

Performance Gap in Entity Knowledge Extraction Across Modalities in Vision Language Models

Ido Cohen Daniela Gottesman Mor Geva Raja Giryes

Tel Aviv University

{idoc@mail, gottesman3@mail, morgeva@tauex, raja@tauex}.tau.ac.il

Abstract

Vision-language models (VLMs) excel at extracting and reasoning about information from images. Yet, their capacity to leverage internal knowledge about specific entities remains under-explored. This work investigates the disparity in model performance when answering factual questions about an entity described in text versus depicted in an image. Our results reveal a significant accuracy drop — reaching 18% for some models — when the entity is presented visually instead of textually. To study this gap we present POPVQA, a dataset which allows separating entity recognition and question answering, and use it to benchmark several models. We hypothesize that this decline arises from limitations in how information flows from image tokens to query tokens. Thus, we use mechanistic interpretability tools to reveal that, although image tokens are preprocessed by the vision encoder, meaningful information flow from these tokens occurs only in the much deeper layers. Furthermore, critical image processing happens in the language model’s middle layers, allowing few layers for consecutive reasoning, highlighting a potential inefficiency in how the model utilizes its layers for reasoning. These insights shed light on the internal mechanics of VLMs and offer pathways for enhancing their reasoning capabilities. POPVQA can be found [at this link](#).

1 Introduction

Over the past several years, vision-language models (VLMs) have made remarkable progress in linking visual inputs with natural language descriptions, enabling them to perform tasks such as image recognition (Radford et al., 2021a; Mirza et al., 2024), video understanding (Lin et al., 2024), and Visual Question Answering (VQA) (Liu et al., 2023b,a). Much of the existing research evaluating these models focuses on their capacity to understand and describe information presented in the input image: counting objects, assessing complex visual arrangements, and understanding interactions between humans and objects (Dang et al.,

2024; Basu et al., 2024). While this has demonstrated that VLMs can integrate visual and textual information effectively, it does not fully address a critical question:

how well can VLMs exploit their internal knowledge about entities presented in an image rather than ones simply described in text?

This question opens a new angle on VLM evaluation. Instead of probing their grasp of broad scenes, we consider a different scenario: retrieving factual knowledge about visually presented entities. For example, if the subject of a question is a notable figure, say Robin Williams, we can ask the model “*Who is Robin William’s spouse?*”, which we refer to as a textual representation. Or we can provide the model with an image of Robin Williams and ask “*Who is the spouse of the subject in this image?*”, which we refer to as a visual representation. The second case may be viewed as a two-hop reasoning task. First, the model must recognize the entity visually—*identification*. Next, it must connect this recognized entity to its stored factual knowledge—*extraction*. Figure 1 provides an example where the model contains the knowledge to answer the question, but fails to extract it when provided with a visual representation. A detailed explanation of this evaluation appears in Section 4.

Our empirical findings reveal a large accuracy drop when relying on visual representations rather than textual ones. Further probing suggests that this stems from the model investing too many layers in identifying the subject visually, leaving fewer layers for the second hop of factual reasoning (Biran et al., 2024). To investigate these dynamics more concretely, we utilize techniques of mechanistic interpretability such as activation patching (Vig et al., 2020) and attention knockout (Geva et al., 2023). They allow us to intervene in the model’s processing mid-inference, altering how and when image information influences the final answer. By tracking the flow of information through the network, we show that entity recognition and knowledge extraction are deeply entangled processes, with identification




Figure 1: Illustration of our motivating observation. The model successfully extracts its knowledge about the entity when the entity is represented in text. It further successfully identifies the entity in the image, but it fails to combine the two.

dominating the representational space longer than what one would expect. Consequently, the model’s internal knowledge is leveraged less efficiently.

The contribution of our work is twofold. First, we release POPVQA, a dataset that facilitates validating entity identification prior to question answering. Second, we provide a better understanding of how VLMs process entities and integrate their internal factual knowledge with image information, highlighting a gap between the processing of visual and textual information. Instead of treating entity-level knowledge extraction as a monolithic process, we reveal that it involves at least two interdependent steps. We show that the initial reliance on visual representations can restrict later factual reasoning compared to textual representations. This understanding highlights the need for more multi-modality specific evaluation protocols.

Our experiments, which measure the performance gap and apply activation-based hidden state interventions, provide a foundation for future efforts to enhance the models’ ability to draw on their underlying knowledge when presented with visual inputs. Code to reproduce our experiments can be found [at this link](#).

2 Problem Setup

Goal. We ask how well VLMs can leverage their internal factual knowledge when presented with visual representations of entities, compared to textual ones. We suggest that VLMs perform two key steps: identifying the entity in the image, followed by connecting it to its internal knowledge to provide factual information. By isolating these steps and examining where models might fall short, we aim to pinpoint the layer-by-layer flow of information and advance our understanding about why models struggle to draw on their internal knowledge when relying on visual inputs.

Preliminaries and Notation. VLMs are architectures that integrate visual and textual modalities to perform tasks such as image captioning, visual question answering and more. Our VLM analysis

in this work is mainly focused on the common architecture of LLaVA (Liu et al., 2023b), and thus, we will follow the same notation.

The VLM is composed of a pre-trained visual encoder g , a projection layer W , and an LLM f . Given an input image \mathbf{X}_v and text query \mathbf{X}_t , \mathbf{X}_v is first processed by g to provide the visual features $\mathbf{Z}_v = g(\mathbf{X}_v)$ then passed through the projection layer to align the visual features into the LLM’s embedding space as a series of tokens $\mathbf{H}_v = W \cdot \mathbf{Z}_v$. In parallel, \mathbf{X}_t is passed through the LLM’s tokenizer and embedding layer to create a series of tokens \mathbf{H}_t . Next the two are concatenated and passed through the LLM $f([\mathbf{H}_v, \mathbf{H}_t])$. Figure 2, which illustrates the experiments we perform in the paper (see more details in Section 5 and 6), can be used also to better understand these notations.

Every token in this input can be written as $\mathbf{h}_{i,v}^0$ or $\mathbf{h}_{i,t}^0$ where the v and t indicate a visual or textual token, the i indicates the i^{th} visual or textual tokens, and the 0 indicates input to layer 0. As the language model processes the entire input at each layer l , it builds a representation $(\mathbf{h}_{1,v}^l \dots \mathbf{h}_{n,v}^l, \mathbf{h}_{1,t}^l \dots \mathbf{h}_{m,t}^l, \mathbf{h}_{1,g}^l \dots \mathbf{h}_{k,g}^l)$, where $\mathbf{h}_{1,g}^l \dots \mathbf{h}_{k,g}^l$ are the generated tokens hidden representations, $|\mathbf{H}_v| = n$, $|\mathbf{H}_t| = m$ and l is the layer. We refer to the tokens once inside the LLM as hidden states.

In our experiments we focus on how information flows from the visual hidden states $(\mathbf{h}_{1,v}^l \dots \mathbf{h}_{n,v}^l)$ to the textual and generated tokens. We do this by applying techniques of mechanistic interpretability.

Dataset and Experiments Overview. To analyze how VLMs integrate visual and textual information, we introduce a new dataset called POPVQA in Section 3, featuring popular entities that VLMs are likely to have internal knowledge about. This dataset focuses on a two-hop reasoning process: (i) identification of an entity based on an image, followed by (ii) extraction of factual information related to that entity. We compare model responses when the entity is given as text (“Robin Williams”) versus as an image, pro-

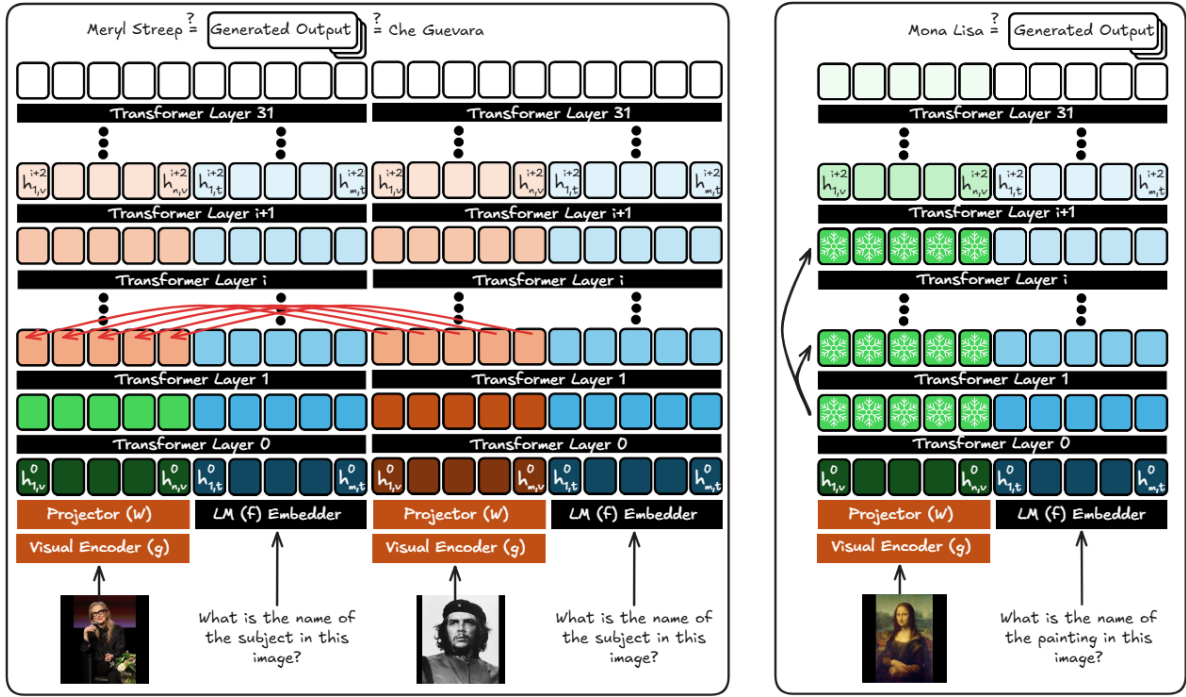


Figure 2: The left pane is an illustration for the experiment in Section 5. The hidden states at the positions of the visual tokens are copied from the forward pass of the injected entity at the source layer (in this case layer 1) replacing the hidden states of the original entity’s forward pass at the same layer and position. The right pane is an illustration of the experiment in Section 6. Starting from the source layer, the hidden states at the positions of the visual tokens are copied and patched instead of the hidden states at the same positions in the subsequent layers until the target layer, effectively “freezing” their processing by the LLM.

viding insight into the gap between visual and textual representations. In these experiments, we first present a motivational case study in Section 4, demonstrating that various VLMs (e.g., LLaVA7B, Qwen2-VL) show a marked accuracy drop when relying on visual inputs. This experiment identifies a gap between the use of visual and language information in VLMs.

Motivated by the first experiment, we turn to further explore how VLMs rely on visual and textual information. We delve deeper into the reasons behind the gap, focusing on LLaVA with different encoders such as CLIP and SigLIP. Specifically, we aim to understand at what point in the inference pass the entity in an image is identified, and this identity is propagated to the query and generated token positions.

Through activation patching (Vig et al., 2020) and other mechanistic interpretability methods, we trace how and where entity recognition occurs, and how it impacts the model’s ability to retrieve the correct facts. Section 5 describes how we employ activation patching to identify the layer at which the model extracts the entity identity, establishing an upper bound on identification layers. Building on these findings, Section 6 applies forward activation patching to investigate whether the entity identity

could be inferred at earlier layers.

We include additional experiments using attention knockouts (Geva et al., 2023) in appendix B, which further validate our conclusions. Throughout the paper, we use POPVQA to systematically evaluate VLM performance on visually driven factual retrieval and demonstrate how visual and textual components interact within a unified transformer architecture.

3 The POPVQA Dataset

In order to isolate the model’s ability to extract knowledge about a given entity from its ability to identify it in the first place, we want to restrict our questions to entities that the model recognizes. Other datasets that have knowledge seeking goals such as infoSeek (Chen et al., 2023) or VeQuAE (Lerner et al., 2022) do not include the entity in the image but rather only the answer to the question, thus, making it difficult to do this separation. Therefore, we built the POPVQA dataset to facilitate this need. It is inspired by PopQA (Mallen et al., 2023) that created a QA dataset for various entities using Wikipedia.

The first part of POPVQA comprises of 15,395 entity-image pairs from Wikidata (Vrandečić and Krötzsch, 2014) along with aliases and popularity

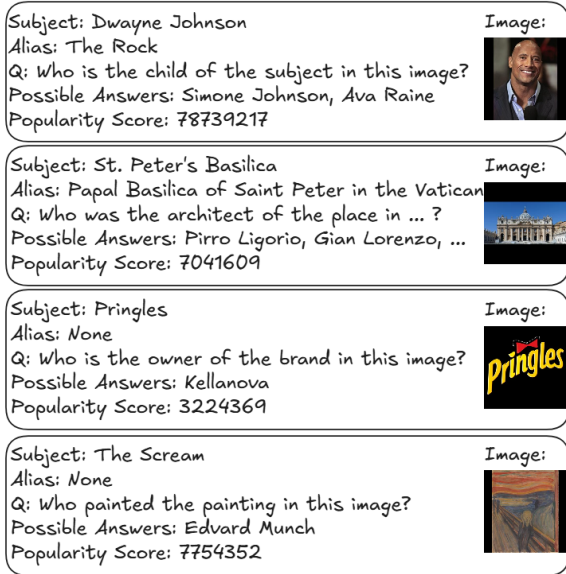



Figure 3: The POPVQA dataset contains various entity-image pairs of celebrities, landmarks, brand logos and painting. For each entity we keep the image, its subject, popularity score, possible aliases and questions and answers related to that entity.

score (number of page views). The vast majority of entities are celebrities (63.6%). To allow ablation of other entity types, we added landmarks (18%), such as famous structures or UNESCO sites, famous paintings (14.6%), and brand logos (3.8%). All images are resized to 336×336 which is a standard input resolution, by adding black borders to create a square, then resizing. This was done in order to maintain the original aspect ratio, without losing any details. The popularity score was used only in the creation of the dataset such that we select only entities with high popularity, which increases the chance that the models will recognize them (Mallen et al., 2023). We do not use this value directly in the experiments.

The second part of POPVQA provides information seeking questions about the entities; at least two per entity. To generate the questions, we convert subject-relation-object triplets from Wikidata into question-answer pairs using manually designed templates (e.g., Petroni et al., 2019; Sciavolino et al., 2021). For example, the triplet (*Robin Williams, spouse, Marsha Garces*) is transformed into the question, “*Who is the spouse of Robin Williams?*”, with the answer “*Marsha Garces*”. We filter out questions where the answer can be

deduced from the image alone without recognizing the entity, such as eye color. We point out that this is not bulletproof as, for example, the native language of an athlete can be guessed if the athlete is wearing a national team jersey.

One legitimate concern is the presence of stale information, particularly for entities like celebrities whose personal circumstances (e.g., marital status, offspring) may change over time. However, this issue affects only a small subset of the entities in POPVQA, and an even smaller fraction of the associated questions. To mitigate this, we allow multiple correct answers per question where applicable. Additionally, we provide code for rebuilding POPVQA using up-to-date snapshots of Wikipedia and Wikidata, enabling users to generate a version that reflects more recent factual knowledge.

In the textual representation case, the model is presented with the question as is. For the visual case, the model receives an image of the entity along with a reformulated query, such as “*Who is the spouse of the subject of this image?*”. This ensures that the model needs to internally identify the subject depicted in the image before reasoning and generating an answer.

4 Motivating Experiment

We start by quantifying the model’s performance gap in entity knowledge extraction when presented with visual versus textual representations of entities. Given that the model recognizes the entity, asking about the visual representation, is not a harder task in essence, but a matter of computational capacity. Our motivating experiment demonstrate that, indeed, using the visual representation is more challenging. The subsequent experiments aim at explaining this gap and suggest that the reason for this bigger challenge is that the model needs to use part of its capacity to first identify the subject in the image and only then it can answer the question with the remaining capacity. Thus, as the model is of a fixed depth, once the model identifies the entity, it becomes a question of extracting the answer within the limits of the ‘remaining compute’ of the model.

Given the above discussion, we relate to the textual representation accuracy as the potential of the model. The lower the drop between the textual and visual accuracy, the better the model is at utilizing the visual input to elicit extraction of its internal knowledge. This could indicate better alignment between the embedding spaces of the LLM and the visual backbone.

To run the experiment, we first identify which

Model	#Identified	Img Acc	Txt Acc	Drop
LLaVA7B	925	0.276	0.453	0.177
LM-CLIP	844	0.216	0.378	0.163
LM-SigLIP	660	0.201	0.377	0.176
LLaVA34B	1286	0.534	0.656	0.121
Qwen2-VL	3143	0.433	0.476	0.043

Table 1: Results on POPVQA. Reported accuracy is first calculated per subject then averaged across the entire dataset.

entities each model is capable of recognizing. This is done by presenting the model with an image of each entity along with a prompt asking it to identify the subject (e.g., “Identify by name the subject in this image.”). Entities for which the model returns the correct name are considered identified. We then evaluate the model’s ability to answer factual questions for these identified entities given the two representations, textual versus visual, using the formulation in Section 3. This setup ensures that performance differences are not due to failures in entity recognition, but reflect differences in how well the model retrieves knowledge once the entity is known.

We show results for the following models; LLaVA-1.5-7B (Liu et al., 2023b) as the canonic base VLM; we refer to it as LLaVA7B. LLaVA-MORE-8B (Cocchi et al., 2025) which uses a larger Llama 3.1 as its LLM and we try both variants using CLIP (Radford et al., 2021b) and SigLIP (Zhai et al., 2023) as backbones allowing for controlled comparison between visual backbones; we refer to these as LM-CLIP and LM-SigLIP. We use LLaVA-1.6-34B to assess the effect of model scale, and refer to it as LLaVA34B. We also compare to Qwen2-VL (Wang et al., 2024) that has a different architecture and training scheme.

Results. Table 1 shows the results for the different models. We measured accuracy per entity as the proportion of correctly answered questions for that entity and calculated the overall mean accuracy across all entities. To verify the significance of the results, we performed a Wilcoxon signed-rank test between the per-entity accuracies for visual versus textual representations. The p-values range from $p=4.02 \times 10^{-29}$ for the lower drop to $p=3.83 \times 10^{-76}$ for the higher drop. A more fine grained analysis by the entity types can be found in appendix A.

Surprisingly, LLaVA7B outperforms both LLaVA-MORE-8B variants, recognizing more entities and reaching higher accuracy. Both LLaVA-MORE-8B variants have similar accuracy, both with visual and textual representations, but the variant that uses

SigLIP as a visual backbone identifies considerably more entities. LLaVA34B has significantly better results with a lower drop showing the benefit of increasing model depth to improve the computational capacity. Qwen2-VL also exhibits strong results, with the lowest drop, pointing to the advantage of more modern training schemes which include unfreezing the visual backbone parameters during training.

Interestingly, a common mistake made by the model when asked to name the subject’s spouse/parent/sibling/child, is to name the subject itself. This observation provides an indication to our hypothesis that by the time information about the entity propagates from the image tokens to later token positions, there is not enough computational capacity left to use this information, and therefore the model outputs the entity’s name.

Note, though, that some entities did exhibit an increase in accuracy with visual inputs. This occurred specifically when textual input alone was insufficient for reasoning, but visual input provided contextually inferable details. For example, a subject wearing the French national team uniform, implicitly informed the model that the language spoken by the subject is French.

5 Late Information Flow

To analyze the information flow in the VLMs, we apply interventions in their inference. We continue to focus on entity-image pairs for which, without intervention, the model correctly identifies the entity in the image. For all the following experiments we focus on the LLaVA architecture due to its simple architecture that allows for controlled probing, leaving other architectures as future work. We use LLaVA7B which successfully recognizes 925 entities from POPVQA. Yet, in order to show comparative results with different backbones we show also results on LM-SigLIP and LM-CLIP as well. These backbones were chosen because their training process inherently brings visual and textual representations closer which makes them a common choice in VLMs.

We begin by identifying the point of critical information flow from the image hidden representation to later positions. This will provide an upper bound on the layer depth at which entity identification happens and point us to where this information is likely to propagate onward. We can subsequently infer that by this point the model mostly finalizes its image-based representation enrichment and focuses on leveraging this information for prediction. We hypothesize that

this happens in deep layers, leaving few layers for factual knowledge extraction, which may begin to explain the performance gap presented in Section 4.

The experiment involves two distinct forward passes through the model, using different entities: an **original entity** and a **injected entity** (as shown in Figure 2). Each forward pass is initiated with a prompt asking the model to identify the subject of an image. We extract the image-token representations $\mathbf{h}_{1,v}^\ell \dots \mathbf{h}_{n,v}^\ell$ from the **injected entity** forward pass at a specific layer. These representations are then patched at the same token positions and same layer of the forward pass of the **original entity**. Hence, we call this technique cross patching. After patching, we allow the model to complete its forward pass and predict which entity is identified. The layer at which the predicted entity switches from being the injected entity to the original entity identifies the critical point where the model has completed sufficient processing of the visual representation such that patching does not affect its prediction. When the patching happens between entities of the same type, we use the original prompt. When we patch between different type entities, we replace the reference word (painting, place, etc.) with the word "subject" so as not to bias the output to a certain type.

Results. Figure 4 shows the effect of patching image hidden states across different layers, indicated by the X axis. The first three graphs show results for patching between entities of the same type, for different models. The three models exhibit different behaviors but with similar trends. The main difference between the models is the layer of the critical point, mentioned earlier (where the blue and orange lines cross). Although for all models this occurs after the middle of the model (layer 16), it varies from 17 for LM-SigLIP to 24 for LM-CLIP. It is important to note the higher standard deviation for the LLaVA-MORE-8B variants. For LLaVA7B the critical point is clearly marked by a sudden change at layer 20, while for LM-CLIP the change from predicting the injected entity to the original is more gradual but still clearly after the middle layer.

In all three cases, when the image tokens are patched at or before the middle layer, the model predominantly predicts the injected entity. This strengthens the hypothesis that the middle layers play a critical role in how the model interprets the image. Our results suggest that despite having 15-20 layers to process and propagate information about the original entity to later positions, the information conveyed within the narrow window of the central layers ultimately determines the outcome. Patching beyond

layer 25 results in the model reliably predicting the target entity, suggesting that image-based entity identification is mostly finalized by this stage.

The last (bottom) graph in Figure 4, shows results for patching between different types of entities for LLaVA7B. Note the almost identical behavior to cross patching between same types (top), albeit a slightly faster decline of the orange line. This provides even stronger evidence that the model’s prediction is compelled by the information propagated in latter layers. Even when having 13 layers to attend to the visual hidden states that encode the identity of, say, LeBron James, new information that comes in at layer 14 encoding the logo of Spotify causes the model to predict the latter. Up until layer 14 this happens with a striking 100% of entities, regardless of type.

The results suggest two possible processing paths:

I. Parallel identification and knowledge extraction:

Generation and query hidden states, can be enriched by the image hidden states, via the attention mechanism, concurrently with entity identification before the middle layers. This would imply that the model incrementally combines these processes over several layers, with information passed in later layers having more impact.

II. Sequential processing:

Identification may occur just after the middle layer, with the remaining layers allocated to factual reasoning and prediction tasks. This would leave a limited subset of late layers for extracting factual knowledge about the entity, potentially leading to degraded question-answering performance.

6 Long Processing is Often Necessary

Having identified the middle layers as the point where critical information propagates from image hidden states to later positions, and having established an upper bound on where identification happens, we will try to answer which of the above mentioned processing paths is more likely, and also answer whether all these layers are truly required for identification. Could it be that the entity was identified in earlier layers, but the model was unable to utilize this until later layers?

To answer these questions we apply activation patching of the image hidden states from an early layer into all subsequent layers until layer 20. In detail, at every run we choose a source layer *source* and set $(\mathbf{h}_{1,v}^\ell \dots \mathbf{h}_{n,v}^\ell) = (\mathbf{h}_{1,v}^{source} \dots \mathbf{h}_{n,v}^{source})$ for values of ℓ ranging from *source* to 20. We repeat this for *source* layers 0 to 19. This effectively "freezes" the hidden states at the image token positions, not allowing them to be updated by neither the attention nor MLP

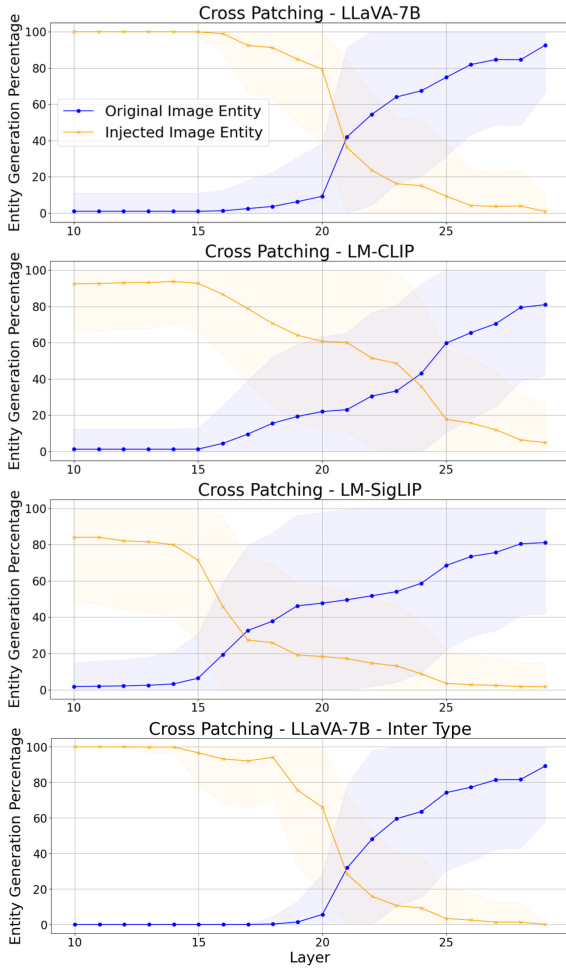


Figure 4: Results for the experiment in Section 5, illustrated in Figure 2. The top three graphs show results when cross patching between forward passes of entities of the same type, for three different models (indicated in the titles). The bottom graph shows results for cross patching between different types, for LLaVA7B. The X axis indicates the injection layer. The orange line indicates the percent of times the injected entity was generated. The blue line is the percent of times the original entity was generated.

modules. Importantly, this does not prevent the query and generated positions to attend to the image tokens, only that they see the same state at all the patched layers. We choose 20 as the end point since this is past the point we identified in the last experiment, where we claim that entity identification is mostly done and propagated to later positions. If identification indeed happens in early layers, "freezing" all the layers until the point of information propagation should retain most of the identification accuracy.

Results. Figure 5 shows the results, where every bar shows the percent of correctly identified entities after applying patching from the *source* layer indicated by the x axis onto all layers until layer 20. There is no dis-

tinction here between different types of entities. From these results we see an interesting difference between the models. For LLaVA7B, processing the image hidden states for more layers is required to reach the same results as LM-SigLIP and LM-CLIP do with between 4 and 10 layers less. All layers seem to take some part in this process, as every layer allows identifying more entities. As expected, this gap diminishes the more layers are allowed for processing. When patching through layers 13 to 20, LLaVA7B catches up to the other models. Interestingly, for both LLaVA-MORE-8B variants, patching through all layers from 0 until 20, still retains almost 50% (LM-SigLIP) and over 30% (LM-CLIP) of recognized entities. This suggests that in these cases, the representation provided by the visual backbone encodes a sufficient signal to identify the entity. This alone would not be enough without the LLM being able to utilize this representation, suggesting better alignment of the visual and textual embedding spaces. These results complement those of the previous experiment showing that LM-SigLIP, which had the earliest critical point in the previous experiments, can identify the most entities with the least amount of processing (manifested in number of layers).

With these results in mind, we turn back to our motivating experiment and ask whether early identification leads to better QA accuracy. As the graph shows, this question is mostly relevant for the LLaVA-MORE-8B variants. In this experiment, we split our dataset into entities where forward patching from early (< 5) layers maintains correct identification. We refer to these as *early-id*, and those requiring late layers, *late-id*.

Table 2 shows the results for both models on both splits. Notice that the early-id entities achieve higher accuracy in both models and this rise is consistent for both the image and textual representation. This consistency between modalities teaches us that better knowledge of an entity (higher accuracy) leads to early recognition, and not the other way around. Had that been the case, we would expect to see a rise in image accuracy for early-id entities without a corresponding rise in text accuracy. This leads to the accuracy drop between image accuracy and text accuracy to actually be higher for the early-id entities. This means that although the model has access to the entity's identity at an early layer, leaving more layers for the second hop, the model does not utilize this advantage. This could indicate that although this information exists, it does not reach the query tokens via the attention mechanism early enough to inform subsequent factual recall. This strengthens the observation in Section 5 that information propagates in the middle layers.




Figure 5: Results for the experiment in Section 6, illustrated in Figure 2. Each bar shows the percent of correctly identified entities after applying patching from the layer indicated by the x axis into all layers up to layer 19.

Model	Split	Img Acc	Txt Acc	Drop
LM-CLIP	Early	0.260	0.442	0.182
LM-CLIP	Late	0.192	0.347	0.154
LM-SigLIP	Early	0.234	0.417	0.183
LM-SigLIP	Late	0.168	0.338	0.170

Table 2: Results on POPVQA, split by early and late identified entities. Early identification is correlated with higher accuracy but not with lower drop.

7 Related Work

While VLMs demonstrate impressive capabilities in tasks such as visual reasoning, object recognition, and multimodal understanding (Basu et al., 2024), their internal mechanisms and interpretability remain active areas of research. Below, we review the related works categorized into general VLM interpretability, cross-modal information flow, attention mechanisms, and mechanistic approaches to visual token processing.

Dang et al. (2024) survey VLMs proposing a framework categorizing interpretability across data, model, and training perspectives. They highlight the alignment and decomposition of modalities as key challenges, emphasizing the need for explainability to ensure trustworthiness in multimodal AI applications. In this context, techniques for token-level, neuron-level, and embedding-level interpretability have been systematically analyzed, offering insights into the mechanisms driving decision-making processes.

Previous research has examined the internal information distribution within these models (Basu et al., 2024), highlighted the limitations arising from pre-trained vision encoders (Tong et al., 2024), and investigated the issue of hallucinations in VLMs (Bai et al., 2024; Jiang et al., 2025). Studies such as Huo et al. (2024) delve into neuron-level interpretations, uncovering how specific neurons contribute to multimodal reasoning. Similarly, Jiang et al. (2024) address hallucination issues in VLMs by examining the influence of vision-language representations. Both works highlight the importance of dissecting model components to improve transparency and robustness.

On the methodological front, attention-based mechanisms remain a cornerstone for interpretability. Tao et al. (2024) explore global and local semantic representations, while Liu et al. (2023b) focus on visual instruction tuning to enhance alignment in cross-modal embeddings. Chefer et al. (2021) provides a technique to highlight the input pixels in an image that affects the decision of a transformer. Chen et al. (2022) further extends this approach for CLIP and Yellinek et al. (2023) suggests using a tree structure to improve the use of HilaCAM (Chefer et al., 2021) in VLMs.

Finally, causal approaches and model manipulation techniques have gained traction for evaluating VLMs. Palit et al. (2023) introduce causal tracing tools to understand information flow in vision-language systems, while Zhang et al. (2024) analyze information redundancy and relevance across reasoning tasks. These works demonstrate the efficacy of manipulating hidden states and attention mechanisms to investigate and improve model performance.

Some datasets probe how well VLMs handle knowledge-intensive queries (Lerner et al., 2024; Zhu et al., 2024; Onoe et al., 2025; Yanuka et al., 2025; Wei et al., 2025). For instance, InfoSeek (Chen et al., 2023) is a VQA dataset tailored for information-seeking questions that cannot be resolved using only common sense knowledge. Their findings reveal that state-of-the-art VLMs often struggle with these questions; however, fine-tuning on InfoSeek enables them to leverage more fine-grained knowledge learned during pre-training. Unlike InfoSeek, which focuses on direct information-seeking queries, our dataset explicitly includes a bridge entity as part of a two-hop reasoning process—requiring both identification of the entity and extraction of the relevant fact. This design allows us to disentangle the efficacy of visual recognition from the subsequent factual retrieval and more directly analyze the gap between textual and visual representations.

Our work is most related to Neo et al. (2025); Kaduri et al. (2024) who investigated processing of information in VLMs. Neo et al. (2025) provided insights into the role of attention layers in capturing visual attributes and object-specific details. Contemporary to us, Kaduri et al. (2024) demonstrated that query tokens store high-level image descriptors. Their findings reveal that middle layers of the model are critical for cross-modal knowledge transfer, while fine-grained visual details are extracted in a spatially localized manner. Yet, both of these works do not study the gap we reveal between the processing of textual and visual representations.

Parcalabescu and Frank (2023) introduce a method that uses Shapley values to quantify how VLMs balance information from each modality, revealing varying degrees of unimodal collapse. Parcalabescu and Frank (2025) examine how VLM decoders rely on visual and textual inputs differently when producing answers versus explanations, noting greater reliance on text overall but an increased role for images in explanation tasks. Unlike these approaches, which focus on measuring how images and text are combined in general, our work studies a specific two-step reasoning process—entity identification followed by factual retrieval—and highlights a distinct performance gap when relying on visual instead of textual cues. Zhu et al. (2024) propose a pipeline to detect and mitigate what they refer to as parametric knowledge conflicts in large vision-language models, introducing a dynamic contrastive decoding method that improves performance on InfoSeek. While their work addresses cross-modal inconsistencies by focusing on parametric conflicts, we contribute a distinct perspective centered on bridging entity identification and factual retrieval.

Interestingly, when compared to human behavior, the behavior of VLMs we find in this paper contrasts with findings from cognitive psychology. Specifically, humans tend to perform better with visual rather than textual information, a phenomenon known as the picture superiority effect (Jenkins et al., 1967). While explanations for this effect are still debated, the dual-coding theory (Paivio and Csapo, 1973) attributes this to images being encoded both visually and verbally, unlike text which is only encoded verbally. These findings suggest that humans benefit more from visual input, in contrast to VLMs, which are likely biased toward text due to its prevalence in their training data. We leave a further exploration of the relation between the two to future research.

8 Conclusion

This work explores how VLMs process visual entities, enhancing understanding of multi-modal architectures. Our findings highlight that reasoning about entities shown in an image is not a straightforward extension of reasoning about entities described in text, rather, it more closely resembles the processing of multi-hop questions. Across different model and visual backbones we see that models dedicate a substantial portion of their computational "budget" to entity identification and information processing, leaving few layers for the subsequent factual reasoning step. Thus, critical image-related processing often concludes relatively

late in the processing flow, limiting the model's capacity to extract factual knowledge in the remaining layers thus limiting its overall performance. We also find that this issue persists even for entities that could be identified earlier, indicating that current models may not be using their representational capacity efficiently.

By highlighting these challenges, our results offer pathways for improvement. Potential directions for tackling the root of this issue include adjusting training regimes to better separate or balance identification and knowledge extraction steps, as well as designing architectural enhancements that reduce entanglement between these stages, which could lead to VLMs that better, and more reliably, utilize their internal knowledge. For instance, this behavior may be mitigated to some extent by using early-fusion models, that aim to create a shared embedding space between the visual and textual modalities. We leave this exploration for future work.

Limitations

The scope of this work focuses on understanding the cause of the observed behavior, and leaves the solutions as future works.

In the second part of our experiments, we focused on the performance of LLaVA style models, as it is a canonical architecture in the field, and allows applying interventions with ease. While we did observe a gap in other architectures, we leave the probing experimentation on these for future work. Indeed, we did not test all type of VLMs, where some may exhibit different behavior to the kind described in this work. We believe that studying the gap we report in this work for a larger variety of VLM types would be an interesting direction for a future research.

When using visual representations, the image may contain features that can steer the model towards or away from the correct answer. For example, a football fan wearing his favorite team's shirt could be mistaken for a football player. This work does not go into disentangling the effects of such confounders on the observed gap.

Lastly, extending our analysis to entities with different properties and of additional ethnicities is an important direction for future work.

Acknowledgments

We thank Omri Kaduri and Moran Yanuka for their input and feedback. This work was supported by the Tel Aviv University Center for AI and Data Science (TAD) and the Israeli Science Foundation.

References

- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*.
- Samyadeep Basu, Martin Grayson, Cecily Morrison, Besmira Nushi, Soheil Feizi, and Daniela Massiceti. 2024. Understanding information storage and transfer in multi-modal large language models. *arXiv preprint arXiv:2406.04236*.
- Eden Biran, Daniela Gottesman, Sohee Yang, Mor Geva, and Amir Globerson. 2024. [Hopping too late: Exploring the limitations of large language models on multi-hop queries](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14113–14130, Miami, Florida, USA. Association for Computational Linguistics.
- Hila Chefer, Shir Gur, and Lior Wolf. 2021. Generic attention-model explainability for interpreting bi-modal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 397–406.
- Peijie Chen, Qi Li, Saad Biaz, Trung Bui, and Anh Nguyen. 2022. gscorecam: What objects is clip looking at? In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 1959–1975.
- Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. 2023. [Can pre-trained vision and language models answer visual information-seeking questions?](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14948–14968, Singapore. Association for Computational Linguistics.
- Federico Cocchi, Nicholas Moratelli, Davide Caffagni, Sara Sarto, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. 2025. [Llava-more: A comparative study of llms and visual backbones for enhanced visual instruction tuning](#). *Preprint*, arXiv:2503.15621.
- Yunkai Dang, Kaichen Huang, Jiahao Huo, Yibo Yan, Sirui Huang, Dongrui Liu, Mengxi Gao, Jie Zhang, Chen Qian, Kun Wang, Yong Liu, Jing Shao, Hui Xiong, and Xuming Hu. 2024. Explainable and interpretable multimodal large language models: A comprehensive survey. *ArXiv*.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Jingwei Huo, Yu Yan, Bo Hu, Yuan Yue, and Xipeng Hu. 2024. [MMNeuron: Discovering Neuron-Level Domain-Specific Interpretation in Multimodal Large Language Model](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6801–6816. Association for Computational Linguistics.
- J. R. Jenkins, D. C. Neale, and S. L. Deno. 1967. [Differential memory for picture and word stimuli](#). *Journal of Educational Psychology*, 58(5):303–307.
- Nick Jiang, Anish Kachinthaya, Suzie Petryk, and Yossi Gandelsman. 2024. Interpreting and editing vision-language representations to mitigate hallucinations. *arXiv preprint arXiv:2410.02762*.
- Nick Jiang, Anish Kachinthaya, Suzie Petryk, and Yossi Gandelsman. 2025. Interpreting and editing vision-language representations to mitigate hallucinations. In *ICLR*.
- Omri Kaduri, Shai Bagon, and Tali Dekel. 2024. [What’s in the image? a deep-dive into the vision of vision language models](#). *Preprint*, arXiv:2411.17491.
- Paul Lerner, Olivier Ferret, and Camille Guinaudeau. 2024. Cross-modal retrieval for knowledge-based visual question answering. In *Advances in Information Retrieval*.
- Paul Lerner, Olivier Ferret, Camille Guinaudeau, Hervé Le Borgne, Romaric Besançon, Jose G Moreno, and Jesús Lovón Melgarejo. 2022. [ViQuAE, a dataset for knowledge-based visual question answering about named entities](#). In *Proceedings of The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’22*, New York, NY, USA. Association for Computing Machinery.
- Bin Lin, Yang Ye, Bin Zhu, Jiayi Cui, Munan Ning, Peng Jin, and Li Yuan. 2024. Video-LLaVA: Learning united visual representation by alignment before projection. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5971–5984.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. [LLaVA-NeXT: Improved reasoning, OCR, and world knowledge](#).
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. [Visual instruction tuning](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916. Curran Associates, Inc.
- Alex Mullen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- M. Jehanzeb Mirza, Leonid Karlinsky, Wei Lin, Sivan Doveh, , Jakub Micorek, Mateusz Kozinski, Hilde Kuhene, and Horst Possegger. 2024. [Meta-Prompting for Automating Zero-shot Visual Recognition with LLMs](#). In *ECCV*.
- Clement Neo, Luke Ong, Philip Torr, Mor Geva, David Krueger, and Fazl Barez. 2025. [Towards interpreting visual information processing in vision-language models](#). In *The Thirteenth International Conference on Learning Representations*.

- Yasumasa Onoe, Sunayana Rane, Zachary Berger, Yonatan Bitton, Jaemin Cho, Roopal Garg, Alexander Ku, Zarana Parekh, Jordi Pont-Tuset, Garrett Tanzer, Su Wang, and Jason Baldridge. 2025. Docci: Descriptions of connected and contrasting images. In *ECCV*, pages 291–309.
- Allan Paivio and Kalman Csapo. 1973. [Picture superiority in free recall: Imagery or dual coding?](#) *Cognitive Psychology*, 5(2):176–206.
- Vaishnavi Palit, Rohit Pandey, Ankit Arora, and Paul Pu Liang. 2023. [Towards vision-language mechanistic interpretability: A causal tracing tool for blip.](#) arXiv:2308.14179.
- Letitia Parcalabescu and Anette Frank. 2023. [MM-SHAP: A performance-agnostic metric for measuring multimodal contributions in vision and language models & tasks.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4032–4059, Toronto, Canada. Association for Computational Linguistics.
- Letitia Parcalabescu and Anette Frank. 2025. [Do vision & language decoders use images and text equally? how self-consistent are their explanations?](#) In *The Thirteenth International Conference on Learning Representations*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021a. [Learning Transferable Visual Models from Natural Language Supervision.](#) In *ICML*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021b. [Learning transferable visual models from natural language supervision.](#) In *International Conference on Machine Learning*.
- Christopher Scovelino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2021. [Simple entity-centric questions challenge dense retrievers.](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mingyu Tao, Qingqiu Huang, Kai Xu, Linjie Chen, Yansong Feng, and Dongyan Zhao. 2024. [Probing Multimodal Large Language Models for Global and Local Semantic Representations.](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13050–13056. ELRA and ICCL.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. [Eyes wide shut? exploring the visual shortcomings of multimodal llms.](#) In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9568–9578.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis.](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*, 57(10):78–85.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. [Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution.](#) *arXiv preprint arXiv:2409.12191*.
- Cong Wei, Yang Chen, Haonan Chen, Hexiang Hu, Ge Zhang, Jie Fu, Alan Ritter, and Wenhui Chen. 2025. [Uniir: Training and benchmarking universal multimodal information retrievers.](#) In *ECCV*.
- Moran Yanuka, Assaf Ben-Kish, Yonatan Bitton, Idan Szpektor, and Raja Giryes. 2025. [Bridging the visual gap: Fine-tuning multimodal models with knowledge-adapted captions.](#) In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10497–10518, Albuquerque, New Mexico. Association for Computational Linguistics.
- Nir Yellinek, Leonid Karlinsky, and Raja Giryes. 2023. [3vl: using trees to teach vision & language models compositional concepts.](#) *Preprint*, arXiv:2312.17345.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. [Sigmoid Loss for Language Image Pre-Training.](#) In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11941–11952, Los Alamitos, CA, USA. IEEE Computer Society.
- Xinyu Zhang, Yuxuan Quan, Chunlei Shen, Xiao Yuan, Shuicheng Yan, Lin Xie, Wenhui Wang, Chuang Gu, Hao Tang, and Jieping Ye. 2024. [From redundancy to relevance: Information flow in vlms across reasoning tasks.](#) arXiv:2406.06579.
- Tinghui Zhu, Qin Liu, Fei Wang, Zhengzhong Tu, and Muhao Chen. 2024. [Unraveling cross-modality knowledge conflicts in large vision-language models.](#) *Preprint*, arXiv:2410.03659.

A Motivating Experiment - Results Breakdown By Type

Tables 3 and 4 show a breakdown of the results in section 4 according to type. The first shows the image and text accuracy and drop, while the second shows a fine grained p-value analysis for the Wilcoxon signed-rank test. This breakdown provides some interesting insights. The drop in the paintings case is often low, with Qwen2-VL even doing better with the visual representation. This seems to be due to the model being able to provide more informed guesses for some of the questions. Other than that, there is no correlation between specific types and accuracy drop. It seems that different models have different strengths and weaknesses, for example LM-SigLIP recognizes significantly less brand logos despite the name of the brand often appearing on the logo.

Type	#Identified	Img Acc	Txt Acc	Drop
LLaVA7B				
Brands	343	0.270	0.432	0.162
Celebs	294	0.342	0.543	0.201
Landmarks	168	0.280	0.502	0.222
Paintings	120	0.128	0.299	0.171
LM-CLIP				
Brands	364	0.238	0.374	0.136
Celebs	128	0.278	0.379	0.101
Landmarks	212	0.250	0.486	0.236
Paintings	140	0.083	0.249	0.165
LM-SigLIP				
Brands	237	0.169	0.406	0.237
Celebs	136	0.301	0.454	0.153
Landmarks	149	0.259	0.459	0.200
Paintings	138	0.111	0.207	0.096
LLaVA34B				
Brands	416	0.546	0.630	0.083
Celebs	415	0.564	0.742	0.179
Landmarks	346	0.534	0.655	0.122
Paintings	109	0.435	0.455	0.020
Qwen2-VL				
Brands	462	0.363	0.411	0.049
Celebs	1240	0.418	0.544	0.125
Landmarks	1059	0.474	0.489	0.015
Paintings	382	0.445	0.276	-0.168

Table 3: Fine-grained results on POPVQA broken down by different entity types. Reported accuracy is first calculated per subject then averaged across the entire dataset.

B Additional Experiment: Critical Information is Propagated Late

To investigate the stages of the model’s processing pipeline where entity identification occurs, we ask the model to identify the subject in an image and employ an attention knockout strategy following the methodology of Geva et al. (2023). This involves selectively prohibiting attention from textual tokens to image to-

Split	LLaVA7B	LM-CLIP	LM-SigLIP
Overall	3.83×10^{-76}	1.42×10^{-49}	6.36×10^{-47}
Brands	1.07×10^{-21}	1.16×10^{-16}	4.78×10^{-21}
Celebs	3.48×10^{-40}	1.34×10^{-8}	6.87×10^{-12}
Landmarks	4.88×10^{-18}	6.22×10^{-22}	2.90×10^{-13}
Paintings	2.51×10^{-7}	1.29×10^{-8}	7.33×10^{-5}

Split	LLaVA-34B	Qwen2-VL
Overall	4.27×10^{-56}	4.02×10^{-29}
Brands	9.92×10^{-11}	7.60×10^{-5}
Celebs	3.42×10^{-50}	1.33×10^{-115}
Landmarks	3.59×10^{-13}	1.17×10^{-3}
Paintings	3.11×10^{-1}	1.00

Table 4: Fine-grained p-value analysis of Wilcoxon signed-rank test between the per-entity accuracies for visual versus textual representations, broken down by different entity types.

kens at various layers of the model, effectively blocking cross-modal information flow. We do this by inserting *inf* in the attention mask at $[x, y]$ where x are the indices of $\mathbf{h}_{1,t}^\ell \dots \mathbf{h}_{m,t}^\ell, \mathbf{h}_{1,g}^\ell \dots \mathbf{h}_{k,g}^\ell$, y are the indices of $\mathbf{h}_{1,v}^\ell \dots \mathbf{h}_{n,v}^\ell$, and ℓ is any layer we want to knock out. We conduct this experiment in two directions:

1. **Top-down:** Knocking out attention from a start layer until the final layer of the model.
2. **Bottom-up:** Knocking out attention from the model’s first layer until an end layer.

Notably, part of this experiment has been done in concurrent work by (Kaduri et al., 2024), but on different data, and asking different questions. This experiment checks whether the VLM trend follows Geva et al. (2023) which analyze the information flow for factual recall given textual inputs.

Results. Figure 6 presents both types of knockout on LLaVA7B:

1. **Top-down:** Knocking out the second half or more of the model’s attention results in complete collapse of model prediction. The sharp change in layers 15-20 indicates that these layers are critical for information flow from the image tokens to the query and generation tokens. Knocking out the last third of layers has a minimal effect on the model’s output, indicating the model no longer attends to the visual tokens after the critical mid-layer processing is complete, demonstrating a transition to text-dominated reasoning.
2. **Bottom-up:** Shows that unlike the top layers, restricting the attention in early layers has an immediate and progressive effect on the prediction, underscoring their importance. The effect increases around the middle layers until again collapsing the model’s prediction. This shows that there could be some redundancy in the first layers, and the model can begin attending to




Figure 6: **Attention Knockout:** query and generated tokens are restricted from attending to the image tokens in all layers from layer 0 until the marker (in orange) or from the marker until the last layer (in blue)

image tokens quite late in the processing pipeline and still achieve mostly correct entity identification.

These trends mostly agree with Geva et al. (2023) that knock out attention edges of subject tokens in LLMs. This may indicate that the model treats the image tokens similarly to subject textual tokens. One must remember that unlike the textual tokens, the image tokens are already processed and enriched within the vision encoder. This means that although one could hope that image tokens would enjoy extraction at much earlier stages, allowing for contextualization by the text tokens through more layers, the model exhibits suboptimal integration of image information, relying heavily on mid-layer processing while under utilizing earlier layers for cross-modal fusion.

Qualitative results. To complement the quantitative results, we examined qualitative examples of this degradation, in bottom-up knockout. For instance, Table 5 shows the models output for different values of ℓ on an image of LeBron James in action. This example aligns with the general observed quantitative trend, showing that entity identification degrades progressively when employing bottom up attention knockout, the model continues to extract useful visual features before ultimately failing and hallucinating. Another such example at the end.

C Additional Qualitative Examples of Attention Knockout

Bellow is another result of bottom up attention knockout. Every line shows the output of the model after knocking out another layer.

1. Donald Trump is the 45th President of the United States, serving from January 2017 to January 2021.

Layer	Output
10	LeBron James is an American professional basketball player who was born on December 30, 1979, in Akron, Ohio
19	The image features a man wearing a yellow shirt and a red hat, holding a basketball in his hands. He appears to be a basketball player
20	The image features a man wearing a yellow shirt, standing in front of a crowd of people.
21	The image features a man wearing a suit and tie, standing in front of a building.
22	The image features a man named LeBron James, who is an American professional basketball player.
29	The image features a man named "Mike" who is standing in a field. He is wearing a black shirt

Table 5: An example of bottom-up knockout on an image of LeBron James. The layer numbers indicate the final knockout layer. Only layers with meaningful change in output were included for brevity. Note that the model briefly recovers accurate entity identification.

2. Donald Trump is an American businessman, politician, and television personality. He was born on June 14, 1946
3. Donald Trump is an American businessman, politician, and television personality. He was born on June 14, 1946
4. Donald Trump is an American businessman, politician, and television personality. He was born on June 14, 1946
5. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be a public speaker or a politician
6. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be giving a speech or a presentation
7. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be giving a speech or a presentation
8. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be giving a speech or a presentation
9. The image features a man with a receding hairline, wearing a suit and tie, and giving a speech.

10. The image features a man with a distinctive hairstyle, sitting at a podium and giving a speech.
11. The image features a man with a distinctive hairstyle, sitting at a podium and giving a speech.
12. The image features a man with short hair, wearing a suit and tie, and giving a speech. He appears to be a professional orator
13. The image features a man with a distinctive hairstyle, wearing a suit and tie.
14. The image features a man with short hair, wearing a suit and tie, and giving a speech. He appears to be a professional orator
15. The image features a man with a beard, wearing a suit and tie, and giving a speech.
16. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be a public speaker
17. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be a public speaker
18. The image features a man wearing a suit and tie, standing in front of a blue background. He appears to be a professional or businessman
19. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be a professional speaker or an announcer
20. The image features a man wearing a suit and tie, standing in front of a microphone. He appears to be a professional speaker or an announcer
21. The image features a man wearing a suit and tie, standing in front of a building. He appears to be a businessman or a professional
22. The image features a man wearing a suit and tie, standing in front of a building. He appears to be a businessman or a professional
23. The image features a man standing in front of a building, possibly a church. He is wearing a suit and tie
24. The image features a man standing in front of a building, possibly a church. He is wearing a suit and tie
25. The image features a man standing in a field, possibly a farmer or a worker. He is wearing a hat
26. The image features a man standing in a field, possibly a farmer or a worker in the field. He is wearing a hat
27. The image features a man standing in a field, possibly a farmer or a worker in the field. He is wearing a hat
28. The image features a man standing in a field, possibly a farmer or a worker in the field. He is wearing a hat
29. The image features a man standing in a field, possibly a farmer or a worker in the field. He is wearing a hat
30. The image features a man named "Mike" who is wearing a suit and tie. He appears to be a businessman
31. The image features a man named "Mike" who is wearing a suit and tie. He appears to be a businessman
32. The image features a man named "Mike" who is wearing a suit and tie. He appears to be a businessman

D Resources

All the experiments were conducted using a single A100 80GB or H100 80GB GPU.

E Ethics of Data Collection

Note that the data was collected with only permissive licenses. Moreover, we have an approval from our institution IRB for "Checking and understanding the performance of deep neural networks".