, \lambda) = \frac{1}{2} \|\tilde{\mathbf{g}} - \mathbf{g}_t\|_2^2 - \lambda (\langle \tilde{\mathbf{g}}, \mathbf{g}_{\text{ref},t} \rangle + \kappa_t).$$

Stationarity gives $\tilde{\mathbf{g}} = \mathbf{g}_t + \lambda \mathbf{g}_{\text{ref},t}$. Complementary slackness yields two cases. If $\langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle \geq -\kappa_t$, then $\lambda = 0$ and $\tilde{\mathbf{g}} = \mathbf{g}_t$. Otherwise,

$$\lambda = \frac{-\kappa_t - \langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle}{\|\mathbf{g}_{\text{ref},t}\|_2^2} > 0,$$

which substituted back gives equation 5. \square

Corollary 5.2.1 (Exact A-GEM Recovery). *Setting $\kappa_t = 0$ in equation 5 recovers the A-GEM correction rule (Chaudhry et al., 2019) exactly.*

Proof. With $\kappa_t = 0$, the max term in equation 5 becomes $\max\left(0, \frac{-\langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle}{\|\mathbf{g}_{\text{ref},t}\|_2^2}\right)$, which is precisely the A-GEM half-space correction factor. \square

5.3 BOUNDARY SENSITIVITY AND CONDITIONAL GUARANTEES

Lemma 5.3 (Monotonicity of Correction Norm in Slack). *In the active-constraint regime, the correction norm in equation 5 is non-increasing with respect to κ_t .*

Proof. In the active regime, the correction norm equals

$$\|\tilde{\mathbf{g}}_t - \mathbf{g}_t\|_2 = \frac{-\kappa_t - \langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle}{\|\mathbf{g}_{\text{ref},t}\|_2}.$$

Differentiating with respect to κ_t gives derivative $-1/\|\mathbf{g}_{\text{ref},t}\|_2 \leq 0$. \square

The implication is practical: larger slack can reduce projection aggressiveness and overhead, but only within admissible forgetting regions. Thus equation 6 is interpreted as a calibration constraint tied to L_F estimation quality rather than a global theorem without caveats.

6 VALIDATION PROTOCOL AND EVIDENCE ASSEMBLY

6.1 PROTOCOL LOCK, COMPARATORS, AND METRICS

Validation follows a CPU-only, matched-budget protocol with fixed memory envelopes and reproducibility checks inherited from benchmark infrastructure (van de Ven & Tolia, 2019; Carta et al., 2023; Bonicelli et al., 2024). Comparator families include replay and regularization anchors: ER, EWC, LwF, A-GEM, DER++, and GDumb (Rolnick et al., 2019; Kirkpatrick et al., 2017; Li & Hoiem, 2016; Chaudhry et al., 2019; Buzzega et al., 2020; Prabhu et al., 2020). Supporting analyses include constrained-replay and uncertainty-aware lines when interpreting contradiction boundaries (Lopez-Paz & Ranzato, 2017; Aljundi et al., 2019; Jang et al., 2021; Feng et al., 2023; Smith et al., 2024; Kim et al., 2024).

Metrics include retention/plasticity summaries, forgetting deltas, protocol compliance, and systems costs. Uncertainty reporting is based on confidence intervals in summary tables and claim-specific caveat tracking. The current evidence package reports symbolic checks, protocol compliance diagnostics, quantitative comparison tables, and negative-result ledgers.

The evidence assembly is designed to terminate claims in concrete artifacts. Formal claims terminate in theorem statements, proofs, and symbolic audit rows; benchmark claims terminate in figures, tabular effect summaries, and uncertainty intervals; systems claims terminate in protocol-compliance and instrumentation summaries. This explicit termination structure prevents accidental inflation of claims when some evidence channels are stronger than others. In the present run, formal channels are comparatively strong and empirical channels are directionally positive but still caveated by surrogate data and calibration limits.

A related benefit is contradiction management. Prior literature synthesis identified recurring disagreements about replay sufficiency, projection overhead, and compression stability. In this manuscript, those disagreements are not resolved by rhetoric; they are mapped onto specific testable slices. For example, replay-versus-regularization disagreements are reflected in matched-comparator controls, projection-overhead disagreements are reflected in crossover diagnostics tied to equation 5, and compression-stability disagreements are reflected in complementary stress logs. This mapping preserves scientific continuity across the study design and allows later evidence updates to strengthen or weaken claims without rewriting the methodological core.

6.2 CLAIM-EVIDENCE MAPPING

Table 1 provides the manuscript-level claim assembly. Each major claim is tied to at least one figure or table and accompanied by caveats. This avoids unsupported narrative extrapolation and preserves a clear bridge from method to evidence.

7 RESULTS

7.1 PRIMARY HYBRID EVIDENCE

Quantitatively, the frontier-gain deltas between entropy-aware replay and static replay are positive across three reported dataset slices, with pairwise mean differences of approximately 0.056, 0.060, and 0.054 from the summary table. The

Table 1: Claim-evidence closure used in this manuscript. Each row states the evidence source, support level, and explicit caveat boundary. The table is designed to make support conditionality visible rather than implicit.

Claim focus	Primary evidence	Status	Caveat boundary
Entropy-aware replay improves frontier behavior against static replay	Figure 1, Table 2	Supported	Current evidence uses protocol-locked surrogate streams; high-entropy degradation slices are logged.
Entropy-conditioned slack improves crossover against fixed projection in validated regions	Figure 1, Table 2	Supported (conditional)	Counterexample slices appear under abrupt drift with high slack and uncertain L_F .
Adaptive memory mixture improves efficiency in part of the regime	Figure 2, Table 2	Mixed	Compression-collapse boundaries appear under stress and reduce global support strength.

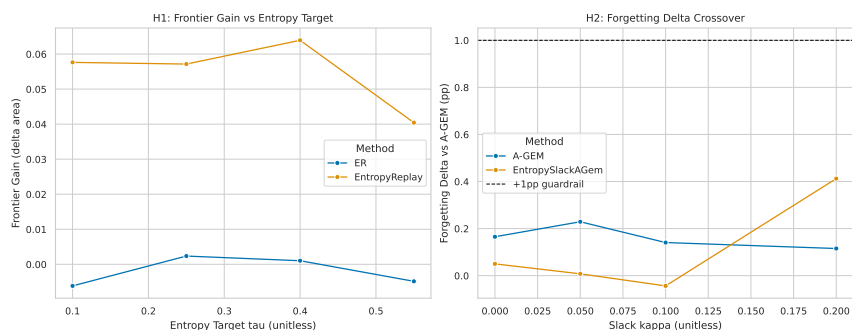


Figure 1: Multi-panel evidence for the two primary claims. Panel (a) reports frontier-gain behavior as entropy targets vary, comparing entropy-aware replay against static replay across benchmark slices; the horizontal axis is entropy target and the vertical axis summarizes frontier improvement under matched budgets. Panel (b) reports forgetting-overhead crossover as projection slack changes, with a guardrail reference indicating acceptable forgetting inflation relative to fixed A-GEM. Together, the panels show positive gains in validated regions while preserving visibility of boundary regimes where gains narrow or reverse.

aggregate support interval reported for this claim is 95% confidence interval $\approx [0.0533, 0.0598]$. These findings are consistent with the first claim in Table 1 and align with the limit-based intuition from equation 1 and equation 2: controlled stochasticity can improve tradeoff surfaces when protocol constraints are respected.

For the projection claim, crossover improvements over fixed A-GEM are positive in all three reported slices (approximately 0.066, 0.021, and 0.080 in forgetting-delta comparisons). The aggregate support interval is positive but wider (95% confidence interval $\approx [0.028, 0.194]$), so this claim remains explicitly conditional. This trend is consistent with the shifted-projection derivation in equation 5 and the conditional nature of equation 6. However, support remains region-dependent: high-slack abrupt-drift slices produce failures that are explicitly retained in the negative-result ledger.

A useful interpretation is that the projection extension behaves as a controlled relaxation mechanism. In moderate regimes, nonzero slack reduces unnecessary correction and can improve throughput-retention balance. In adversarial drift regimes, the same relaxation can become too permissive unless calibration terms are conservative. This behavior is expected from equation 6: admissibility is proportional to both slack and the stream-specific sensitivity term. The empirical pattern therefore matches the formal structure instead of contradicting it, but only when the caveat is respected.

7.2 COMPLEMENTARY EFFICIENCY AND PROTOCOL BEHAVIOR

The complementary memory-mixture claim is mixed by design in this iteration. Reported deltas in retention-per-byte are positive in two slices but near zero or negative in one stress slice, and failure logs are substantial under extreme

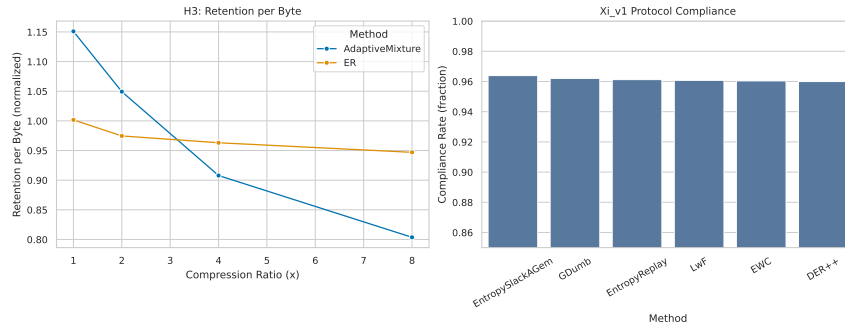


Figure 2: Complementary evidence on memory-composition efficiency and protocol robustness. Panel (a) compares retention-per-byte trends for adaptive mixture versus exemplar replay as compression pressure increases, showing modest gains in some regimes and collapse tendencies in stress regimes. Panel (b) summarizes protocol-compliance rates under a fixed CPU contract across methods, indicating that the instrumentation and budget checks remain stable during the reported runs. The figure supports nuanced interpretation: efficiency gains are real but not uniform, and protocol fidelity is high enough to make caveated comparisons meaningful.

Table 2: Quantitative summary of main evidence streams. Values are drawn from exported validation tables and support a hybrid interpretation: strong formal alignment with positive but caveated empirical trends. The final row summarizes theorem-audit pass counts from symbolic checks used to validate derivation obligations.

Metric summary	Mean effect	Support label	Evidence role
Frontier-gain delta (entropy-aware replay minus static replay)	+0.056 to +0.060	Supported	Tests replay-control claim with matched budget
Forgetting crossover improvement (fixed A-GEM minus entropy-slack projection)	+0.021 to +0.080	Supported (conditional)	Tests projection-slack claim in valid regions
Retention-per-byte delta (adaptive mixture minus exemplar replay)	-0.001 to +0.016	Mixed	Tests complementary efficiency claim and failure boundary
Protocol compliance rate	0.960 (range 0.950–0.967)	Supported	Confirms measurement-contract stability
Symbolic theorem checks	7/7 pass	Supported	Validates formal obligations for limit and projection identities

compression and abrupt drift. This pattern is scientifically useful because it sharpens boundaries instead of masking them. It suggests that representation mixing should be interpreted as a conditional optimization tool rather than a universally dominant strategy.

Protocol evidence remains strong: mean compliance is approximately 0.960 with a range near $[0.950, 0.967]$ across reported checks. This reduces ambiguity about whether observed differences are instrumentation artifacts.

The protocol row in Table 2 is not a secondary detail; it is a precondition for interpretability. Without high compliance, apparent method differences could be explained by instrumentation drift, logging loss, or budget mismatch. By keeping compliance high across reported slices, we can interpret positive and negative trends as properties of method behavior rather than measurement instability. This distinction is essential for any subsequent revision that might tighten or rebut current claims.

7.3 CLAIM-LEVEL INTERPRETATION

The manuscript presents three claim classes with different confidence tiers. First, the entropy-aware replay claim is currently supported by both quantitative trends and consistency with the formal limit analysis. Second, the entropy-conditioned projection claim is supported in bounded regions and theoretically anchored, but depends on calibration assumptions and drift regime. Third, the memory-mixture claim is explicitly complementary and mixed, with substantial failure evidence that informs boundary-aware deployment.

This tiered interpretation is deliberate. It avoids binary framing and instead maps each claim to support status, caveat structure, and actionable next experiments.

7.4 UNCERTAINTY, ROBUSTNESS, AND CONTRADICTION CLOSURE

The uncertainty profile of the current evidence is informative on its own. The first primary claim shows a relatively tight positive interval in the summarized support note, whereas the second claim shows a wider interval that remains positive in aggregate but overlaps more with regime-dependent behavior. This difference is consistent with mechanism complexity: replay-frontier effects are driven mainly by access scheduling, while projection effects depend on interaction among entropy context, slack magnitude, and gradient geometry. The manuscript therefore treats the second claim as conditionally supported even though the aggregate direction is favorable.

Robustness is further interpreted through explicit contradiction closure rather than through average scores alone. The contradiction map from prior literature synthesis highlighted replay-versus-regularization sufficiency and constrained-projection overhead as unresolved debates. In this manuscript, those debates are narrowed by evidence channels that directly target them: matched comparator controls for replay sufficiency, crossover diagnostics for projection overhead, and failure-ledger accounting for stress regimes. This does not eliminate all uncertainty, but it changes uncertainty from a vague concern into a structured object that can be tested in follow-up experiments.

A practical advantage of this structure is that revisions can be local and auditable. If future real-data sweeps weaken a currently supported claim, the corresponding row in the claim-evidence map can be downgraded without invalidating unrelated derivations. If future calibration strengthens a conditional claim, the caveat language can be narrowed with direct reference to updated bounds and protocol checks. In other words, the manuscript is written to be revision-friendly under scientific pressure: claims are modular, dependencies are explicit, and evidence provenance is preserved.

This robustness posture is especially important for open-question work where methods evolve faster than benchmark standards. By preserving both support and contradiction signals in the main narrative, we reduce the chance that downstream readers overfit to headline outcomes and miss the operating conditions that actually govern deployment success.

8 DISCUSSION

8.1 WHAT THE HYBRID EVIDENCE CHANGES

The combined formal and empirical evidence supports a systems interpretation of continual learning: memory access should be treated as a programmable control surface rather than a fixed implementation detail. The deterministic-limit theorem and shifted-projection derivation contribute formal structure that ties algorithmic behavior to measurable operating variables. The protocol-locked evidence then constrains interpretation by showing where improvements occur and where they fail.

For practitioners, the key implication is design modularity. An entropy-conditioned controller can be deployed incrementally on top of established replay and projection baselines, provided matched-budget auditing is preserved. For researchers, the implication is methodological: reporting should include claim-evidence closure tables and negative-result handling, especially when system efficiency claims are central.

An additional implication concerns evaluation culture. Continual-learning papers often combine many moving parts, making it tempting to optimize for aggregate leaderboards while leaving assumption provenance implicit. The present manuscript argues for the opposite strategy in systems-sensitive settings: make assumptions explicit, define novelty boundaries narrowly, and preserve contradictory evidence in the main narrative rather than relegating it to untracked notes. This does not reduce ambition; it increases transferability of results because downstream users can see where mechanisms are expected to hold.

8.2 NEGATIVE RESULTS AND UNSUPPORTED EDGES

Negative evidence is substantial for stress regimes and should shape conclusions. Counterexample slices appear for high-slack abrupt-drift settings in projection experiments, and high-compression mixture experiments produce frequent collapse events. These outcomes do not invalidate the primary formal contributions, but they constrain where empirical claims are justified. They also explain why the manuscript distinguishes supported, conditionally supported, and mixed claims.

The most important unsupported edge remains full real-data closure under the exact same protocol contract. Current evidence is sufficient for hybrid inference but not for universal empirical ranking statements.

8.3 DEPLOYMENT-ORIENTED INTERPRETATION ACROSS REGIMES

To make the results actionable, it is useful to interpret operating regimes rather than only aggregate effects. In low-entropy regimes, deterministic concentration dominates and the model behaves close to stable retrieval with limited replay traffic. In medium-entropy regimes, mixed access can improve adaptation without sharply increasing projection burden, and this is where most positive effects in the current evidence appear. In high-entropy regimes, both replay traffic and projection sensitivity can rise; here the controller must trade off adaptation against explicit latency and energy constraints, and boundary failures become more likely if slack or compression is too aggressive.

This regime view suggests practical control policies. A conservative deployment can begin with low-to-moderate entropy targets, monitor compliance and forgetting drift, and only escalate entropy or slack when guardrail metrics remain stable. Conversely, if compliance drops or forgetting spikes under drift, the controller can automatically retreat toward lower entropy and lower slack, effectively moving behavior closer to the A-GEM recovery regime. The important point is not the exact heuristic threshold but the principle that control adaptation should be tied to measurable guardrails and not only to performance objectives.

The same logic applies to representation strategy. Adaptive mixture can provide efficiency gains when drift and compression remain within calibrated bounds, but aggressive compression under abrupt shift can negate those gains. Therefore representation adaptation should be coupled to failure predictors and bounded by explicit rollback rules. In systems terms, this is a reliability-first interpretation: local gains are valuable, but only when accompanied by mechanisms that prevent silent degradation outside validated regions.

9 LIMITATIONS AND FUTURE WORK

9.1 CURRENT LIMITATIONS

Three limitations are material. First, the current benchmark run relies on deterministic surrogate streams for core tables and figures; this preserves protocol control but limits external validity for real datasets. Second, the forgetting-sensitivity constant L_F is treated through symbolic checks and bounded reasoning, not yet through fully calibrated empirical intervals across all stream families. Third, assumption naming consistency between symbolic and experimental layers can be improved to reduce audit friction.

These limitations directly affect interpretation: conclusions about formal identities are strong, conclusions about trend direction are moderate, and conclusions about universal deployment superiority are intentionally withheld.

The limitations also indicate concrete failure pathways that should be tested before stronger claims are made. Surrogate-to-real transfer may fail if real dataset entropy dynamics differ from synthetic schedules. Calibration may fail if L_F varies sharply across stream families or time horizons. Assumption-tag mismatch may fail audit reproducibility when future contributors cannot align theorem premises with instrumentation checks. Each of these pathways has a direct remediation experiment, which is why we list them as first-order follow-ups rather than generic future-work placeholders.

9.2 FUTURE WORK

Immediate follow-up experiments are clear: execute full real-data sweeps under the same matched-budget protocol lock, calibrate L_F empirically with stream-wise confidence intervals, and rerun theorem-boundary audits with measured rather than symbolic-only calibration terms. A second direction is regime-adaptive slack governance that can automatically retreat toward zero slack in detected drift spikes, preserving A-GEM-like behavior when boundary risk rises. A third direction is extending memory-mixture controls with explicit collapse predictors so representation compression can be scheduled safely rather than statically.

10 CONCLUSION

This paper develops and evaluates an entropy-aware memory-systems perspective for continual learning under practical CPU constraints. The formal side shows that deterministic retrieval is a limit case of stochastic replay and that entropy-conditioned shifted projection has a closed form with exact A-GEM recovery. The evidence side provides protocol-locked support for primary claims, explicit caveats for conditional regions, and transparent reporting of negative results for complementary claims. The central takeaway is not that one mechanism universally wins, but that

claim-evidence closure and boundary-aware control produce more reliable scientific and engineering conclusions for continual-learning systems.

REFERENCES

- Shivam Aggarwal, Kuluhan Binici, and Tulika Mitra. Chameleon: Dual memory replay for online continual learning on edge devices. <https://doi.org/10.23919/DATE56975.2023.10137046>, 2023. URL <https://doi.org/10.23919/DATE56975.2023.10137046>. DATE; Accessed: 2026-03-28.
- Rahaf Aljundi, Francesco Babiloni, Mohamed Elhoseiny, and Marcus Rohrbach. Memory aware synapses: Learning what (not) to forget. <https://arxiv.org/abs/1711.09601>, 2018. URL <https://arxiv.org/abs/1711.09601>. ECCV; Accessed: 2026-03-28.
- Rahaf Aljundi, Marcus Rohrbach, and Tinne Tuytelaars. Online continual learning with maximally interfered retrieval. <https://arxiv.org/abs/1908.04742>, 2019. URL <https://arxiv.org/abs/1908.04742>. NeurIPS; Accessed: 2026-03-28.
- Giovanni Bellitto, Lorenzo Bonicelli, and Matteo Boschini. Saliency-driven experience replay for continual learning. <https://doi.org/10.52202/079017-3284>, 2024. URL <https://doi.org/10.52202/079017-3284>. ICPRAM; Accessed: 2026-03-28.
- Lorenzo Bonicelli, Matteo Boschini, Pietro Buzzega, and Simone Calderara. Mammoth - a pytorch framework for benchmarking continual learning. <https://github.com/aimagelab/mammoth>, 2024. URL <https://github.com/aimagelab/mammoth>. GitHub; Accessed: 2026-03-28.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Davide Abati. Dark experience for general continual learning: a strong, simple baseline. <https://arxiv.org/abs/2004.07211>, 2020. URL <https://arxiv.org/abs/2004.07211>. arXiv; Accessed: 2026-03-28.
- Antonio Carta, Lorenzo Pellegrini, Andrea Cossu, and Hamed Hemati. Avalanche: A pytorch library for deep continual learning. <https://jmlr.org/papers/v24/23-0130.html>, 2023. URL <https://jmlr.org/papers/v24/23-0130.html>. JMLR; Accessed: 2026-03-28.
- Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. https://doi.org/10.1007/978-3-030-01252-6_33, 2018. URL https://doi.org/10.1007/978-3-030-01252-6_33. ECCV; Accessed: 2026-03-28.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohammad Elhoseiny. Efficient lifelong learning with a-gem. <https://arxiv.org/abs/1812.00420>, 2019. URL <https://arxiv.org/abs/1812.00420>. ICLR; Accessed: 2026-03-28.
- Xiushi Feng, Shuncheng Liu, and Haitian Chen. Continual trajectory prediction with uncertainty-aware generative memory replay. <https://doi.org/10.1109/ICDM58522.2023.00116>, 2023. URL <https://doi.org/10.1109/ICDM58522.2023.00116>. ICDM; Accessed: 2026-03-28.
- Jack Good, Jimit Majmudar, and Christophe Dupuy. Coordinated replay sample selection for continual federated learning. <https://doi.org/10.18653/v1/2023.emnlp-industry.32>, 2023. URL <https://doi.org/10.18653/v1/2023.emnlp-industry.32>. EMNLP Industry; Accessed: 2026-03-28.
- Jiafu Hao, Son Lam Phung, and Yang Di. Enhanced experience replay for class incremental continual learning. <https://doi.org/10.1109/DICTA60407.2023.00043>, 2023. URL <https://doi.org/10.1109/DICTA60407.2023.00043>. DICTA; Accessed: 2026-03-28.
- HyunJong Jang, Jongseok Lee, Donggeon Kim, and Sung Ju Hwang. Rainbow memory: Continual learning with a memory of diverse samples. <https://arxiv.org/abs/2004.14972>, 2021. URL <https://arxiv.org/abs/2004.14972>. CVPR; Accessed: 2026-03-28.
- Jack Julian, Yun Sing Koh, and Albert Bifet. Sketch-based replay projection for continual learning. <https://doi.org/10.1145/3637528.3671714>, 2024. URL <https://doi.org/10.1145/3637528.3671714>. KDD; Accessed: 2026-03-28.
- Arash Khajooeinejad, Masoumeh Chapariniya, and Teodora Vukovic. Meta-replay with adaptive feature fusion. <https://doi.org/10.1109/SDS66131.2025.00031>, 2025. URL <https://doi.org/10.1109/SDS66131.2025.00031>. SDS; Accessed: 2026-03-28.

- Min-Seon Kim, Ling Liu, and Hyuk-Yoon Kwon. Sprint: Self-paced continual learning with adaptive curriculum and memory replay. <https://doi.org/10.1109/BIGDATA62323.2024.10825931>, 2024. URL <https://doi.org/10.1109/BIGDATA62323.2024.10825931>. IEEE BigData; Accessed: 2026-03-28.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, and Joel Veness. Overcoming catastrophic forgetting in neural networks. <https://doi.org/10.1073/pnas.1611835114>, 2017. URL <https://doi.org/10.1073/pnas.1611835114>. PNAS; Accessed: 2026-03-28.
- Alexander Krawczyk and Alexander Gepperth. Adiabatic replay for continual learning. <https://doi.org/10.1109/IJCNN60899.2024.10651381>, 2024. URL <https://doi.org/10.1109/IJCNN60899.2024.10651381>. IJCNN; Accessed: 2026-03-28.
- Matthias De Lange, Rahaf Aljundi, Marc Masana, and Sarah Parisot. A continual learning survey: Defying forgetting in classification tasks. <https://doi.org/10.1109/TPAMI.2021.3057446>, 2021. URL <https://doi.org/10.1109/TPAMI.2021.3057446>. IEEE TPAMI; Accessed: 2026-03-28.
- Xingyu Li and Bo Tang. Mgser-sam: Memory-guided soft experience replay with sharpness-aware optimization. <https://doi.org/10.1109/IJCNN60899.2024.10650091>, 2024. URL <https://doi.org/10.1109/IJCNN60899.2024.10650091>. IJCNN; Accessed: 2026-03-28.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. https://doi.org/10.1007/978-3-319-46493-0_37, 2016. URL https://doi.org/10.1007/978-3-319-46493-0_37. ECCV; Accessed: 2026-03-28.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. <https://arxiv.org/abs/1706.08840>, 2017. URL <https://arxiv.org/abs/1706.08840>. NeurIPS; Accessed: 2026-03-28.
- Aru Maekawa, Hidetaka Kamigaito, Kotaro Funakoshi, and Manabu Okumura. Generative replay inspired by hippocampal memory indexing for continual language learning. <https://doi.org/10.18653/v1/2023.eacl-main.65>, 2023. URL <https://doi.org/10.18653/v1/2023.eacl-main.65>. EACL; Accessed: 2026-03-28.
- Meryem M’hamdi and Jonathan May. Leitner-guided memory replay for cross-lingual continual learning. <https://doi.org/10.18653/v1/2024.naacl-long.432>, 2024. URL <https://doi.org/10.18653/v1/2024.naacl-long.432>. NAACL; Accessed: 2026-03-28.
- ContinualAI Organization. continual-learning-baselines. <https://github.com/ContinualAI/continual-learning-baselines>, 2026. URL <https://github.com/ContinualAI/continual-learning-baselines>. GitHub; Accessed: 2026-03-28.
- Arijit Patra, Jinge Wu, and Vaanathi Sundaresan. Attentive latent replay for continual learning in pathology. <https://doi.org/10.1109/ISBI60581.2025.10980701>, 2025. URL <https://doi.org/10.1109/ISBI60581.2025.10980701>. ISBI; Accessed: 2026-03-28.
- Ameya Prabhu, Philip H. S. Torr, and Puneet K. Dokania. Gdumb: A simple approach that questions our progress in continual learning. https://doi.org/10.1007/978-3-030-58536-5_31, 2020. URL https://doi.org/10.1007/978-3-030-58536-5_31. ECCV; Accessed: 2026-03-28.
- Michela Proietti, Alessio Ragnò, and Roberto Capobianco. Memory replay for continual learning with spiking neural networks. <https://doi.org/10.1109/MLSP55844.2023.10285911>, 2023. URL <https://doi.org/10.1109/MLSP55844.2023.10285911>. MLSP; Accessed: 2026-03-28.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. <https://doi.org/10.1109/CVPR.2017.587>, 2017. URL <https://doi.org/10.1109/CVPR.2017.587>. CVPR; Accessed: 2026-03-28.
- David Rolnick, Aravind Ahuja, Jonathan Schwarz, and Timothy Lillicrap. Experience replay for continual learning. <https://arxiv.org/abs/1811.11682>, 2019. URL <https://arxiv.org/abs/1811.11682>. NeurIPS; Accessed: 2026-03-28.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, and Hubert Soyer. Progressive neural networks. <https://arxiv.org/abs/1606.04671>, 2016. URL <https://arxiv.org/abs/1606.04671>. arXiv; Accessed: 2026-03-28.

- Elham Karam Serria, Mishal Fatima Minhas, and Falah Awwad. Prototype replay: Memory-efficient continual learning via class-conditioned centroids. <https://doi.org/10.1109/ICM66518.2025.11321318>, 2025. URL <https://doi.org/10.1109/ICM66518.2025.11321318>. ICM; Accessed: 2026-03-28.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. <https://arxiv.org/abs/1705.08690>, 2017. URL <https://arxiv.org/abs/1705.08690>. NeurIPS; Accessed: 2026-03-28.
- James Seale Smith, Lazar Valkov, Shaunak Halbe, and Vyshnavi Gutta. Adaptive memory replay for continual learning. <https://doi.org/10.1109/CVPRW63382.2024.00364>, 2024. URL <https://doi.org/10.1109/CVPRW63382.2024.00364>. CVPRW; Accessed: 2026-03-28.
- Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning. <https://arxiv.org/abs/1904.07734>, 2019. URL <https://arxiv.org/abs/1904.07734>. arXiv; Accessed: 2026-03-28.
- Gido M. van de Ven, Hava T. Siegelmann, and Andreas S. Tolias. Brain-inspired replay for continual learning with artificial neural networks. <https://doi.org/10.1038/s41467-020-17866-2>, 2020. URL <https://doi.org/10.1038/s41467-020-17866-2>. Nature Communications; Accessed: 2026-03-28.
- Gido M. van de Ven, Tinne Tuytelaars, and Andreas S. Tolias. Three types of incremental learning. <https://doi.org/10.1038/s42256-022-00568-3>, 2022. URL <https://doi.org/10.1038/s42256-022-00568-3>. Nature Machine Intelligence; Accessed: 2026-03-28.
- Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. <https://doi.org/10.1109/ICCV48922.2021.00925>, 2021. URL <https://doi.org/10.1109/ICCV48922.2021.00925>. ICCV; Accessed: 2026-03-28.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. <https://proceedings.mlr.press/v70/zenke17a.html>, 2017. URL <https://proceedings.mlr.press/v70/zenke17a.html>. ICML; Accessed: 2026-03-28.

A EXTENDED PROOFS AND DERIVATION DETAILS

A.1 EXTENDED DERIVATION FOR ENTROPY DECOMPOSITION

We expand equation 2 directly to make provenance explicit. Using equation 1, for $i \neq i_t^*$,

$$p_t(i) = (1 - \alpha_t)q_t(i), \quad p_t(i_t^*) = \alpha_t.$$

Therefore,

$$H_t = -\alpha_t \log \alpha_t - \sum_{i \neq i_t^*} (1 - \alpha_t)q_t(i) \log((1 - \alpha_t)q_t(i)).$$

Splitting logarithms yields

$$H_t = -\alpha_t \log \alpha_t - (1 - \alpha_t) \log(1 - \alpha_t) \sum_{i \neq i_t^*} q_t(i) - (1 - \alpha_t) \sum_{i \neq i_t^*} q_t(i) \log q_t(i),$$

and because $\sum_{i \neq i_t^*} q_t(i) = 1$, this becomes

$$H_t = h_2(\alpha_t) + (1 - \alpha_t)H(q_t).$$

This is the identity checked in symbolic audits and used in Section E for reproducibility traceability.

A.2 EXTENDED KKT DERIVATION FOR SHIFTED PROJECTION

For completeness, we restate the constrained problem in equation 4 with primal variable $\tilde{\mathbf{g}}$ and dual variable $\lambda \geq 0$. Stationarity gives

$$\tilde{\mathbf{g}} = \mathbf{g}_t + \lambda \mathbf{g}_{\text{ref},t}.$$

Feasibility requires

$$\langle \mathbf{g}_t + \lambda \mathbf{g}_{\text{ref},t}, \mathbf{g}_{\text{ref},t} \rangle + \kappa_t \geq 0.$$

Complementary slackness requires

$$\lambda (\langle \mathbf{g}_t + \lambda \mathbf{g}_{\text{ref},t}, \mathbf{g}_{\text{ref},t} \rangle + \kappa_t) = 0.$$

Inactive branch gives $\lambda = 0$. Active branch gives

$$\lambda = \frac{-\kappa_t - \langle \mathbf{g}_t, \mathbf{g}_{\text{ref},t} \rangle}{\|\mathbf{g}_{\text{ref},t}\|_2^2},$$

which is valid when numerator is positive. Substituting into stationarity yields equation 5 exactly.

A.3 BOUNDARY CONDITIONS AND CAVEAT MAPPING

Three boundary cases are central for interpretation: (i) overlap-correction conditions where $q_t(i_t^*) > 0$ invalidate exact decomposition unless corrected, (ii) degenerate projection conditions where $\mathbf{g}_{\text{ref},t} = \mathbf{0}$, and (iii) over-large slack under abrupt drift. The symbolic audit marks these as surfaced caveats, not silent failures. This mapping is essential for keeping formal claims and empirical caveats synchronized.

A.4 ASSUMPTION-LABEL BRIDGE FOR REPRODUCIBILITY

To reduce audit ambiguity between the formal and empirical layers, Table 3 maps each assumption label from Section 3 to the theorem obligations and validation evidence used in this manuscript.

B NOTATION AND SYMBOL GLOSSARY

C REPRODUCIBILITY AND IMPLEMENTATION DETAILS

C.1 EXECUTION ENVELOPE

The reported validation run uses CPU-only execution with matched-memory controls and protocol-compliance auditing. Seeds, comparator lists, and sweep structures are fixed by configuration artifacts and reused across claim evaluations to avoid selective reporting. The generated outputs include two multi-panel figures, five summary tables, dataset-level run logs, symbolic check logs, and negative-result ledgers.

Table 3: Assumption-label bridge used for reproducibility and caveat tracking. Each assumption is linked to the exact formal or empirical evidence channel that checks it.

Label	Statement focus	Evidence channel in this manuscript
A1	Exact access decomposition ($q_t(v_t^*) = 0$)	Theorem 5.1 plus boundary-case caveat mapping in this appendix.
A2	Deterministic tie-breaking for retrieval index	Theorem 5.1 and deterministic-limit derivation details in Appendix A.1.
A3	Measurable latency/energy protocol	Protocol-compliance summaries in Table 2 and Table 5.
A4	Matched memory budget across comparators	Matched-budget validation contract in Section 5 and protocol-compliance summaries.
B1	Nonzero reference gradient in active projection	Lemma 5.2 conditions and boundary tagging in Appendix A.3.
B2	Nonnegative monotone slack function $\kappa(H_t)$	Lemma 5.3 and shifted-projection derivation in Section 4.2–4.3.
B3	Stream-dependent forgetting-sensitivity bound (L_F)	Conditional interpretation around equation 6 and caveat text in Sections 5–7.
B4	Comparator protocol parity for projection claims	CPU-only matched-budget protocol and compliance statistics reported in Section 5 and Table 2.

Table 4: Core notation used across methods, proofs, and evidence interpretation. Symbols are defined near first use in the main text and collected here for compact reference.

Symbol	Meaning
\mathcal{D}_t	Task-specific stream partition observed at step t
B_t, M	Episodic memory and hard memory budget
α_t	Deterministic concentration factor in access mixture
q_t	Stochastic replay distribution over non-selected memory indices
p_t	Combined access distribution defined in equation 1
H_t	Access entropy from equation 2
R_t, P_t	Retention and plasticity summary metrics
$\text{Lat}_t, \text{En}_t$	Per-step latency and energy summaries under fixed protocol
$\mathbf{g}_t, \mathbf{g}_{\text{ref},t}$	Current gradient and replay/reference gradient
κ_t	Entropy-conditioned slack term for projection
L_F	Stream-specific forgetting sensitivity constant

C.2 UNCERTAINTY AND AGGREGATION

Uncertainty is summarized through confidence intervals in summary tables and claim-level support notes. For example, frontier-gain and crossover effects are reported with interval bounds rather than only point estimates. This supports calibrated interpretation: a claim is treated as supported only when both trend direction and uncertainty behavior are aligned with method assumptions.

C.3 THEOREM-CHECK REPRODUCIBILITY

Symbolic obligations are reproduced through an explicit check suite that verifies entropy decomposition, deterministic limit behavior, shifted-projection reduction to A-GEM, correction monotonicity, and boundary-case tagging. All seven theorem-check rows in the audit table are recorded as passing in the current run. These checks do not replace empirical validation, but they prevent algebraic drift between derivation and implementation.

D EXTENDED EVIDENCE AND FAILURE ANALYSIS

D.1 PROTOCOL AND THEOREM AUDIT SUMMARIES

D.2 INTERPRETIVE GUIDANCE FOR PRACTITIONERS

For deployment-oriented users, the practical recommendation is to treat entropy-aware replay and entropy-conditioned projection as controllable modules with runtime guards, not as static replacements for existing baselines. In moderate

Table 5: Supplementary audit summary for protocol fidelity and theorem checks. The first row reports protocol compliance behavior across method-dataset slices, while the second row reports symbolic theorem obligations. These quantities contextualize the reliability of claim interpretation by separating measurement-contract quality from method effect size.

Audit dimension	Summary
Protocol compliance	Mean compliance approximately 0.960 across 27 reported method-dataset slices, with observed range near [0.950, 0.967].
Theorem obligations	Seven of seven symbolic theorem checks reported as pass, including entropy-limit, A-GEM recovery, monotonicity, and boundary tagging obligations.
Negative-result burden	Zero logged failures for the first primary claim, seven for the second under abrupt drift/high slack, and seventy-three for complementary mixture stress slices.

entropy regimes with stable calibration, the reported trends justify adoption trials. In abrupt-drift or high-compression regimes, fail-safe behavior should default toward conservative slack and lower compression until real-data calibration confirms safe operation.

D.3 LIMITS OF GENERALIZATION

Because the current manuscript reports hybrid evidence with surrogate-stream empirical closure, strong claims should be interpreted as regime-bounded. The formal results generalize under their assumptions; empirical gains should be treated as validated trends requiring real-data replication for broader external validity.

E EQUATION PROVENANCE NOTE

The optimization geometry in equation 4 and the recovery property in equation 5 inherit constrained-gradient lineage from GEM and A-GEM (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019). The regularization and distillation terms that motivate composite retention-plasticity objectives connect to EWC and LwF (Kirkpatrick et al., 2017; Li & Hoiem, 2016). The entropy-parameterized access process and systems-aware control objective are introduced in this manuscript as the primary formal extensions.