

# GRAB: A Challenging Graph Analysis Benchmark for Large Multimodal Models

Jonathan Roberts<sup>1</sup>, Kai Han<sup>2</sup>, Samuel Albanie

<sup>1</sup>University of Cambridge, <sup>2</sup>The University of Hong Kong

jdr53@cam.ac.uk, kaihhanx@hku.hk, samuel.albanie.academic@gmail.com

<https://grab-benchmark.github.io>

## Abstract

Large multimodal models (LMMs) have exhibited proficiencies across many visual tasks. Although numerous well-known benchmarks exist to evaluate model performance, they increasingly have insufficient headroom. As such, there is a pressing need for a new generation of benchmarks challenging enough for the next generation of LMMs. One area that LMMs show potential is graph analysis, specifically, the tasks an analyst might typically perform when interpreting figures such as estimating the mean, intercepts or correlations of functions and data series. In this work, we introduce GRAB, a graph analysis benchmark, fit for current and future frontier LMMs. Our benchmark is predominantly synthetic, ensuring high-quality, noise-free questions. GRAB is comprised of 3284 questions, covering five tasks and 23 graph properties. We evaluate 20 LMMs on GRAB, finding it to be a challenging benchmark, with the highest performing model attaining a score of just 21.0%. Finally, we conduct various ablations to investigate where the models succeed and struggle. We release GRAB and a lightweight GRAB-Lite to encourage progress in this important, growing domain.

## 1. Introduction

Driven by the union of increased compute resources, a rapidly growing research community, and competition fuelled by burgeoning commercial opportunities, the capabilities of frontier large multimodal models are swiftly improving and expanding. Furthermore, the timelines for advances, especially among closed-source models, appear to be shortening, with the time between model releases diminishing. This has important implications for evaluation and benchmarking, which are critical to gauge the relative strengths and weaknesses of models. In particular, as performance increases, common benchmarks become saturated and stale – for example, GPT-4o [31] scores 90.5%, 90.2%, and 88.7% on the widely used MGSM, HumanEval, MMLU benchmarks, respectively. The usefulness of a benchmark in distinguishing model capability diminishes when there is little

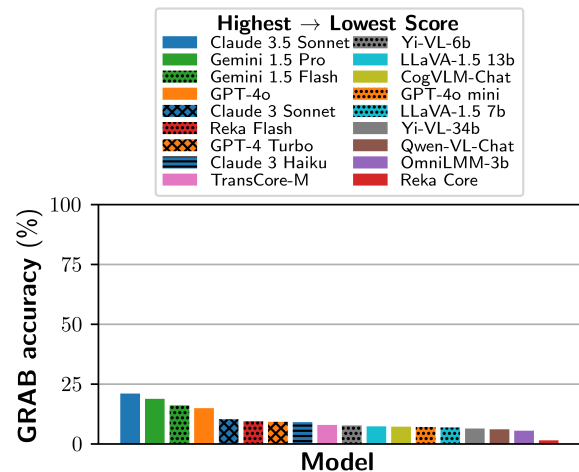


Figure 1. **Overall performance on GRAB.** Our benchmark proves challenging for frontier LMMs. The highest performing LMM attains an accuracy of just 21.0% on GRAB.

room for improvement. Moreover, label inaccuracies are reportedly pervasive [29] even in the most used benchmarks, further reducing possible headroom.

One particular task of interest and importance is the interpretation of scientific and mathematical figures and charts, which is central to many analytical endeavours. Specifically, we target use cases where the underlying data of a figure is not accessible, for example in documents, sketches, or other image types, and quantities can only be deduced from visual interpretation. Given their prevalence, the insights gained from understanding and reasoning over these types of figures make this task an important and valuable capability for LMMs. Various forms of chart evaluation have been present in machine learning literature for quite some time through established benchmarks [18, 28] and more recent works [45, 52]. However, these benchmarks need an uplift for the current and future generations of frontier models.

To this end, we introduce Graph Analysis Benchmark (GRAB), consisting of 3284 questions. Almost all figures in GRAB are generated synthetically using the Matplotlib library [16]. Although some realism might be lost when




Figure 2. The **GRAB Analysis Benchmark** consists of 3284 graph analysis questions that prove challenging for frontier LMMs. The questions cover 23 graph properties organised into five core tasks: (i) **Properties** focuses on the analysis of features of individual functions and series; (ii) **Functions** and (iii) **Series** require computing the mean of properties across multiple functions and series; (iv) **Transforms** involves determining the properties of a function after it has undergone a series of transforms; and, (v) **Real** encompasses question styles from the other tasks in more realistic formats, including sketched on whiteboard or paper, embedded in digital contexts or with added noise.

focusing on synthetically generated data, there are numerous benefits to taking this approach: (1) **Difficulty** – the complexity of questions can be controlled and tuned; (2) **Noise** – ground truth is determined automatically during curation, avoiding the need for potentially erroneous post hoc annotation; (3) **Composition** – we can directly select the characteristics of questions in our benchmark and arbitrarily control the scale; and, (4) **Contamination** – while the same figure style will almost certainly have been seen during the pretraining of most LMMs, it is unlikely that the exact figures and questions generated for our benchmark have. The inclusion of a 1114-question ‘Real’ task with figures that are hand-drawn, have added noise or are placed in digital screenshots increases the realism of the benchmark by introducing artefacts, distractors and camera effects.

Leveraging the control enabled by our synthetic generation, we optimise GRAB for usefulness and longevity. Specifically, we include challenging questions that frontier LMMs struggle to answer (see Fig. 1) and we control the aesthetic parameters of each graph to ensure all relevant features are clear and questions are answerable. To balance streamlined evaluation and robustness, we include 3284 questions in GRAB. We also provide a smaller 500-question version, called GRAB-Lite, as a lighter-weight evaluation.

We aim for a comprehensive evaluation of the abilities of frontier models to analyse figures by including a broad range of chart types and question formats. The tasks in GRAB involve determining key data properties, such as mean, variance, and interquartile range, as well as function properties like gradients, stationary points, and function parameters. Question complexity is increased by incorporating multiple functions and series (see Fig. 2) and increasing the required answer precision. We evaluate 20 LMMs on our benchmark, which proves extremely challenging for the current generation of models, with the best performing attaining a score of 21.0%. While some models exhibit a modicum of success on the easiest question set, the vast majority of our bench-

mark is beyond the capabilities of all models. We conduct a detailed error analysis and set of ablations to gain insights into which tasks and categories were the most challenging, as well as investigations of how question format and plotting library impact performance.

We summarise the major contributions of this work as follows: (i) We introduce GRAB, a challenging 3284 question benchmark for graph analysis, along with a smaller 500-question GRAB-Lite version (ii) We comprehensively characterise the capabilities of 20 current generation frontier LMMs on GRAB, and (iii) We provide insights into model strengths and limitations through numerous ablations.

## 2. Related Work

### 2.1. Figure interpretation benchmarks

Motivated by the last decade of progress in computer vision, a range of prior work has sought to evaluate figure interpretation capabilities of leading models. One family of work has focused on *synthetic* data generation pipelines, profiting from the reliable ground truth and fine-grained control enabled by this approach. This has yielded figure binary classification tasks [18], captioning [8], and various structural understanding, retrieval and reasoning tasks with limited vocabularies [17]. Synthetic figure generation has also been combined with corpora of online statistics [7, 28, 47] and Kaggle datasets [52] to achieve greater figure diversity. A second line of work has sought to construct datasets by curating ‘real’ figures directly from online sources and parsing their metadata. For this purpose, figures have been sourced from research papers [45, 46] and online sources such as *Our World in Data* [26]. GRAB draws particular inspiration from the family of synthetic benchmarks, to which it also belongs. However, a key distinction between our tasks and prior works is GRAB’s significantly greater difficulty, as evidenced by our empirical results illustrating that even the strongest frontier LMMs struggle to make headway (§4).

## 2.2. LMM evaluation and benchmarks

Since the emergence of LMMs, numerous benchmarks have been created to evaluate and characterise their capabilities across many domains. Interpreting and answering questions of natural images is a key task, with benchmarks such as MMBench [22], SEEDBench [19], SEEDBench-2 [20], MME [10], MM-Vet [53], LVLM-eHub [51], ScienceQA [24] and MMMU [54] finding prominence. However, as outlined in [9], although widely used, some of these most popular benchmarks contain problematic questions, including those that do not require visual content to be answered. Unlike older datasets developed to cater to the supervised learning paradigm that typically required much larger scales, more recent LMM-focused benchmarks contain significantly smaller test sets of a few thousand questions or a few hundred (e.g., [40]). Various benchmarks can be found in more niche domains, such as geographic and geospatial evaluation [44], hallucinations [15] and OCR [23]. A closely related work to GRAB is MathVista [25], which evaluates mathematical reasoning in visual contexts. MathVista is a metadataset (similar to [43]) consisting of a broad range of mostly multiple-choice questions from 31 multimodal datasets covering a diverse set of figure types and mathematical tasks ranging from table interpretation to puzzles. Though similar, the focus of MathVista on general mathematical reasoning, is distinct from ours, which is centered on the more narrow task of deriving and analysing graph properties. Moreover, our benchmark is considerably more challenging: GPT-4o scores >60% on MathVista<sup>1</sup> compared to 13.6% on GRAB.

## 3. GRaph Analysis Benchmark

Our GRaph Analysis Benchmark consists of 3284 questions covering five core tasks and 23 graph properties. In the following section, we provide an overview of the benchmark, including descriptions of the categories and tasks, example questions and answers, and an overview of our synthetic curation process. High-level statistics can be found in Tab. 1.

### 3.1. Categories

Questions in our GRAB benchmark cover 23 graph properties from which analytical insights can commonly be derived from data or functions in graphs and charts. We group these properties into the categories shown in Tab. 2. We consider this set of properties to cover most real use cases and require some measure of visual analytical reasoning. We refrain from including ‘easier’ OCR-centered questions – such as legend, title, axis-label interpretation – in our benchmark to keep the focus on visual analytical reasoning, and instead only include questions of this type as an ablation. Further-

<sup>1</sup>[https://huggingface.co/spaces/opencompass/open\\_vlm\\_leaderboard](https://huggingface.co/spaces/opencompass/open_vlm_leaderboard)

more, as outlined in §2, numerous prior benchmarks have included this type of question.

### 3.2. Tasks

We structure GRAB into the following tasks covering different aspects of graph analysis. Example questions and answers for each task can be found in Figs. 2 and 4.

#### 3.2.1. Properties

This initial task aims to provide a broad evaluation of the capability to understand and derive properties of graphs, including both basic fundamental properties like *number of series* or more advanced properties requiring calculations, like *total area bounded*. Aside from *number of series* questions, the graphs involve just a single function or data series.

#### 3.2.2. Functions

*Functions* builds on *Properties* with increased complexity by adding up to a total of 10 functions per graph. Questions require an additional step of averaging across all functions. For example, ‘Estimate the mean total area bounded by the function and the x-axis of the functions shown in the plot’. Added difficulty is introduced by functions overlapping.

#### 3.2.3. Series

Similar to *Functions*, *Series* builds on the initial task by adding up to 10 data series per graph, with questions requiring estimates of the mean of specific properties across the observable series, for example, ‘Estimate the mean Pearson correlation coefficient of the data series shown in the plot’. The additional series introduces challenges such as overlapping and clustered points.

#### 3.2.4. Transforms

Questions in the *Transforms* task include a single function and require computing a graph property after performing a sequence of up to 10 transforms (rotations, translations, scales and reflections).

#### 3.2.5. Real

Organised into four splits, the *Real* task modifies the other tasks with increased realism. This task includes *Properties*-style questions with figures hand-drawn on (1) **Paper** and (2) **Whiteboards**, and questions styled on all the other tasks with figures placed in (3) **Screenshots** (e.g., in emails, slides, documents and video calls) or with arbitrary (4) **Noise** added to the figures (e.g., artefacts, flips and blurs).

### 3.3. Curation

We leverage the Matplotlib [16] and Seaborn [49] libraries to synthetically create the majority of the graphs in GRAB and hand-draw only a small subset of the *Real* task. For each task, we initially created a longlist of 250 questions per graph property. We then performed downsampling to select a final question set with uniformly distributed answers. This

Statistic	Number
Questions	3284
Tasks	5
Graph properties	23
Properties questions	660
Functions questions	710
Series questions	490
Transforms questions	310
Real questions	1114
10s precision questions	46
Integer precision questions	2450
1 d.p. precision questions	787

Table 1. GRAB statistics

Category	Properties
Intercepts & Gradients	x-intercept, y-intercept, gradient
Stationary Points	stationary point(s) coordinates
Trigonometric	amplitude, vertical shift, period
Functions	function equations
Counting	number of points, number of series
Correlation	{Pearson’s, Spearman’s rank, Kendall rank} correlation coefficient
Area Bounded	total area bounded, net area bounded
Measures of Spread	mean, median, interquartile range, variance
Range & Extrema	min/max values, domain length, range

Table 2. GRAB categories and graph properties




Figure 3. GRAB categories

downselection aimed to increase question diversity, reduce bias in the generation process and avoid an overabundance of questions with answers converging towards 0. We associate each graph property (Tab. 2) with a set of data generation functions and function parameter ranges. For categories such as *Intercepts & Gradients* or *Functions*, we sample values for the property of interest and then fit the generation parameters to produce functions/series that match the value. For the other categories, such as *Correlation* or *Area Bounded*, we randomly sample the generation parameters and calculate the property value from the data. In each case, we ensure questions are unique. For each graph in the *Properties* task, aesthetic parameters (e.g. figure/font size and line style) are randomly sampled from a defined set of options. For the *Series*, *Functions* and *Transforms* tasks, we limit the variability of the graphs by configuring the appearance to a defined schema: all graphs have the same x and y-limits, grid lines, tick labels, font and resolution (2251x2171). A single function type is used for each question, irrespective of the number of series. Synthetically generated questions from the *Real* task leveraged the generation procedure from the other tasks to create new figures. These figures were then embedded in digital screenshots (**Screenshots** split) or manipulated with visual noise (**Noise** split). Figures in the **Paper** and **Whiteboard** splits were hand-drawn using a range of functions and graph styles and subsequently photographed.

For each task, we include sufficient questions to represent different combinations of graph properties and question complexity. We aim for an approximate balance per graph property, resulting in the category distribution shown in Fig. 3. Most questions require answers to integer precision, though some categories (e.g., *Area Bounded*) are evaluated to nearest 10 precision. To increase the difficulty, ~25% of questions require 1 decimal place precision.

### 3.3.1. Quality control

The curation process underwent numerous iterations of human review that sought to tune the hyperparameters involved in creating the graphs. This activity was primarily focused on jointly ensuring (i) the questions were answerable from a technical perspective, (ii) the ground-truth answers were cor-

rect and that (iii) the rendered figures were suitably readable. Similar spot-checks were carried out on the final dataset.

## 4. Experiments

### 4.1. Baselines

We evaluate 20 LMMs on GRAB, covering a broad range of both open- and closed-source models. We use chat or instruction-tuned variants rather than base models, of each model evaluated, as they are typically better at following task and output instructions. We benchmark 4 families of frontier closed-source LMMs, specifically the GPT-4 [30, 31], Gemini [11, 42] Claude 3 [3], and Reka [39] series. Additionally, we select the following open-source models for evaluation: Qwen-VL [6], TransCore-M [41], OmniLMM [37], Yi [1], CogVLM [48], and LLaVA-1.5 [21].

### 4.2. Experimental settings

#### 4.2.1. Inference

We evaluate the closed-source models via the Vertex AI API [12] {Gemini, Claude}, OpenAI API [35] {GPT}, and Reka API [2] {Reka}. Following benchmarking conventions, we set model hyperparameters that result in deterministic and reproducible generation. This entails utilising a greedy search decoding strategy in which the most probable token is selected from the model vocabulary  $V$  at each step, conditional on the preceding tokens *i.e.*,  $w_{n+1} = \arg \max_{w \in V} P(w|w_1, w_2, \dots, w_n)$ . We achieve this by setting the *temperature* and *topk* parameters to zero and 1, respectively, and specifying random seeds. We run the open-source models using HuggingFace Transformers [50] and the OpenCompass toolkit [38].

#### 4.2.2. Prompting

We adopt a simple prompting strategy throughout our evaluation, similar to that used by the popular repository *simple-evals* [36]. Concretely, we use a single basic *user* prompt consisting of the question and output format instructions. We do not modify the system prompt or tailor the prompt for each model. Although model performance is known to be sensitive to these input configurations, our investigation

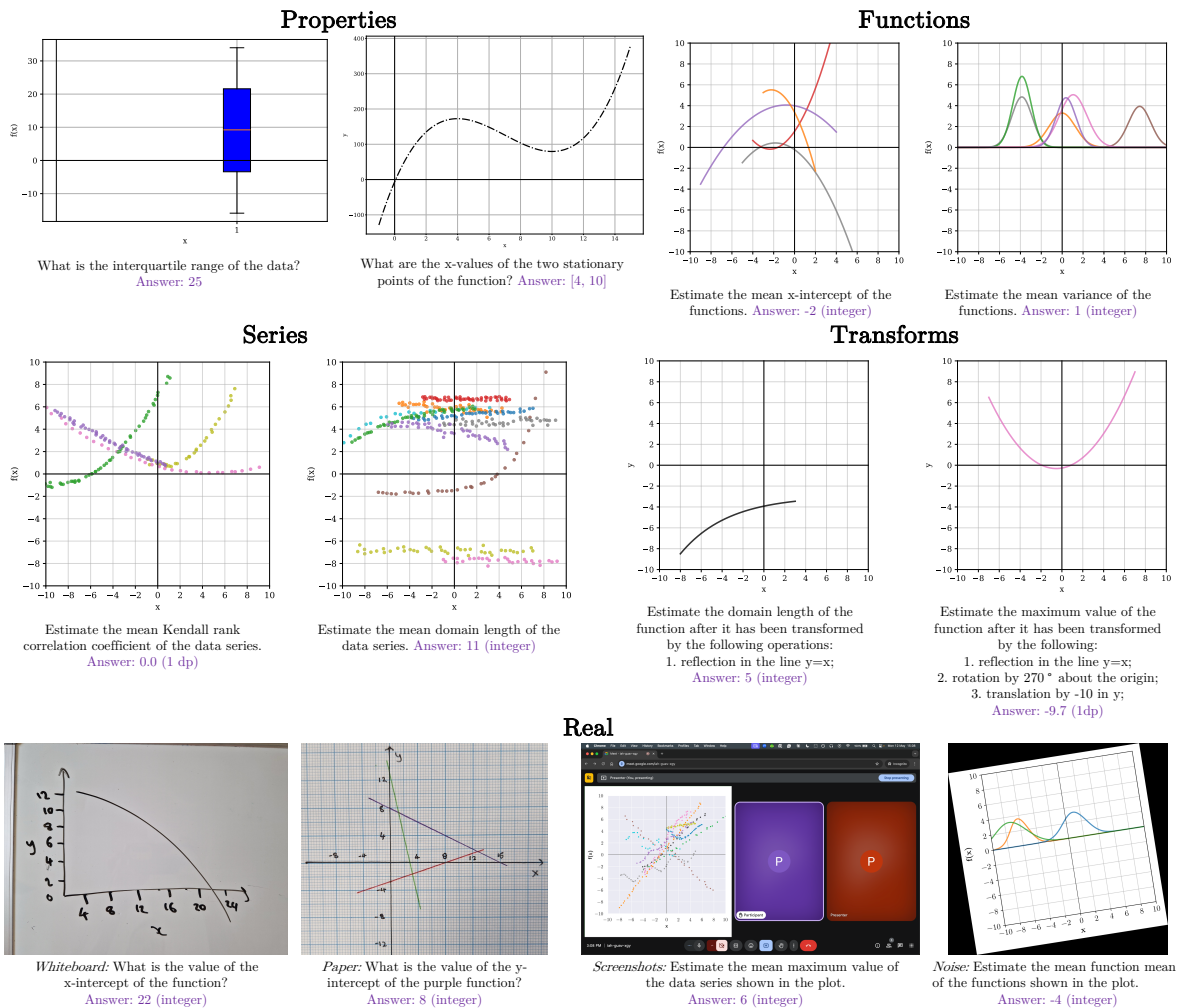


Figure 4. Example questions for the five tasks in the GRAB benchmark. All questions include synthetically rendered graphs apart from the Whiteboard and Paper splits of the *Real* task, which are hand-drawn and photographed.

focuses on chat or instruction-tuned models, which are optimised for instruction following: *if a model can only perform a task when a specific phrase is included in the prompt, this is a clear weakness*. Being able to answer questions in a realistic setting (*i.e.* as a human would be presented with the questions) offers a more generalised evaluation.

### 4.2.3. Evaluation

Our core evaluation protocol is exact matching between model output and the ground truth answers. The only post-processing step we take is removing leading whitespace. An answer is only scored as correct if the output characters at each position match those of the expected answer. This means that although verbose answers of the form ‘*The answer is...*’ or ‘<Reasoning> <Answer>’ might contain the correct answer, they will be marked incorrect. While this may lower model scores compared to human or LLM evaluation (*e.g.* [40, 45]), we consider this a better evaluation as

it jointly assesses task proficiency and instruction-following ability. Furthermore, this methodology does not necessitate (potentially) timely human scoring or the means to run a strong LLM evaluator. As mentioned previously, when considering potential downstream usage, a model is not useful if it can perform a given reasoning task perfectly but fails to reliably follow output instructions. As an ablation, we compare this strict exact matching protocol to automatic evaluation using both closed-source and open-source LLMs.

### 4.3. Overall results

We present the results on GRAB in Tab. 3. A clear observation is that **our benchmark proves challenging for frontier models**, with low overall performances from all models (see also Fig. 1). **The highest performing model, Claude 3.5 Sonnet, achieves a score of just 21.0%**. As expected, the highest scores are attained on the *Properties* task, which includes questions that just require the derivation of a single

Model	Properties (660)	Functions (710)	Accuracy (%)			Real (1114)	Overall (3284)
			Series (490)	Transforms (310)			
Closed-source LMMs							
Claude 3 Haiku [3]	14.2	6.6	8.8	3.9	9.2	9.1	
Claude 3 Sonnet [3]	15.3	8.6	4.5	4.8	12.4	10.3	
Claude 3.5 Sonnet [4]	<b>41.8</b>	<b>15.5</b>	11.0	<b>10.0</b>	19.6	<b>21.0</b>	
GPT-4 Turbo [30]	18.5	8.5	4.9	3.5	7.5	9.2	
GPT-4o mini [32]	15.8	6.8	5.7	2.9	4.0	7.1	
GPT-4o [31]	24.7	10.8	9.2	3.5	17.3	14.9	
Gemini 1.0 Pro Vision [11]	20.2	5.8	6.9	6.1	*	*	
Gemini 1.5 Flash [42]	28.5	11.5	8.4	9.0	17.1	16.1	
Gemini 1.5 Pro [42]	34.2	11.4	<b>13.3</b>	6.5	<b>20.3</b>	18.8	
Reka Edge [39]	11.8	8.7	11.6	1.9	*	*	
Reka Flash [39]	13.2	10.1	6.3	3.9	10.0	9.5	
Reka Core [39]	1.7	0.0	4.3	0.3	1.3	1.5	
Open-source LMMs							
CogVLM-Chat [48]	7.0	4.9	5.1	3.9	<b>10.5</b>	7.2	
Qwen-VL-Chat [6]	<b>10.2</b>	6.6	5.1	2.9	4.6	6.1	
OmniLMM-3b [37]	6.7	4.9	4.1	4.5	6.2	5.5	
TransCore-M [41]	7.9	<b>9.2</b>	7.6	3.9	8.2	<b>7.9</b>	
Yi-VL-6b [1]	5.6	8.6	7.1	4.2	9.7	7.7	
Yi-VL-34b [1]	7.6	5.9	5.5	2.3	7.5	6.4	
LLaVA-1.5 7b [21]	4.7	7.5	6.5	<b>4.8</b>	8.5	6.9	
LLaVA-1.5 13b [21]	5.0	7.7	<b>8.4</b>	3.9	8.9	7.3	

Table 3. **Results on GRAB.** The highest open and closed-source model scores for each task are **highlighted** and **bold**. \* unable to perform full evaluation due to model deprecation. Overall scores are weighted averages based on task size.

graph property. Lower scores are achieved on the *Functions* and *Series* tasks, likely due to the added complexity of computing the mean of a given property over multiple functions/series. There is a macro-trend of slightly higher accuracy on the functions task compared to the series task, possibly due to the irregularity introduced by the noise in the series data points. Significantly lower accuracy is attained on the *Transforms* task. Reasonable scores are attained on the *Real* task, which includes questions from the other tasks, that approximately correspond to the average scores on the other tasks. This indicates that the *Real* questions are of a similar difficulty.

A contributing factor to these scores is the degree to which the models follow output format instructions. Given our exact string matching evaluation protocol, any deviation from the expected output is deemed incorrect. Despite potentially having superior mathematical reasoning abilities, models that struggle to precisely follow instructions or have proclivities to generate verbose answers suffer severe performance degradation. With the exception of Reka Core (which showed poor output format instruction following capabilities), all closed-source models outperformed the leading open-source model. However, the open-source models did consistently generate answers in the correct format. In this low-scoring regime where even the best models struggle, a comparison of model performance should be made with the caveat that all models can barely perform the tasks in the benchmark. **We anticipate our benchmark will yield most benefit when evaluating the next generation of LMMs.**

#### 4.4. Real splits analysis

In Tab. 4, we report model scores on questions in each split in the *Real* task. The highest scores are attained on the

Model	Accuracy (%)				Overall (1114)
	Whiteboard (41)	Paper (41)	Screenshots (516)	Noise (516)	
Closed-source LMMs					
Claude 3 Haiku [3]	9.8	4.9	9.5	9.3	9.2
Claude 3 Sonnet [3]	9.8	14.6	12.8	12.0	12.4
Claude 3.5 Sonnet [4]	14.6	17.0	18.6	<b>21.1</b>	19.6
GPT-4 Turbo [30]	9.8	4.9	6.6	8.3	7.5
GPT-4o mini [32]	7.3	2.4	4.7	3.3	4.0
GPT-4o [31]	29.3	9.8	18.0	16.3	17.3
Gemini 1.5 Flash [42]	31.7	17.1	16.3	16.7	17.1
Gemini 1.5 Pro [42]	<b>34.4</b>	<b>24.4</b>	<b>21.9</b>	17.2	<b>20.3</b>
Reka Flash [39]	19.5	7.3	7.9	11.4	10.0
Reka Core [39]	4.9	0.0	1.0	1.4	1.3
Open-source LMMs					
CogVLM-Chat [48]	<b>14.6</b>	7.3	<b>11.0</b>	9.9	<b>10.5</b>
Qwen-VL-Chat [6]	12.2	4.9	2.9	5.6	4.6
OmniLMM-3b [37]	7.3	2.4	6.6	6.0	6.2
TransCore-M [41]	2.4	<b>9.8</b>	7.6	9.1	8.2
Yi-VL-6b [1]	7.3	4.9	9.9	<b>10.1</b>	9.7
Yi-VL-34b [1]	4.9	4.9	7.8	7.8	7.5
LLaVA-1.5 7b [21]	4.9	4.9	9.5	8.1	8.5
LLaVA-1.5 13b [21]	7.3	4.9	8.9	9.3	8.9
Gemini 2.0 Flash [13]	22.0	17.0	18.8	20.5	19.6
Gemini 2.5 Flash [14]	<b>36.6</b>	<b>22.0</b>	<b>29.5</b>	<b>29.3</b>	<b>29.4</b>

Table 4. **Performance on the GRAB ‘Real’ task splits.** Overall scores are weighted averages based on task size.

Whiteboard and Paper splits, which include only less complex *Properties*-style questions. Notably, Claude 3.5 Sonnet struggles, scoring higher on the other splits. In general, model scores on both the Screenshots and Noise splits are comparable to the overall GRAB scores. This suggests the models are well-suited to dealing with real-world images, as the addition of visual distortions due to the noise artefacts or distractors and additional context in the screenshots does not significantly degrade performance. We also report the performance of two newer releases of the Gemini 1.5 Flash model (2.0 and 2.5), each showing subsequent improvements.

#### 4.5. Category analysis

A decomposition of model performance into question category is provided in Tab. 5. Significant variance is shown across categories, with some models much better suited for certain categories than others. Overall, the highest accuracies are attained on the *Counting* category, which encompasses simple questions determining the number of series shown in each plot, or the approximate number of data points. On the other hand, **no model correctly answered a single question from the Functions category.** This challenging category requires correct derivation of multiple parameters (*e.g.* the gradient and intercept for linear functions) and correct formatting of the output. Inspection of the model output suggests adherence to the strict output format was not an issue: the models were unable to correctly determine the function parameters. However, in a number of cases, proposed functions were close to the correct answers, *e.g.*, with all parameters correct apart from one that was off by one. Minor perturbations to the graph appearance did result in a small number of correct answers.

Model	Category	Accuracy (%)									
		Range (360)	Measures of Spread (370)	Area Bounded (150)	Intercepts & Gradients (290)	Stationary Points (160)	Functions (30)	Correlation (120)	Extrema (360)	Counting (60)	Trigonometric (270)
Closed-source LLMs											
Claude 3 Haiku [3]		4.7	11.9	2.7	11.4	10.6	0.0	20.8	5.8	13.3	10.0
Claude 3 Sonnet [3]		4.7	10.0	4.0	13.8	8.1	0.0	9.2	8.9	18.3	11.9
Claude 3.5 Sonnet [4]		<b>13.6</b>	<b>28.6</b>	2.7	<b>25.5</b>	<b>25.6</b>	0.0	9.2	<b>24.7</b>	30.0	<b>29.3</b>
GPT-4 Turbo [30]		6.7	13.0	0.0	13.1	11.9	0.0	5.0	8.6	18.3	14.8
GPT-4o mini [32]		3.3	10.8	2.0	10.7	8.1	0.0	2.5	8.1	26.7	15.6
GPT-4o [31]		9.2	18.6	2.0	14.1	16.2	0.0	15.8	8.1	30.0	21.5
Gemini 1.0 Pro Vision [11]		4.7	16.2	1.3	7.9	13.1	0.0	12.5	13.1	23.3	10.4
Gemini 1.5 Flash [42]		11.7	18.9	2.7	16.2	16.9	0.0	8.3	22.2	25.0	16.3
Gemini 1.5 Pro [42]		9.4	23.2	<b>4.7</b>	17.9	24.4	0.0	26.7	19.4	<b>33.3</b>	19.3
Reka Edge [39]		2.8	9.7	3.3	10.0	8.1	0.0	<b>38.3</b>	4.7	18.3	13.3
Reka Flash [39]		6.7	11.6	3.3	15.5	7.5	0.0	7.5	5.3	18.3	12.6
Reka Core [39]		0.0	0.0	0.0	0.0	1.2	0.0	25.0	0.3	0.0	0.0
Open-source LLMs											
CogVLM-Chat [48]		6.9	5.4	2.7	4.1	5.6	0.0	0.8	4.7	16.7	7.4
Qwen-VL-Chat [6]		3.6	<b>8.6</b>	1.3	6.6	<b>10.6</b>	0.0	1.7	<b>6.7</b>	<b>18.3</b>	10.4
OmniLMM-3b [37]		8.9	3.2	2.7	5.9	3.1	0.0	0.0	5.0	11.7	6.7
TransCore-M [41]		8.6	4.1	2.7	<b>10.0</b>	8.8	0.0	15.8	5.8	8.3	10.4
Yi-VL-6b [1]		5.0	6.5	<b>3.3</b>	7.9	5.6	0.0	12.5	5.6	1.7	<b>11.5</b>
Yi-VL-34b [1]		2.8	7.8	2.7	8.3	1.9	0.0	8.3	4.2	6.7	10.0
LLaVA-1.5 7b [21]		<b>9.7</b>	4.9	<b>3.3</b>	4.8	5.6	0.0	9.2	5.3	5.0	6.3
LLaVA-1.5 13b [21]		8.3	4.6	2.7	7.9	4.4	0.0	<b>18.3</b>	5.3	8.3	5.2

Table 5. **GRAB per-category performance.** The highest open and closed-source model scores for each category are highlighted and bold. To isolate per-category performance and avoid the effects of noise and distractors, the results in this table do not include the *Real* task.

#### 4.6. Complexity analysis

The *Functions*, *Series* and *Transforms* tasks were designed to include questions spanning a controllable complexity scale. For *Functions* and *Series*, complexity refers to the number of additional functions or series present on the graph, while for *Transforms*, it represents the number of transformations applied to the function. In both cases, the complexity score ranges from 0 (single function/series and no transformations) to 9. In Fig. 5, we present performance decomposed along this complexity scale for a representative selection of models. Intuitively, we expect performance to decrease as complexity increases due to the added challenge of determining the quantity of interest from additional functions/series or after additional transforms. However, the complexity plots for the evaluated models follow two distinct patterns. For the better performing models, such as Claude 3.5 Sonnet and Gemini 1.5 Pro, performance does clearly decrease as complexity increases from 0 to 3, before remaining around 10%. For the weaker models, however, the results fluctuate around 10% across the entire complexity domain. In these cases, even the lowest complexity questions are too challenging. As expected, the questions requiring decimal precision proved much harder than integer precision questions.

#### 4.7. Evaluation protocols

In addition to the exact matching evaluation protocol (§4.2.3), we explore model performance with more lenient automatic evaluation protocols. Specifically, we utilised an LLM to parse the generated output from the LLMs and extract the answer text, before evaluating it via exact matching as before. We carried this out using both a frontier LLM

(Gemini 1.5 Pro) and a small open-source LLM (LLaMA 3.1 8b-instruct [27]). Details of the prompts used can be found in the Appendix. We found the LLMs could extract specific answers from verbose output with high accuracy but performed at determining the correctness of a model response given the ground truth answer. A comparison of the different evaluation protocols can be found in Fig. 6. **On the whole, there is almost no difference between the exact matching and automatic evaluation protocols.** Similarly, there is little difference between extraction with a frontier closed-source and a small open-source model, suggesting this approach can feasibly be conducted with little compute. However, **we advocate for the adoption of exact matching without the use of an LLM as the standard evaluation approach going forwards** as most models are clearly able to sufficiently follow output instructions and it provides a better signal of model capabilities. Cases where there are differences between the performance protocols indicate models that are the poorest instruction followers.

#### 4.8. Additional ablations

Additional ablations and analyses can be found in the Appendix, including a comparison of the Matplotlib and Seaborn plotting libraries, alternative evaluation metrics, results on an OCR-based question set, and an analysis of performance with different question formats.

#### 4.9. GRAB-Lite

Recently, there has been a shift towards *reasoning* models that leverage inference-time scaling and spend more effort “thinking” [33]. While this strategy can lead to better-

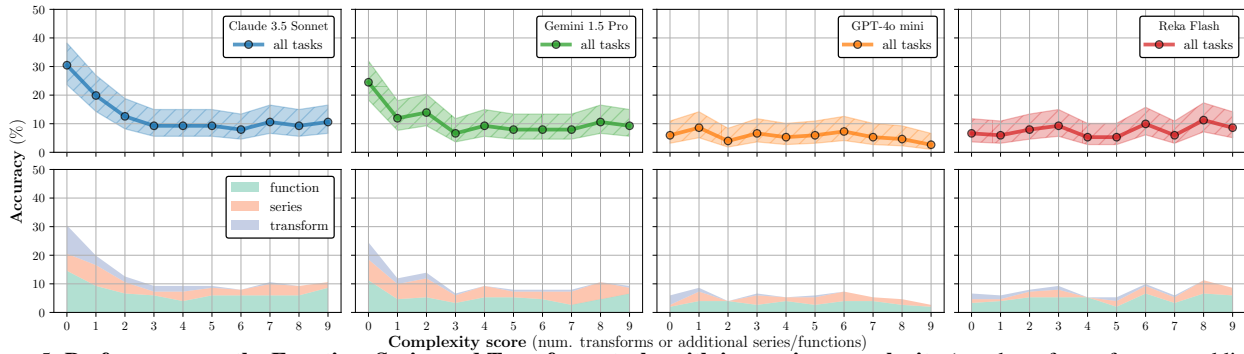


Figure 5. **Performance on the Function, Series and Transforms tasks with increasing complexity** (number of transforms or additional series/functions). *Top row*: accuracy across all 3 tasks with shaded region displaying 95% Wilson confidence intervals (each point on the plot represents 151 questions). *Bottom row*: accuracy per task decomposition.




Figure 6. **Overall accuracy on the Properties, Functions, Series and Transforms tasks with different automatic evaluation protocols.** Scores are shown for *exact matching* and *answer extraction (with different LLMs) and then exact matching*. All accuracy differences relative to the exact matching protocol are annotated – in nearly every case, there is either minor or no difference in score.

Model	Accuracy (%)					
	Properties (100)	Functions (100)	Series (100)	Transforms (100)	Real (100)	Overall (500)
Claude 3.5 Sonnet [4]	39.0	15.0	11.0	13.0	20.0	19.6
Gemini 1.5 Pro [42]	23.0	10.0	14.0	7.0	25.0	15.8
GPT-4o [31]	21.0	7.0	10.0	6.0	19.0	12.6
Reasoning Models						
Claude 3.7 Sonnet [5]	36.0	13.0	13.0	11.0	10.0	16.6
Gemini 2.5 Flash [14]	34.0	27.0	29.0	22.0	30.0	28.4
o1 [34]	27.0	15.0	28.0	26.0	25.0	24.2

Table 6. **Results on GRAB-Lite**, a lightweight 500-question subset of GRAB designed to evaluate costly reasoning models.

reasoned output, it requires more output tokens, which increases the cost of inference significantly and constrains the scale of evaluation. This creates a pressing need for *lightweight* benchmarks. To address this, we introduce a small 500-question version of our GRAB, called GRAB-Lite. GRAB-Lite is an approximately category, complexity, and task-balanced composition of GRAB with synthetic graphs generated as an equal split between Matplotlib and Seaborn. We evaluate 3 strong GRAB baselines and some frontier reasoning models on GRAB-Lite and report the results in Tab. 6. The baselines attain similar overall scores on both GRAB and GRAB-Lite, with Claude 3.5 Sonnet scoring highest on both. The reasoning models also perform relatively well, with Gemini 2.5 Flash achieving the highest overall score. o1 averaged >5k completion tokens (\$0.33) per question, highlighting the need for such a lightweight eval.

## 5. Conclusions

We introduce GRAB, a GRaph Analysis Benchmark for LMMs featuring five tasks covering a diverse range of graph configurations and settings. We leverage synthetic data generation to create high-quality, difficult questions that prove challenging for frontier LMMs, with the best-performing model attaining a score of 21.0%. On GRAB, we evaluate 20 closed and open-source models, attaining a broad profile of the abilities of the current generation of LMMs. During inference, we use lightweight prompts with clear output format instructions and assess answer correctness by exact string matching. This approach jointly assesses the abilities of models to perform both the visual perception and reasoning in our tasks and reliably follow output instructions. We find little difference in performance using exact matching evaluation vs. automatic LLM evaluation, suggesting most current models can sufficiently follow output instructions, validating exact matching as a suitable evaluation protocol. We also curate a lightweight GRAB-Lite question set designed to evaluate reasoning models. Finally, we conduct studies revealing the categories and splits the models struggle with, and the impacts of different question types and plotting libraries. We release our datasets along with generation and evaluation code for the community to use and hope our work encourages research towards the next generations of LMMs.

## Acknowledgements

This work was supported by the UKRI Centre for Doctoral Training in Application of Artificial Intelligence to the study of Environmental Risks (reference EP/S022961/1), an Isaac Newton Trust grant, a research gift from Google, an EPSRC HPC grant, the Hong Kong Research Grant Council - General Research Fund (Grant No. 17211024), and HKU Seed Fund for Basic Research. Samuel would like to acknowledge the support of Z. Novak and N. Novak in enabling his contribution.

## References

- [1] 01-ai. Yi. <https://github.com/01-ai/Yi>, 2023. 4, 6, 7
- [2] Reka AI. Reka AI API. <https://platform.reka.ai/dashboard>, 2024. 4
- [3] Anthropic. Introducing the next generation of Claude. <https://www.anthropic.com/news/claude-3-family>, 2024. 4, 6, 7
- [4] Anthropic. Claude 3.5 Sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024. 6, 7, 8
- [5] Anthropic. Claude 3.7 Sonnet System Card. <https://assets.anthropic.com/m/785e231869ea8b3b/original/claude-3-7-sonnet-system-card.pdf>, 2025. 8
- [6] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-VL: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 4, 6, 7
- [7] Ritwick Chaudhry, Sumit Shekhar, Utkarsh Gupta, Pranav Maneriker, Prann Bansal, and Ajay Joshi. Leaf-qa: Locate, encode and attend for figure question answering. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3512–3521, 2020. 2
- [8] Charles Chen, Ruiyi Zhang, Eunye Koh, Sungchul Kim, Scott Cohen, Tong Yu, Ryan Rossi, and Razvan Bunescu. Figure captioning with reasoning and sequence-level training. *arXiv preprint arXiv:1906.02850*, 2019. 2
- [9] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are We on the Right Way for Evaluating Large Vision-Language Models? *arXiv preprint arXiv:2403.20330*, 2024. 3
- [10] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 3
- [11] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 4, 6, 7
- [12] Google. Vertex AI. <https://cloud.google.com/vertex-ai/>, 2024. 4
- [13] Google. Gemini 2.0 Flash Model Card. <https://storage.googleapis.com/model-cards/documents/gemini-2-flash.pdf>, 2025. 6
- [14] Google. Gemini 2.5 Flash Model Card. <https://storage.googleapis.com/model-cards/documents/gemini-2.5-flash.pdf>, 2025. 6, 8
- [15] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoub, and Dinesh Manocha, Tianyi Zhou. HALLUSION-BENCH: An Advanced Diagnostic Suite for Entangled Language Hallucination & Visual Illusion in Large Vision-Language Models. *arXiv e-prints*, pages arXiv–2310, 2023. 3
- [16] J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. 1, 3
- [17] Kushal Kafle, Brian Price, Scott Cohen, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2018. 2
- [18] Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. Figureqa: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*, 2017. 1, 2
- [19] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. SEED-Bench: Benchmarking multimodal LLMs with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023. 3
- [20] Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. SEED-Bench: Benchmarking Multimodal Large Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13299–13308, 2024. 3
- [21] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024. 4, 6, 7
- [22] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023. 3
- [23] Yuliang Liu, Zhang Li, Hongliang Li, Wenwen Yu, Mingxin Huang, Dezhi Peng, Mingyu Liu, Mingrui Chen, Chunyuan Li, Lianwen Jin, et al. On the hidden mystery of ocr in large multimodal models. *arXiv preprint arXiv:2305.07895*, 2023. 3
- [24] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022. 3
- [25] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathemati-

- cal reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023. 3
- [26] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022. 2
- [27] Meta. Meta-Llama-3.1-8B-Instruct. <https://huggingface.co/meta-llama/Meta-Llama-3.1-8B-Instruct>, 2024. 7
- [28] Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536, 2020. 1, 2
- [29] Curtis G Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. *arXiv preprint arXiv:2103.14749*, 2021. 1
- [30] OpenAI. GPT-4V(ision) System Card. [https://cdn.openai.com/papers/GPTV\\_System\\_Card.pdf](https://cdn.openai.com/papers/GPTV_System_Card.pdf), 2023. 4, 6, 7
- [31] OpenAI. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. 1, 4, 6, 7, 8
- [32] OpenAI. GPT-4o mini: advancing cost-efficient intelligence. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>, 2024. 6, 7
- [33] OpenAI. Learning to reason with LLMs. <https://openai.com/index/learning-to-reason-with-llms/>, 2024. 7
- [34] OpenAI. OpenAI o1: A Large Language Model for Complex Reasoning. OpenAI website, 2024. <https://openai.com/o1/>. 8
- [35] OpenAI. API Reference. <https://platform.openai.com/docs/api-reference>, 2024. 4
- [36] OpenAI. simple-evals. <https://github.com/openai/simple-evals>, 2024. Accessed: 15-05-2024. 4
- [37] OpenBMB. OmniLMM. <https://github.com/OpenBMB/OmniLMM>, 2024. 4, 6, 7
- [38] OpenCompass Contributors. OpenCompass: A Universal Evaluation Platform for Foundation Models. <https://github.com/open-compass/opencompass>, 2023. 4
- [39] Aitor Ormazabal, Che Zheng, Cyprien de Masson d’Autume, Dani Yogatama, Deyu Fu, Donovan Ong, Eric Chen, Eugenie Lamprecht, Hai Pham, Isaac Ong, et al. Reka Core, Flash, and Edge: A Series of Powerful Multimodal Language Models. *arXiv preprint arXiv:2404.12387*, 2024. 4, 6, 7
- [40] Piotr Padlewski, Max Bain, Matthew Henderson, Zhongkai Zhu, Nishant Relan, Hai Pham, Donovan Ong, Kaloyan Aleksiev, Aitor Ormazabal, Samuel Phua, et al. Vibe-Eval: A hard evaluation suite for measuring progress of multimodal language models. *arXiv preprint arXiv:2405.02287*, 2024. 3, 5
- [41] PCIResearch. TransCore-M. <https://github.com/PCIResearch/TransCore-M>, 2023. 4, 6, 7
- [42] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024. 4, 6, 7, 8
- [43] Jonathan Roberts, Kai Han, and Samuel Albanie. SATIN: A multi-task metadataset for classifying satellite imagery using vision-language models. *arXiv preprint arXiv:2304.11619*, 2023. 3
- [44] Jonathan Roberts, Timo Lüddecke, Rehan Sheikh, Kai Han, and Samuel Albanie. Charting New Territories: Exploring the geographic and geospatial capabilities of multimodal LLMs. *arXiv preprint arXiv:2311.14656*, 2023. 3
- [45] Jonathan Roberts, Kai Han, Neil Houlsby, and Samuel Albanie. SciFIBench: Benchmarking Large Multimodal Models for Scientific Figure Interpretation. *arXiv preprint arXiv:2405.08807*, 2024. 1, 2, 5
- [46] Noah Siegel, Zachary Horvitz, Roie Levin, Santosh Divvala, and Ali Farhadi. Figureseer: Parsing result-figures in research papers. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 664–680. Springer, 2016. 2
- [47] Hrituraj Singh and Sumit Shekhar. STL-CQA: Structure-based transformers with localization and encoding for chart question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3275–3284, 2020. 2
- [48] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. CogVLM: Visual Expert for Pretrained Language Models. *arXiv preprint arXiv:2311.03079*, 2023. 4, 6, 7
- [49] Michael L. Waskom. seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021. 3
- [50] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. "Transformers: State-of-the-Art Natural Language Processing". In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, 2020. Association for Computational Linguistics. 4
- [51] Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. LVLm-eHub: A comprehensive evaluation benchmark for large vision-language models. *arXiv preprint arXiv:2306.09265*, 2023. 3
- [52] Zhengzhuo Xu, Sinan Du, Yiyan Qi, Chengjin Xu, Chun Yuan, and Jian Guo. Chartbench: A benchmark for complex visual reasoning in charts. *arXiv preprint arXiv:2312.15915*, 2023. 1, 2
- [53] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang.

Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023. 3

- [54] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. *arXiv preprint arXiv:2311.16502*, 2023. 3