

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CONSTRUCTIVE DISTORTION: IMPROVING MLLMs WITH ATTENTION-GUIDED IMAGE WARPING

Anonymous authors

Paper under double-blind review

ABSTRACT

Multimodal large language models (MLLMs) often miss small details and spatial relations in cluttered scenes, leading to errors in fine-grained perceptual grounding. We introduce AttWarp, a lightweight method that allocates more resolution to query-relevant content while compressing less informative areas, all while preserving global context. At test time, the approach uses an MLLM’s cross-modal attention to perform rectilinear warping of the input image, reallocating spatial resolution toward regions the model deems important, without changing model weights or architecture. This attention-guided warping preserves all original image information but redistributes it non-uniformly, so small objects and subtle relationships become easier for the same model to read while the global layout remains intact. [Across ten benchmarks \(TextVQA, GQA, DocVQA, POPE, MMMU, MIA-Bench, MMVP, VQAv2, RealWorldQA, BLINK\)](#) and four MLLMs (LLaVA, Qwen-VL, InternVL, and InstructBLIP), AttWarp consistently improves accuracy, strengthens compositional reasoning, and reduces hallucinations, outperforming four competitive baselines that manipulate raw images at test time. Together, these results show that attention-guided warping prioritizes information relevant to the query while preserving context, and that the same MLLMs perform better when given such warped inputs.

1 INTRODUCTION

Humans perceive certain regions of a scene by dynamically allocating high-resolution resources to areas of interest. This behavior is described as the interplay between *foveal vision*, which provides detailed perception at the center of attention, and *peripheral vision*, which rapidly scans the broader scene in lower resolution (Carrasco, 2011). This warped way of perceiving our surroundings is dynamic and dependent on task demands. As Gibson argued (Gibson, 1966), perceptual systems actively restructure their input, sampling dense information where it’s most needed. This introduces a form of distortion, not to obscure, but to enhance relevance.

While advanced deep learning models incorporate some aspect of this through attention mechanisms, they leave significant issues in fine-grained perceptual grounding. Multimodal LLMs often fail to identify small details, distinguish between similar objects, and understand complex spatial relationships in cluttered scenes, leading to misclassification and incorrect reasoning (Yang et al., 2024b; He et al., 2025; Kim & Ji, 2024). In this work, we investigate the benefits of this principle of warped perception in the context of modern multimodal LLMs. Particularly, we investigate the research questions: *what is an effective method for warping images that preserve global context while expanding task-relevant regions? Would existing multimodal LLMs perform better with warped images?*

To answer the first research question, we devise a lightweight recipe for warping images that preserves global context while expanding task-relevant regions. Our method, Attention-Guided Image Warping (AttWarp), operates as a plug-and-play enhancement, requiring no modifications to the underlying MLLM architecture. As illustrated in Fig. 1, given an input image and a query, we first extract cross-modal attention maps from the MLLM’s language decoder. These attention maps are aggregated into an *Attention Score Matrix* and further condensed into 1D *marginal attention profiles* along both horizontal and vertical axes (Fig. 1, middle). These profiles quantify the importance of each row and column in the image, with taller bars indicating regions that should receive more visual real estate.

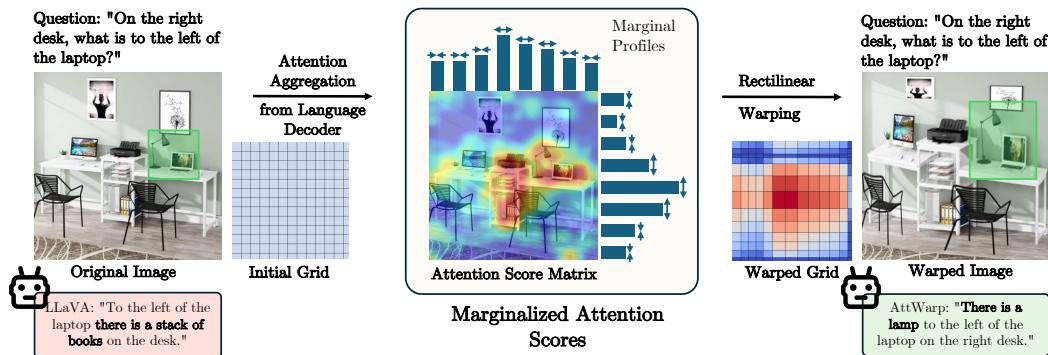


Figure 1: AttWarp overview. Given a query, our method extracts cross-modal attention maps from the MLLM’s language decoder, aggregates them into marginal attention profiles, and uses rectilinear warping to expand high-attention regions while compressing low-attention areas. The warped image is then processed by the same MLLM, which now produces the correct answer.

These marginal profiles guide our novel *rectilinear warping* process, which non-uniformly resamples the image grid, expanding high-attention regions and compressing low-attention areas, as shown by the red-to-blue gradient in Fig. 1 (right). Crucially, our choice of warping ensures that all original image information is preserved, maintaining global context unlike methods that crop or mask parts of the image. In the warped image, task-relevant objects such as the lamp and laptop highlighted by the green bounding box, are visually enlarged, making fine-grained details and spatial relationships more accessible to the MLLM. Towards the second research question, we find that such warped images when processed by the same MLLM lead to improved performance across ten multimodal benchmarks, and the idea generalizes to multiple MLLM backbones. We empirically validate that our rectilinear design is crucial to this improvement without changing a single parameter of the MLLM. We can extend this framework in two directions: multiple calls to AttWarp lead to further improvements (AttWarp-Chain), and we can learn a distilled model to directly predict marginal profiles instead of estimating them from MLLM attention maps (AttWarp-Distill).

Conceptually, AttWarp utilizes attention to modify the input image itself, different from the typical use of attention to reweight latent features. Importantly, our contribution is complementary to research that improves attention mechanisms in MLLMs, such as refining attention heads (Bi et al., 2024; Kang et al., 2025), adding auxiliary objectives (Yan et al., 2024), or redesigning cross-modal fusion layers (He et al., 2025). Finally, note that we intervene *before* feature extraction, while the above methods operate *after* the image has already been encoded, often from features that have already lost critical spatial detail (Pantazopoulos et al., 2024).

In summary, our key contributions are: 1) A *lightweight method* AttWarp that addresses the fine-grained perceptual grounding challenges identified in multimodal LLMs by modifying input images before feature extraction, requiring no MLLM finetuning. 2) Consistent empirical gains over *four competitive baselines*, across *ten standard vision-language benchmarks* that test diverse capabilities such as fine-grained multimodal understanding, compositional reasoning, and hallucination mitigation. 3) Demonstrating generalization across *multiple MLLM backbones* and attention sources, underscoring its plug-and-play compatibility. 4) *Rigorous analysis* validating that our warps indeed expand task-relevant regions and AttWarp’s rectilinear design helps preserve the original data distribution.

2 RELATED WORK

Perception challenges in MLLMs. MLLMs such as Flamingo (Alayrac et al., 2022), LLaVA (Liu et al., 2023b), Qwen (Yang et al., 2024a), MiniGPT-4 (Zhu et al., 2023), and GPT-4V (OpenAI, 2023) have advanced image-grounded dialogue and reasoning (Mitra et al., 2024; Dong et al., 2024; Zhang et al., 2025a; 2024b). Yet they still struggle with *fine-grained* perception—missing small attributes (e.g. object (Zhang et al., 2024a)), misclassifying sub-categories (Geigle et al., 2024; Kim & Ji, 2024; Yu et al., 2025), and confusing geometric primitives (Zhang et al., 2025b). These limitations motivate our work, which aims to enhance the fine-grained perception of the MLLMs by improving query-specific spatial resolution prior to feature extraction.

108 **Different Approaches for Fine-Grained Visual Understanding.** We group prior efforts into the
 109 following categories: (i) Bounding-box methods (Peng et al., 2023; Chen et al., 2023a; Zhang et al.;
 110 Lu et al., 2024) steer attention by feeding cropped regions obtained from bounding boxes. (ii) Mask-
 111 based methods (Chen et al., 2024b; Yuan et al., 2024; You et al., 2023; Zhang et al., 2024b) supply
 112 pixel-accurate masks—often from Segment-Anything (Kirillov et al., 2023). (iii) Cascade methods
 113 (He et al., 2025; Yang et al., 2023a; Yu et al., 2024) overlay detector cues or saliency heatmaps to
 114 bias the input. (iv) Reasoning methods (Surís et al., 2023; Wei et al., 2022) decompose queries into
 115 low-level visual steps. Our approach achieves stronger grounding while avoiding extra detectors,
 116 masks, or multi-step reasoning chains.

117 **Pixel-Level Warping Techniques for Saliency Emphasis.** Classical work includes seam carving
 118 (Rubinstein et al., 2010), saliency-aware warps (Wolf et al., 2007; Recasens et al., 2018), energy
 119 minimisation (Karni et al., 2009), finite-element grids (Kaufmann et al., 2013), mesh parametrisation
 120 (Guo et al., 2009), and seam & scale methods (Zhang et al., 2009). Learning-based variants explore
 121 adaptive resizing (Talebi & Milanfar, 2021), saliency enhancement (Ghosh et al., 2019; Miangoleh
 122 et al., 2023), and domain adaptation (Zheng et al., 2025), with contemporaneous magnification
 123 work (Mao et al., 2025). Many existing approaches, such as those employing energy minimization
 124 or seam carving, are optimization-based. Consequently, processing each input sample can take
 125 several minutes. In contrast, our proposed method leverages a single forward pass of a Cumulative
 126 Distribution Function (CDF), enabling near-instantaneous processing. Here, we build on saliency-
 127 aware sampling by introducing a *query-conditioned, rectilinear* warp that preserves the image’s
 128 regular grid structure, ensuring compatibility with the MLLM’s vision encoder.

129 3 ATTENTION-GUIDED IMAGE WARPING

130 We propose a simple and effective test-time technique to improve visual grounding of MLLMs.
 131 Instead of feeding the original image directly, we apply a spatial transformation guided by the model’s
 132 internal attention, reshaping image regions based on their relevance to the query. Below, we provide
 133 a high-level overview of AttWarp, followed by a detailed description of its components.

134 **Overview.** As illustrated in Fig. 1, AttWarp uses cross-modal attention maps from deeper layers of
 135 the MLLM (see Sec. 3.4) to guide a distribution-preserving non-uniform resampling of the original
 136 image. This resampling operation, termed **rectilinear warping** (Sec. 3.1), redistributes pixel density
 137 across the image: regions with high attention are spatially expanded, while less relevant regions are
 138 compressed. Relevance is always defined with respect to the specific query, making the warping
 139 adaptive to task semantics. Crucially, the warped image retains a regular grid structure, ensuring
 140 compatibility with standard vision encoders. Next, in Sec. 3.2, we introduce AttWarp-Chain an
 141 extension of AttWarp that iteratively applies multiple warps, grounding each step in the model’s
 142 evolving attention and improving performance on complex queries. Finally, in Sec. 3.3 we introduce
 143 AttWarp-Distill, a computationally efficient version optimized for inference speed, which runs
 144 3× faster than prior methods by shifting additional computation to training time through learned
 145 warping functions.

146 3.1 RECTILINEAR IMAGE WARPING

147 Given an input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ and an attention score matrix $A \in \mathbb{R}^{H \times W}$ (from Sec. 3.4), our
 148 goal is to obtain a function F that transforms \mathbf{I} to a warped image $\mathbf{W} = F(\mathbf{I}; A)$. The warping
 149 function F is designed to magnify important regions (high attention) and compress less relevant ones.

150 First, we compute marginal attention profiles (PDFs) along rows and columns to decompose the 2D
 151 attention matrix into 1D score vectors:

$$152 \text{Horizontal Attention Profile: } m_x(j) = \sum_{i=1}^H A_{ij}, \text{ Vertical Attention Profile: } m_y(i) = \sum_{j=1}^W A_{ij}. \quad (1)$$

153 Here, $i \in (1, 2 \dots H)$ and $j \in (1, 2 \dots W)$. This decomposition facilitates rectilinear warping,
 154 enabling independent transformations along the horizontal and vertical axes while preserving the grid
 155 structure. Subsequently, we convert these marginals into cumulative distribution functions (CDFs):

$$156 M_x(j) = \frac{\sum_{k=1}^j m_x(k)}{\sum_{k=1}^W m_x(k)}, \quad M_y(i) = \frac{\sum_{k=1}^i m_y(k)}{\sum_{k=1}^H m_y(k)}. \quad (2)$$

162 These resulting cumulative functions (M_x, M_y) are monotonically increasing and therefore invertible.
 163 We define the warping functions using their inverses, known as the *Inverse Distribution Functions*:

$$164 \quad f_X^{\text{Warp}}(j) = W \cdot M_x^{-1}(j/W), \quad f_Y^{\text{Warp}}(i) = H \cdot M_x^{-1}(j/H) \quad (3)$$

166 where $i \in (1, 2 \dots H)$ and $j \in (1, 2 \dots W)$. Together, these inverse mappings f_X^{Warp} and f_Y^{Warp}
 167 constitute the overall warping transformation F , yielding the warped image. The final warped image
 168 is computed through bilinear sampling, applied along all three channels:
 169

$$170 \quad \mathbf{W}[i, j] = \text{Interpolate}(\mathbf{I}, \text{Bilinear})(f_Y^{\text{Warp}}(i), f_X^{\text{Warp}}(j)). \quad (4)$$

172 3.2 ITERATIVE IMAGE WARPING (ATTWARP-CHAIN)

174 The optimal degree of warping depends on the query and the image. Queries focusing on small
 175 details benefits from strong warping, while broader scene may require minimal warping. A naive way
 176 is to use a superlinear (or sublinear) transformation over the attention score matrix to upweigh (or
 177 downweigh) the attention-guided warp. In this section, we introduce a more intuitive and nuanced
 178 scheme that performs better called AttWarp-Chain.

179 We build an iterative warping method based on an empirical observation i.e., warping improves
 180 MLLMs attention (See 4.3) and enhanced attention maps subsequently yield better warping. Leverag-
 181 ing this insight, we develop AttWarp-Chain which after each iteration, extracts an updated attention
 182 map, and progressively refines the warp applied in the previous iteration. Formally,

$$183 \quad \text{Initialization: } \mathbf{W}^{(0)} = \mathbf{I}, \quad \text{Chain step: } \mathbf{W}^{(d)} = F(\mathbf{W}^{(d-1)}; A^{(d-1)}). \quad (5)$$

185 Here, $A^{(d-1)}$ denotes the attention map computed from the warped visual input $\mathbf{W}^{(d-1)}$.

186 A practical question left to answer is when to terminate the chain of iterative warping steps? Instead of
 187 encoding this as a hyperparameter, we propose a more adaptive route. As the relevant query-specific
 188 region expands, the attention map spreads more uniformly over the image. Eventually, the attention
 189 distribution stabilizes, indicated by minimal changes between successive attention maps. We quantify
 190 this stability through the following stopping criterion:

$$191 \quad \mathcal{D}_{\text{KL}} \left(P^{(d)} \middle| P^{(d-1)} \right) < \epsilon_{\text{KL}}, \quad (6)$$

192 where $P^{(d)}$ and $P^{(d-1)}$ are normalized attention probability distributions from iterations d and $d-1$,
 193 respectively. This termination ensures AttWarp-Chain achieves optimal spatial emphasis while
 194 mitigating the risks associated with noisy or overly aggressive warping. We quantitatively show
 195 effectiveness of AttWarp-Chain and termination criteria in App. D.3.

197 3.3 LEARNING TO PREDICT MARGINAL ATTENTION PROFILES (ATTWARP-DISTILL)

199 Many applications (e.g., edge AR and embodied agents) need fast, precise grounding. Masking,
 200 cropping, or re-running attention adds latency; even AttWarp-Chain requires multiple passes. We
 201 therefore learn a *single-pass* predictor that outputs the horizontal and vertical marginals, $m_x(j)$ and
 202 $m_y(i)$, directly from an image–text pair. This neural functional approximation removes attention
 203 retrieval at inference and keeps the warping semantics of Sec. 3.1.

204 **Teacher: generating marginal targets.** We use the base MLLM and the attention-extraction pipeline
 205 from Sec. 3.4 to produce training targets. For each image–text pair (\mathbf{I} , text) we compute the attention
 206 score matrix A , derive axis-wise marginals via equation 1, and normalize them to unit mass, yielding
 207 (m_x, m_y) . This defines the dataset $\mathcal{D} = \{(\mathbf{I}_n, \text{text}_n, m_{x,n}, m_{y,n})\}_{n=1}^N$. Constructing these targets is
 208 intentionally done once offline to train the student; at inference the student replaces this pipeline,
 209 amortizing cost and enabling single-pass usage.

210 **Student: AttWarp-Distill.** Using these offline targets, we train a compact network to predict
 211 (m_x, m_y) from an image–text pair (architecture in Fig. 2). We encode the image with CLIP ViT-L/14
 212 to obtain vision tokens \mathbf{Z} and obtain text tokens \mathbf{q} from a tokenizer E_t applied to the query. Text
 213 conditions the vision tokens through Feature-wise Linear Modulation (FiLM) (Perez et al., 2018).
 214 A small MLP maps the text tokens \mathbf{q} to per-channel scale and shift parameters (a, b) , and applies
 215 them channel-wise to obtain the modulated tokens $\tilde{\mathbf{Z}} = a \odot \mathbf{Z} + b$. We then upsample $\tilde{\mathbf{Z}}$ to
 (H, W) and average along one axis at a time to obtain two 1D summaries (horizontal and vertical).

216 Two light Conv1D heads turn these summaries into logits, which a SoftMax converts to valid
 217 marginals (\hat{m}_x, \hat{m}_y) . Training minimizes the expected L_1 discrepancy over \mathcal{D} , i.e., the average of
 218 $\|\hat{m}_x - m_x\|_1 + \|\hat{m}_y - m_y\|_1$ across samples.

219 **Single-pass, fast inference.** At inference, given $(\mathbf{I}, \text{text})$,
 220 AttWarp-Distill outputs (\hat{m}_x, \hat{m}_y) in one forward
 221 pass. We convert them to CDFs (\hat{M}_x, \hat{M}_y) via Eq. equation 2, invert to coordinates using Eq. equation 3, and
 222 bilinearly sample the image as in Eq. equation 4 to obtain
 223 \mathbf{W} . This retains the semantics of Sec. 3.1 while reducing
 224 cost. Training details appear in App. G.3.

227 3.4 ATTENTION SCORE MATRIX IMPLEMENTATION

228 Here, we describe a general recipe for constructing the
 229 attention score matrix \mathbf{A} for any base MLLM. The pro-
 230 cedure has two steps: (1) *Attention retrieval*, which reads
 231 raw cross-attention weights, and (2) *Attention aggregation*,
 232 which collapses those weights into a single spatial map for
 233 the given image–text pair.

234 **Attention map retrieval.** The image \mathbf{I} is tokenized into
 235 $n_{\text{img}} = h_{\text{feat}} w_{\text{feat}}$ vision tokens on a $h_{\text{feat}} \times w_{\text{feat}}$ grid.
 236 After processing \mathbf{I} and the text, the MLLM produces n_{out}
 237 output tokens. From selected decoder layers \mathcal{L} and all
 238 attention heads (n_{heads}), we obtain cross-attention matrices
 239 $a^{(\ell,h)} \in \mathbb{R}^{n_{\text{out}} \times n_{\text{img}}}$; entry $a_{m,t}^{(\ell,h)}$ is the weight from output
 240 token m to image token t .

241 **Attention map aggregation.** We now average over out-
 242 put tokens, heads, and the chosen layers to form a single
 243 spatial map. Each image token index $t \in \{1, \dots, n_{\text{img}}\}$ corresponds to grid location (i, j) via
 244 $t = (i-1)w_{\text{feat}} + j$. The aggregated score at (i, j) is

$$\tilde{A}_{i,j} = \frac{1}{n_{\text{out}} \cdot n_{\text{heads}} \cdot |\mathcal{L}|} \sum_{\ell \in \mathcal{L}} \sum_{m=1}^{n_{\text{out}}} \sum_{h=1}^{n_{\text{heads}}} a_{m,t}^{(\ell,h)}. \quad (7)$$

245 Here, $a_{m,t}^{(\ell,h)}$ denotes the weight from output token m to image token t . We upsample \tilde{A} from $h_{\text{feat}} \times$
 246 w_{feat} to the image resolution $H \times W$ using Lanczos, smooth with a $k \times k$ AvgPool, and optionally
 247 apply a scalar transform \mathcal{T} to control sharpness. The final attention score matrix is $\mathcal{T}(\tilde{A}_{ij})$, $\mathcal{T} \in$
 248 $\{x, x^2, \sqrt{x}, \dots\}$. Here, $i \in \{1, 2, \dots, H\}$, $j \in \{1, 2, \dots, W\}$. Sharper transforms like x^2
 249 emphasize high-attention regions more than linear ones, which is useful for fine-grained queries. For
 250 LLaVA, we choose the 20th layer i.e. $\mathcal{L} = \{20\}$. For Qwen, we use the 16th layer i.e. $\mathcal{L} = \{16\}$.
 251 The strategy for choosing layer(s) \mathcal{L} and attention aggregation method is explained quantitatively in
 252 App. B. For more details on implementation and design choices, refer App. G.1.

253

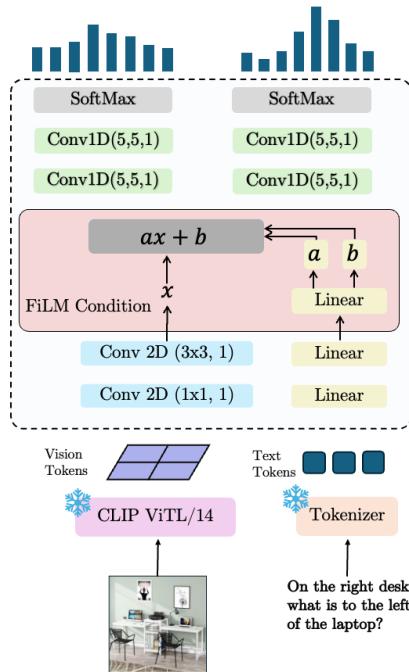
254 4 EXPERIMENTS

255 In this section, we detail our experimental framework and findings. We begin by describing the
 256 evaluation benchmarks and the baseline models used for comparison (Sec. 4.1). Next, we present
 257 our key quantitative results across multiple multimodal benchmarks (Sec. 4.2). Finally, we include
 258 ablations and analyses to share insights into AttWarp (Sec. 4.3).

259

260 4.1 BENCHMARKS AND BASELINES

261 **Benchmarks.** We evaluate AttWarp on *ten diverse benchmarks* designed to assess key multimodal
 262 capabilities, including general multimodal understanding, compositional reasoning, spatial relationships,
 263 visual hallucination, and fine-grained visual understanding (Fig. 3 shows qualitative results).
 264 • *GQA (visual reasoning):* Reasoning about objects, attributes, and relations in real-world im-
 265 ages (Hudson & Manning, 2019).



266 **Figure 2: AttWarp-Distill architecture:**
 267 CLIP vision tokens are FiLM-modulated by
 268 text and projected to 1D marginal predictors.
 269 corresponds to grid location (i, j) via

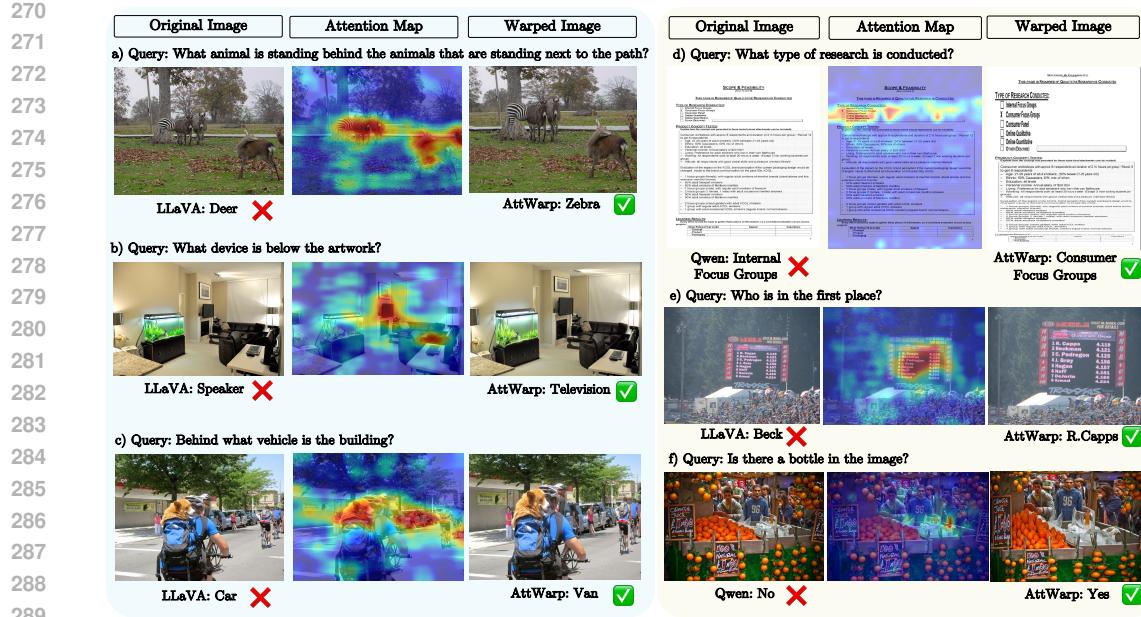


Figure 3: AttWarp improves compositional & spatial reasoning *e.g.* from GQA dataset (a) correctly identifying zebra behind the path, (b) television below the artwork; text understanding in documents *e.g.* from DocVQA (d) consumer focus groups, fine-grained recognition of small/occluded objects *e.g.* from POPE (f) detecting a bottle

- **TextVQA (scene text understanding):** Answering questions that require reading and grounding text in natural images (Singh et al., 2019).
- **DocVQA (document image understanding):** Extracting and reasoning over textual and structural information in scanned documents (Mathew et al., 2021).
- **POPE (robustness evaluation):** Probing fine-grained hallucination and reliability in vision–language models (Li et al., 2023).
- **MMMU (general multimodal understanding):** Broad multi-disciplinary evaluation across STEM, humanities, and social sciences (Yue et al., 2024).
- **VQAv2 (general visual QA):** Large-scale visual question answering requiring joint reasoning over image content, natural language, and commonsense knowledge (Goyal et al., 2017).
- **RealWorldQA (spatial reasoning in the wild):** Spatial reasoning and relative localization in complex real-world scenes captured from vehicles and other environments (xAI, 2024).
- **BLINK (fine-grained perception):** Core visual perception skills via fine-grained tasks (*e.g.*, relative depth estimation, object localization, counting) cast into a VQA-style interface (Fu et al., 2024).
- **MMVP (fine-grained perception):** Evaluation on *CLIP-blind* image pairs that stress robustness to systematic changes in visual patterns such as orientation and viewpoint (Tong et al., 2024).
- **MIA-Bench (instruction following):** Multimodal instruction following under layered constraints on style, length, and content while remaining visually grounded (Qian et al., 2024).

We report the result of AttWarp on GQA, TextVQA, DocVQA, POPE, and MMMU in Tab. 1 and results on MMVP, BLINK, RealWorldQA, MIA and VQAv2 in Tab. 3.

Baselines. For a rigorous evaluation, beyond the base MLLMs, we also compare AttWarp to *four representative baselines for test-time visual intervention*. Baselines span strategies for editing input image to guide MLLM attention directly at inference (see Fig. 4 for visual examples).

- FGVP (Yang et al., 2023b) (Region Isolation via Masking): applies semantic masks on the target region and reversely blurs the background (*or* applies a green mask) outside the target region.
- SoM (Yang et al., 2023a) (Visual Grounding with Explicit Markers): segments the input image semantically using an off-the-shelf model, and labels each segment with a unique visual marker.
- APIPrompting (Yu et al., 2024) (Attention-Modulated Image Representation): computes an attention heatmap using an auxiliary VLM (LLaVA or CLIP), and overlays it onto the input image.
- ViCrop (Zhang et al., 2025a) (Context Reduction through Attention-Guided Cropping): crops images around regions of high saliency, based on the model’s attention map. Notably, ViCrop takes two image input *i.e.* both the original and the cropped image.

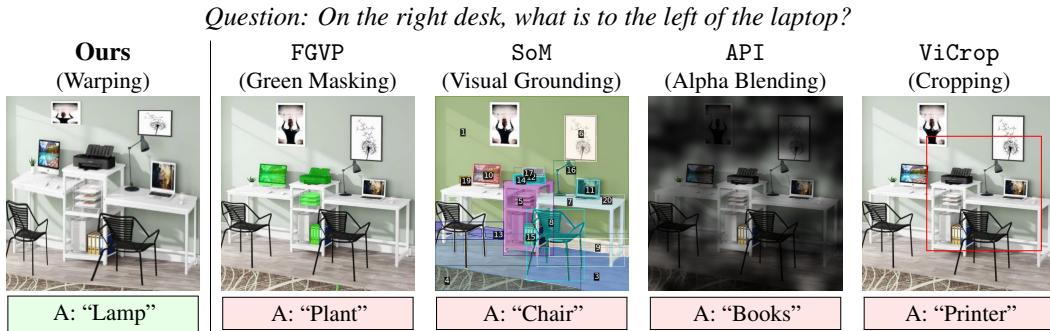
324 **Table 1:** Main results on TextVQA, GQA, MMMU, POPE, and DocVQA datasets in accuracy (%). The Δ
 325 Accuracy row reports the absolute improvement of AttWarp-Chain over the base MLLM.

#	Methods	Key Technique	TextVQA	GQA	MMMU	POPE	DocVQA
LLaVA (Liu et al., 2024a) (<i>MLP vision-language connector & open data</i>)							
1	Base MLLM		49.3	60.5	36.9	85.3	18.1
2	FGVP-mask (Yang et al., 2023b)	Green mask overlay	39.4	59.2	36.1	85.3	19.0
3	FGVP-blur (Yang et al., 2023b)	Blur background	33.9	59.5	35.0	83.1	18.6
4	SoM (Yang et al., 2023a)	Grounded segments	18.8	54.5	35.6	78.5	15.8
5	API (Yu et al., 2024)	Alpha channel fade	49.9	60.6	36.9	85.9	17.4
6	ViCrop (Zhang et al., 2025a)	Add object crop	56.3	60.9	37.2	87.0	22.5
7	AttWarp (<i>ours</i>)	Rectilinear warping	58.1	63.7	40.4	87.5	25.5
8	AttWarp-Distill (<i>ours</i>)	Efficient inference	57.2	62.7	38.8	87.4	22.4
9	AttWarp-Chain (<i>ours</i>)	Adaptive Chains	60.3	64.4	41.6	88.2	27.6
10	Δ Accuracy		+11.0	+3.9	+4.7	+2.9	+9.5
Qwen (Yang et al., 2024a) (<i>Cross-attention VL adapter & partially closed data</i>)							
11	Base MLLM		81.0	62.4	47.3	86.1	77.3
12	FGVP-mask (Yang et al., 2023b)	Green mask overlay	77.3	55.8	46.0	84.4	56.6
13	FGVP-blur (Yang et al., 2023b)	Blur background	72.3	55.8	46.5	81.3	38.6
14	SoM (Yang et al., 2023a)	Grounded segments	61.5	47.8	45.1	75.8	57.4
15	API (Yu et al., 2024)	Alpha channel fade	81.6	61.1	47.4	85.8	68.4
16	ViCrop (Zhang et al., 2025a)	Add object crop	83.8	60.6	47.1	86.7	82.5
17	AttWarp (<i>ours</i>)	Rectilinear warping	84.7	64.0	50.4	87.4	84.1
18	AttWarp-Distill (<i>ours</i>)	Efficient inference	84.1	63.1	48.9	87.2	81.8
19	AttWarp-Chain (<i>ours</i>)	Adaptive Chains	85.9	64.8	51.0	88.0	85.3
20	Δ Accuracy		+4.9	+2.4	+3.7	+1.9	+8.0

348 4.2 QUANTITATIVE RESULTS

350 In Tab. 1, we present results for three methods introduced in this work, *i.e.*, AttWarp,
 351 AttWarp-Chain, and AttWarp-Distill, evaluated on two MLLMs and five diverse benchmarks.

352 **AttWarp outperforms all baselines on five benchmarks (Tab. 1).** AttWarp achieve state-of-the-art
 353 results on tasks including text recognition and understanding (TextVQA: LLaVA +8.8%, Qwen +3.7%);
 354 DocVQA: LLaVA +7.4%, Qwen +6.8%), compositional and spatial reasoning (GQA: LLaVA +3.2%,
 355 Qwen +1.6%), general multimodal question answering (MMMU: LLaVA +3.5%, Qwen +3.1%),
 356 and fine-grained understanding and hallucination reduction (POPE: LLaVA +2.2%, Qwen +1.3%).
 357 These consistent improvement across diverse tasks arises from AttWarp’s capability to highlight
 358 task-relevant objects while preserving *global context* and spatial relationships, thus delivering strong
 359 performance on global queries in GQA and fine-grained queries in POPE. Moreover, these results
 360 illustrate that AttWarp’s *is agnostic to image type* – effective across natural scenes (GQA, TextVQA,
 361 POPE), documents (DocVQA), and dense diagrams (MMMU). We provide per-category performance
 362 in App. C.1 and study on the extent of warping in App. C.3. Overall, the superior performance of
 363 AttWarp underscores the significance of enhancing perception for question answering.



376 **Figure 4:** AttWarp and prior works of image manipulation on the running example. While plausible, prior
 377 works are unable to answer the question correctly.

378 **AttWarp is plug-and-play with standard MLLMs (Tab. 1).** By default, we evaluate AttWarp
 379 using LLaVA, consistent with prior works (Zhang et al., 2025a; Yu et al., 2024; Yang et al., 2023a).
 380 We further demonstrate AttWarp versatility on a stronger MLLM, Qwen2.5-VL. Qwen2.5-VL uses a
 381 distinct vision-language fusion strategy with a dedicated cross-modal attention module to integrate
 382 visual and textual features, contrasting LLaVA’s direct projection of visual features into the language
 383 model’s input space. The consistent performance gains of AttWarp across Qwen2.5-VL and LLaVA-
 384 v1.5-7b highlight AttWarp’s robust, plug-and-play compatibility with diverse MLLM architectures.
 385 Additionally, AttWarp achieves similar gains with two other architectures: a dynamic-resolution and
 386 pixel-unshuffling approach (InternVL3), and an instruction-aware Q-Former for vision–language
 387 alignment (InstructBLIP). See App. D.1 for results.

388 **Chaining provides consistent performance gains over**
 389 **standard AttWarp (Tab. 1, rows 7 & 9, rows 17 &**
 390 **19).** The multi-step AttWarp-Chain further boosts per-
 391 formance by iteratively refining attention maps. Exam-
 392 ining LLaVA results (rows 7 & 9), we observe consistent
 393 improvements over AttWarp across all five benchmarks
 394 (TextVQA +2.2%, GQA +0.7%, MMMU +1.2%, POPE
 395 +0.7 %, and DocVQA +2.1%). A similar trend is evident
 396 when evaluated on the stronger base model Qwen (rows
 397 17 & 19). Qualitatively, Fig. 5 demonstrates how AttWarp-Chain adaptively adjusts the warping
 398 extent, resulting in enhanced visual grounding and improved task performance.

399 **AttWarp-Distill balances accuracy and speed**
 400 **(Tab. 1, Tab. 2).** AttWarp-Distill consistently out-
 401 performs the base MLLM and matches or exceeds the
 402 performance of ViCrop (Tab. 2). Optimized spec-
 403 ifically for inference efficiency, AttWarp-Distill re-
 404 quires only a single MLLM forward pass, making it ap-
 405 proximately $3\times$ faster and $2.8\times$ more computationally
 406 efficient than ViCrop; aligning closely with the com-
 407 putational cost of the base MLLM (8.7 vs. 8.5 TFLOPs).
 408 In Tab. 2, we compare our method against ViCrop as it
 409 achieves comparable performance to AttWarp-Distill. Other baselines, such as FGVP, SoM, and
 410 APIPrompting, not only perform substantially worse (see Tab. 1), but also incur additional overhead
 411 due to multiple inference steps. For more details and cost analysis of AttWarp refer App. G.4.

411 **AttWarp-Distill is generalizable (Tab. 1).** We train AttWarp-Distill on the standard
 412 training splits of widely used open-source datasets (TextVQA, GQA, and DocVQA), which form part
 413 of the training corpora of most base MLLMs. The substantial improvements on TextVQA (LLaVA
 414 +7.9%, Qwen +3.1%), GQA (LLaVA +2.2%, Qwen +0.7%), and DocVQA (LLaVA +4.3%, Qwen
 415 +4.5%), demonstrate its strong in-domain generalization capability, whereas its robust performance
 416 on POPE (LLaVA +2.1%, Qwen +1.1%), and MMMU (LLaVA +1.9%, Qwen +1.6%), highlights
 417 its effective out-of-domain generalization. Further details on the training procedure and marginal
 418 prediction in App. G.3.

419 **Table 3: AttWarp performance (in %) on VQAv2 and visual-centric benchmarks.** We report individual accuracy
 420 for MMVP, as it aligns with the format of other evaluations. We use LLaVA as base MLLM.

LLaVA	VQAv2	MMVP	BLINK	RealWorldQA	MIA
Base MLLM	75.6	48.3	38.3	49.3	63.9
AttWarp-Distilled (ours)	77.1	49.3	39.7	51.1	65.4
AttWarp (ours)	78.9	50.7	40.4	52.1	67.2
AttWarp-Chains (ours)	79.7	51.0	41.2	52.9	68.8
Δ Accuracy	+4.1	+2.7	+2.9	+3.6	+4.9

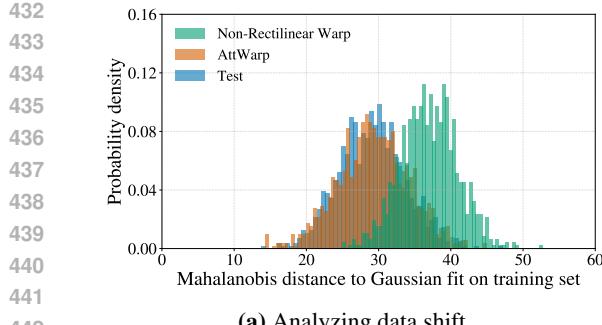
427 **Results on VQAv2 and visual-centric benchmarks:** We further evaluate AttWarp on five additional
 428 benchmarks: four visual-centric datasets (RealWorldQA, BLINK, MMVP, MIA-Bench) and the
 429 general-purpose VQAv2 benchmark. Across all five, AttWarp consistently outperforms the base
 430 MLLM (Tab. 3), indicating stronger fine-grained visual understanding, more accurate spatial reason-



431 **Figure 5: AttWarp-Chain improves on**
 432 **AttWarp**

433 **Table 2: Comparison of computational over-**
 434 **head.** Base MLLM used is LLaVA. Metrics are
 435 TFLOPs, peak VRAM (in GB), and number of
 436 MLLM passes. Values in brackets show rela-
 437 tive cost compared to ViCrop.

	TFLOPs ↓	Peak VRAM ↓	MLLM passes ↓
ViCrop	24.2	22	3
AttWarp-Distill	8.7 (0.4×)	15	1
Base MLLM	8.5	15	1



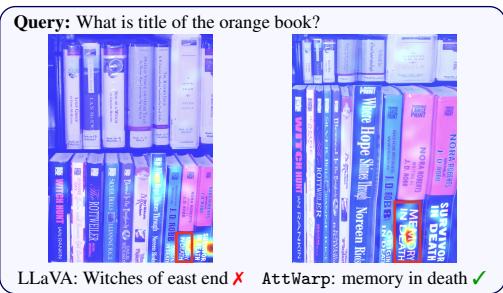
(a) Analyzing data shift.

(b) Distribution metrics (\downarrow = lower is better).

Distribution ν under:	KID (μ_{train}, ν) \downarrow	FID (μ_{train}, ν) \downarrow
AttWarp	31.5	49.8
Non-Rectilinear Warp	174.9	73.9
Test Set	19.3	56.6

449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

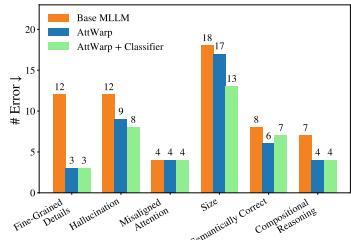
Figure 6: (a) Mahalanobis distance to the train Gaussian. (b) Train→Test, AttWarp, Non-Rect. FID/KID summary. (c–d) Attention–redistribution GT alignment. See Appx. I for setup and additional plots.



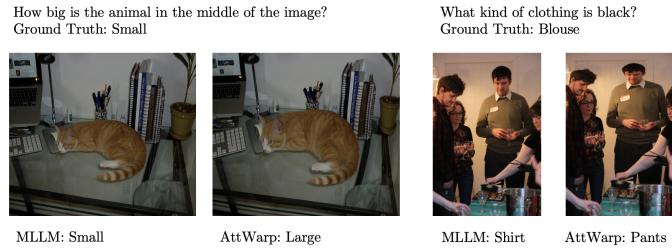
(c) Attention alignment with Ground truth boxes.

(d) Attention–GT alignment metrics on TextVQA.

Metric	No warp	With AttWarp
Pointing Game Accuracy \uparrow	37.4%	42.4%
Proportion \uparrow	0.117	0.155



(a) Quantitative error plot.



(b) Examples of failure cases of AttWarp.

473
474
475
476
477
478
479
480
481
482
483
484
485

Figure 7: Error Characterization and Failure Cases.

ing, and improved visually grounded instruction following. AttWarp-Distill, retains most of these gains with lower computational cost, while AttWarp-Chain, leverages warping at multiple layers for further improvements. As summarized in Tab. 3, all three variants, AttWarp, AttWarp-Distill, and AttWarp-Chain, improve performance across these five benchmarks, supporting rectilinear attention-guided warping as an effective and general mechanism for enhancing fine-grained visual grounding in MLLMs.

4.3 ABLATIONS AND ANALYSIS

Warping Improves MLLMs Attention. To gauge how AttWarp reshapes the *spatial faithfulness* of the internal attention of the model, we adopt two widely-adopted localization metrics that rely only on ground-truth boundary boxes (bbox) and the heat map itself: 1) **Pointing Game Accuracy (top-1 attention in bbox):** checks whether the *single* most salient pixel of the attention map falls within the GT bbox (Zhang et al., 2016). 2) **Proportion (fraction attention within box):** the fraction of *total* attention mass that lands inside the bbox (Wang et al., 2020). As observed in Tab. 6d, the +5% relative jump in *Pointing Game Accuracy* and the +3.8% boost in *Proportion* on TextVQA confirm that our rectilinear warp *tightens* the focus of the model: the attention mass and dominant peak shifts to the task relevant region. These findings indicate that AttWarp improves the MLLM performance because of better attention distribution (see Fig. F and App.F). We further extend this study and empirically verify that AttWarp *expands the correct image regions*. We quantify this by checking how often AttWarp expands the bboxes of the relevant regions (refer App. E for more details).

AttWarp’s rectilinear design preserves the image distribution. Pixel-level transforms risk shifting test images away from the manifold on which the model was trained. We probe this risk for LLaVA by

486 first extracting ViT-L/14 CLIP features (Radford et al., 2021) for $12k$ randomly sampled GQA-train
 487 images and fitting a full-covariance Gaussian to those embeddings. Figure 6a compares the resulting
 488 *Mahalanobis distance* histogram for three image sets: (i) Test images (blue); (ii) images warped by
 489 AttWarp (orange); and (iii) a non-rectilinear warping baseline, inspired by (Recasens et al., 2018)
 490 (green). AttWarp almost exactly overlaps the unmodified test distribution (both peak at $\approx 29\sigma$),
 491 whereas distribution based on Non-Rectilinear Warp shifts rightward (peak $\approx 37\sigma$) and exhibits
 492 a heavier tail, indicating a measurable distribution drift. We also test this aspect using standard
 493 two-sample metrics, *FID* (Heusel et al., 2017) and *KID* (Bińkowski et al., 2018) between the training
 494 distribution and each test variant. As evident from Table 6b, AttWarp achieves significantly better
 495 alignment with the training distribution (**31.5** vs. 174.9 KID and **49.8** vs. 73.9 FID), effectively
 496 reproducing the train–test baseline gap. Alongside this, as can be seen in Figure 6a, our approach
 497 closely matches the baseline metrics indicating that AttWarp preserves the underlying image manifold
 498 and introduces negligible distribution shift (see App. I for graphs and details, and App. C.5 for
 499 perceptual geometry analysis).

500 **Error Analysis.** We randomly sampled 150 VQA tasks from the GQA and TextVQA evaluation.
 501 A total of 61 were incorrect for the base LLaVA model, and 42 were incorrect for AttWarp. We
 502 bin these into the closest failure modes: *Fine-Grained Details*, where the answer is very small in
 503 size; *Hallucination* when the answer includes details not present in the image; *Misaligned Attention*,
 504 when focus shifts to the wrong object; *Size*, for questions involving object scale; *Semantically*
 505 *Correct*, when the answer is correct but phrased differently; and *Compositional Reasoning*, involving
 506 multiple objects and relationships. Fig. 7a shows fewer errors in fine-grained and compositional
 507 cases. However, we note that warping can sometimes suppress peripheral context needed for global
 508 reasoning, and performance may degrade if the underlying attention is noisy. While absolute object
 509 sizes are changed, relative proportions are preserved, limiting errors in size-related tasks.

510 **Reducing the Error Modes.** The error analysis shows that AttWarp reduces errors across all
 511 categories compared to the base MLLM (e.g., hallucination: 12 \rightarrow 9, fine-grained details: 12 \rightarrow 3).
 512 However, AttWarp remains susceptible to errors related to object size and misaligned attention,
 513 motivating the development of methods that further enhance robustness.

514 To explicitly target these failure cases, we design a simple classifier that decides whether to apply
 515 AttWarp. Concretely, we reuse AttWarp-Distill’s weights and network, and replace its last two layers
 516 with a binary classifier head. We create a training set by evaluating AttWarp on the training split
 517 of AttWarp-Distill (App. G), and use these outcomes to train the classifier. We denote the resulting
 518 classifier-gated variant by AttWarp † , i.e., AttWarp † = Classifier + AttWarp.

519 We then perform the same error analysis for AttWarp † , shown in Fig. 7a (green bars). AttWarp †
 520 further reduces errors in size and hallucination while keeping other categories essentially unchanged.
 521 Misaligned-attention errors remain unchanged, as both the base MLLM and AttWarp fail on the
 522 same underlying attention misalignment. Designing mechanisms that directly correct attention
 523 misalignment is therefore a productive direction for future work.

524 5 CONCLUSION

525 In this work, we introduced AttWarp, a plug-and-play, test-time image warping that uses an MLLM’s
 526 cross-modal attention to rectilinearly resample the input, expanding query-relevant regions while
 527 preserving global context. Without changing model weights or architecture, it consistently improves
 528 accuracy, spatial grounding, and hallucination rates across ten benchmarks and four MLLMs. By
 529 intervening before feature extraction, AttWarp complements internal attention refinements and shows
 530 that input-level, information-preserving transformations can help the same models see better.

540

6 REPRODUCIBILITY STATEMENT

541
 542 We will release a complete, lightweight reference implementation of AttWarp under an open-
 543 source license, including: (1) attention-extraction adapters for LLaVA, Qwen-VL, InternVL, and
 544 InstructBLIP; (2) rectilinear warping code with evaluation scripts and exact configs for all ten
 545 benchmarks; and (3) precomputed marginal profiles and warped images for the reported splits. For
 546 the distilled variant, we will publicly release trained weights and training code. Links to the repository
 547 and any model weights will be provided in the camera-ready. Also, we provide a detailed discussion
 548 of the implementation details in Appendix App. G, where we outline the hyperparameters, as well as
 549 the training and inference procedures used.

550

551 REFERENCES

552 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel
 553 Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language
 554 model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736,
 555 2022.

556 Jing Bi, Junjia Guo, Yunlong Tang, Lianggong Bruce Wen, Zhang Liu, and Chenliang Xu. Unveiling
 557 visual perception in language models: An attention head analysis approach. *arXiv preprint
 558 arXiv:2412.18108*, 2024.

559 Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd
 560 gans. In *ICLR*, 2018.

561 Marisa Carrasco. Visual attention: The past 25 years. *Vision research*, 51(13):1484–1525, 2011.

562 Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing
 563 multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023a.

564 Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang.
 565 An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-
 566 language models. In *European Conference on Computer Vision*, pp. 19–35. Springer, 2024a.

567 Mengzhao Chen, Mingbao Lin, Ke Li, Yunhang Shen, Yongjian Wu, Fei Chao, and Rongrong
 568 Ji. Cf-vit: A general coarse-to-fine method for vision transformer. In *Proceedings of the AAAI
 569 conference on artificial intelligence*, volume 37, pp. 7042–7052, 2023b.

570 Mengzhao Chen, Wenqi Shao, Peng Xu, Mingbao Lin, Kaipeng Zhang, Fei Chao, Rongrong Ji,
 571 Yu Qiao, and Ping Luo. Diffrate: Differentiable compression rate for efficient vision transformers.
 572 In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 17164–17174,
 573 2023c.

574 Yi-Chia Chen, Wei-Hua Li, Cheng Sun, Yu-Chiang Frank Wang, and Chu-Song Chen. Sam4mllm:
 575 Enhance multi-modal large language model for referring expression segmentation. In *European
 576 Conference on Computer Vision*, pp. 323–340. Springer, 2024b.

577 Wenliang Dai, Junnan Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. Instructblip:
 578 Towards general-purpose vision-language models with instruction tuning. *arXiv 2023*. *arXiv
 579 preprint arXiv:2305.06500*, 2, 2023.

580 Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkang Yang, Winston Hu, Yongming Rao, and Ziwei Liu.
 581 Insight-v: Exploring long-chain visual reasoning with multimodal large language models. *arXiv
 582 preprint arXiv:2411.14432*, 2024.

583 Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith,
 584 Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not
 585 perceive. In *European Conference on Computer Vision (ECCV)*, Cham, 2024. Springer Nature
 586 Switzerland.

587 Gregor Geigle, Radu Timofte, and Goran Glavaš. African or european swallow? benchmarking large
 588 vision-language models for fine-grained object classification. *arXiv preprint arXiv:2406.14496*,
 589 2024.

594 Sanjay Ghosh, Ruturaj G. Gavaskar, and Kunal N. Chaudhury. Saliency guided image detail
 595 enhancement. In *Proceedings of the National Conference on Communications (NCC)*, pp. 1–5.
 596 IEEE, 2019.

597

598 James J. Gibson. *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston, 1966.
 599 ISBN 9780395044940.

600 Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V
 601 in VQA matter: Elevating the role of image understanding in Visual Question Answering. In
 602 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

603

604 Yanwen Guo, Feng Liu, Jian Shi, Kun Zhou, and Baining Guo. Image retargeting using mesh
 605 parametrization. *IEEE Transactions on Multimedia*, 11(5):856–867, 2009.

606

607 Hulingxiao He, Geng Li, Zijun Geng, Jinglin Xu, and Yuxin Peng. Analyzing and boosting the
 608 power of fine-grained visual recognition for multi-modal large language models. *arXiv preprint*
 609 *arXiv:2501.15140*, 2025.

610

611 Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common
 612 corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

613

614 Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-
 615 to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.
 616 URL <https://arxiv.org/abs/2208.01626>.

617

618 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans
 619 trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural*
 620 *Information Processing Systems*, volume 30, pp. 6626–6637, 2017.

621

622 Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning
 623 and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer*
 624 *Vision and Pattern Recognition (CVPR)*, pp. 6700–6709, 2019. doi: 10.1109/CVPR.2019.00686.

625

626 Seil Kang, Jinyeong Kim, Junhyeok Kim, and Seong Jae Hwang. Your large vision-language model
 627 only needs a few attention heads for visual grounding. *arXiv preprint arXiv:2503.06287*, 2025.

628

629 Zachi Karni, Daniel Freedman, and Craig Gotsman. Energy-based image deformation. *Computer*
 630 *Graphics Forum*, 28(5):1257–1268, 2009.

631

632 Peter Kaufmann, Olga Wang, Alexander Sorkine-Hornung, Aljoscha Smolic, and Markus Gross.
 633 Finite element image warping. *Computer Graphics Forum*, 32(2pt1):31–39, 2013.

634

635 Jeonghwan Kim and Heng Ji. Finer: Investigating and enhancing fine-grained visual concept
 636 recognition in large vision language models. *arXiv preprint arXiv:2402.16315*, 2024.

637

638 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
 639 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings*
 640 *of the IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026, 2023.

641

642 Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning
 643 segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer*
 644 *Vision and Pattern Recognition*, pp. 9579–9589, 2024.

645

646 Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluat-
 647 ing object hallucination in large vision-language models. In *Proceedings of the 2023 Confer-
 648 ence on Empirical Methods in Natural Language Processing*, pp. 292–305, Singapore, 2023.
 649 Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.20. URL
 650 <https://aclanthology.org/2023.emnlp-main.20/>.

651

652 Zhihang Lin, Mingbao Lin, Luxi Lin, and Rongrong Ji. Boosting multimodal large language models
 653 with visual tokens withdrawal for rapid inference. In *Proceedings of the AAAI Conference on*
 654 *Artificial Intelligence*, volume 39, pp. 5334–5342, 2025.

648 Chang Liu, Henghui Ding, and Xudong Jiang. Gres: Generalized referring expression segmentation.
 649 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.
 650 23592–23601, 2023a.

651

652 Haotian Liu, Chunyuan Li, Qingsong Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in*
 653 *neural information processing systems*, 36:34892–34916, 2023b.

654 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 655 tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 656 pp. 26296–26306, 2024a.

657

658 Haotian Liu, Chunyuan Li, Qingsong Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in*
 659 *neural information processing systems*, 36, 2024b.

660 Jinghui Lu, Haiyang Yu, Yanjie Wang, Yongjie Ye, Jingqun Tang, Ziwei Yang, Binghong Wu, Qi Liu,
 661 Hao Feng, Han Wang, et al. A bounding box is worth one token: Interleaving layout and text in a
 662 large language model for document understanding. *arXiv preprint arXiv:2407.01976*, 2024.

663

664 Shunqi Mao, Chaoyi Zhang, and Weidong Cai. Through the magnifying glass: Adaptive perception
 665 magnification for hallucination-free vlm decoding. *arXiv preprint arXiv:2503.10183*, 2025.

666

667 Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. Docvqa: A dataset for vqa on document
 668 images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*
 669 (WACV), pp. 2199–2208, 2021. doi: 10.1109/WACV48630.2021.000225.

670

671 Sadegh Aliakbarian Miangoleh, Emre Yaksoy, Miika Aittala, Sylvain Paris, and Frédo Durand.
 672 Realistic saliency guided image enhancement. In *Proceedings of the IEEE/CVF Conference on*
 673 *Computer Vision and Pattern Recognition (CVPR)*, pp. 186–195, 2023.

674

675 Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. Compositional chain-of-thought
 676 prompting for large multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer*
 677 *Vision and Pattern Recognition*, pp. 14420–14431, 2024.

678

679 OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. doi: 10.48550/arXiv.2303.
 680 08774. URL <https://arxiv.org/abs/2303.08774>.

681

682 Georgios Pantazopoulos, Alessandro Suglia, Oliver Lemon, and Arash Eshghi. Lost in space: Probing
 683 fine-grained spatial understanding in vision and language resamplers. In *Proceedings of the 2024*
 684 *Conference of the North American Chapter of the Association for Computational Linguistics: Human*
 685 *Language Technologies (Volume 2, Short Papers)*, 2024. The authors show that the
 686 compression performed by popular resampler modules leaves most spatial information “largely
 687 absent,” evidencing a lossy visual encoding stage.

688

689 Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties
 690 in gan evaluation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
 691 *recognition*, pp. 11410–11420, 2022.

692

693 Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu
 694 Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint*
 695 *arXiv:2306.14824*, 2023.

696

697 Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. FiLM: Visual
 698 reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial*
 699 *Intelligence*, volume 32. AAAI Press, 2018.

700

701 Yusu Qian, Hanrong Ye, Jean-Philippe Fauconnier, Peter Grasch, Yinfeng Yang, and Zhe Gan. MIA-
 702 Bench: Towards better instruction following evaluation of multimodal LLMs. *arXiv preprint*
 703 *arXiv:2407.01509*, 2024.

704

705 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 706 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 707 models from natural language supervision. In *ICML*, 2021.

702 Adria Recasens, Petr Kellnhofer, Simon Stent, Wojciech Matusik, and Antonio Torralba. Learning
 703 to zoom: A saliency-based sampling layer for neural networks. In *Proceedings of the European*
 704 *Conference on Computer Vision (ECCV)*, pp. 51–66, 2018.

705

706 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 707 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Confer-*
 708 *ence on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

709 Michael Rubinstein, Diego Gutierrez, Olga Sorkine, and Ariel Shamir. A comparative study of image
 710 retargeting. *ACM Transactions on Graphics (TOG)*, 29(6):1–10, 2010.

711

712 Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh,
 713 and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF*
 714 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8317–8326, 2019.

715 Yizheng Sun, Hao Li, Chang Xu, Chenghua Lin, Riza Batista-Navarro, and Jingyuan Sun. Silent
 716 hazards of token reduction in vision-language models: The hidden impact on consistency. *arXiv*
 717 *e-prints*, pp. arXiv–2503, 2025.

718

719 Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for
 720 reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp.
 721 11888–11898, 2023.

722

723 Hossein Talebi and Peyman Milanfar. Learning to resize images for computer vision tasks. In
 724 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 497–506,
 2021.

725

726 Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide
 727 shut? exploring the visual shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF*
 728 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9568–9578, June 2024.

729

730 Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and
 731 Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In
 732 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*
 (CVPR W), 2020.

733

734 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 735 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 736 *neural information processing systems*, 35:24824–24837, 2022.

737

738 Lihui Wolf, Michael Guttmann, and Daniel Cohen-Or. Non-homogeneous content-driven video-
 739 retargeting. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 1–6. IEEE, 2007.

740

741 xAI. Realworldqa, 2024. URL <https://huggingface.co/datasets/xai-org/RealworldQA>.

742

743 Siming Yan, Min Bai, Weifeng Chen, Xiong Zhou, Qixing Huang, and Li Erran Li. Vigor: Improving
 744 visual grounding of large vision language models with fine-grained reward modeling. In *European*
 745 *Conference on Computer Vision*, pp. 37–53. Springer, 2024.

746

747 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 748 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
 749 *arXiv:2412.15115*, 2024a.

750

751 Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark
 752 prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*,
 753 2023a.

754

755 Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in
 756 space: How multimodal large language models see, remember, and recall spaces. *arXiv preprint*
 757 *arXiv:2412.14171*, 2024b.

758

Lingfeng Yang, Yueze Wang, Xiang Li, Xinlong Wang, and Jian Yang. Fine-grained visual prompting.
 759 *Advances in Neural Information Processing Systems*, 36:24993–25006, 2023b.

756 Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao,
 757 Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity.
 758 *arXiv preprint arXiv:2310.07704*, 2023.

759

760 Hong-Tao Yu, Xiu-Shen Wei, Yuxin Peng, and Serge Belongie. Benchmarking large vision-language
 761 models on fine-grained image tasks: A comprehensive evaluation. *arXiv preprint arXiv:2504.14988*,
 762 2025.

763 Runpeng Yu, Weihao Yu, and Xinchao Wang. Attention prompting on image for large vision-language
 764 models. In *European Conference on Computer Vision*, pp. 251–268. Springer, 2024.

765

766 Yuqian Yuan, Wentong Li, Jian Liu, Dongqi Tang, Xinjie Luo, Chi Qin, Lei Zhang, and Jianke Zhu.
 767 Osprey: Pixel understanding with visual instruction tuning. In *Proceedings of the IEEE/CVF
 768 Conference on Computer Vision and Pattern Recognition*, pp. 28202–28211, 2024.

769

770 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruqi Liu, Ge Zhang, Samuel Stevens, Dongfu
 771 Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin,
 772 Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen.
 773 Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert
 774 agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
 (CVPR)*, 2024.

775

776 Jianming Zhang, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural
 777 attention by excitation backprop. In *Proceedings of the 14th European Conference on Computer
 778 Vision (ECCV)*, pp. 543–559, 2016.

779

780 Jiarui Zhang, Jinyi Hu, Mahyar Khayatkhoei, Filip Ilievski, and Maosong Sun. Exploring perceptual
 781 limitation of multimodal large language models. *arXiv preprint arXiv:2402.07384*, 2024a.

782

783 Jiarui Zhang, Mahyar Khayatkhoei, Prateek Chhikara, and Filip Ilievski. Mllms know where to
 784 look: Training-free perception of small visual details with multimodal llms. *arXiv preprint
 785 arXiv:2502.17422*, 2025a.

786

787 Liang Zhang, Meng Zhang, Bo Wang, and Ruigang Wang. Optimized image resizing using seam
 788 carving and scaling. In *Proceedings of the 17th ACM International Conference on Multimedia*, pp.
 789 147–156. ACM, 2009.

790

791 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 792 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on
 793 computer vision and pattern recognition*, pp. 586–595, 2018.

794

795 Shan Zhang, Aotian Chen, Yanpeng Sun, Jindong Gu, Yi-Yu Zheng, Piotr Koniusz, Kai Zou, Anton
 796 van den Hengel, and Yuan Xue. Open eyes, then reason: Fine-grained visual mathematical
 797 understanding in mllms. *arXiv preprint arXiv:2501.06430*, 2025b.

798

799 Shilong Zhang, Peize Sun, Shoufa Chen, Minn Xiao, Wenqi Shao, Wenwei Zhang, Yu Liu, Kai Chen,
 800 and Ping Luo. Gpt4roi: Instruction tuning large language model on regionof-interest, 2024. In
 801 *URL https://openreview.net/forum*.

802

803 Zhehao Zhang, Ryan Rossi, Tong Yu, Franck Dernoncourt, Ruiyi Zhang, Jiuxiang Gu, Sungchul
 804 Kim, Xiang Chen, Zichao Wang, and Nedim Lipka. Vipact: Visual-perception enhancement via
 805 specialized vlm agent collaboration and tool-use. *arXiv preprint arXiv:2410.16400*, 2024b.

806

807 Shen Zheng, Anurag Ghosh, and Srinivasa G. Narasimhan. Instance-warp: Saliency guided image
 808 warping for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference
 809 on Applications of Computer Vision (WACV)*, pp. 8186–8195, 2025.

810

811 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: En-
 812 hancing vision-language understanding with advanced large language models. *arXiv preprint
 813 arXiv:2304.10592*, 2023.

814

815 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
 816 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
 817 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

810 APPENDIX
811812 This appendix provides additional material to support and extend the findings in the main paper. Each
813 entry is clickable.
814

815 **A FAQ** Clarifications on distributional integrity, layer usage, failure modes with poor attention,
816 and multi-object focus.

817 **B Ablation Study on Attention Score Matrix** Empirical analysis of where the best attention
818 comes from and how we use it.

819 **B.1 Layer Selection Across MLLM Layers & Head Aggregation** – Best layers to use
820 and why averaging heads is robust.

821 **B.2 Pixel-Space vs. Feature-Space Warping** – Why rectilinear pixel warping is stable
822 and model-agnostic.

823 **B.3 Effect of the Transform Hyperparameter (\mathcal{T})** – Sensitivity across
824 sqrt/identity/square/cube.

825 **B.4 Impact of Attention Bias, Corruptions, and Adversarial Perturbations** – Robust-
826 ness under ImageNet-C noise and adversarial misdirection.

827 **C Task-Specific Analysis of AttWarp** Detailed breakdown of where AttWarp helps and the
828 few places it does not.

829 **C.1 Task Categories, Benefits, and Limitations** – Fine-grained perception, spatial reason-
830 ing, and hallucination mitigation.

831 **C.2 Distortion and Upper-Bound Concerns** – Geometry preservation and semantic cate-
832 gory results.

833 **C.3 Global Context and Warping Intensity** – Adaptive strength via Jacobian-based
834 intensity.

835 **C.4 Task-wise Comparison with ViCrop** – Category-level head-to-head on GQA.

836 **C.5 Perceptual Geometry Preservation** – LPIPS-based geometry preservation experiment.

837 **D Additional Experiments** Broader evaluations showing generality and stability.

838 **D.1 Extended Generalization** – Results on InstructBLIP and InternVL-3.

839 **D.2 Fine-Grained and Category-Wise DocVQA** – Breakdown across document struc-
840 tures.

841 **D.3 Stability and Termination in AttWarp-Chains** – Fixed vs. adaptive depth with KL
842 stopping.

843 **D.4 External Attention Maps** – Using Stable Diffusion and Qwen-VL attention to help
844 LLaVA.

845 **E Beyond Standard VQA: Extended Experiments** Forward-looking pilots demonstrating
846 versatility.

847 **E.1 AttWarp Expands All Query-Relevant Regions** – Single/multi-object bbox expansion
848 on TextVQA and gRef.

849 **E.2 Open-Vocabulary Object Detection** – Integration with LISA-LLaVA.

850 **E.3 Leveraging Smaller Models to Improve Larger Models** – 7B attention boosting 34B.

851 **F Attention Redistribution: Reduction in Model Uncertainty** Visual and quantitative
852 evidence of sharper, less diffused attention.

853 **G Implementation Details** Reproducibility and configuration.

854 **G.1 Quantitative Evaluation Setup** – Models, layers, resolutions, thresholds, and base-
855 lines.

856 **G.2 Stable Diffusion Experimental Details** – Attention extraction and schedules.

857 **G.3 Training Setup for AttWarp-Distill** – Teacher marginals and optimization.

858 **G.4 Computational Cost Calculation and Analysis** – FLOPs/tokens/VRAM vs. ViCrop.

859 **H Qualitative Examples** Case studies of AttWarp and AttWarp-Chain.

860 **H.1 Examples of AttWarp** – Corrections over baseline.

861 **H.2 Examples of AttWarp-Chain** – Iterative refinement and convergence.

862 **I FID and KID Analysis** Train–test alignment in CLIP feature space (ViT-B/32 and ViT-
863 L/14).

864
865

A FAQ

866
867
1. How is distributional integrity maintained during warping?868
Ans. Two primary reasons underlie this preservation.869
870
871
872
873
874
875
First, the rectilinear warping preserves spatial attributes, positional relationships, and
structural integrity. At the tokenization stage, objects maintain their semantic identities.
Second, the rectilinear warping directly aligns with the data augmentation strategy
used during CLIP pre-training. Specifically, CLIP pre-training uses RandomResized-
Crop—randomly sampling 8–100% of the image with aspect ratios 3/4–4/3, then
resizing to 224×224 px—scaling rows and columns independently yet axis-aligned,
thus exposing the model to varied scales and positions for robust, scale-invariant
representations Radford et al. (2021). As the warp is defined as

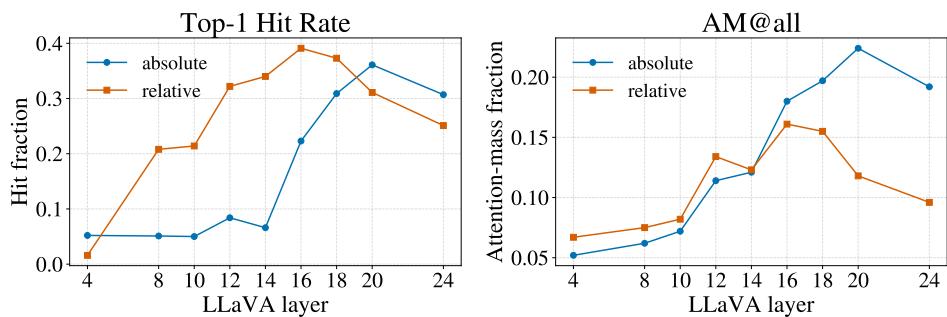
876
877
$$(x', y') = (F_x(x), F_y(y)),$$

878
879
880
881
882
883
with monotone 1-D CDFs, its Jacobian is strictly diagonal. Consequently, (1) each ViT
patch is subjected only to axis-wise scaling ($F'_x(x)$ along x , $F'_y(y)$ along y), and (2)
borders remain orthogonal. Hence, what the encoder encounters at the token level in
AttWarp is identical in form to the resize-then-crop augmentation used during CLIP
training, albeit with locally different scale factors. Thus, rectilinear warping keeps
token statistics on the same manifold learned during pre-training.884
885
2. Did you use a single layer or multiple layers? The formulation you presented appears to be
for a multi-layer approach.886
887
888
889
Ans. While the formulation demonstrates a general multi-layer capability, our quantitative
and qualitative analyses indicated that a single layer (specifically the 16th layer for
the Qwen and the 20th layer for the LLaVA model) performed better than an average
across all layers in this instance.890
3. If the attention is highly inaccurate, then what?891
892
893
894
895
896
Ans. If the attention map is highly inaccurate, both the base MLLM and the MLLM
with AttWarp tend to fail and produce incorrect answers. However, in most practical
scenarios, the attention maps are only moderately suboptimal. In these cases, AttWarp
is particularly effective—its warping mechanism enhances the attention distribution,
leading to improved performance. Detailed experimentation focused on robustness are
presented in Appendix B.5.897
4. What happens when there are multiple objects of focus? Does AttWarp work in that case?898
899
900
901
902
903
Ans. As demonstrated in Appendix E.1, we conducted a study on cases with multiple
objects of focus—ranging from two up to five distinct regions. Our results show that
AttWarp consistently expands the target regions, with 89% of the annotated bounding
boxes exhibiting increased area after transformation. In many instances, all objects of
interest (up to five per example) are effectively expanded, highlighting the method’s
robustness in handling complex, multi-object referential expressions.904
5. What if image warping largely distorts the image?905
906
907
908
909
910
911
Ans. The extent of distortion introduced by image warping is governed by the
choice of transformation function \mathcal{T} . For the transformations we utilize— $\mathcal{T} \in \{\text{sqrt, identity, square}\}$ —significant distortion is not observed in practice. Extreme
warping only occurs if the attention map is highly concentrated on a single, small
region, which is rare and typically corresponds to extremely fine-grained queries. In
such cases, the resulting distortion is often advantageous, as it further magnifies the
most relevant region, thereby improving the model’s ability to answer the query.912
913
914
915
916
917

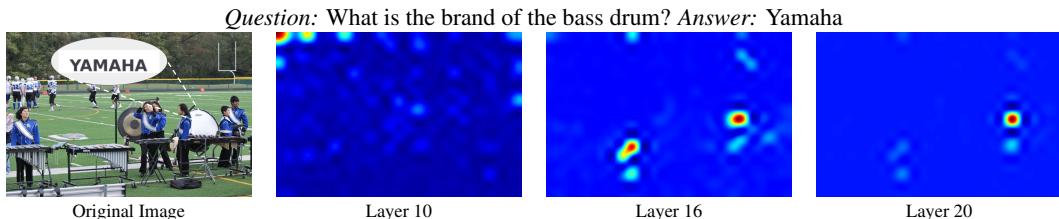
918 **B ABLATION STUDY ON ATTENTION SCORE MATRIX**
919920 **B.1 LAYER SELECTION ACROSS MLLM LAYERS**
921

922 To select the attention map that best captures the visual semantic signal of fine-grained image
923 details with respect to the query, we compared two normalization schemes—*absolute* and *relative*
924 attention—through every cross attention layer of *LLaVA-1.5-7B*. The absolute map is the raw cross-
925 attention weight assigned to each image token when the model answers the question, whereas the
926 relative map divides this weight by a *caption-only* baseline obtained from an auxiliary forward pass
927 that asks the model to “describe the image briefly.” The intuition is that relative normalization
928 suppresses static scene priors (e.g. sky or grass that invariably attract some attention) and instead
929 highlights query-specific regions. We benchmark both variants on 2000 TextVQA validation images
930 using **Pointing-Game Top-1 Hit** (Zhang et al., 2016) and **AM@all** (Wang et al., 2020). Fig. 8
931 presents a *layer-wise* localization analysis for LLaVa 7b: scores increase through the mid-network
932 for both schemes, peak with the *absolute* map from layer 20 (TOP-1 ≈ 0.36 , AM@ALL ≈ 0.22),
933 and then plateau or decline; relative attention helps in earlier layers but never exceeds the absolute
934 variant beyond layer 18 and likewise weakens past layer 20. These findings indicate that the deeper
935 blocks progressively concentrate the model’s “attention budget” on the object of interest, but the final
936 layers begin to divert attention to token generation. Consequently, all LLaVA-based warps in our
937 work AttWarp employ the absolute attention of layer 20. An analogous sweep on *Qwen-VL* reveals
938 its optimum at layer 16. The qualitative evidence in Fig. 9 mirrors the quantitative trend: attention
939 tightens around the “yamaha” logo on the bass drum from layers 10 to 20, underscoring the utility of
940 this layer-wise study – previously unexplored in prior works and guiding our choice of layer 20 for
941 LLaVA-based models and layer 16 for *Qwen-VL* as the most reliable attention sources.

942 **Practical Considerations against Adaptive Layer Selection -** Dynamically evaluating multiple
943 layers per query would add significant runtime overhead, undermining the low-latency goals of
944 test-time adaptation. Our focus in this study was to validate the core AttWarp mechanism under the
945 best static choices.



947 **Figure 8:** Layer-wise localisation quality of LLaVA-1.5-7B cross-attention maps on TextVQA Zhang et al.
948 (2025a) images. Curves report **Top-1 Hit Rate** (Zhang et al., 2016) (left) and **AM@all** (Wang et al., 2020)
949 (right) for absolute (blue) and relative (orange) attention.



950 **Figure 9:** Below (left to right): (a) Original image, zoomed in at yamaha (answer), (b) to (d) attention maps
951 captured from layers 10, 16, 20 respectively. As can be seen, the attention localization improves drastically from
952 layer 10 to 20, indicating improved query-specific spatial understanding.

LLaVA					
Method	TextVQA	GQA	MMMU	POPE	DocVQA
Base MLLM	49.3	60.5	36.9	85.3	18.1
AttWarp (Layer 20)	58.1	63.7	40.4	87.5	25.5
AttWarp (Layer 24)	58.4	62.8	39.1	87.1	24.9
Qwen					
Method	TextVQA	GQA	MMMU	POPE	DocVQA
Base MLLM	81.0	62.4	47.3	86.1	77.3
AttWarp (Layer 16)	84.7	64.0	50.4	87.4	84.1
AttWarp (Layer 24)	84.2	63.6	49.8	87.6	83.9

Table 4: Performance of AttWarp across different attention layers on LLaVA and Qwen backbones.

Robustness to Attention Layer Choice (Table 4): These results reveal a simple and robust rule for selecting attention maps in MLLMs. As previously observed in CNNs and ViTs, MLLMs too exhibit a similar trait: deeper layers produce attention maps that become increasingly centered on objects of interest. Specifically, we find that attention maps from layers with depth ≥ 15 in MLLMs consistently yield task-appropriate, region-of-interest-aligned maps for AttWarp. To validate robustness to the specific choice of layer, we re-run AttWarp using a new, much deeper layer (24) for both LLaVA and Qwen and evaluate on five benchmarks. Across all datasets, the layer-24 variants closely match or slightly exceed the gains obtained with layers 16/20, demonstrating that AttWarp is largely insensitive to the precise layer index as long as it is sufficiently deep. In practice, we therefore recommend using an earlier deep layer (e.g., 16–20) to reduce computational overhead while preserving the full performance benefits.

B.2 ATTENTION AGGREGATION STUDY

Robustness of Attention Head Aggregation Each attention head captures a different aspect of the multimodal interaction (e.g., color, shape, texture, positional cues). Beyond selecting the optimal layer, we analyzed the strategy for aggregating attention across heads. Our analysis confirms that uniformly averaging attention from all heads provides the most robust and informative signal for warping. As shown in Table 5, this approach significantly outperforms alternatives like max-pooling or using random subsets of heads, validating our data-driven design choice.

Table 5: Effect of attention head aggregation strategy on TextVQA accuracy (%). Our method of averaging all heads is the most effective and robust.

Aggregation Methodology	TextVQA Accuracy
Mean over all 32 heads (Ours)	58.1
Max-pooling across 32 heads (token-wise)	55.3
Mean over 8 randomly selected heads	54.6
Random single head (re-sampled per run)	51.9

B.3 PIXEL-SPACE VS. FEATURE-SPACE WARPING

An alternative to our proposed image warping is to directly manipulate the internal feature space of the multimodal LLM. We explored this option by injecting a bias into the hidden states after normalization in the first cross-attention block:

$$\mathbf{h}' = \text{LayerNorm}(\mathbf{h}) + \lambda \cdot \mathbf{b},$$

where \mathbf{b} is constructed from attention weights and λ controls the bias strength. Although conceptually attractive, this approach proved unstable in practice: performance gains were marginal (only **+0.6%** on TextVQA). Moreover, the results of the feature space warping were highly sensitive to the choice of λ . As highlighted in previous research Sun et al. (2025), direct manipulations at the internal feature level inherently risk causing significant distribution shifts that interact unpredictably with architectural components such as pre-or post-normalization layers.

In contrast, our rectilinear warping *pixel-space* avoids these pitfalls and offers several concrete benefits.

- **Stability and robustness:** The internal computation of the model remains unchanged, ensuring consistent behavior across datasets and queries.
- **Interpretability:** The warp is visually transparent: the expanded regions correspond directly to areas of high attention, allowing intuitive inspection and debugging.
- **Architecture agnostic:** No architectural modifications are needed, making the approach compatible across diverse MLLMs and even with external attention sources.
- **Geometry preservation:** The axis-wise CDF warp expands relevant areas while maintaining global structure and relative spatial layout, safeguarding spatial reasoning ability.

B.4 EFFECT OF THE TRANSFORM HYPERPARAMETER (\mathcal{T}) IN ATTWARP

To assess the role of the transform function \mathcal{T} , we performed an ablation across several functional forms. Across both TextVQA and GQA (Tables 6 and 7), results are stable for simple choices such as square root, identity, square, and cubic. The identity and square transforms consistently achieve the best accuracy (58.1–58.3 on TextVQA; 63.3–63.5 on GQA), while alternatives yield only slightly lower scores. These findings show that AttWarp is robust to the choice of \mathcal{T} and that the identity transform serves as a strong default—delivering state-of-the-art accuracy without dataset-specific tuning and outperforming competitive baselines such as ViCrop.

Table 6: Effect of the transform function \mathcal{T} on TextVQA accuracy (%).

Transform	Accuracy (%)
LLaVA Baseline	49.3
\mathcal{T} = sqrt	56.8
\mathcal{T} = Identity	58.1
\mathcal{T} = square	58.3
\mathcal{T} = cube	57.8

Table 7: Effect of the transform function \mathcal{T} on GQA accuracy (%).

Method	LLaVA	Qwen-VL
Baseline Model	60.5	62.4
ViCrop	60.9	60.6
AttWarp (Identity)	63.3	63.5
AttWarp (sqrt)	63.7	64.0

B.5 IMPACT OF ATTENTION BIAS AND ROBUSTNESS TO CORRUPTIONS AND ADVERSARIAL PERTURBATIONS

A natural concern is that the performance of AttWarp depends on the underlying pretrained model. If the attention is biased—for example, by over-focusing on high-frequency features—this may limit the generalization of our method. Similar to other test-time adaptation (TTA) approaches for multimodal LLMs, such as ViCrop, our method leverages attention maps to highlight query-relevant regions and is therefore subject to the same dependency. It is important to emphasize, however, that AttWarp is designed as a TTA mechanism to enhance accuracy for visual question answering (VQA), rather than a debiasing approach to correct distribution shift in the model parameters.

To examine robustness under conditions where attention may be less reliable, we conducted additional experiments using synthetic corruptions following the standard IMAGENET-C protocol Hendrycks & Dietterich (2019). Specifically, we evaluated impulse noise, Gaussian noise, and shot noise to depict aggressive high-frequency injections applied to the TextVQA dataset. All corrupted images were resized to the standard 512×512 resolution for evaluation, and no retraining or adaptation beyond test-time warping was performed.

Table 8: Accuracy (%) of LLaVA and LLaVA+AttWarp under different image corruptions, following the IMAGENET-C protocol.

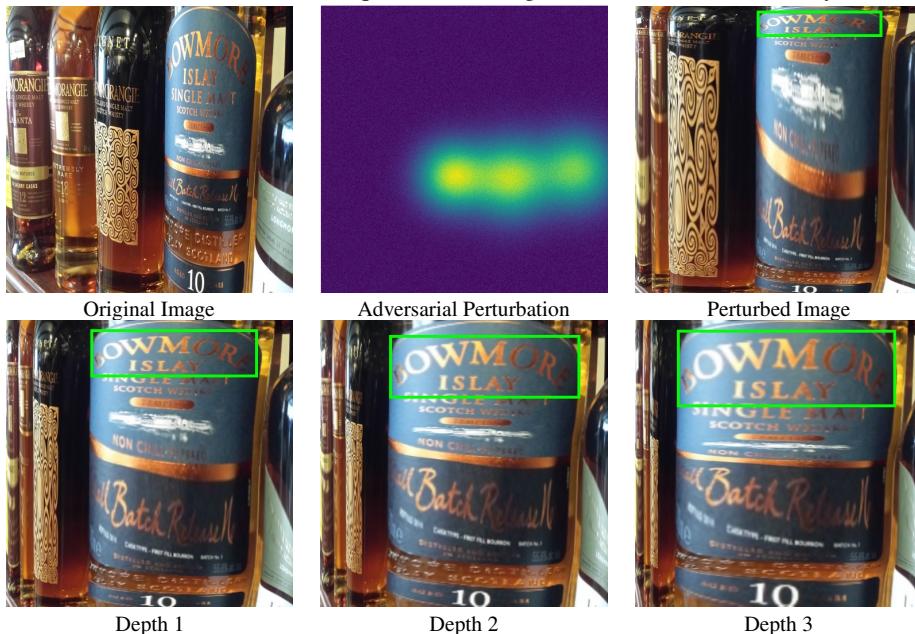
Corruption	LLaVA	LLaVA + AttWarp
Impulse noise	36.8	40.4
Gaussian noise	37.6	41.0
Shot noise	36.0	39.8

These results demonstrate that while AttWarp inherits the biases of the attention maps it uses, it still provides consistent accuracy gains even under corruption. This analysis shows that AttWarp is not a

1080 debiasing method but remains a reliable and effective TTA mechanism to improve VQA performance
 1081 in both clean and noisy settings.
 1082

1083 **Robustness to Adversarial Perturbations.** To further probe resilience, we introduce targeted
 1084 adversarial perturbations specifically engineered to misdirect the MLLM’s initial attention distribution
 1085 away from semantically pertinent regions (Fig. 10, top). This adversarial setting provides a stringent
 1086 test of AttWarp’s iterative refinement capability (Sec. 3.2). As illustrated in Fig. 10 (bottom), while
 1087 the initial adversarial attack successfully perturbs attention and induces a faulty warp, subsequent
 1088 iterations progressively correct the transformation, re-aligning focus with the relevant visual content.
 1089 This behavior highlights a critical self-correction mechanism: iterative AttWarp is able to recover
 1090 from compromised attention signals, ensuring more reliable grounding even under adversarial
 1091 conditions.
 1092
 1093
 1094

1095 *Question:* What brand liquor is on the right? *Answer:* Bowmore Islay



1106 **Figure 10:** (top row) the original image, the adversarial perturbation designed to corrupt attention, and the
 1107 resulting perturbed image in which the task-relevant region is shrunk; (bottom row) outputs of AttWarp at Depth
 1108 1–3, illustrating how our method progressively overcomes the interference to refocus on the relevant region.
 1109
 1110
 1111
 1112
 1113
 1114

1124 C TASK SPECIFIC ANALYSIS OF ATTWARP

1129 C.1 TASK CATEGORIES, BENEFITS, AND LIMITATIONS

1132 AttWarp magnifies the query-relevant region while preserving relative object positions through its
 1133 rectilinear design. This leads to clear gains in three categories: **fine-grained perception**, **spatial
 reasoning**, and **hallucination mitigation**. Table 9 reports the improvements across datasets.
 1134

1134
1135**Table 9:** Category-level improvements of AttWarp. Accuracy in %. Best results in **bold**.

Task Category	Dataset	LLaVA	AttWarp	Δ
Fine-grained perception	TextVQA (fine-grained)	49.9	59.6	+9.3
	DocVQA (table/list/form/handwritten)	13.6	19.5	+5.9
Spatial reasoning	TextVQA (spatial)	54.5	64.9	+10.4
	DocVQA (layout)	29.4	37.6	+8.2
	MMMU (spatial)	38.2	44.8	+6.6
	GQA (relation)	51.5	56.4	+4.9
Hallucination mitigation	POPE	85.3	87.5	+2.2
	GQA (object)	86.1	89.4	+3.3

1145

1146

1147 Across GQA’s 15 semantic sub-categories, AttWarp improves performance consistently. The only
 1148 exception is a marginal drop of **0.2%** in the *compare-attribute* sub-category involving direct size
 1149 comparisons. This remains the sole degradation observed, underscoring the reliability of the approach.
 1150

1151 C.2 DISTORTION AND UPPER-BOUND CONCERNs

1152 A concern is whether warping distorts shapes or positional reasoning. By design, AttWarp preserves
 1153 grid alignment and relative geometry. This is reflected in GQA’s semantic dimensions shown in
 1154 Table 10, where accuracy improves across all categories.
 1155

1156

1157

Table 10: Performance on GQA semantic categories. Accuracy in %.

Category	LLaVA	AttWarp
Relation	51.5	56.4 (+4.9)
Attribute	67.8	69.3 (+1.5)
Category	51.7	55.1 (+3.4)
Object	86.1	89.4 (+3.3)
Global	62.5	65.5 (+3.0)

1166

1167

1168 Manual sub-sampling from other datasets confirms similar gains. AttWarp boosts positional reasoning
 1169 and object shape sensitivity across TextVQA, DocVQA, and MMMU as shown in Table 11.
 1170

1171

1172

Table 11: Cross-dataset improvements on positional reasoning and object shape. Accuracy in %.

Dataset	Category	LLaVA	AttWarp	Δ
TextVQA	Positional reasoning	54.5	64.9	+10.4
DocVQA	Layout (positional)	29.4	37.6	+8.2
MMMU	Object shape	36.6	40.7	+4.1
DocVQA	Diagram (object shape)	18.8	22.2	+3.4

1178

1179

1180

1181 These results show that AttWarp enhances rather than constrains reasoning involving spatial relations,
 1182 positions, and shapes.
 1183

1184

1185

1186

1187

C.3 GLOBAL CONTEXT AND WARPING INTENSITY

1188

1189

1189 For questions requiring global context, AttWarp adapts the strength of warping to the input. We
 1190 quantify this with a warping intensity metric, defined as the mean log-change of grid-cell Jacobians.
 1191 Fine-grained queries undergo stronger magnification, while global queries are only mildly affected.
 1192

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Table 12: Warping intensity by question type. Values are mean log-change of Jacobian determinants.

Category	log_change_mean
Attribute	0.19
Category	0.18
Global	0.16
Object	0.19
Relation	0.22
Fine-grained	0.26

Fine-grained questions produce a $\sim 30\%$ area change, amplifying small detail-rich regions. Global questions produce only a $\sim 17\%$ change, preserving scene layout while still enabling mild emphasis. This explains performance increase on global tasks.

C.4 TASK-WISE COMPARISON WITH ViCROP

ViCrop uses dual inputs (original and cropped images). A natural concern is whether this dual-image design provides an advantage. Table 13 shows that AttWarp consistently outperforms ViCrop across semantic categories on GQA.

Table 13: Comparison with ViCrop on GQA semantic categories. Accuracy in %. Best results in **bold**.

Category	LLaVA	ViCrop	AttWarp
Relation	51.5	51.9	56.4
Attribute	67.8	68.2	69.3
Category	51.7	52.0	55.1
Object	86.1	86.3	89.4
Global	62.5	63.4	65.5

This analysis shows that AttWarp not only avoids the pitfalls of dual-image inputs but also provides consistent improvements across all semantic dimensions.

C.5 PERCEPTUAL GEOMETRY PRESERVATION

LPIPS (Zhang et al., 2018) measures distances between image pairs in a deep feature space extracted from a pretrained CNN (we use VGG16), with channel-wise weights calibrated on human perceptual judgments; as a result, Euclidean distances in this feature space approximate geodesic distances on a perceptual manifold, i.e., they reflect meaningful changes in structure, texture, and semantics (Zhang et al., 2018). For our rectilinear warp (AttWarp), we obtain a significantly lower (i.e., better) LPIPS score than the non-rectilinear warp baseline (0.14 vs. 0.38), demonstrating that AttWarp effectively preserves perceptual geometry.

D ADDITIONAL EXPERIMENTS

D.1 EXTENDED GENERALIZATION EXPERIMENTS

We further examined the generalizability of AttWarp by applying it to two additional multimodal LLMs: the recently released **InternVL-3 8B** (opensourced in April 2025) Zhu et al. (2025) and **InstructBLIP** Dai et al. (2023). For both models, attention maps were extracted following the same protocol used for LLaVA and Qwen.

InstructBLIP - On TEXTVQA and GQA, AttWarp improves over both the baseline InstructBLIP and the ViCrop baseline. Moreover, the chained variant (AttWarp-Chain) yields additional gains, confirming that iterative refinement can further sharpen attention-driven warping.

InternVL-3 8B - On TEXTVQA and GQA, AttWarp consistently outperforms both the strong baseline InternVL-3 and ViCrop. These results indicate that the benefits of our approach transfer effectively even to cutting-edge architectures trained at scale.

Table 14: Generalization of AttWarp across **InstructBLIP** and **InternVL-3**. Evaluation metric: accuracy (%). Best results in **bold**.

Method	InstructBLIP		InternVL-3 8B	
	TextVQA	GQA	TextVQA	GQA
Baseline	35.2	49.4	80.2	61.4
ViCrop	46.6	49.7	82.7	63.9
AttWarp	47.8	51.3	84.6	65.9
AttWarp-Chain	49.6	52.4	—	—

Together, these extended evaluations reaffirm the robustness of AttWarp across a diverse set of multimodal LLMs—spanning both established architectures (InstructBLIP) and the latest state-of-the-art models (InternVL-3). The consistent improvements over strong baselines and ViCrop highlight AttWarp as a broadly applicable, model-agnostic test-time adaptation mechanism.

D.2 FINE-GRAINED AND CATEGORY-WISE RESULTS ON DOCVQA

To further evaluate the capacity of AttWarp for both fine-grained recognition and global structural reasoning, we analyze its performance on the DocVQA dataset across diverse structural categories. These include *Free Text*, *Table*, *Layout*, *Form*, *Handwritten*, and *Diagram*, which together capture the spectrum of challenges in document question answering. All experiments are conducted on images resized to 512×512 , since original DocVQA images are much larger. This resizing ensures computational feasibility across methods. In particular, for dual-image methods such as ViCrop, processing full-resolution images creates prohibitively many tokens, making the resized setting especially relevant for fair comparison.

Table 15: Category-wise accuracy (%) on DocVQA using Qwen2.5-VL-7B as the base MLLM. All results are reported on images resized to 512×512 to ensure computational feasibility across methods. AttWarp consistently outperforms all baselines across every document structure category except Yes/No, where ViCrop performs best. Gains span both fine-grained (Form, Hand-written, Table) and global (Layout, Diagram) reasoning categories.

Method	Overall	Free Text	Table	Layout	Form	Hand-written	Diagram	Others	Image	Y/N
Qwen	77.8	78.1	76.3	85.6	75.8	63.2	79.6	80.0	71.4	82.1
API	68.4	68.9	60.8	79.5	68.9	60.9	72.1	80.0	66.1	64.3
ViCrop	82.3	83.0	78.8	87.2	77.9	68.3	80.8	84.8	76.5	96.4
AttWarp	84.1	86.2	79.1	89.9	83.5	71.4	81.5	86.7	82.1	92.9

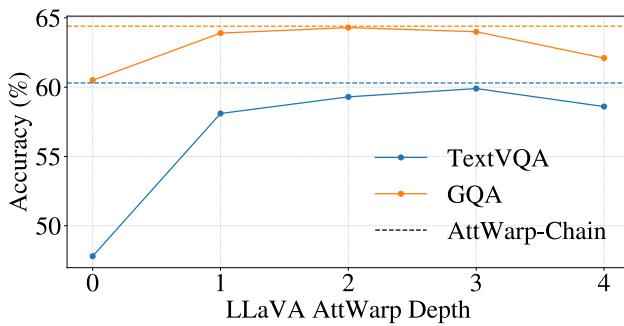
Table 15 shows that AttWarp delivers consistent and significant improvements over all baselines, not only in overall accuracy but also within every category. The method achieves particularly strong gains in challenging fine-grained settings such as *Form* (+5.6) and *Handwritten* (+8.2), which require precise localization and interpretation. At the same time, improvements in global categories such as *Layout* (+2.7) and *Diagram* (+0.7) highlight its ability to capture holistic structural cues. The only exception is the Yes/No category, where ViCrop slightly outperforms our approach. Importantly, these results were computed in the resized 512×512 setting, which avoids the prohibitive token explosion of large original images for dual-image methods like ViCrop, thereby ensuring fair and tractable comparisons. Overall, these findings demonstrate the robustness and versatility of AttWarp for document question answering tasks spanning both local detail and global structure.

D.3 STABILITY AND TERMINATION IN ATTWARP-CHAINS

This subsection provides additional results complementing the discussion in Section 3.2, where we introduced an adaptive stopping criterion for iterative warping. Since AttWarp-chains operates

1296 recursively, each warped input influences subsequent attention maps. While this iterative refinement
 1297 initially sharpens focus and improves accuracy, excessive iterations can lead to over-expansion, noise,
 1298 and performance degradation. This underscores the importance of a principled termination rule.
 1299

1300 Figure 11 illustrates this phenomenon on TEXTVQA (blue curve) and GQA (orange curve). Accuracy
 1301 improves significantly in the first few iterations—rising from $\sim 47\%$ to $\sim 60\%$ on TEXTVQA, and
 1302 from $\sim 60\%$ to $\sim 64\%$ on GQA—but then plateaus and declines once the depth exceeds two to three
 1303 iterations. This decline reflects the recursive instability of fixed-length warping. By contrast,
 1304 the adaptive AttWarp-Chain (dashed lines) consistently outperforms all fixed-depth settings by halting
 1305 precisely when attention distributions stabilize.
 1306
 1307



1318
 1319 **Figure 11:** Impact of iterative warping depth on accuracy for TEXTVQA and GQA. Fixed-length warping
 1320 initially improves performance but degrades with excessive iterations due to recursive instability. The adaptive
 1321 AttWarp-Chain (dashed lines), guided by the KL divergence stopping criterion, consistently achieves the best
 1322 accuracy while avoiding over-warping.
 1323

1324 The adaptive stopping rule monitors changes in successive attention maps and terminates the chain
 1325 once the distributions converge, as formalized by the KL divergence threshold in Eq. 6. This ensures
 1326 that refinement halts exactly when attention has sufficiently concentrated on query-relevant regions.
 1327 In practice, AttWarp-Chain not only prevents instability but also achieves higher accuracy than
 1328 any fixed-depth configuration while being computationally more efficient by avoiding unnecessary
 1329 iterations.
 1330
 1331

1332 D.4 EXTERNAL ATTENTION MAPS

1333
 1334 **AttWarp is compatible with attention from external models.** All prior experiments applied
 1335 AttWarp using *internal* cross-attention maps, *i.e.* extracted from the same MLLM used for down-
 1336 stream inference. Here we probe another axis of generalization: can model A (e.g., LLaVA) benefit
 1337 from an image warp constructed from the attention maps of model B? Demonstrating such flexibil-
 1338 ity would establish AttWarp as a model-agnostic, general-purpose mechanism independent of the
 1339 substrate attention provider.
 1340

1341 To investigate this, we source attention maps from two different classes of external models: (i) the
 1342 text-to-image generative backbone Stable Diffusion 2.1 Rombach et al. (2022), and (ii) a strong
 1343 multimodal LLM, Qwen-VL Yang et al. (2024a). The warped images are then processed by LLaVA for
 1344 downstream tasks. Results in Table 16 show that AttWarp consistently improves performance across
 1345 both TEXTVQA and GQA, even when relying on external substrate attention. Notably, MLLMs serve
 1346 as better sources than generative vision models: attention from Qwen-VL yields the strongest gains
 1347 (+10% on TEXTVQA, +3.4% on GQA), while Stable Diffusion also provides improvements though
 1348 of smaller magnitude (+6.7%, +2.2%). Interestingly, even external MLLM attention outperforms the
 1349 base LLaVA, underscoring the benefit of stronger substrate models. Complementary experiments
 transferring attention in the opposite direction *i.e.* from a weaker LLaVA-7B to strengthen a larger
 LLaVA-34B—are reported in Appendix E.3.

1350 **Table 16:** Performance of AttWarp using attention from internal vs. external models. Evaluation metric:
 1351 accuracy (%). Best results in **bold**.

Method (Source of Attention)	TextVQA	GQA
Base LLaVA	49.3	60.5
+ AttWarp (Internal: LLaVA)	58.1	63.7
+ AttWarp (Stable Diffusion Rombach et al. (2022))	56.0	62.7
+ AttWarp (Qwen-VL Yang et al. (2024a))	59.3	63.9

1360 These findings demonstrate that AttWarp is not restricted to internal attention but generalizes robustly
 1361 across external sources as well, with stronger multimodal LLMs providing the most effective substrate
 1362 attention maps.

E BEYOND STANDARD VQA: EXTENDED EXPERIMENTS

1366 In this section, we present several forward-looking experiments to probe the broader applicability
 1367 and potential of AttWarp. We demonstrate its effectiveness in expanding all query-relevant regions,
 1368 generalizing to open-vocabulary object detection, and even improving the performance of larger
 1369 models using attention maps from smaller ones. Collectively, these results highlight the versatility
 1370 and extensibility of our approach beyond standard VQA settings.

E.1 ATTWARP EXPANDS ALL QUERY-RELEVANT REGIONS

1374 We evaluate whether AttWarp effectively enlarges query-relevant regions using ground-truth bound-
 1375 ing boxes from prior datasets. Results show that our method consistently expands the salient regions,
 1376 both for single- and multi-object queries.

1378 **Table 17:** Expansion of query-relevant bounding boxes under AttWarp.

Dataset	% Boxes Expanded	Mean Area Increase
TextVQA (single-region)	94.0	+76%
gRef (multi-region)	88.6	+39%

1386 For TextVQA (Zhang et al., 2025a), AttWarp expanded **94%** of salient bounding boxes, with an
 1387 average area increase of **76%**, confirming its ability to magnify relevant image regions.

1388 To test cases with multiple objects of focus, we used the gRef dataset (Liu et al., 2023a), which
 1389 provides multiple ground-truth bounding boxes per query. Here, AttWarp expanded **88.6%** of target
 1390 boxes with a mean area increase ratio of **1.39 (+39% area)**.

1391 These results demonstrate that AttWarp reliably enlarges query-relevant regions, including complex
 1392 multi-object settings, while preserving their spatial grounding.

E.2 OPEN-VOCABULARY OBJECT DETECTION

1397 We next evaluate the versatility of AttWarp on Open-Vocabulary Object Detection (OVOD), a setting
 1398 that extends beyond VQA. In OVOD, the goal is to localize objects in images based on free-form
 1399 referring expressions rather than a fixed label set. We use a dataset of 10,000 examples to test whether
 1400 our approach improves localization under this challenging setup.

1401 Our pipeline applies AttWarp with LLaVA-7B, guided by each referring expression. The warped
 1402 image is then processed by the LISA-LLaVA-7B framework Lai et al. (2024) to predict bounding
 1403 boxes, which are mapped back to the original image space using the inverse warp. This enables direct
 IoU-based evaluation against ground-truth annotations.

Table 18: OVOD performance with LISA-LLaVA-7B. Accuracy in %.

Model	Original Images	Warped Images	Δ
LISA-LLaVA-7B	54.0	61.0	+7.0

Representative qualitative examples are shown in Fig. 12. These results confirm that AttWarp extends effectively to OVOD, improving localization accuracy by +7%.

Referring expression: **Man in black facing us**



Referring expression: **Girl with yellow shirt**

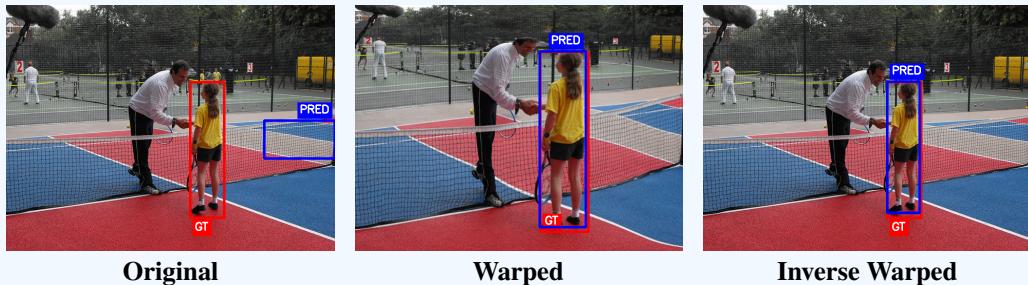


Figure 12: Qualitative OVOD results with LISA-LLaVA-7B. Each example shows the original, warped, and inverse-warped images. Predicted bounding boxes are in [blue](#), ground-truth annotations in [red](#).

E.3 LEVERAGING SMALLER MODELS TO IMPROVE LARGER MODELS

We test whether attention maps from a smaller model can enhance the performance of a larger model. Specifically, we extract attention maps from **LLaVA-1.5-7B** to warp input images, which are then processed by the stronger **LLaVA-1.6-34B** model.

Table 19: TextVQA accuracy (%) of LLaVA-1.6-34B with and without image warping using attention maps from LLaVA-1.5-7B.

Model	Without Warping	With Warping (7B maps)	Δ
LLaVA-1.6-34B	72.6	74.1	+1.5

This result shows that weaker models' attention maps can be repurposed to improve the performance of stronger models, highlighting a novel direction for leveraging model complementarities.

F ATTENTION REDISTRIBUTION: REDUCTION IN MODEL UNCERTAINTY

The qualitative examples below exemplify the attention-redistributive effect of AttWarp. Our approach consistently improves and redirects the focus of the model toward query-specific image regions. This can be seen by the increased coverage of relevant bounding boxes and the sharper alignment of attention peaks with target regions, in contrast to the diffused or off-target patterns often

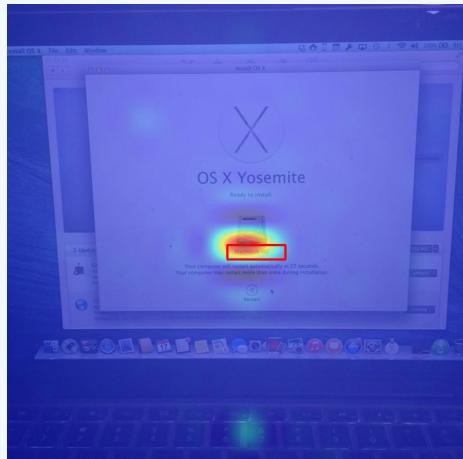
1458 observed in the original attention maps. These visual trends support the quantitative gains reported
 1459 in the main paper as discussed in Sec. 4.3 and Tab. 6d, where we showed improved localization on
 1460 TEXTVQA.

1461 **Localization precision on GQA.** To further validate query-specific focus on more datasets, we
 1462 extend this study to the GQA dataset, which provides bounding-box annotations. Both Pointing
 1463 Game accuracy and AM@all improve after applying AttWarp. Post-warp, the Pointing Game score
 1464 rises from 0.412 to **0.419**, and AM@all from 0.139 to **0.154**. These gains align with the improvements
 1465 observed on TEXTVQA and confirm that the rectilinear warp sharpens spatial localization across
 1466 different benchmarks.

1467 Taken together, the evidence from both TEXTVQA and GQA confirms that AttWarp not only pre-
 1468 serves global distributional fidelity but also reduces model uncertainty by consistently redistributing
 1469 attention mass toward query-relevant regions.

1470 **Question:** What type of hard drive does the computer have?

1471 **Answer:** Macintosh



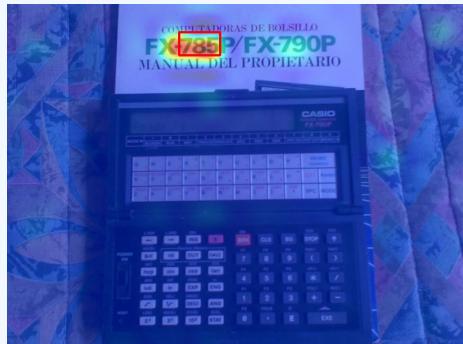
1472 **Original**



1473 **Warped Output**

1474 **Question:** What number follows the first *fx* at the top?

1475 **Answer:** 785



1476 **Original**



1477 **Warped Output**



G IMPLEMENTATION DETAILS

This section provides the implementation details for the experiments presented in Sec. 4.2 and Sec. 4.3. We first describe the setup for the quantitative results and experimental studies shown in Sec. 4.2 and Tab. 1. We then detail the configuration for the ablation studies for Sec. 4.3 (external models). We then include extended results of Tab. 1 and additional ablation analysis.

G.1 QUANTITATIVE EVALUATION SETUP

As described in Sec. 3.4, we extract attention maps from the 20th layer for LLaVA-1.5-7b Liu et al. (2024b) and the 16th layer for Qwen2.5-VL-7b Yang et al. (2024a). The extracted attention maps are resized to the original image resolution for all downstream transformations and evaluation. For dataset-specific transforms in AttWarp, we apply the identity transform to TextVQA Singh et al. (2019), DocVQAMathew et al. (2021), POPELi et al. (2023), and MMMUYue et al. (2024), while for GQA Hudson & Manning (2019), the SQRT transform is used.

All experiments are conducted using four NVIDIA Tesla V100-SXM2 GPUs (16GB each). For LLaVA-1.5-7b, we use the official implementation¹ with an input image resolution of 336×336 and a patch grid of 24×24 . For Qwen2.5-VL-7b, all images are resized to 512×512 due to memory constraints and inference time constraints (e.g. ViCrop for DocVQA takes 90hours of inference time on one H100), which results in a slightly reduced baseline accuracy compared to the numbers reported in the Qwen2.5-VL-7b official release.

We set the KL divergence threshold in AttWarp-Chain to 0.2 (and a compulsory stopping condition of 5 iterations), based on empirical observations. For all baselines (APIYu et al. (2024), FGVPYang et al. (2023b), SoMYang et al. (2023a), and ViCropZhang et al. (2025a)), we use their official implementations. For SoM, each image is segmented with Semantic-SAM (Swin-L). For FGVP, the MLLM (LLaVA/Qwen) is first prompted to output query-specific relevant objects; then SAM (ViT-H) and CLIP (ViT-L/14) masks those regions and produce two inputs: one with the background Gaussian-blurred and one with the masked region highlighted green. Hyperparameters of FGVP and SoM are selected by grid search over 100 randomly sampled images per dataset to select the optimal values. Hyperparameters for SoM: granularity = 3, α = 0.4, text-size = 640; FGVP: blur σ = 5, α = 0.5, IoU threshold 0.86, stability threshold 0.92, and minimum mask area 400. In the case of API and ViCrop, we exactly use the official implementation (API² and ViCrop³) with the same hyperparameters.

Accuracy is used as the evaluation metric for all datasets: for each question-answer pair (q, a) , the prediction is considered correct if the model output exactly matches a , with accuracy computed as the average over all examples. In TextVQA and DocVQA, we do not provide any OCR-extracted tokens

¹<https://github.com/haotian-liu/LLaVA>

²github.com/yu-rp/apiprompting

³github.com/saccharomyces/mllms_know

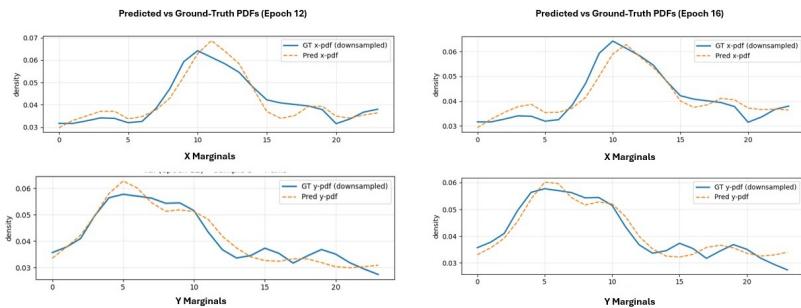
1566 to the MLLMs—only the image and question are given, following the evaluation prompt format
 1567 outlined in the respective papers. Finally, to ensure reproducibility and transparency, we will release
 1568 all code, configurations, and analysis scripts required to reproduce our experiments and results.
 1569

1570 G.2 STABLE DIFFUSION EXPERIMENTAL DETAILS 1571

1572 In Appendix Sec. G.2, we use Stable Diffusion for external attention; here, we detail its experimental
 1573 design. We adapt AttWarp to Stable Diffusion 2.1 (SD-2.1) by first inverting the input image I
 1574 into the model’s latent space with a five-step truncated DDIM schedule [1000, 800, 600, 400,
 1575 200], conditioning on the question q at every step. Each inverted latent is then forwarded through
 1576 ten denoising iterations while recording the cross-attention tensors of the final two UNet layers —
 1577 previously identified by Hertz et al. (2022) as exhibiting particularly sharp token-level groundings.
 1578 Summing these tensors over heads and space yields token-wise energies from which we retain the
 1579 top- $k=20$ tokens, average their channels into a single low-resolution attention map, apply a square-
 1580 root contrast stretch, and—together with its inverse on the one-dimensional marginals—use it to
 1581 drive a single \sqrt{C} warp at 500×500 resolution. As reported in Table 16, this external attention
 1582 improves the vanilla LLava baseline by **+6.7%** on TextVQA, yet remains below the gains from
 1583 internal LLava attention (+8.8%) and from the stronger Qwen-VL (+10.0%). We attribute this gap to
 1584 task mis-alignment: diffusion models are trained for literal scene synthesis (e.g. “*a giraffe in a misty*
 1585 *forest*”), whereas visual question answering centres on self-referential queries (e.g. “*what animal is*
 1586 *in the forest?*”), so the resulting diffusion attention maps, although sharp, do not always highlight
 1587 regions most informative for answering such questions.

1588 G.3 TRAINING SETUP FOR ATTWARP-DISTILL 1589

1590 We use Qwen-2.5VL-7B as the teacher model to extract attention maps and derive corresponding
 1591 marginals for training the student model. Specifically, attention maps from the teacher are first
 1592 extracted, after which category-specific marginals are computed for each dataset. To adaptively scale
 1593 attention distributions according to task granularity, we apply distinct transformations per dataset: a
 1594 square-root transform for DocVQA, a cube-root transform for GQA, and an identity transform for
 1595 TextVQA and fine-grained datasets such as POPE. The student model is trained for a maximum of 20
 1596 epochs, employing early stopping based on validation-set performance. Optimization uses AdamW
 1597 with a batch size of 16, learning rate of 3×10^{-4} , weight decay of 1×10^{-4} , and gradient clipping
 1598 with a maximum norm of 1.0. We show the predicted marginals in Fig. 13.



1609 **Figure 13:** The marginal distributions predicted by the student networks closely align with the ground truth
 1610 marginals, demonstrating robust knowledge transfer.
 1611

1612 G.4 COMPUTATIONAL COST CALCULATION AND ANALYSIS 1613

1614 We present a detailed cost analysis to emphasize the efficiency and practical advantages of AttWarp
 1615 compared to the competitive baseline (ViCrop (Zhang et al., 2025a)). Here, we specifically demon-
 1616 strate results for LLava-1.5v-7b (similar trends hold true for other models, such as Qwen). Our
 1617 approach significantly surpasses the FGVP (Yang et al., 2023b), SoM (Yang et al., 2023a), and API
 1618 (Yu et al., 2024) in accuracy. Therefore, we focus this analysis on ViCrop, the competitive baseline,
 1619 which achieves comparable yet inferior performance to AttWarp across all datasets (see Sec. 4.2 and

1620 Tab. 1). Computational complexity for all MLLMs is calculated following the methodology described
 1621 in (Lin et al., 2025; Chen et al., 2023b;c; 2024a).

1622 Standard LLaVA-1.5 processes a single 336×336 image through a ViT-L/14 encoder, producing 576
 1623 visual tokens and incurring a computational cost of 8.5 TFLOPs. In contrast, AttWarp requires two
 1624 MLLM passes per query: the first pass extracts attention maps up to a specific layer (e.g., the 20th
 1625 layer of LLaVA, where we insert a hook), while the second pass generates the final output from the
 1626 warped image based on these attention maps. The total cost for AttWarp—including both passes and
 1627 the lightweight warping operation—corresponds to 1,152 vision tokens (2×576) and 13.8 TFLOPs.
 1628 The cost of the warping step itself is negligible.

1629 In comparison, ViCrop adopts a more resource-intensive pipeline, requiring three MLLM passes:
 1630 two initial passes to obtain relative attention maps (up to the 14th layer), followed by a final forward
 1631 pass processing both the original and cropped image. This results in a significantly larger number of
 1632 vision tokens (2,304) and a total computational load of 24.2 TFLOPs.

1633 As summarized in Tab. 2, AttWarp delivers substantial computational efficiency. Specifically, it
 1634 reduces the total number of vision tokens by a factor of two relative to ViCrop (1,152 vs. 2,304)
 1635 and requires 1.5 times fewer MLLM passes (2 vs. 3), translating to a $1.8 \times$ reduction in compute
 1636 cost (13.8 TFLOPs for AttWarp vs. 24.2 TFLOPs for ViCrop per query). In terms of memory,
 1637 AttWarp achieves a peak VRAM usage of 15.5 GB, identical to the base LLaVA-1.5v-7b model and
 1638 within the capacity of standard GPUs (e.g., V100s). The peak VRAM required for ViCrop is 22 GB,
 1639 significantly higher (by a factor of 1.46) compared to AttWarp.

1640 These results strongly motivate the use of AttWarp, demonstrating that it not only achieves superior
 1641 accuracy (see Sec. 4.2 and Tab. 1), but also yields substantial savings in computational resources
 1642 compared to approaches like ViCrop.

1649 H QUALITATIVE EXAMPLES

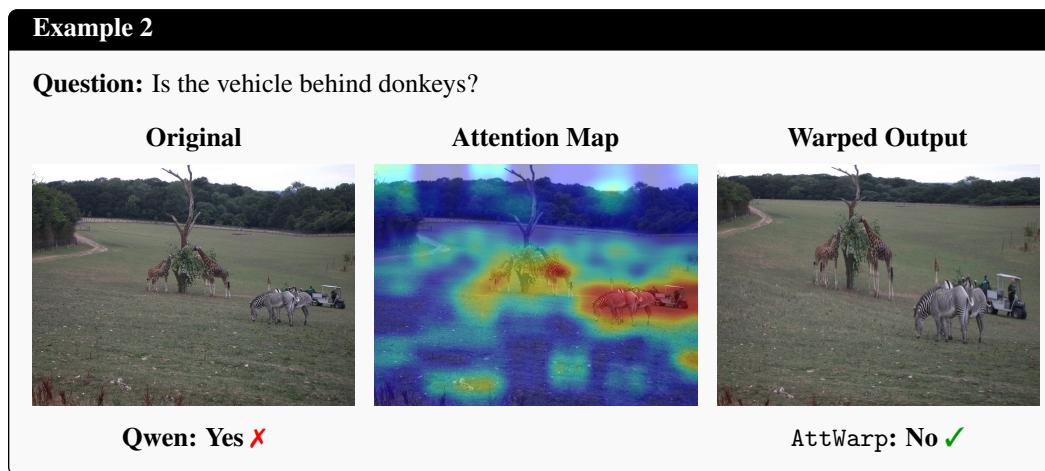
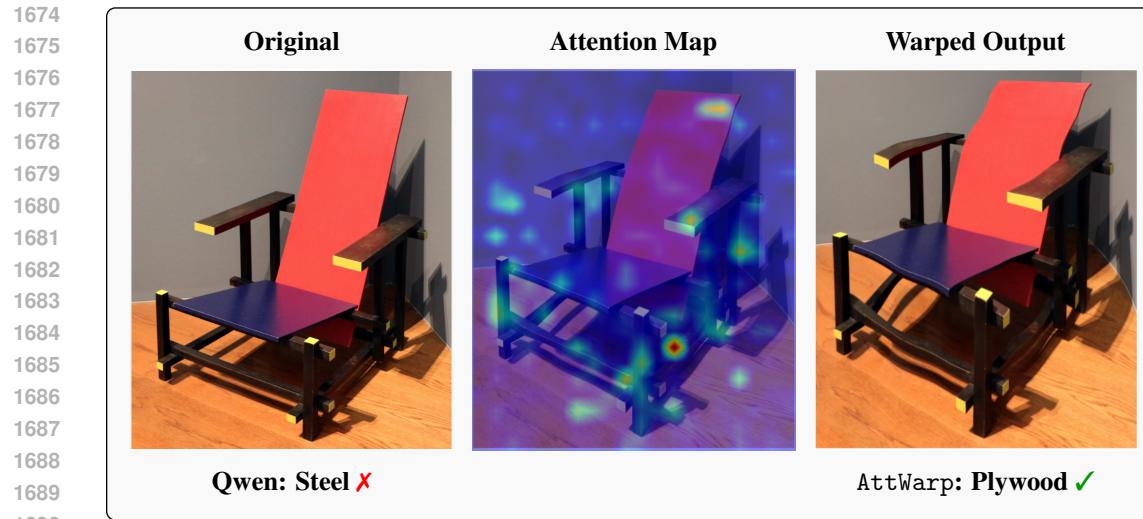
1650
 1651
 1652
 1653
 1654 Here, we present examples demonstrating the effectiveness of both AttWarp and AttWarp-Chain.
 1655 For each case, the predicted answers are displayed beneath the corresponding images. These
 1656 examples highlight scenarios where the base MLLM initially produces an incorrect answer, which is
 1657 subsequently corrected by AttWarp.

1658 For AttWarp-Chain, results are at varying depths (as expected), illustrating how our method adaptively
 1659 refines warping based on the specific query and image content. This depth-wise adaptation
 1660 underscores the robustness and flexibility of our approach, enabling consistently strong performance
 1661 across diverse queries and visual contexts.

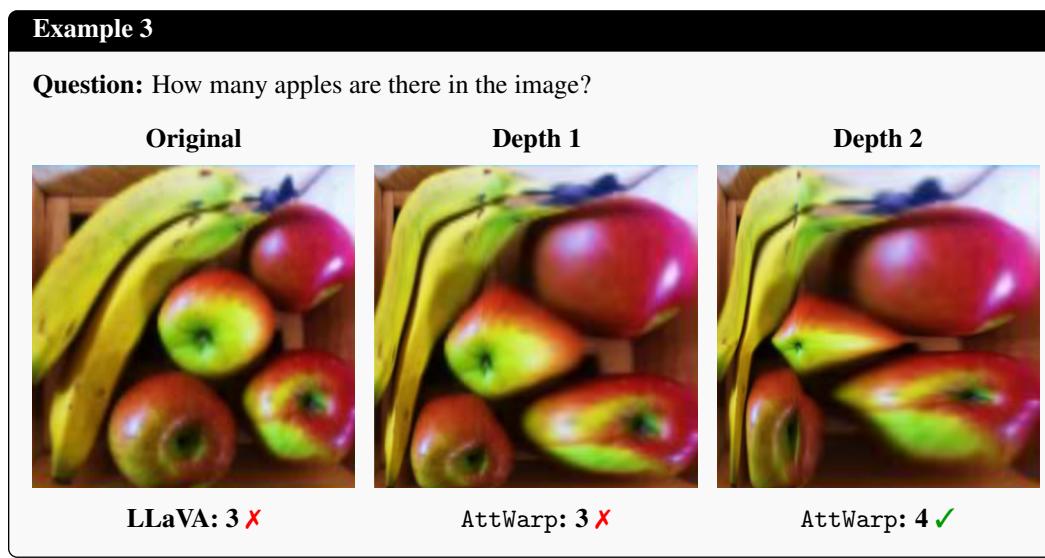
1662 H.1 ATTWARP

1663 Example 1

1664 **Question:** What material is the chair made of?



H.2 ATTWARP-CHAIN



1728
1729**Example 4**

1730

Question: What is a player's number?

1731

1732

1733

1734

1735

1736

1737

1738

1739

1740

1741

1742

1743

1744

1745

1746

Example 5

1747

1748

Question: What year is this whiskey from?

1749

1750

1751

1752

1753

1754

1755

1756

1757

1758

1759

1760

1761

1762

1763

1764

I FID AND KID ANALYSIS

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

1775

1776

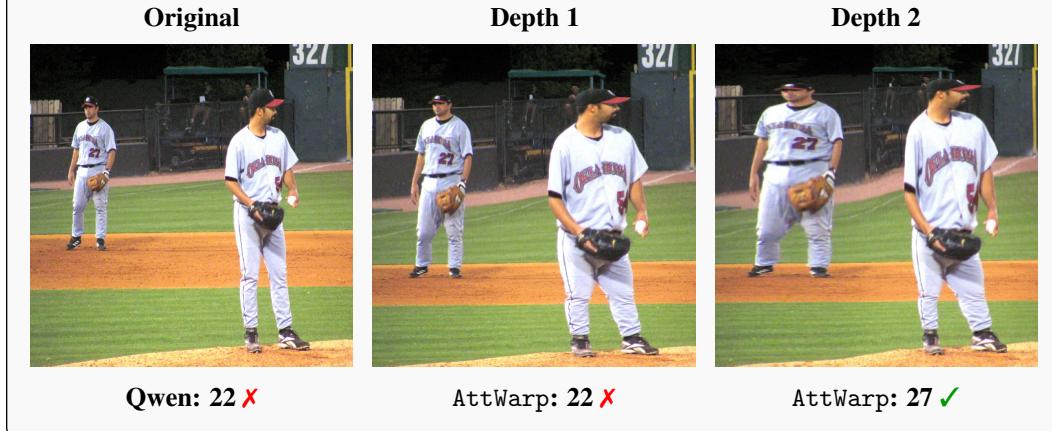
1777

1778

1779

1780

1781

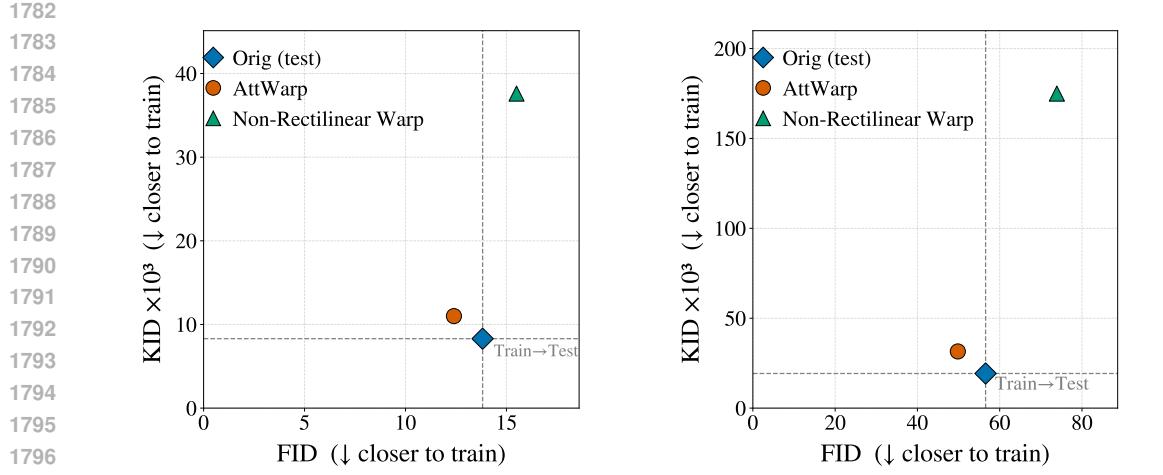


We quantify how closely each test variant remains on the training manifold using two complementary, CLIP-feature based distances: **Fréchet Inception Distance (FID)** and **Kernel Inception Distance (KID)**. Lower values indicate closer alignment to the training distribution.

Why rectilinear warping preserves the distribution. Our warp is axis-aligned and monotone,

$$(x', y') = (F_x(x), F_y(y)),$$

so its Jacobian is diagonal. At the token level this induces only per-axis rescaling of ViT patches while keeping grid orthogonality. This mirrors CLIP's *resize/crop* pre-training augmentation (e.g., RandomResizedCrop), which likewise applies axis-wise scaling before tokenization. Consequently, the pooled CLIP image embeddings for rectilinearly warped images are expected to stay on (or near) the same feature manifold as training images.



(a) FID–KID/ $\times 10^3$ distances between the training distribution and each test variant (Original, AttWarp, Non-Rectilinear Warp) in CLIP **ViT-B/32** space. Points nearer the lower-left are closer to the training manifold.

(b) FID–KID/ $\times 10^3$ distances between the training distribution and each test variant (Original, AttWarp, Non-Rectilinear Warp) in CLIP **ViT-L/14** space.

Figure 14: Comparison of Train→{Test, AttWarp, Non-Rectilinear} FID–KID across two ViT backbones.

Setup. We reuse 12,000 GQA-TRAIN images to define the training reference used throughout the distributional analyses (also used for the Mahalanobis study in Section 4.3). Image embeddings are obtained with CLIP **ViT-B/32** and **ViT-L/14** backbones as follows:

1. **One embedding per image.** We use CLIP’s *global pooled image embedding* returned by `encode_image` (i.e., the projected [CLS] / pooled token). This is the representation CLIP trains to align with text; it is therefore the correct summary statistic for distributional tests. We do *not* mean-pool patch tokens.
2. **Reference model.** Fit a *single* full-covariance Gaussian $\mathcal{N}(\mu, \Sigma)$ on the training embeddings. A full covariance is required because CLIP feature dimensions are correlated; diagonal covariances overstate distances along correlated axes. We do *not* use a Gaussian mixture.
3. **Mahalanobis diagnostic (used for histograms referenced in Fig. 6a).** For each embedding x in a split $S \in \{\text{Orig (test)}, \text{AttWarp}, \text{Non-RectilinearWarp}\}$, compute a single distance to the *training* Gaussian:

$$d_M(x) = \sqrt{(x - \mu)^\top \Sigma^{-1} (x - \mu)}.$$

The histogram is built from $\{d_M(x) : x \in S\}$. It is *not* a 12k-choose-2 pairwise computation. With population parameters, $d_M(x)^2$ follows χ_d^2 ; with estimated parameters it approximates Hotelling’s T^2 .

4. **FID/KID.** Using `cleanfid` Parmar et al. (2022) with default settings (64 batch, 8 workers, CLIP 224×224 preprocessing), compute Train→ S distances for each backbone. The Gaussian in step 2 is *fit once* on train; all eval splits are compared to that same reference to reveal shift (we never refit on test/warped sets).

Results. Fig. 14 shows Train→Test distances for both backbones. In **ViT-B/32** (Fig. 14a) and **ViT-L/14** (Fig. 14b), AttWarp (orange circle) stays within the natural Train→Test gap, reducing FID (13.8→12.4 for B/32; 56.6→49.8 for L/14) with only mild KID changes (8.3→11.0; 19.3→31.5). In contrast, the Non-Rectilinear warp (green triangle) produces a clear shift—FID rises to 15.5 and 73.9 and KID inflates markedly (37.6 and 174.9)—consistent with the Mahalanobis drift observed in Fig. 6a. These results quantitatively substantiate the design rationale above: rectilinear, axis-aligned warping preserves CLIP-feature distributional integrity, whereas free-form distortion does not.