The Narrow Gate: Localized Image-Text Communication in Native Multimodal Models

Alessandro Pietro Serra^{1,2*} Francesco Ortu^{1,3*} Emanuele Panizon^{1*}

Lucrezia Valeriani^{1,3} Lorenzo Basile^{1,3} Alessio Ansuini¹

Diego Doimo^{1†} Alberto Cazzaniga^{1†}

¹ Area Science Park, Trieste, Italy ² SISSA, Trieste, Italy ³ University of Trieste, Trieste, Italy {diego.doimo, alberto.cazzaniga}@areasciencepark.it

Abstract

Recent advances in multimodal training have significantly improved the integration of image understanding and generation within a unified model. This study investigates how vision-language models (VLMs) handle image-understanding tasks, focusing on how visual information is processed and transferred to the textual domain. We compare native multimodal VLMs, models trained from scratch on multimodal data to generate both text and images, and non-native multimodal VLMs, models adapted from pre-trained large language models or capable of generating only text, highlighting key differences in information flow. We find that in native multimodal VLMs, image and text embeddings are more separated within the residual stream. Moreover, VLMs differ in how visual information reaches text: non-native multimodal VLMs exhibit a distributed communication pattern, where information is exchanged through multiple image tokens, whereas models trained natively for joint image and text generation tend to rely on a single post-image token that acts as a *narrow gate* for visual information. We show that ablating this single token significantly deteriorates image-understanding performance, whereas targeted, token-level interventions reliably steer image semantics and downstream text with fine-grained control.

1 Introduction

The rise of foundation models [1] trained on vast amounts of text has transformed natural language processing (NLP), showing that a single large language model (LLM) [2] can handle many different linguistic tasks [3–5]. The rich set of features encoded in LLM embeddings has been then used as an effective prior knowledge both for text-conditional image generation [6–8] and image understanding [9–13]. Recently, the availability of large open datasets [14, 15] and improved techniques to align text and image embeddings [16] have also enabled the creation of multimodal models that can understand and generate visual content within a single architecture [17–20]. This unification allows a deeper understanding of the visual world, as generative tasks often require insight into the fundamental concepts and relationships within the data [21]. For example, a model that generates images from text descriptions must grasp the semantic content of those images to ensure that they faithfully reflect the details and intent of the text [22, 23]. As a result, research has rapidly progressed in

^{*}Equal contribution

[†]Equal supervision.

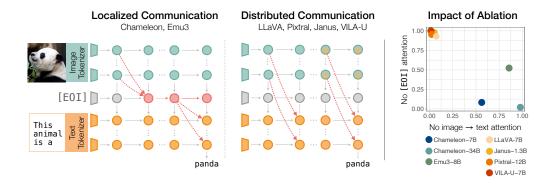


Figure 1: **Image-Text Communication in Vision Language Models.** The figure compares how different VLM architectures handle image-text information flow. **Left**: Native multimodal models (Chameleon, Emu3) process visual and textual tokens separately with information transfer occurring primarily through a single *narrow gate* - the end-of-image token ([E0I]). **Center**: In non-native multimodal models (LLaVA, Pixtral, Janus, VILA-U), communication is distributed across many internal image tokens; visual tokens in deeper layers align more strongly with text tokens. **Right**: The figure shows the relative performance on MS-COCO image captioning after performing different ablations, compared to no ablation (see section 3.4). The results show that for narrow gate models, removing attention to the [E0I] token (y-axis) impacts performance more significantly than removing attention to all image tokens (x-axis).

integrating multiple modalities into a unified framework with an increasingly deeper multimodal fusion. Early approaches used cross-attention modules between modality-specific encoders [17]. Further advancements highlighted the importance of using a pre-trained LLM backbone combined with lightweight projection layers, often fine-tuning the LLM to refine multimodal representations [19, 24, 25]. More recently, research has shown success in training vision language models (VLM) from scratch [20, 26], achieving performance similar to text-only output VLMs on visual understanding tasks. In this work, we use the term *native multimodal* to refer to VLMs that are trained from scratch on both text and image modalities, without relying on a pretrained LLM backbone and CLIP-based text-aligned vision encoder, and capable of generating both images and text.

Although recent studies have examined the internal mechanisms of unimodal output VLMs [27–30], it remains largely unknown how modalities interact in the hidden representations of native multimodal models. In this work, we focus on some of the latest native VLMs, Chameleon [20] and Emu3 [26], and compare them with a range of *non-native* ones, including both text-only VLMs such as LLaVA [11] and multimodal generators that rely on a pretrained language backbone like VILA-U [31], on how they transfer information from the visual domain to the textual domain in various image understanding tasks.

We observe that in native multimodal models (Chameleon and Emu3), image and text representations remain well-separated in different clusters from input to output, while for non-native models they tend to mix in late layers (section 3.1). Additionally, in this second class of models, visual information flows to text through multiple tokens. In contrast, native multimodal VLMs channel the global image information into a single token placed immediately after the image, the *end-of-image token*, denoted by [EOI](section 3.2). The [EOI] token behaves like a memory token [32] or *narrow gate* through which the image information must pass to guide text generation.

We establish the functional role of [EOI] by showing that (1) blocking attention from text tokens to [EOI] causes a sharp drop in performance across classification, VQA, and captioning tasks (section 3.4) and that (2) editing the representation of [EOI] alters the semantic content of the generated text (section 4.1), confirming its causal role in mediating image semantics. These properties do not appear in non-native multimodal VLMs, where the information between modalities flows through many distributed image tokens. Finally, we propose a fine-tuning strategy to remove the narrow gate in native multimodal models by redistributing visual-to-text transfer across many image tokens (section 4.2).

2 Background and Methods

Model architectures. We analyze six VLMs differing in their training objectives and multimodal integration. Two of them, Chameleon [20] and Emu3 [26], are native multimodal models trained from scratch to jointly generate images and text using a unified discrete tokenizer with a VQ-GAN image encoder. Since the released Emu3 variants were specialized either for image generation (Emu3-Gen) or for understanding (Emu3-Chat), we fine-tune the Emu3-Gen checkpoint on a balanced mixture of image—text datasets so that it can handle both tasks within a single model (see appendix B.3). This enables a direct comparison with Chameleon, which natively supports the joint generation of images and text. As non-native multimodal VLMs, we analyze two text-output VLMs: LLaVA-onevision-7B [33] (LLaVA) and Pixtral-12B [13] (Pixtral), fine-tuned on Qwen2-7B and Mistral Nemo 12B, respectively. Additionally, we also analyze Janus-1.3B [34] (Janus) and VILA-U [31], which are both multimodal in output and fine-tuned on a pre-trained LLM, DeepSeek-LLM-1.3B [35] and LLaMA-2-7B, respectively [36]. Janus employs a SigLIP vision encoder and separate pathways for image understanding and generation, while VILA-U combines CLIP-based visual features with a VQ-VAE discretization stage, yielding semantically rich visual embeddings. All the models are decoder-only and trained with next-token prediction loss. In Chameleon, Emu3, Janus, and VILA-U the loss is computed on both image and text tokens, while in LLaVA and Pixtral, it applies only to text tokens. In all the cases, a special end-of-image token ([EOI]) marks the transition from image to text. Additional architectural details are provided in table A2, and related multimodal training strategies are discussed in appendix A.

Datasets. We study visual understanding tasks where the model is asked to answer questions from VQAv2 [37], generate image captions for Flickr30k [38] and MS-COCO [39], and complete simple prompts about ImageNet images [40]. We provide the experimental setup, including dataset composition, prompt usage, and data selection criteria, in the appendix B.

2.1 Analytical Tools

Analyzing information flow in VLMs. To understand how VLMs process information, we analyze both attention maps and token representations across layers. We denote the *residual stream* representation of the token at position i at layer l as x_i^l . In our setup, images are typically followed by textual instructions. To study how these modalities interact, we focus on the cross-modal attention: let $N_{\texttt{[EOI]}}$ be the position of the end-of-image token, and $A_{i,j}$ the attention from token i to token j. We define $A_{\texttt{text} \to \texttt{img}} = A_{i,j}$ with $N_{\texttt{[EOI]}} < i \le N$, $0 < j < N_{\texttt{[EOI]}}$, which captures how much text tokens attend to visual tokens in the sequence.

Quantifying cross-modal attention. We construct a metric that quantifies the average attention that all the text tokens give to a token at position j within the image part of the prompt, which we assume to span the first $N_{\texttt{[EOI]}}$ tokens. Formally, we define the (relative) cross-modal attention f_j^l as

$$f_j^l = \frac{1}{C} \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} \sum_{i > N_{\text{[EGI]}}} A_{i,j}^{l,h} \tag{1}$$

Where $j=0,1,\ldots,N_{\texttt{[EDI]}}$, l,h identify, respectively, the layer and head. C is a normalization factor such that $\sum_{j\leq N_{\texttt{[EDI]}}} f_j^l = 1$.

Blocking Cross-Modal communication with Attention Knockout. We use the *attention knockout* [41] to selectively block attention between tokens. Given a set of source tokens \mathcal{S} and a set of target tokens \mathcal{T} , we zero out the attention weights $A_{i,j}$, where $i \in \mathcal{S}$ and $j \in \mathcal{T}$ across *all heads* in selected layers to prevent the target tokens from attending to the source tokens.

Probing the semantic information with Neighborhood Overlap. To evaluate how well the residual stream encodes an abstract property of the image, like the identity of the main object represented (ImageNet) or the important parts of the scene (MS-COCO), we use the *neighborhood overlap* [42]. The neighborhood overlap measures the consistency between the k-nearest neighbors of a datapoint i in the residual stream at a layer l, denoted by $\mathcal{N}_k^l(i)$, and the nearest neighbors in a reference data representation, $\mathcal{N}_k^{gt}(i)$, which encodes the abstract property of interest. In the case of ImageNet, all the points belonging to the same class are considered nearest neighbors [42] for the

purpose of computing the reference $\mathcal{N}_k^{gt}(i)$. The *neighborhood overlap*, denoted by $\chi_k^{l,gt}$, can be then written as:

$$\chi_k^{l,gt} = \frac{1}{nk} \sum_{i=1}^n \left| \mathcal{N}_k^l(i) \cap \mathcal{N}_k^{gt}(i) \right| \tag{2}$$

where $|\cdot|$ denotes the cardinality of a set, and n the dataset size. In our experiments, we set k=30.

Modifying the semantic information with localized interventions. To evaluate the influence of specific components on model outputs, we use activation patching [43, 44]. This approach involves two forward passes on two different inputs: We first collect the activations \hat{x}_i^l from a target input, then replace the corresponding activations x_i^l in a base input during a second forward pass. The impact is assessed by comparing the resulting output distribution $q_{\text{base}}^{\text{patched}}$ to the target output distribution p_{target} . We use a variant of the Jaccard index [45] to quantify the similarity between $q_{\text{base}}^{\text{patched}}$ and p_{target} :

Similarity
$$(q_{\text{base}}^{\text{patched}}, p_{\text{target}}) = \sum_{i} \min(q_i, p_i)$$
 (3)

The score ranges from 0 (no overlap) to 1 (identical distributions), measuring how much the patched model mimics the target behavior.

Reproducibility. We used the HuggingFace implementations Chameleon [46, 47], Emu3-Gen [48], LLaVA [49], Pixtral [50], VILA-U [31] and Janus [51] models. We run all the experiments on a single NVIDIA A100 GPUs with 40GB VRAM. For the Chameleon 34B, we used two 40GB GPUs. The code needed to reproduce the experiments is available at: ritareasciencepark.github.io/Narrow-gate.

3 Results

3.1 Modality Gap in Native Multimodal VLMs

Native multimodal VLMs generate both image and text tokens using a shared transformer backbone, enabling unified next-token prediction across modalities within a common representation space. In contrast, text-output VLMs produce only textual tokens, using image tokens solely as contextual input. Although both VLM types encode image and text tokens in the residual stream, their role in the generation processes differ significantly. This motivates our first analysis of whether native multimodal models develop modality-specific subspaces in the residual stream that separate image and text representations.

Cross-modal angular separation in VLMs. To address this question, we randomly select 10,000 image-caption pairs from the Flickr30k dataset and extract the residual stream representations of all hidden layers at a randomly chosen *text token position* (see appendix B for the complete prompt structure). To compare the geometric organization of the different modalities, we first measure the median cosine similarity between representations of image and text at each hidden layer.

As shown in figure 2-left, this analysis reveals different trends across architectures. In the Chameleon-7B (blue) and Chameleon-34B (green), the representative vectors of image and text tokens remain nearly orthogonal throughout the hidden layers, with median cosine similarity values consistently below 0.10. The cosine similarity in LLaVA (yellow) and Emu3 (green) is higher, starting at 0.2 in the initial layers of the network. But while in Emu3, the cosine similarity decreases to 0 in the last layers of the network, in LLaVA it rises to 0.5. This implies that in Emu3, the two modalities remain separated in late layers, and in LLava, they increasingly mix. Similarly, other non-native multimodal models (Janus, Pixtral, VILA-U) show the same qualitative behavior of LLaVA and are reported to appendix C.1.

Clustering of visual and textual embeddings. To characterize the degree of mixture of the modalities from a different perspective, taking into account the norm of the embeddings, we cluster the embeddings with the Advanced Density Peaks algorithm [52] (see appendix B.4 for a brief introduction) and analyze the composition of each cluster. Figure 2-right shows the homogeneity score [53] of the clusters, measuring how the clustering of tokens aligns with their modality. In Chameleon and Emu3 models, the homogeneity score is always 1, meaning that each cluster contains

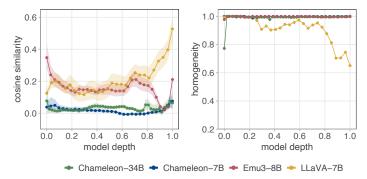


Figure 2: **Modality Gap in VLMs.** (**left**) Cosine similarity between text and image token embeddings shows that LLaVA achieves increasing alignment with depth, while Emu3 shows little alignment and Chameleon maintains orthogonality. Points represent median cosine similarity, with shaded areas indicating the inter-quartile range. (**right**) Homogeneity score assesses how well token clusters (via Advanced Density Peaks) correspond to their original modality.

embeddings from a single modality. In contrast, LLaVA's homogeneity score decreases from almost 1 to approximately 0.6 in the later layers of the network, indicating a tendency for the text and image embeddings to mix. In the appendix C.1, we show that, similarly to LLaVA, also in Janus, Pixtral, and VILA-U, the clusters of late layers contain embeddings from both modalities.

We conclude that textual and visual representation spaces are largely separated in the Chameleon and Emu3 models, which are the only ones trained from scratch on multimodal data and can generate both images and textual samples. This brings forth the problem of finding how and where communication between the two modalities is performed.

3.2 Analysis of Cross-Modal Attention

As shown in section 3.1, hidden representations of visual and textual tokens mix in the late layers of non-native multimodal models, whereas they occupy well-separated regions of the residual stream in native multimodal models. Despite this structural difference, all model families achieve strong performance on image-understanding tasks, suggesting that cross-modal attention effectively transfers semantic information from visual to textual tokens. In this process, attention matrices play a central role: while enabling information exchange across modalities, they must bridge the modality gap, which is particularly pronounced in native multimodal models such as Chameleon and Emu3. These observations motivate the definition of two key properties for tokens that mediate cross-modal communication: (i) having a high weight in text-to-image attention and (ii) encoding rich semantic information about the visual input. In the following, we examine whether Chameleon-7B, Emu3, and LLaVA (the latter representing non-native multimodal models) contain tokens that satisfy both criteria, identifying candidates that act as communication gates between modalities. Corresponding analyses for other models are provided in figure A3.

Cross-modal attention patterns. To quantify the role of each token in image-to-text semantic communication, we use the ImageNet dataset and select 100 images per class for 100 animal classes [40]. We construct 10,000 prompts of images followed by text of the form " $\langle image \rangle$ This animal is a __". The sentence was chosen to guide the model towards generating a class-relevant response.

Figure 3 illustrates the relative text-on-image attention, as defined in section 2.1, at different token positions for Chameleon-7B (left), Emu3 (center), and LLaVA (right). We single out tokens with an average value larger than 1%. The remaining tokens are aggregated as *internal image*. Both in Chameleon and Emu, the special [E0I] token captures a large amount of cross-modal attention (shown in green in the figure). In Chameleon, this token alone receives 40% to 50% of the total attention from the textual tokens between layer 2 and 6, remaining above 15 to 20% in the second half of the network. Similarly, in Emu3 [E0I] receives more than 30 to 40% of text-on-image attention after layer 10. Other highly attended tokens are, in Chameleon, the 32nd image token from layer 5 to layer 10 (orange) and the last token of the image receiving between 66% and 86% in the middle and final layers (shown in pink). In Emu instead, a small group of tokens describing the size of the image

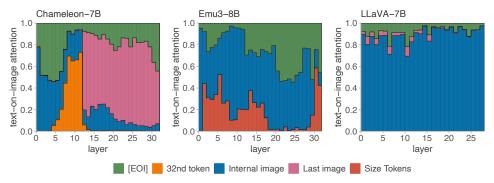


Figure 3: Cross-Modal Attention Contributions of Image Tokens. Contribution of different image token positions to the total text-on-image attention across layers in Chameleon (left), Emu (center), and LLaVA (right), computed on ImageNet data. Tokens with an average contribution larger than 1% are singled out. The remaining tokens are aggregated as "internal image".

receive from 10 to 40% of the total cross-modal attention before layer 17 and after layer 30 (shown in dark orange). The remaining attention is distributed among the other 1024 image tokens. In the right panel, we present the distribution of average attention among tokens in LLaVA. The [E0I] token accounts for approximately 10-20% of the attention between layers 0 and 13, and less than 10% for the remaining layers. All the remaining attention is distributed across the other image tokens. These results suggest the emergence of two distinct communication regimes. Native multimodal models tend to concentrate text-on-image attention on a few single tokens (such [E0I]), which may act as a privileged channel for visual information. In contrast, non-native multimodal models allocate little attention to [E0I], with attention instead dominated by internal image tokens, suggesting a more diffuse and distributed communication between modalities.

3.3 Probing the Semantic Content of Visual Tokens

To determine whether the tokens that receive the highest text-on-image attention also encode abstract information about the visual input – the second condition for serving as communication channels – we probe the semantic content of the hidden representations. For this purpose, we use the neighborhood overlap ($\chi_k^{l,gt}$, see section 2.1) with two different reference ground truths: the labels of the ImageNet dataset and the caption embeddings of a Qwen2-7B-GTE text encoder for the MS-COCO.

Alignment to ImageNet class labels. First, we measure the neighborhood overlap $\chi_k^{l,gt}$ between image-token embeddings and ImageNet class labels (section 2.1), averaging across positions when multiple tokens are analyzed. As shown in figure 4, the <code>[EOI]</code> token in Chameleon-7B and Emu3 exhibits a sharp rise in $\chi_k^{l,gt}$ after the first few layers, exceeding 0.4 and 0.25 respectively, and retaining high values through the model's depth. In contrast, the 32nd, final, and size tokens maintain values near zero, indicating that they attract attention but do not encode meaningful semantic content. The internal image tokens (blue curves) initially capture visual semantics in early layers but gradually lose this information beyond layer 10, leaving <code>[EOI]</code> as the dominant semantic carrier. This pattern holds across scales: in Chameleon-34B, <code>[EOI]</code> exceeds $\chi_k^{l,gt} > 0.4$ after layer 16, while other tokens remain below 0.25 (figure A3). For LLaVA, $\chi_k^{l,gt}$ for <code>[EOI]</code> starts around 0.2 and decreases below 0.1 in later layers, whereas internal image tokens maintain high and stable overlap values (> 0.4). A similar trend is observed in other non-native multimodal, as is shown in figure A3 models such as Janus, Pixtral, and VILA-U: internal image tokens retain the highest $\chi_k^{l,gt}$, while <code>[EOI]</code> contributes little semantic signal. These results suggest two distinct encoding regimes. In native multimodal models (Chameleon, Emu3), the <code>[EOI]</code> token both attracts the majority of text-on-image attention and carries the richest visual semantics, fulfilling the conditions for a narrow communication gate. In non-native multimodal models, semantic content remains broadly distributed across internal image tokens, indicating a diffuse, multi-token communication pattern.

Alignment to the embedding image captions. To verify that this behavior is not specific to classification, we repeated the analysis on MS-COCO captioning data. We compute the neighborhood

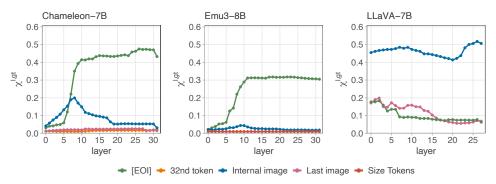


Figure 4: **Localization of Visual Semantic Information.** The figure compares how visual semantic information is localized in Chameleon, Emu3, and LLaVA, by measuring the neighborhood overlap between image tokens and ImageNet labels. The blue curves show the average overlap across all image tokens, excluding the 32nd token for the case of Chameleon.

overlap between the visual-token representations and their corresponding captions. To define a metric in the "space of captions", we use their semantic embeddings through a Qwen2-7B-GTE text encoder (see appendix C.2). The results (figures A2 and A3) mirror those on ImageNet: in native multimodal models, [E0I] maintains the strongest alignment with caption embeddings, whereas in non-native multimodal models, semantic information remains distributed across internal image tokens.

3.4 Ablation Experiments: The Effect of Localized Communication on Downstream Tasks

Building on the previous analyses, where we identified the tokens most involved in visual–textual communication, we now test their functional importance through attention ablations. Specifically, we evaluate how blocking communication at selected token positions affects performance on three standard vision–language benchmarks: VQAv2 [37], Flickr30k [38], and MS-COCO [39]. Details on evaluation metrics are provided in appendix B.

Ablation setup. For each model, we randomly sample 2,000 examples per dataset and measure baseline performance before applying targeted attention knockouts. The performance for visual question answering VQAv2 are evaluated with the VQA metric [54], for the captioning benchamarks instead we used CIDEr [55]. We block (i) communication between the [E0I] and text tokens by zeroing out $A_{\text{text}\to [\text{E0I}]}$, and (ii) all text-on-image attention by zeroing out $A_{\text{text}\to \text{img}}$ (section 2.1). This isolates the impact of the [E0I] token relative to the broader visual stream. Results for Chameleon, Emu3, and LLaVA are summarized in table 1; additional models (Pixtral, Janus, VILA-U) and further token positions are reported in appendix C.4.

Table 1: **Effect of Attention Knockout on Image Understanding Tasks**. Performance of the models on VQAv2, Flickr-30k, MS-COCO, and ImageNet under different ablation settings. Numbers in bold mark the worst performance for each model and task.

Model	Ablation	VQAv2	MS-COCO	Flickr	ImageNet
Chameleon-7B	$\begin{array}{l} \text{-} \\ \text{text} \rightarrow \texttt{[EOI]} \\ \text{text} \rightarrow \text{img} \end{array}$	0.51 0.25 0.40	0.48 0.04 0.27	0.34 0.04 0.17	0.46 0.01 0.47
Emu3	$\begin{array}{c} \text{-} \\ \text{text} \rightarrow \texttt{[EOI]} \\ \text{text} \rightarrow \text{img} \end{array}$	0.57 0.48 0.42	0.63 0.33 0.54	0.29 0.13 0.21	0.35 0.24 0.30
LLaVA	$\begin{array}{c} \text{-} \\ \text{text} \rightarrow \texttt{[EOI]} \\ \text{text} \rightarrow \text{img} \end{array}$	0.80 0.80 0.00	0.98 0.97 0.01	0.70 0.71 0.02	0.5 0.45 0.05

Results. In native multimodal models, ablating the [E0I] token causes a substantial performance collapse across all tasks. For instance, in Chameleon-7B, accuracy drops from 0.51 to 0.25 on

VQAv2 and from 0.34 to 0.04 on Flickr30k, with similar reductions exceeding 90% on MS-COCO and near-zero overlap on ImageNet. Emu3-8B shows comparable degradation, with performance falling by roughly 50% across tasks. In both models, blocking all text-on-image attention ($A_{\text{text}\to \text{img}}$) produces smaller declines than ablating the single [E01] token, confirming its dominant role. Similar trends hold for Chameleon-34B (appendix C.4). In contrast, non-native multimodal models rely on distributed communication. In LLaVA, performance is unaffected by the ablation of [E01] but collapses entirely when all text-on-image attention is removed (0.00 on VQAv2, 0.02 on Flickr30k, 0.01 on MS-COCO, 0.05 on ImageNet). Pixtral, Janus, and VILA-U exhibit the same pattern, indicating that their cross-modal information flow is mediated by many internal image tokens rather than a single position. These ablation results reinforce the distinction between communication regimes. Native multimodal models (Chameleon, Emu3) depend on a localized, token-level bottleneck—the [E01] token—for transferring visual semantics to text. By contrast, non-native multimodal models (LLaVA, Pixtral, Janus, VILA-U) exhibit distributed communication: information flows through numerous image tokens and cannot be disrupted by a localized ablation.

4 Implication of Narrow Gate for Model Editing and Robustness

4.1 Steering Image Semantics Through Activation Patching

The previous analyses showed that in native multimodal models (Chameleon, Emu3), the [E0I] token encodes global visual semantics (section 3.2) and acts as the main communication channel to text (section 3.4). We now test whether this localized structure also enables *constructive* interventions, specifically, whether modifying the representation at [E0I] can steer the model's interpretation of an image toward a desired semantic class.

Activation patching setup. To do so, we select 20 ImageNet animal classes and sample 100 images per class, forming 10 class pairs. For each pair, we perform activation patching (section 2.1) with the following procedure: (i) extract the [EOI] representation at each layer l from the target class images ([EOI] l,target); (ii) inject this representation into the residual stream of the base class images by replacing [EOI] $^{l^*,base}$ with [EOI] $^{l^*,target}$ at a chosen layer l^* ; and (iii) evaluate the model's output distribution for the final text token in the prompt " $\langle image \rangle$ This animal is a __". We measure the similarity between output probabilities (see equation (3)) and the fraction of cases where the target class probability exceeds that of the base class. For further details on experimental setup see appendix B.

Results. Figure 5-left reports the similarity measure as a function of the patched layer for Chameleon-7B, Chameleon-34B, Emu3, and LLaVA. In native multimodal models, patching [E0I]

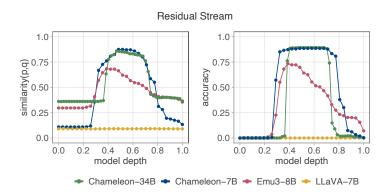


Figure 5: **Impact of Activation Patching at** [EOI]. Results of patching the [EOI] representation from a target class onto the [EOI] representation on a base class. The experiment is performed at each layer. Two metrics are evaluated: (**left**) The similarity measure (defined in equation (3)) quantifies how the probability distribution over the vocabulary of the patched image aligns with that of the target. (**right**) The accuracy, defined as the fraction of patched images where the probability of the target class token is larger than the probability of the base class token.

representations from the target class induces strong semantic steering: the similarity rises sharply after a relative depth of ~ 0.3 , peaking at 0.86 in Chameleon-7B/34B and 0.70 in Emu3. The right panel quantifies the success rate of semantic transfer: patching changes the predicted class from base to target in approximately 90% of cases for Chameleon and 75% for Emu3. In contrast, LLaVA, representative of non-native multimodal models, shows no measurable effect. [E0I] patching leaves both similarity and class probabilities unchanged: the base class consistently dominates the output, indicating that cross-modal information is distributed across multiple image tokens rather than localized in a single position. These results demonstrate that in native multimodal models, the [E0I] token provides a causal handle for controlling image semantics: altering its representation is sufficient to steer model predictions. In non-native multimodal models, by contrast, semantic information is distributed throughout the image representation, rendering localized interventions ineffective.

4.2 Improving Model Robustness with Localized Fine-tuning

The presence of a narrow gate in native multimodal models can constrain how the model integrate visual and linguistic features, making its representations overly dependent on a specific internal pathway. To promote a richer and more distributed exchange between modalities, in this section, we show a fine-tuning strategy that masks the [EOI] token during training, forcing the model to encode and communicate visual information stored in [EOI] through alternative internal tokens. Details about the training setup are reported in appendix B.3.

Masked fine-tuning. Figure 6 shows how model performance evolves during training on 2000 test samples from the MS-COCO and VQAv2 tasks. Solid lines report fine-tuning runs with [E0I] masked, while dashed lines correspond to standard fine-tuning where all input tokens are attended. Purple profiles show the performance when [E0I] is ablated, while green profiles show the test performance when all tokens are visible. At the beginning of training, ablating [E0I] causes a large performance drop, especially for the COCO captioning task (see section 3.4). However, when [E0I] is masked during fine-tuning, performance on both MS-COCO and VQAv2 (solid-purple profiles) increases and approaches the "unablated" performance (green profiles), which remains approximately constant during training. In contrast, when [E0I] is not masked during training (dashed purple profiles), the performance gap and the narrow gate phenomenon remains, since the model continues to rely mostly on [E0I] to store and communicate visual information.

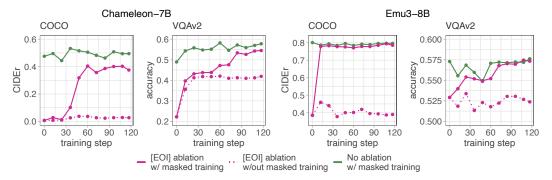


Figure 6: **Performance Dynamics During Masked** [E01] **Fine-tuning.** Masking the [E01] token during fine-tuning reduces the model's reliance on a single narrow gate and encourages a more distributed communication between image and text. Solid lines indicate runs where [E01] is masked; dashed lines correspond to standard fine-tuning. Green curves report test accuracy with all tokens visible, while purple curves show performance when [E01] is ablated. When trained with masked [E01], the performance under ablation (solid purple) progressively recovers and approaches the unablated baseline (green), indicating improved robustness and reduced dependence on [E01].

5 Discussion and Conclusion

The Role of [E01] **in Cross-Modal Communication.** In this work, we analyzed information flow within transformer-based VLMs using geometric and mechanistic interpretability techniques. Our study focused on native multimodal models - those trained from scratch to generate both images and

text - filling a gap in the existing literature, mainly focused on text-output or fine-tuned multimodal systems [56].

Our analyses reveal a *narrow gate* communication mechanism in native multimodal VLMs. These models exhibit a pronounced modality gap: visual and textual representations remain well separated throughout the network (section 3.1), yet semantic transfer occurs effectively through a *single token*, the end-of-image ([EOI]). The [EOI] token accumulates visual semantics in early and mid layers and subsequently channels this information to textual tokens (section 3.2). Ablation (section 3.4) and activation-patching (section 4.1) experiments demonstrate that disrupting or modifying this token has a direct causal impact on model predictions, confirming that [EOI] serves as the central conduit for image-to-text communication.

In contrast, non-native multimodal models (e.g., LLaVA, Janus, Pixtral, VILA-U) display a more *distributed communication* pattern, in which visual semantics are shared across many image tokens. These results suggest that the emergence of localized, token-level communication is a specific property of native multimodal architectures and training regimes.

The Role of Architecture in the Emergence of the Narrow Gate The presence of a narrow gate is linked to the modality gap between image and text representations. Based on our results we hypothesized that three architectural and training factors jointly influence its emergence. The multimodal output objective: a training objective that requires generating both image and text tokens, encourages the model to maintain distinct representational pathways for each modality, enforcing their separation. The training history: models trained from scratch on multimodal data, develop different subspaces for each modality subspaces, while models derived from pre-trained text-only LLMs and later fine-tuned for multimodal use (Janus, VILA-U) retain a text-aligned internal geometry and exhibit distributed communication. The type of image encoder: low-level visual tokenizers such as VQ-GAN (used in Chameleon), represent images through local features rather than high-level semantic embeddings. This increases the abstraction gap between image and text tokens and promotes the formation of a distinct communication interface centered on [EOI]. Taken together, these factors show that the narrow gate is not an incidental artifact but a structural consequence of jointly training a multimodal generator with low-level image tokenization.

Implications and Future Directions The localized communication via [E01] introduces both limitations and opportunities. On one hand, the reduced communication bandwidth may constrain multimodal reasoning and expose the model to potential vulnerabilities. For example small perturbations or targeted manipulations of a single token can strongly influence output semantics. Encouraging a more distributed information flow, for instance by masking text-on-[E01] attention during late-stage fine-tuning, could mitigate these issues. On the other hand, the narrow gate mechanism provides clear interpretability and practical benefits. It enables controlled steering of visual semantics, simplifies tracing how visual information shapes text generation, and facilitates efficient fine-tuning or alignment by localizing cross-modal transfer to a single token. Moreover, the [E01] token acts as a natural *memory token* [32], summarizing global image information and allowing early image tokens to be discarded thereby reducing memory and computation costs [57, 58].

Overall, our study identifies the architectural and representational conditions under which narrow gate communication emerges in multimodal transformers. By bridging geometric analysis, causal interventions, and training design, our findings highlight how architectural choices determine whether visual—textual communication is concentrated through a single token or distributed across many, clarifying in this way an essential dimension of modern vision—language modeling.

Limitations

Our analysis is limited to the image-to-text direction and to a subset of possible VLM architectures. We do not evaluate the communication in the text-to-image direction, or native VLMs that rely on higher-level, continuous encoding, and models with diffusion-based decoders. Additionally, our study focuses only on native vision-language models. Extending these analyses to native multimodal models trained across a broader range of modalities remains an important open direction for future research.

Acknowledgments

The authors acknowledge the AREA Science Park supercomputing platform ORFEO made available for conducting the research reported in this paper, and the technical support of the Laboratory of Data Engineering staff. Alessandro Pietro Serra, Francesco Ortu, Emanuele Panizon, Lorenzo Basile, Alessio Ansuini, Diego Doimo and Alberto Cazzaniga were supported by the project "Supporto alla diagnosi di malattie rare tramite l'intelligenza artificiale" CUP: F53C22001770002 and "Valutazione automatica delle immagini diagnostiche tramite l'intelligenza artificiale", CUP: F53C22001780002. Alessio Ansuini and Alberto Cazzaniga were supported by the European Union – NextGenerationEU within the project PNRR "PRP@CERIC" IR0000028 - Mission 4 Component 2 Investment 3.1 Action 3.1.1. Lucrezia Valeriani was supported by the project "QuB - Quantum Behavior in Biological Function" CUP: J95F21002820001. Lorenzo Basile was supported by the European Union – NextGenerationEU within the project PNRR "Finanziamento di progetti presentati da giovani ricercatori" - Mission 4 Component 2 Investment 1.2, CUP: J93C25000440001.

References

- [1] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen A. Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir P. Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R'e, Dorsa Sadigh, Shiori Sagawa, Keshay Santhanam, Andy Shih, Krishna Parasuram Sriniyasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. ArXiv, 2021.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [3] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. Semantic parsing on Freebase from question-answer pairs. In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard, editors, *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA, October 2013. Association for Computational Linguistics.
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate, 2016.
- [5] Mehdi Allahyari, Seyedamin Pouriyeh, Mehdi Assefi, Saeid Safaei, Elizabeth D. Trippe, Juan B. Gutierrez, and Krys Kochut. Text summarization techniques: A brief survey, 2017.
- [6] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation, 2021.
- [7] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 36479–36494. Curran Associates, Inc., 2022.

- [8] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.
- [9] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 15638–15650, June 2022.
- [10] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikoł aj Bińkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, Advances in Neural Information Processing Systems, volume 35, pages 23716–23736. Curran Associates, Inc., 2022.
- [11] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916. Curran Associates, Inc., 2023.
- [12] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli VanderBilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Sarnat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Jen Dumas, Crystal Nam, Sophie Lebrecht, Caitlin Wittlif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, and Aniruddha Kembhavi. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models, 2024.
- [13] Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, Soham Ghosh, Amélie Héliou, Paul Jacob, Albert Q. Jiang, Kartik Khandelwal, Timothée Lacroix, Guillaume Lample, Diego Las Casas, Thibaut Lavril, Teven Le Scao, Andy Lo, William Marshall, Louis Martin, Arthur Mensch, Pavankumar Muddireddy, Valera Nemychnikova, Marie Pellat, Patrick Von Platen, Nikhil Raghuraman, Baptiste Rozière, Alexandre Sablayrolles, Lucile Saulnier, Romain Sauvestre, Wendy Shang, Roman Soletskyi, Lawrence Stewart, Pierre Stock, Joachim Studnia, Sandeep Subramanian, Sagar Vaze, Thomas Wang, and Sophia Yang. Pixtral 12b, 2024.
- [14] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022.
- [15] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. Obelics: An open web-scale filtered dataset of interleaved image-text documents, 2023.
- [16] Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, and Ying Shan. Making LLaMA SEE and draw with SEED tokenizer. In The Twelfth International Conference on Learning Representations, 2024.
- [17] Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. Any-to-any generation via composable diffusion. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, Advances in Neural Information Processing Systems, volume 36, pages 16083–16099. Curran Associates, Inc., 2023.
- [18] Gemini Team. Gemini: A family of highly capable multimodal models, 2024.
- [19] Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, Hang Yan, Jie Fu, Tao Gui, Tianxiang Sun, Yu-Gang Jiang, and Xipeng Qiu. AnyGPT: Unified multimodal LLM with discrete sequence modeling. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 9637–9662, Bangkok, Thailand, August 2024. Association for Computational Linguistics.
- [20] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. arXiv preprint arXiv:2405.09818, 2024.

- [21] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, and Li Yi. DreamLLM: Synergistic multimodal comprehension and creation. In *The Twelfth International Conference on Learning Representations*, 2024.
- [22] Hong Chen, Xin Wang, Yuwei Zhou, Bin Huang, Yipeng Zhang, Wei Feng, Houlun Chen, Zeyang Zhang, Siao Tang, and Wenwu Zhu. Multi-modal generative ai: Multi-modal llm, diffusion and beyond, 2024.
- [23] Shenghao Xie, Wenqiang Zu, Mingyang Zhao, Duo Su, Shilong Liu, Ruohua Shi, Guoqi Li, Shanghang Zhang, and Lei Ma. Towards unifying understanding and generation in the era of vision foundation models: A survey from the autoregression perspective, 2024.
- [24] Quan Sun, Qiying Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Emu: Generative pretraining in multimodality. In The Twelfth International Conference on Learning Representations, 2024.
- [25] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. NExt-GPT: Any-to-any multimodal LLM. In Forty-first International Conference on Machine Learning, 2024.
- [26] Xinlong Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Zhen Li, Qiying Yu, Yingli Zhao, Yulong Ao, Xuebin Min, Tao Li, Boya Wu, Bo Zhao, Bowen Zhang, Liangdong Wang, Guang Liu, Zheqi He, Xi Yang, Jingjing Liu, Yonghua Lin, Tiejun Huang, and Zhongyuan Wang. Emu3: Next-token prediction is all you need, 2024.
- [27] Clement Neo, Luke Ong, Philip Torr, Mor Geva, David Krueger, and Fazl Barez. Towards interpreting visual information processing in vision-language models, 2024.
- [28] Samyadeep Basu, Martin Grayson, Cecily Morrison, Besmira Nushi, Soheil Feizi, and Daniela Massiceti. Understanding information storage and transfer in multi-modal large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [29] Yuxin Wen, Qingqing Cao, Qichen Fu, Sachin Mehta, and Mahyar Najibi. Efficient vision-language models by summarizing visual tokens into compact registers, 2024.
- [30] Matthew Trager, Pramuditha Perera, Luca Zancato, Alessandro Achille, Parminder Bhatia, and Stefano Soatto. Linear spaces of meanings: Compositional structures in vision-language models. In *ICCV*, pages 15349–15358, 2023.
- [31] Yecheng Wu, Zhuoyang Zhang, Junyu Chen, Haotian Tang, Dacheng Li, Yunhao Fang, Ligeng Zhu, Enze Xie, Hongxu Yin, Li Yi, Song Han, and Yao Lu. VILA-U: a unified foundation model integrating visual understanding and generation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [32] Mikhail S. Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V. Sapunov. Memory transformer, 2021.
- [33] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer, 2024.
- [34] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. *arXiv* preprint arXiv:2410.13848, 2024.
- [35] DeepSeek-AI. Deepseek llm: Scaling open-source language models with longtermism. arXiv preprint arXiv:2401.02954, 2024.
- [36] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023.
- [37] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [38] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.

- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, Computer Vision ECCV 2014, pages 740–755, Cham, 2014. Springer International Publishing.
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- [41] Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual associations in auto-regressive language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12216–12235. Association for Computational Linguistics, 2023.
- [42] Diego Doimo, Aldo Glielmo, Alessio Ansuini, and Alessandro Laio. Hierarchical nucleation in deep neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- [43] Atticus Geiger, Hanson Lu, Thomas Icard, and Christopher Potts. Causal abstractions of neural networks. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pages 9574–9586, 2021.
- [44] Arthur Conmy, Augustine N. Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023.
- [45] Paul Jaccard. Distribution de la flore alpine dans le bassin des dranses et dans quelques régions voisines. Bull Soc Vaudoise Sci Nat, 37:241–272, 1901.
- [46] Facebook AI. Chameleon-7b. https://huggingface.co/facebook/chameleon-7b, 2024.
- [47] Facebook AI. Chameleon-30b. https://huggingface.co/facebook/chameleon-30b, 2024.
- [48] BAAI. Emu3-gen. https://huggingface.co/BAAI/Emu3-Gen-hf, 2025.
- [49] Llava-hf. Llava-onevision-7b. https://huggingface.co/llava-hf/llava-onevision-qwen2-7b-ov-hf, 2024.
- [50] Mistral Community. Pixtral-12b. https://huggingface.co/mistral-community/pixtral-12b, 2024
- [51] DeepSeek AI. Janus-1.3b. https://huggingface.co/deepseek-ai/Janus-1.3B, 2024.
- [52] Maria d'Errico, Elena Facco, Alessandro Laio, and Alex Rodriguez. Automatic topography of highdimensional data sets by non-parametric density peak clustering. *Information Sciences*, 560:476–492, 2021.
- [53] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In Jason Eisner, editor, Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL), pages 410–420, Prague, Czech Republic, June 2007. Association for Computational Linguistics.
- [54] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [55] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [56] Zihao Lin, Samyadeep Basu, Mohammad Beigi, Varun Manjunatha, Ryan A Rossi, Zichao Wang, Yufan Zhou, Sriram Balasubramanian, Arman Zarei, Keivan Rezaei, et al. A survey on mechanistic interpretability for multi-modal foundation models. arXiv preprint arXiv:2502.17516, 2025.
- [57] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In The Twelfth International Conference on Learning Representations, 2024.

- [58] Yuxin Wen, Qingqing Cao, Qichen Fu, Sachin Mehta, and Mahyar Najibi. Efficient vision-language models by summarizing visual tokens into compact registers. *arXiv preprint arXiv:2410.14072*, 2024.
- [59] Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models, 2024.
- [60] David M. Chan, Rodolfo Corona, Joonyong Park, Cheol Jun Cho, Yutong Bai, and Trevor Darrell. Analyzing the language of visual tokens, 2024.
- [61] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pretraining with frozen image encoders and large language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the* 40th International Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pages 19730–19742. PMLR, 23–29 Jul 2023.
- [62] Peter Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Adithya Jairam Vedagiri IYER, Sai Charitha Akula, Shusheng Yang, Jihan Yang, Manoj Middepogu, Ziteng Wang, Xichen Pan, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, Advances in Neural Information Processing Systems, volume 37, pages 87310–87356. Curran Associates, Inc., 2024.
- [63] Yang Jin, Kun Xu, Kun Xu, Liwei Chen, Chao Liao, Jianchao Tan, Quzhe Huang, Bin CHEN, Chengru Song, dai meng, Di ZHANG, Wenwu Ou, Kun Gai, and Yadong MU. Unified language-vision pretraining in LLM with dynamic discrete visual tokenization. In *The Twelfth International Conference on Learning Representations*, 2024.
- [64] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 14398–14409, June 2024.
- [65] Yuying Ge, Yixiao Ge, Ziyun Zeng, Xintao Wang, and Ying Shan. Planting a seed of vision in large language model, 2023.
- [66] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [67] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, Computer Vision – ECCV 2022, pages 89–106, Cham, 2022. Springer Nature Switzerland.
- [68] Shengqiong Wu, Hao Fei, Xiangtai Li, Jiayi Ji, Hanwang Zhang, Tat-Seng Chua, and Shuicheng YAN. Towards semantic equivalence of tokenization in multimodal LLM. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [69] Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, Advances in Neural Information Processing Systems, 2022.
- [70] Lucrezia Valeriani, Diego Doimo, Francesca Cuturello, Alessandro Laio, Alessio Ansuini, and Alberto Cazzaniga. The geometry of hidden representations of large transformer models. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023.
- [71] Emily Cheng, Diego Doimo, Corentin Kervadec, Iuri Macocco, Jade Yu, Alessandro Laio, and Marco Baroni. Emergence of a high-dimensional abstraction phase in language transformers. *arXiv preprint arXiv:2405.15471*, 2024.
- [72] Diego Doimo, Alessandro Serra, Alessio Ansuini, and Alberto Cazzaniga. The representation landscape of few-shot learning and fine-tuning in large language models, 2024.
- [73] Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning, 2023.

A Related Works

Special Tokens, Memory Tokens, Registers. The importance of special tokens to store and redistribute global information was emphasized by Burtsev et al. [32]. In their work, they add these special tokens at the beginning of the sequence to serve as memory units accessible through the self-attention mechanism, improving performance on various language processing tasks. In vision transformers, Darcet et al. [57] found that high-norm tokens in low-informative areas of the images are used to store global information while discarding the local one. Adding learnable "register tokens" dedicated to global information processing helped eliminate the high-norm artifacts, improving performance. Similarly, Wen et al. [29] used learnable registers to summarize salient visual information and remove the image tokens altogether, improving the efficiency of vision-language models.

Information Flow in Text-Only VLMs. In text-only VLMs, this kind of summarization is particularly useful since much of the attention to image tokens is concentrated in the initial layers, likely on a few "anchor tokens" [28, 59]. Basu et al. [28] demonstrated with multimodal causal tracing that, in VQA tasks, a subset of strongly attended *late image tokens* in early layers transfer information to the text. However, a later ablation study by Neo et al. [27] showed that important information about objects is often localized in tokens corresponding to the object's position in the image. Interestingly, they also observed that image tokens, which may have different statistical properties compared to textual tokens [60], in the deeper layers encode vocabulary words that describe the objects they belong to.

Multimodal Vision Language Models. Many approaches have been proposed to adapt decoderonly LLMs trained with autoregressive loss for image understanding [10-13, 61, 62] and multimodal generation tasks [16, 19–21, 24–26, 31, 34, 63, 64]. These models consist of an LLM, an image encoder, an image decoder (for multimodal-output VLMs), and adapters between the LLM and visual components. The success of diffusion-based image generation has led to the widespread use of stable diffusion as the decoding architecture in multimodal-output models [16, 19, 21, 24, 25, 64]. For the image encoding part, most architectures use a vision transformer (ViT), often that of CLIP [16, 19, 21, 24, 25, 34, 63, 64], due to the effectiveness of ViT embeddings in encoding image semantics [65]. The outputs from the visual encoder are either directly projected into the LLM embedding space [21, 24, 25, 65] or, in the case of multimodal generation, often mapped to a discrete set of tokens [16, 19, 63]. A drawback of these methods is the complexity of the encoding architecture. Moreover, the image-text contrastive loss can make the image representations resemble textual ones excessively [62], with text often favored as the preferred modality [17, 19, 25]. Alternative approaches use VQ-VAEs [66, 67] to encode and decode the images and train the LLMs from scratch on multimodal datasets [20, 26]. This native multimodal approach fuses the visual and textual representations early in training instead of using complex pipelines to bridge the modality gap [61, 65, 68, 69] and eliminates the need for adapters, greatly simplifying the architecture design.

B Experimental Setup

B.1 Models Architectures

Table A2 shows the detailed configurations of each model employed in the experiments.

Parameters	Chameleon-7B	Chameleon-34B	Emu3-8B	LLaVA-7B	Pixtral-12B	Janus-1.3B	VILA-U-7B
N. Parameters	7B	34B	8B	7B	12B	1.3B	7B
Backbone	-	-	-	Qwen2-7B	Mistral Nemo 12B	DeepSeek-LLM-1.3B	LLaMa-2-7B
Visual Encoder	VQ-GAN	VQ-GAN	VQ-GAN	SigLIP	Pixtral-ViT	VQ	RQ-VAE+SigLIP
Image Generation	Yes	Yes	Yes	No	No	Yes	Yes
N. Layers	32	48	32	28	40	24	32
N. Attention Heads	32	48	32	28	32	16	32
Hidden Size	4096	8192	4096	3584	5120	2048	4096

Table A2: **Comparison of Model Architectures.** Detailed configurations of each model, including parameters, visual encoder type, generation capabilities, layer counts, attention heads, and hidden sizes.

B.2 Datasets details

Captioning prompt constructed from Flickr30K and MSCOCO. The Flickr-30k [38] dataset and MSCOCO dataset [39] contain respectively 30,000 and 330,000 images. Each sample of the first and approximately 200,000 of the latter is associated with five captions. For the captioning benchmarks presented in tables 1 and A3 we constructed the prompts with the following structure: "(image) Provide a one-sentence caption for the provided image. __"

Image-text prompt constructed from Flickr-30k. To generate prompts used in section 3.1 that alternate between text-first and image-first formats, we structured the inputs by concatenating all five captions for each image, either preceding or following the image itself, depending on the intended sequence.

Question answering prompt constructed from VQAv2. The dataset VQAv2 [37] consists of 204,721 images from the MSCOCO dataset each associated with a variable number multiple choice question answering. For this benchmark presented in tables 1 and A3 we constructed prompts concatenating the image, the first of the associated question, and the instruction "Answer the question using a single word or phrase."

Image-text prompt constructed from ImageNet. For the experiment described in sections 3.2 and 3.4 we use the following prompt construction "<image> This animal is a". For the activation patching experiment described in section 4.1 we use the same prompt for Chameleon-7B. For Chameleon-34B, we adopt a slightly modified prompt: "<image> Answer the question using a single word, number, or short phrase. This animal is a" as Chameleon-34B refused to respond to the original prompt.

For the activation patching experiment, described in section 4.1, we manually select 20 class from the list above to ensure semantic diversity in the animal represented, and then we random sample 100 images for each class, obtaining a total of 2000 images. The selected pairs are:

```
• (american_alligator.n.01, arabian_camel.n.01)
```

```
• (bald_eagle.n.01, barn_spider.n.01)
```

- (bee_eater.n.01, cheetah.n.01)
- (flamingo.n.01, great_grey_owl.n.01)
- (green_mamba.n.01, grey_whale.n.01)
- (hippopotamus.n.01, jaguar.n.01)
- (king_penguin.n.01, kit_fox.n.01)
- (lionfish.n.01, macaw.n.01)
- (proboscis_monkey.n.01, siberian_husky.n.01)
- (tailed_frog.n.01, trilobite.n.01)

B.3 Finetuning details

Emu3-Gen finetuning. We fine-tuned Emu3 on a mixture of datasets using 37.5k samples from the VQAv2 [37] and MS-COCO-2014 [39] training sets and 150k samples the LLaVA-instruct-150K using the Hugging Face implementation. We fine-tuned Emu3 for one epoch with a batch size of 64 using LoRA with a rank of 64, an alpha of 16, Adam optimizer without weight decay, and cosine annealing scheduler starting with a learning rate of 5e-5.

Removing the narrow gate with fine-tuning. We fine-tuned Chameleon-7b and Emu3 on a dataset of 15k samples, 7.5k samples from the LLaVA-instruct-150K training set, and 7.5k from the VQAv2 training set. We fine-tuned the models for one epoch with a batch size of 128, using LoRA with a rank of 128, alpha of 256, Adam optimizer with weight decay 0.1, cosine annealing learning scheduler with a starting learning rate of 2e-5.

B.4 Advanced density peaks clustering.

The Advanced Density Peaks (ADP) clustering [52] is a mode-seeking, density-based clustering algorithm that finds the modes of the probability density on the data's low-dimensional data without performing any explicit dimensional reduction. We summarize the main steps below and refer the interested reader to the original paper [52].

The algorithm consists of three steps: the estimation of the data's intrinsic dimension, the estimation of the local density around each point, and a final density-based clustering of the data. Following previous works [70–72], we estimate the intrinsic dimension with Gride, a neighbor-based intrinsic dimension estimator, setting the rank k of the nearest neighbor involved in the estimate to 16 (see [72] for more details). We then measure the local density around each data point with a kNN approach: $\rho_{i,k} = \frac{k}{NV_k}$. Here, N is the number of data points, and V is the volume of the ball, which has a radius equal to the distance between the point i and its kth nearest neighbor. Importantly, we measure the volume on the intrinsic manifold using the intrinsic dimension value estimated in the first step.

The third step is the density-based clustering. With the knowledge of the ρ_i , we find a collection of density peaks $\mathcal{C} = \{c^1, ... c^n\}$, assign the data points around them, and find the density $\rho^{\alpha,\beta}$ of saddle points between a pair of clusters c_α c_β with the procedure described in [52]. The statistical reliability of the peaks is assessed with a t-test on $\log \rho^\alpha - \log \rho^{\alpha,\beta}$, where ρ^α is the maximum density of peak c_α , and $\rho^{\alpha,\beta}$ the density of the saddle point between c_α and c_β . Once the confidence level Z is fixed, all the clusters that do not pass the t-test are merged since the value of their density peaks are considered indistinguishable from the nearby saddle point. The process is repeated until all the peaks satisfy the t-test and are statistically robust with a confidence Z [52]. In the analysis reported in section 3.1, we remove clusters with size lower than 50.

C Additional Results

C.1 Modality Gap in VLMs

We extend the results of section 3.1 to Janus, Pixtral and VILA-U. As illustrated in Figure A1, all of the additional models exhibit an alignment between modalities that grows along the model depth, similar to the behavior observed in LLaVA. Figure A1-left confirms this alignment across modalities for all the models. Figure A1-right reveals that the homogeneity within modality clusters is lower compared to the Chameleon models and Emu3. This indicates that LLaVA, Janus, Pixtral and VILA-U have a smaller modality gap.

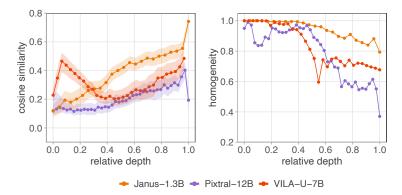


Figure A1: **Modality Gap in VLMs.** (**left**) Cosine similarity between text and image token embeddings as a function of model depth reflects the orthogonality of modalities in Janus, Pixtral and VILA-U models. Points represent median cosine similarity, with shaded areas indicating the interquartile range. (**right**) Homogeneity score of token clusters generated via Advanced Density Peaks with respect to their original modality.

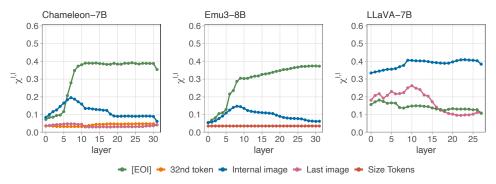


Figure A2: **Localization of Visual Semantic Information.** Neighborhood overlap between visual token embeddings and caption embeddings on the MS-COCO dataset for the models Chameleon-7B, Emu3-8B and LLaVA-7B. The results closely resemble those of figure 4

C.2 A comparison of the semantic content of MS-COCO images in [EOI] and Qwen-GTE captions embeddings.

In section 3.3 we investigated the alignment between token representations across layers and the classes of ImageNet. In this section, we will now extend those findings to the MS-COCO captioning task to investigate if the results generalize to tasks different from classification. We will proceed as follows: first, for each image in the MS-COCO dataset, we extract the activations of the relevant visual tokens at each layer of every model analyzed in our study; Concurrently, we embedded the captions for these images using the Qwen2-7B-GTE text encoder [73]. To measure the alignment between these two sets of representations we used neighborhood overlap described at section 2.1 where in this context the reference data representations are the embeddings from the text encoder. In other words for each sample i, we computed the fraction of shared k-nearest neighbors between the visual token embedding and the caption embedding according to equation (2). This approach allowed us to perform a direct comparison of the semantic alignment across all models and token types.

The results for the models of figure 4 are presented in figure A2, while the profiles for the remaining models are depicted in figure A3-right. We can notice that both group plots closely mirror the findings from our initial ImageNet experiments. This strong correspondence indicates that the localization of semantic information within specific tokens is a consistent characteristic of models like Chameleon and Emu, and that this phenomenon is not task-dependent. Furthermore, the high degree of alignment between the representations of these specialized visual tokens and the caption embeddings suggests that they encode high-level semantic abstractions that are comparable to those captured by sophisticated text-based sentence transformers.

C.3 Cross-Modal Attention and Semantic Content of Visual Tokens

The results for Chameleon-34B closely mirror those of Chameleon-7B and Emu3, as discussed in section 3.2. Figure A3-left visualizes the cross-modal attention contributions of image tokens, where tokens contributing more than 1% individually are shown explicitly, while the rest are grouped as "internal image". In Chameleon-34B, cross-modal attention is primarily concentrated on three tokens—[E01], the first image token, and the last. In contrast, Janus, Pixtral, and VILA-U exhibit a pattern similar to LLaVA (section 3.2), where attention to special tokens remains low. At the same time, a significant portion of the information is retained within internal image tokens. Figure A3-center further reveals that in Chameleon-34B as in the other native multimodal VLMs, [E0I] consistently maintains high $\chi_k^{l,gt}$ values across layers, indicating its dominant role in encoding visual information. Meanwhile, all other tokens exhibit near-zero overlap values, suggesting minimal contribution to meaningful visual representation. For Janus, both [E01] and internal image tokens carry strong information from the earliest layers, implying that visual encoding may be directly influencing their representations. Similarly in Pixtral, both the [E0I] and the last image token exhibit a $\chi_h^{l,gt}$ higher than that of LLaVA, although still below Janus. However, as shown in appendix C.4, the information encoded in these special tokens do not play a significant role in information flow within the backbone models. Among the non-native multimodal VLMs, VILA-U most closely resembles LLaVA in terms

of semantic alignment, as most of the information is concentrated within the internal image tokens, while all other special tokens achieving negligible overlap scores.

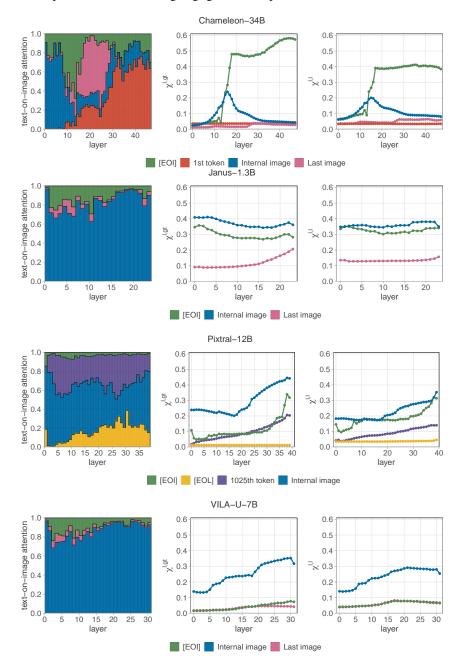


Figure A3: Localization of Information in Different Models. Each row corresponds to a model. The left panel shows the contribution of image token positions to total text-on-image attention. The center panel presents the neighborhood overlap between selected image tokens and ImageNet labels. The right panel finally presents the neighborhood overlap between visual token embeddings and caption embeddings on the MS-COCO dataset.

C.4 Effect of attention knockout on vision-language tasks

Table A3 extends the results of table 1, showing the ablation effect on all tokens identified for their high cross-modal attention. The analysis further confirms that none of the special tokens, except for [E0I], play a significant role, as their ablation does not noticeably affect performance on any benchmarks.

Model	Ablation	VQAv2	Flickr	MS-COCO	ImageNet $(\chi_k^{out,gt})$
	-	0.51	0.34	0.48	0.46
Chameleon-7B	$text \rightarrow [EOI]$	0.25	0.04	0.04	0.01
	$text \rightarrow Last-Image$	0.42	0.20	0.32	0.46
	$\text{text} \rightarrow 32^{nd}$	0.51	0.34	0.49	0.46
Chameleon-34B	-	0.59	0.41	0.47	0.43
	$text \rightarrow [EOI]$	0.39	0.01	0.01	0.04
	$text \rightarrow Last-Image$	0.58	0.39	0.46	0.41
	$text \rightarrow First\text{-Image}$	0.57	0.38	0.46	0.42
	-	0.57	0.29	0.63	0.35
Emu3	$\text{text} \rightarrow [\texttt{EOI}]$	0.48	0.13	0.32	0.24
	$text \rightarrow img$	0.42	0.22	0.54	0.30
LLaVA	-	0.8	0.70	0.98	0.5
	$text \rightarrow [EOI]$	0.8	0.71	0.97	0.45
	$text \rightarrow img$	0.00	0.00	0.01	0.05
	-	0.79	0.59	0.68	0.63
Pixtral	$text \rightarrow [EOI]$	0.77	0.62	0.68	0.61
rixuai	$text \rightarrow [EOL]s$	0.76	0.54	0.61	0.61
	$text \rightarrow Last-Image$	0.77	0.57	0.68	0.62
	text $\rightarrow 1025^{th}$	0.77	0.55	0.62	0.61
	$text \rightarrow img$	0.37	0.00	0.01	0.06
Janus	-	0.51	0.32	0.43	0.65
	$text \rightarrow [EOI]$	0.51	0.32	0.42	0.65
	$text \rightarrow img$	0.00	0.01	0.01	0.06
VILA-U	-	0.71	0.42	0.61	0.64
	$text \rightarrow \texttt{[EOI]}$	0.72	0.43	0.61	0.65
	$text \rightarrow Last-Image$	0.72	0.42	0.61	0.64
	$text \rightarrow img$	0.39	0.01	0.02	0.08

Table A3: Effect of Attention Knockout on Image Understanding Tasks. Performance of the models on visual question answering (VQAv2), image captioning (Flickr-30k and MS-COCO), and image classification (ImageNet) under different ablation settings. We ablated the special tokens with high cross-modality attention by removing their communication with the text tokens. Numbers in bold mark the worst performance for each model and task.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly state the claims made, including the contributions made in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors? Answer: [Yes]

Justification: Limitations are discussed in the dedicated section towards the end of the paper. Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Inside the "Background and Methods" section we provide a detailed description of the methods and the experiments presented in this work.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code to reproduce the results, along with further implementation details, is included in the supplementary material. The experiments are conducted with publicly available datasets.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Experimental setting, full implementation details and prompt structure necessary for reproducibility are included in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars are included when the reported quantity has the sufficient statistical variability.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Information about the compute resources used is provided in the "Reproducibility" section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Yes, the paper conforms, in every respect, with the NeurIPS Code of Ethics Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In the "Discussion and Conclusions" section we warn against potential malicious or unintended uses of the results presented in the paper, as well as discuss potential positive outcomes of our findings.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes],

Justification: All assets (code, data, and models) used in the paper have been properly cited. Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer:

[NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

• For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.