

DARE: Diverse Visual Question Answering with Robustness Evaluation

Hannah Sterz¹ Jonas Pfeiffer² Ivan Vulić¹

¹Language Technology Lab, University of Cambridge, UK

²Google DeepMind, Switzerland

Abstract

Vision Language Models (VLMs) extend remarkable capabilities of text-only large language models and vision-only models, being able to learn from and process multi-modal vision-text input. While modern VLMs perform well on a number of standard image classification and image-text matching tasks, they still struggle with a number of crucial vision-language (VL) reasoning abilities such as counting and spatial reasoning. Moreover, while they might be very brittle to small variations in instructions and/or evaluation protocols, existing benchmarks fail to evaluate their robustness (or rather the lack of it). In order to couple challenging VL scenarios with comprehensive robustness evaluation, we introduce **DARE**, **D**iverse **V**isual **Q**uestion **A**nswering with **R**obustness **E**valuation, a carefully created and curated multiple-choice VQA benchmark. DARE evaluates VLM performance on five diverse categories and includes four robustness-oriented evaluations based on the variations of prompts, the subsets of answer options, the output format, and the number of correct answers. Among a spectrum of other findings, we report that state-of-the-art VLMs still struggle with questions in most categories and are unable to consistently deliver their peak performance across the tested robustness evaluations. Consequently, our work calls for the systematic addition of robustness evaluations in future VLM research.

1 Introduction and Motivation

Building on the recent groundbreaking accomplishments of text-only Large Language Models (LLMs) across a wide range of (text-based) NLP tasks (Chiang et al., 2023; Touvron et al., 2023), there has been a growing interest in Vision-Language Models (VLMs) (Liu et al., 2023; Zhu et al., 2024; Gemini Team Google, 2023; Chen et al., 2023). Such VLMs expand the input to comprise images as well as text, and enable dealing

with cross-modal and multi-modal tasks and reasoning (Alayrac et al., 2022; Achiam et al., 2023; Liu et al., 2023).

VLMs have been proven to perform well in standard image classification and image-text matching and reasoning tasks (Radford et al., 2021). For instance, performance on the VQAv2 benchmark (Goyal et al., 2017), which targets visual understanding and commonsense reasoning, is approaching human performance. The same also holds for benchmarks that require parsing and understanding of document images and text embedded into images such as TextVQA (Singh et al., 2019) and DocVQA (Mathew et al., 2021).

However, there still exist particular vision-language (VL) reasoning problems which pose a challenge for the modern VLMs, such as counting and spatial reasoning (Onoe et al., 2024; Li et al., 2024). Furthermore, standard benchmarks fail to evaluate *model robustness* to input and question variation,¹ a chosen inference protocol, as well as to specifications of the output and its format. High robustness to these aspects would ensure that the answer is not based on biases learned during training but due to a good understanding of the questions as well as of the instructions provided to the VLM. Put simply, we define robustness as the *ability to perform consistently across such variations*.

To study and address these gaps, especially emphasizing the robustness aspects, we carefully construct a new and challenging VQA benchmark, termed **D**iverse **V**isual **Q**uestion **A**nswering with **R**obustness **E**valuation (DARE). It spans multiple-choice questions across five diverse, challenging scenarios/categories (e.g., conditional counting, visual commonsense, among others,

¹In simple words, these benchmarks only include each question once and in a particular form, whereas a robust VLM should be able to answer correctly irrespective of different variations of the question where the correct answer can be clearly and unambiguously denoted (e.g., counting or determining spatial relations of objects).

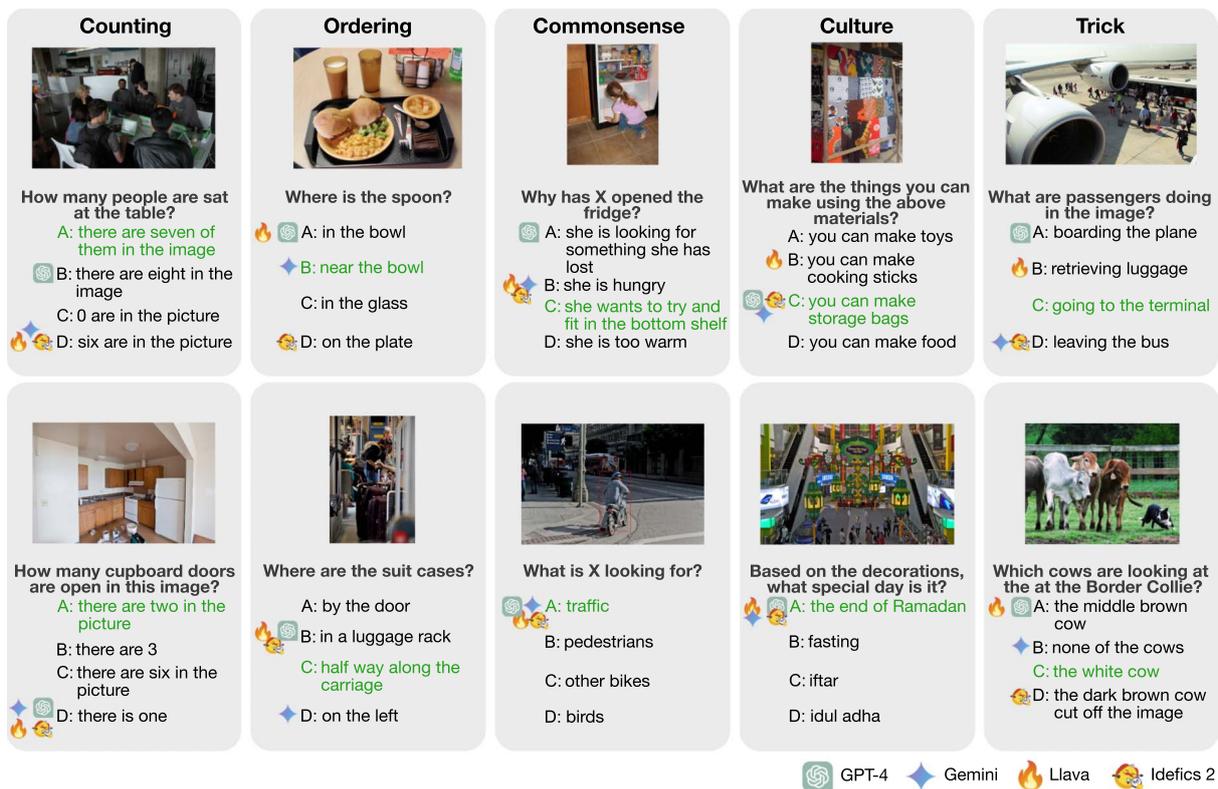


Figure 1: Examples of single-correct questions for each of the five categories covered in DARE.

see later in §3). DARE comes with the following key features. **1)** Diverse scenarios cover a range of crucial vision-language reasoning abilities needed by VLMs. **2)** Since current standard benchmarks are already largely saturated, they cannot (anymore) provide a good and insightful overview of the VLMs’ abilities. Therefore, DARE comprises challenging and carefully curated evaluation instances from all five diverse scenarios. **3)** DARE provides the opportunity to study robustness within VQA evaluation.

We study robustness across multiple axes, which can roughly be grouped into **1)** variations of the prompt/instruction, **2)** variations of the set of answer options, **3)** variations of the output format, and **4)** variations of the number of correct answers; see later in §3.2. These axes of robustness, while typically neglected and discarded in current evaluations of VLMs, can unveil biases learned during pretraining and robustness to variations of the task, and potential inconsistencies of VLMs.

Unlike prior work on complex VL datasets where the primary purpose has been to evaluate complex visual understanding and language grounding (Gupta et al., 2019; Thrush et al., 2022; Kamath et al., 2023), the scenarios covered by

DARE do not rely on visually challenging images. Instead, they target reasoning abilities that are inherently easy for humans but require understanding of the question in the text form *coupled* with image understanding; see Figure 1. First, we include *conditional counting* and *spatial reasoning* as well-defined visual understanding tasks that VLMs are known to struggle with, and they require no additional (world) knowledge.² Two other scenarios require additional knowledge. *Visual commonsense reasoning* questions require an understanding of the image and knowledge about the impact on the thoughts and actions of an individual. The *culture* scenario requires knowledge about different cultures including their religions, dishes, and customs. The final, so-called *trick* scenario is model-based and creates challenging questions by identifying mistakes in image descriptions created by most powerful, state-of-the-art VLMs, GPT-4 and Gemini. We believe that this combination of scenarios, while definitely non-exhaustive, covers a wide range

²Conditional counting has the additional challenge that it requires multiple steps. First, all objects of a certain class need to be identified and filtered by some condition then the remaining objects need to be counted.

of skills and can provide detailed insights into VLM performance when coupled with additional robustness challenges.

We then use DARE to examine and profile state-of-the-art closed-source and open-weights VLMs across the defined scenarios and robustness aspects, also aiming to trace their performance ‘chronologically’ by contrasting performance of most recent versus earlier checkpoints of the same VLM. As some main findings, we report that VLMs still struggle with the ‘simple-for-humans’ vision understanding tasks of conditional counting and spatial reasoning, and they are not robust to variations in answer options. Even if the correct answer has the same content across all variations and is just phrased differently, the worst-case performance can be up to 33.6% lower than the one observed with ‘vanilla’ evaluation. All VLMs perform worse on questions with a varying number of possible correct answers. Especially LLaVA 1.6 and Idefics2 show a strong bias towards marking exactly one answer as correct. We also find that the generation and extraction strategies of the correct answers heavily impact performance: e.g., Gemini benefits from predicting the answers in JSON format, while LLaVA and Idefics2 do not perform well when prompted to provide JSON-formatted answers.

In hopes to guide future developments of VLMs, we share DARE at <https://huggingface.co/datasets/cambridge1tl/DARE>.

2 Background and Related Work

VQA: Preliminaries. VQA is a task of providing the correct answer(s) given an image and a question about the image. In the standard *multiple-choice* VQA variant, n possible answer options to choose the correct answer(s) from are also provided, where one can differentiate between (a) the ‘*single-correct*’ setup (where it is known that there is always a single correct answer among the provided options) and (b) the ‘*multi-correct*’ setup³ (or *0-to-n* setup, where any number of correct answers from the provided options is allowed, including 0 and all n of them). We follow this format and definitions of different setups in this work.

³This differs substantially from the multiple-questions multiple-answers setup (Tang et al., 2024), where multiple questions are processed at the same time and the model generates one answer for each question.

VQA Tasks and Datasets. The VQA task was introduced in the VQA benchmark (Antol et al., 2015), where it consisted of short questions about the image that required visual understanding and commonsense. This dataset contains language bias,⁴ later addressed by VQAv2 (Goyal et al., 2017), which is still a standard benchmark despite the fact that it approaches human-level performance (Chen et al., 2022). GQA (Hudson and Manning, 2019) provides questions about images covering targets addressing the semantic compositionality of scenes. This dataset is commonly treated via a classification task where all possible answers are possible classes. Consequently, this way the model can only answer questions where the correct answer is represented by one of the classes.

Several VQA datasets cover specific, finer-grained scenarios. VCR (Zellers et al., 2019) targets visual commonsense reasoning, which is the ability to make assumptions about plausible explanations for the motivation and next actions and thoughts. It uses movie scenes as images which results in a more limiting license of the dataset. CVQA (Romero et al., 2024) and CulturalVQA (Nayak et al., 2024) are VQA datasets that cover culturally diverse questions, limited to single-correct setups, and without any robustness analyses. TextVQA (Singh et al., 2019), DocVQA (Mathew et al., 2021) and InfographicVQA (Mathew et al., 2022) target the ability to extract information of images of text, documents or infographics. VizWiz (Gurari et al., 2018) collects images and questions asked by blind people to interact with their environment. This includes object identification, colour detection, or reading of texts. The model-based MMVP benchmark (Tong et al., 2024) targets weaknesses of current models by (i) identifying image pairs which have similar CLIP-computed embeddings (Radford et al., 2021) but different DINOv2-based embeddings (Oquab et al., 2024), and then (ii) annotating them with questions about the differences.

Some benchmarks target multiple scenarios. The BLINK benchmark (Fu et al., 2024) focuses on visually challenging tasks such as finding

⁴Depending on the start of the question some answers are more likely: e.g., for a question starting with ‘what sport’, ‘tennis’ is the most likely answer, and correct in 41% cases.

bounding boxes for objects, image similarity, and camera movement between images. MMBench (Liu et al., 2024c) covers perception and reasoning questions. The majority of the samples focus on perception. *Why DARE?* Specialized datasets can only provide detailed insights into one single aspect of the VLMs performance. While benchmarks such as MMBench and BLINK target scenarios that are difficult for current VLMs, they focus on (direct) perception. With DARE we aim to cover scenarios that require reasoning such as filtering objects by a condition, spatial reasoning, and common sense reasoning. In the creation of DARE our goal is also to avoid the use of already established and (potentially and likely) saturated datasets, and to manually validate and curate all the included data instances to ensure high data quality.

Robustness Evaluation. Robustness refers to the ability to perform the same task across variations of the instances (Jia and Liang, 2017; Dhole et al., 2023; Liang et al., 2023; Sclar et al., 2024). These variations can be performed over a variety of aspects. In the scope of multiple-choice questions, Wang et al. (2024) and Pezeshkpour and Hruschka (2024) show that the correct answer position within the options influences the performance of the model. Contrary to that, Zheng et al. (2024) find that not the order but rather the tokens indicating the options (e.g., A–D, 1–4) impact prediction.

Another axis considered in some existing *text-only* datasets is variations to the input. Variations such as deliberate typos, perturbations, and synonymy substitutions provide insights into the ability of the LLMs to respond to variations of the same question correctly (Jia and Liang, 2017; Dhole et al., 2023; Liang et al., 2023). Such robustness tests then hint to worst-case model performance across the variations. While there are text-only language understanding benchmarks with such robustness measures, vision-language benchmarks are just starting to include robustness. MMBench (Liu et al., 2024c) performs a circular evaluation that prompts the model multiple times with the same question and shuffled options to reduce the impact of the answer position. However, this only covers a small set of possible variations of the VQA task. To the best of our knowledge, DARE is the first vision-language benchmark to include robustness across several key axes as its crucial feature.

3 DARE: Dataset Overview

As introduced in §2, current VL benchmarks critically lack robustness evaluations: They only evaluate the task for specific instances. However, the VLM task performance may be heavily impacted by model biases and might thus not stem from actually understanding the provided text and images. We therefore create DARE, a multiple-choice VQA benchmark with multiple robustness evaluations. In a nutshell, the DARE dataset combines 5 different categories requiring reasoning and visual understanding, where all categories are framed as multiple-choice questions.⁵ We first briefly introduce the 5 categories (§3.1) followed by aspects of robustness evaluation (§3.2). A detailed description of DARE creation is provided in §3.3.

3.1 Categories in DARE

The five categories in DARE cover a range of reasoning and vision understanding tasks. We include conditional counting and ordering to test vision understanding, one-hop reasoning, and spatial reasoning. As categories that need (world) knowledge we include visual commonsense reasoning (VCR) and culture-based VQA. The VCR questions provide insights into the ability to understand the scene and infer the thoughts and motivations of a person in the scene. The culture category tests knowledge about different cultures. Finally, inspired by recent work (Tong et al., 2024), the trick category targets shortcomings of the (currently) best-performing VLMs make in image descriptions. Some examples per each category are illustrated in Figure 1.

Conditional Counting. These questions require counting objects based on an additional condition; it can be color, position in the image, material, etc. Therefore, answering the questions requires not only identifying all objects of one type in the image, but also applying an additional filter on them before counting: e.g., the first example in Figure 1 requires identifying all people and then filtering only for people who are sitting.

Ordering. It targets the ability to understand the order of objects and their spatial relation; e.g.,

⁵The multiple-choice format allows easy evaluation with accuracy metrics while still enabling coverage of a wide variety of questions with different topics and complexity.

the first example question in Figure 1 requires to identify the relation between the spoon and the other objects, whereas the second example shows a challenging data instance: One of the wrong answer options, ‘*In a luggage rack*’, is where one would expect luggage a priori without seeing the image.

Visual Commonsense Reasoning/VCR. In order to answer the questions, it is required to grasp the scene and make commonsense derivations from it about a person marked in the image, which might cover plausible motivation, thoughts, or the most likely next action of the person.

Culture. Many objects and concepts encountered both in text and images are rooted in their culture and/or differ across different cultures (Liu et al., 2024a). This category thus requires answering questions about images of important concepts in different cultures, where we (incorrectly, for simplicity) approximate the culture by the language spoken by people. This requires the model to reason over concepts from a wide set of cultures. For instance, the examples in Figure 1 require knowledge about holidays, and use of certain objects.

Trick. Here, we directly target mistakes in descriptions generated by state-of-the-art VLMs, where the key assumption is that the mistakes in the description point to incorrect understanding of that aspect of the image. To this end, we use GPT-4 and Gemini to obtain descriptions that we present to the annotators so that they can identify misconceptions and write questions targeting them. We note that through this design principle the questions are tied to the model used to generate the description. Therefore, this category should be used only for evaluation and comparison of models which were not used to generate the descriptions. The examples in this category cover a wide range of challenging questions (e.g., see examples in Figure 1).

3.2 Overview of Robustness Evaluation Scenarios in DARE

Robustness. Task performance of LLMs and VLMs (as assessment of their abilities) may be substantially impacted by the prompt, examples, variations of the same text, and other factors (Wei et al., 2022; Liang et al., 2023; Lu et al., 2022;

Zheng et al., 2024). Robustness is the ability to perform consistently across those variations.

Variations in different aspects of the task require different skills from the model. Therefore, it is important to evaluate them in a variety of variants of the task to get a comprehensive understanding of model robustness. These variants can cover all aspect of the task. We include variants that are easy for a human as they do not change the content of the prompt or question. These variations cover the central parts of the task: prompt, answer options, and output to determine the robustness to changes in these aspects. Moreover, we look at the robustness to the number of correct options which is a more challenging version of the task to evaluate the model’s ability to adapt to changes in the structure of the task. Each of the four robustness axes requires different skills. This enables us to get a broad view on the model’s robustness and what type of variations the model fails on.

Single-Correct Answer is the standard, default scenario which serves as the reference point for all robustness evaluations. Each question has four answer options (by default marked with capital letters A–D), out of which exactly one is correct. The prompt provides the information that there is exactly one correct answer (see Appendix A.1 for the full prompt). This setup is comparable to other multiple-choice datasets. The information allows one to pick the most likely answer without the need to decide the correctness of each option independently. Starting from this basic setup, DARE currently includes four axes of variations that test model robustness:

(1) Varying the Sets of Answer Options. After collecting up to four correct and four incorrect options (see Appendix 3.3), we sample three sets of four options where each set has a single correct option. The challenge in this setup is to answer the same question correctly more than once and to identify different correct answer options (among varying incorrect options as well). We report accuracy over all three variations. This provides the worst-case performance of the model for ‘unfortunate variations’ of the data instances. Again, we provide the information that there is always a single correct answer as part of the prompt. We refer to this setup as *Single-Correct Answer + Variations*. Table 1 shows the number of samples for this setup (columns *Single+Var*).

	Validation			Test		
	<i>Single</i>	<i>Single+Var</i>	<i>Multi</i>	<i>Single</i>	<i>Single+Var</i>	<i>Multi</i>
Count	250	250	250	250	250	250
Order	231	228	250	237	237	250
VCR	250	–	–	258	–	–
Culture	–	–	–	223	178	261
Trick	232	232	250	229	229	250
Total	963	710	750	1197	894	1011

Table 1: The number of questions/samples across different evaluation scenarios (*Single*: single correct answer, *Single+Var*: single correct answer plus variations of the set of answer options (see §3.2), *Multi*: multiple (0 to n) correct answers possible).

(2) Multiple Correct Answers (0-to- n). Here, we remove the requirement that there is only a single correct option compared to the *Single-Correct Answer* setup. This requires the models to determine themselves the number of correct options: it is inherently a more challenging task as there are more options one can answer with, and there is fewer meta-information that can be used to rule out options. As a result, this is also more challenging for humans. For questions with multiple correct answer options, the performance on questions with one or more correct answers is different from the performance on questions with zero correct answers. Averaging the accuracy thus can lead to misleading results. For instance, a model that would simply label each question as not having an answer would get an accuracy $\sim 20\%$. To avoid this conflation, we report results separately for question sets with $i = 0, \dots, 4$ correct answers.

(3) Varying the Prompt. We use three different prompts in the *Single-Correct Answer* and the *Multiple-Correct Answer* scenarios (see Appendix A.1 for the prompts). We report the mean and standard deviation of the results on these prompts.

The first prompt is a variation of the MMLU prompt (Hendrycks et al., 2021), adapted to VQA multiple-choice questions. The second prompt introduces a scenario (*Imagine you are a student (. . .)*) in which it is crucial to answer the questions correctly.⁶ The third prompt describes the

⁶This prompt draws inspiration from the work of Frohmann et al. (2024), which suggests and then empirically validates that using a prompt which defines a scenario that has high stakes for the model improves performance.

task factually, aiming to also minimise the lexical overlap with the first, default prompt.

(4) Varying the Evaluation Protocol and Output Format.

In order to make the answers usable for other software for systems, the model must be able to provide them in a specified format. We test the ability of the model to provide the answer in structured formats such as JSON and CSV, and compare it against standard evaluation protocols for multiple-choice questions such as (i) ‘running text’ generation and (ii) direct comparison of logits. In the first, ‘vanilla’ setup we prompt the model to provide the answer as a single character (i.e., choosing from A–D) or a list of characters, depending on the number of correct options expected (single-correct versus multi-correct setups). Given that VLMs are typically also not very robust to formatting even this simple output according to the provided instructions (i.e., we can encounter variants of the output such as *B*, *B.*, *B:* etc.), we manually define a regular expression aimed to regularize this variation and extract the final, normalised answer (see Appendix A.2). This variant is termed OUT-GEN.

The second variant uses logit probabilities of output characters denoting the answer options. We limit the set of tokens the model can output to the tokens enumerating the options, that is, A, B, C, and D. For the single-correct setup, we let the model generate exactly one of these tokens. To extend this to the multi-correct setup, we include the *end-of-generation* token. In this scenario, the output length is not limited to one. We let the model greedily generate tokens until the end-of-generation token obtains the highest probability or the maximum number of tokens is reached. Each option corresponding to generated tokens is labelled as correct. This variant, termed, OUT-LOG, requires access to logit probabilities, which is not available for many closed-source API-gated models.

We can also prompt the models to provide the answers in a valid JSON format (e.g., an example output is $\{\text{“answer”}: \text{“B”}\}$) - this variant is referred to as OUT-JSON. As a side experiment, we further probe robustness by experimenting with a small selection of other possible output formats (including instructions which are very easy for humans to comprehend) later in §4. The exact prompts for all the variants are listed in Appendix A.

Impact on Task Difficulty. Robustness evaluations change the circumstances of the task to test whether the model can maintain performance. These changes might have consequences on task complexity and difficulty. In the following, we discuss the impact of the introduced robustness variations on task difficulty.

All used *prompts* contain the same information; thus the prompt does not affect task difficulty. The *output format* needs to be structured or constrained so that a parsing algorithm can extract the options from the generated text. We choose the formats so they are easy to interchange for humans. For instance, the prompts for the JSON format include an example that enables someone without knowledge of JSON to provide the correct answer(s). However, these variations, even though they are easy for humans, could still result in changed task difficulty.

To create the *sets of answer options*, the options are drawn randomly from the annotations. Therefore, we assume that the different subsets have the same difficulty on average. The *varying number of correct options* variation increases difficulty notably. The chance of correctly guessing is 25% for single-correct setups. However, by allowing multiple or no correct answers, the set of correct answers is the power set of A, B, C, D with a probability of a correct guess of 6.25%.

3.3 Data Creation and Annotation Process

Image Selection. For the *conditional counting*, *ordering*, *VCR*, and *trick* categories, we manually select 500 images from the COCO 2017 dataset (Lin et al., 2014), with 250 images taken from its development set, and another 250 from the test set. The main (albeit subjective) selection criterion is that the images must show scenes that provide the grounds for challenging questions that are aligned with the corresponding categories.⁷ As a necessary extra step in *VCR*, we use DETR (Carion et al., 2020) to mark the subject of the question with a bounding box (denoted by X in the question, see Figure 1). The boxes are manually validated to ensure that they mark a clearly visible person. For *trick*, we generate descriptions with

⁷For instance, close-ups of individual objects typically *do not allow* for the construction of challenging questions that require counting, spatial reasoning or VCR.

Gemini (Gemini Team Google, 2023) and GPT-4 (Achiam et al., 2023), each model describing one half of the development set and one half of the test set.

The *culture* category comprises images sampled from the MaRVL dataset (Liu et al., 2021), following a similar procedure as with COCO. We approximate cultures by languages and select, again manually, a diverse subset of images for each of the five languages, representing a variety of concepts. We annotate the same number of images for each language. After quality control this results in 61 ‘Chinese’, 38 ‘Swahili’, 57 ‘Tamil’, 60 ‘Turkish’, and 45 ‘Indonesian’ samples. Since COCO and MaRVL images are available online, we cannot rule out that VLMs for which their training data are not disclosed have seen the images during pretraining. To prevent contamination, we do not publish the annotations of the annotated test set of DARE.

Question Creation. We present human annotators⁸ with an image and ask them to write a question for one of the categories, where each category comes with its own, customized set of annotation instructions with some examples (see Appendix §B for details). For the conditional counting category, we ask them to provide the question and the correct number as the answer. This way, we can generate correct and incorrect answer options from templates to get a multiple-choice setup consistent with the other categories. For the ordering and trick categories, we ask the annotators to provide the question as well as four correct answers and four incorrect answers. We employ a similar setup for the culture category, but we require the annotators to be fluent speakers of the language the culture is approximated by. This ensures that they are familiar with the concepts shown in the image. The commonsense questions ask for the most plausible answer. We ask annotators to provide the question, a single, most plausible answer and

⁸All annotators are hired through Prolific and compensated with \$12 per hour. We require annotators to have English as their first language and have a high-school diploma or higher education. The language requirement for cultural questions (during question creation, quality control) is adapted to being fluent in English and the language associated with the culture, to ensure annotators are familiar with the culture.

three other plausible answers, and take the most plausible option as the single correct answer.⁹

Quality Control. To ensure high quality in DARE, we run a quality control stage, where we ask another set of annotators to validate the annotated samples. For *conditional counting*, we ask two annotators to answer the questions with the correct number. If both agree with the (original) annotation we include it in the dataset and generate answer options using predefined templates to obtain a multiple-choice question. If there is a tie between the annotators we ask a third annotator to answer the question and include the data instance into DARE iff they also agree with the original correct answer and one of the two competing annotators. When dealing with categories which can have multiple correct (and incorrect) answers, we present each option separately with the question and the corresponding image and ask the annotator to decide whether the answer is correct or not for the given question and image. We get the answers from two annotators and include a third one if they disagree. We keep only questions where two control annotators agree with the original annotation.

For *commonsense*, the notion of a correct answer is different than for the other categories: All answer options should be plausible. The answer that should be selected is then the most plausible one. We employ majority voting to determine which answer option is considered the most plausible by most annotators. This can be different from what the annotator intended as the most plausible option.

Human Performance. We evaluate human performance on the DARE test set by sampling 100 samples for each category and assigning humans¹⁰ to provide answers. As the multi-correct setup is more difficult than the single-correct setup, we assess the human baseline for both scenarios separately. Human performance in the single-correct setup is shown in Table 2. Humans seem to find

⁹Since for VCR we consider the most plausible option as the correct answer, this category cannot be easily transferred into the scenario with multiple correct answers: i.e., the plausibility threshold that determines correct answers would be very subjective and we would expect humans to perform poorly on that task as well. Therefore, we skip the robustness test with multiple correct answers for VCR, cf. Table 1.

¹⁰These annotators are also hired through Prolific with the same requirements as data creation and quality control.

	Count	Order	VCR	Culture	Trick
Human Acc.	92	96	70	82	92

Table 2: Human performance (accuracy) in the single-correct answer setup.

	0	1	2	3	4	F-1
Count	77	86	81	89	79	89
Order	67	87	68	59	59	88
Culture	59	67	44	18	41	73
Trick	82	74	54	57	41	85

Table 3: Human performance in the multi-correct setup; accuracy per question group (based on the number of correct answers) and averaged F-1 scores reported.

categories *conditional counting*, *order*, and *trick* easy, with corresponding accuracies $> 90\%$. The lowest accuracy is achieved for the VCR category. Here, the correct answer is the most plausible out of a set of plausible options, which makes it more subjective and correspondingly more challenging than the other categories where the answers are much more clear-cut.

Human performance in the more challenging multi-correct setup is shown in Table 3. As expected, absolute performance drops compared to the single-correct setup. Foreshadowing, human performance for *conditional counting*, *order*, and *trick* categories is clearly stronger than performances of the VLMs in our evaluation (cf. §4 later). For the *culture* category, annotators show a tendency to only select one correct answer and not consider the other options, which is the main reason behind slightly lower scores for that category.

Final Dataset. We note that we apply only the prompt variations and a subset of the output formats to the questions in the multi-correct scenario.¹¹

The final data statistics over different categories and setups are provided in Table 1. Unless stated otherwise, we always report the results on the

¹¹Varying the set of answer options is not possible while keeping the number of correct options fixed for each question. This would introduce multiple variations at once, which is why we focus only on variations in prompt and output format for multi-correct setups, and we look at the performance for the different number of correct options separately.

	Count			Order			VCR			Culture			Trick		
	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON
GPT-4	46.1 \pm 4	40.9 \pm 4	46.3 \pm 1	52.8 \pm 2	52.0 \pm 1	51.9 \pm 2	63.1 \pm 2	60.6 \pm 2	63.2 \pm 1	78.3 \pm 6	83.4 \pm 2	86.4 \pm 1	53.7 \pm 1	55.7 \pm 2	55.0 \pm 2
Gemini	55.5 \pm 2	-	53.7 \pm 9	71.6 \pm 11	-	74.3 \pm 6	60.7 \pm 3	-	61.5 \pm 1	83.8 \pm 7	-	87.9 \pm 2	70.0 \pm 5	-	71.5 \pm 7
Llava 1.6	42.0 \pm 8	40.9 \pm 4	30.8 \pm 20	48.7 \pm 9	45.9 \pm 2	38.1 \pm 2	51.2 \pm 1	51.2 \pm 4	29.6 \pm 23	72.2 \pm 3	75.4 \pm 0	51.9 \pm 4	51.7 \pm 3	51.2 \pm 3	45.6 \pm 4
Idefics2	46.8 \pm 2	44.1 \pm 1	0.0 \pm 0	59.1 \pm 1	59.4 \pm 1	0.0 \pm 0	60.9 \pm 2	60.3 \pm 2	0.0 \pm 0	78.8 \pm 9	84.2 \pm 1	0.0 \pm 0	59.5 \pm 1	57.3 \pm 1	0.0 \pm 0

Table 4: Average accuracy and standard deviation on the questions with a single correct answer for the five categories across three different instruction prompts, and across three different output formats (GEN, LOG, JSON, see §3.2).

test portion, for which the correct answers are not publicly available to prevent data contamination (Dong et al., 2024). The validation portion, with the correct answers available, is intended for hyper-parameter tuning and local evaluation. Along with the final dataset, we also release the annotations from which we sample the setups from Table 1; this might enable creation of new robustness evaluations by other researchers as well, beyond the ones introduced in our work.

4 Experiments and Results

Model Selection. We select a representative combination of closed API-gated and open-sourced models, focusing on the models considered state-of-the-art and with the strongest performance on previous VL benchmarks. The choice has been further motivated by the aims of 1) getting a diverse perspective on their performance and robustness on diverse scenarios of DARE, and also 2) the ability to trace model performance and progress ‘historically’ by assessing the most up-to-date checkpoints of the models as well as their earlier checkpoints. We choose the following models: Gemini Flash, GPT-4, LLaVA 1.6, and Idefics 2 (see Appendix A.3 for technical details and hyper-parameters).

Evaluation Metrics. Unless stated otherwise, we report performance as *accuracy* scores, capturing the proportion of correctly answered questions. For the multi-correct setup, to also take partially correct answers into account, we additionally provide averaged F-1 scores: They are computed on the option level based on the recall and precision of identifying correct options.

4.1 Main Results and Discussion

We now discuss the results across different categories and robustness aspects discussed in §3.

Single-Correct Answer Setup. The results, summarized in Table 4, indicate that even the

simplest multiple-choice scenario is challenging for the state-of-the-art VLMs, with ample room for improvement. Gemini Flash shows the highest scores on average, followed by GPT-4 and Idefics2 which achieves similar results with OUT-GEN and OUT-LOG as Gemini Flash and GPT-4. However, Idefics2 fails completely with OUT-JSON.

While the scores for all categories reveal substantial gaps to human-level performance except for *culture*, *conditional counting* seems especially challenging for all the VLMs in our evaluation, and across different evaluation protocols and formats. It is possible to attain higher absolute peak scores for the *order* and *trick* categories, but this is achieved with a subset of models coupled with specific evaluation protocols and output formats, suggesting some fundamental issues with robustness.

Robustness to Different Sets of Answer Options. We now test the proportion of questions for which the models can correctly answer all variations in the sets of answer options (i.e., the *Single Correct Answer + Variations* setup). This provides an approximation of the worst-case model performance conditioned on the options provided to the model, with the results shown in Table 5.

Across all models, categories as well as output formats, performance in this setup is notably below the previous scores in the basic setup (cf. Table 4). The gap is very substantial and this holds even for the *conditional counting* category, where the correct answer is always the same number, which only gets expressed with different templates (e.g., 3 versus *there are 3*, see some examples in Figure 1 again). For the other scenarios, the different variations can also provide different correct answers (e.g., a position relative to different objects in an image for the *Ordering* category), which also adds to task complexity. Absolute scores and gaps, as expected, also depend on the model and the chosen output format

	Count			Order			Culture			Trick		
	GEN	LOG	JSON									
GPT-4	46.9 _{-1.5}	15.2 _{-22.4}	18.4 _{-29.2}	26.0 _{-22.4}	23.2 _{-30.0}	31.6 _{-21.6}	67.3 _{-15.7}	66.7 _{-16.3}	83.7 _{-4.2}	34.1 _{-20.9}	38.4 _{-18.3}	41.5 _{-14.8}
Gemini	36.4 _{-20.8}	-	36.0 _{-23.6}	53.2 _{-24.9}	-	60.7 _{-17.8}	41.1 _{-34.2}	-	76.6 _{-11.3}	38.8 _{-31.1}	-	61.1 _{-14.9}
Llava 1.6	16.8 _{-32.4}	11.6 _{-33.6}	15.6 _{-29.2}	23.6 _{-30.4}	21.5 _{-25.8}	17.7 _{-19.4}	63.5 _{-7.8}	67.4 _{-8.3}	37.6 _{-18.5}	23.6 _{-30.1}	34.5 _{-20.5}	30.4 _{-11.1}
Idefics2	18.8 _{-26.0}	18.8 _{-26.0}	0.0 _{-0.0}	31.0 _{-28.5}	32.9 _{-25.7}	0.0 _{-0.0}	51.7 _{-32.7}	83.1 _{-1.7}	0.0 _{-0.0}	37.5 _{-22.3}	41.0 _{-16.2}	0.0 _{-0.0}

Table 5: Accuracy when models answer all variations on sets of answer options for the same question in the single-correct answer setup (i.e., *Single-Correct + Variations* from §3.2) correctly, indicating worst-case performance and by default being lower or equal to the scores in the single-correct setup without variations. All the scores are reported with the default prompt, without any prompt variation. We also report the difference to the main results in the basic single-correct setup as the subscript, again with the same prompt.

	0			1			2			3			4			F1			
	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	
Count	GPT 4	19.6	35.5	11.6	41.9	0.7	38.1	37.0	16.1	51.6	33.1	40.2	54.2	10.0	0.0	31.3	58.4	57.8	62.3
	Gemini	28.6	-	14.0	48.5	-	51.6	54.6	-	43.6	51.9	-	86.0	65.4	-	93.9	74.9	-	82.1
	LLaVA 1.6	0.0	0.0	66.0	29.3	39.3	9.3	0.7	0.0	2.7	0.0	0.0	0.0	2.6	25.0	0.0	41.8	55.3	14.6
	Idefics 2	0.0	0.0	100.0	42.0	40.7	0.0	0.0	1.3	0.0	2.0	6.8	0.0	4.6	28.1	0.0	41.9	60.2	0.0
Order	GPT 4	16.7	22.7	18.7	33.3	2.0	36.7	20.7	4.0	24.0	5.3	24.0	8.7	0.0	0.7	2.7	55.3	59.6	58.0
	Gemini	24.0	-	14.0	50.0	-	56.0	39.3	-	46.0	10.0	-	25.3	14.7	-	11.3	66.1	-	72.8
	LLaVA 1.6	0.0	0.0	80.7	30.7	44.7	9.3	0.0	0.0	1.3	0.0	0.0	0.0	1.3	18.7	0.0	38.8	49.6	13.9
	Idefics 2	0.0	0.0	100.0	52.0	40.7	0.0	2.7	6.0	0.0	0.0	7.3	0.0	0.0	22.0	0.0	47.2	56.4	0.0
Culture	GPT 4	54.2	62.2	35.8	55.9	21.4	67.9	42.7	20.2	54.9	36.4	53.3	41.0	19.0	5.1	17.7	71.0	71.5	78.1
	Gemini	56.9	-	20.8	60.2	-	75.6	43.5	-	52.9	39.0	-	62.9	28.4	-	53.4	78.7	-	84.9
	LLaVA 1.6	0.0	0.0	77.1	66.2	73.1	8.3	0.6	0.0	5.7	0.0	14.8	0.0	15.1	36.4	0.0	51.2	61.4	14.7
	Idefics 2	11.3	2.7	100.0	72.5	63.0	0.0	5.7	14.4	0.0	4.5	21.1	0.0	6.9	30.7	0.0	51.1	66.3	0.0
Trick	GPT 4	21.7	30.7	11.5	38.0	6.0	44.7	35.2	13.2	34.6	12.7	24.0	24.7	8.7	1.3	10.7	63.1	59.8	61.1
	Gemini	30.5	-	2.7	46.0	-	56.0	32.0	-	41.7	20.0	-	36.7	12.0	-	28.7	64.3	-	72.1
	LLaVA 1.6	0.0	0.0	71.4	44.0	46.0	10.0	1.3	0.0	3.3	0.0	0.0	0.0	2.0	26.0	0.0	42.2	51.1	14.6
	Idefics 2	0.7	0.0	100.0	62.7	48.7	0.0	4.6	4.6	0.0	2.7	8.7	0.0	2.0	23.3	0.0	48.0	58.8	0.0

Table 6: Accuracy scores in the Multi-Correct Setup, showing the proportion of questions for which the models detected all the answers correctly, split over the questions based on the number of gold correct answers (from 0 to 4); F1 scores are also reported to account for partially correct answers as well. We omit standard deviation over 3 prompts for clarity of presentation: for that, see Table 8.

(e.g., Idefics2 has a smaller performance drop with OUT-LOG than with OUT-GEN).

For the *Culture* category, the performance gap is smaller and for specific configurations such as Idefics2 with OUT-LOG the gap is even less than 2%. However, this does not hold in general, but only for some models and configurations, e.g., Idefics2 with OUT-GEN does show a substantial gap. This observation further highlights the importance of evaluating robustness along with performance, to ensure that VL understanding and reasoning capabilities are not limited only to specific scenarios, evaluation protocols, or task instances.

Multi-Correct Answer Setup. This setup requires the model to reason separately over the answer options as the number of correct answer

is not limited to a single option, and can also be zero. The results, split over groups of questions based on the number of correct answers expected, are provided Table 6, with standard deviation due to varying the prompt provided in Table 8 in Appendix C.

First, we note that the results for the questions with one correct answer (i.e., comparable to the previous single-correct setup) are now lower in this, more challenging setup as the model lacks the prior information on the number of correct answers.

Looking into performance of specific models, Gemini Flash seems to perform the best overall, followed by GPT-4. However, a closer inspection again reveals major problems with robustness, as peak absolute scores are achieved only with specific configurations and can vary considerably

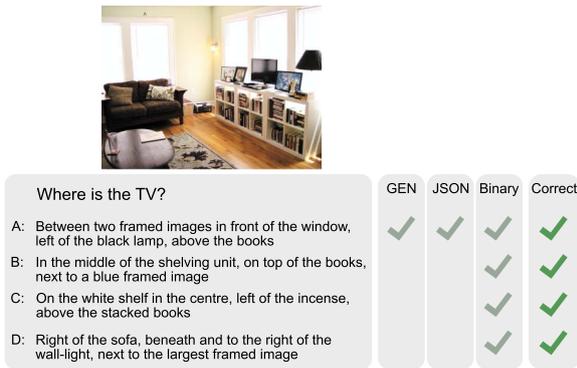


Figure 2: An example question from the *order* category that Gemini 1.5 Flash answers incorrectly with the OUT-GEN and the OUT-JSON format with all options presented in a single prompt, but gets correctly when framed as a binary task for each individual option.

over different output formats and question types. For instance, OUT-LOG substantially outperforms OUT-GEN. For OUT-LOG, the structure of the answer is integrated with the evaluation protocol, which might positively impact performance. Furthermore, for some models in some categories, OUT-JSON yields the highest absolute scores (e.g., both Gemini Flash and GPT-4 display the highest F-1 scores with OUT-JSON). We also observe that LLaVA 1.6 and Idefics2 struggle with questions that do not have exactly one correct answer. We hypothesize that this is due to their pretraining bias, where they are skewed towards single-correct setups.

We further explore different, alternative setups for the multi-correct questions, aiming to address some of the previously detected gaps in multi-correct setups. First, adding ‘All’ and ‘None of the Above’ directly as extra answer options can mostly help with questions that comprise 0 or 4 correct answers; however, the models remain brittle and inconsistent behavior over models and groups of questions has been observed; see Table 9 in the Appendix. Next, we particularly note the results with the so-called *binary decisions* setup: it relies on making ‘local’ decisions per each option. Put simply, VLMs are prompted to determine whether the option is correct as a binary decision for each option individually without information about other options. The results of this variant are provided in Table 7 in the Appendix. While independently querying the VLM for each option increases inference costs, it also improves performance across the board, and can reduce the

bias to pick one correct option (see Figure 2 for an example).

Importantly, even when the gains with this binary-decisions setup are achieved, there is still ample room for improvement with all the VLMs in the multi-correct setup. On average, for *counting*, *ordering* or *culture* less than a third of the questions is completely solved by the state-of-the-art VLMs.

Robustness to Different Prompts. We measure this by the standard deviation of the scores across three prompts, both in single-correct (see Table 4 again) as well as in multi-correct setups (see Table 8 in Appendix C). In the single-correct setup, the differences in accuracy due to prompt variation range from 0 to 20.4, depending on other factors (i.e., the underlying model, category, output format). Gemini Flash, while reaching the highest average accuracy in most categories, at the same time tends to have a higher standard deviation than the other VLMs and is thus less robust to the chosen prompt. The highest standard deviation in this setup is with LLaVA 1.6 on *conditional counting* with OUT-JSON. Standard deviations seem even higher in the multi-correct setup (we relied on the ‘global’ multiple-choice approach), again indicating that the scores are very volatile, and the VLMs seem even less robust to prompt variation here.

Robustness to Output Formats. Further, based on the results from Tables 4–6, we observe that the performance of the VLMs varies substantially conditioned on the chosen evaluation protocol and the output format. As indicated before, Gemini Flash performs best with a JSON-formatted output (i.e., the OUT-JSON variant), and it also suffers from inconsistent formatting with OUT-GEN. Therefore, in order to investigate to what extent the detected lower performance is due to answer formatting not picked up by the manually defined post-processing regular expression (see §3.2), we randomly sample 50 answers that are labeled as incorrect, and find that only 3/50 of the questions are false negatives.

On the other hand, LLaVA’s performance with OUT-JSON drops notably compared to other evaluation protocols and formats, while Idefics2 is completely unable to provide JSON-formatted answers. This indicates a bias picked up during training of these VLMs. Namely, a general-purpose

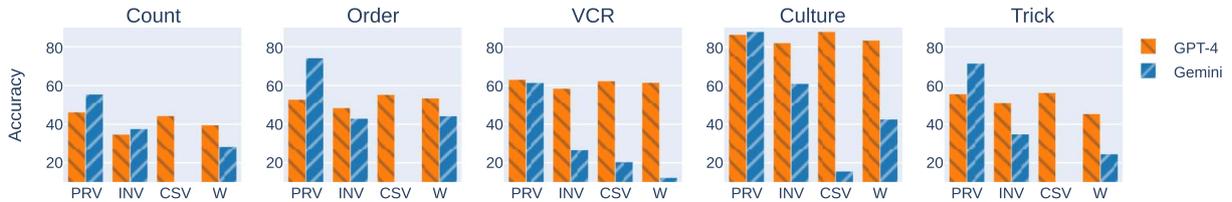


Figure 3: Evaluation across additional output formats (OUT- $\{\text{INV}, \text{CSV}, \text{W}=\text{WORD}\}$, see §3.2) in the basic single-correct answer setup without variations. We do not report the results with LLaVA 1.6 and Idefics2 as they are straight zeros in all the evaluation runs for the three additional formats. The PRV column refers to the peak score for each model and category with any of the previously tested formats (OUT- $\{\text{GEN}, \text{LOG}, \text{JSON}\}$ from Table 4).

VLM with a true understanding of the prompt should be able to provide the answer in any format specified. However, LLaVA 1.6 and Idefics2 have been instruction-tuned to provide the letter indicating the correct answer option. This explains why they perform best with OUT-GEN in most cases.

Probing Additional Output Formats. In order to further explore the robustness and flexibility of the VLMs to different output formats, we investigate three additional formats within the simplest, single-correct setup without variations: **1)** we ‘inversely’ instruct the VLMs to generate a list of incorrect answers instead of a correct answer (labeled OUT-INV); **2)** similar to OUT-JSON, we expect answers as valid CSV-formatted output (labeled OUT-CSV); **3)** to further ‘stress-test’ semantic understanding of the provided instructions, we propose a simple transformation where we expect the VLM to output any word from the English vocabulary that starts with the letter of the correct answer (OUT-WORD). The results with these additional output formats are summarised in Figure 3.

LLaVA 1.6 and Idefics2 are again completely unable to follow these output formats, which aligns with the previous observation concerning the JSON-formatted output. Gemini Flash struggles with the CSV format, whereas GPT-4 reaches peak scores with that format for 3/5 categories. We further observe substantial drops with Gemini Flash with the other two formats, while GPT-4 seems most robust when faced with these three output formats, but again with drops compared to the previous three formats (the PRV column) from Table 4, except for CSV. In sum, these results further point to general robustness and instruction-following issues with all the VLMs in our evaluation and again emphasise the gap to expected human-level performance.

4.2 Further Discussion and Analyses

As one general finding, we observe that different models perform well on different categories and in different configurations. Whereas Idefics2 shows less variation across different prompts, it struggles on the other three robustness evaluations. GPT-4 seems as the most robust model when varying the output format, while Gemini Flash outperforms other models in general in the multi-correct setup. Put simply, there is no general-purpose or ‘one-size-fits-all’ solution, and all the models suffer from issues with robustness and flexibility, questioning their general semantic understanding.

Furthermore, we observe a gap between the API-gated models GPT-4 and Gemini Flash and the two open models. The former can (at least to some extent) handle different output formats and the multiple-correct setup substantially better than the latter. This hints that current open-weights models are still limited by tasks and output formats they observe during pretraining.

Tracing Progress of VLMs. To assess whether performance on the challenging questions from DARE improves over time and over newer VLM checkpoints, which would hint that the VLMs are indeed gradually improving, we evaluate different checkpoints of the same model in the single-correct answer setup (relying on the OUT-GEN format, similar trends have been observed with the other output formats). The results are plotted in Figure 4, with further details on the evaluated checkpoints available in Appendix A.4. There are some trends visible across all categories. Every model family shows improvements in the majority of categories, with only small or anecdotal drops (e.g., GPT-4 in VCR, LLaVA in the *culture* category). LLaVA displays the smallest improvement while Idefics shows the steepest improvement; however, Idefics also starts from

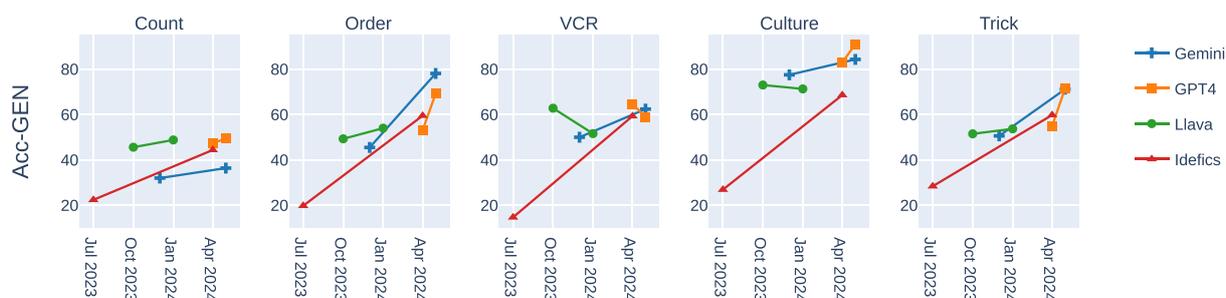


Figure 4: Tracing progress on the five categories of DARE. We assess (i) Gemini 1 and Gemini Flash 1.5; (ii) GPT-4 and GPT-4o, (iii) LLaVA 1.5 and LLaVA 1.6; (iv) Idefics1 and Idefics2.

the lowest initial results. Gemini Flash 1.5 and GPT-4o show salient gains for the *spatial ordering* and *conditional counting* questions over their predecessors Gemini 1 and GPT4.

These results indicate that DARE with its coverage of different categories, evaluation setups, and challenging scenarios could be used for benchmarking progress (or the lack of it as in the case of LLaVA) of different VLM families in the future: (i) while we detect gains with the newest checkpoints over the previous releases, (ii) there is still a large gap to human-level performance across the board, along with the wide array of robustness issues, as discussed in this paper.

5 Conclusion and Outlook

We have introduced DARE, a novel VQA benchmark for vision-language models (VLMs) which targets image understanding and multi-modal reasoning across five diverse scenarios, offering challenging multiple-choice questions, which were carefully selected and manually annotated, and support evaluations in single-correct and multi-correct answer setups. DARE puts a special focus on *evaluating various robustness aspects of the VLMs*, and it includes variations of the questions and evaluation protocols across several crucial axes (e.g., input prompt, output format, answer options).

Our extensive experiments on DARE highlight that modern VLMs still struggle with the robustness aspects, and their performance varies wildly and depends on multiple factors such as the evaluation setup (single-correct vs multi-correct), category complexity, evaluation protocol, specified output format, etc. VLMs that are constructed from visual encoders aligned to text-only LLMs with visual instruction tuning struggle the most

in our robustness evaluations. Our work highlights the importance of including datasets that target task setups that are not directly covered in pretraining into their construction. In general, motivated by our experiments on tracing progress of VLMs over time (see Figure 4 again), we hope that DARE will help guide future developments of VLMs, with a particular focus on increasing their flexibility to different variations covered in DARE and consequently mitigating their critical issues with robustness.

Acknowledgments

Hannah Sterz thanks the Cambridge Trust for their support via the International Scholarship. This work has been supported by a Royal Society University Research Fellowship ‘*Inclusive and Sustainable Language Technology for a Truly Multilingual World*’ (no 221137) awarded to Ivan Vulić. We thank Aishwarya Kamath and Jeremiah Harmsen for thoughtful comments on initial drafts of the paper.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, 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 Fulford, 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, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L. Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. 2022. Flamingo: A visual language model for few-shot learning. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 – December 9, 2022*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*. <https://doi.org/10.1109/ICCV.2015.279>
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer. https://doi.org/10.1007/978-3-030-58452-8_13
- Xi Chen, Xiao Wang, Soravit Changpinyo, A. J. Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Nan Ding, Keran Rong, Hassan Akbari, Gaurav Mishra, Linting Xue, Ashish V. Thapliyal, James Bradbury, and Weicheng Kuo. 2023. Pali: A jointly-scaled multilingual language-image model. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1–5, 2023*.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Nan Ding, Keran Rong, Hassan Akbari, Gaurav Mishra, Linting Xue, Ashish Thapliyal, James Bradbury, Weicheng Kuo, Mojtaba Seyedhosseini, Chao Jia, Burcu Karagol Ayan, Carlos Riquelme, Andreas Steiner, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. 2022. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.
- Kaustubh Dhole, Varun Gangal, Sebastian Gehrmann, Aadesh Gupta, Zhenhao Li, Saad Mahamood, Abinaya Mahadiran, Simon Mille, Ashish Shrivastava, Samson Tan, Tongshang Wu, Jascha Sohl-Dickstein, Jinho Choi, Eduard Hovy, Ondřej Dušek, Sebastian Ruder, Sajant Anand, Nagender Aneja, Rabin Banjade, Lisa Barthe, Hanna Behnke, Ian Berlot-Attwell, Connor Boyle, Caroline Brun, Marco Antonio Sobrevilla Cabezudo, Samuel Cahyawijaya, Emile Chapuis, Wanxiang Che, Mukund Choudhary, Christian Clauss, Pierre Colombo, Filip Cornell, Gautier Dagan, Mayukh Das, Tanay Dixit, Thomas Dopierre, Paul-Alexis Dray, Suchitra Dubey, Tatiana Ekeinhor,

- Marco Di Giovanni, Tanya Goyal, Rishabh Gupta, Louanes Hamla, Sang Han, Fabrice Harel-Canada, Antoine Honoré, Ishan Jindal, Przemysław Joniak, Denis Kleyko, Venelin Kovatchev, Kalpesh Krishna, Ashutosh Kumar, Stefan Langer, Seungjae Ryan Lee, Corey James Levinson, Hualou Liang, Kaizhao Liang, Zhexiong Liu, Andrey Lukyanenko, Vukosi Marivate, Gerard de Melo, Simon Meoni, Maxine Meyer, Afnan Mir, Nafise Sadat Moosavi, Niklas Meunnighoff, Timothy Sum Hon Mun, Kenton Murray, Marcin Namysl, Maria Obedkova, Priti Oli, Nivranshu Pasricha, Jan Pfister, Richard Plant, Vinay Prabhu, Vasile Pais, Libo Qin, Shahab Raji, Pawan Kumar Rajpoot, Vikas Raunak, Roy Rinberg, Nicholas Roberts, Juan Diego Rodriguez, Claude Roux, Vasconcellos Samus, Ananya Sai, Robin Schmidt, Thomas Scialom, Tshephisho Sefara, Saqib Shamsi, Xudong Shen, Yiwen Shi, Haoyue Shi, Anna Shvets, Nick Siegel, Damien Sileo, Jamie Simon, Chandan Singh, Roman Sitelew, Priyank Soni, Taylor Sorensen, William Soto, Aman Srivastava, Aditya Srivatsa, Tony Sun, Mukund Varma, A. Tabassum, Fiona Tan, Ryan Teehan, Mo Tiwari, Marie Tolkiehn, Athena Wang, Zijian Wang, Zijie Wang, Gloria Wang, Fuxuan Wei, Bryan Wilie, Genta Indra Winata, Xinyu Wu, Witold Wydmanski, Tianbao Xie, Usama Yaseen, Michael Yee, Jing Zhang, and Yue Zhang. 2023. NL-augmenter: A framework for task-sensitive natural language augmentation. *Northern European Journal of Language Technology*, 9. <https://doi.org/10.3384/nejlt.2000-1533.2023.4725>
- Yihong Dong, Xue Jiang, Huanyu Liu, Zhi Jin, Bin Gu, Mengfei Yang, and Ge Li. 2024. Generalization or memorization: Data contamination and trustworthy evaluation for large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12039–12050, Bangkok, Thailand. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-acl.716>
- Markus Frohmann, Igor Sterner, Ivan Vulić, Benjamin Minixhofer, and Markus Schedl. 2024. Segment any text: A universal approach for robust, efficient and adaptable sentence segmentation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11908–11941, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.665>
- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. 2024. BLINK: Multimodal large language models can see but not perceive. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XXIII*, volume 15081 of *Lecture Notes in Computer Science*, pages 148–166. Springer. https://doi.org/10.1007/978-3-031-73337-6_9
- Gemini Team Google. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2017.670>
- Agrim Gupta, Piotr Dollár, and Ross B. Girshick. 2019. LVIS: A dataset for large vocabulary instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*, pages 5356–5364. Computer Vision Foundation / IEEE. <https://doi.org/10.1109/CVPR.2019.00550>
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617. <https://doi.org/10.1109/CVPR.2018.00380>
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*.

- Drew A. Hudson and Christopher D. Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6700–6709. <https://doi.org/10.1109/CVPR.2019.00686>
- Robin Jia and Percy Liang. 2017. Adversarial examples for evaluating reading comprehension systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D17-1215>
- Aishwarya Kamath, Sara Price, Jonas Pfeiffer, Yann LeCun, and Nicolas Carion. 2023. TRICD: Testing robust image understanding through contextual phrase detection. Technical Report.
- Chengzu Li, Caiqi Zhang, Han Zhou, Nigel Collier, Anna Korhonen, and Ivan Vulić. 2024. TopViewRS: Vision-language models as top-view spatial reasoners. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1786–1807, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.106>
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel J. Orr, Lucia Zheng, Mert Yüksesgönül, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. Holistic evaluation of language models. *Transactions on Machine Learning Research*, 2023.
- 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. https://doi.org/10.1007/978-3-319-10602-1_48
- Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2024a. Culturally aware and adapted NLP: A taxonomy and a survey of the state of the art. *arXiv preprint arXiv:2406.03930*.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. Visually grounded reasoning across languages and cultures. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.emnlp-main.818>
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 – 16, 2023*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024c. Mmbench: Is your multi-modal model an all-around player? In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part VI*, volume 15064 of *Lecture Notes in Computer Science*, pages 216–233. Springer. https://doi.org/10.1007/978-3-031-72658-3_13
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. Fantastically ordered prompts and where to find them:

- Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.556>
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and C. V. Jawahar. 2022. InfographicVQA. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1697–1706. <https://doi.org/10.1109/WACV51458.2022.00264>
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. DocVQA: A dataset for VQA on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2200–2209. <https://doi.org/10.1109/WACV48630.2021.00225>
- Shravan Nayak, Kanishk Jain, Rabiul Awal, Siva Reddy, Sjoerd Van Steenkiste, Lisa Anne Hendricks, Karolina Stanczak, and Aishwarya Agrawal. 2024. Benchmarking vision language models for cultural understanding. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5769–5790, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.329>
- 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. 2024. DOCCI: Descriptions of connected and contrasting images. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LX*, volume 15118 of *Lecture Notes in Computer Science*, pages 291–309. Springer. https://doi.org/10.1007/978-3-031-73027-6_17
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024.
- Pouya Pezeshkpour and Estevam Hruschka. 2024. Large language models sensitivity to the order of options in multiple-choice questions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017, Mexico City, Mexico. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-naacl.130>
- 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. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- David Romero, Chenyang Lyu, Haryo Akbarianto Wibowo, Teresa Lynn, Injy Hamed, Aditya Nanda Kishore, Aishik Mandal, Alina Dragonetti, Artem Abzaliev, Atnafu Lambebo Tonja, Bontu Fufa Balcha, Chenxi Whitehouse, Christian Salamea, Dan John Velasco, David Ifeoluwa Adelani, David Le Meur, Emilio Villa-Cueva, Fajri Koto, Fauzan Farooqui, Frederico Belcavello, Ganzorig Batnasan, Gisela Vallejo, Grainne Caulfield, Guido Ivetta, Haiyue Song, Henok Biadgign Ademtew, Hernán Maina, Holy Lovenia, Israel Abebe Azime, Jan Christian Blaise Cruz, Jay Gala, Jiahui Geng, Jesus-German Ortiz-Barajas, Jinheon Baek, Jocelyn Dunstan, Laura Alonso Alemany, Kumaranage Ravindu Yasas Nagasinghe, Luciana Benotti, Luis Fernando D’Haro, Marcelo Viridiano, Marcos Esteche-Garitagaitia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Joutiteau, Mihail Mihaylov, Naome Etori, Mohamed Fazli Mohamed Imam, Muhammad Farid Adilazuarda, Munkhjargal Gochoo, Munkh-Erdene Otgonbold, Olivier Niyomugisha,

- Paula Mónica Silva, Pranjal Chitale, Raj Dabre, Rendi Chevi, Ruochen Zhang, Ryandito Diandaru, Samuel Cahyawijaya, Santiago Góngora, Soyeong Jeong, Sukannya Purkayastha, Tatsuki Kuribayashi, Teresa Clifford, Thanmay Jayakumar, Tiago Timponi Torrent, Toqeer Ehsan, Vladimir Araujo, Yova Kementchedjhieva, Zara Burzo, Zheng Wei Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Tamar Solorio, and Alham Fikri Aji. 2024. Cvqa: Culturally-diverse multilingual visual question answering benchmark. In *Advances in Neural Information Processing Systems*, volume 37, pages 11479–11505. Curran Associates, Inc.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying language models’ sensitivity to spurious features in prompt design or: How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*.
- Amanpreet Singh, Vivek Natarjan, Meet Shah, Yu Jiang, Xinlei Chen, Devi Parikh, and Marcus Rohrbach. 2019. Towards VQA models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8317–8326. <https://doi.org/10.1109/CVPR.2019.00851>
- Peng Tang, Srikar Appalaraju, R. Manmatha, Yusheng Xie, and Vijay Mahadevan. 2024. Multiple-question multiple-answer text-VQA. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 73–88, Mexico City, Mexico. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.naacl-industry.7>
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022*, pages 5228–5238. IEEE. <https://doi.org/10.1109/CVPR52688.2022.00517>
- 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16–22, 2024*, pages 9568–9578. IEEE. <https://doi.org/10.1109/CVPR52733.2024.00914>
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024. Large language models are not fair evaluators. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11–16, 2024*, pages 9440–9450. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-long.511>
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 – December 9, 2022*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2019.00688>
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*.

A Evaluation: Technical Details and Reproducibility

A.1 Prompts

VLMs results heavily depend on the used prompt: For reproducibility, we report all prompts used during evaluation. For the **single-correct answer** setup we use the following prompt for OUT-GEN and OUT-LOG:

Main prompt for multiple-choice questions with a single correct answer (OUT-GEN and OUT-LOG)

The following are multiple choice questions about <QUESTION TYPE>. You should directly answer the question by choosing the correct option given the image and the question. Give only the letter indicating the correct answer e.g. "A"

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answer:

Prompt 2 for multiple-choice questions with a single correct answer (OUT-GEN and OUT-LOG)

Imagine you are a student in an exam consisting of the following multiple choice question about <QUESTION TYPE>. [You need to do well in this exam in order to not fail the class.] Provide the answer by choosing the correct option, e.g. 'B'.

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answer:

Prompt 3 for multiple-choice questions with a single correct answer (OUT-GEN and OUT-LOG)

You will be presented with a multiple choice question about <QUESTION TYPE>. Please return the answer as the letter corresponding to the correct option e.g. 'A'.

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answer:

To get the VLM to return JSON-formatted answers we need to specify that requirement in the prompt. We use the following prompts:

Prompt 1 for multiple-choice questions with a single correct answer and output in json-format (OUT-JSON)

The following are multiple choice questions about <QUESTION TYPE>. You should directly answer the question by choosing the correct option given the image and the question. Provide the answer in json format e.g. {"answer": "A"}.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answer:

Prompt 2 for multiple-choice questions with a single correct answer and output in json-format (OUT-JSON)

Imagine you are a student in an exam consisting of the following multiple choice question about <QUESTION TYPE>. [You need to do well in this exam in order to not fail the class.] Provide the answer by choosing the correct option as json e.g. {"answer": "B"}.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answer:

Prompt 3 for multiple-choice questions with a single correct answer and output in json-format (OUT-JSON)

You will be presented with a multiple choice question about <QUESTION TYPE>. Please return the answer as json for instance {"answer": "D"}.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answer:

For the **multi-correct** setup the prompt needs to specify that there can be a varying number of correct answer options. We use the following prompts:

Prompt 1 for multiple-choice questions with multiple correct answers (OUT-GEN and OUT-LOG)

The following are multiple choice questions about <QUESTION TYPE>. You should directly answer the question by choosing the correct options given the image and the question. There can be zero to four correct answers. If no answer is correct answer NONE otherwise provide a list of the correct answer options.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answers:

Prompt 2 for multiple-choice questions with a multiple correct answer (OUT-GEN and OUT-LOG)

Imagine you are a student in an exam consisting of the following multiple choice question about <QUESTION TYPE>. [You need to do well in this exam in order to not fail the class.] Provide the answer by choosing the correct option. Your teacher decided to make this exam extra hard. There can be zero to four correct answers. If there is no correct answer, answer with NONE otherwise return the list of correct options e.g. "A, C".

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answers:

Prompt 3 for multiple-choice questions with a multiple correct answer (OUT-GEN and OUT-LOG)

You will be presented with a multiple choice question about <QUESTION TYPE>. Provide the correct answer given the question and the image as the number corresponding to the answer. Attention: There can be an arbitrary number of correct options. Return all that you identify as correct, e.g. "B, D". If no option is correct answer with NONE.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answers:

The JSON-formatted output again requires a slightly different prompt:

Prompt 1 for multiple-choice questions with a multiple correct answer with output in json-format (OUT-JSON)

The following are multiple choice questions about <QUESTION TYPE>. You should directly answer the question by choosing the correct options given the image and the question. There can be zero to four correct answers. Provide the answer in json format e.g. {"answers": ["A", "B"]} or if there is no correct answer {"answers": []}.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answers:

Prompt 2 for multiple-choice questions with a multiple correct answer with output in json-format (OUT-JSON)

Imagine you are a student in an exam consisting of the following multiple choice question about <QUESTION TYPE>. [You need to do well in this exam in order to not fail the class.] Provide the answer by choosing the correct option. Your teacher decided to make this exam extra hard. There can be zero to four correct answers. There can be zero to four correct answers. Provide the answer by choosing the correct option as json e.g. {"answer": ["B", "D"]}.

Question: <QUESTION>

Options:

- A. <ANSWER A>
- B. <ANSWER B>
- C. <ANSWER C>
- D. <ANSWER D>

Answers:

Prompt 3 for multiple-choice questions with a multiple correct answer with output in JSON-format (OUT-JSON)

You will be presented with a multiple choice question about <QUESTION TYPE>. 0-4 answers can be correct. Please return the answer as json, for instance {"answer": ["A", "D"]}

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answers:

Prompt for multiple-choice questions with a single correct answer with output in inverse-format (OUT-INV)

The following is a question about <QUESTION TYPE>. Provide only the three incorrect answers to the question as the list of the corresponding letters e.g. 'A, C, D'.

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answers:

Prompt for multiple-choice questions with a single correct answer with output in csv-format (OUT-CSV)

The following are multiple choice questions about <QUESTION TYPE>. You should directly answer the question by choosing the correct option given the image and the question. Provide the answer in csv format e.g. 'answer, A'.

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answers:

Prompt 3 for multiple-choice questions with a single correct answer with output in word-format (OUT-WORD)

The following are multiple choice questions about <QUESTION TYPE>. You should answer the question by giving the answer as a word that starts with the letter of corresponding to the correct option e.g. if A is correct 'Apple'. Your answer has to start with that word.

Question: <QUESTION>

Options:

A. <ANSWER A>

B. <ANSWER B>

C. <ANSWER C>

D. <ANSWER D>

Answer:

A.2 Regex for Extracting Answers

To extract the correct answer from the text we use regular expressions. The generated text is parsed for: $(\hat{|\backslash s}[ABCD](, [ABCD])+(\$|\backslash n|,|\backslash.|\backslash s)$ to find a list of letters. To find single-letter answers the expression is parsed for: $(\langle\text{OPT}\rangle:)|(\langle\text{OPT}\rangle\backslash.)|(\langle\text{OPT}\rangle\backslash n)|(\langle\text{OPT}\rangle\$)$ for all answer options: A, B, C, and D, where $\langle\text{OPT}\rangle$ refers to the four options.

A.3 Models & Hyper-Parameters

GPT-4 (Achiam et al., 2023), the API-gated model by OpenAI, supports text as well as image input. We use the *gpt-4-turbo-2024-04-09* variant, unless stated otherwise.

Gemini (Gemini Team Google, 2023) is another API-gated model, created by Google, which supports textual and visual input. We use *gemini-1.5-flash*, a lightweight variant targeting speed and efficiency.

Idefics2 is an open-weights model that aligns OpenCLIP (Radford et al., 2021) with Llama (Touvron et al., 2023) with a Vision Language Connector. We use *HuggingFaceM4/idefics2-8b*.

LLaVA (Liu et al., 2024b) is an open-weights model that aligns the CLIP ViT-L/14 (Radford et al., 2021) embeddings with the Vicuna LLM (Chiang et al., 2023) via a projection of the image embeddings. This allows the model to handle text and image input jointly. We use *llava-hf/llava-v1.6-vicuna-7b-hf*.

For all the models, we opt for their default, suggested hyper-parameters: e.g., temperature is set to 0 with GPT-4 and Gemini, while we use the default generation configuration of Idefics2 and LLaVa, corresponding to greedy decoding.

A.4 Model Progress

With frequent new versions of each model family, an interesting question is whether new versions already improve the score on our benchmark. To investigate this, we report the score of a previous or more recent model in the same family for the *conditional counting* and *order* categories. The development is illustrated in Figure 4. We use *gemini-1.0-pro-001* for Gemini 1.0, *llava-hf/llava-1.5-7b-hf* for LLaVA 1.5, *HuggingFaceM4/idefics-9b* for Idefics 1 and *gpt-4o-2024-05-13* for GPT-4o. For the date of the model, we use the date of the first commit of

Question Answering with Conditional Counting

As humans, we can easily identify and filter objects in images. It is easy to count all dogs sitting on a couch or all yellow mugs. We want to create a benchmark measuring the ability of vision and language models, such as GPT-4, to perform such tasks.

Task

You will be presented with an image. Your task is to create questions that require to count and filter the objects in the image. The questions can include:

- colours of objects
- the position of objects e.g. on the couch
- the normal usage of objects e.g. furniture intended for sitting
- grouping several objects by a hypernym e.g. all musical instruments

Make sure that the question:

- requires to apply **at least one filtering condition** to the image (e.g. which windows are open, which dogs are on the couch, ...)
- use different types of filtering conditions. Don't use colour (or any other type) in all your questions. Be creative!
- has only one correct answer

It is important that the questions are **easy to answer for a human**.

So frame the question in an unambiguous way that only allows one correct answer!

Examples



Question: How many windows are open?

Correct answer: 3

Explanation: The question requires not just counting all the windows, but also identifying the subset of windows that is open and counting that.



Question: How many horses are shown in the mirror?

Correct answer: 2

Explanation: It is not enough to count the horses in the image. Instead, the question requires the additional step to identify which horses are in the mirror.

Figure 5: Guidelines for the *conditional counting* category.

the model to HuggingFace for models available on the HuggingFace hub. For API-gated models we use the dates specified in their documentation.

B Annotation Guidelines

We use guidelines customised for each category in DARE, with specific instructions and examples, to ensure that annotators pay attention to the most important aspects associated with each category. To illustrate the structure of the guidelines we report the guidelines for two of the categories. Figure 5 shows our guidelines for *conditional counting*.

C Additional Results

		0	1	2	3	4	F1
Counting	GPT 4	86 ₊₅₀	16 ₋₂₆	8 ₋₄₄	16 ₋₃₈	20 ₋₁₁	37.9 ₋₂₄
	Gemini	13 ₋₁₁	18 ₋₃₈	14 ₋₃₂	50 ₊₂₅	65 ₊₅₁	69.3 ₋₃
	LLaVA 1.6	0 ₋₆₆	0 ₋₃₉	6 ₊₃	24 ₊₂₄	92 ₊₆₇	69.6 ₊₁₅
	Idefics2	32 ₋₆₈	16 ₋₂₆	20 ₊₁₉	24 ₊₁₇	44 ₊₁₆	64.4 ₊₄
Order	GPT 4	72 ₋₄₉	28 ₋₁₁	16 ₋₈	12 ₊₅	12 ₊₉	56.2 ₋₃
	Gemini	44 ₊₂₀	40 ₋₁₆	54 ₊₈	50 ₊₂₅	44 ₊₂₉	80.2 ₊₇
	LLaVA 1.6	20 ₋₆₁	18 ₋₂₇	26 ₊₂₅	32 ₊₃₂	60 ₊₄₁	72.1 ₊₂₃
	Idefics2	24 ₋₇₆	10 ₋₄₂	22 ₊₁₆	32 ₊₂₅	42 ₊₂₀	69.1 ₊₁₃
Culture	GPT 4	94 ₊₃₂	58 ₋₁₀	36 ₋₁₀	29 ₋₂₄	32 ₊₁₅	83.0 ₊₅
	Gemini	84 ₊₂₇	63 ₋₁₃	55 ₊₂	50 ₋₁₃	64 ₊₂₈	82.9 ₋₂
	LlaVA 1.6	45 ₋₃₂	39 ₋₃₄	21 ₊₁₃	35 ₊₂₀	75 ₊₃₉	64.1 ₊₃
	Idefics2	80 ₋₂₀	37 ₋₃₆	19 ₊₅	25 ₊₄	36 ₊₅	64.1 ₋₂
Trick	GPT 4	88 ₊₅₆	20 ₋₂₅	20 ₋₁₅	10 ₋₁₅	16 ₊₅	47.4 ₋₁₆
	Gemini	59 ₊₂₈	64 ₊₈	35 ₋₇	38 ₋₀	38 ₊₉	74.6 ₊₃
	LLaVA 1.6	22 ₋₄₉	14 ₋₃₂	16 ₊₁₃	14 ₊₁₂	66 ₊₄₀	68.3 ₊₁₇
	Idefics2	61 ₋₃₉	24 ₋₃₉	20 ₊₁₅	10 ₊₁₀	20 ₋₃	57.3 ₋₂

Table 7: Results in the multi-correct setup where each answer option is considered ‘locally’ with the question, that is, individually with only binary outcome possible (correct or incorrect). Cf. Table 6 for the results with the global multiple-choice approach. The numbers in the subscript are deltas to the best corresponding result as reported in Table 6.

		0			1			2			3			4			F1		
		GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON	GEN	LOG	JSON
Counting	GPT 4	19 \pm 7.3	36 \pm 16.2	12 \pm 16.6	42 \pm 3.7	1 \pm 1.2	38 \pm 15.3	37 \pm 18.2	16 \pm 1.5	52 \pm 21.7	33 \pm 32.1	40 \pm 32.3	54 \pm 36.2	10 \pm 17.3	0 \pm 0.0	31 \pm 23.6	58.4 \pm 2.0	57.8 \pm 6.3	62.3 \pm 24.4
	Gemini	24 \pm 7.3	-	14 \pm 22.5	50 \pm 11.1	-	56 \pm 7.2	39 \pm 12.1	-	46 \pm 6.0	10 \pm 8.7	-	25 \pm 2.3	14 \pm 22.0	-	11 \pm 1.2	66.1 \pm 5.2	-	72.8 \pm 4.0
	Llava 1.6	0 \pm 0.0	0 \pm 0.0	66 \pm 7.2	29 \pm 4.6	39 \pm 9.0	9 \pm 16.2	1 \pm 1.2	0 \pm 0.0	3 \pm 4.6	0 \pm 0.0	0 \pm 0.0	0 \pm 0.0	3 \pm 4.5	25 \pm 25.5	0 \pm 0.0	41.8 \pm 9.5	55.3 \pm 3.1	14.6 \pm 25.3
	Idefics 2	0 \pm 0.0	0 \pm 0.0	100 \pm 0.0	42 \pm 2.0	41 \pm 4.2	0 \pm 0.0	0 \pm 0.0	1 \pm 3.5	0 \pm 0.0	2 \pm 3.5	7 \pm 1.2	0 \pm 0.0	0 \pm 0.0	28 \pm 29.5	0 \pm 0.0	41.9 \pm 5.7	60.2 \pm 5.0	0 \pm 0.0
Order	GPT 4	17 \pm 12.7	23 \pm 16.0	19 \pm 19.4	33 \pm 5.0	2 \pm 2.0	37 \pm 8.3	21 \pm 3.1	4 \pm 2.0	24 \pm 6.0	5 \pm 1.2	24 \pm 0.0	9 \pm 2.3	0 \pm 0.0	1 \pm 1.2	3 \pm 3.1	55 \pm 3.3	60 \pm 0.9	58 \pm 5.4
	Gemini	24 \pm 12.5	-	14 \pm 22.5	50 \pm 11.1	-	56 \pm 7.2	39 \pm 12.1	-	46 \pm 6.0	10 \pm 8.7	-	25 \pm 2.3	15 \pm 22.0	-	11 \pm 1.2	66 \pm 5.2	-	73 \pm 0.4
	Llava 1.6	0 \pm 0.0	0 \pm 0.0	81 \pm 28.4	31 \pm 3.1	45 \pm 10.1	9 \pm 16.2	0 \pm 0.0	0 \pm 0.0	1 \pm 2.3	0 \pm 0.0	0 \pm 0.0	0 \pm 0.0	1 \pm 1.2	19 \pm 17.2	0 \pm 0.0	39 \pm 3.1	50 \pm 7.3	14 \pm 22.7
	Idefics 2	0 \pm 12.7	0 \pm 16.0	100 \pm 19.4	52 \pm 5.0	41 \pm 2.0	0 \pm 8.3	3 \pm 3.1	6 \pm 2.0	0 \pm 6.0	0 \pm 1.2	7 \pm 0.0	0 \pm 2.3	0 \pm 0.0	22 \pm 11.2	0 \pm 3.1	47 \pm 3.3	56 \pm 0.9	0 \pm 5.4
Culture	GPT 4	54 \pm 8.2	62 \pm 20.5	36 \pm 15.3	56 \pm 7.2	21 \pm 7.7	68 \pm 3.9	43 \pm 5.8	20 \pm 6.6	55 \pm 8.7	36 \pm 6.6	53 \pm 3.0	41 \pm 9.0	19 \pm 5.3	5 \pm 1.0	18 \pm 7.3	71 \pm 6.9	72 \pm 5.5	78 \pm 3.3
	Gemini	57 \pm 23.1	-	21 \pm 14.1	60 \pm 12.1	-	76 \pm 6.8	44 \pm 21.4	-	53 \pm 1.9	39 \pm 12.3	-	63 \pm 3.9	28 \pm 18.4	-	53 \pm 4.8	79 \pm 5.3	-	85 \pm 0.8
	Llava 1.6	0 \pm 0.0	0 \pm 0.0	77 \pm 39.6	66 \pm 4.3	73 \pm 8.4	8 \pm 14.4	1 \pm 1.1	0 \pm 0.0	6 \pm 9.8	0 \pm 0.0	15 \pm 20.7	0 \pm 0.0	15 \pm 1.9	36 \pm 27.4	0 \pm 0.0	51 \pm 1.7	61 \pm 8.6	15 \pm 25.4
	Idefics 2	11 \pm 8.2	3 \pm 20.5	100 \pm 15.3	73 \pm 7.2	63 \pm 7.7	0 \pm 3.9	6 \pm 5.8	14 \pm 6.6	0 \pm 8.7	4 \pm 6.6	21 \pm 3.0	0 \pm 9.0	7 \pm 5.3	31 \pm 1.0	0 \pm 7.3	51 \pm 6.9	66 \pm 5.5	0 \pm 3.3
Trick	GPT 4	22 \pm 6.7	31 \pm 9.3	12 \pm 5.9	38 \pm 2.0	6 \pm 4.0	45 \pm 11.7	35 \pm 7.0	13 \pm 4.7	35 \pm 2.9	13 \pm 4.6	24 \pm 6.9	25 \pm 10.1	9 \pm 2.3	1 \pm 1.2	11 \pm 11.7	63 \pm 4.9	60 \pm 4.9	61 \pm 2.3
	Gemini	30 \pm 31.9	-	3 \pm 1.6	46 \pm 10.4	-	56 \pm 5.3	32 \pm 21.2	-	42 \pm 4.2	20 \pm 9.2	-	37 \pm 4.6	12 \pm 8.7	-	29 \pm 5.0	64 \pm 10.4	-	72 \pm 0.7
	Llava 1.6	0 \pm 0.0	0 \pm 0.0	71 \pm 47.8	44 \pm 5.3	46 \pm 5.3	10 \pm 12.5	1 \pm 2.3	0 \pm 0.0	3 \pm 5.7	0 \pm 0.0	0 \pm 0.0	0 \pm 0.0	2 \pm 3.5	26 \pm 22.3	0 \pm 0.0	42 \pm 2.8	51 \pm 7.9	13 \pm 23.3
	Idefics 2	1 \pm 6.7	0 \pm 9.3	100 \pm 5.9	63 \pm 2.0	49 \pm 4.0	0 \pm 11.7	5 \pm 7.0	5 \pm 4.7	0 \pm 2.9	3 \pm 4.6	9 \pm 6.9	0 \pm 10.1	2 \pm 2.3	23 \pm 1.2	0 \pm 11.7	48 \pm 4.7	59 \pm 4.9	0 \pm 2.3

Table 8: Accuracy scores in the Multi-Correct Setup, showing the proportion of questions for which the models detected all the answers correctly, split over the questions groups based on the number of gold correct answers (from 0 to 4); F1 scores are also reported to account for partially correct answers as well. Standard deviation due to variation of the prompt (over three different prompts, see Appendix A.1) is also reported in smaller font.

		0	1	2	3	4	F1			0	1	2	3	4	F1
		Counting	GPT-4	20	42	37	33			10	58	GPT-4	54	56	43
+all	4		39	59	60	92	78	+all	41	67	43	40	45	80	
+none	8		39	51	67	35	71	+none	60	69	47	45	20	80	
+all+none	6		35	49	67	43	71	+all+none	49	65	45	46	30	79	
Gemini	24		50	39	10	15	66	Gemini	57	60	44	39	28	79	
+all	2		45	27	13	100	78	+all	51	58	49	44	87	85	
+none	22		51	31	65	84	79	+none	82	57	49	52	42	83	
+all+none	14		53	41	67	86	81	+all+none	82	62	49	58	53	85	
LLaVA 1.6	0		29	1	0	3	42	LLaVA 1.6	0	66	1	0	15	51	
+all	0		34	0	0	4	37	+all	0	60	0	0	19	53	
+none	0		36	0	0	0	36	+none	24	65	0	0	0	45	
+all+none	0		40	0	0	0	36	+all+none	27	67	0	0	6	46	
Idefics 2	0	42	0	2	5	42	Idefics 2	11	73	6	5	7	51		
+all	0	48	0	0	0	40	+all	2	81	0	0	40	63		
+none	2	44	0	0	0	39	+none	35	81	0	0	0	50		
+all+none	2	44	0	0	0	40	+all+none	27	79	0	0	26	60		
Order	GPT-4	17	33	21	5	0	55	GPT-4	22	38	35	13	9	60	
	+all	4	42	22	16	18	63	+all	12	40	41	18	22	65	
	+none	6	36	22	12	6	61	+none	16	44	33	12	8	62	
	+all+none	4	38	22	10	2	61	+all+none	12	46	35	12	10	63	
	Gemini	24	50	39	10	15	66	Gemini	31	46	32	10	12	64	
	+all	14	58	46	20	36	75	+all	16	50	37	20	60	73	
	+none	64	46	54	18	4	71	+none	53	36	37	24	22	70	
	+all+none	54	44	52	20	14	72	+all+none	55	40	31	24	30	69	
	LLaVA 1.6	0	31	0	0	1	39	LLaVA 1.6	0	44	1	0	2	42	
	+all	0	34	0	0	0	37	+all	0	38	0	0	8	42	
	+none	0	30	0	2	0	38	+none	4	42	0	0	0	39	
	+all+none	0	34	0	2	0	38	+all+none	24	46	0	0	6	38	
Idefics 2	0	52	3	0	0	47	Idefics 2	1	63	5	3	2	48		
+all	0	56	0	0	8	48	+all	0	58	0	0	30	55		
+none	6	58	0	0	0	44	+none	13	58	0	0	0	44		
+all+none	2	58	0	0	8	48	+all+none	13	56	0	0	24	53		

Table 9: Accuracy and F1 score of the models with OUT-GEN on the multi-correct setup compared to adding the additional option ‘All of the above’ (+all), ‘None of the above’ (+none), and both (+all+none). This illustrates that explicit options for these cases can improve performance.