

Not Just What’s There: Enabling CLIP to Comprehend Negated Visual Descriptions Without Fine-tuning

Junhao Xiao¹, Zhiyu Wu², Hao Lin³, Yi Chen^{1*}, Yahui Liu^{4*},
Xiaoran Zhao⁵, Zixu Wang⁵, Zejiang He⁵

¹Central China Normal University,

²Fudan University,

³Huazhong University of Science and Technology,

⁴Kuaishou Technology,

⁵National University of Defense Technology
chenyi30@mail.ccnu.edu.cn, yahui.cvrs@gmail.com

Abstract

Vision-Language Models (VLMs) like CLIP struggle to understand negation, often embedding affirmatives and negatives similarly (e.g., matching “no dog” with dog images). Existing methods refine negation understanding via fine-tuning CLIP’s text encoder, risking overfitting. In this work, we propose CLIPGLASSES, a plug-and-play framework that enhances CLIP’s ability to comprehend negated visual descriptions. CLIPGLASSES adapts a dual-stage design: a *Lens* module disentangles negated semantics from text embeddings, and a *Frame* module predicts context-aware repulsion strength, which is integrated into the modified similarity computation to penalize alignment with negated semantics, thereby reducing false positive matches. Experiments show that CLIP equipped with CLIPGLASSES achieves competitive in-domain performance and outperforms state-of-the-art methods in cross-domain generalization. Its superiority is especially evident under low-resource conditions, indicating stronger robustness across domains.

Code — <https://github.com/Codecode-X/CLIPGlasses.git>

Introduction

Recent advances in large-scale pretraining have advanced Vision-Language Models (VLMs), with CLIP (Radford et al. 2021) emerging as a foundational model. CLIP enables cross-modal alignment by projecting images and texts into a shared embedding space, underpinning core tasks such as retrieval, captioning, and text-conditioned generation.

However, CLIP exhibits notable limitations in modeling negation semantics. As shown in Figure 1, CLIP fails to properly handle negation cues like “no” or “without” in text inputs, incorrectly matching negated concepts with corresponding visual content rather than recognizing their absence. This failure stems from the sparsity of negation expressions in pretraining corpora (i.e., only 0.7% (Park et al. 2025)), which prevents contrastive learning from effectively capturing semantic polarity reversals.

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

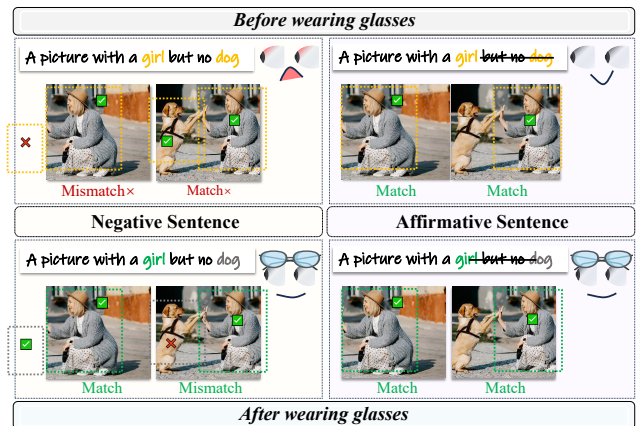


Figure 1: CLIPGLASSES enhances CLIP’s capacity for negation understanding by introducing a dynamic repulsion mechanism that suppresses image-text similarity for negated concepts, thus enabling inverse matching while preserving alignment with affirmed content.

Several existing approaches have attempted to adapt CLIP models to negation-sensitive tasks through parameter fine-tuning (Yuksekgonul et al. 2022; Alhamoud et al. 2025; Park et al. 2025; Singh et al. 2025), but these approaches pose two critical drawbacks. First, constructing large-scale negation-annotated datasets is time-consuming and resource-intensive. Second, fine-tuning introduces the risk of catastrophic forgetting, whereby enhanced negation understanding comes at the expense of deteriorating general-purpose performance. This represents a fundamental trade-off between specialized capability and broad applicability.

To address these limitations, we draw inspiration from two key observations. First, although affirmative and negative semantics lie close in CLIP’s feature space, visualization analysis (Figure 2) reveals structured separability enabled by layer-specific encoding (Quantmeyer, Mosteiro, and Gatt 2024), indicating the feasibility of disentangling negation-related information from CLIP-encoded text embeddings. Second, cognitive studies show that humans process negation in two stages, first identifying the negated con-

cept, then inverting its meaning (Kaup, Zwaan, and Lüdtkke 2007; Orenes, Beltrán, and Santamaría 2014; Zuanazzi et al. 2024). Guided by these insights, we design a plug-and-play framework CLIPGLASSES that leverages CLIP’s latent negation representations and emulates this two-stage process.

Specifically, CLIPGLASSES extends CLIP through two lightweight modules. The *Lens* module disentangles negated semantics from text embeddings, while the *Frame* module predicts context-dependent repulsion strength. These components jointly enable a modified similarity computation: negated content identified by *Lens* is explicitly penalized through repulsion vectors whose magnitude is dynamically scaled by *Frame*’s strength predictions. This integrated process is visually summarized in Figure 3. To coordinate these components, we employ a progressive three-stage regimen with frozen CLIP parameters: first training *Lens* to disentangle negated text representations, then training *Frame* to model dynamic repulsion using ground-truth negation features, and finally joint optimization of both modules with *Lens* outputs to enhance synergy. In this way, instead of modifying the model’s “eyes”, we simply let it wear “glasses” to better perceive the negation.

Experimental results establish CLIPGLASSES’s balanced advantages through systematic comparisons. While fine-tuning baselines overfit negation datasets for marginal in-domain gains (CoN-CLIP: 99.70% vs. our 96.56% on CC-Neg-val), this incurs unwilling costs: degraded cross-domain generalization (25.70% vs. our 34.51% on Neg-COCO-MCQ). Under low-resource constraints (5K images), these limitations intensify—our approach surpasses CoN-CLIP by 27.45 points on CC-Neg-val and 5.29 points on Neg-COCO-MCQ. Moreover, unlike fine-tuning approaches that impair CLIP’s native zero-shot performance on standard non-negation benchmarks, our non-invasive architecture effectively preserves these intrinsic capabilities. These results substantiate the efficacy of CLIPGLASSES in achieving transferable, semantically grounded negation understanding without sacrificing model robustness or generalization. We summarize our key contributions as follows:

- We present CLIPGLASSES, a non-intrusive framework enhancing CLIP’s negation modeling via human-inspired two-stage processing without parameter modification.
- We design a novel architecture, including a syntax-semantic *Lens* for disentangling negation semantics, and a *Frame* for modeling context-aware repulsion, and a modified similarity computation that explicitly reverses alignment with negated content.
- Our method attains state-of-the-art trade-offs between in-domain accuracy and cross-domain generalization, without compromising CLIP’s native zero-shot abilities.

Related Work

Vision-Language Models and Their Limitations. Vision-language models like CLIP (Radford et al. 2021) demonstrate strong performance on broad cross-modal tasks, including retrieval (Qin et al. 2025; Caffagni et al. 2025; Fang et al. 2024; Zhang et al. 2025a; Li et al. 2025b; Zhang et al.

2025b; Li et al. 2025a), captioning (Kim et al. 2025; Nam et al. 2025; Li et al. 2025c) and VQA (Xing et al. 2024; Sun et al. 2024; Jiang et al. 2025) and generation (Zhang et al. 2025d; Feng et al. 2025; Rombach et al. 2022; Zhang et al. 2025c; Fang et al. 2025; Zhang et al. 2025e), leveraging representations learned from large-scale noisy image-text pairs. However, they exhibit critical deficiencies in fine-grained semantic understanding, manifesting in four key areas: over-reliance on shallow statistical cues (Zhang et al. 2023; Geirhos et al. 2020), failures in compositional reasoning (Yuksekonul et al. 2022; Thrush et al. 2022), weak attribute grounding, and bag-of-words-like text processing that ignores syntactic structure (Koishigarina, Uselis, and Oh 2025). Negation understanding poses a particularly challenging case, with VLMs frequently failing to distinguish negated and affirmative inputs (Morante and Blanco 2021; Ma et al. 2023) due to affirmation bias and scarcity of negation in pretraining data (occurring in $\leq 0.7\%$ of captions (Park et al. 2025)), significantly impairing negation-sensitive applications, such as medical or clinical scenarios (Ko and Park 2025; Su et al. 2025).

Prior Work on Negation Understanding. Existing methods universally fine-tune CLIP’s text encoder and fall into two categories based on data strategy: (1) Compositional Perturbation, typified by NegCLIP (Yuksekonul et al. 2022), which uses word-order shuffling to simulate syntactic variation in negation; and (2) Paraphrastic Augmentation, including CoN-CLIP (Singh et al. 2025), Negation-CLIP (Park et al. 2025), and NegBench (Alhamoud et al. 2025), which leverage LLM-generated negatives to enrich contrastive training. While effective in-domain, these methods overwrite CLIP’s pretrained representations, risking overfitting, catastrophic forgetting (Zhai et al. 2024; Jha, Gong, and Yao 2024; Li et al. 2024), and poor cross-domain generalization due to the lack of explicit negation modeling. In contrast, CLIPGLASSES introduces dedicated negation reasoning via architectural extensions, preserving CLIP’s zero-shot capabilities without modifying its parameters.

Preliminary: Two-Stage Modeling for Negation Understanding

Although affirmative and negative semantics are often embedded in close proximity within CLIP’s feature space, visualization analysis (See Figure 2) reveals a promising potential for semantic disentanglement. This separability arises from layer-specific negation encoding mechanisms (Quantmeyer, Mosteiro, and Gatt 2024) and forms a solid basis for targeted extraction of negation-related information.

Additionally, findings from cognitive neuroscience indicate that humans typically process negation in two distinct stages (Kaup, Zwaan, and Lüdtkke 2007; Orenes, Beltrán, and Santamaría 2014; Zuanazzi et al. 2024): First, by identifying the target object or concept being negated; Second, by inverting its semantic implication to derive the negated meaning. Inspired by this two-stage cognitive mechanism, we propose a corresponding computational modeling strategy to better handle negation in vision-language alignment. In the first stage, the model explicitly identifies and extracts nega-

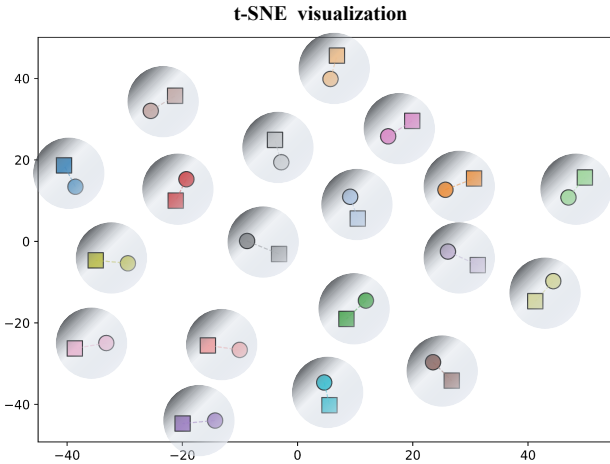


Figure 2: t-SNE (Cai and Ma 2022) visualization of CLIP text features for multiple positive-negative sentence pairs (e.g., “there is a woman” vs. “there is not a woman”). Circles and squares denote positive and negative forms, colors distinguish different pairs. Feature clusters across pairs are well-separated, showing CLIP’s strong instance-level discrimination. However, positive and negative features within individual pair remain closely positioned, indicating that while CLIP has limited negation modeling capabilities, there exists clear potential for semantic disentanglement.

tion semantics from textual features. In the second stage, it modulates the image-text similarity by explicitly suppressing alignment with negated concepts.

Methodology

To instantiate the two-stage modeling paradigm outlined above, we propose CLIPGLASSES, a modular extension to CLIP for negation-aware vision-language alignment. It comprises two components: *Lens*, which disentangles negated semantics from text embeddings, and *Frame*, which predicts context-dependent repulsion strength. These components are integrated into a modified similarity computation that penalizes alignment with negated content. An overview of the architecture is shown in Figure 3. We describe the architecture and functionality of each module in the following and then present the final matching strategy.

Lens: Syntax–Semantic Dual-Stream Architecture

While CLIP’s text encoder effectively models affirmative semantics, it lacks explicit mechanisms for capturing the structural and contextual characteristics of negation. As a result, its output features often fail to distinguish the semantic differences inherent in negated expressions. To address this limitation, we propose the Lens module, a syntax-semantic dual-stream architecture that explicitly models negation-aware text representations via hierarchical attention.

Syntactic Stream. Natural language negation frequently manifests through syntactic patterns (e.g., auxiliary constructions like “do not”, adverbials like “never”) that exhibit local structural dependencies. To capture these cues, we

extract features T_1, T_2, T_3 from the first three layers of the CLIP text encoder, which encode low-level syntactic information (Quantmeyer, Mosteiro, and Gatt 2024). Each layer’s features are projected into a shared latent space via the transformation:

$$\mathcal{P}(X) = \text{LN}(\text{GELU}(WX)) \quad (1)$$

where W denotes a learnable weight matrix unique to each projection instance, GELU is the Gaussian Error Linear Unit (Hendrycks and Gimpel 2016), and LN represents Layer Normalization (Ba, Kiros, and Hinton 2016). The resulting query vectors are:

$$Q_i = \mathcal{P}_{q,i}(T_i), \quad i \in \{1, 2, 3\} \quad (2)$$

Semantic Stream. While syntactic structures signal negation expression, precise interpretation requires global semantic comprehension. For instance, in the sentence “*He claimed he finished the work, but actually hadn’t*”, the negation scope of “*hadn’t*” depends critically on the antecedent context (“*finished the work*”) and the contrastive conjunction (“*but*”). To capture such context-dependent negation, we utilize the final-layer output T_{clip} of the CLIP text encoder (Lin et al. 2024; Jing et al. 2024), which provides globally contextualized representations for generating:

$$K = \mathcal{P}_k(T_{\text{clip}}), \quad V_i = \mathcal{P}_{v,i}(T_{\text{clip}}), \quad i \in \{1, 2, 3\} \quad (3)$$

Hierarchical Attention Fusion. Syntactic representations from different layers encode linguistic features at varying granularities: Q_1 captures local token interactions (e.g., negation particles and adjacent verbs), while the deeper Q_3 encodes phrasal-level dependencies. To dynamically integrate these hierarchical patterns with global semantics, we compute attention weights between each Q_i and the semantic key K , modulated by learnable scalars α_i that adaptively calibrate each syntactic level’s contribution:

$$T_{\text{attn}} = \sum_{i=1}^3 \text{softmax} \left(\frac{Q_i K^\top}{\sqrt{D}} + \log \alpha_i \right) V_i \quad (4)$$

Residual Gating. While hierarchical attention enriches the representation with negation-relevant structure, over-reliance on it may cause semantic drift or loss of core content. To address this, we apply a residual gate that adaptively blends the attended representation with the original CLIP features:

$$g = \sigma(W_g T_{\text{attn}} + b_g) \quad (5)$$

$$T_{\text{neg}} = \text{FFN}(g \odot T_{\text{attn}} + (1 - g) \odot T_{\text{clip}}) \quad (6)$$

where σ is the sigmoid and \odot denotes Hadamard multiplication. The gate g acts as a soft selector, allowing the model to amplify structural adjustments only when necessary. The bias b_g is initialized to a negative value to favor original features during early training. Finally, a feed-forward neural network (FFN) refines the output representation.

Frame: Cross-Modal Dynamic Repulsion Weight Generator

Negation in language varies in intensity and contextual nuance (e.g., “not” vs. “may not”), directly affecting the degree

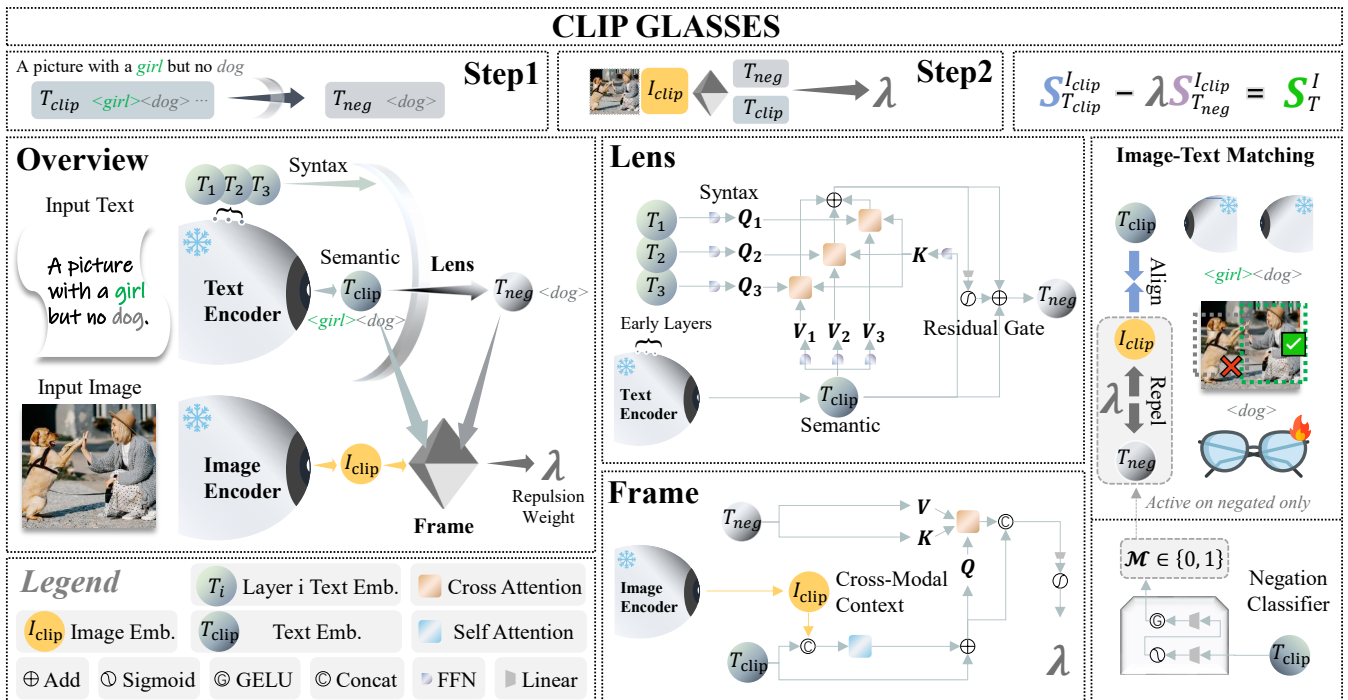


Figure 3: CLIPGLASSES enhances CLIP’s capability to model negative semantics by introducing two modules: *Lens* and *Frame*. *Lens* disentangles negated concepts (e.g., “dog” in “no dog”) from the text embedding T_{clip} . *Frame* dynamically predicts a repulsion strength λ based on cross-modal context. The final similarity score is computed as $S = S_{I2T} - \lambda \cdot S_{I2T}^{\text{neg}}$, aligning images with affirmed content while repelling from negated concepts when negation is present in the text.

to which negated concepts influence vision-language alignment. To address this, the *Frame* module dynamically estimates a repulsion weight λ based on joint image-text representations, which governs the strength of semantic inversion. **Cross-Modal Context.** Negation is inherently linguistic, yet its interpretation is often grounded in visual context (Sun et al. 2021; Janssens et al. 2024). To estimate λ in a context-aware manner, we first encode cross-modal interactions by allowing each modality to attend to the other.

Before fusing the textual and visual features, we apply L2 normalization to eliminate scale discrepancies that could hinder effective attention. For numerical stability, a small constant ϵ is added in the denominator when normalizing the negated text feature:

$$\hat{I}_{\text{clip}} = \frac{I_{\text{clip}}}{\|I_{\text{clip}}\|_2}, \quad \hat{T}_{\text{clip}} = \frac{T_{\text{clip}}}{\|T_{\text{clip}}\|_2}, \quad \hat{T}_{\text{neg}} = \frac{T_{\text{neg}}}{\|T_{\text{neg}}\|_2 + \epsilon} \quad (7)$$

To capture bidirectional cross-modal dependencies, we employ a joint self-attention mechanism (Vaswani et al. 2017) over concatenated text and image features. Compared to directional cross-attention, this symmetric early fusion enables both modalities to attend to each other jointly, avoiding representational bias. Formally:

$$F_{T\&I} = \text{SelfAttn} \left([\hat{T}_{\text{clip}}; \hat{I}_{\text{clip}}] \right) \in \mathbb{R}^{2 \times D} \quad (8)$$

Since λ modulates textual semantics, its estimation should be grounded in a representation that is text-centric yet contextually enriched by visual cues. We thus retain only the

text-side output $F_{T\&I}^{(1)}$ from the joint self-attention. To preserve original semantics while enabling adaptive fusion, we apply a residual connection:

$$T_{\text{fuse}} = \alpha \hat{T}_{\text{clip}} + F_{T\&I}^{(1)} \quad (9)$$

where $\alpha \in [0, 1]$ is a learnable scalar (initialized to 0.1) that controls the balance between the original and the context-enhanced text features.

Dynamic Repulsion Weight. Estimating the repulsion strength λ hinges on highlighting negated semantics relevant to the fused text-visual context. Since negation inherently involves contrasting the original semantics with negated concepts, a mechanism that attends to the negated features conditioned on the enriched textual context is essential.

To this end, we first compute a cross-attention output C , where the fused representation T_{fuse} acts as the query, and the negated semantics \hat{T}_{neg} serve as keys and values. This allows the model to dynamically align and weight the negated information relevant to the current context:

$$C = \text{CrossAttn} \left(\underbrace{T_{\text{fuse}}}_Q, \underbrace{\hat{T}_{\text{neg}}}_K, \underbrace{\hat{T}_{\text{neg}}}_V \right) \quad (10)$$

Next, the repulsion weight λ is estimated by projecting the concatenation of the fused feature T_{fuse} and the cross-attention output C through a learnable linear layer followed by a sigmoid activation:

$$\lambda = \sigma (W_\lambda [T_{\text{fuse}}; C] + b_\lambda) \quad (11)$$

Image-Text Matching

With the negated text representation T_{neg} generated by the *Lens* module and the context-sensitive repulsion weight λ estimated by the *Frame* module, we now describe how these components jointly contribute to the final image-text matching objective. The core idea is to preserve the original alignment ability of CLIP while penalizing false positives caused by visual-textual agreement with negated semantics.

The final similarity score thus combines a standard matching term with a negation-aware repulsion component. For generality across both text-to-image and image-to-text retrieval, we adopt a unified notation where q denotes the query modality (either text or image) and k denotes the key modality (the corresponding paired instance).

Basic Similarity Calculation. We begin with the original CLIP similarity as the base score. This term reflects the semantic alignment between the image and the text:

$$S_{\text{base}}(q, k) = \exp(\theta_T) \cdot \frac{q^\top k}{\|q\|_2 \|k\|_2} \quad (12)$$

where θ_T is the temperature parameter of the pretrained CLIP model.

Negation-aware Similarity Adjustment. While the base similarity captures standard semantic alignment, it is insufficient for handling negation. If the text explicitly negates certain visual content, the model must reduce the matching score accordingly. To this end, we introduce a negation-aware repulsion term that reflects how well the image aligns with the negated semantics. The repulsion term is computed by:

$$\mathcal{R}_{\text{neg}} = \lambda \cdot \max\left(\exp(\theta_T) \cdot \frac{q^\top k_{\text{neg}}}{\|q\|_2 \|k_{\text{neg}}\|_2}, 0\right) \quad (13)$$

where the $\max(\cdot, 0)$ operation ensures that the repulsion term only penalizes positive alignment with negated concepts, preventing undesirable score inflation.

Final Similarity Score. To avoid unnecessary interference in affirmative cases and better preserve CLIP’s performance on general (non-negated) tasks, we further introduce a conditional mask that selectively activates the repulsion term based on the presence of negation.

Concretely, we define a binary decision variable $\mathcal{M} \in \{0, 1\}$ that determines whether the similarity score should be corrected. The decision is made by a lightweight negation classifier $\mathcal{G} : \mathbb{R}^D \rightarrow [0, 1]$, which predicts the probability of negation from the original text representation:

$$\mathcal{M} = \mathbb{I}[\mathcal{G}(T_{\text{clip}}) > \tau_{\text{neg}}] \quad (14)$$

where τ_{neg} is a threshold (default 0.5). The classifier \mathcal{G} is implemented as a two-layer MLP.

The repulsion term \mathcal{R}_{neg} is applied only when $\mathcal{M} = 1$, ensuring that the correction is restricted to negated inputs while maintaining standard CLIP behavior elsewhere.

The final similarity score integrates the base CLIP similarity with the negation-aware repulsion term under the control of the conditional mask:

$$S = S_{\text{base}} - \mathcal{M} \cdot \mathcal{R}_{\text{neg}} \quad (15)$$

To ensure stable optimization, the mask is fixed to $\mathcal{M} = 1$ during training, allowing gradients to propagate consistently. It is only activated at inference time to selectively suppress negated alignments.

Training Strategy

We adopt a staged training strategy to progressively model cross-modal negation semantics. The process consists of three phases: (1) training the *Lens* module to capture fine-grained representations of negated text; (2) training the *Frame* module to establish dynamic interactions between negation and image features; and (3) joint optimization to enhance collaboration between the two modules.

Stage 1: Independent Training of the Lens Model. In this stage, we extract the negated object from the input text and generate a short prompt (e.g., “This image shows a {negobj}”). This prompt is then passed to the CLIP text encoder to extract the ground truth negation feature $T_{\text{obj}}^{\text{neg}} \in \mathbb{R}^D$ as the supervision signal.

We design a training objective that combines the semantic similarity loss \mathcal{L}_{sim} and the cross-modal alignment loss $\mathcal{L}_{\text{align}}$ to ensure the learned negation features are both accurate and aligned with image features:

$$\mathcal{L} = \mathcal{L}_{\text{sim}} + \delta \mathcal{L}_{\text{align}} \quad (16)$$

The semantic similarity loss \mathcal{L}_{sim} ensures semantic consistency between the predicted negation feature T_{neg} and the ground truth negation feature $T_{\text{obj}}^{\text{neg}}$:

$$\mathcal{L}_{\text{sim}} = 1 - \frac{1}{B} \sum_{j=1}^B \frac{T_{\text{neg}}^{(j)} \cdot T_{\text{obj}}^{\text{neg},(j)}}{\|T_{\text{neg}}^{(j)}\|_2 \|T_{\text{obj}}^{\text{neg},(j)}\|_2} \quad (17)$$

The cross-modal alignment loss $\mathcal{L}_{\text{align}}$ constrains the alignment between the negation feature T_{neg} and the image feature I , preventing semantic drift:

$$\mathcal{L}_{\text{align}} = \frac{1}{B} \sum_{j=1}^B \left| \frac{T_{\text{neg}}^{(j)} \cdot I^{(j)}}{\|T_{\text{neg}}^{(j)}\|_2 \|I^{(j)}\|_2} - \frac{T_{\text{obj}}^{\text{neg},(j)} \cdot I^{(j)}}{\|T_{\text{obj}}^{\text{neg},(j)}\|_2 \|I^{(j)}\|_2} \right| \quad (18)$$

Additionally, we implement a dynamic loss balancing strategy, starting with a higher emphasis on semantic similarity ($\delta = 0.5$) and gradually shifting weight toward the cross-modal constraint ($\delta = 1.0$) to stabilize negation representation learning before enforcing image alignment.

Stage 2: Independent Training of the Frame. In this stage, the ground-truth negation feature $T_{\text{obj}}^{\text{neg}} \in \mathbb{R}^D$ is directly used as input to the *Frame* module, serving as the predicted negation feature from the *Lens*. Training is based on a generalized InfoNCE loss:

$$\mathcal{L}_{\text{gen}}(q, k^+, \mathcal{N}) = -\log \frac{e^{S(q, k^+)}}{e^{S(q, k^+)} + \sum_{k^- \in \mathcal{N}} e^{S(q, k^-)}} \quad (19)$$

where q denotes the query feature, k^+ the corresponding positive key, and \mathcal{N} the set of negative keys k^- . The similarity function $S(\cdot, \cdot)$ is defined as in Eq. (15).

Method	Training Set	Training Size	Testing			
		(Images/Captions)	Testing Set	Accuracy	Testing Set	Accuracy
NegCLIP	COCO-caption	82K/31K	CC-Neg-val	62.63%	Neg-COCO-MCQ	10.20%
CoN-CLIP	CC-Neg	188K/376K	CC-Neg-val	99.70%	Neg-COCO-MCQ	25.70%
CLIPGlasses	CC-Neg	188K/376K	CC-Neg-val	96.56% ^{-3.14%}	Neg-COCO-MCQ	34.51% ^{+8.81%}
CoN-CLIP	Neg-COCO-R	5K/22K	CC-Neg-val	65.91%	Neg-COCO-MCQ	30.61%
CLIPGlasses	Neg-COCO-R	5K/22K	CC-Neg-val	93.36% ^{+27.45%}	Neg-COCO-MCQ	35.90% ^{+5.29%}

Table 1: Performance comparison highlighting the trade-off between in-domain fitting and cross-domain generalization.

Depending on the availability of hard negatives, the loss is applied in two ways. When explicit hard negatives are available (e.g., from the CCNeg dataset), training is performed using constructed triplets (q, k^+, k^-) . Otherwise, in-batch contrastive learning is adopted, treating all other batch samples as implicit negatives.

Stage 3: Joint Training of Lens and Frame. After training the Lens and Frame separately, we further design a joint training process to optimize their collaborative effect. Specifically, based on the training strategy of Stage 2, the ground-truth negated object feature $T_{obj}^{neg} \in \mathbb{R}^D$ in the input of the Frame is replaced with the output T_{neg} of the Lens module, while the rest of the settings remain the same as in Stage 2.

Comparative Experiments

To validate our method, we conduct systematic comparisons with state-of-the-art baselines, evaluate the retention of CLIP’s zero-shot capabilities, and perform ablations to assess the contributions of individual components.

Comparative Analysis

Existing work typically fine-tunes CLIP’s text encoder without architectural changes. Given their structural homogeneity—differing primarily in training data—we focus on two fully open-source, paradigm-representative baselines: Neg-CLIP (Yuksekgonul et al. 2022) and CoN-CLIP (Singh et al. 2025). Comparative results are shown in Table 1.

Trained on the CC-Neg dataset (Singh et al. 2025) (188K images, 376K captions), CLIPGlasses achieves 96.56% accuracy on the in-domain CC-Neg-val set. While slightly below CoN-CLIP’s 99.70%, this reflects a deliberate design choice to prioritize generalization over overfitting. The benefit becomes evident on the cross-domain Neg-COCO-MCQ benchmark (Alhamoud et al. 2025), where CLIPGlasses surpasses CoN-CLIP by 8.81 percentage points (34.51% vs. 25.70%).

Under low-resource conditions using the Neg-COCO-R subset (5K images, 22K captions; Alhamoud et al. 2025), this advantage becomes more pronounced: CLIPGlasses outperforms CoN-CLIP by 27.45 points on CC-Neg-val (93.36% vs. 65.91%) and by 5.29 points on Neg-COCO-MCQ (35.90% vs. 30.61%).

These results highlight two core limitations of fine-tuning-based approaches: reliance on large-scale data, as

Method	Training Set	ImageNet	caltech101
Vanilla CLIP	-	53.87	<u>90.96</u>
CoN-CLIP	CC-Neg	50.98	88.91
CLIPGLASSES	Neg-COCO-R	<u>53.51</u>	90.54
CLIPGLASSES	CC-Neg	53.28	90.97

Table 2: Zero-shot performance on standard non-negation benchmarks for models trained on negation datasets.

seen in CoN-CLIP’s drop under constraints, and the absence of explicit negation modeling, as reflected in NegCLIP’s persistently low performance. In contrast, our non-invasive architecture maintains consistent performance across settings, validating its strength in capturing transferable, semantically grounded negation.

Inherent Ability Retention Analysis

A critical concern in enhancing negation understanding is whether such improvements come at the cost of a model’s inherent strengths. To assess this, we compare zero-shot performance on standard non-negation benchmarks (ImageNet (Deng et al. 2009) and Caltech101 (Li et al. 2022)) after training on negation datasets.

As shown in Table 2, CLIPGLASSES retains near-original performance, matching or even surpassing the vanilla CLIP on both datasets. In contrast, the CoN-CLIP shows notable degradation, particularly on ImageNet. These results confirm that our non-invasive design preserves CLIP’s general visual-language alignment ability while enhancing negation understanding.

Ablation Analysis

To assess the individual contributions of key components, we conduct hierarchical ablation experiments on the CC-Neg dataset. Using the full model as our reference standard, we evaluate two key metrics: classification accuracy (Accuracy) and the average confidence margin between positive and negative pairs, termed false alignment rate (FAR), which is defined as:

$$\text{FAR} = \frac{1}{N} \sum_{i=1}^N |s_i^+ - s_i^-| \quad (20)$$

where s_i^+ and s_i^- denote the similarity scores between the image and the positive and negative captions in the i -th sam-

Ablation Setting	Acc	Δ Acc	FAR	Δ FAR
- <i>Lens Module</i> -				
w/o Syntax Stream	94.09%	-2.47%	4.4927	-3.3326
w/o Semantic Stream	94.93%	-1.63%	6.0925	-1.7328
w/o Residual Gating	68.93%	-27.63%	0.7875	-7.0378
- <i>Frame Module</i> -				
w/o Cross-modal	93.79%	-2.77%	7.3051	-0.5202
w/o Repulsion weight	63.74%	-32.82%	0.5684	-7.2569
Full Model	96.56%	-	7.8253	-

Table 3: Ablation on CC-Neg evaluating the *Lens* and *Frame* modules. “Acc” and “FAR” refer to Accuracy and false alignment rate respectively, Δ denotes the performance drop relative to the full model.

ple, and N is the total number of samples. Higher FAR indicates stronger discriminative ability between affirmations and negations. The results of the ablation experiments are shown in Table 3.

Effect of Syntactic Stream. Removing the syntactic stream results in a 2.47% decrease in accuracy and a more substantial 3.33-point decline in FAR, demonstrating its critical role in enhancing the model’s structural sensitivity to negation. In particular, the reduction in FAR indicates diminished confidence in negation detection, underscoring the essential role of syntactic grounding in achieving reliable alignment.

Effect of Semantic Stream. Excluding the semantic stream results in a 1.63% accuracy drop and a 1.73-point reduction in FAR. These results indicate that while syntactic patterns effectively locate negation cues, accurate interpretation requires access to global semantic context. The semantic stream, derived from the final CLIP layer, enables the model to identify the target of negation by capturing sentence-level meaning. Without this semantic grounding, the model exhibits reduced performance in disambiguating negation, particularly when polarity lacks explicit syntactic markers.

Effect of Residual Gating. Disabling the residual gating leads to a substantial 27.63% drop in accuracy and a 7.04-point reduction in FAR, demonstrating severe degradation in both classification and discriminative confidence. This validates the critical importance of balancing syntactic attention with the original semantic representation. Without residual gating, the model becomes susceptible to overfitting on structurally salient yet semantically irrelevant patterns, particularly in linguistically ambiguous contexts. The gating mechanism enables the model to preserve core sentence meaning when structural alignment is uncertain, thereby enhancing the robustness and stability of negation reasoning.

Effect of Cross-modal Context. Disabling cross-modal interaction results in a 2.77% accuracy decline and a 0.52-point reduction in FAR. Under this configuration, the repulsion weight λ is estimated solely from textual features, thereby limiting the model’s ability to adjust repulsion strength based on visual content—a critical component for context-aware negation understanding.

Effect of Dynamic Repulsion Weight. We observe substantial performance degradation upon removing this module

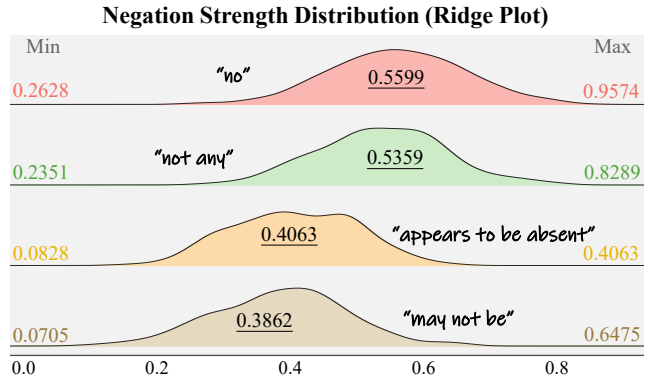


Figure 4: Distribution of predicted repulsion weight λ under varying negation strengths. Stronger negations (e.g., “no”) yield higher λ , confirming the model’s ability to adaptively modulate semantic repulsion.

(32.82% accuracy decrease) underscores its critical importance. This ablation result motivates a deeper investigation into the underlying mechanisms driving its contribution. We hypothesize that the module’s primary function is to dynamically calibrate repulsion strength according to the linguistic intensity of negation. To test this hypothesis, we evaluate the predicted λ values across a controlled spectrum of negation intensities. Using Qwen2.5-72B (Yang et al. 2024), we generate four recaptioned variants for 500 randomly sampled CC-Neg examples, each corresponding to four distinct levels of negation strength: strong (“no”), moderate (“not any”), weak (“appears to be absent”), and weakest (“may not be”).

As shown in Figure 4, the predicted λ distributions exhibit a consistent decreasing trend aligned with the reduction in negation strength. These results demonstrate that the module adaptively modulates repulsion intensity in response to linguistic cues.

Conclusion

In this paper, we present CLIPGLASSES, a novel framework that addresses CLIP’s limitations in negation understanding. Our non-intrusive design introduces cognitively inspired modules: a *Lens* for disentangling negated semantics and a *Frame* for modeling context-aware repulsion, together enabling negation-sensitive similarity computation without modifying CLIP’s pretrained parameters. Experiments show that our method achieves competitive in-domain accuracy, state-of-the-art cross-domain generalization and low-resource robustness, all while preserving CLIP’s native zero-shot capabilities. Nevertheless, a shared limitation across current methods remains in handling non-visual negations (e.g., “not authentic”). Future work will explore integrating commonsense knowledge to address such cases.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (General Program, No. 62377024, 2024–2027).

References

- Alhamoud, K.; Alshammari, S.; Tian, Y.; Li, G.; Torr, P.; Kim, Y.; and Ghassemi, M. 2025. Vision-language models do not understand negation. *arXiv preprint arXiv:2501.09425*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Caffagni, D.; Sarto, S.; Cornia, M.; Baraldi, L.; and Cucchiara, R. 2025. Recurrence-Enhanced Vision-and-Language Transformers for Robust Multimodal Document Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cai, T. T.; and Ma, R. 2022. Theoretical Foundations of t-SNE for Visualizing High-Dimensional Clustered Data. *Journal of Machine Learning Research*, 23(301): 1–54.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Fang, S.; Wang, Y.; Zhang, J.; Li, Z.; and Wang, Y. 2025. Adaptive Language-Aware Image Reflection Removal Network. In Kwok, J., ed., *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, 973–981. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Fang, X.; Liu, D.; Fang, W.; Zhou, P.; Xu, Z.; Xu, W.; Chen, J.; and Li, R. 2024. Fewer steps, better performance: Efficient cross-modal clip trimming for video moment retrieval using language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1735–1743.
- Feng, Z.; Guo, Q.; Xiao, X.; Xu, R.; Yang, M.; and Zhang, S. 2025. Unified Video Generation via Next-Set Prediction in Continuous Domain. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19427–19438.
- Geirhos, R.; Jacobsen, J.-H.; Michaelis, C.; Zemel, R.; Brendel, W.; Bethge, M.; and Wichmann, F. A. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11): 665–673.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Janssens, R.; Wolfert, P.; Demeester, T.; and Belpaeme, T. 2024. Integrating visual context into language models for situated social conversation starters. *IEEE Transactions on Affective Computing*, 16(1): 223–236.
- Jha, S.; Gong, D.; and Yao, L. 2024. Clap4clip: Continual learning with probabilistic finetuning for vision-language models. *Advances in neural information processing systems*, 37: 129146–129186.
- Jiang, Z.; Xu, J.; Zhang, S.; Shen, T.; Li, J.; Kuang, K.; Cai, H.; and Wu, F. 2025. FedCFA: Alleviating Simpson’s Paradox in Model Aggregation with Counterfactual Federated Learning. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 17662–17670.
- Jing, D.; He, X.; Luo, Y.; Fei, N.; Wei, W.; Zhao, H.; Lu, Z.; et al. 2024. Fineclip: Self-distilled region-based clip for better fine-grained understanding. *Advances in Neural Information Processing Systems*, 37: 27896–27918.
- Kaup, B.; Zwaan, R. A.; and Lüdtke, J. 2007. The experiential view of language comprehension: How is negation represented. *Higher level language processes in the brain: Inference and comprehension processes*, 255–288.
- Kim, T.; Lee, S.; Kim, S.-W.; and Kim, D.-J. 2025. Vip-cap: Retrieval text-based visual prompts for lightweight image captioning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Ko, H.; and Park, C.-M. 2025. Bringing CLIP to the Clinic: Dynamic Soft Labels and Negation-Aware Learning for Medical Analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 25897–25906.
- Koishigarina, D.; Uselis, A.; and Oh, S. J. 2025. CLIP Behaves like a Bag-of-Words Model Cross-modally but not Uni-modally. *arXiv preprint arXiv:2502.03566*.
- Li, B.; Dong, H.; Zhang, D.; Zhao, Z.; Gao, J.; and Li, X. 2025a. Exploring Efficient Open-Vocabulary Segmentation in the Remote Sensing. *arXiv preprint arXiv:2509.12040*.
- Li, B.; Zhang, D.; Zhao, Z.; Gao, J.; and Li, X. 2025b. Stitchfusion: Weaving any visual modalities to enhance multimodal semantic segmentation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 1308–1317.
- Li, F.-F.; Andreeto, M.; Ranzato, M.; and Perona, P. 2022. Caltech 101.
- Li, H.; Ding, L.; Fang, M.; and Tao, D. 2024. Revisiting catastrophic forgetting in large language model tuning. *arXiv preprint arXiv:2406.04836*.
- Li, K.; Jiang, Z.; Shen, Z.; Wang, Z.; Lv, C.; Zhang, S.; Wu, F.; and Wu, F. 2025c. MadaKV: Adaptive Modality-Perception KV Cache Eviction for Efficient Multimodal Long-Context Inference. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, 13306–13318. Association for Computational Linguistics.
- Lin, Y.; Chen, M.; Zhang, K.; Li, H.; Li, M.; Yang, Z.; Lv, D.; Lin, B.; Liu, H.; and Cai, D. 2024. Tagclip: A local-to-global framework to enhance open-vocabulary multi-label classification of clip without training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3513–3521.
- Ma, Z.; Hong, J.; Gul, M. O.; Gandhi, M.; Gao, I.; and Krishna, R. 2023. Crepe: Can vision-language foundation models reason compositionally? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Morante, R.; and Blanco, E. 2021. Recent advances in processing negation. *Natural Language Engineering*, 27(2): 121–130.
- Nam, J.; Im, J.; Kim, W.; and Kil, T. 2025. Extract free dense misalignment from CLIP. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.

- Orenes, I.; Beltrán, D.; and Santamaría, C. 2014. How negation is understood: Evidence from the visual world paradigm. *Journal of memory and language*, 74: 36–45.
- Park, J.; Lee, J.; Song, J.; Yu, S.; Jung, D.; and Yoon, S. 2025. Know” No” Better: A Data-Driven Approach for Enhancing Negation Awareness in CLIP. *arXiv preprint arXiv:2501.10913*.
- Qin, X.; Zhang, P.; Yang, J. J. O.; Zeng, G.; Li, Y.; Wang, Y.; Zhang, W.; and Dai, P. 2025. CLIP is Almost All You Need: Towards Parameter-Efficient Scene Text Retrieval without OCR. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Quantmeyer, V.; Mosteiro, P.; and Gatt, A. 2024. How and where does CLIP process negation? *arXiv preprint arXiv:2407.10488*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on Machine Learning (ICML)*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Singh, J.; Shrivastava, I.; Vatsa, M.; Singh, R.; and Bharati, A. 2025. Learning the Power of ”No”: Foundation Models with Negations. In *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*.
- Su, X.; Mao, Q.; Wu, Z.; Lin, X.; You, S.; Liao, Y.; and Xu, C. 2025. Large language models driven neural architecture search for universal and lightweight disease diagnosis on histopathology slide images. *npj Digital Medicine*, 8(1): 682.
- Sun, L.; Wang, J.; Zhang, K.; Su, Y.; and Weng, F. 2021. RpBERT: a text-image relation propagation-based BERT model for multimodal NER. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 13860–13868.
- Sun, Z.; Fang, Y.; Wu, T.; Zhang, P.; Zang, Y.; Kong, S.; Xiong, Y.; Lin, D.; and Wang, J. 2024. Alpha-clip: A clip model focusing on wherever you want. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13019–13029.
- Thrush, T.; Jiang, R.; Bartolo, M.; Singh, A.; Williams, A.; Kiela, D.; and Ross, C. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Xing, F.; Li, M.; Wang, Y.-G.; Zhu, G.; and Cao, X. 2024. Clipvqa: Video quality assessment via clip. *IEEE Transactions on Broadcasting*.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2024. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Yuksekgonul, M.; Bianchi, F.; Kalluri, P.; Jurafsky, D.; and Zou, J. 2022. When and why vision-language models behave like bags-of-words, and what to do about it? *arXiv preprint arXiv:2210.01936*.
- Zhai, Y.; Tong, S.; Li, X.; Cai, M.; Qu, Q.; Lee, Y. J.; and Ma, Y. 2024. Investigating the catastrophic forgetting in multimodal large language model fine-tuning. In *Conference on Parsimony and Learning*, 202–227. PMLR.
- Zhang, S.; Xu, Y.; Usuyama, N.; Xu, H.; Bagga, J.; Tinn, R.; Preston, S.; Rao, R.; Wei, M.; Valluri, N.; et al. 2023. Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. *arXiv preprint arXiv:2303.00915*.
- Zhang, T.; Liu, P.; Lu, Y.; Cai, M.; Zhang, Z.; Zhang, Z.; and Zhou, Q. 2025a. Cwnet: Causal wavelet network for low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8789–8799.
- Zhang, T.; Liu, P.; Zhang, Z.; and Zhou, Q. 2025b. SPJFNet: Self-Mining Prior-Guided Joint Frequency Enhancement for Ultra-Efficient Dark Image Restoration. *arXiv preprint arXiv:2508.04041*.
- Zhang, T.; Liu, P.; Zhong, Z.; Zhang, Z.; and Zhou, Q. 2025c. Beyond Illumination: Fine-Grained Detail Preservation in Extreme Dark Image Restoration. *arXiv preprint arXiv:2508.03336*.
- Zhang, X.; Li, J.; Zhang, J.; Dang, Z.; Ren, J.; Bo, L.; and Tu, Z. 2025d. SemTalk: Holistic Co-speech Motion Generation with Frame-level Semantic Emphasis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13761–13771.
- Zhang, X.; Li, J.; Zhang, J.; Ren, J.; Bo, L.; and Tu, Z. 2025e. EchoMask: Speech-Queried Attention-based Mask Modeling for Holistic Co-Speech Motion Generation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 10827–10836.
- Zuanazzi, A.; Ripollés, P.; Lin, W. M.; Gwilliams, L.; King, J.-R.; and Poeppel, D. 2024. Negation mitigates rather than inverts the neural representations of adjectives. *PLoS biology*, 22(5): e3002622.