EVIDENCE-GUIDED CONSTRAINT RELIABILITY FOR PARTIALLY RELEVANT VIDEO RETRIEVAL

Anonymous authors
Paper under double-blind review

ABSTRACT

Partially relevant video retrieval learns local query–video correspondence from video-level supervision. Unpaired relations induce uncertain training constraints because pairing labels leave local compatibility unresolved. Two challenges complicate constraint calibration: early optimization relies on immature relation evidence, and later responses conflate relation difficulty with positive-like local compatibility. These observations suggest that constraint calibration should evolve with relation maturity, using retriever-independent semantic evidence early and difficulty-conditioned task evidence as relations become informative. We propose Gradual Reliability-Aware Constraint Evolution (GRACE) to instantiate this principle. GRACE coordinates semantic prior evidence and evolving task evidence within a unified contrastive objective. Semantic Evidence Allocation Recalibration (SEAR) reuses each candidate video's paired queries to recalibrate candidate allocation through set-level semantic support. Positive-reference Attribution for Intervention Refinement (PAIR) compares paired and unpaired relations at similar difficulty to obtain relation-specific evidence. It then uses evidence above the conditional reference mean to refine constraint strength continuously. Together, the two mechanisms adjust candidate allocation and effective training pressure as task relations develop. Experiments on Charades-STA, ActivityNet Captions, and TVR demonstrate consistent improvements across retrieval settings and support the complementary roles of semantic prior and task-derived evidence.

1 INTRODUCTION

Partially relevant video retrieval (PRVR) retrieves untrimmed videos that contain local content relevant to a natural-language query (Dong et al., 2022). A video may contain several scenes, actions, and events, while a query describes only a subset of its content. PRVR therefore requires local query–video correspondence, even when training annotations provide only video-level pairs. Models must develop this correspondence during optimization, making relation learning under coarse supervision central to effective retrieval.

Existing methods improve local correspondence through multiple-instance aggregation, multiscale matching, and knowledge transfer (Dong et al., 2022; Wang et al., 2024b; Dong et al., 2023). These approaches combine clip- and frame-level responses, model temporal structure, and transfer pretrained vision–language knowledge into retrieval models. Ambiguity-aware relation learning further uses uncertainty and similarity to adjust supervision for potentially compatible pairs (Cho et al., 2025). These advances motivate a complementary question: how should the available relation evidence determine the strength and scope of constraint adjustment throughout training?

Two difficulties arise when training assigns constraints to unpaired candidates. **Evidence–Intervention Mismatch** occurs when immature local responses immediately determine candidate contributions and feed their early preferences back into optimization; semantic priors can instead provide a reference independent of the evolving retrieval model. **Difficulty–Reliability Entanglement** persists as task relations mature: a high response may indicate either a difficult comparison or positive-like local compatibility, while video-level pairing does not resolve the latter for unpaired candidates. Constraint adjustment therefore requires evidence that is both independent when relations are immature and difficulty-conditioned when task evidence becomes informative. Figure 1 summarizes these uncertain relations and complementary evidence sources.

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.

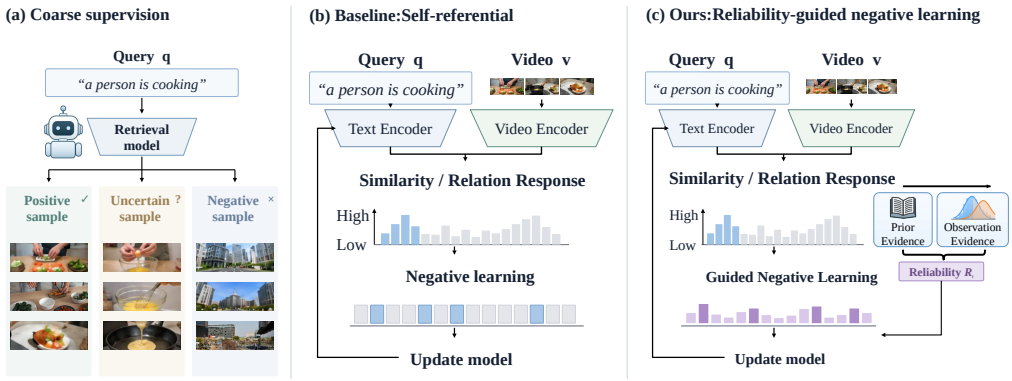

Figure 1: **Motivation of Evidence-Guided Constraint Reliability.** **(a)** Coarse supervision yields heterogeneous relations, where nominally unpaired videos may correspond to uncertain or reliable negatives. **(b)** Conventional learning directly feeds model-induced responses back into optimization. **(c)** GRACE progressively coordinates semantic and task-derived evidence to assess constraint reliability and calibrate unpaired relations.

To address these issues, we propose **Gradual Reliability-Aware Constraint Evolution (GRACE)**, which coordinates evidence sources with the constraint adjustments they can support. GRACE decomposes each candidate's contrastive gradient into a total coefficient and a normalized allocation, separating where training pressure is assigned from how much is applied. **Semantic Evidence Allocation Recalibration (SEAR)** uses the queries paired with each candidate video as semantic prior evidence to recalibrate allocation through set-level support, while its influence decreases as task statistics mature. **Positive-reference Attribution for Intervention Refinement (PAIR)** compares paired and unpaired local responses at similar difficulty and converts evidence above the conditional reference mean into continuous relation-specific multipliers. Together, the two mechanisms evolve constraint calibration from semantic prior evidence toward task-derived evidence while leaving the inference-time retrieval model unchanged.

Our contributions are as follows:

- **A formulation of evolving constraint reliability.** We identify Evidence–Intervention Mismatch and Difficulty–Reliability Entanglement as complementary challenges in learning from unpaired relations under video-level supervision.
- **A unified framework for constraint calibration.** SEAR reallocates candidate contributions using semantic prior evidence, and PAIR refines their effective weights using difficulty-conditioned positive-reference evidence. Their gradient decomposition separates changes in allocation from changes in total training pressure.
- **Evaluation across three PRVR benchmarks.** Experiments on Charades-STA, ActivityNet Captions, and TVR, together with component ablations, support the effectiveness of GRACE across datasets and retrieval settings.

2 Related Work

2.1 Partially Relevant Video Retrieval

PRVR learns query–moment correspondence from video-level supervision by strengthening local relation modeling and handling ambiguous associations. MS-SL, GMMFormer, and ProtoPRVR improve local matching through multiscale aggregation, temporal interaction, and semantic prototypes (Dong et al., 2022; Wang et al., 2024b; Moon et al., 2025a). GMMFormerV2 and ARL introduce uncertainty-aware matching and relational supervision (Wang et al., 2024a; Cho et al., 2025), whereas RAL and Holmes use probabilistic modeling or hierarchical evidence (Zhang et al., 2025; Li et al., 2026a). Collectively, these methods improve how local evidence is represented and

supervised, but leave open how evidence reliability should determine the scope of constraint adjustment.

2.2 Relation Uncertainty in Contrastive Learning

Contrastive learning under incomplete supervision must separate hard but valid negatives from potentially compatible pairs. Hard-negative sampling and MoCHi emphasize informative comparisons through difficulty control or representation mixing (Robinson et al., 2021; Kalantidis et al., 2020), whereas Debiased Contrastive Learning and False Negative Cancellation reduce false-negative contamination through objective correction or candidate treatment (Chuang et al., 2020; Huynh et al., 2022). Difficulty-Based Sampling further distinguishes difficulty from false-negative status using relative difficulty (Jang & Wang, 2023). From a statistical perspective, non-negative PU risk estimation and kernel-based mixture proportion estimation address learning with incomplete labels and reference-component estimation (Kiryo et al., 2017; Ramaswamy et al., 2016), while their assumptions caution against equating reference attribution with true relevance. These studies motivate difficulty-conditioned, continuous calibration rather than binary reassignment of uncertain relations.

2.3 Auxiliary Evidence for Weakly Supervised Learning

Auxiliary evidence can complement coarse pairing labels. CLIP provides transferable cross-modal semantics (Radford et al., 2021); Cap4Video and DL-DKD inject generated-caption or distillation knowledge (Wu et al., 2023; Dong et al., 2023), while semantic-collapse mitigation preserves pretrained text correlations during PRVR training (Moon et al., 2025b). Model-derived supervision instead evolves during optimization: Mean Teacher and ALBEF use weight-averaged or momentum teachers (Tarvainen & Valpola, 2017; Li et al., 2021), and uncertainty-guided self-training transfers model confidence to temporal grounding under video-level supervision (Huang et al., 2023). Together, these works show that prior semantic structure and evolving task evidence can both supplement weak supervision; our focus is how reliability should govern distinct forms of constraint adjustment under weak supervision.

3 Method

3.1 Problem Formulation

For a query q and an untrimmed video v, the baseline model scores local correspondence over temporal units $\mathcal{M}(v) = \{m_1, \dots, m_T\}$:

$$r_t = s_\theta(q, m_t), \qquad z(q, v) = \mathrm{Agg}(r_1, \dots, r_T). \tag{1}$$

Here, s_θ denotes the local response function, and Agg retains the baseline aggregation rule. We use a to index a query, z_a^+ for its paired logit, and $\mathcal{N}_a$ for its unpaired candidate indices. The logit z_{aj} scores candidate j; any temperature scaling is implicit in the logits. Video-level annotations leave these candidates' local compatibility unresolved. The baseline noise-contrastive estimation (NCE) objective assigns these candidates denominator terms:

$$\ell_a^{(0)} = -\log \frac{e^{z_a^+}}{e^{z_a^+} + M_a}, \qquad M_a = \sum_{j \in \mathcal{N}_a} e^{z_{aj}}. \tag{2}$$

Define the normalized candidate allocation $p_{aj} = e^{z_{aj}}/M_a$ and total coefficient $\beta_a = M_a/(e^{z_a^+} + M_a)$. Then

$$\frac{\partial \ell_a^{(0)}}{\partial z_{aj}} = \beta_a p_{aj}. \tag{3}$$

p_a distributes gradient mass across candidates, while β_a sets the total magnitude. This decomposition exposes two intervention axes. Replacing p_a redistributes a fixed query-level gradient budget, whereas changing β_a changes the total negative pressure applied to the query. GRACE keeps these effects explicit so that each evidence source only supports the intervention it can justify. This factorization separates candidate-wise redistribution from query-level modulation, allowing each evidence source to operate only on the component supported by its reliability signal.

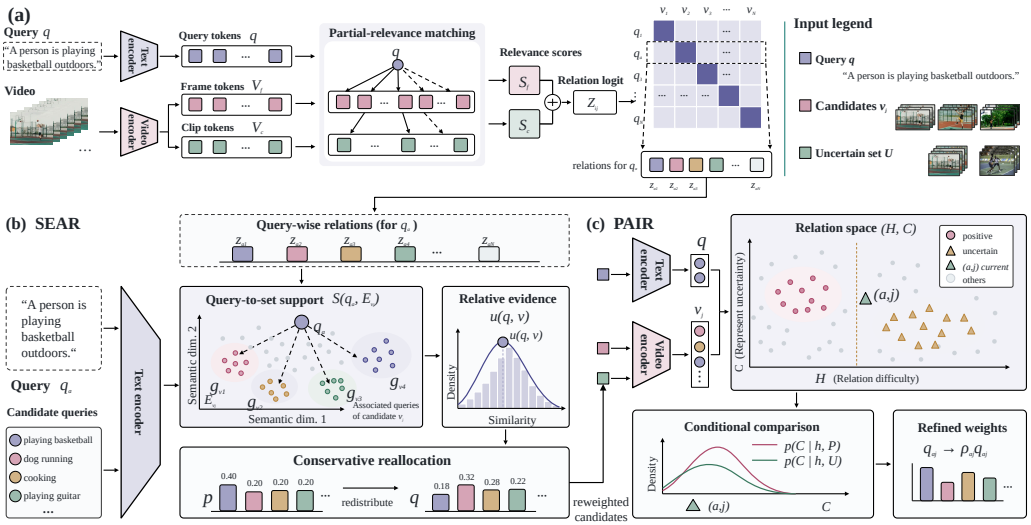

Figure 2: **Overview of the proposed framework.** (a) The baseline performs partial-relevance matching to obtain relation logits and the batch relation matrix $Z = [z_{ij}]$, from which a query-wise relation profile is extracted for SEAR. (b) SEAR uses set-level semantic support and relative evidence to conservatively reallocate candidate contributions. (c) The reweighted candidates are further processed by PAIR, which compares paired and unpaired reference relations in the (H, C) space and refines their effective weights.

3.2 Overview of GRACE

GRACE calibrates the allocation and effective weight of unpaired candidates through complementary evidence sources. Figure 2 summarizes the framework. The baseline produces a batch relation matrix $Z = [z_{ij}]$; its row for query a provides the relation profile for subsequent calibration (Figure 2(a)). SEAR extracts semantic prior evidence from video-associated query sets and transforms p_a into a normalized allocation q_a. PAIR uses evolving local responses to identify additional positive-reference evidence among relations with similar difficulty. It converts this evidence into multipliers ρ_{aj} that refine the effective weights $\rho_{aj} q_{aj}$. A single weighted contrastive objective integrates both adjustments during training. Both modules act on the original relation logits and return training coefficients; neither introduces an alternative inference branch. The separation therefore attributes changes to constraint calibration rather than to an additional scorer.

3.3 Semantic Evidence Allocation Recalibration

SEAR uses semantic prior evidence to recalibrate candidate allocation while task relations are still developing (Figure 2(b)). A video can have complementary or overlapping query descriptions. Set-level support therefore accounts for both query matching and correlations among descriptions.

Set-level semantic support. Let $\mathcal{D}_{\text{tr}}$ contain the training query–video pairs. For each candidate video v, collect its paired queries into

$$E_v = \{q' \mid (q', v) \in \mathcal{D}_{\text{tr}}\} = \{e_{v1}, \ldots, e_{vm_v}\}. \tag{4}$$

SEAR reuses training annotations to construct this set. A frozen text encoder f_{sem} maps the current query and each description to unit vectors:

$$h_q = \frac{f_{\text{sem}}(q)}{\|f_{\text{sem}}(q)\|_2}, \qquad g_{vr} = \frac{f_{\text{sem}}(e_{vr})}{\|f_{\text{sem}}(e_{vr})\|_2}. \tag{5}$$

A Gaussian kernel with bandwidth $\sigma > 0$ measures query-set and within-set similarity:

$$k_\sigma(u, w) = \exp\left(-\frac{\|u - w\|_2^2}{2\sigma^2}\right), \qquad [K_v]_{rs} = k_\sigma(g_{vr}, g_{vs}), \qquad [\mathbf{k}_{qv}]_r = k_\sigma(h_q, g_{vr}). \tag{6}$$

Here, K_v records within-set correlations, while $\mathbf{k}_{qv}$ records query-to-description similarity.

Definition 3.1 (Regularized set support)**.** For $\lambda_{\mathrm{rk}} > 0$, define the semantic support of E_v for q as

$$S(q, E_v) = \mathbf{k}_{qv}^\top (K_v + \lambda_{\mathrm{rk}} I)^{-1} \mathbf{k}_{qv}. \tag{7}$$

Let ϕ denote the Gaussian kernel's mapping into a reproducing kernel Hilbert space $\mathcal{H}$. The support measures regularized reconstruction of the query from the description set:

$$\min_{\alpha \in \mathbb{R}^{m_v}} \left\| \phi(h_q) - \sum_r \alpha_r \phi(g_{vr}) \right\|_{\mathcal{H}}^2 + \lambda_{\mathrm{rk}} \|\alpha\|_2^2 = 1 - S(q, E_v). \tag{8}$$

The normal equations give $\alpha^\star = (K_v + \lambda_{\mathrm{rk}} I)^{-1} \mathbf{k}_{qv}$; substitution yields Eq. (8) because $k_\sigma(h_q, h_q) = 1$. The inverse matrix adjusts contributions according to correlations within the set. This construction combines complementary semantic directions while accounting for redundancy among overlapping descriptions within the set.

Query-conditioned allocation. Absolute support values are query dependent. Let $\widehat{F}_q^{\mathrm{ref}}$ be the empirical cumulative distribution function of support over training videos unpaired with q. The resulting query-relative semantic support is defined as

$$u(q, v) = \widehat{F}_q^{\mathrm{ref}}\big(S(q, E_v)\big) \in [0, 1]. \tag{9}$$

Larger values indicate stronger support relative to the query's own unpaired reference. With $u_{aj} = u(q, v_j)$, SEAR exponentially tilts the baseline allocation:

$$q_{aj} = \frac{p_{aj} \exp[-\eta_s(t) u_{aj}]}{\sum_{r \in \mathcal{N}_a} p_{ar} \exp[-\eta_s(t) u_{ar}]}, \qquad \eta_s(t) \geq 0. \tag{10}$$

Here t indexes training progress, and KL denotes Kullback–Leibler divergence. The allocation minimizes $\mathrm{KL}(r \| p_a) + \eta_s(t) \sum_j r_j u_{aj}$ over the probability simplex. The divergence term anchors the solution to p_a, while the semantic term reduces average support under the new allocation. Thus, for $\eta_s(t) > 0$, higher support gives a smaller q_{aj}/p_{aj} within each query. Because q_a remains normalized, SEAR changes where the baseline gradient mass is assigned without using semantic similarity alone to suppress the query-level coefficient.

3.4 Positive-reference Attribution for Intervention Refinement

PAIR uses task responses to refine constraints after accounting for relation difficulty (Figure 2(c)). It compares local compatibility statistics between paired and unpaired relations at similar difficulty, then converts excess positive-reference evidence into continuous multipliers. Conditioning on global hardness isolates local specificity that cannot be explained by relation difficulty alone, providing the statistical basis for relation-specific refinement.

Difficulty and local specificity. A lagged model snapshot $\bar{\theta}$ fixes the response reference within each statistics-update interval. For relation (a, j), the baseline selection rule chooses a representative temporal unit $\bar{m}_{aj}$. Write its response as $\bar{s}(a, j) = s_{\bar{\theta}}(q, \bar{m}_{aj})$. Relative to the query's unpaired reference $\mathcal{C}_a^{\mathrm{ref}}$, we define the query-relative difficulty percentile

$$H_{aj} = \frac{1}{|\mathcal{C}_a^{\mathrm{ref}}|} \sum_{k \in \mathcal{C}_a^{\mathrm{ref}}} \mathbf{1}\{\bar{s}(a, k) \leq \bar{s}(a, j)\}. \tag{11}$$

The percentile H_{aj} measures how difficult the candidate is relative to other candidates for the same query. To measure query specificity, hold $\bar{m}_{aj}$ fixed. Let $\mathcal{Q}_{\bar{m}_{aj}}^{\mathrm{ref}}$ denote the reference queries for that unit, and define specificity as follows:

$$C_{aj} = \frac{1}{|\mathcal{Q}_{\bar{m}_{aj}}^{\mathrm{ref}}|} \sum_{q' \in \mathcal{Q}_{\bar{m}_{aj}}^{\mathrm{ref}}} \mathbf{1}\{s_{\bar{\theta}}(q', \bar{m}_{aj}) \leq s_{\bar{\theta}}(q, \bar{m}_{aj})\}. \tag{12}$$

Thus, H compares candidates for a fixed query, and C compares queries for a fixed temporal unit. Every reference query uses the same selected unit when computing C. The pair (H, C) separates global competition from local query specificity: H defines comparable difficulty neighborhoods, while C tests whether a relation is unusually selective within such a neighborhood.

Conditional positive-reference evidence. Let P and U denote the paired and unpaired training relation populations. Both use the same temporal-unit selection and response-reference rules. Because specificity can vary with difficulty, PAIR compares $p(C \mid H = h, P)$ and $p(C \mid H = h, U)$. Partition the specificity range into common bins $\{B_\ell\}_{\ell=1}^{L}$. With a nonnegative local kernel K, bandwidth $b_H > 0$, and smoothing constant $\epsilon > 0$, estimate

$$\widehat{p}_{D,\ell}(h) = \frac{\sum_{i \in D} w_i(h) \mathbf{1}\{C_i \in B_\ell\} + \epsilon}{\sum_{i \in D} w_i(h) + L\epsilon}, \quad w_i(h) = K\left(\frac{H_i - h}{b_H}\right), \quad D \in \{P, U\}. \tag{13}$$

Local weighting lets relations with similar difficulty support the conditional comparison. Smoothing ensures positive bin probabilities for the following ratios.

Definition 3.2 (Positive-reference attribution)**.** Let $\widehat{F}_D^{(h)}$ denote the categorical distribution given by $\widehat{p}_{D,\ell}(h)$. Its maximal compatible positive-reference proportion satisfies

$$\begin{aligned} \widehat{\kappa}(h) &= \min\left\{1, \min_{\ell} \frac{\widehat{p}_{U,\ell}(h)}{\widehat{p}_{P,\ell}(h)}\right\}, \\ \widehat{F}_U^{(h)} &= \widehat{\kappa}(h)\widehat{F}_P^{(h)} + [1 - \widehat{\kappa}(h)]\widehat{G}_h. \end{aligned} \tag{14}$$

For the bin $\ell(a, j)$ containing C_{aj}, define the reference responsibility

$$R_{aj} = \widehat{\kappa}(H_{aj}) \frac{\widehat{p}_{P,\ell(a,j)}(H_{aj})}{\widehat{p}_{U,\ell(a,j)}(H_{aj})}. \tag{15}$$

Each bin bounds the removable reference proportion so that the residual distribution $\widehat{G}_h$ remains nonnegative. When $\widehat{\kappa}(h) = 1$, the two distributions coincide and the residual carries zero weight. Equation (15) follows Bayes' rule for this fitted mixture.

Proposition 3.3 (Centered reference evidence)**.** *For each fixed h, the fitted distributions satisfy $0 \le R \le 1$ and $\mathbb{E}_{\widehat{F}_U^{(h)}}[R] = \widehat{\kappa}(h)$. We therefore define the centered excess evidence by*

$$\xi_{aj} = [R_{aj} - \widehat{\kappa}(H_{aj})]_+, \qquad [x]_+ = \max(x, 0). \tag{16}$$

If $\widehat{F}_P^{(h)} = \widehat{F}_U^{(h)}$, every relation at h has $\xi = 0$.

The bound follows from Eq. (14); summing $\widehat{p}_{U,\ell}(h)R(h, \ell)$ over bins gives the mean identity. Centering retains evidence above the conditional unpaired mean, emphasizing specificity regions with greater positive-reference density.

Remark 3.4 (Interpretation). The responsibility measures attribution to the fitted paired-reference component. Interpreting it as a true relevance probability requires additional identifiability assumptions. GRACE uses the excess ξ to calibrate constraints continuously.

Bounded constraint refinement. PAIR maps excess evidence to a relation-specific multiplier using a query-level coefficient $0 \le \lambda_a < 1$:

$$\rho_{aj} = 1 - \lambda_a \xi_{aj}. \tag{17}$$

Because $0 \le \xi_{aj} \le 1$, the multiplier satisfies $0 < \rho_{aj} \le 1$. For $\lambda_a > 0$, larger excess evidence reduces the direct multiplier; zero excess preserves $\rho_{aj} = 1$. The resulting weight $\rho_{aj} q_{aj}$ refines SEAR's allocation in the unified objective.

3.5 Unified Objective and Training Evolution

GRACE integrates both adjustments through the evidence-derived coefficient

$$\omega_{aj} = \rho_{aj} \frac{q_{aj}}{p_{aj}}, \qquad \ell_a^{\mathrm{GRACE}} = -\log \frac{e^{z_a^+}}{e^{z_a^+} + \sum_{j \in \mathcal{N}_a} \omega_{aj} e^{z_{aj}}}. \tag{18}$$

For the gradient analysis, ω_{aj} is held fixed with respect to the current retrieval logits; the corresponding implementation is given in Appendix A.4. Define $A_a = \sum_j \rho_{aj} q_{aj}$ and

$$\widetilde{q}_{aj} = \frac{\rho_{aj} q_{aj}}{A_a}, \qquad \beta'_a = \frac{M_a A_a}{e^{z_a^+} + M_a A_a}. \tag{19}$$

Proposition 3.5 (Unified gradient decomposition)**.** *For positive normalized allocations p_a, q_a and $0 < \rho_{aj} \leq 1$, Eq.* (18) *gives*

$$\frac{\partial \ell_a^{\mathrm{GRACE}}}{\partial z_{aj}} = \beta'_a \widetilde{q}_{aj}, \qquad \beta_a p_a \xrightarrow{\mathrm{SEAR}} \beta_a q_a \xrightarrow{\mathrm{PAIR}} \beta'_a \widetilde{q}_a. \tag{20}$$

With fixed ω, substituting $e^{z_{aj}} = M_a p_{aj}$ gives the decomposition. For SEAR, $\rho_{aj} = 1$ implies $A_a = 1$ and $\beta'_a = \beta_a$. PAIR permits $A_a \leq 1$, changing the final allocation while keeping the total coefficient at $\beta'_a \leq \beta_a$.

Evidence-dependent intervention scope. Equation (20) makes the two evidence sources non-interchangeable. SEAR preserves $A_a = 1$ and $\beta'_a = \beta_a$, so semantic evidence redistributes negative-gradient mass without changing the aggregate coefficient. Once conditional task statistics receive sufficient support, PAIR changes both $\widetilde{q}_a$ and β'_a through ρ_{aj}, enabling bounded attenuation for relations with excess positive-reference evidence. Semantic evidence therefore controls where negative supervision is allocated, whereas task-derived evidence can additionally modulate the strength of that supervision throughout optimization.

GRACE uses the same objective throughout training. The semantic strength $\eta_s(t)$ starts positive and decreases, while the lagged snapshot and conditional statistics are refreshed periodically with a fixed response reference in each interval. PAIR defaults to $\rho_{aj} = 1$ when statistical support is insufficient. At inference, only the baseline query–video score is used; the semantic encoder and reference statistics are training-only.

4 Experiments

4.1 Experimental Setup

Setup. We evaluate on Charades-STA (CHA) (Gao et al., 2017), ActivityNet Captions (ACT) (Krishna et al., 2017), and TVR (Lei et al., 2020) under the PRVR protocol (Dong et al., 2022), reporting R@K for $K \in \{1, 5, 10, 100\}$ and Rsum. Controlled comparisons use GMMFormerV2 (Wang et al., 2024a) with its official pre-extracted features and three paired seeds (3407, 3408, 3409). Each GRACE run is matched to the corresponding baseline seed, so reported differences are not confounded by a different initialization. SEAR uses a frozen CLIP ViT-L/14 encoder, while PAIR refreshes its lagged snapshot every five epochs; neither changes the inference-time features or scoring function. All controlled comparisons retain the same backbone, pre-extracted features, and inference scoring function. Detailed implementation settings are provided in Appendix A.

4.2 Main Results

Table 1 compares GRACE with representative PRVR methods under the same recall metrics. We use Rsum as the aggregate measure while retaining each cutoff to distinguish improvements in early ranking from broader candidate coverage. GRACE improves Rsum by 1.3, 1.5, and 2.3 points on CHA, ACT, and TVR, respectively, attaining the highest Rsum among the listed methods.

Table 1: Comparison with representative PRVR methods on three benchmarks.

Method	CHA					ACT					TVR				
	1	5	10	100	**Sum**	1	5	10	100	**Sum**	1	5	10	100	**Sum**
MS-SL (2022)	1.8	7.1	11.8	47.7	68.4	7.1	22.5	34.7	75.8	140.1	13.5	32.1	43.4	83.4	172.4
PEAN (2023)	**2.7**	8.1	13.5	50.3	74.7	7.4	23.0	35.5	75.9	141.8	13.5	32.8	44.1	83.9	174.2
GMMFormer (2024b)	2.1	7.8	12.5	50.6	72.9	8.3	24.9	36.7	76.1	146.0	13.9	33.3	44.5	84.9	176.6
BGM-Net (2024)	1.9	7.4	12.2	50.1	71.6	7.2	23.8	36.0	76.9	143.9	14.1	34.7	45.9	85.2	179.9
HLFormer (2025)	2.6	8.5	13.7	**54.0**	78.7	8.7	27.1	40.1	79.0	154.9	15.7	37.1	48.5	86.4	187.7
DreamPRVR (2026b)†	2.2	8.4	14.4	53.6	78.4	8.7	27.3	40.4	**79.3**	155.7	16.4	38.1	49.4	86.4	190.2
GMMFormerV2 (2024a)	2.5	8.6	13.9	53.2	78.2	8.9	27.1	40.2	78.7	154.9	16.2	37.6	48.8	86.4	189.1
+ GRACE	2.4	**8.8**	**14.5**	53.8	**79.5**	**9.0**	**27.7**	**41.0**	78.8	**156.4**	**16.5**	**38.4**	**49.9**	**86.6**	**191.4**

† denotes our three-seed reproduction.

The gains are consistent across all reported cutoffs on ACT and TVR. On CHA, R@1 changes by -0.1, whereas R@5/10/100 increase by 0.2/0.6/0.6, so the aggregate improvement is not driven by a single retrieval depth. Because GRACE uses the same backbone, pre-extracted features, and inference score as GMMFormerV2, the controlled difference isolates the effect of training-time constraint calibration.

GRACE also exceeds the strongest listed non-GRACE Rsum by 0.8, 0.7, and 1.2 points on CHA, ACT, and TVR, respectively.

4.3 Ablation Studies

Components. Table 2(A) compares the base NCE objective, SEAR alone ($\rho = 1$), PAIR alone ($\mathbf{q} = \mathbf{p}$), and their combination. Each module improves retrieval, while their combination exceeds the stronger individual module by 0.5, 0.5, and 0.9 Rsum points, supporting complementary allocation and relation-specific refinement.

Table 2: Structural ablations in Rsum. (A) Contributions of SEAR and PAIR across benchmarks. (B) Design controls on CHA; all other settings follow full GRACE.

(A) Components	CHA	ACT	TVR
Baseline NCE	78.2	154.9	189.1
+ SEAR	78.8	155.6	190.0
+ PAIR	79.0	155.9	190.5
GRACE	**79.5**	**156.4**	**191.4**

(B) Design control	CHA
Max / mean similarity	79.1 / 78.8
W/o hardness conditioning	78.7
Raw R	78.9
Constant semantic influence	79.2

Key designs. Table 2(B) evaluates the main design choices on CHA by replacing set support with maximum/mean similarity, removing hardness conditioning, using raw responsibility R, or keeping semantic influence constant. All controls in panel (B) remain below full GRACE in panel (A), supporting set-level evidence, conditional centering, and gradual calibration.

4.4 Evidence and Mechanism Analysis

Semantic evidence tracks candidate difficulty. Figure 3(a) compares the candidate rankings induced by semantic support u and NCE logits. At the top 5%, their overlap rises from 12.7% after epoch 1 to 26.3% after epoch 9, against a 4.7% random expectation. The alignment concentrates in the ranking head, where allocation changes matter most, supporting semantic evidence as a guide for conservative candidate redistribution.

Task evidence retains local structure beyond difficulty. Figures 3(b)–(c) compare paired and unpaired relations after matching the empirical distribution of global hardness H. The retained early and late snapshots emphasize the transition from immature to mature relation geometry: paired

relations preserve stronger local specificity C, and by the later snapshot 63.1% of paired mass versus 19.2% of matched unpaired mass lies at $C = 1$. Thus, C retains local query–video compatibility information that global difficulty leaves unresolved.

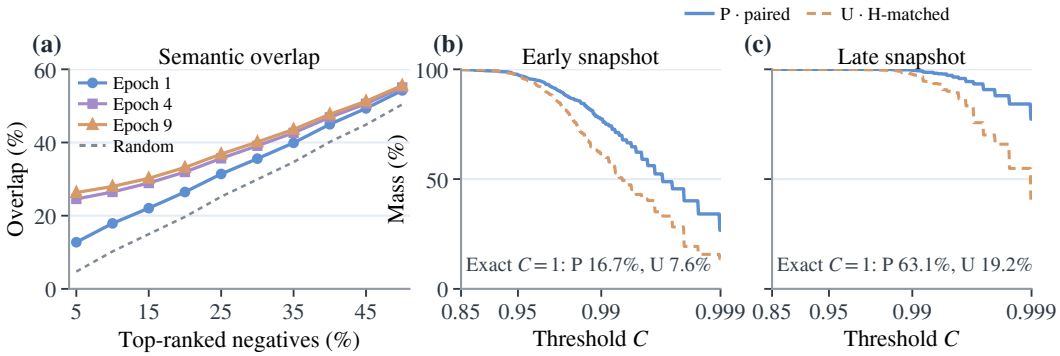

Figure 3: Evidence and mechanism diagnostics. (a) Semantic evidence increasingly aligns with retrieval difficulty; the gray dashed line denotes random ranking. (b) Early PAIR snapshot after hardness matching. (c) Later PAIR snapshot after hardness matching. Solid blue denotes paired relations P, and dashed warm curves denote hardness-matched unpaired relations U.

4.5 Efficiency

Table 3 shows that GRACE increases training time from 0.551 to 0.959 min/epoch and peak memory from 7.859 to 8.498 GB, while corpus scoring latency, inference memory, and parameter count are unchanged. Query encoding latency changes only from 0.0124 to 0.0148 ms/query, indicating that the additional cost is primarily on the training side.

Table 3: Training overhead and inference cost under matched settings.

	Training		Inference			
Method	Time (min/epoch)	Memory (GB)	Query enc. (ms/query)	Corpus score (ms/query)	Memory (GB)	Params. (M)
Baseline	0.551	7.859	0.0124	0.1159	3.349	30.791
GRACE	0.959	8.498	0.0148	0.1159	3.349	30.791

5 Conclusion and Limitation

We presented GRACE, which evolves constraint calibration from semantic allocation to difficulty-conditioned relation refinement. SEAR redistributes candidate contributions according to semantic evidence, while PAIR applies bounded relation-specific attenuation within the same contrastive objective. Together, they provide complementary mechanisms for improving the reliability of training constraints under partial relevance. Results on three PRVR benchmarks and controlled analyses consistently support their complementary roles and effectiveness, with no change to the inference-time retriever or scoring function. These findings highlight constraint reliability as a useful perspective for learning under partially relevant supervision. Future work will study more challenging settings with weaker semantic evidence, sparser conditional references, or greater relevance ambiguity.

GRACE relies on auxiliary semantic evidence and conditional references to calibrate training constraints, whose reliability may decrease in highly ambiguous or sparsely referenced scenarios. In addition, our current study focuses on established PRVR benchmarks, and its generalization to broader retrieval settings remains to be explored.

Reproducibility Statement

Appendix A provides the implementation protocol and training procedure, Appendix B reports seed-level results and controlled ablations, and Appendix C states the assumptions and gives the theoretical arguments underlying GRACE. Controlled baseline/GRACE comparisons use matched seeds together with the same backbone, pre-extracted features, and evaluation protocol.

AI Use Statement

Generative AI tools were used for language polishing and presentation refinement, as well as for retrieval and discovery of potentially relevant literature and related materials. They were not used for research ideation or execution, drafting substantive sections, generating synthetic datasets, or proving mathematical claims. All AI-assisted edits and retrieved materials were reviewed by the authors, who take full responsibility for the final content of the paper.

References

C.-H. Cho, W. Moon, W. Jun, M. Jung, and J.-P. Heo. Ambiguity-restrained text-video representation learning for partially relevant video retrieval. In *AAAI*, volume 39, pp. 2500–2508, 2025.

C.-Y. Chuang, J. W. Robinson, Y.-C. Lin, A. Torralba, and S. Jegelka. Debiased contrastive learning. In *NeurIPS*, volume 33, 2020.

J. Dong, X. Chen, M. Zhang, X. Yang, S. Chen, X. Li, and X. Wang. Partially relevant video retrieval. In *ACM MM*, pp. 246–257, 2022.

J. Dong, M. Zhang, Z. Zhang, X. Chen, D. Liu, X. Qu, X. Wang, and B. Liu. Dual learning with dynamic knowledge distillation for partially relevant video retrieval. In *ICCV*, pp. 11302–11312, 2023.

J. Gao, C. Sun, Z. Yang, and R. Nevatia. TALL: Temporal activity localization via language query. In *ICCV*, pp. 5267–5275, 2017.

Y. Huang, L. Yang, and Y. Sato. Weakly supervised temporal sentence grounding with uncertainty-guided self-training. In *CVPR*, pp. 18908–18918, 2023.

T. Huynh, S. Kornblith, M. R. Walter, M. Maire, and M. Khademi. Boosting contrastive self-supervised learning with false negative cancellation. In *WACV*, pp. 2785–2795, 2022.

T. Jang and X. Wang. Difficulty-based sampling for debiased contrastive representation learning. In *CVPR*, pp. 24039–24048, 2023.

X. Jiang, Z. Chen, X. Xu, F. Shen, Z. Cao, and X. Cai. Progressive event alignment network for partial relevant video retrieval. In *ICME*, pp. 1973–1978, 2023.

Y. Kalantidis, M. B. Sariyildiz, N. Pion, P. Weinzaepfel, and D. Larlus. Hard negative mixing for contrastive learning. In *NeurIPS*, volume 33, 2020.

R. Kiryo, G. Niu, M. C. du Plessis, and M. Sugiyama. Positive-unlabeled learning with non-negative risk estimator. In *NeurIPS*, volume 30, 2017.

R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles. Dense-captioning events in videos. In *ICCV*, pp. 706–715, 2017.

J. Lei, L. Yu, T. L. Berg, and M. Bansal. TVR: A large-scale dataset for video-subtitle moment retrieval. In *ECCV*, pp. 447–463, 2020.

J. Li, R. Selvaraju, A. Gotmare, S. Joty, C. Xiong, and S. C. H. Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, volume 34, 2021.

J. Li, J. Wang, C. Tan, N. Lian, L. Chen, Y. Wang, M. Zhang, S.-T. Xia, and B. Chen. Enhancing partially relevant video retrieval with hyperbolic learning. In *ICCV*, pp. 23074–23084, 2025.

J. Li, P. Lai, X. Lou, J. Wang, Y. Wang, K. Chen, Y. Wang, and S.-T. Xia. Revisiting uncertainty: On evidential learning for partially relevant video retrieval. In *ICML*, 2026a.

J. Li, X. Lou, J. Wang, Y. Wang, Y. Wang, S.-T. Xia, and B. Chen. Imagine before concentration: Diffusion-guided registers enhance partially relevant video retrieval. In *CVPR*, pp. 9710–9721, 2026b.

W. Moon, C.-H. Cho, W. Jun, T. Kim, I. Lee, D. Wee, M. Shim, and J.-P. Heo. Prototypes are balanced units for efficient and effective partially relevant video retrieval. In *ICCV*, pp. 21789–21799, 2025a.

W. Moon, M. Jung, G. Park, T.-Y. Kim, C.-H. Cho, W. Jun, and J.-P. Heo. Mitigating semantic collapse in partially relevant video retrieval. In *NeurIPS*, volume 38, 2025b.

A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, volume 139, pp. 8748–8763, 2021.

H. Ramaswamy, C. Scott, and A. Tewari. Mixture proportion estimation via kernel embeddings of distributions. In *ICML*, volume 48, pp. 2052–2060, 2016.

J. Robinson, C.-Y. Chuang, S. Sra, and S. Jegelka. Contrastive learning with hard negative samples. In *ICLR*, 2021.

A. Tarvainen and H. Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, volume 30, 2017.

Y. Wang, J. Wang, B. Chen, T. Dai, R. Luo, and S.-T. Xia. GMMFormer V2: An uncertainty-aware framework for partially relevant video retrieval. *arXiv preprint arXiv:2405.13824*, 2024a.

Y. Wang, J. Wang, B. Chen, Z. Zeng, and S.-T. Xia. GMMFormer: Gaussian-mixture-model based transformer for efficient partially relevant video retrieval. In *AAAI*, volume 38, pp. 5767–5775, 2024b.

W. Wu, H. Luo, B. Fang, J. Wang, and W. Ouyang. Cap4Video: What can auxiliary captions do for text-video retrieval? In *CVPR*, pp. 10704–10713, 2023.

S. Yin, S. Zhao, H. Wang, T. Xu, and E. Chen. Exploiting instance-level relationships in weakly supervised text-to-video retrieval. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 20(10):316:1–316:21, 2024.

L. Zhang, P. Song, J. Dong, K. Li, and X. Yang. Enhancing partially relevant video retrieval with robust alignment learning. In *Findings of EMNLP*, pp. 4615–4629, 2025.

APPENDIX

This appendix provides additional material supporting the implementation, empirical evaluation, and theoretical analysis of GRACE. Appendix A details the implementation and reproducibility protocol, including the training settings and the realization of SEAR and PAIR. Appendix B reports extended and per-seed experimental results, component and design-control analyses, and the DreamPRVR three-seed reproduction on CHA, ACT, and TVR. Appendix C presents the theoretical motivation and derivations underlying evolving constraint reliability, semantic evidence allocation, and positive-reference refinement.

A IMPLEMENTATION AND REPRODUCIBILITY DETAILS

This section complements the main paper with the implementation protocol used for GRACE and the controlled comparisons, while omitting dataset-specific engineering details that do not alter the method. GRACE augments GMMFormerV2 only during training. SEAR uses fixed semantic evidence to recalibrate the allocation of unpaired candidates, whereas PAIR uses lagged task responses to refine their effective constraint strength. At inference, all auxiliary evidence is discarded and the original retrieval model is used unchanged.

A.1 REPRODUCIBILITY PROTOCOL

Retrieval setting. We directly use the official pre-extracted video and text features released with GMMFormerV2 and retain its retrieval architecture and standard data-processing pipeline. Unless otherwise stated below, model optimization and architecture settings follow the official GMMFormerV2 configuration. This keeps the comparison focused on the training-time constraint calibration introduced by GRACE rather than on changes to the underlying retriever.

Training runs. All matched baseline/GRACE comparisons use three paired random seeds, 3407, 3408, and 3409. Each run is trained for **100 epochs**. Within each seed, the baseline and GRACE share the same retrieval inputs, sampling procedure, optimization setting, and evaluation protocol. The auxiliary evidence used by GRACE is constructed exclusively from the training split; evaluation queries and labels are never used to build semantic or relation-reference statistics. Table A.1 summarizes the core settings shared by the reported runs.

Table A.1: Core implementation settings used by the current GRACE configuration.

Setting	Configuration
Retrieval model	GMMFormerV2
Retrieval inputs	Official GMMFormerV2 pre-extracted video/text features
Semantic encoder	Frozen CLIP ViT-L/14 text encoder
Random seeds	3407, 3408, 3409
Training duration	100 epochs
PAIR snapshot/statistics refresh	Every 5 epochs
Evidence source	Training split only
Inference	Original GMMFormerV2 retrieval model only

GRACE-specific settings. SEAR uses a Gaussian kernel over normalized semantic query embeddings, with its bandwidth estimated separately for each dataset using a median-distance heuristic on the training queries. The ridge coefficient is set to 10^{-3}, and the semantic influence starts at 0.1 and follows a half-cosine decay to zero by epoch 10. PAIR discretizes local specificity into 16 bins and performs hardness conditioning with a triangular kernel of bandwidth 0.05, using a pseudocount of 0.5 for density smoothing. Its attenuation strength is determined adaptively for each query or video anchor. Except for the dataset-adaptive semantic bandwidth, these settings are shared across all three benchmarks.

A.2 SEAR IMPLEMENTATION

For every training video v, SEAR reuses its paired training-query set

$$E_v = \{q' : (q', v) \in \mathcal{D}_{\mathrm{tr}}\}. \tag{A.1}$$

The current query and all descriptions in E_v are encoded with a frozen CLIP ViT-L/14 text encoder and ℓ_2 normalized. These semantic embeddings are fixed throughout retrieval training and are used only to construct SEAR evidence.

Set-level semantic support. On the normalized semantic vectors, SEAR uses the Gaussian kernel and regularized set-support construction defined in the main paper,

$$k_\sigma(x, y) = \exp\left(-\frac{\|x - y\|_2^2}{2\sigma^2}\right), \tag{A.2}$$

$$S(q, E_v) = \mathbf{k}_{qv}^\top (K_v + \lambda_{\mathrm{rk}} I)^{-1} \mathbf{k}_{qv}. \tag{A.3}$$

The kernel bandwidth is determined from training semantic representations and then held fixed during a run; the ridge term is likewise fixed and is used only to stabilize the set-level reconstruction. We do not retune these quantities across random seeds.

Query-relative calibration and allocation. Raw semantic support is converted to the query-relative evidence coordinate $u(q, v)$ using the query's unpaired training reference set, as described in the main paper. Given the baseline within-negative allocation p_{aj}, SEAR produces

$$q_{aj} = \frac{p_{aj} \exp[-\eta_s(t) u_{aj}]}{\sum_{r \in \mathcal{N}_a} p_{ar} \exp[-\eta_s(t) u_{ar}]}. \tag{A.4}$$

The semantic strength $\eta_s(t)$ starts positive and decreases with training progress. The schedule is fixed before the three-seed evaluation and is shared by all matched runs. SEAR does not introduce a separate training stage or modify the retrieval objective; it only changes the normalized allocation within the same objective. If the semantic evidence required for an anchor cannot be constructed, SEAR falls back to the original allocation, $\mathbf{q}_a = \mathbf{p}_a$.

A.3 PAIR IMPLEMENTATION

PAIR constructs task-derived evidence from a lagged copy of the retrieval model. A snapshot $\bar{\theta}$ fixes the response reference within each statistics-update interval. The first reference statistics are initialized from a lagged model state, and the snapshot together with its associated statistics is subsequently refreshed every **five epochs**. Between refreshes, the reference model and statistics remain fixed.

Reference relations. For each relation (a, j), PAIR uses the same baseline temporal-unit selection rule as the main retriever. The selected temporal unit is kept fixed when constructing the query-relative difficulty H_{aj} and the reciprocal specificity C_{aj}. Paired and unpaired reference relations are drawn only from the training split, with observed positive relations excluded from the unpaired population.

Difficulty-conditioned statistics. PAIR estimates the paired and unpaired conditional distributions

$$p(C \mid H = h, P), \qquad p(C \mid H = h, U), \tag{A.5}$$

using the kernel-smoothed histogram estimator defined in the main paper. A common binning scheme, local hardness kernel, bandwidth, and smoothing rule are used within each run and are kept fixed across the three random seeds. These estimator settings control numerical resolution only; they do not introduce an additional supervision source.

From the estimated distributions, PAIR computes the maximal compatible paired-reference proportion $\widehat{\kappa}(h)$, the relation responsibility R_{aj}, and the centered excess evidence

$$\xi_{aj} = \left[R_{aj} - \widehat{\kappa}(H_{aj})\right]_+. \tag{A.6}$$

The excess evidence is mapped to the bounded multiplier

$$\rho_{aj} = 1 - \lambda_a \xi_{aj}, \qquad 0 \le \lambda_a < 1, \tag{A.7}$$

exactly as in the main paper. PAIR therefore performs continuous attenuation rather than assigning hard pseudo-labels. When the local paired/unpaired reference population is insufficient for a stable conditional estimate, PAIR uses identity refinement, $\xi_{aj} = 0$ and hence $\rho_{aj} = 1$. We intentionally describe this fallback at the level of statistical support rather than introducing implementation-specific support thresholds into the method definition.

A.4 Unified Training Procedure

GRACE uses a single training objective throughout all 100 epochs. SEAR changes where negative-gradient mass is allocated, while PAIR can additionally reduce the effective pressure assigned to relations with excess positive-reference evidence. The two mechanisms are integrated through the stop-gradient weight

$$\omega_{aj} = \mathrm{sg}\left(\frac{\rho_{aj} q_{aj}}{p_{aj}}\right). \tag{A.8}$$

Algorithm A.1 summarizes the resulting training procedure.

Algorithm A.1 GRACE training procedure

Require: Retrieval model f_θ; training set $\mathcal{D}_{\mathrm{tr}}$; 100 training epochs
1: Load the official GMMFormerV2 retrieval features and initialize f_θ.
2: Precompute the fixed CLIP ViT-L/14 semantic embeddings used by SEAR.
3: Initialize the lagged snapshot $\bar{\theta}$ and its PAIR reference statistics.
4: **for** epoch $e = 1, \ldots, 100$ **do**
5: **for** mini-batch $\mathcal{B}$ **do**
6: Compute live retrieval logits and baseline negative allocation p.
7: Compute SEAR evidence and obtain the normalized allocation q.
8: Query the fixed PAIR statistics to obtain ξ and ρ.
9: Use identity calibration whenever the required auxiliary evidence is unavailable.
10: Form $\omega = \mathrm{sg}(\rho q / p)$ and optimize the weighted contrastive objective.
11: **end for**
12: **if** e is a five-epoch refresh boundary **then**
13: Update the lagged snapshot and recompute the PAIR reference statistics.
14: **end if**
15: **end for**

Objective integration. GRACE modifies only the negative allocation and relation-specific attenuation in the contrastive term. The original positive terms and the remaining GMMFormerV2 training components are retained. The auxiliary semantic and PAIR quantities are detached from the live retrieval logits when constructing ω, implementing the fixed-coefficient treatment used in the main-paper gradient analysis.

Inference and efficiency reporting. At inference, GRACE loads only the trained GMMFormerV2 retrieval model. The CLIP semantic encoder, semantic evidence, lagged snapshots, and PAIR statistics are not required. Efficiency comparisons therefore use identical inference inputs and retrieval architecture for the baseline and GRACE. Training and inference measurements are collected under matched hardware, precision, batch-size, and corpus settings; latency is measured after warm-up with synchronized repeated runs. This protocol supports the training-overhead and inference-cost measurements reported in the main paper without introducing dataset-specific implementation details into the method specification.

B Extended and Per-seed Results

This appendix reports run-to-run stability, paired improvements, complete retrieval metrics for structural ablations and design controls, and the reproduced DreamPRVR results on CHA, ACT, and TVR

used in the main comparison. Sample standard deviation is used for all three-run summaries. Rsum is computed from the unrounded run-level recall values, while manuscript-level values are rounded after aggregation.

B.1 Per-seed Results on the Main Benchmarks

Tables B.1–B.3 expand the matched GMMFormerV2 baseline and GRACE into the three paired runs. Because run-level recalls are displayed to two decimal places, the displayed entries need not sum exactly to Rsum after rounding; Rsum is computed from the underlying unrounded values.

Table B.1: Per-seed retrieval results on CHA.

Method	Run	R@1	R@5	R@10	R@100	Rsum
Baseline	3407	2.58	8.44	14.11	52.93	78.06
	3408	2.61	8.41	13.90	53.49	78.41
	3409	2.28	8.90	13.82	53.15	78.15
	Mean $\pm$ SD	2.49 ± 0.18	8.58 ± 0.27	13.94 ± 0.15	53.19 ± 0.28	78.21 ± 0.18
GRACE	3407	2.45	8.58	14.49	53.68	79.19
	3408	2.47	8.98	14.46	53.98	79.89
	3409	2.39	8.82	14.52	53.68	79.41
	Mean $\pm$ SD	2.44 ± 0.04	8.79 ± 0.20	14.49 ± 0.03	53.78 ± 0.17	$\mathbf{79.50 \pm 0.36}$

Table B.2: Per-seed retrieval results on ACT.

Method	Run	R@1	R@5	R@10	R@100	Rsum
Baseline	3407	8.92	26.88	40.46	78.32	154.58
	3408	9.00	27.11	40.16	78.84	155.11
	3409	8.72	27.37	39.93	78.99	155.01
	Mean $\pm$ SD	8.88 ± 0.14	27.12 ± 0.24	40.18 ± 0.26	78.72 ± 0.35	154.90 ± 0.28
GRACE	3407	8.89	27.72	40.78	78.46	155.85
	3408	9.05	27.73	41.37	78.73	156.88
	3409	9.06	27.56	40.73	79.12	156.47
	Mean $\pm$ SD	9.00 ± 0.10	27.67 ± 0.09	40.96 ± 0.36	78.77 ± 0.33	$\mathbf{156.40 \pm 0.52}$

Table B.3: Per-seed retrieval results on TVR.

Method	Run	R@1	R@5	R@10	R@100	Rsum
Baseline	3407	16.24	37.99	48.55	86.07	188.85
	3408	16.40	37.52	48.83	86.54	189.29
	3409	16.08	37.41	49.13	86.54	189.16
	Mean $\pm$ SD	16.24 ± 0.16	37.64 ± 0.31	48.84 ± 0.29	86.38 ± 0.27	189.10 ± 0.23
GRACE	3407	16.48	38.10	49.49	86.86	190.92
	3408	16.45	38.48	50.32	86.52	191.76
	3409	16.51	38.69	49.83	86.49	191.52
	Mean $\pm$ SD	16.48 ± 0.03	38.42 ± 0.30	49.88 ± 0.42	86.62 ± 0.20	$\mathbf{191.40 \pm 0.43}$

B.2 Paired Improvements Across Seeds

Because the baseline and GRACE use the same seeds, paired Rsum differences directly show whether the aggregate improvement is consistent across matched runs. Table B.4 reports these paired gains.

Table B.4: Paired Rsum improvements of GRACE over the matched GMMFormerV2 baseline.

Dataset	Seed 3407	Seed 3408	Seed 3409	Mean gain
CHA	1.13	1.48	1.26	**1.29**
ACT	1.27	1.76	1.46	**1.50**
TVR	2.07	2.47	2.36	**2.30**

B.3 Full Component Ablations

Table B.5 first reports per-seed Rsum for the four structural variants. Tables B.6–B.8 then expand each variant to R@1/5/10/100.

Table B.5: Per-seed Rsum for the structural component ablations corresponding to Table 2(A) in the main paper.

Dataset	Variant	Seed 3407	Seed 3408	Seed 3409	Mean ± SD
CHA	Baseline NCE	78.06	78.41	78.15	78.21 ± 0.18
	+ SEAR	78.52	79.06	78.84	78.81 ± 0.27
	+ PAIR	78.74	79.27	78.98	79.00 ± 0.27
	GRACE	79.19	79.89	79.41	79.50 ± 0.36
ACT	Baseline NCE	154.58	155.11	155.01	154.90 ± 0.28
	+ SEAR	155.28	155.96	155.56	155.60 ± 0.34
	+ PAIR	155.55	156.33	155.82	155.90 ± 0.40
	GRACE	155.85	156.88	156.47	156.40 ± 0.52
TVR	Baseline NCE	188.85	189.29	189.16	189.10 ± 0.23
	+ SEAR	189.58	190.20	190.22	190.00 ± 0.36
	+ PAIR	190.03	190.75	190.72	190.50 ± 0.41
	GRACE	190.92	191.76	191.52	191.40 ± 0.43

Table B.6: Full component-ablation metrics on CHA.

Variant	R@1	R@5	R@10	R@100	Rsum
Baseline NCE	2.49 ± 0.18	8.58 ± 0.27	13.94 ± 0.15	53.19 ± 0.28	78.21 ± 0.18
+ SEAR	2.46 ± 0.06	8.72 ± 0.13	14.12 ± 0.16	53.50 ± 0.18	78.81 ± 0.27
+ PAIR	2.51 ± 0.11	8.74 ± 0.13	14.23 ± 0.29	53.52 ± 0.31	79.00 ± 0.27
GRACE	2.44 ± 0.04	8.79 ± 0.20	14.49 ± 0.03	53.78 ± 0.17	**79.50 ± 0.36**

Table B.7: Full component-ablation metrics on ACT.

Variant	R@1	R@5	R@10	R@100	Rsum
Baseline NCE	8.88 ± 0.14	27.12 ± 0.24	40.18 ± 0.26	78.72 ± 0.35	154.90 ± 0.28
+ SEAR	8.95 ± 0.05	27.32 ± 0.24	40.46 ± 0.10	78.87 ± 0.41	155.60 ± 0.34
+ PAIR	8.98 ± 0.17	27.39 ± 0.26	40.58 ± 0.17	78.95 ± 0.35	155.90 ± 0.40
GRACE	9.00 ± 0.10	27.67 ± 0.09	40.96 ± 0.36	78.77 ± 0.33	**156.40 ± 0.52**

Table B.8: Full component-ablation metrics on TVR.

Variant	R@1	R@5	R@10	R@100	Rsum
Baseline NCE	16.24 ± 0.16	37.64 ± 0.31	48.84 ± 0.29	86.38 ± 0.27	189.10 ± 0.23
+ SEAR	16.31 ± 0.15	37.80 ± 0.33	49.14 ± 0.45	86.75 ± 0.32	190.00 ± 0.36
+ PAIR	16.38 ± 0.08	37.93 ± 0.13	49.34 ± 0.27	86.85 ± 0.32	190.50 ± 0.41
GRACE	16.48 ± 0.03	38.42 ± 0.30	49.88 ± 0.42	86.62 ± 0.20	$\mathbf{191.40 \pm 0.43}$

B.4 Full Design-control Results on CHA

The main paper reports only Rsum for the design controls. Tables B.9 and B.10 provide the three individual runs and the complete rank-wise metrics, respectively.

Table B.9: Per-seed Rsum for design controls on CHA corresponding to Table 2(B) in the main paper.

Design control	Seed 3407	Seed 3408	Seed 3409	Mean $\pm$ SD
Max similarity	78.82	79.44	79.06	79.10 ± 0.31
Mean similarity	78.55	79.11	78.74	78.80 ± 0.29
W/o hardness conditioning	78.39	79.03	78.71	78.71 ± 0.32
Raw R	78.68	79.14	78.87	78.90 ± 0.23
Constant semantic influence	78.92	79.54	79.11	79.19 ± 0.32
GRACE	79.19	79.89	79.41	$\mathbf{79.50 \pm 0.36}$

Table B.10: Full design-control metrics on CHA.

Design control	R@1	R@5	R@10	R@100	Rsum
Max similarity	2.44 ± 0.07	8.76 ± 0.19	14.39 ± 0.14	53.51 ± 0.38	79.10 ± 0.31
Mean similarity	2.42 ± 0.11	8.69 ± 0.11	14.26 ± 0.07	53.43 ± 0.27	78.80 ± 0.29
W/o hardness conditioning	2.46 ± 0.11	8.67 ± 0.08	14.22 ± 0.27	53.35 ± 0.29	78.71 ± 0.32
Raw R	2.47 ± 0.07	8.72 ± 0.24	14.32 ± 0.09	53.39 ± 0.38	78.90 ± 0.23
Constant semantic influence	2.45 ± 0.07	8.76 ± 0.22	14.42 ± 0.16	53.57 ± 0.06	79.19 ± 0.32
GRACE	2.44 ± 0.04	8.79 ± 0.20	14.49 ± 0.03	53.78 ± 0.17	$\mathbf{79.50 \pm 0.36}$

B.5 Three-seed DreamPRVR Reproduction

The main comparison marks the reproduced DreamPRVR results on CHA, ACT, and TVR with a dagger. Tables B.11–B.13 provide the corresponding three-run results.

Table B.11: Three-seed DreamPRVR reproduction on CHA.

Run	R@1	R@5	R@10	R@100	Rsum
3407	2.04	8.44	14.09	53.60	78.17
3408	2.20	8.20	14.60	53.66	78.66
3409	2.20	8.41	14.38	53.41	78.41
Mean $\pm$ SD	2.15 ± 0.09	8.35 ± 0.13	14.35 ± 0.26	53.56 ± 0.13	$\mathbf{78.41 \pm 0.24}$

Table B.12: Three-seed DreamPRVR reproduction on ACT.

Run	R@1	R@5	R@10	R@100	Rsum
3407	8.62	27.17	40.25	79.18	155.22
3408	8.76	27.39	40.51	79.39	156.05
3409	8.72	27.34	40.44	79.33	155.83
Mean $\pm$ SD	8.70 ± 0.07	27.30 ± 0.12	40.40 ± 0.13	79.30 ± 0.11	$\mathbf{155.70 \pm 0.43}$

Table B.13: Three-seed DreamPRVR reproduction on TVR.

Run	R@1	R@5	R@10	R@100	Rsum
3407	16.46	37.93	49.17	86.28	189.84
3408	16.42	38.32	49.28	86.45	190.47
3409	16.20	37.93	49.63	86.53	190.28
Mean $\pm$ SD	16.36 ± 0.14	38.06 ± 0.23	49.36 ± 0.24	86.42 ± 0.13	$\mathbf{190.20 \pm 0.32}$

C Theoretical Motivation and Design Rationale

This appendix develops a structural analysis of the problem addressed by GRACE and of the design choices underlying SEAR and PAIR. The purpose is not to model the full nonconvex dynamics of a neural retriever, nor to claim that an external semantic encoder is universally correct. Instead, the results make explicit the conditions under which (i) semantic evidence is useful when task relations are immature, (ii) a persistent external prior becomes undesirable after task relations mature, and (iii) task-derived relation evidence can support a broader intervention scope. Implementation-specific estimators and hyperparameters are deliberately excluded from the analysis.

Notation. Fix an anchor query and suppress its index when no ambiguity arises. Let $\Delta_n = \{r \in \mathbb{R}^n_+ : \sum_{j=1}^n r_j = 1\}$ be the probability simplex over its unpaired candidates. The baseline candidate allocation is $p \in \mathrm{int}(\Delta_n)$, and $u \in [0,1]^n$ denotes semantic support, with larger u_j indicating stronger evidence that candidate j should be treated conservatively as a negative. For $\eta \geq 0$, define the semantic tilt

$$T_\eta(p;u)_j = \frac{p_j \exp(-\eta u_j)}{Z_\eta(p,u)}, \qquad Z_\eta(p,u) = \sum_{r=1}^n p_r \exp(-\eta u_r). \tag{C.1}$$

This is the abstract form of the allocation used by SEAR.

C.1 Why Constraint Reliability Must Evolve

The central difficulty is not merely that relation evidence is imperfect, but that the relative reliability of available evidence changes as the task representation develops. We formalize the two opposing requirements and then analyze the stability of task-derived rank evidence.

C.1.1 Early-stage allocation error and semantic correction

Let $r^\star \in \Delta_n$ denote a latent ideal allocation: the distribution of negative-gradient mass that would be used if the unresolved relation structure were known. It is an analysis device and is not required by the algorithm. Holding the aggregate gradient coefficient fixed isolates the allocation error between the current model and this ideal target.

Proposition C.1 (A sufficient condition for beneficial semantic correction)**.** *For $q_\eta = T_\eta(p;u)$, define*

$$\mathcal{R}(\eta) = \mathrm{KL}(r^\star \| q_\eta). \tag{C.2}$$

Then

$$\mathcal{R}'(\eta) = \mathbb{E}_{r^\star}[u] - \mathbb{E}_{q_\eta}[u], \qquad \mathcal{R}''(\eta) = \mathrm{Var}_{q_\eta}(u) \geq 0. \tag{C.3}$$

Consequently, if

$$\mathbb{E}_p[u] > \mathbb{E}_{r^\star}[u], \tag{C.4}$$

then there exists $\eta_0 > 0$ *such that*

$$\mathrm{KL}(r^\star \| T_\eta(p;u)) < \mathrm{KL}(r^\star \| p) \qquad \textit{for all } \eta \in (0, \eta_0]. \tag{C.5}$$

Proof. Using $\log q_{\eta,j} = \log p_j - \eta u_j - \log Z_\eta$,

$$\mathcal{R}(\eta) = \mathrm{KL}(r^\star \| p) + \eta\, \mathbb{E}_{r^\star}[u] + \log Z_\eta(p,u). \tag{C.6}$$

Differentiating $\log Z_\eta$ gives $\frac{\mathrm{d}}{\mathrm{d}\eta} \log Z_\eta = -\mathbb{E}_{q_\eta}[u]$, which yields the first identity in Eq. (C.3). Moreover, $\frac{\mathrm{d}}{\mathrm{d}\eta}\mathbb{E}_{q_\eta}[u] = -\operatorname{Var}_{q_\eta}(u)$, which yields the second identity. At $\eta = 0$, $q_0 = p$, so condition (C.4) implies $\mathcal{R}'(0) < 0$. Continuity then guarantees a nonempty interval on which $\mathcal{R}(\eta) < \mathcal{R}(0)$. □

Proposition C.1 makes the early-stage requirement precise. If the current allocation places more mass on semantically supported candidates than the ideal negative allocation does, a sufficiently mild semantic tilt improves the allocation in KL divergence. The statement is deliberately conditional: semantic evidence is useful when it is directionally aligned with unresolved relation error, rather than by assumption in every training state.

C.1.2 Late-stage bias of a persistent external prior

The same semantic tilt need not remain desirable after the task model has learned an appropriate relation geometry. The next result shows that a nonvanishing external tilt generally leaves a residual allocation distortion even if the baseline allocation itself becomes task consistent.

Proposition C.2 (Persistent-prior distortion)**.** *Let* $p_t \to r^\star$ *and* $\eta_t \to \eta_\infty$*. Then*

$$T_{\eta_t}(p_t;u) \longrightarrow T_{\eta_\infty}(r^\star;u). \tag{C.7}$$

If $\eta_\infty > 0$ *and* u *is not constant* $r^\star$*-almost surely, then*

$$T_{\eta_\infty}(r^\star;u) \neq r^\star \tag{C.8}$$

and

$$0 < \mathrm{KL}(r^\star \| T_{\eta_\infty}(r^\star;u)) = \eta_\infty \mathbb{E}_{r^\star}[u] + \log \mathbb{E}_{r^\star}[e^{-\eta_\infty u}] \le \frac{\eta_\infty^2}{8}. \tag{C.9}$$

If instead $\eta_t \to 0$*, then* $T_{\eta_t}(p_t;u) \to r^\star$*.*

Proof. Continuity of Eq. (C.1) gives the first and last convergence statements. Setting $p = r^\star$ in the KL expansion used in Proposition C.1 yields the equality in Eq. (C.9). Strict positivity follows from strict Jensen inequality when $\eta_\infty > 0$ and u is nonconstant on the support of $r^\star$. Since $u \in [0,1]$, Hoeffding's lemma applied to $-\eta_\infty(u - \mathbb{E}_{r^\star} u)$ gives the upper bound. □

Thus, a decreasing semantic strength has two complementary roles. It retains the possibility of an early corrective effect, while controlling the worst-case residual distortion once the task allocation becomes reliable. Exact convergence to zero is sufficient for asymptotic removal of the external tilt; more generally, Eq. (C.9) quantifies how the remaining distortion scales with its final strength.

C.1.3 Stability of lagged task-derived rank evidence

PAIR relies on query-relative and moment-relative rank statistics. Their usefulness depends on whether a lagged model provides a sufficiently stable relation reference. The following result isolates the effect of representation drift without committing to a particular finite-sample estimator.

Lemma C.3 (Rank stability under bounded parameter drift)**.** *Let every score entering a rank comparison satisfy*

$$|s_\theta(x) - s_{\theta'}(x)| \le L\|\theta - \theta'\| \tag{C.10}$$

for some $L > 0$. Consider a rank statistic

$$\Psi_\theta = \frac{1}{N}\sum_{r=1}^{N} \mathbf{1}\{d_r(\theta) \geq 0\}, \tag{C.11}$$

where each $d_r(\theta)$ is the difference of two such scores and all compared temporal units are held fixed. For a lagged parameter $\bar{\theta}$,

$$|\Psi_\theta - \Psi_{\bar{\theta}}| \leq \frac{1}{N}\sum_{r=1}^{N} \mathbf{1}\left\{|d_r(\bar{\theta})| \leq 2L\|\theta - \bar{\theta}\|\right\}. \tag{C.12}$$

In particular, if all comparison margins satisfy $|d_r(\bar{\theta})| > 2L\|\theta - \bar{\theta}\|$, then $\Psi_\theta = \Psi_{\bar{\theta}}$.

Proof. Because each difference contains two L-Lipschitz scores,

$$|d_r(\theta) - d_r(\bar{\theta})| \leq 2L\|\theta - \bar{\theta}\|. \tag{C.13}$$

An indicator can change only if the corresponding comparison changes sign. Such a sign change requires $|d_r(\bar{\theta})| \leq |d_r(\theta) - d_r(\bar{\theta})|$, which is possible only for terms counted on the right-hand side of Eq. (C.12). Averaging the changed indicators proves the result. □

Both difficulty and local specificity are averages of score-comparison indicators, so Lemma C.3 applies to their rank structure when the response reference is fixed. As optimization stabilizes and $\|\theta - \bar{\theta}\|$ decreases, only increasingly small-margin comparisons can change. Under a continuous margin distribution, the fraction of such comparisons vanishes with the drift radius. This does not eliminate finite-sample error, but it shows why one important source of unreliability—staleness of the task-derived reference—contracts as relation learning matures.

Taken together, Propositions C.1 and C.2 and Lemma C.3 formalize the motivating asymmetry. A stable external prior can reduce allocation error when the task model is immature, yet a persistent prior generically distorts a mature task allocation. Conversely, task-derived rank evidence is most defensible after model drift has become small relative to relation margins. This is the theoretical basis for evolving evidence authority rather than assigning one evidence source a fixed role throughout training.

C.2 Theoretical Rationale of SEAR

SEAR requires two design choices: a set-level semantic support that respects correlations among video-associated descriptions, and a conservative transformation from that support to candidate allocation. We analyze them in turn.

C.2.1 Set-level support as regularized RKHS reconstruction

Let $E_v = \{e_{v1}, \ldots, e_{vm}\}$ be the descriptions associated with video v, let h_q and g_{vr} be normalized semantic embeddings, and let k be a positive-definite kernel with feature map ϕ into an RKHS $\mathcal{H}$. Define

$$[K_v]_{rs} = k(g_{vr}, g_{vs}), \qquad [\mathbf{k}_{qv}]_r = k(h_q, g_{vr}), \tag{C.14}$$

and, for $\lambda > 0$,

$$S(q, E_v) = \mathbf{k}_{qv}^\top (K_v + \lambda I)^{-1}\mathbf{k}_{qv}. \tag{C.15}$$

Proposition C.4 (Regularized set reconstruction)**.** *If $k(h_q, h_q) = 1$, then*

$$1 - S(q, E_v) = \min_{\alpha \in \mathbb{R}^m} \left\| \phi(h_q) - \sum_{r=1}^{m} \alpha_r \phi(g_{vr}) \right\|_{\mathcal{H}}^2 + \lambda\|\alpha\|_2^2. \tag{C.16}$$

Consequently, $0 \leq S(q, E_v) \leq 1$, and S is invariant to any permutation of the descriptions in E_v.

Proof. Expanding the objective in Eq. (C.16) gives

$$1 - 2\alpha^\top \mathbf{k}_{qv} + \alpha^\top (K_v + \lambda I)\alpha. \tag{C.17}$$

Since $K_v + \lambda I$ is positive definite, the unique minimizer is

$$\alpha^\star = (K_v + \lambda I)^{-1}\mathbf{k}_{qv}. \tag{C.18}$$

Substitution yields $1 - S(q, E_v)$. Nonnegativity of the minimized objective gives $S \leq 1$, while positive definiteness of $(K_v + \lambda I)^{-1}$ gives $S \geq 0$. For a permutation matrix P, reordering the set gives $K_v' = PK_vP^\top$ and $\mathbf{k}_{qv}' = P\mathbf{k}_{qv}$. Hence

$$(\mathbf{k}_{qv}')^\top (K_v' + \lambda I)^{-1}\mathbf{k}_{qv}' = \mathbf{k}_{qv}^\top (K_v + \lambda I)^{-1}\mathbf{k}_{qv}. \tag{C.19}$$

□

The proposition gives a precise set-level interpretation. The query is compared with the span of all associated semantic directions, while the Gram matrix couples their contributions according to within-set correlation. This differs from maximum similarity, which retains only one description, and mean similarity, which combines descriptions without explicitly accounting for their dependence.

C.2.2 Minimum-change semantic reallocation

Proposition C.5 (Unique minimum-KL reallocation)**.** *For $p \in \mathrm{int}(\Delta_n)$, $u \in \mathbb{R}^n$, and $\eta \geq 0$, the optimization problem*

$$\min_{r \in \Delta_n} \{\mathrm{KL}(r\|p) + \eta\langle r, u\rangle\} \tag{C.20}$$

has the unique solution

$$q_j = T_\eta(p; u)_j = \frac{p_j e^{-\eta u_j}}{\sum_r p_r e^{-\eta u_r}}. \tag{C.21}$$

Moreover,

$$\frac{q_i/p_i}{q_j/p_j} = \exp[-\eta(u_i - u_j)], \tag{C.22}$$

so candidates with greater semantic support receive strictly smaller relative allocation whenever $\eta > 0$.

Proof. The KL term is strictly convex on the simplex interior because $p_j > 0$, and the semantic term is linear. The Lagrangian stationarity condition is

$$\log \frac{r_j}{p_j} + 1 + \eta u_j + \nu = 0, \tag{C.23}$$

which gives $r_j \propto p_j e^{-\eta u_j}$. Normalization yields the stated solution, and strict convexity gives uniqueness. Equation (C.22) follows directly by taking the ratio of two normalized weights. □

Several consequences clarify why the reallocation is conservative. First, adding a query-dependent constant to every semantic score leaves q unchanged, because only relative evidence matters. Second, $q = p$ at $\eta = 0$, and a first-order expansion gives

$$q_j = p_j \left[1 - \eta\big(u_j - \mathbb{E}_p[u]\big)\right] + O(\eta^2). \tag{C.24}$$

Thus, weak semantic intervention perturbs the baseline allocation according to centered evidence rather than replacing it with an independent semantic distribution.

Corollary C.6 (Conservation of aggregate negative-gradient mass)**.** *Suppose the baseline negative-logit gradient is $g_j^{(0)} = \beta p_j$ for some $\beta > 0$. Replacing p with the SEAR allocation q gives*

$$\sum_j \beta q_j = \beta = \sum_j \beta p_j. \tag{C.25}$$

Hence semantic evidence changes where negative-gradient mass is assigned, but not its aggregate coefficient.

This conservation property is the theoretical expression of low-authority intervention: external evidence corrects allocation while preserving the total negative pressure imposed by the task objective.

C.3 Theoretical Rationale of PAIR and Progressive Refinement

PAIR is based on a conditional comparison between paired and unpaired relation populations at the same difficulty. We formulate the argument at the distribution level; empirical estimation is a separate implementation issue.

C.3.1 Maximal compatible positive-reference proportion

Fix a difficulty value h. Let P_h and U_h be the conditional distributions of local specificity C for paired and unpaired relations, respectively. Assume they admit densities p_h and u_h with respect to a common dominating measure μ. Define

$$\kappa_h^\star = \sup\{\kappa \in [0,1] : U_h - \kappa P_h \text{ is a nonnegative measure}\}. \tag{C.26}$$

Proposition C.7 (Maximal compatible mixture)**.** *With the convention that $u_h(c)/p_h(c) = 0$ whenever $p_h(c) > 0$ and $u_h(c) = 0$,*

$$\kappa_h^\star = \operatorname*{ess\,inf}_{c:p_h(c)>0} \frac{u_h(c)}{p_h(c)}. \tag{C.27}$$

If $\kappa_h^\star < 1$, then

$$G_h = \frac{U_h - \kappa_h^\star P_h}{1 - \kappa_h^\star} \tag{C.28}$$

is a probability distribution and

$$U_h = \kappa_h^\star P_h + (1 - \kappa_h^\star) G_h. \tag{C.29}$$

If $\kappa_h^\star = 1$, then $U_h = P_h$ almost everywhere.

Proof. The measure $U_h - \kappa P_h$ is nonnegative if and only if $u_h(c) - \kappa p_h(c) \geq 0$ μ-almost everywhere. Equivalently, every feasible κ satisfies

$$\kappa \leq \frac{u_h(c)}{p_h(c)} \qquad \mu\text{-a.e. on } \{p_h > 0\}. \tag{C.30}$$

The largest feasible value is therefore the essential infimum in Eq. (C.27). Since both densities integrate to one, this value cannot exceed one. For $\kappa_h^\star < 1$, the residual measure is nonnegative and has total mass $1 - \kappa_h^\star$, so its normalization is a probability distribution. If $\kappa_h^\star = 1$, then $u_h \geq p_h$ almost everywhere and equality of their integrals implies $u_h = p_h$ almost everywhere. □

The proposition identifies the largest paired-reference component compatible with the observed unpaired distribution, without asserting that this component is the true prevalence of relevant relations. It is a distributional compatibility quantity.

C.3.2 Responsibility centering identifies positive-reference enrichment

For $u_h(c) > 0$, define the responsibility under the compatible mixture by

$$R_h(c) = \kappa_h^\star \frac{p_h(c)}{u_h(c)}. \tag{C.31}$$

On null sets where the ratio is undefined, set $R_h = 0$.

Proposition C.8 (Responsibility background and centered excess)**.** *The responsibility satisfies*

$$0 \leq R_h(c) \leq 1 \quad U_h\text{-a.e.}, \qquad \mathbb{E}_{U_h}[R_h] = \kappa_h^\star. \tag{C.32}$$

If $\kappa_h^\star > 0$, then

$$R_h(c) - \kappa_h^\star = \kappa_h^\star \frac{p_h(c) - u_h(c)}{u_h(c)} \quad U_h\text{-a.e.} \tag{C.33}$$

Therefore the centered excess

$$\xi_h(c) = [R_h(c) - \kappa_h^\star]_+ \tag{C.34}$$

is positive exactly on regions where $p_h(c) > u_h(c)$, up to null sets. If $P_h = U_h$, then $\xi_h = 0$ almost everywhere.

Proof. Feasibility of $\kappa_h^\star$ gives $\kappa_h^\star p_h(c) \le u_h(c)$ almost everywhere, proving $R_h \le 1$; nonnegativity is immediate. Moreover,

$$\mathbb{E}_{U_h}[R_h] = \int u_h(c)\, \kappa_h^\star \frac{p_h(c)}{u_h(c)}\, \mathrm{d}\mu(c) = \kappa_h^\star. \tag{C.35}$$

Subtracting $\kappa_h^\star$ from Eq. (C.31) gives Eq. (C.33). For $\kappa_h^\star > 0$, its sign is the sign of $p_h - u_h$. Finally, if $P_h = U_h$, Proposition C.7 gives $\kappa_h^\star = 1$ and $R_h = 1$, hence $\xi_h = 0$. □

The conditional mean in Eq. (C.32) is the difficulty-specific responsibility background. Subtracting it is therefore not an arbitrary normalization: it removes attribution that is expected under the entire unpaired population at the same difficulty and retains only positive-reference enrichment. By contrast, raw responsibility can be nonzero even in specificity regions that are not enriched relative to the conditional unpaired distribution.

Remark C.9 (Scope of the responsibility). $R_h(c)$ is a Bayes responsibility for the fitted decomposition in Eq. (C.29). Identifying it with the probability that an individual relation is truly relevant would require additional assumptions such as mixture identifiability and a correct component model. PAIR uses only the weaker, justified interpretation: excess attribution to a compatible paired-reference component.

C.3.3 Bounded relation-specific attenuation

Let $\lambda \in [0, 1)$ and define

$$\rho_h(c) = 1 - \lambda \xi_h(c). \tag{C.36}$$

Corollary C.10 (Bounded refinement)**.** *For every fixed h,*

$$0 \le \xi_h(c) \le 1 - \kappa_h^\star \tag{C.37}$$

and

$$1 - \lambda(1 - \kappa_h^\star) \le \rho_h(c) \le 1. \tag{C.38}$$

In particular, $0 < \rho_h(c) \le 1$: PAIR can attenuate a negative constraint, but cannot reverse it or amplify it beyond its incoming weight.

Proof. Proposition C.8 gives $0 \le R_h \le 1$, so $0 \le [R_h - \kappa_h^\star]_+ \le 1 - \kappa_h^\star$. Substituting this bound into Eq. (C.36) yields Eq. (C.38). Strict positivity follows from $\lambda < 1$. □

C.3.4 Distinct intervention scopes in the unified objective

The final result connects the two design rationales at the gradient level. Let the baseline negative mass be $M = \sum_j e^{z_j}$, let $p_j = e^{z_j}/M$, let $q \in \Delta_n$ be the SEAR allocation, and let $0 < \rho_j \le 1$ be the PAIR multiplier. Consider the weighted objective

$$\ell = -\log \frac{e^{z^+}}{e^{z^+} + \sum_j \omega_j e^{z_j}}, \qquad \omega_j = \mathrm{sg}\left(\rho_j \frac{q_j}{p_j}\right), \tag{C.39}$$

where sg denotes stop-gradient. Define

$$A = \sum_j \rho_j q_j, \qquad \widetilde{q}_j = \frac{\rho_j q_j}{A}, \qquad \beta' = \frac{MA}{e^{z^+} + MA}. \tag{C.40}$$

Proposition C.11 (Progressive constraint authority)**.** *Holding the evidence coefficients fixed during differentiation,*

$$\frac{\partial \ell}{\partial z_j} = \beta' \widetilde{q}_j. \tag{C.41}$$

With SEAR alone, $\rho_j = 1$ gives $A = 1$, $\widetilde{q} = q$, and $\beta' = \beta = M/(e^{z^+} + M)$. With PAIR, $A \le 1$ and hence $\beta' \le \beta$. Therefore

$$\underbrace{\beta p}_{\text{baseline}} \xrightarrow{\text{SEAR}} \underbrace{\beta q}_{\text{allocation correction}} \xrightarrow{\text{PAIR}} \underbrace{\beta' \widetilde{q}}_{\text{allocation and bounded attenuation}}. \tag{C.42}$$

Proof. In the forward pass,

$$\sum_j \omega_j e^{z_j} = \sum_j \rho_j \frac{q_j}{p_j} M p_j = MA. \tag{C.43}$$

Since ω_j is held fixed in the backward pass,

$$\frac{\partial \ell}{\partial z_j} = \frac{\omega_j e^{z_j}}{e^{z^+} + MA} = \frac{M \rho_j q_j}{e^{z^+} + MA} = \beta' \widetilde{q}_j. \tag{C.44}$$

If $\rho_j = 1$, normalization of q gives $A = 1$. In general, $0 < \rho_j \leq 1$ gives $0 < A \leq 1$. The map $A \mapsto MA/(e^{z^+} + MA)$ is increasing, so $\beta' \leq \beta$. □

Equation (C.42) summarizes the theoretical division of labor. Semantic evidence is admitted through a minimum-change transformation that preserves aggregate negative-gradient mass. Task-derived positive-reference enrichment is admitted only after conditioning on difficulty and is mapped to a multiplier with an explicit lower bound; it can therefore adjust both the relative allocation and the aggregate strength of negative supervision. This establishes a principled progression of intervention authority rather than an interchangeable use of two evidence sources.