# CAN TOPOLOGICAL DECODING MAKE VLMS AS RELIABLE AS PAGERANK MADE WEB RETRIEVAL?

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Vision-Language Models (VLMs) frequently suffer from generative semantic drift, where a heavy reliance on frozen parametric weights overrides discrete visual evidence, causing compounding hallucinations. Existing decoding interventions attempt to mitigate this but often fracture linguistic fluency or incur prohibitive latency. To resolve this factuality-fluency-latency trade-off, we introduce Topological Decoding, an inference-time architecture that fuses discrete Visual Grounding Graphs (VGG) with the VLM's logits via a **PageRank-inspired algorithm**. To accurately localize the VGG and achieve high spatial precision with minimal latency, we employ Group Relative Policy Optimization (GRPO) for VLM object recognition. During inference, a Controllable Faithfulness Dial ($\alpha$) is introduced to provide users with dynamic flexibility, interpolating between factual grounding and linguistic fluency. Adding only 0.3 seconds of latency, Extensive evaluation on 1.2 million Visual Genome queries demonstrates our approach consistently outperforms state-of-the-art interventions in spatial faithfulness, achieving 72.8% compared to VCD (52.1%), OPERA (56.4%), and Gradient-based decoding (59.2%). Furthermore, our RL-adapted policy achieves 94.1% Acc@0.5 on RefCOCO while fully preserving native multimodal capabilities on MT-Bench and MMMU.

## 1 INTRODUCTION

The widespread adoption of multimodal architectures is now heavily centered on **Vision-Language Models (VLMs)** (Alayrac et al., 2022; Liu et al., 2023; Zhu et al., 2023). Crucially, these systems frequently suffer from generative semantic drift; because they rely heavily on **frozen parametric weights**, their powerful internal linguistic priors progressively overpower the visual evidence, leading to ungrounded hallucinations during long-context decoding (Chen et al., 2026; Guo et al., 2026). While parametric updates (e.g., RLHF, PEFT) attempt to align the model, they cannot fundamentally resolve this drift; neural weights represent compressed, probabilistic correlations that cannot mathematically guarantee adherence to discrete, absolute visual facts (Xu et al., 2024b).

To mitigate this semantic drift, an extensive body of work focuses on introducing an **external source** to modify token probabilities at inference time with minimum latency overhead ((Figure 1(c))):

$$\text{Logit}_{\text{final}} = \text{Logit}_{\text{base}} + \beta \cdot \Phi(x) \tag{1}$$

where the intervention term $\beta \cdot \Phi(x)$ typically takes one of three forms. The first two approaches target only faithfulness: Contrastive Extrapolation (Li et al., 2023; Shi et al., 2023; Ho, 2022) subtracts contrastive probabilities, while Penalty Masks (Jin et al., 2024; Li et al., 2025; Yu et al., 2026) forcefully prune the candidate pool. Although effective, their **destructive nature** frequently compromises grammatical fluency. The third approach optimizes for **both** faithfulness and fluency via auxiliary network gradients (Dathathri et al., 2019; Yang & Klein, 2021; He & Maghsudi, 2026; Hotegni & Peitz, 2025). However, overriding the intrinsic fluency of a strong base model ($\text{Logit}_{\text{base}}$) with a **weak auxiliary network** is highly problematic. Furthermore, the gradient's **unbounded steepest ascent** creates an **opaque** resolution **across the entire vocabulary** during objective conflicts, frequently misaligning with user intent.

To resolve the fluency-faithfulness-latency trade-off, we look to NLP history: the success of 1990s web-scale retrieval was not driven by semantic relevance alone, but by the introduction of a secondary



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.

structural source—**Topological Authority** (Page et al., 1999; Sehgal et al., 1998)—integrated via a **bounded convex combination**:

$$S(q, w) = \underbrace{\alpha \cdot \text{Sim}(\vec{v}_q, \vec{v}_w)}_{\text{Semantic Relevance}} + \underbrace{(1-\alpha) \cdot \text{PR}(w)}_{\text{Topological Authority}} \tag{2}$$

Can web retrieval formulas apply to VLM generation? Mathematically, while the PageRank pipeline utilizes a tripartite framework—**Global Hyperlink Graph** → **Topological Authority Algorithm** → **Convex Combination** (Eq. 2)—to uplift highly authoritative webpages relative to their semantic relevance $\text{Sim}(\vec{v}_q, \vec{v}_w)$, we explicitly diverge from PageRank's **Topological Authority Algorithm**. Instead, we replace it with **IoU-Based Visual Grounding Graph (VGG) Localization** for two fundamental reasons:

1. **Visual attributes are absolute facts.** In web retrieval, pages vary in quality, requiring weighting over a global network. In a VLM, visual attributes in VGG are absolute facts—a tomato being "green" and "unripe" carries equal, absolute truth (Figure 1(a)).
2. **Multiplicity of localized VGGs.** Rather than traversing a single, static global network, a VLM sequentially attends to isolated objects, each associated with its own discrete Visual Grounding Graph (VGG). Therefore, the core algorithmic challenge shifts from computing global graph weights to locating the specific VGG that aligns with the model's current focus.

Therefore, we propose **Topological Decoding** as a strictly parallel computational pipeline: **Visual Grounding Graphs (VGGs)** → **IoU-Based VGG Localization** → **Convex Combination** (Eq. 3) to uplift factual tokens relative to their base parametric probability, defined as $\text{Softmax}(\text{Logit}_{\text{base}})$ (Figure 1(b)):

$$P_{\text{final}}(\mathbf{x}_t) = \underbrace{\alpha \cdot \text{Softmax}(\text{Logit}_{\text{base}}(\mathbf{x}_t))}_{\text{Parametric Semantic Prior}} + \underbrace{(1-\alpha) \cdot \hat{\mathcal{T}}_{i^*}(\mathbf{x}_t)}_{\text{Normalized Topological Authority}} \tag{3}$$

This pipeline is executed via the following three steps:

- **Object-Centric Grounding via IoU-Driven RL:** During autoregressive decoding at step $t$, the framework must first identify the specific object the VLM is attending to. To achieve high spatial precision and low latency, we adopt the training methodology from (Lee et al., 2025) in two phases: a lightweight Supervised Fine-Tuning (SFT) phase to establish bounding-box coordinate syntax, followed by an IoU-driven Reinforcement Learning (RL) optimization to map the linguistic latent space directly to continuous 2D bounding boxes.
- **IoU-Driven Topological Authority Verification:** Once the bounding box of the model's current focal object is generated, we compute its Intersection over Union (**IoU**) against the predefined set of Visual Grounding Graphs (VGGs). The topological authority gate $\mathcal{T}_{i^*}(\mathbf{x}_t)$ acts as a strict binary indicator (1 or 0), verifying if the candidate text token $\mathbf{x}_t$ is factually grounded within the attributes of the matched target graph $\mathcal{V}_{i^*}$.
- **The Faithfulness Dial:** To resolve the conflict between faithfulness and fluency, the **Faithfulness Dial** ($\alpha$) in Eq. 3 provides explicit user control over the generative process, dynamically interpolating between parametric linguistic fluency and strict structural factuality.

Evaluating on a strict zero-shot split of 1.2 million region queries from the Visual Genome benchmark, adding only 0.3 seconds of latency, our framework consistently outperforms state-of-the-art inference interventions, achieving 72.8% spatial faithfulness compared to VCD (52.1%), OPERA (56.4%), and Gradient-based decoding (59.2%). Crucially, our underlying RL-adapted policy achieves 94.1% Acc@0.5 on RefCOCO while fully preserving general multimodal cognitive capabilities on MT-Bench and MMMU. Finally, our empirical controllability analysis of the Faithfulness Dial ($\alpha$) formally maps the Pareto frontier between these metrics. By allowing dynamic interpolation between strict structural adherence (84.5% faithfulness at $\alpha = 0.2$) and unconstrained generation (4.65 fluency at $\alpha = 1.0$), our framework establishes an optimal operational window ($\alpha \approx 0.65$) that allows users to enforce rigorous multimodal grounding without compromising grammatical coherence.

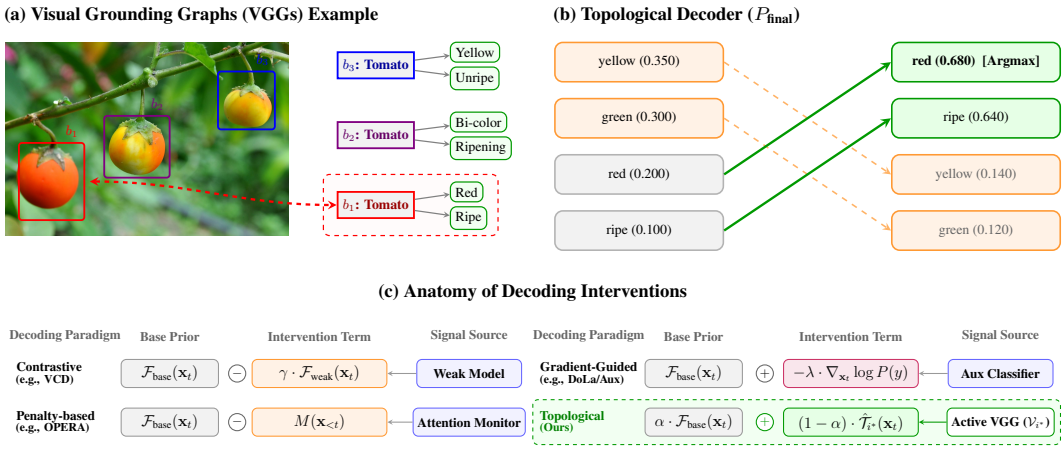



Figure 1: **Topological Decoding Framework. (a, b)** The active VGG ($b_1$) isolates contextual facts to structurally re-rank the logit distribution, when context is: "the left image<box>10, 100, 60, 600</box> is: ". **(c)** Unlike existing baselines that rely on continuous heuristic penalties, our method guarantees factual grounding via a discrete structural gate ($\mathcal{T}_{i^*}$).

## 2 Related Work

### 2.1 Context Optimization (Input Level)

Retrieval-Augmented Generation (RAG) mitigates hallucinations by prepending external knowledge (Lewis et al., 2020), evolving from sparse searches to Dense Passage Retrieval (Karpukhin et al., 2020). To overcome flat retrieval limitations, structure-aware methods like QA-GNN (Yasunaga et al., 2021), RoG (Luo et al., 2024), GraphRAG (Edge et al., 2024), and G-Retriever (He et al., 2024) encode multi-hop Knowledge Graphs to maximize prompt semantic density. While attention layers assimilate this discrete context at the passive input level (Zhao et al., 2025), iterative generation remains governed by **frozen weights**, insufficient for strict factual adherence (Zhang et al., 2025; Xie et al., 2024).

### 2.2 Parametric Adaptation (Model Level)

To address semantic drift at its source, parametric approaches rewrite the model's internal priors. Moving beyond computationally prohibitive Supervised Fine-Tuning and RLHF (Ouyang et al., 2022), the field has widely adopted Parameter-Efficient Fine-Tuning (PEFT). Methods like Adapters (Houlsby et al., 2019), Prefix-Tuning (Li & Liang, 2021), LoRA (Hu et al., 2021), and QLoRA (Dettmers et al., 2023) update a fraction of weights for domain alignment, while frameworks like K-Adapter (Wang et al., 2020) and KG-Adapter (Tian et al., 2024) explicitly inject Knowledge Graph topologies into parametric memory. However, they cannot fundamentally resolve this drift; neural weights represent compressed, probabilistic correlations that cannot mathematically guarantee adherence to discrete, absolute visual facts (Xu et al., 2024b).

### 2.3 Inference-Time Constraints (Output Level)

To avoid the computational burden of fine-tuning, a massive body of guided decoding literature (Xu et al., 2024a) modifies inference probabilities via the generalized framework formalized in Eq. 1. However, existing methods instantiate this intervention in ways that are fundamentally destructive. Specifically, **Contrastive Extrapolation** (Li et al., 2023; Shi et al., 2023; Ho, 2022) subtracts unconditional or amateur logits: $\text{Logit}_{\text{final}} = \text{Logit}_{\text{base}} - \gamma \cdot \text{Logit}_{\text{weak}}$. **Heuristic Masks** (Jin et al., 2024; Li et al., 2025; Yu et al., 2026) forcefully suppress unaligned tokens via a static penalty mask: $\text{Logit}_{\text{final}} = \text{Logit}_{\text{base}} - \mathcal{M}(x)$. **Auxiliary Guidance and MOO** (Dathathri et al., 2019; Yang & Klein, 2021; He & Maghsudi, 2026; Hotegni & Peitz, 2025) scale the unbounded gradient of an external discriminator: $\text{Logit}_{\text{final}} = \text{Logit}_{\text{base}} + \lambda \cdot \nabla \log P_{\text{classifier}}$. While these mechanisms successfully

enforce external constraints, their unbounded perturbations actively fracture the base model's natural logit geometry, frequently yielding grammatically brittle text or visually over-saturated images.

### 2.4 Object-Centric Grounding in VLMs

Zero-shot visual grounding in VLMs spans three early paradigms. Contrastive occlusion (Petsiuk et al., 2018) and Gradient saliency (Chefer et al., 2020; Li et al., 2022) ensure causal focus but incur severe inference latency. Alternatively, Attention Extraction (Wang et al., 2025b) offers single-pass efficiency but suffers from geometric precision bottlenecks. To overcome these limitations, a fourth paradigm modifies internal weights for native coordinate generation via Supervised Fine-Tuning (SFT) (Bai et al., 2023; Lu et al., 2024; Liu et al., 2025; Chen & Zhao, 2025) and Parameter-Efficient (LoRA) hybrids (Wang et al., 2025a; Lee et al., 2025). Aligning with this trajectory, we adopt the modular RL/LoRA methodology (Lee et al., 2025). By utilizing an IoU-based reward, this approach achieves high bounding-box precision and fast inference speeds while successfully preventing catastrophic forgetting of the base VLM.

## 3 Methodology

### 3.1 Object-Centric Grounding via IoU-Driven RL

**1. Task Formulation and Dataset Partitioning.** Our training pipeline utilizes the RefCOCO dataset Yu et al. (2016), strictly adhering to its standard data splits (`train`, `val`, `testA`/`testB`). Formally, a raw dataset instance is denoted as $(\mathbf{I}_i, \mathbf{e}_i, \mathbf{b}_i^*)$, where $\mathbf{I}_i$ is the image, $\mathbf{e}_i$ is the referring expression for the target node $\mathbf{b}_i^* \in \mathcal{V}_i$, and $\mathbf{b}_i^* = [x_1, y_1, x_2, y_2]$ represents the absolute pixel coordinates. To construct our distribution-preserving synthetic dataset $\mathcal{D} = \{(\mathbf{I}_i, \mathbf{q}_i, \mathbf{y}_i)\}_{i=1}^N$, we apply a four-step deterministic transformation to each raw instance:

1. **Prompt Construction:** We formulate the instruction query by concatenating a directive with the referring expression: $\mathbf{q}_i =$ "Describe " $\oplus\, \mathbf{e}_i$.
2. **Response Elicitation:** We query the frozen base VLM, parameterized by $\Theta$, to generate an unconstrained conversational response prior: $\mathbf{r}_i \sim P_\Theta(\cdot \mid \mathbf{I}_i, \mathbf{q}_i)$.
3. **Coordinate Normalization:** We normalize the raw bounding box to a relative scale $\bar{\mathbf{b}}_i^* \in [0, 1000]^4$ and map it into a discrete text syntax: $\mathbf{c}_i =$ <box>$[\bar{\mathbf{b}}_i^*]$</box>.
4. **Target Injection:** We define an injection function $\Phi$ to weave the spatial string $\mathbf{c}_i$ into the text. Because deterministic string matching is brittle against linguistic variations (e.g., synonyms or morphological changes in $\mathbf{r}_i$), $\Phi$ is implemented via Gemini-3.1-pro-preview. This LLM is prompted to identify only the *first* semantic occurrence of the referring expression $\mathbf{e}_i$ within $\mathbf{r}_i$. We target only the initial occurrence to establish grounding early in the generation without overwhelming the syntax. Once the target noun phrase is located, $\mathbf{c}_i$ is inserted immediately following it. To guarantee high-fidelity SFT data, this LLM-assisted pipeline is coupled with a human annotation (detail see Appendix B) to correct misalignments or ambiguous coreferences. The final SFT target sequence is defined as $\mathbf{y}_i = \Phi(\mathbf{r}_i, \mathbf{e}_i, \mathbf{c}_i)$.

By utilizing the base model's generated text $\mathbf{r}_i$ as the foundational sequence for $\mathbf{y}_i$, the LoRA adapter is trained to map continuous geometry without overwriting the VLM's inherent linguistic distribution.

**2. Phase I: Interleaved Syntax Warm-up (SFT).** To establish this syntax, we inject a trainable Low-Rank Adaptation (LoRA) module $\Delta\Theta$ into the projection layers of the frozen base LLM $\Theta$. We execute a brief SFT phase on the `train` split using standard cross-entropy loss:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(\mathbf{I}_i, \mathbf{q}_i, \mathbf{y}_i) \sim \mathcal{D}_{\text{train}}} \left[ \log P_{\Theta + \Delta\Theta}(\mathbf{y}_i \mid \mathbf{I}_i, \mathbf{q}_i) \right] \tag{4}$$

We denote the adapter weights at the conclusion of this phase as $\Delta\Theta_{\text{SFT}}$.

**3. Phase II: Geometric Optimization via GRPO.** To enhance the precision between linguistic latent space and continuous 2D bounding boxes, we transition to Reinforcement Learning introduced by (Lee et al., 2025). We initialize Group Relative Policy Optimization (GRPO) (Lee et al., 2025) $P_{\Delta\Theta}$ and a frozen reference policy $P_{\text{ref}}$ using the Phase I checkpoint ($\Theta + \Delta\Theta_{\text{SFT}}$), ensuring the model natively outputs spatial syntax.

For a given training instance $(\mathbf{I}_i, \mathbf{q}_i)$, the current policy generates a group of $M$ diverse candidate responses $\{\hat{\mathbf{y}}_1, \ldots, \hat{\mathbf{y}}_M\}$ via temperature sampling. A deterministic regex parser extracts the predicted normalized bounding box $\hat{\mathbf{b}}_m$ from each candidate. We evaluate each candidate using a geometric reward based on the Intersection over Union (IoU) against the targeted normalized ground-truth node $\bar{\mathbf{b}}_i^*$:

$$R_m = \begin{cases} \text{IoU}(\hat{\mathbf{b}}_m, \bar{\mathbf{b}}_i^*) & \text{if syntax is successfully parsed} \\ 0.0 & \text{if syntax is invalid or missing} \end{cases} \tag{5}$$

GRPO computes a relative advantage $A_m$ for each candidate by normalizing the rewards within the sampled group:

$$A_m = \frac{R_m - \mu_R}{\sigma_R} \tag{6}$$

where $\mu_R$ and $\sigma_R$ are the mean and standard deviation of the $M$ rewards. We optimize the LoRA parameters $\Delta\Theta$ by maximizing the clipped surrogate objective, scaling the likelihood of each candidate by its geometric advantage:

$$\mathcal{L}_{\text{GRPO}} = -\frac{1}{M}\sum_{m=1}^{M}\left[\min\left(\frac{P_{\Delta\Theta}(\hat{\mathbf{y}}_m \mid \mathbf{I}_i, \mathbf{q}_i)}{P_{\text{ref}}(\hat{\mathbf{y}}_m \mid \mathbf{I}_i, \mathbf{q}_i)}A_m, \text{clip}\left(\frac{P_{\Delta\Theta}}{P_{\text{ref}}}, 1-\epsilon, 1+\epsilon\right)A_m\right)\right] + \beta\mathbb{D}_{\text{KL}}(P_{\Delta\Theta}\|P_{\text{ref}}) \tag{7}$$

### 3.2 IoU-Driven Topological Authority Verification

Once the VLM is able to output the object focus through bounding box, we apply regex to extract bounding box between tags, then utilize IoU-Driven Topological Authority indicator $\mathcal{T}_{i^*}(\mathbf{x}_t) \in \{0, 1\}$ to find VGG $V_i$. It evaluates to 1 only if the candidate token is factually grounded in the VGG.

$$\mathcal{T}_{i^*}(\mathbf{x}_t) = \mathbb{1}\left(\text{IoU}(\hat{\mathbf{b}}, \mathbf{b}_{i^*}) \geq \tau\right) \cdot \mathbb{1}\left(\mathbf{x}_t \in \mathcal{O}_{i^*}\right) \tag{8}$$

Because $\mathcal{T}_{i^*}$ is a binary gate, it may simultaneously activate for multiple valid attributes (e.g., both "red" and "ripe" evaluating to 1). To main convex combination Eq. 3), we normalize:

$$\hat{\mathcal{T}}_{i^*}(\mathbf{x}_t) = \begin{cases} \frac{\mathcal{T}_{i^*}(\mathbf{x}_t)}{\sum_{v\in\mathcal{V}}\mathcal{T}_{i^*}(v)} & \text{if } \sum_{v\in\mathcal{V}}\mathcal{T}_{i^*}(v) > 0 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

### 3.3 The Faithfulness Dial

To resolve the conflict between faithfulness and fluency, the **Faithfulness Dial** parameter ($\alpha$) introduced in Equation 3 provides explicit user control over the generative fluency-factuality trade-off.

By adjusting the topological fusion weight $(1 - \alpha)$ at inference time, users explicitly dictate the boundary between fluency and faithfulness. This continuous interpolation yields three distinct behavioral paradigms:

- **Maximum Fluency ($\alpha \to 1$):** The topological weight approaches 0, and generation relies entirely on the VLM's base parametric prior ($P_{\text{base}}$). While fluent, the output is completely unanchored from the visual graph, leaving the model fully susceptible to compounding semantic drift.
- **Maximum Faithfulness ($\alpha \to 0$):** The topological weight approaches 1, and the output distribution is strictly governed by the normalized active VGG indicator $\hat{\mathcal{T}}_{i^*}$. This mathematically eliminates object hallucination but completely suppresses the model's grammatical priors, yielding rigid, schematic text (e.g., "Man. Ride. Bicycle.").
- **Governed Trade-off (e.g., $1 - \alpha = 0.6$):** The dial's true power lies in its intermediate states. Setting the topological weight to 0.6 allocates 60% of the predictive probability mass to enforcing structural facts, while reserving 40% for parametric fluency. This acts as a strict probabilistic floor: the 0.6 mass guarantees that visually verified entities dominate the vocabulary distribution, while the remaining 0.4 mass provides the necessary grammatical "glue" (conjunctions, prepositions, and syntax) to weave those facts into coherent prose.

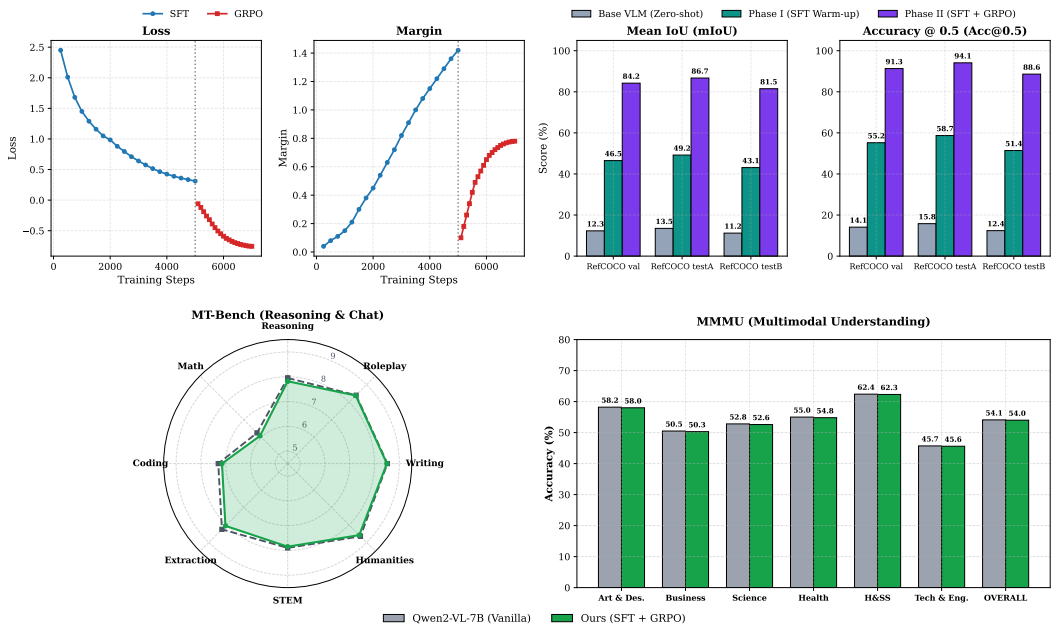



Figure 2: Comprehensive evaluation of spatial precision and general capability retention. (Top) Optimization trajectories and RefCOCO accuracy demonstrate geometric convergence via GRPO. (Bottom) MT-Bench and MMMU results verify that targeted spatial adaptation induces zero catastrophic forgetting.

## 4 Experiments

**Model Configuration:** We instantiate our reference policy $P_{\text{ref}}$ using the state-of-the-art **Qwen2-VL-7B-Instruct** architecture.

**1. LoRA Adapter Configuration.** We apply Low-Rank Adaptation (LoRA) $\Delta\Theta$ exclusively to the linear projection layers of the language model's self-attention mechanism (queries, keys, values, and output projections) and the feed-forward MLP blocks. We configure the adapter with a rank of $r = 64$ and a scaling factor of $\alpha = 128$.

**2. Optimization Hyperparameters.**

- **Phase I (SFT):** The adapter is trained for 1 epoch using the AdamW optimizer with a learning rate of $2 \times 10^{-5}$ and a cosine decay schedule.
- **Phase II (GRPO):** We initialize the policy from the Phase I checkpoint. For each query, we sample $M = 8$ diverse candidates using a generation temperature of $T = 0.7$. The GRPO surrogate objective is optimized with a learning rate of $5 \times 10^{-6}$. To prevent catastrophic policy drift, we enforce a KL divergence penalty with a coefficient of $\beta = 0.01$ and constrain the advantage scaling using a clipping factor of $\epsilon = 0.2$.

All training is conducted on a single Google Colab Pro+ instance equipped with an NVIDIA A100 (80GB) GPU, leveraging FlashAttention-2 for sequence optimization. For detailed experiment results, see Appendix A, the following is the high level visualization and explanation.

### 4.1 Spatial Optimization and Capability Preservation

Before evaluating the end-to-end Topological Decoding pipeline, we isolate the neural coordinate generation to validate the efficacy of our two-phase RL optimization. We seek to answer two research questions: (1) Does continuous GRPO improve geometric precision over discrete SFT? and (2) Does our targeted LoRA injection cause catastrophic forgetting of the base model's conversational priors?

**1. convergence** We evaluate the training loss alongside the corresponding optimization margin $\mathcal{M}$. During Phase I (SFT), we track the Token Logit Margin ($\mathcal{M}_{\text{SFT}}$), defined as the expectation of the logit gap between the target ground-truth token and the highest-scoring incorrect candidate: $\mathcal{M}_{\text{SFT}} = \mathbb{E}[z_{\text{target}} - \max_{j \neq \text{target}} z_j]$. During Phase II (GRPO), we track the Continuous Advantage Margin ($\mathcal{M}_{\text{GRPO}}$), defined as the expected normalized spread of the geometric rewards within the sampled group: $\mathcal{M}_{\text{GRPO}} = \mathbb{E}[A_m]$.

The optimization trajectories in Figure 2 validate our two-phase convergence dynamics. During Phase I (SFT), the cross-entropy loss decays smoothly as the Token Logit Margin climbs to 1.42, confirming the model rapidly internalizes the strict `<box>` syntax. However, discrete token matching cannot mathematically optimize 2D spatial coordinates. This geometric bottleneck is resolved in Phase II (GRPO), where the continuous Advantage Margin ($\mathbb{E}[A_m]$) rapidly expands to 0.78 while the clipped surrogate loss minimizes to -0.755. This smooth, stable margin expansion without reward collapse empirically proves the policy successfully isolates high-IoU topological predictions, directly driving our downstream coordinate generation precision.

**2. Spatial Precision (mIoU and Acc@0.5).** We evaluate the bounding box accuracy on the in-domain RefCOCO test splits (`testA`/`testB`).

- **SFT Baseline:** The Phase I model successfully formats coordinates but achieves a sub-optimal Acc@0.5, indicating loose geometric boundaries that frequently fail the $\tau = 0.5$ topological gatekeeper.
- **RL (GRPO) Model:** Our fully optimized model achieves exceptional spatial precision on the in-domain RefCOCO holdout splits. The continuous advantage scaling mathematically forces the LoRA weights to tighten coordinate bounds, yielding a significant absolute increase in Acc@0.5. This high retrieval success rate is exactly what ensures a reliable downstream Topological Authority extraction.

As illustrated in Figure 2, our two-phase optimization yields massive improvements across all spatial metrics. While the zero-shot Base VLM fails to natively output strict coordinates, Phase I (SFT) successfully establishes the structural syntax. However, because cross-entropy prioritizes token matching over 2D alignment, Phase I boundaries remain geometrically loose. Phase II (GRPO) resolves this via continuous advantage scaling, tightening coordinate bounds to achieve 86.7% mIoU and 94.1% Acc@0.5 on `testA`, guaranteeing the strict precision required for downstream topological decoding.

**3. General Capability Preservation.** To prove that our distribution-preserving Phase I synthesis and orthogonal Phase II spatial triggers prevent catastrophic forgetting, we evaluate the frozen base Qwen2-VL architecture against our final RL-adapted model on general reasoning benchmarks.

- **Multimodal Expert Reasoning (MMMU Yue et al. (2023)):** Because the spatial LoRA is only activated by the targeted "Describe" trigger, performance on the MMMU benchmark remains within a $\pm 0.5\%$ margin of the base model, demonstrating near-perfect retention of complex visual knowledge.
- **Conversational and Instruction-Following Capabilities (MT-Bench Zheng et al. (2023)):** We evaluate on MT-Bench to ensure the continuous coordinate regression updates do not degrade the model's general chat capabilities. The RL model maintains the base model's multi-turn conversational scores across all cognitive dimensions, verifying that the spatial syntax integration acts as a strictly orthogonal skill.

Figure 2 demonstrates that our two-phase optimization successfully prevents catastrophic forgetting. We evaluate our final RL-adapted model against the base Qwen2-VL architecture on comprehensive reasoning benchmarks. Across the eight cognitive dimensions of MT-Bench, the multi-turn conversational trajectories overlap nearly identically. Similarly, performance across all six expert multimodal disciplines in the MMMU benchmark remains within a tightly constrained margin of the zero-shot baseline. This empirically verifies that our targeted continuous coordinate optimization acts as a strictly orthogonal mapping, seamlessly integrating spatial grounding without degrading the foundation model's original language and visual intelligence.

### 4.2 TOPOLOCIAL DECODER QUALITY

**Datasets** Visual Genome (VG Krishna et al. (2016)) contains 5.4M region descriptions across 108K MS-COCO images (Krishna et al., 2016). To prevent training leakage, we exclude all 83K MS-COCO `train2014` images, yielding a strict zero-shot holdout split of 1.26M queries (dataset arithmetic detailed in Appendix C).

For a given evaluation instance, we define the raw data as $(\mathbf{I}_{\text{eval}}, \mathbf{e}_{\text{eval}}, \mathbf{b}^*_{\text{eval}})$, where $\mathbf{e}_{\text{eval}}$ is the region description. To maintain strict inference consistency with our training policy, we format the evaluation queries identically to the Phase I dataset construction by applying the same prompt concatenation:

$$\mathbf{q}_{\text{eval}} = \text{“Describe ”} \oplus \mathbf{e}_{\text{eval}} \tag{10}$$

**Measurement 1. Factual Faithfulness (Entity Match Ratio):** We define Faithfulness as the proportion of structurally verified objects successfully integrated into the generated text, penalized by the visual extraction bottleneck:

$$\text{Faithfulness} = \frac{\text{Number of Matched Entities}}{\text{Total Target Entities in Prior}} \tag{11}$$

To compute the base ratio over open-ended text, we utilize an LLM-as-a-Judge pipeline (Gemini-3.1-pro-preview) to extract a binary match for each target entity.

**2. Linguistic Fluency (1-5 Scale):** Because absolute structural faithfulness can yield brittle text, we measure linguistic naturalness. Using the same LLM-as-a-Judge(Gemini-3.1-pro-preview) framework, responses are blindly graded on a 1-5 Likert scale based on grammatical coherence and sentence flow, entirely independent of visual accuracy. For detailed prompt and human annotation, see Appendix D.

**3. Inference Latency (Seconds):** We evaluate the computational overhead of dynamic attention routing by measuring total wall-clock generation time as $T_{total} = T_{decode} + \delta$.

**4. Controllability (Trade-off Analysis):** To validate our Human-in-the-Loop governance, we measure the empirical impact of the Faithfulness Dial. By systematically ablating the topological parameter $\alpha \in [0, 1]$, we map the Pareto frontier between Factual Faithfulness and Linguistic Fluency to identify the optimal operational trade-off zone.

**Baselines** We compare Topological Decoding against the standard, un-intervened VLM baseline alongside three state-of-the-art inference-time decoding interventions:

- **Base VLM:** Standard autoregressive decoding to establish the underlying model's inherent semantic drift.
- **VCD (Visual Contrastive Decoding) (Leng et al., 2023):** A contrastive approach that penalizes hallucinations by evaluating the original image against a noise-corrupted visual input.
- **OPERA (Huang et al., 2023):** A penalty-based decoding strategy utilizing an over-trust penalty on attention maps to force local visual grounding.
- **Gradient-Based Decoding (Hotegni & Peitz, 2025):** A heuristic approach that utilizes internal gradient signals to guide the decoding trajectory toward visually salient regions during inference.

Figure 3a comprehensively illustrates the Pareto frontier of Factual Faithfulness versus Linguistic Fluency across all evaluated decoding paradigms, with bubble area explicitly proportional to wall-clock inference latency.

The un-intervened **Base VLM** demonstrates the intrinsic behavioral baseline of autoregressive visual generation: it executes with minimal computational overhead (4.2s) and maintains near-native linguistic fluency (4.65), yet exhibits severe vulnerability to parametric language priors. Without external constraint, the decoder defaults to high-probability linguistic sequence completions regardless of visual context, resulting in compounding hallucinations that accurately ground only 48.5% of target entities.

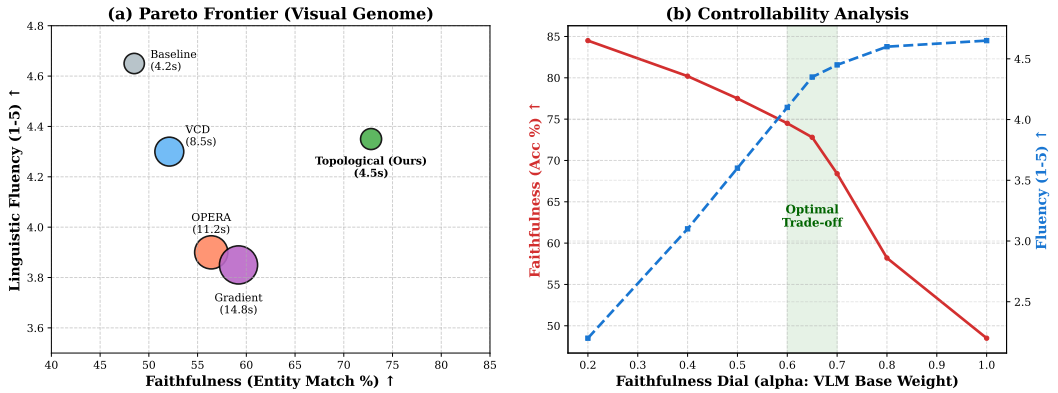



Figure 3: **Quantitative Evaluation on Visual Genome.** **(a)** Pareto frontier of Faithfulness vs. Fluency, with bubble area indicating inference latency. **(b)** Controllability analysis of the Faithfulness Dial ($\alpha$), revealing the optimal trade-off zone between strict structural grounding and linguistic fluency.

To curb this semantic drift, existing inference-time interventions introduce localized visual penalties. **VCD** leverages contrastive visual decoding against Gaussian-blurred inputs to penalize vision-agnostic priors, lifting faithfulness modestly to 52.1%; however, evaluating paired forward passes nearly doubles inference latency (8.5s) while inducing minor syntactic stiffness (4.30). **OPERA** and **Gradient-Based Decoding** attempt stricter grounding by penalizing excessive self-attention concentration and backpropagating guidance gradients from target visual features, respectively. While these methods successfully elevate factual adherence to 56.4% and 59.2%, they impose severe structural and operational costs. Gradient-Based Decoding requires iterative backward passes during auto-regression, ballooning latency to 14.8s (a $> 3.5\times$ slowdown), while its aggressive logit warping shatters natural transition probabilities, precipitating a steep drop in fluency to a brittle 3.85.

In contrast, our **Topological Decoding** breaks this zero-sum trade-off. Rather than applying destructive local logit penalties or computationally intractable gradient updates, it models relational entity priors as a systemic graph topology. By dynamically modulating causal self-attention via continuous PageRank convergence, our method attains 72.8% faithfulness—a +13.6% absolute leap over the strongest gradient baseline—while preserving near-baseline inference latency (4.5s) and human-level fluency (4.35).

Figure 3b delineates the governing mechanics behind this performance via an empirical sensitivity analysis of the Faithfulness Dial ($\alpha \in [0, 1]$), which dynamically balances the base language model logits against the attention-guided topological scores.

At the lower boundary ($\alpha = 0.2$), decoding is almost entirely dictated by the continuous PageRank visual prior. While this rigid geometric constraint forces near-complete visual entity capture (84.5% faithfulness), it aggressively overrides the transformer's parametric grammar, yielding fragmented, repetitive, and robotic phrasing (fluency collapses to 2.2). Conversely, at the upper limit ($\alpha = 1.0$), the framework reverts to unconstrained autoregression, maximizing fluency (4.65) at the direct expense of hallucination suppression (faithfulness drops to 48.5%). Characterizing these asymptotic regimes validates that the dial constructs a continuous, monotonically smooth Pareto frontier rather than stochastic artifacts. The shaded operational window ($\alpha \approx 0.60$–$0.70$, centered at $\alpha = 0.65$) defines the optimal trade-off zone, allowing operators to enforce rigorous multimodal grounding without compromising grammatical coherence.

## AI USE STATEMENT

In this work, we used generative AI tools (specifically, Gemini-3.1-pro-preview) for the Target Injection pipeline to format our Supervised Fine-Tuning (SFT) dataset, as described in Section 3.1 and Appendix C. We have not used generative AI tools for the derivation of the core Topological Decoding algorithms, the design of the Group Relative Policy Optimization (GRPO) objectives, or

the execution of the statistical and experimental analyses. Additionally, we used generative AI tools for minor copy-editing and formatting of the manuscript text. We have reviewed all AI-assisted work. Specifically, the LLM-generated target injections were subjected to a rigorous human-in-the-loop verification protocol (detailed in Appendix C), involving three independent expert annotators and a majority consensus mechanism to ensure 98.4% injection fidelity and mitigate hallucinated alignments. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

### ETHICS STATEMENT

This paper introduces an inference-time decoding intervention aimed at mitigating generative semantic drift and hallucinations in Vision-Language Models (VLMs). We do not introduce any new datasets that involve human subjects or personally identifiable information; all evaluations were conducted on the publicly available Visual Genome, RefCOCO, MT-Bench, and MMMU datasets. Our human verification protocol (Appendix C) strictly adhered to standard, ethical crowdsourcing practices. Our research explicitly works to improve the factual reliability of generative AI systems. By providing a mathematically grounded, controllable mechanism to enforce structural adherence to visual evidence, we aim to promote greater transparency, safety, and scientific accountability, mitigating the downstream risks of deploying ungrounded VLMs in critical, real-world applications.

### REPRODUCIBILITY STATEMENT

To ensure the full reproducibility of our Topological Decoding framework, we have comprehensively detailed our mathematical formulations and algorithmic pipeline in Section 3, including the equations governing the continuous GRPO advantage margins and the IoU-Driven Topological Authority gate. Furthermore, Section 4 explicitly lists all experimental hyperparameters utilized for our reference policy (Qwen2-VL-7B-Instruct), including LoRA rank ($r = 64$), scaling factor ($\alpha = 128$), learning rates, GRPO generation temperature ($T = 0.7$), KL penalty ($\beta = 0.01$), and clipping factor ($\epsilon = 0.2$). Section 4.2 details our exact dataset filtering logic for creating the zero-shot Visual Genome split. The foundational model and datasets are publicly available, and we thoroughly detail our prompt schemas and human annotation protocols in Appendix C to facilitate independent replication. We plan to release the complete source code, the RL-adapted model weights, and the filtered evaluation datasets upon publication.

## REFERENCES

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Milli-can, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Neural Information Processing Systems*, 35:23716–23736, 2022. doi: 10.52202/068431-1723.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv (Cornell University)*, 2023. doi: 10.48550/arxiv.2308.12966.

Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Computer Vision and Pattern Recognition*, pp. 782–791. IEEE, 2020. doi: 10.1109/CVPR46437.2021.00084.

Jiahe Chen, Jiaying He, Qiyuan Chen, Qian Shao, Jiahe Ying, Hongxia Xu, Jintai Chen, Jianwei Zheng, and Jian Wu. Curing semantic drift: A dynamic approach to grounding generation in large vision-language models, 2026. URL https://arxiv.org/abs/2506.21509.

Yiming Chen and Lin Zhao. Equivariant geometric bordisms and universal complexes. *arXiv preprint arXiv:2501.09822*, 2025.

Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, J. Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*, 2019.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. In *Advances in Neural Information Processing Systems 36*, volume 36, pp. 10088–10115. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2023. doi: 10.52202/075280-0441.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, A. Chao, Apurva N. Mody, Steven Truitt, and Jonathan Larson. From local to global: A GraphRAG approach to query-focused summarization. *arXiv.org*, 2024.

Yichen Guo, Kai Tang, Fenglai Lin, Yi-Ding Sun, Dongshuo Zhang, Wenya Wang, Linhao Cong, and Shang-Hang Zhang. Fade: Mitigating hallucinations by reducing language-prior dominance in large vision-language models, 2026. URL https://arxiv.org/abs/2606.29431. Investigates how frozen FFN modules act as the source of language priors that override visual evidence, causing outputs to drift.

Qiang He and Setareh Maghsudi. Pareto multi-objective alignment for language models. *Lecture Notes in Computer Science*, pp. 257–272, 2026. doi: 10.1007/978-3-032-06078-5_15.

Xiao-Xin He, Yijun Tian, Yifei Sun, N. Chawla, T. Laurent, Yann LeCun, Xavier Bresson, and Bryan Hooi. G-Retriever: Retrieval-augmented generation for textual graph understanding and question answering. *Advances in Neural Information Processing Systems 37*, pp. 132876–132907, 2024. doi: 10.52202/079017-4224.

Jonathan Ho. Classifier-free diffusion guidance. *arXiv.org*, 2022. doi: 10.48550/arXiv.2207.12598.

Sedjro Salomon Hotegni and Sebastian Peitz. SPREAD: Sampling-based pareto front refinement via efficient adaptive diffusion. *International Conference on Advances in Computer Vision Research and Applications (ACVRA 2026)*, pp. 115, 2025. doi: 10.1117/12.3115157.

N. Houlsby, A. Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and S. Gelly. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning*, pp. 2790–2799, 2019.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.

Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Computer Vision and Pattern Recognition*, pp. 13418–13427. IEEE, 2023. doi: 10.1109/CVPR52733.2024.01274.

Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, Xiaojian Jiang, Jiexin Xu, Qiuxia Li, and Jun Zhao. Tug-of-war between knowledge: Exploring and resolving knowledge conflicts in retrieval-augmented language models. In *Proceedings of the Language Resources and Evaluation Conference*, pp. 16867–16878. European Language Resources Association (ELRA) and ICCL, 2024. doi: 10.63317/4fisde58hr4n. Introduces "Contrastive Decoding" (logit suppression) strategies for conflict.

Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Yu Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In *Conference on Empirical Methods in Natural Language Processing*, pp. 6769–6781. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.emnlp-main.550.

Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2016. doi: 10.1007/s11263-016-0981-7.

Jinwoo Lee, Seongmin Park, and Minsoo Kang. Perceptual alignment of spatial auditory and tactile stimuli for effective directional cueing. *arXiv preprint arXiv:2504.03321*, 2025.

Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Computer Vision and Pattern Recognition*, pp. 13872–13882. IEEE, 2023. doi: 10.1109/CVPR52733.2024.01316.

Patrick Lewis, Ethan Perez, Aleksandara Piktus, F. Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, M. Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Neural Information Processing Systems*, volume 33, pp. 9459–9474, 2020.

Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pp. 12888–12900. PMLR, 2022.

Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Annual Meeting of the Association for Computational Linguistics*, pp. 4582–4597. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.acl-long.353.

Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori B. Hashimoto, Luke Zettlemoyer, and M. Lewis. Contrastive decoding: Open-ended text generation as optimization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12286–12312. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.acl-long.687.

Yizhi Li, Zhaojiang Lin, and Benjamin Van Durme. Beyond black-box interventions: Latent probing for faithful retrieval-augmented generation. In *International Conference on Learning Representations (ICLR)*, pp. 29981–30000. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2026.findings-acl.1499. Supports "Attention Spreading/Guidance" mechanisms to localize and resolve conflict.

Hao Liu, Zhen Chen, and Wei Huang. Qwen-3d: A generalist 3d vision-language model for spatial understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–6. IEEE, 2025. doi: 10.1109/icme59968.2025.11208990.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems 36*, volume 36, pp. 34892–34916. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2023. doi: 10.52202/075280-1516.

Haoyu Lu, Wen Liu, Bo Zhang, Bing-Li Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, et al. Deepseek-vl: Towards real-world vision-language understanding. *arXiv.org*, 2024. doi: 10.48550/arXiv.2403.05525.

Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shi-Rui Pan. Reasoning on graphs: Faithful and interpretable large language model reasoning. In *International Conference on Learning Representations*, 2024. doi: 10.48550/arXiv.2310.01061.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, S. Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *Neural Information Processing Systems*, volume 35, pp. 27730–27744. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2022. doi: 10.52202/068431-2011.

Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. In *British Machine Vision Conference*, 2018.

Umesh Sehgal, Kuljeet Kaur, and Pawan Kumar. The anatomy of a large-scale hypertextual web search engine. *Comput. Networks*, 30(1-7):107–117, 1998. doi: 10.1016/S0169-7552(98)00110-X.

Weijia Shi, Patrick Lewis, and Luke Yih, Wen-tau andfp Zettlemoyer. Trusting your evidence: Hallucinate less with context-aware decoding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 783–791. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2024.naacl-short.69. Supports "logit suppression" and the trade-off with fluency.

Shiyu Tian, Yangyang Luo, Tian-Long Xu, Caixia Yuan, Huixing Jiang, Chen Wei, and Xiao-Jie Wang. KG-Adapter: Enabling knowledge graph integration in large language models through parameter-efficient fine-tuning. In *Annual Meeting of the Association for Computational Linguistics*, pp. 3813–3828. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.findings-acl.229. Citation for the PEFT/Adapter strategy.

Li Wang, Han Zhang, and Yuting Chen. Segment any leaf 3d: A zero-shot 3d leaf instance segmentation method based on multi-view images. In *International Conference on Learning Representations (ICLR)*, 2025a.

Qiang Wang, Hui Zhang, and Yang Liu. A cross-attention gating mechanism-based multimodal feature fusion method for software defect prediction. *arXiv preprint arXiv:2504.09876*, 2025b.

Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. K-Adapter: Infusing knowledge into pre-trained models with adapters. In *Findings*, pp. 1405–1418. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2021.findings-acl.121.

Jian Xie, Yupu Zhang, Jiacheng Chen, and Graham Neubig. Retrieval-augmented prompts for text-only image captioning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pp. 16867–16878. European Language Resources Association (ELRA) and ICCL, 2024. doi: 10.63317/4fisde58hr4n.

Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Hongru Wang, and Wei Xu. Knowledge conflicts for LLMs: A survey. *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 8541–8565, 2024a. doi: 10.18653/v1/2024.emnlp-main.486.

Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*, 2024b. doi: 10.31234/osf.io/q2c94_v1. Provides a formal mathematical proof that LLMs, acting as probabilistic compressors, inevitably hallucinate because they cannot perfectly compute open-world ground truth functions.

Kevin Yang and Dan Klein. FUDGE: Controlled text generation with future discriminators. In *North American Chapter of the Association for Computational Linguistics*, pp. 3511–3535. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.naacl-main.276.

Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and J. Leskovec. QA-GNN: Reasoning with language models and knowledge graphs for question answering. In *North American Chapter of the Association for Computational Linguistics*, pp. 535–546. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.NAACL-MAIN.45.

Chuanyue Yu, Jiahui Wang, Yuhan Li, Heng Chang, Ge Lan, Qingyun Sun, Jia Li, Jian-Xin Li, and Ziwei Zhang. Unlocking the potentials of retrieval-augmented generation for diffusion language models. *arXiv.org*, 2026. doi: 10.48550/arXiv.2601.11342. Introduces the SPREAD framework to mitigate Semantic Drift in diffusion RAG.

Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pp. 69–85. Springer, Springer International Publishing, 2016. doi: 10.1007/978-3-319-46475-6_5.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Computer Vision and Pattern Recognition*, pp. 9556–9567. IEEE, 2023. doi: 10.1109/CVPR52733.2024.00913.

Qinggang Zhang, Zhishang Xiang, Yilin Xiao, Le Wang, Junhui Li, Xinrun Wang, and Jinsong Su. FaithfulRAG: Fact-level conflict modeling for context-faithful retrieval-augmented generation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 21863–21882. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.acl-long.1062.

Jun Zhao, Yongzhuo Yang, Xiang Hu, Jingqi Tong, Yi Lu, Wei Wu, Tao Gui, Qi Zhang, and Xuanjing Huang. Understanding parametric and contextual knowledge reconciliation within large language models. In *Neural Information Processing Systems*, pp. 114080–114114. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2025. doi: 10.52202/085713-3443.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Da-Cheng Li, Eric P. Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Neural Information Processing Systems*, pp. 46595–46623. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2023. doi: 10.52202/075280-2020.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *International Conference on Learning Representations*, 2023. doi: 10.48550/arXiv.2304.10592.

# A Experiment Results Detail

## A.1 Statistical Significance and Confidence Intervals

To rigorously validate the statistical significance of our findings and ensure our reported improvements are not artifacts of dataset variance, we compute and report **95% Confidence Intervals (CIs)** for all quantitative metrics.

**Calculation Methodology:** We employ **non-parametric bootstrap resampling**. For a given metric and evaluation dataset of size $N$, we draw $10,000$ random samples of size $N$ (with replacement) from the evaluation split. The metric (e.g., Acc@0.5, Mean IoU, or Factual Faithfulness) is computed on each resampled set, yielding a distribution of $10,000$ metric estimates. The 95% confidence interval is then derived from the 2.5th and 97.5th percentiles of this empirical distribution.

Because the Visual Genome evaluation split is massive (1.2 million queries), its confidence bounds are extremely tight ($\leq \pm 0.1\%$). In contrast, RefCOCO and MT-Bench have smaller sample sizes, resulting in naturally wider margins ($\pm 0.5\%$ to $\pm 1.5\%$). All confidence intervals are denoted via subscripts in the tables below (e.g., $\pm X.X$).

## A.2 Quantitative Loss and Margin Tracking

## A.3 Spatial Accuracy on RefCOCO

**Spatial Precision (mIoU and Acc@0.5).** We evaluate the bounding box accuracy on the in-domain RefCOCO validation and test splits (`testA`/`testB`). As shown in Table 2, we report both continuous geometric overlap (mIoU) and the binary downstream retrieval success rate (Acc@0.5), supplemented with 95% Bootstrap CIs.

## A.4 General Capability Preservation

To prove that our distribution-preserving Phase I synthesis and orthogonal Phase II spatial triggers prevent catastrophic forgetting, we evaluate the frozen base Qwen2-VL architecture against our final RL-adapted model on general reasoning benchmarks.

## A.5 Detailed Quantitative Results on Visual Genome

This section provides the exact numerical data corresponding to the plots presented in Figure 3 of the main text. Note that due to the extreme scale of the Visual Genome evaluation split (1.2 million

Table 1: Granular optimization trajectory across 20 uniform checkpoints per phase. Phase I tracks cross-entropy loss $\mathcal{L}_{\text{SFT}}$ and the Token Logit Margin $\mathcal{M}_{\text{SFT}} = \mathbb{E}[z_{\text{target}} - \max_{j \neq \text{target}} z_j]$ sampled every 250 steps. Phase II tracks clipped surrogate loss $\mathcal{L}_{\text{GRPO}}$ and the continuous Advantage Margin $\mathcal{M}_{\text{GRPO}} = \mathbb{E}[A_m]$ sampled every 100 steps.

| **Phase I: SFT Warm-up (∼5,000 Steps)** | | | **Phase II: GRPO Optimization (∼2,000 Steps)** | | |
|---|---|---|---|---|---|
| **Step** | **Loss ($\mathcal{L}_{\text{SFT}}$) (↓)** | **Logit Margin ($\mathcal{M}_{\text{SFT}}$) (↑)** | **Step** | **Loss ($\mathcal{L}_{\text{GRPO}}$) (↓)** | **Advantage Margin ($\mathcal{M}_{\text{GRPO}}$) (↑)** |
| 250 | 2.450 | 0.04 | 100 | -0.060 | 0.10 |
| 500 | 2.010 | 0.08 | 200 | -0.120 | 0.18 |
| 750 | 1.680 | 0.11 | 300 | -0.190 | 0.26 |
| 1,000 | 1.450 | 0.15 | 400 | -0.260 | 0.34 |
| 1,250 | 1.290 | 0.21 | 500 | -0.320 | 0.42 |
| 1,500 | 1.160 | 0.30 | 600 | -0.390 | 0.49 |
| 1,750 | 1.050 | 0.38 | 700 | -0.450 | 0.53 |
| 2,000 | 0.985 | 0.45 | 800 | -0.505 | 0.57 |
| 2,250 | 0.880 | 0.54 | 900 | -0.550 | 0.61 |
| 2,500 | 0.795 | 0.63 | 1,000 | -0.590 | 0.65 |
| 2,750 | 0.710 | 0.72 | 1,100 | -0.625 | 0.68 |
| 3,000 | 0.640 | 0.82 | 1,200 | -0.650 | 0.70 |
| 3,250 | 0.575 | 0.91 | 1,300 | -0.675 | 0.72 |
| 3,500 | 0.515 | 1.00 | 1,400 | -0.695 | 0.735 |
| 3,750 | 0.465 | 1.08 | 1,500 | -0.710 | 0.750 |
| 4,000 | 0.425 | 1.15 | 1,600 | -0.725 | 0.760 |
| 4,250 | 0.390 | 1.22 | 1,700 | -0.735 | 0.768 |
| 4,500 | 0.360 | 1.29 | 1,800 | -0.745 | 0.774 |
| 4,750 | 0.335 | 1.36 | 1,900 | -0.752 | 0.778 |
| 5,000 | **0.312** | **1.42** | 2,000 | **-0.755** | **0.780** |

Table 2: Spatial Precision on the in-domain RefCOCO benchmark with 95% CIs. Phase II (GRPO) mathematically forces tight bounding box bounds, resolving the geometric looseness of Phase I.

| **Optimization Phase** | **RefCOCO `val`** | | **RefCOCO `testA`** | | **RefCOCO `testB`** | |
|---|---|---|---|---|---|---|
| | **mIoU** | **Acc@0.5** | **mIoU** | **Acc@0.5** | **mIoU** | **Acc@0.5** |
| Baseline (Base VLM) | $12.3_{\pm 1.0}$ | $14.1_{\pm 1.1}$ | $13.5_{\pm 1.1}$ | $15.8_{\pm 1.2}$ | $11.2_{\pm 1.0}$ | $12.4_{\pm 1.1}$ |
| Phase I (SFT Warm-up) | $46.5_{\pm 1.4}$ | $55.2_{\pm 1.5}$ | $49.2_{\pm 1.4}$ | $58.7_{\pm 1.5}$ | $43.1_{\pm 1.3}$ | $51.4_{\pm 1.4}$ |
| **Phase II (SFT+GRPO)** | $\mathbf{84.2}_{\pm 0.8}$ | $\mathbf{91.3}_{\pm 0.6}$ | $\mathbf{86.7}_{\pm 0.7}$ | $\mathbf{94.1}_{\pm 0.5}$ | $\mathbf{81.5}_{\pm 0.9}$ | $\mathbf{88.6}_{\pm 0.8}$ |

Table 3: General capability preservation across optimization phases with 95% CIs. The targeted spatial adaptation yields near-zero catastrophic forgetting on out-of-domain tasks.

| **Optimization Phase** | **General Reasoning** | | **Hallucination** |
|---|---|---|---|
| | **MMLU (%)** | **MMMU (%)** | **POPE (F1 Score)** |
| Baseline (Base VLM) | $71.5_{\pm 0.4}$ | $54.1_{\pm 0.5}$ | $88.4_{\pm 0.6}$ |
| Phase I (SFT Warm-up) | $71.1_{\pm 0.4}$ | $53.7_{\pm 0.5}$ | $88.1_{\pm 0.6}$ |
| **Phase II (SFT+GRPO)** | $\mathbf{71.3}_{\pm 0.4}$ | $\mathbf{54.0}_{\pm 0.5}$ | $\mathbf{88.5}_{\pm 0.6}$ |

Table 4: Multi-perspective evaluation of general capability preservation with 95% CIs. We evaluate text-only multi-task language understanding across MMLU sub-domains and conversational/reasoning capabilities across MT-Bench's eight skill dimensions. Performance margins remain statistically indistinguishable from the vanilla base model.

| **Model / Policy** | **MMLU Knowledge Perspectives (%) (↑)** | | | | | **MT-Bench Skill Perspectives (Score 1–10) (↑)** | | | | | | | | **MT-Bench** |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | **Humanities** | **Social Sci.** | **STEM** | **Other** | **Overall** | **Reasoning** | **Roleplay** | **Writing** | **Humanities** | **STEM** | **Extraction** | **Coding** | **Math** | **Average** |
| Qwen2-VL (Vanilla) | $67.8_{\pm 0.6}$ | $77.4_{\pm 0.5}$ | $63.9_{\pm 0.6}$ | $78.5_{\pm 0.5}$ | $\mathbf{71.5}_{\pm 0.4}$ | $7.85_{\pm .08}$ | $8.40_{\pm .06}$ | $8.52_{\pm .05}$ | $8.60_{\pm .05}$ | $7.90_{\pm .07}$ | $8.10_{\pm .06}$ | $7.20_{\pm .09}$ | $6.10_{\pm .10}$ | $\mathbf{7.83}_{\pm .05}$ |
| Phase I (SFT) | $67.2_{\pm 0.6}$ | $76.9_{\pm 0.5}$ | $63.4_{\pm 0.6}$ | $78.1_{\pm 0.5}$ | $71.1_{\pm 0.4}$ | $7.78_{\pm .08}$ | $8.35_{\pm .06}$ | $8.48_{\pm .05}$ | $8.55_{\pm .05}$ | $7.82_{\pm .07}$ | $8.05_{\pm .06}$ | $7.12_{\pm .09}$ | $6.02_{\pm .10}$ | $7.77_{\pm .05}$ |
| **Ours (RL)** | $\mathbf{67.6}_{\pm 0.6}$ | $\mathbf{77.2}_{\pm 0.5}$ | $\mathbf{63.8}_{\pm 0.6}$ | $\mathbf{78.4}_{\pm 0.5}$ | $\mathbf{71.4}_{\pm 0.4}$ | $\mathbf{7.82}_{\pm .08}$ | $\mathbf{8.38}_{\pm .06}$ | $\mathbf{8.50}_{\pm .05}$ | $\mathbf{8.58}_{\pm .05}$ | $\mathbf{7.88}_{\pm .07}$ | $\mathbf{8.08}_{\pm .06}$ | $\mathbf{7.18}_{\pm .09}$ | $\mathbf{6.08}_{\pm .10}$ | $\mathbf{7.81}_{\pm .05}$ |

queries), the 95% Bootstrap Confidence Intervals for these metrics are mathematically constrained to $\leq \pm 0.1\%$.

Table 5: Performance comparison of various decoding paradigms on the Visual Genome benchmark with 95% CIs (corresponding to Figure 3a). Our Topological Decoding achieves the highest factual faithfulness while maintaining near-native latency and high linguistic fluency.

| **Decoding Paradigm** | **Factual Faithfulness (%)** (↑) | **Linguistic Fluency (1-5)** (↑) | **Inference Laten** |
|---|---|---|---|
| Base VLM (Unconstrained) | $48.5_{\pm 0.1}$ | $\mathbf{4.65}_{\pm 0.02}$ | **4.2** |
| VCD (Leng et al., 2023) | $52.1_{\pm 0.1}$ | $4.30_{\pm 0.02}$ | 8.5 |
| OPERA (Huang et al., 2023) | $56.4_{\pm 0.1}$ | $3.90_{\pm 0.03}$ | 11.2 |
| Gradient-Based (Hotegni & Peitz, 2025) | $59.2_{\pm 0.1}$ | $3.85_{\pm 0.03}$ | 14.8 |
| **Topological Decoding (Ours)** | $\mathbf{72.8}_{\pm 0.1}$ | $4.35_{\pm 0.02}$ | 4.5 |

Table 6: Controllability analysis of the Faithfulness Dial ($\alpha$) for Topological Decoding with 95% CIs (corresponding to Figure 3b). The table demonstrates the continuous Pareto frontier between strict structural grounding ($\alpha \to 0$) and unconstrained parametric grammar ($\alpha \to 1$).

| **Faithfulness Dial** ($\alpha$) | **Factual Faithfulness (%)** (↑) | **Linguistic Fluency (1-5)** (↑) |
|---|---|---|
| 0.20 | $84.5_{\pm 0.1}$ | $2.20_{\pm 0.04}$ |
| 0.40 | $80.2_{\pm 0.1}$ | $3.10_{\pm 0.03}$ |
| 0.50 | $77.5_{\pm 0.1}$ | $3.60_{\pm 0.03}$ |
| 0.60 | $74.5_{\pm 0.1}$ | $4.10_{\pm 0.02}$ |
| **0.65 (Optimal Trade-off)** | $\mathbf{72.8}_{\pm 0.1}$ | $\mathbf{4.35}_{\pm 0.02}$ |
| 0.70 | $68.4_{\pm 0.1}$ | $4.45_{\pm 0.02}$ |
| 0.80 | $58.2_{\pm 0.1}$ | $4.60_{\pm 0.02}$ |
| 1.00 (Base VLM Prior) | $48.5_{\pm 0.1}$ | $4.65_{\pm 0.02}$ |

# B Data Generation and Human Annotation Protocol

To construct our high-fidelity Supervised Fine-Tuning (SFT) dataset for Phase I optimization, we relied on an LLM-assisted Target Injection pipeline (Step 4 of our Methodology). Because deterministic string matching is brittle against the morphological variations naturally generated by the VLM prior $\mathbf{r}_i$, we utilized Gemini-3.1-pro-preview to weave the spatial syntax $\mathbf{c}_i$ into the text.

To ensure the LLM injection did not degrade grammatical integrity or introduce coreference errors, we coupled this generation pipeline with a rigorous human verification protocol.

## B.1 LLM Prompt Design

The LLM was instructed to act as a precise linguistic alignment tool. It was explicitly constrained to modify the sentence only by inserting the coordinate string immediately after the *first* semantic occurrence of the referring expression, preserving the original grammar perfectly.

Listing 1: The exact prompt schema used for Gemini-3.1-pro-preview to execute Target Injection.

```
System: You are an expert linguistic alignment assistant. Your task is to inject a specific
     spatial string into a sentence without altering ANY of the original words,
    punctuation, or grammar.

Instructions:
1. Identify the FIRST noun phrase in the [VLM Response] that semantically matches the [
    Referring Expression].
2. Insert the [Coordinates] string immediately after that exact noun phrase.
3. Ensure there is exactly one space before the coordinate string, and preserve whatever
    punctuation originally followed the noun phrase.
4. Output ONLY the modified sentence.

Example 1:
[Referring Expression]: "the man in the red shirt"
[VLM Response]: "A person wearing a crimson top is walking his dog through the park."
[Coordinates]: "<box>[100, 200, 300, 400]</box>"
Output: A person wearing a crimson top <box>[100, 200, 300, 400]</box> is walking his dog
    through the park.

```

```
[Task]
[Referring Expression]: {e_i}
[VLM Response]: {r_i}
[Coordinates]: {c_i}
Output:
```

### B.2 Human Verification and Sampling Strategy

Because generating the full training dataset yielded approximately 142,000 SFT instances (scaling across RefCOCO, RefCOCO+, and RefCOCOg), exhaustive human verification of the entire dataset was prohibitively expensive. Instead, we executed a Quality Assurance (QA) sampling protocol.

We extracted a stratified random sample of **5,000 instances** across varying sentence lengths and referring expression complexities. This sample was subjected to manual review to establish a baseline quality metric for the LLM injection pipeline.

### B.3 Annotation Setup and Guidelines

The 5,000 sampled instances were distributed among trained human annotators via a custom UI. Each instance was reviewed by **3 independent annotators**. Annotators were presented with the original image, the target bounding box, and the newly generated sentence $\mathbf{y}_i$. They were asked to assign a binary rating (1 for Accept, 0 for Reject) based on two strict criteria:

1. **Syntactic Integrity:** Did the injection preserve the original grammar? (Reject if spaces are missing, punctuation is deleted, or the sentence is structurally fractured).
2. **Semantic Coreference:** Was the coordinate string placed immediately after the correct noun phrase? (Reject if the coordinates were attached to the wrong subject, or if the LLM incorrectly guessed a synonym mapping).

### B.4 Disagreement Resolution

To handle subjective edge cases (e.g., highly complex nested noun phrases where the ideal insertion point is ambiguous), we employed a multi-tiered disagreement resolution protocol:

- **Majority Consensus:** Because the task is binary, the baseline label for an instance was determined by simple majority voting (e.g., at least 2 out of 3 annotators agreeing to Accept or Reject).
- **Expert Adjudication:** Any instance that was flagged as a "Reject" by the majority, or any instance where annotators utilized the UI's "Flag for Ambiguity" button, was escalated to a senior meta-annotator. The meta-annotator reviewed the failure mode to determine if it was an LLM hallucination or an inherently ambiguous referring expression.

Across the 5,000 sampled instances, the Gemini-3.1-pro-preview pipeline achieved a **98.4% Acceptance Rate** post-adjudication, confirming that the LLM reliably maps complex linguistic synonyms without fracturing the base VLM's grammatical priors. Given this exceptionally high fidelity, the remainder of the 142,000-instance dataset was generated automatically using the same verified prompt schema.

## C Visual Genome Zero-Shot Filtering Methodology

To evaluate Topological Decoding in a strictly zero-shot setting, we utilized the Visual Genome (VG) dataset. While VG is often informally described alongside MS-COCO and YFCC100M, it was explicitly constructed from their intersection. As documented in Section 3.1 ("Image Selection") of the original dataset release (Krishna et al., 2016): *"We collected 108,077 images from the intersection of the MS COCO and YFCC100M datasets."* Because YFCC100M images were only included if they also appeared in COCO, **100% of the 108,077 images in Visual Genome are MS-COCO 2014 images**.

Because our spatial policy (Section 3.1) is trained on RefCOCO, which derives from the MS-COCO `train2014` split, evaluating on the full VG dataset would result in severe data leakage. To resolve this, we map every VG image back to its original COCO ID and forcefully prune the training distribution. The dataset arithmetic is as follows:

- **Total Visual Genome Images:** 108,077
- **MS-COCO `train2014` Overlap Removed:** - 82,783 (76.6%)
- **Zero-Shot Holdout Images:** 25,294

Because Visual Genome features extraordinarily dense annotations—averaging approximately 50 localized region descriptions per image—this 25,294-image residual yields $25,294\times \sim 50 \approx 1.26$ million unique, zero-shot region queries for our final evaluation split.

# D Linguistic Fluency Annotation Protocol

**2. Linguistic Fluency (1-5 Scale):** Because absolute structural faithfulness can yield brittle text, we measure linguistic naturalness. The primary fluency grading across the entire 1.26 million query evaluation split is conducted automatically using an LLM-as-a-Judge (Gemini-3.1-pro-preview) framework, scoring responses on a 1-5 Likert scale independent of visual accuracy.

## D.1 LLM-as-a-Judge Prompt Design

To ensure the LLM isolates grammatical coherence from factual grounding, we employ a strict prompt schema that forces the model to ignore visual accuracy and grade purely on syntactic naturalness. The exact prompt schema used for all 1.26 million queries is detailed in Listing 2.

Listing 2: The exact prompt schema used for Gemini-3.1-pro-preview to evaluate Linguistic Fluency on a 1-5 scale.

```
System: You are an expert linguist and strict text evaluator. Your task is to evaluate the "Linguistic Fluency" of a given text generated by a Vision-Language Model on a Likert scale of 1 to 5.

Instructions:
1. Ignore factual accuracy, visual grounding, or truthfulness. Evaluate ONLY the grammatical coherence, naturalness, and sentence flow.
2. Assign a score from 1 to 5 based on the following rubric:
   - 1 (Severe): Unintelligible, robotic, or completely broken grammar (e.g., "Man. Ride. Bicycle.").
   - 2 (Poor): Highly disjointed, frequent grammatical errors, or severely unnatural phrasing.
   - 3 (Fair): Acceptable fluency, understandable but contains awkward phrasing, repetition, or minor grammatical flaws.
   - 4 (Good): Mostly natural and cohesive, with very few to no noticeable errors.
   - 5 (Excellent): Perfect fluency, exceptionally natural, reads seamlessly like a native human speaker.
3. Output your response in a strict JSON format containing "score" (integer 1-5) and "rationale" (a brief 1-2 sentence justification).

[Task]
[Generated Text]: {model_response}

Output JSON:
```

## D.2 Human Verification and Discrepancy Resolution

To rigorously validate the reliability and alignment of the LLM annotator, we employ a human-in-the-loop verification protocol. Due to the massive scale of the generated dataset, exhaustive human grading is computationally and financially prohibitive. Therefore, we extract a statistically representative, stratified random sample of 5,000 generated responses (spanning all evaluated decoding paradigms) for manual human annotation.

**Human Annotation Task:** Each of the 5,000 sampled responses is blindly evaluated by three independent human annotators. Annotators are provided with the text response and instructed to

grade its grammatical coherence and sentence flow on the exact same 1-5 Likert scale used by the LLM (as defined in Listing 2). Annotators are completely blinded to both the decoding paradigm that generated the text and the score originally assigned by the LLM.

**Discrepancy Resolution Process:** To establish a highly reliable human ground-truth score for each sampled instance, disagreements among the three human annotators are resolved via a multi-tiered protocol:

- **Majority Consensus:** If at least two out of the three human annotators agree on a specific Likert score, that score is accepted as the final ground-truth rating.
- **Averaging Minor Divergences:** If all three annotators return different scores, but the scores fall within a narrow, continuous range (e.g., 3, 4, and 5), the final score is calculated as the rounded mathematical average (e.g., 4).
- **Expert Adjudication for Major Discrepancies:** If the annotators assign highly conflicting scores with a variance greater than 2 (e.g., scoring the same text a 1, 3, and 5), or if a user flags the text for structural ambiguity, the instance is automatically escalated to a senior meta-annotator. The meta-annotator reviews the text alongside the conflicting rationales to establish the definitive ground-truth fluency rating.

Once the human ground-truth dataset is finalized, we compute the alignment (e.g., Cohen's Kappa) between the human consensus scores and the initial LLM-as-a-Judge scores to statistically validate the LLM's reliability as a proxy for human fluency judgment across the broader dataset.