

# Losing Visual Needles in Image Haystacks: Vision Language Models are Easily Distracted in Short and Long Contexts

Aditya Sharma\*<sup>G</sup> Michael Saxon\* William Yang Wang  
University of California, Santa Barbara  
[LoCoVQA.github.io](https://github.com/LoCoVQA)

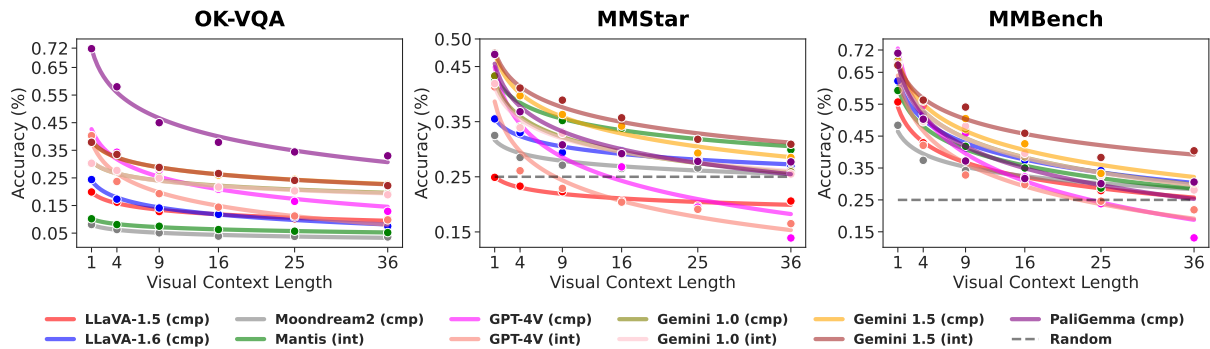


Figure 1: The impact of visual context on vision-language models (VLMs) in our modified, multi-image versions of the OK-VQA, MMStar, and MMBench evaluation benchmarks. *Distractor images* around the target image increase the *visual context length* needed to answer the questions. VLM performance exhibits a logarithmic decay against distractor count, evident in both single composed (**cmp**) and multiple interleaved (**int**) input image configurations.

## Abstract

We present LOCoVQA, a dynamic benchmark generator for evaluating long-context extractive reasoning in vision language models (VLMs). LOCoVQA augments test examples for mathematical reasoning, VQA, and character recognition tasks with increasingly long *visual contexts* composed of both in-distribution and out-of-distribution *distractor images*.

Across these tasks, a diverse set of VLMs rapidly lose performance as the visual context length grows, often exhibiting a striking logarithmic decay trend. This test assesses how well VLMs can ignore irrelevant information when answering queries—a task that is quite easy for language models (LMs) in the text domain—demonstrating that current state-of-the-art VLMs lack this essential capability for many long-context applications.

## 1 Introduction

Long-context vision-language models (VLMs) which accept multiple image inputs are pushing the boundaries of multimodal reasoning applications, but evaluation methods have not kept up. As open-weight long-context language models (LMs) have been around for the last few years (Dong et al.,

2023), researchers have developed robust evaluations to compare their capabilities against those of proprietary LMs (Reid et al., 2024; Jiang et al., 2024a). However, until very recently (Jiang et al., 2024b), the small set of long-context VLMs supporting multi-image input have been inaccessible to the public (OpenAI et al., 2023), so evaluation practices for long-context VLMs (Hong et al., 2023; Li et al., 2023b; Dai et al., 2024) are in their infancy (Zhang et al., 2024). Most reasoning-based VLM evals have analogues in text-domain LM tasks. For example, image captioning (Li et al., 2022; Wang et al., 2022; Nguyen et al., 2022) can be likened to text summarization, and visual question answering (VQA) (Chen et al., 2022; Wang et al., 2023; Xu et al., 2023a) to textual question answering (QA).

We aim to bridge gaps in multi-image VLM evaluation by drawing analogies to established long-context LM tasks. Long-document QA (Fan et al., 2019) is an easy-to-evaluate test that directly showcases model performance on a useful application, and hints at its potential in retrieval-augmented generation (RAG) more broadly. Central to these evaluations is the model’s ability **to recognize which elements in a lengthy input are necessary for answering a query and which are not**, a skill we dub *extractive reasoning*. Do long-context VLMs also exhibit these extractive reasoning capabilities?

\*Equal contribution. <sup>G</sup> Now at Google.

Image ( $X$ )	Question ( $Q$ )	Answer ( $A$ )
Randomly chosen from dataset $\mathcal{D}$	What cartoon character is on this cake?	Winnie the Pooh



Figure 2: Example of **Image ( $X$ )** corresponding to question-answer pair ( $Q$ ,  $A$ ) under increasing visual context lengths in the *composed* and *interleaved* settings. The green box for illustration only, not included in model inputs.

While current VLM benchmarks align with long-document summarization tasks, such as whole-video question answering (Huang et al., 2021) or summarization (Dilawari and Khan, 2019), no existing benchmarks effectively capture a VLM’s ability to **filter out irrelevant images in a long context to reason or answer a query**—essentially, perform extractive reasoning over images.

Measuring extractive reasoning over image sequences is crucial. Just as extractive textual reasoning facilitates text-domain RAG, multi-modal RAG demands that VLMs efficiently reason over and extract information from documents featuring multiple images. Similarly, video QA necessitates that models focus solely on frames containing relevant information, much like long-document QA.

We introduce LOCoVQA, a benchmark generator for **Long Context Visual Question Answering (VQA) with distractors**. LOCoVQA enables the creation of long-context image understanding evaluations from any image comprehension dataset by presenting a sequence with the question-relevant image alongside a configurable set of visual distractors. This allows us to accurately assess how effectively VLMs extract **only the pertinent information to a query within a cluttered context**.

We find that even top proprietary VLMs struggle with this capability, even over short visual contexts, likely due to fundamental deficiencies in their training objectives. By unveiling this LOCoVQA identifies a crucial area for performance improvement in future VLMs. In summary, we:

- Evaluate long-context extractive reasoning across a wide array of VLMs using LOCoVQA.
- Find that all existing VLMs exhibit significant fragility with increasing distractor context sizes.

- Use LOCoVQA to create a *visual needle in a haystack* test, revealing substantial positional biases in SOTA VLMs in extractive reasoning.

## 2 Related Work

Text-based long-context tasks such as long-document question-answering (QA) (Chen et al., 2023), summarization (Phang et al., 2022), and retrieval-augmented generation (RAG) (Lewis et al., 2020; Xu et al., 2023b) offer analogues to long-context reasoning in VLMs. Many of these tasks (e.g., QA, RAG) require natural language extractive reasoning, as discussed above.

Few VLM long-context evals measure extractive reasoning. Existing VQA benchmarks involving *distractor information* (thereby extraction) do not focus on long contexts, and long-context VQA benchmarks do not involve distractors. MultipanelVQA (Fan et al., 2024) includes distractor information, but only within single input images. Some VLM evaluations focused on hallucinations indirectly capture a notion of distraction (Ghosh et al., 2024; Favero et al., 2024) but do not explicitly measure it alone. MILEBench (Song et al., 2024) evaluates VQA in long contexts, but only on non-distractor tasks such as video summarization or difference detection, where all inputs are relevant.

*Needle-in-a-haystack* evaluation tasks (asking models to recover hidden passphrases at various positions) do require long-context extraction, but reasoning is not involved. Gemini 1.5 (Reid et al., 2024) extended this task to video comprehension by hiding a simple passphrase in many positions within a video. Wang et al. (2024) present a static benchmark concurrently to ours centered on multi-modal needle-in-a-haystack.

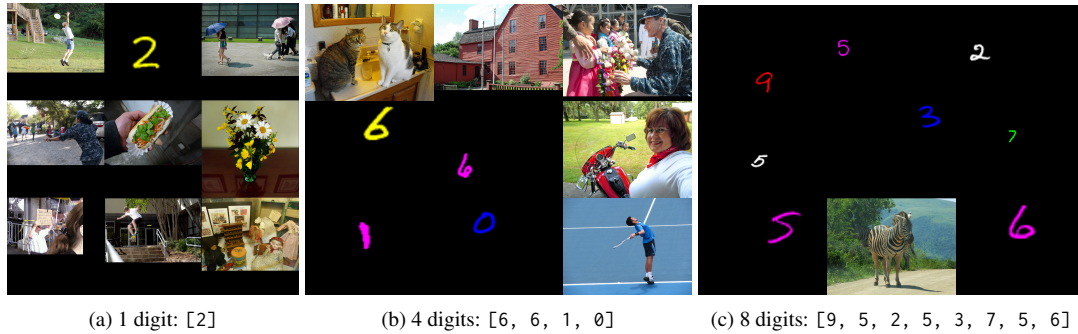


Figure 3: Each subfigure represents a variable number of hidden MNIST digits in a 9 image composed context.

### 3 LOCOVQA Generation Method

Samples synthesized through LOCOVQA contain one or more *content* images  $X$  corresponding to a question and answer pair  $(Q, A)$ . Content images can be sampled from various image comprehension benchmark, such as OK-VQA (Marino et al., 2019), MMStar (Chen et al., 2024), MNIST (LeCun et al., 1998), and others. Alongside the content image(s), each sample includes up to 35 *distractor images*, which are either sampled in-distribution from the content image set (ensuring no content image collisions to prevent ambiguity, as described in §3.1.3) or out-of-distribution from other image sets such as MS COCO (Lin et al., 2014).

Samples can be arranged as **interleaved** image sequences for VLMs that accept multi-image inputs or as **composite** images arranged in a grid, as depicted in Figure 2. For all models capable of ingesting interleaved sequences, we evaluate both interleaved and composite examples.

*Visual context* refers to the visual elements within an image or sequence of images that are relevant to answering a question. This includes the *content* image and distractor images. The challenge is to identify and focus on the pertinent details while ignoring irrelevant information. We can sample images both in-distribution to create challenging visual contexts or out-of-distribution to isolate extraction capabilities in the simplest setting.

As a dynamic generator LOCOVQA can convert any vision-language dataset into a long-context benchmark (Saxon et al., 2024a). We describe the three datasets we used below.

#### 3.1 Single-image Reasoning Tasks

First, we discuss the visual reasoning benchmarks which we expand into long-context samples containing one content image per sample. For most question answering and visual reasoning tasks, this

is the only LOCOVQA expansion that makes sense: few VQA samples can be combined such that a plausible new QA pair requiring information from multiple images. Since most interleaved models we evaluate support 36 or fewer images as sequential inputs without modification, we do not evaluate any models with context lengths beyond 36. However, LOCOVQA scales to any size. For the single-context reasoning tasks, we exclusively sample the distractors in-distribution.

##### 3.1.1 OK-VQA

OK-VQA (Marino et al., 2019) is a single-image visual question answering dataset containing 5,072 question-answer-image triples. It requires external knowledge to reason beyond the image. LOCOVQA generates in-distribution long-context OK-VQA samples, ensuring that no content images have concept collisions that may complicate evaluation. For instance, the question about a character on top of a cake, as shown in Figure 2, is sampled from OK-VQA.

Since OK-VQA is an open-domain, free-form answer dataset, we score the samples using three metrics: exact match (full score if the model’s response contains any ground truth answer as a substring), and continuous text generation metrics BERTScore (Zhang et al., 2020) and ROUGE<sub>L</sub> (Lin, 2004) between candidates and references.

##### 3.1.2 MMStar

MMStar (Chen et al., 2024) is a multi-domain visual question answering dataset combining examples from various existing datasets: 424 questions from SEEDBench (Li et al., 2023a), 366 questions from MMBench (Liu et al., 2023b), 100 questions from MMMU (Yue et al., 2023), 345 questions from MathVista (Lu et al., 2023), 69 examples from ScienceQA (Lu et al., 2022), and 196 examples from AI2D (Kembhavi et al., 2016). MMStar con-

tains 1,500 high-quality multiple-choice questions that *require visual information from the images* to answer, a filtering step not initially performed on the source datasets. For example, over 50% of ScienceQA questions can be solved by a text-only LLM (Chen et al., 2024). Similar to OK-VQA, we generate LOCoVQA samples for MMStar using the collision filtering technique to produce pseudo-documents composed of multiple example images.

As a multiple choice dataset, scoring MMStar is more straightforward. Full details on how we faithfully extract multiple choice answers from the models is provided in Appendix D.

### 3.1.3 Filtering Collisions in LOCoVQA

To address the problem of *content-distractor collisions*—where multiple similar in-distribution images in the visual context make the QA pair ambiguous—we implement a robust LM-based filtering method. For each visual context image, we prompt GPT-4 to list the top five entities; if there is overlap, we consider the question potentially ambiguous. Detailed implementations and examples of our filtering method are provided in Appendix C. To validate this approach, we manually assessed a subset of LOCoVQA generations and found it to be consistent, with no such collisions.

## 3.2 Multi-image Reasoning Tasks

In §3.1, we explored tests designed to evaluate whether VLMs can extract a single relevant image from a sequence to answer a query, thereby probing their long-context reasoning capabilities. Extending this to test how well VLMs can extract information from a multi-image sequence is a natural progression. However, VQA examples cannot be easily combined in a way that requires multiple images to answer a single question. Therefore, we turn to constructing “sequential VQA” sets using synthetic tasks. Optical character recognition (OCR) is a straightforward task to convert into multi-image question answering by including multiple OCR examples as interleaved images and asking the VLM to list all the text.

### 3.2.1 MNIST

We use MNIST (LeCun et al., 1998) as it is a canonical dataset for OCR. For a desired visual context length, we sample between 1 and 8 randomly-colored digits from the MNIST training set of around 60K images, resizing them to between 1/6 and 1/2 of the maximum height of other context

Vision-Language Model	Context	Base LM	↓sample
<i>Single-Image Input VLMs</i>			
Moondream2-1.6B	2K	Phi-1.5	✓
LLaVA-1.5-7B	4K	Vicuna	✓
LLaVA-1.6-7B	8K	Mistral	✓
PaliGemma-3B	8K	Gemma	✓
<i>Multi-Image Input VLMs</i>			
Mantis-Bakllava-7B	8K	Mistral	✓
Mantis-Idefics2-8B	8K	Mistral	✓
Gemini 1.0 Pro Vision	32K	Gemini	-
GPT-4 Vision	128K	GPT-4	-
Gemini 1.5 Flash	1M	Gemini	-

Table 1: Overview of ■ open-source and ■ proprietary vision-language models (VLMs). *Downsampling* models (↓sample) rescale input images in the single-image setting rather than take as many ViT inputs as necessary.

images. The remaining distractor images are randomly sampled from a subset of 5K high-quality MS COCO (Lin et al., 2014) validation images. The VLM is then prompted to *list all handwritten digits present in the sequence*.

By varying the number of digits in the sequence, we can dynamically adjust the difficulty level of the multi-image distractor OCR task. Figure 3 illustrates examples with 1, 4, and 8 digits in a 9-image context. An output is considered correct only if the stored string of generated digits exactly matches the ground truth, with no partial credit.

## 4 Experiments

We evaluate the performance of nine current vision-language models on our LOCoVQA-generated benchmarks. The open-weight models tested are Moondream2 (Moondream, 2024), LLaVA-1.5 (Liu et al., 2023a), LLaVA-1.6 (Liu et al., 2024), PaliGemma-3b (Google, 2024), and two Mantis variants (Jiang et al., 2024b). Additional details are provided in Appendix A. Both Mantis variants are VLMs further tuned on interleaved multi-image instruction following. The three proprietary models we evaluate are GPT-4V (Achiam et al., 2023), Gemini 1.0 Pro Vision (Team et al., 2023), and Gemini 1.5 Flash (Reid et al., 2024). Table 1 shows cases all the models along with the model context sizes and image downsampling.

**Interleaved image support.** All closed-source models and Mantis support multi-image interleaved inputs, while the others only support single images. For multi-image models, we evaluate both the **composed (cmp)** & **interleaved (int)** settings. For single-image models, we only test the **composed (cmp)** setting.

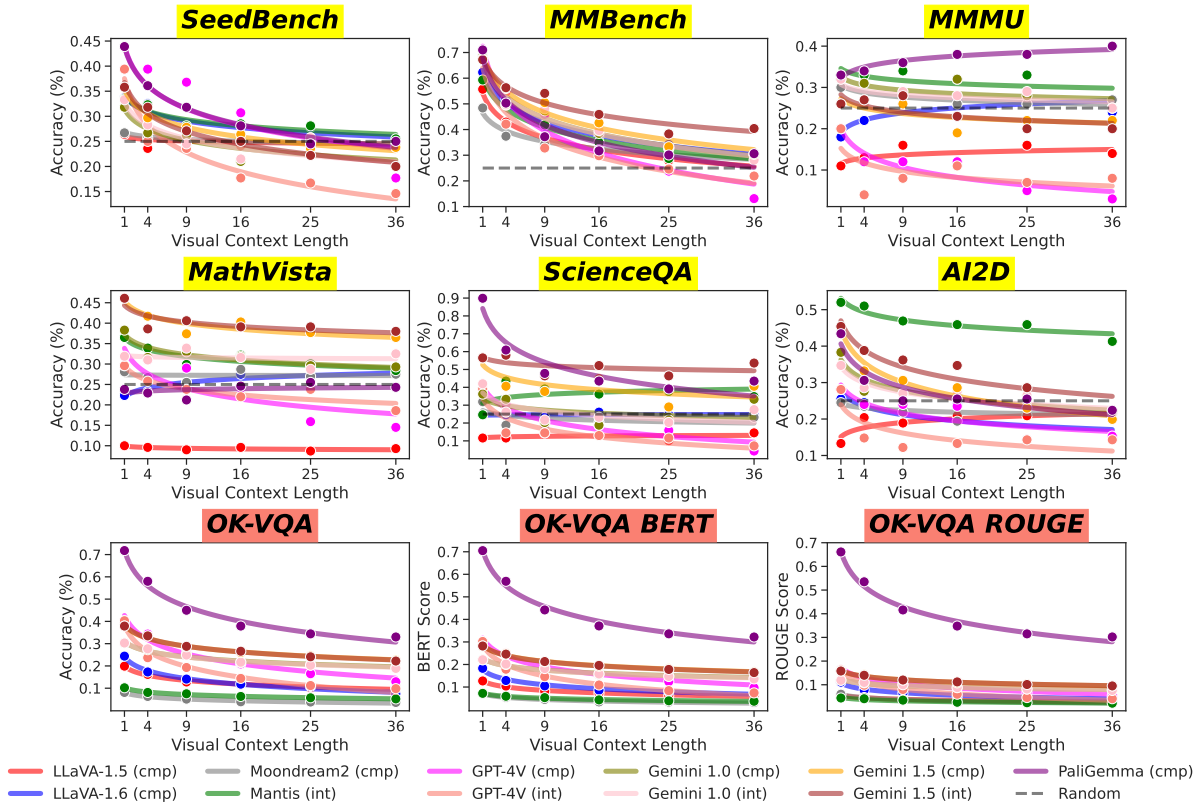


Figure 4: VLM Performance on MMStar and OK-VQA. Note the clearly declining logarithmic fit trends for many of the models. The (model, task) pairs for which these trends do not hold by and large are *below the random baseline*.

**Downsampling images.** Some models accept images of arbitrary resolution, while others automatically downscale inputs. This difference is crucial, especially in the **cmp** setting, as increasing the visual context can lead to information loss. For the downsampling models that support both **cmp** and **int** settings, we assess both. Any performance differences between these settings would highlight the impact of downsampling on the **cmp** setting.

## 5 Results

Figure 4 illustrates how model performance (across 10 experiments, in both composed and interleaved settings) changes with increasing visual context lengths on single-image LOCoVQA tasks. The first two rows display results from the MMStar dataset, which contains filtered-image-necessary examples from six other datasets (Chen et al., 2024), with titles in **yellow** and random guess thresholds indicated by dotted black lines. The bottom row presents OK-VQA scores using three scoring metrics, with titles in **salmon**.

Across all models on OK-VQA and the majority on MMStar subsets for SeedBench, MMBench, ScienceQA, and AI2D, we observe *striking logarithmic decay trends* in model performance

with increasing visual context length. Correlation coefficients,  $r^2$ , and  $p$ -values for each trendline are reported in Table 3, Appendix D. In general, the closed-source models outperform their open-weight counterparts, especially on multiple-choice tasks. Among the open-weight models, PaliGemma performs the best, likely due to its substantially more extensive training on vision-language tasks compared to its counterparts like LLaVA or Mantis.

On several multiple-choice tasks, Mantis (the interleaved open-weight model) outperforms the other open-weight models that rely solely on composite inputs. This is likely due to the image subsampling required by several of the other open-weight models negatively impacting performance when using composite inputs. However, on the OK-VQA task, Mantis and Moondream are the least performant, even at low context lengths. These models were likely not trained on visual question-answering tasks to the same extent as LLaVA variants during their instruction tuning steps.

The most noteworthy point is this: **the logarithmic decay trend holds equally well in interleaved and composed settings**. This indicates that the performance-visual context length trend is fundamental and **cannot be attributed solely**




Figure 5: Radar plots of VLM performance across 8 multimodal benchmarks with varied visual context lengths.

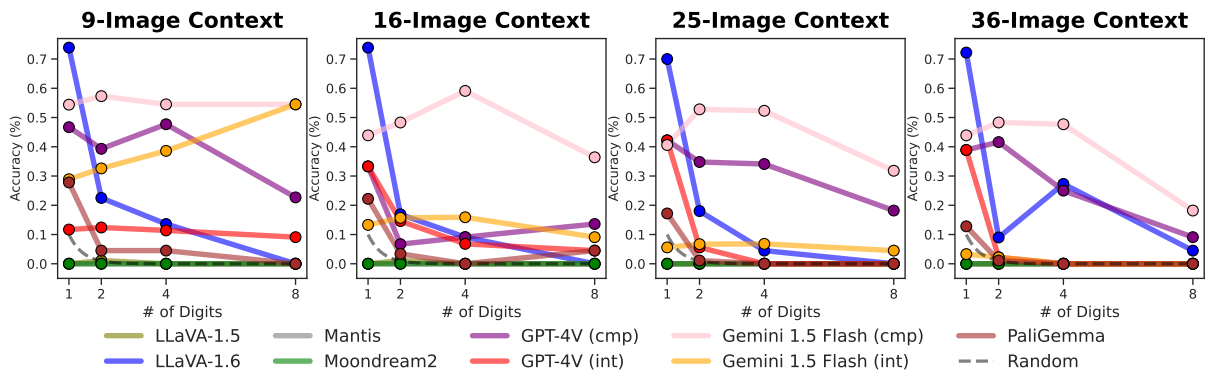


Figure 6: VLM Performance on the MNIST-Digits transcription task as a function of # of digits to transcribe. *These plots have a different x-axis than the plots in Figure 1 and Figure 4: rather than the relationship between context size and performance, we are assessing the relationship between "task difficulty" and performance, at four context sizes.*

**to downsampling effects in the composite setting.** Furthermore, for the models tested in both conditions (GPT-4V and Gemini), the same trends are observed in both interleaved and composite settings. Sub-chance performance on the multiple-choice question-answering datasets, particularly for the closed-weight models, occur due to high refusal rates. This may illustrate how “alignment hampering” can hinder performance.

**Taskwise Performance by context length.** Figure 5 illustrates the models’ performance at context lengths of 1, 9, and 25 for each task, to compare the relative advantages different models have on various tasks. For example, PaliGemma excels on OK-VQA compared to the other models, while Mantis performs well on AI2D. These differences are likely due to variations in training tasks.

**Performance on Multi-image Tasks.** Figure 6 presents model performance on the MNIST transcription task by number of digits to transcribe.

**Note the x-axis in this figure is changed; it is checking for a fixed context length the relationship to number of digits.** While there is a trend of decreasing performance with an increasing number of images, it does not follow a simple pattern like the context-length trend. For example, Gemini 1.5 experiences minimal performance degradation even as the target digit length increases.

Characterizing difficulty as a function of digit length is complex—as digit counts increase in a fixed context window, *the the output label search space grows, while the ratio of relevant images to irrelevant images increases.* This may explain why some models have *consistent* or even *increasing* performance as # digits increases. Analyzing why different models handle these axes of difficulty differently is an interesting future direction.

However, akin to the single-image tasks, increasing the overall *visual context length* (seen in same-color, x value points across plots), makes the task more difficult, albeit without as clear a correlation.

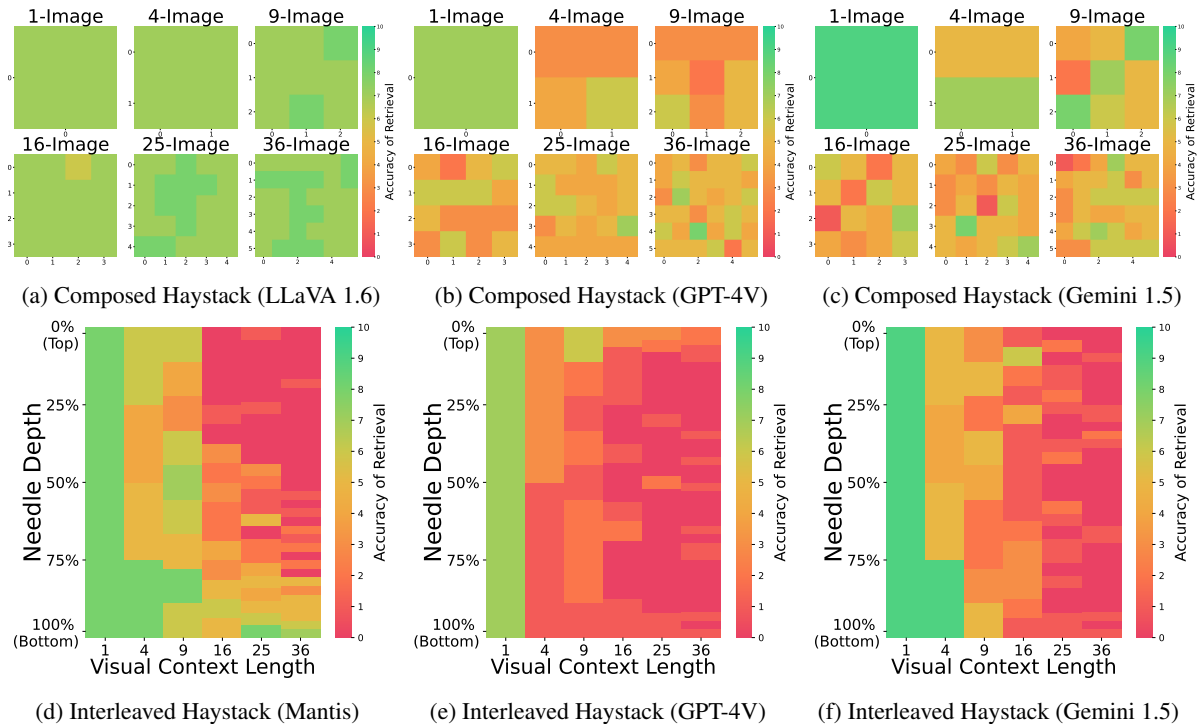


Figure 7: Evaluation of the Visual Needle in a Haystack task using GPT-4V, best-performing VLM, conducted under both composed and interleaved haystack settings. Retrieval accuracy measures the frequency of correct answers produced by the VLM, in this case identifying the MNIST digit. In the composed setting, retrieval accuracy is measured by placing the needle at each cell. In the interleaved setting, needle depth signifies the position within the image sequence, with 0% representing the first image and 100% representing the last. Our evaluation highlights a consistent decline in retrieval accuracy as the visual context length increases.

## 6 Ablation: Needle in a Haystack

As performance within context windows decreases, a natural question arises: **Are performance failures equally distributed across the visual context range?** To investigate this, we adapt our MNIST-based visual context OCR task into a **visual needle in a haystack** task. Needle in a haystack tests are a common minimal test of long-context capabilities in LMs (Kamradt, 2023), involving hiding a passphrase, such as the word “needle,” in various positions within a long document and asking the model to retrieve it. To adapt this concept, we modified our single-digit MNIST task by sampling a set of 10 colored MNIST digits (one for each number, e.g. blue 3, green 7, etc.) and hiding them in each possible position within both interleaved and composed visual context sequences.

By assessing a model’s single-digit recovery rate at each position, we can identify any systematic bias in extraction capabilities by position. Figure 7 presents the performance of a subset of the tested VLMs on the composed and interleaved MNIST visual needle in a haystack tasks. We test GPT-4V and Gemini 1.5 Flash in both settings, and treat

LLaVA 1.6 and Mantis as analogous methods between the interleaved and composed settings.

Across the three composed tests (subfigures a, b, and c), we do not find a systematic bias toward any particular position on the single-digit recognition task. Surprisingly, LLaVA 1.6 performs the best, probably due to more MNIST-like OCR data in its training mix compared to the others.

However, we find strong systematic biases with respect to position in the interleaved setting for all models (subfigures d, e, and f). Mantis shows a preference for late positions in sequences of any length, while GPT-4V and Gemini exhibit weak biases towards early positions. Given that the biases are evident even in the relatively simple MNIST-based OCR task, it is likely that this effect plays a significant role in the performance penalty these models experience with longer contexts.

These results are supported by similar findings in from the image-needle test from Wang et al. (2024). They also found that—contrary to Google’s claims to significant needle-in-a-haystack performance for Gemini Reid et al. (2024)—for non-trivial hidden visual information tasks current SOTA VLMs are woefully non-performant.

## 7 Discussion

Why do we observe such striking logarithmic decay in performance with increasing visual context length? The fact that the trend is consistent in both interleaved and composite settings that down-sampling effects in the latter are not the primary cause of performance decay. There may be an information-theoretic interpretation of this behavior: an increasing signal-to-noise ratio. As the visual context length increases, the amount of signal (relevant information) remains constant, while the amount of noise (distractor information) increases.

However, this signal-to-noise ratio problem also exists in text-only tasks, yet long-context LLMs perform well on needle in a haystack and long-context QA and summarization tasks, underlying their performance in retrieval-augmented generation. When it comes to visual contextual tasks that involve distractor information, VLMs struggle to perform even on the easiest long-context tasks, such as transcribing MNIST characters. They are even less capable on more complex, real-world tasks like distractor VQA. The core issue appears to be that both short-context and long-context VLMs are not trained on tasks requiring attention over multiple context images to retrieve information. In contrast, for long passages of text, information from throughout the passage remains relevant—referring back to a specific sentence early in a document is intrinsic to the LM training objective.

However, the same may not necessarily be true for the sequential image language modeling task. In the existing interleaved vision-language training corpora, images are likely to be followed immediately by their relevant text. As a result, attending to much earlier images in the documents may not be crucial for achieving low LM loss. Additionally, none of the tested open VLMs are trained on image *generation* tasks conditioned on text while updating the core transformer weights. Overlooking this objective may prevent VLMs from robustly modeling the relationship between images and text, which could be crucial for performing well on these tasks.

**Interesting examples of poor performance** Our analysis unveiled a few particularly surprising results. For example, contrary to our expectation that models capable of handling interleaved input would perform better with it than with composite input, we found that GPT-4V actually favors composite input on many subtasks. Additionally, to our surprise LLaVA 1.6 performed much better

than GPT-4V on the composite haystack task at all sizes. This outcome was driven by multiple factors: GPT-4V tends to refuse some tasks rather than guessing, while LLaVA always provides a guess. Furthermore, GPT-4V often “hallucinates” multiple numbers instead of just one.

**Possibility of memorization** Some of the surprising performance results are driven by the training mix. For example, Mantis’s significant lead on AI2D, and LLaVA’s strong performance on the MNIST task may result from having more relevant data or those specific training sets included in their data. However, even considering that some models were trained on these tasks, the pronounced drops in generalization performance as the context length increases are even more striking. This illustrates that current VLMs fundamentally struggle to attend image sequences as well as they do with text.

**Upper bound for video QA** A more construct-valid test for real-world visual extractive reasoning would be QA over a long video where only a sub-element is required. E.g., putting the entire film *Star Wars* as input for QA pair (Q: “which character shoots a green alien in the Mos Eisley cantina?”, A: Han Solo). Information from a single scene must be extracted to perform this task, thereby requiring extractive reasoning. However, it is plausible that sequences of related images, unlike sequences of unrelated images, are more in-distribution for long-context VLMs and extractive tasks involving them could be easier for them to solve. By providing completely unrelated images, our task may be harder than video-based VQA, and may represent an upper-bound for visual extractive reasoning.

## 8 Conclusion

Vision-language models struggle with long-context visual extractive reasoning. Many models fail to extract necessary information from even short contexts of 36, 25, or 16 images, resulting in near- or sub-baseline performance across tasks. LO-CoVQA presents a simple benchmark-generating process applicable to any VQA or retrieval evaluation dataset, making it easy to assess extractive reasoning in VLMs. Our findings suggest that training tasks requiring attention across multiple context images to extract information—rather than simple single-image tasks—should be included in VLM training. By measuring this capacity it offers an appealing direction for future work.



## Limitations

Although LOCOVQA is a generalized process applicable to any VLM benchmark, we only evaluated it on three tasks. While the strong logarithmic decay trends between visual context length and performance observed across all three are compelling, expansion of LOCOVQA to additional tasks would render the findings clearer. While LOCOVQA samples distractor images from the same datasets for open-domain and multiple-choice VQA, and our process appears to accurately filter out collisions, a small number of collisions still occur—it is inherently difficult to ensure no failures in an automated generating process (Saxon et al., 2024b). This may lead to a natural ceiling on VLM performance at each context length.

Other important long-context capabilities likely exist that are not captured by LOCOVQA or prior work such as MILEBench (Song et al., 2024). Augmenting these evaluations with tests that capture additional orthogonal VLM long-context capabilities is an important direction for future work.

## Acknowledgements & Contributions

Thanks to Yujie Lu, Luke Yoffe, Tejas Srinivasan, Weixi Feng, and Deepak Nathani for discussion and comments. This work was supported in part by the National Science Foundation Graduate Research Fellowship under Grant No. 1650114, and CAREER Award under Grant No. 2048122. AS wrote code, ran experiments, and produced figures. MS was the primary writer and editor of the manuscript. Both designed experiments and conducted analysis.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *ArXiv preprint*, abs/2303.08774.
- Hung-Ting Chen, Fangyuan Xu, Shane A. Arora, and Eunsol Choi. 2023. [Understanding retrieval augmentation for long-form question answering](#). *ArXiv preprint*, abs/2310.12150.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. 2024. [Are we on the right way for evaluating large vision-language models?](#) *ArXiv preprint*, abs/2403.20330.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. 2022. [Pali: A jointly-scaled multilingual language-image model](#). *ArXiv preprint*, abs/2209.06794.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. 2024. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#). *Advances in Neural Information Processing Systems*, 36.
- Aniqa Dilawari and Muhammad Usman Ghani Khan. 2019. [Asovs: abstractive summarization of video sequences](#). *IEEE Access*, 7:29253–29263.
- Zican Dong, Tianyi Tang, Lunyi Li, and Wayne Xin Zhao. 2023. [A survey on long text modeling with transformers](#). *ArXiv preprint*, abs/2302.14502.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Yue Fan, Jing Gu, Kaiwen Zhou, Qianqi Yan, Shan Jiang, Ching-Chen Kuo, Xinze Guan, and Xin Eric Wang. 2024. [Muffin or chihuahua? challenging large vision-language models with multipanel vqa](#). *ArXiv preprint*, abs/2401.15847.
- Alessandro Favero, Luca Zancato, Matthew Trager, Siddharth Choudhary, Pramuditha Perera, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. 2024. [Multi-modal hallucination control by visual information grounding](#). *ArXiv preprint*, abs/2403.14003.
- Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Utkarsh Tyagi, Oriol Nieto, Zeyu Jin, and Dinesh Manocha. 2024. [Vdgd: Mitigating lvlm hallucinations in cognitive prompts by bridging the visual perception gap](#). *ArXiv preprint*, abs/2405.15683.
- Google. 2024. [Paligemma](#). <https://ai.google.dev/gemma/docs/paligemma>.
- Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. 2023. [Cogagent: A visual language model for gui agents](#). *ArXiv preprint*, abs/2312.08914.
- Luyang Huang, Shuyang Cao, Nikolaus Parulian, Heng Ji, and Lu Wang. 2021. [Efficient attentions for long document summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1419–1436, Online. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego

- de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. **Mistral 7b**. *ArXiv preprint*, abs/2310.06825.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024a. **Mixtral of experts**. *ArXiv preprint*, abs/2401.04088.
- Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max Ku, Qian Liu, and Wenhu Chen. 2024b. **Mantis: Interleaved multi-image instruction tuning**. *ArXiv preprint*, abs/2405.01483.
- Gregory Kamradt. 2023. Needle in a haystack - pressure testing llms. [https://github.com/gkamradt/LLMTest\\_](https://github.com/gkamradt/LLMTest_).
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. A diagram is worth a dozen images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. **What matters when building vision-language models?** *ArXiv preprint*, abs/2405.02246.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. **Retrieval-augmented generation for knowledge-intensive NLP tasks**. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023a. **Seed-bench: Benchmarking multimodal llms with generative comprehension**. *ArXiv preprint*, abs/2307.16125.
- Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo Ye, He Chen, Guohai Xu, Zheng Cao, Ji Zhang, Songfang Huang, Fei Huang, Jingren Zhou, and Luo Si. 2022. **mPLUG: Effective and efficient vision-language learning by cross-modal skip-connections**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7241–7259, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. **Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models**. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023c. **Textbooks are all you need ii: phi-1.5 technical report**. *ArXiv preprint*, abs/2309.05463.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. **Improved baselines with visual instruction tuning**.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. **Llava-next: Improved reasoning, ocr, and world knowledge**.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023b. **Mmbench: Is your multi-modal model an all-around player?** *ArXiv preprint*, abs/2307.06281.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2023. **Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts**. *ArXiv preprint*, abs/2310.02255.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. **OK-VQA: A visual question answering benchmark requiring external knowledge**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3195–3204. Computer Vision Foundation / IEEE.
- Moondream. 2024. tiny vision language model. <https://github.com/vikhyat/moondream>.
- Van-Quang Nguyen, Masanori Suganuma, and Takayuki Okatani. 2022. **Grit: Faster and better image captioning transformer using dual visual features**. In

*European Conference on Computer Vision*, pages 167–184. Springer.

- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Alvenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fullford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reichihiro Nakano, Rajeesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. [Gpt-4 technical report](#). *ArXiv preprint*, abs/2303.08774.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with gpt-4](#). *ArXiv preprint*, abs/2304.03277.
- Jason Phang, Yao Zhao, and Peter J. Liu. 2022. [Investigating efficiently extending transformers for long input summarization](#). *ArXiv preprint*, abs/2208.04347.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *ArXiv preprint*, abs/2403.05530.
- Michael Saxon, Ari Holtzman, Peter West, William Yang Wang, and Naomi Saphra. 2024a. [Benchmarks as microscopes: A call for model metrology](#). In *First Conference on Language Modeling*.
- Michael Saxon, Yiran Luo, Sharon Levy, Chitta Baral, Yezhou Yang, and William Yang Wang. 2024b. [Lost in translation? translation errors and challenges for fair assessment of text-to-image models on multilingual concepts](#). *ArXiv preprint*, abs/2403.11092.
- SkunkworksAI. 2023. [Bakllava](#). <https://huggingface.co/llava-hf/bakllava-v1-hf>.
- Dingjie Song, Shunian Chen, Guiming Hardy Chen, Fei Yu, Xiang Wan, and Benyou Wang. 2024. [Milebench: Benchmarking mllms in long context](#). *ArXiv preprint*, abs/2404.18532.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu,

- Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. [Gemini: a family of highly capable multimodal models](#). *ArXiv preprint*, abs/2312.11805.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. [Gemma: Open models based on gemini research and technology](#). *ArXiv preprint*, abs/2403.08295.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. 2022. [Git: A generative image-to-text transformer for vision and language](#). *ArXiv preprint*, abs/2205.14100.
- Weiyun Wang, Shuibo Zhang, Yiming Ren, Yuchen Duan, Tiantong Li, Shuo Liu, Mengkang Hu, Zhe Chen, Kaipeng Zhang, Lewei Lu, et al. 2024. [Needle in a multimodal haystack](#). *ArXiv preprint*, abs/2406.07230.
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. 2023. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19175–19186.
- Haiyang Xu, Qinghao Ye, Ming Yan, Yaya Shi, Jiabo Ye, Yuanhong Xu, Chenliang Li, Bin Bi, Qi Qian, Wei Wang, et al. 2023a. [mplug-2: A modularized multi-modal foundation model across text, image and video](#). In *International Conference on Machine Learning*, pages 38728–38748. PMLR.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence C. McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2023b. [Retrieval meets long context large language models](#). *ArXiv preprint*, abs/2310.03025.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. 2023. [Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi](#). *ArXiv preprint*, abs/2311.16502.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Gengyuan Zhang, Yurui Zhang, Kerui Zhang, and Volker Tresp. 2024. Can vision-language models be a good guesser? exploring vlms for times and location reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 636–645.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

## A Evaluated Models

We evaluate: Moondream2-1.6b (Moondream, 2024) based on Phi-1.5 (Li et al., 2023c), LLaVA-1.5 (Liu et al., 2023a) with Vicuna-7b (Peng et al., 2023) as the LLM backbone, LLaVA-1.6 (LLaVA-Next) (Liu et al., 2024), an improvement over LLaVA-1.5 with higher image resolution and better visual reasoning, uses Mistral-7b (Jiang et al., 2023), PaliGemma-3b (Google, 2024), based on open components from the SigLip (Zhai et al., 2023) image encoder and the Gemma (Team et al., 2024) language model, Mantis-bakllava-7b (Jiang et al., 2024b) fine-tuned from BakLLaVA (SkunkworksAI, 2023) and derived from LLaVA but using Mistral-7b (Jiang et al., 2023), Mantis-Idetics2-8b (Jiang et al., 2024b), the current state-of-the-art Mantis variant, based on Idetics2-8b (Laurençon et al., 2024).

## B LoCoVQA Release Information

Full source code for generating distract datasets is available at [locovqa.github.io](https://locovqa.github.io). LoCoVQA is released under the Apache v2.0 license.

## C Filtering Collisions in LoCoVQA

Figure 8 shows an example method used to filter for collisions in 4-Image Context Length input from Figure 2b. The LLM Query ( $Q$ ) was used to prompt GPT-4 LLM. The second figure shows an example of a collision occurring when creating a 4-image context. In this case, the two images contain oranges, so we resample one of the images to avoid collision. The OK-VQA questions in this case are “What type of plant do these fruits grow from?” for the content image and “In which US states are these fruits commonly grown?” for the distractor image. Both require visual reasoning step that the fruit described in this image is an orange.

## D Model-level Scoring Details

**OK-VQA.** For the OK-VQA free-form ground truth answers, if any of the ground truth candidate answers is a substring of the model-generated answer, we award full points for the exact-matching setting. The other two metrics we used were

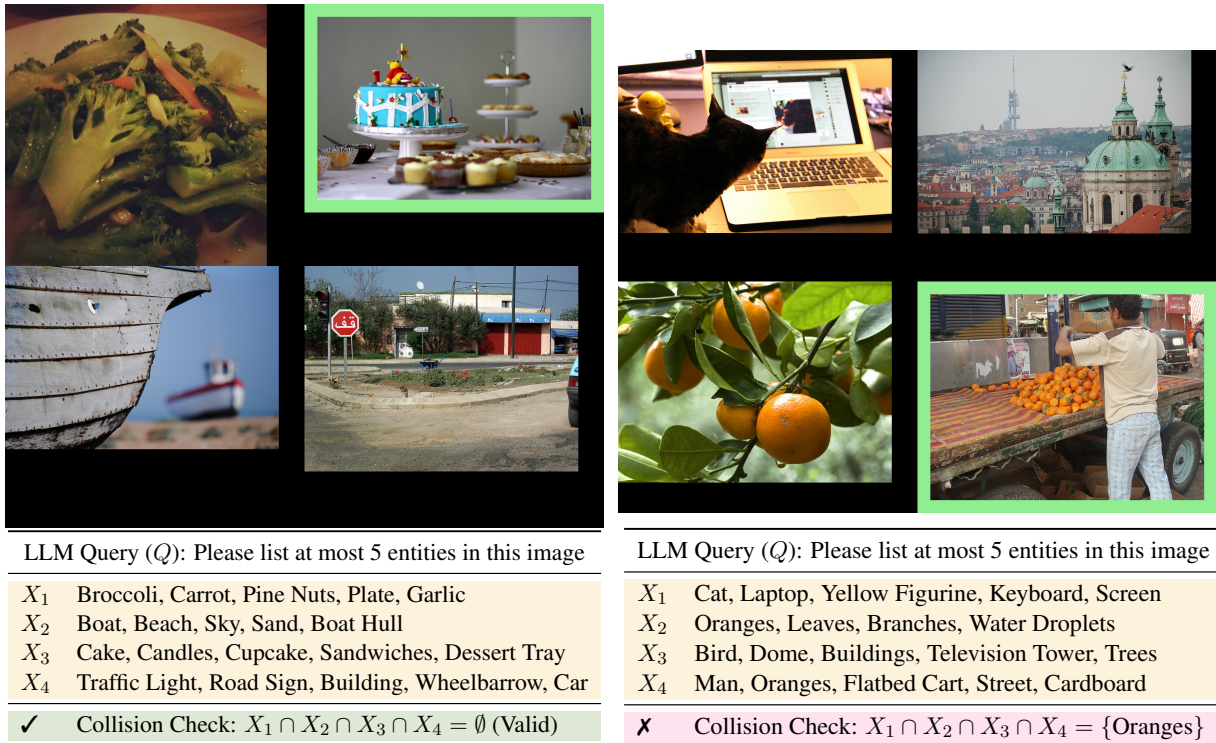


Figure 8: Collision Filtering Method for LOCoVQA by prompting LLM with query  $Q$  to identify entities. Cell with   represents the entities for each image  $X_i$ . If there are no entities in common, there are no collisions so we mark cell with   indicating this is a valid construction of images for LOCoVQA.

BERTScore and ROUGE scores. BERTScore (Zhang et al., 2020) is a robust text comparison metric which matches candidate and reference based on the cosine similarities of the embeddings. Using the Sentence Transformer package all-MiniLM-L6-v2 model, we calculated the maximum BERTScore between all the ground truth answers against the model answer. For the ROUGE-score, we compared the ground truth answer and model answer using the default settings with ROUGE-L, which measure the longest common subsequence at the sentence-level

**MMStar.** For multiple-choice answer form in MMStar, when prompting the VLM, we ask it to produce the answer in the following format: “Please provide the answer in format  $\langle \text{answer} \rangle \langle / \text{answer} \rangle$  where the answer is between the tags. For example if the result is apple, your answer should be  $\langle \text{answer} \rangle \text{apple} \langle / \text{answer} \rangle$ .” This format ensure some grammatical structure and requires the answer to be enclosed within the  $\langle \text{answer} \rangle$  tags. GPT4V, Gemini 1.0, and Gemini 1.5 under both composed and interleaved settings produced answers following this format. However, we did not observe consistent behavior from LLaVA-based and Mantis-based variants. Moondream

and Gemma consistently responded with a single-choice answer, making the evaluation of these two models the easiest. Due to the variance in VLM responses, we adopted a robust evaluation procedure. First, we checked if the answer was between the  $\langle \text{answer} \rangle$  tags and, if so, extracted the MCQ choice directly from the tag. If not, we noted that outputs often followed the format, “Answer: *choice*,” where the choice follows directly after. We also checked edge cases, including instances where the first letter in the string is a multiple-choice followed by a colon, choices provided in parentheses, only the answer text without the corresponding letter, and a single letter provided.

**MNIST.** For MNIST evaluation, we ensure that the model response contains a list of digits separated by commas. If the output is separated by spaces, we parse it into an array to provide a completely fair evaluation. The response must contain exactly the same number of digits as those in the MNIST digits. We sort both the candidate and reference lists and compare them to check for equality.

**API model considerations.** Open-weight models consistently attempt to provide an answer when prompted, even if it means occasionally producing




Figure 9: Plots of non-response rate and non-template responses of closed-source VLMs.




Figure 10: Radar plots of VLM performance across 8 multimodal benchmarks with context lengths  $k = 4, 16, 36$ .

hallucinated responses. In contrast, API-based models like Gemini and GPT-4V sometimes refuse to answer questions under certain circumstances—such as low confidence, potentially unsafe language in the response, or failure to adhere to the response template. We categorize these instances into two primary error types: non-response rate and non-template response rate. Figure 9 illustrates these error types for both Gemini and GPT-4V models on OK-VQA dataset of size 5072.

A non-response case is identified when the model returns an empty string. The Gemini models (1.0 and 1.5) exhibit the highest non-response rates, recording 72 non-responses in visual context with a size of 36, but fewer than 10 for all other context lengths. This pattern may be attributed to the safety mechanisms employed by Gemini or the model’s uncertainty due to the presence of multiple distractors. Meanwhile, the non-response rate for GPT-4V in the interleaved setting shows a distinct increasing trend, possibly because GPT-4V struggles with processing a set of 36 interleaved images. The non-response rate for GPT-4V interleaved is 425 non-response in visual context of 36.

A non-template response occurs when the model fails to enclose its output with `<answer>` tags. This is not penalized in our evaluations—where we still calculate exact match, BERT, ROUGE scores—it is crucial to monitor the trend. As visual context length increases VLMs often struggle to adhere to the prompted response formatting. Among the models tested, Gemini 1.5 demonstrates the best adherence to requested formatting, with only a single error. Gemini 1.0 exhibits slightly more issues, with no more than 100 errors in template adherence. GPT-4V shows the most significant struggle, with up to 600 formatting errors. We classify model refusals, possibly due to low confidence, as failures.

## E Supplementary Results

Figure 10 extends Figure 5 for context lengths  $k = 4, 16, 36$ . Table 2 and Table 3 display the  $r^2$  values for Figure 4 fits, with  $p$ -val asterisks denoting statistical significance. Red highlighting signifies subchance performance (more than half of scores below random choice). Down arrows indicate positive correlation (improved performance with increasing context counters overall trend).

Methods	LLaVA (c)	LLaVA-1.6 (c)	Moondream (c)	PaliGemma(c)	Mantis (i)
MMStar	.924***	.957***	.903***	.947***	.967***
SEEDBench	.119	.825***	.825***	.988***	.885***
MMBench	.959***	.992***	.937***	.946***	.972***
MMMU	.367 <sup>↓</sup> **	.691 <sup>↓</sup> ***	.586**	.923 <sup>↓</sup>	.282***
MathVista	.513*	.907 <sup>↓</sup> ***	.004	.132 <sup>↓</sup>	.848***
ScienceQA	.080 <sup>↓</sup>	.020	.279*	.907***	.209 <sup>↓</sup> *
AI2D	.658 <sup>↓</sup> **	.604**	.375	.882***	.857***
OK-VQA	.985***	.989***	.979***	.968***	.991***
BERT	.986***	.973***	.981***	.968***	.986***
ROUGE	.988***	.973***	.985***	.968***	.918***
Haystack-9	.067	.850	.000	.895*	.000
Haystack-16	.000	.790	.000	.535	.000
Haystack-25	.000	.804*	.000	.651	.000
Haystack-36	.000	.596	.000	.668	.000

Table 2: Logarithmic curve fit  $r^2$  values are reported for each open-weight model, followed by a symbol denoting the  $p$ -value for statistical significance. The symbols represent: \* for  $p \leq 0.05$ , \*\* for  $p \leq 0.01$ , & \*\*\* for  $p \leq 0.001$ . Light pink cells represent correlations for sets of values that fall below chance performance.

Methods	GPT-4V (c)	GPT-4V (i)	Gemini 1.0 (c)	Gemini 1.0 (i)	Gemini 1.5 (c)	Gemini 1.5 (i)
MMStar	.880***	.937***	.972***	.972***	.985***	.981***
SEEDBench	.737***	.956***	.765***	.765***	.977***	.980***
MMBench	.951***	.959***	.977***	.977***	.964***	.948***
MMMU	.861***	.372***	.493**	.493**	.487**	.533**
MathVista	.676***	.858***	.832***	.832***	.840***	.690***
ScienceQA	.932***	.820***	.268*	.268*	.546**	.309*
AI2D	.743***	.635**	.880***	.880***	.969***	.882***
OK-VQA	.963***	.975***	.962***	.962***	.991***	.989***
BERT	.969***	.977***	.955***	.955***	.989***	.993***
ROUGE	.958***	.979***	.955***	.955***	.981***	.989***
Haystack-9	.505	.632	-	-	.052	.895*
Haystack-16	.366	.866*	-	-	.025	.256
Haystack-25	.863*	.704	-	-	.118	.146
Haystack-36	.841*	.645	-	-	.489	.896

Table 3: Logarithmic curve fit  $r^2$  values are reported for each closed-source model, followed by a symbol denoting the  $p$ -value for statistical significance. The symbols represent: \* for  $p \leq 0.05$ , \*\* for  $p \leq 0.01$ , & \*\*\* for  $p \leq 0.001$ . Light pink cells represent correlations for sets of values that fall below chance performance.

Models	Image Context	1 <sup>st</sup> Quartile	2 <sup>nd</sup> Quartile	3 <sup>rd</sup> Quartile	4 <sup>th</sup> Quartile
<b>Mantis</b>	$k = 4$	10.17%	10.73%	10.88%	10.96%
	$k = 9$	8.98%	10.28%	8.79%	10.04%
	$k = 25$	8.16%	8.21%	8.37%	8.64%
	$k = 36$	7.17%	7.72%	9.22%	7.75%
<b>GPT-4V</b>	$k = 4$	27.68%	23.03%	20.27%	23.90%
	$k = 9$	24.65%	15.96%	14.92%	18.92%
	$k = 25$	14.57%	10.51%	8.87%	10.12%
	$k = 36$	15.92%	9.06%	7.17%	6.96%
<b>Gemini 1.5</b>	$k = 4$	33.28%	31.78%	33.28%	34.07%
	$k = 9$	28.37%	28.72%	29.93%	28.60%
	$k = 25$	25.69%	22.82%	23.07%	24.36%
	$k = 36$	23.09%	21.91%	21.91%	21.81%

Table 4: *Impact of ordering on non-haystack tests*—to assess if a positional bias akin to that in the *needle in a haystack* test is present for non-trivial tasks, we separately-averaged the OK-VQA scores for three models by the quartile the content image appears in the sequence. This confirms the same positional biases observed for MNIST haystack tests hold in the more challenging VQA setting.

## F Case Study

### List of Case Study Figures

#### I OK-VQA – Example 1

Visual Context Length ( $k = 1$ ,  $k = 4$ ,  $k = 9$ ,  $k = 25$ ,  $k = 36$ )

#### II MMStar (SeedBench) – Example 2


Visual Context Length ( $k = 1$ ,  $k = 4$ ,  $k = 9$ ,  $k = 25$ ,  $k = 36$ )

#### III MMStar (MMBench) – Example 3

Visual Context Length ( $k = 1$ ,  $k = 4$ ,  $k = 9$ ,  $k = 25$ ,  $k = 36$ )

#### IV MMStar (MathVista) – Example 4

Visual Context Length ( $k = 1$ ,  $k = 4$ ,  $k = 9$ ,  $k = 25$ ,  $k = 36$ )

Example 1 – OK-VQA	
<b>Question:</b> What activity is this?	
	
<image>	
Visual Context Length ( $k = 1$ )	
LLaVA 1.5 (cmp):	This is a horseback riding activity
LLaVA 1.6 (cmp):	Horseback riding
Moondream2 (cmp):	Horseback riding
Mantis (int):	Horseback riding
GPT-4V (cmp):	horseback riding
GPT-4V (int):	horseback riding
Gemini 1.0 (cmp):	horseback riding
Gemini 1.0 (int):	horseback riding
Gemini 1.5 (cmp):	Horseback riding
Gemini 1.5 (int):	Horseback riding
PaliGemma (cmp):	horseback ride
Visual Context Length ( $k = 4$ )	
LLaVA 1.5 (cmp):	a girl riding a horse
LLaVA 1.6 (cmp):	Photography
Moondream2 (cmp):	biking
Mantis (int):	walking
GPT-4V (cmp):	horseback riding
GPT-4V (int):	horseback riding
Gemini 1.0 (cmp):	Horseback Riding
Gemini 1.0 (int):	Horseback Riding
Gemini 1.5 (cmp):	tennis
Gemini 1.5 (int):	tennis
PaliGemma (cmp):	horseback ride



Visual Context Length ( $k = 9$ )	
LLaVA 1.5 (cmp):	woman walking her horse
LLaVA 1.6 (cmp):	collage
Moondream2 (cmp):	biking
Mantis (int):	woman preparing to go surfing
GPT-4V (cmp):	bowling
GPT-4V (int):	horseback riding
Gemini 1.0 (cmp):	horseback riding
Gemini 1.0 (int):	horseback riding
Gemini 1.5 (cmp):	Bowling
Gemini 1.5 (int):	bowling
PaliGemma (cmp):	horseback riding
Visual Context Length ( $k = 16$ )	
LLaVA 1.5 (cmp):	This is a collage of various activities
LLaVA 1.6 (cmp):	Collage
Moondream2 (cmp):	Collage
Mantis (int):	a person riding a horse
GPT-4V (cmp):	surfing
GPT-4V (int):	Riding a horse
Gemini 1.0 (cmp):	surfing
Gemini 1.0 (int):	surfing
Gemini 1.5 (cmp):	horseback riding
Gemini 1.5 (int):	Surfing
PaliGemma (cmp):	collage
Visual Context Length ( $k = 25$ )	
LLaVA 1.5 (cmp):	This is a collage of photos, showcasing various activities.
LLaVA 1.6 (cmp):	Collage
Moondream2 (cmp):	Collage
Mantis (int):	Horseback riding
GPT-4V (cmp):	skiing
GPT-4V (int):	playing catch with a frisbee
Gemini 1.0 (cmp):	This is a photo collage
Gemini 1.0 (int):	This is a photo collage
Gemini 1.5 (cmp):	horseback riding
Gemini 1.5 (int):	frisbee
PaliGemma (cmp):	skiing
Visual Context Length ( $k = 36$ )	
LLaVA 1.5 (cmp):	a man eating a piece of cake
LLaVA 1.6 (cmp):	motorcycle
Moondream2 (cmp):	colleage
Mantis (int):	small garden with a shed
GPT-4V (cmp):	surfing
GPT-4V (int):	Unable to provide the activities
Gemini 1.0 (cmp):	This is a game of Where's Waldo.
Gemini 1.0 (int):	This is a game of Where's Waldo.
Gemini 1.5 (cmp):	Surfing
Gemini 1.5 (int):	life
PaliGemma (cmp):	collage

### Example 2 – MMStar (SeedBench)

**Question:** What is the color of the socks detected in the attribute detection of the <image>?

- (A) Black
- (B) Red
- (C) White
- (D) Striped



<image>

#### Visual Context Length ( $k = 1$ )

LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	B
Moondream2 (cmp):	A
Mantis (int):	B
GPT-4V (cmp):	C
GPT-4V (int):	C
Gemini 1.0 (cmp):	B
Gemini 1.0 (int):	B
Gemini 1.5 (cmp):	A
Gemini 1.5 (int):	A
PaliGemma (cmp):	C

#### Visual Context Length ( $k = 4$ )

LLaVA 1.5 (cmp):	A
LLaVA 1.6 (cmp):	B
Moondream2 (cmp):	A
Mantis (int):	B
GPT-4V (cmp):	C
GPT-4V (int):	C
Gemini 1.0 (cmp):	D
Gemini 1.0 (int):	C
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	A

Visual Context Length ( $k = 9$ )	
LLaVA 1.5 (cmp):	A
LLaVA 1.6 (cmp):	B
Moondream2 (cmp):	A
Mantis (int):	A: Black, B: Red, C: White
GPT-4V (cmp):	There are no visible socks in the image.
GPT-4V (int):	None of the images provided.
Gemini 1.0 (cmp):	C
Gemini 1.0 (int):	C
Gemini 1.5 (cmp):	Blue
Gemini 1.5 (int):	There are no socks in the image.
PaliGemma (cmp):	A
Visual Context Length ( $k = 16$ )	
LLaVA 1.5 (cmp)	A
LLaVA 1.6 (cmp)	B
Moondream2 (cmp)	A
Mantis (int)	A: Black, B: Red, C: White, D: Striped
GPT-4V (cmp)	There are no socks visible in the image
GPT-4V (int)	None of the above.
Gemini 1.0 (cmp)	C
Gemini 1.0 (int)	C
Gemini 1.5 (cmp)	C
Gemini 1.5 (int)	There are no socks in the image.
PaliGemma (cmp)	C
Visual Context Length ( $k = 25$ )	
LLaVA 1.5 (cmp)	C
LLaVA 1.6 (cmp)	B
Moondream2 (cmp)	A
Mantis (int)	B
GPT-4V (cmp)	C
GPT-4V (int)	Option C is not provided
Gemini 1.0 (cmp)	C
Gemini 1.0 (int)	C
Gemini 1.5 (cmp)	There are no socks in the image.
Gemini 1.5 (int)	There are no socks in the image.
PaliGemma (cmp)	B
Visual Context Length ( $k = 36$ )	
LLaVA 1.5 (cmp)	C
LLaVA 1.6 (cmp)	B
Moondream2 (cmp)	B
Mantis (int)	C
GPT-4V (cmp)	I am unable to provide assistance with this request
GPT-4V (int)	None of these
Gemini 1.0 (cmp)	A
Gemini 1.0 (int)	C
Gemini 1.5 (cmp)	There are no socks in the image.
Gemini 1.5 (int)	There are no socks in the image.
PaliGemma (cmp)	B

**Example 3 – MMStar (MMBench)**

**Question:** Which is the main topic of the <image>?

- (A) A woman surfing
- (B) A man skating
- (C) A man surfing
- (D) A woman skating



<image>

**Visual Context Length ( $k = 1$ )**

LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	C
Moondream2 (cmp):	A
Mantis (int):	C
GPT-4V (cmp):	C
GPT-4V (int):	C
Gemini 1.0 (cmp):	C
Gemini 1.0 (int):	C
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	C

**Visual Context Length ( $k = 4$ )**

LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	C
Moondream2 (cmp):	A
Mantis (int):	C
GPT-4V (cmp):	C
GPT-4V (int):	C
Gemini 1.0 (cmp):	C
Gemini 1.0 (int):	C
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	C

Visual Context Length ( $k = 9$ )	
LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	C
Moondream2 (cmp):	A
Mantis (int):	C: A man surfing, D: A woman skiting
GPT-4V (cmp):	C
GPT-4V (int):	-
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	C
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	C

Visual Context Length ( $k = 16$ )	
LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	C
Moondream2 (cmp):	A
Mantis (int):	C: A man surfing, D: A woman skiting
GPT-4V (cmp):	C
GPT-4V (int):	-
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	A
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	C

Visual Context Length ( $k = 25$ )	
LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	A
Mantis (int):	B
GPT-4V (cmp):	-
GPT-4V (int):	-
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	A
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	B
PaliGemma (cmp):	D

Visual Context Length ( $k = 36$ )	
LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	A
Mantis (int):	D
GPT-4V (cmp):	C
GPT-4V (int):	-
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	A
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	C

**Example 4 – MMStar (MathVista)**

**Question:** What is the age gap between these two people in <image>? (Unit: years)

- (A) 4
- (B) 3
- (C) 2
- (D) 1



<image>

**Visual Context Length ( $k = 1$ )**

LLaVA 1.5 (cmp):	E
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	A
Mantis (int):	A
GPT-4V (cmp):	I'm sorry, but I cannot provide information
GPT-4V (int):	I'm sorry, but I cannot provide information
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	A
Gemini 1.5 (cmp):	A
Gemini 1.5 (int):	A
PaliGemma (cmp):	A

**Visual Context Length ( $k = 4$ )**

LLaVA 1.5 (cmp):	D
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	B
Mantis (int):	B
GPT-4V (cmp):	I'm sorry, but I cannot provide information
GPT-4V (int):	I'm sorry, but I cannot provide information
Gemini 1.0 (cmp):	B
Gemini 1.0 (int):	B
Gemini 1.5 (cmp):	D
Gemini 1.5 (int):	C
PaliGemma (cmp):	B

Visual Context Length ( $k = 9$ )	
LLaVA 1.5 (cmp):	E
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	B
Mantis (int):	B
GPT-4V (cmp):	I'm sorry, but I cannot provide information
GPT-4V (int):	I'm sorry, but I cannot provide information
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	D
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	B

Visual Context Length ( $k = 16$ )	
LLaVA 1.5 (cmp):	C
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	B
Mantis (int):	A
GPT-4V (cmp):	I'm sorry, but I cannot provide information
GPT-4V (int):	I'm sorry, but I cannot provide information
Gemini 1.0 (cmp):	B
Gemini 1.0 (int):	B
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	D
PaliGemma (cmp):	B

Visual Context Length ( $k = 25$ )	
LLaVA 1.5 (cmp):	E
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	B
Mantis (int):	B
GPT-4V (cmp):	I'm sorry, but I cannot provide information
GPT-4V (int):	-
Gemini 1.0 (cmp):	D
Gemini 1.0 (int):	D
Gemini 1.5 (cmp):	C
Gemini 1.5 (int):	C
PaliGemma (cmp):	A

Visual Context Length ( $k = 36$ )	
LLaVA 1.5 (cmp):	D
LLaVA 1.6 (cmp):	A
Moondream2 (cmp):	B
Mantis (int):	B
GPT-4V (cmp):	-
GPT-4V (int):	Impossible to determine
Gemini 1.0 (cmp):	A
Gemini 1.0 (int):	A
Gemini 1.5 (cmp):	D
Gemini 1.5 (int):	C
PaliGemma (cmp):	B