

Centurio: On Drivers of Multilingual Ability of Large Vision-Language Model

Gregor Geigle^{12*} Florian Schneider^{3*} Carolin Holtermann⁴
Chris Biemann³ Radu Timofte² Anne Lauscher⁴ Goran Glavaš¹

¹WüNLP, ²Computer Vision Lab, CAIDAS, University of Würzburg

³Language Technology Group, ⁴Data Science Group, University of Hamburg
(gregor.geigle|florian.schneider-1)@uni-(wuerzburg|hamburg).de

gregor-ge.github.io/Centurio

Abstract

Most Large Vision-Language Models (LVLMs) to date are trained predominantly on English data, which makes them struggle to understand non-English input and fail to generate output in the desired target language. Existing efforts mitigate these issues by adding multilingual training data, but do so in a largely ad-hoc manner, lacking insight into how different training mixes tip the scale for different groups of languages. In this work, we present a comprehensive investigation into the training strategies for massively multilingual LVLMs. First, we conduct a series of multi-stage experiments spanning 13 downstream vision-language tasks and 43 languages, systematically examining: (1) the number of training languages that can be included without degrading English performance and (2) optimal language distributions of pre-training as well as (3) instruction-tuning data. Further, we (4) investigate how to improve multilingual text-in-image understanding, and introduce a new benchmark for the task. Surprisingly, our analysis reveals that one can (i) include as many as 100 training languages simultaneously (ii) with as little as 25-50% of non-English data, to greatly improve multilingual performance while retaining strong English performance. We further find that (iii) including non-English OCR data in pre-training and instruction-tuning is paramount for improving multilingual text-in-image understanding. Finally, we put all our findings together and train Centurio, a 100-language LVLM, offering highly competitive performance in an evaluation covering 14 tasks and 56 languages.

1 Introduction

Large Vision-Language Models (LVLMs) extend Large Language Models (LLMs) (Brown et al., 2020) to natively understand images as input (Li et al., 2023; Liu et al., 2023). This leverages the

impressive language generation and reasoning abilities of recent LLMs (Llama Team, 2024; Yang et al., 2024) for vision-language tasks like image captioning or visual question answering.

While there exists a plethora of LVLMs (Zhang et al., 2024), most models are trained with just English data (Liu et al., 2024a; Dai et al., 2023; Liu et al., 2024b). This limits the access for speakers of other languages as the resulting models have several limitations even if the underlying LLMs exhibit multilingual capabilities: the models fail to understand non-English instructions (Schneider and Sitaram, 2024), struggle with non-English text content in images (Tang et al., 2024), and often fail to reply in the correct language, i.e., they have problems with *language fidelity* (Hinck et al., 2024). To ameliorate these issues, LVLMs need to be trained on a multilingual data composition. As the amount of data one can train on is always limited—by time, computing resources, financial costs, or other constraints—an effective distribution of the data across different languages is paramount. Existing multilingual LVLM work has, however, given minimal consideration to this central question of optimal training data composition (e.g., Geigle et al., 2023; Sun et al., 2024; Maaz et al., 2025).

In this work, we comprehensively investigate the space of language distributions of LVLM training mixes, focusing on the presumed trade-off between the number of included languages and performance across languages—grouped by the amount of data available for them—under a *fixed training budget*. We train several models with different data compositions obtained by machine-translating high-quality English data and benchmark them across 13 downstream tasks covering 43 diverse languages—from low-resource languages like Igbo to high-resource languages like German. We focus on four research questions, each building on the previous one, designed to identify an optimal multilingual training mix: **RQ1:** What is the optimal

* Equal contribution.

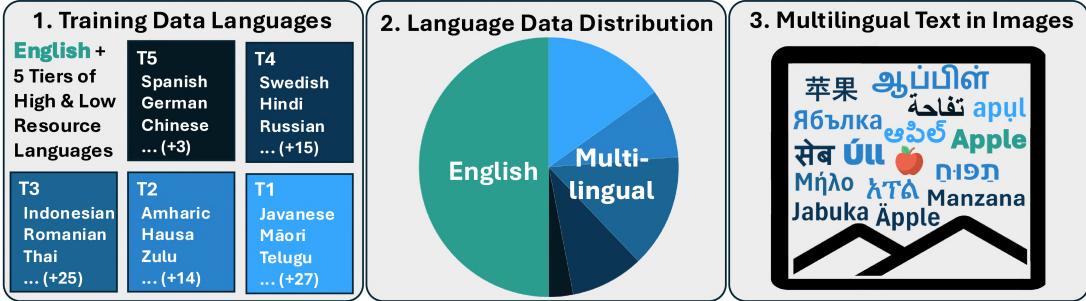


Figure 1: Exploring drivers of multilingual ability: (1) Languages in the training data; (2) Distribution of languages in the training data; (3) Incorporating multilingual OCR samples to understand non-English text in images.

number of training languages? **RQ2 & RQ3:** What is the optimal distribution of data across languages in (RQ3) pre-training data and (RQ2) instruction-tuning? **RQ4:** How to improve the understanding of multilingual text in images? To measure progress for RQ4, we introduce SMPQA (Synthetic Multilingual Plot Question Answering), a novel dataset for testing multilingual OCR capabilities, spanning 11 languages and 7 scripts.

Our findings are encouraging, albeit surprising.

1. We do not observe the infamous “curse of multilinguality” (Conneau et al., 2020; Pfeiffer et al., 2022b) and find that gradually increasing the number of languages incurs only a negligible “performance tax”: scaling from 7 to 100 languages greatly improves performance for languages newly added to the training data, especially with respect to language fidelity, while largely retaining performance levels for all previously added languages.

2. We find that exposure to a language matters more than increasing the training portion of the language or, in particular, that the majority of the training data can still be in English, which lowers the cost of acquiring training data in other languages (e.g., via machine translation). Concretely, we find that turning 25 to 50% of training data multilingual yields strong performance, with more data sometimes even degrading performance; in *pre-training*, having a larger share of multilingual data is more beneficial, but it also saturates after 50%.

3. We obtain mixed results for text-in-image problems: while incorporating (synthetic) OCR data with 5k samples per language rapidly boosts the performance for Latin-script languages, the same does not hold for languages with other scripts.

Finally, to demonstrate the practical impact of our findings, we train Centurio, a massively multilingual LVLM with 100 languages, following what we found to be an “optimal” data distribution across languages for both training stages. Centurio

achieves state-of-the-art results over 14 tasks , matching the performance of popular multilingual open-weight LVLMs like Qwen2-VL (Wang et al., 2024b), InternVL 2.5 (Chen et al., 2024d) and Pangea (Yue et al., 2025) on English and other high-resource languages while outperforming them on low(er)-resource languages.

2 Drivers of Multilingual Ability

The design space for training (multilingual) LVLMs is extensive, ranging from the choice of the image encoder and the alignment module between the image encoder and LLM to the selection of training data. (Karamchetti et al., 2024; Laurençon et al., 2025; Tong et al., 2024). Exhaustively searching through the cross-product of all choices is not feasible. In this work, we focus on extensive evaluation of language distributions of training data in both pre-training and instruction-tuning. Intuitively, this should be a major factor affecting the multilingual ability of an LVLM. Figure 1 illustrates the scope of our analysis. We keep adding groups of languages—from highest-to lowest-resourced, following the “resourceness” tiers of (Joshi et al., 2020)—into the training mix while keeping the data size fixed. Besides the number of languages, our main focus is on the division of the training budget between English and all other languages. Finally, we posit that, besides understanding instructions and generating outputs in different languages, truly multilingual LVLMs must be able to “understand” multilingual text in images. We thus pay special attention to training adaptions for multilingual text-in-image problems.

2.1 Experimental Setup

Architecture. For our experiments, we adopt the popular LLaVA architecture (Liu et al., 2023, 2024a): An image encoder encodes images into

a sequence of visual tokens, which are projected with a 2-layer MLP into the LLM embedding space. These projected visual tokens are then concatenated with the text tokens and fed through the LLM to generate textual output.

For all our experiments, we chose SigLIP S0400¹ (Zhai et al., 2023) as the image encoder because it seems to be the best-performing “CLIP-like”, i.e., language-supervised, embedding model, as shown by Tong et al. (2024). Further, we decided against other current image encoders for the following two reasons: 1) The multilingual language understanding and text generation capabilities evaluated in RQ1, RQ2, and RQ3 are mainly attributed to the LLM, with the image encoder expected to play only a minor role. 2) When it comes to “text-in-image” understanding capabilities assessed in RQ4, where the image encoder is thought to play a more important role, all relevant image encoders are trained primarily with English image captions. Hence, we expect all of them to struggle with low-resource languages and especially with non-Latin scripts due to the lack of exposure to respective data during training.

We chose Phi 3.5 (Abdin et al., 2024b) as our LLM because it exhibits strong multilingual performance while its small size (3.8B parameters) allows for more computationally efficient experimentation while still exhibiting strong multilingual performance. To show that our findings generalize to other LLMs, we repeat a subset of the analysis experiments with Llama 3 (8B) (Llama Team, 2024) as the LLM backbone (see Appendix D.1).

Training Setup. Following previous work (Liu et al., 2024a; Tong et al., 2024), we split the training into two phases : 1) *pre-training*: the model is trained only on image captioning, with dense image captions; 2) *instruction tuning*: the model is trained on a mix of diverse vision-language tasks using several public datasets. While pre-training benefits downstream performance, it is not strictly necessary for the LVLM to perform well on the downstream tasks (Karamcheti et al., 2024). To reduce the computational cost of our analysis (i.e., to avoid coupling each language distribution over pre-training data with every language distribution of instruction-tuning data), we skip pre-training while searching for an optimal language distribution for instruction tuning. Then, with instruction-tuning data fixed, we search for an optimal language dis-

tribution for pre-training data. In both stages, we freeze the image encoder and only update the MLP and LLM (with LoRA (Hu et al., 2022)) weights. We provide further details in Appendix A.

Training Data. Our controlled experiments require comparable data over a wide range of languages. Existing multilingual datasets, available only for some tasks and in a handful of languages² thus do not meet our needs. Instead, we resort to machine translation (MT) and use the open NLLB model (Costa-jussà et al., 2022)³ to translate readily available English datasets.⁴ While MT results in lower data quality, especially for lower-resource languages, it is the only option to obtain multilingual vision-language training data at scale. Moreover, gains from “low-quality” MT data are guaranteed to be met or even surpassed with higher-quality translations (e.g., commercial MT or human translators). Our instruction-tuning data is adapted from LLaVA-Next (Liu et al., 2024b) and contains 0.77M samples. For pre-training, we use the 1.3M dense captions from ShareGPT4v (Chen et al., 2024b). We provide further details in Appendix B.

Evaluation. We curate an extensive test suite of 13 tasks covering 43 languages to assess the multilingual abilities of our models. Following Joshi et al. (2020), we cluster the tested languages into five *tiers*, with *T5* encompassing the high-resource languages (e.g., German, Chinese) and *T1* extremely low-resource languages (e.g., Maori, Telugu). The tasks contained in our test suite are twofold: (1) *discriminative* tasks with questions that require binary (“yes/no”) or multiple-choice answers and (2) *open generation* tasks, where the models need to generate output in the target language (e.g., an image caption or a free-form answer). Generative tasks additionally evaluate a model’s *language fidelity*, i.e., the ability to generate the answer in the language of the instruction. The full list of evaluation tasks and languages, along with further details, is in Appendix C. We report the results for language tiers (*T1–T5*), averaging the scores over all tasks and all tier languages.⁵ We separately report English performance and **exclude** it from *T5*.

²See, for example, datasets collected by Yue et al. (2025)

³facebook/nllb-200-distilled-1.3B

⁴We do **not** translate text-in-image datasets as that would result in mismatches between the instruction/output language and the English text in the image.

⁵While tasks use different measures, all are on the 0-100% scale, so no task skews the average.

¹google/siglip-so400m-patch14-384

2.2 RQ1: Number of Training Languages

We first investigate on *how many* languages to actually train with: does training on few high-resource languages and (zero-shot) cross-lingual transfer to unseen languages suffice, as suggested, e.g., by Shaham et al. (2024); Chen et al. (2024c); Kew et al. (2024), or do we need to explicitly include each targeted languages? Conversely, does training with more languages harm the per-language performance, with a smaller portion of the training data now allocated to each language?

Setup. We focus on the instruction-tuning step: 50% of the data is kept in English⁶, while the other 50% split between N other languages equally, i.e., each language gets $\frac{50}{N}\%$ of the data budget. We gradually increase N , starting with the highest-resource tier (T5) and then including tiers of lower-resource languages (T4 to T1), one at a time. This results in the following setups: **T5** ($N = 6$), **T5-T4** ($N = 24$), **T5-T3** ($N = 52$), **T5-T2** ($N = 69$), and finally **L100** ($N = 99$). In **L100**, in addition to languages from **T5-T2**, we include T1 languages to cover XM3600 (Thapliyal et al., 2022) and otherwise randomly to reach 99 languages.

Results. Table 1 summarizes the results. Expectedly, we find that including a language (tier) in instruction-tuning improves their performance (Table 1a, top half). Nevertheless, the (negative) effect of adding new languages on performance of previously included languages is negligible, if at all present. This makes training massively multilingual LVLMs feasible with only minor performance drawbacks for any given language. In-language training leads to dramatic improvements in language fidelity (i.e., the model producing the output in the correct language), as shown in Table 1b. Interestingly, the more multilingual the training, the larger the fidelity gains also for languages not included in training; explicit in-language training, expectedly, then further improves fidelity for any given language (see Table 27 in the Appendix for detailed per language results). Even when excluding tasks where language fidelity plays a role (Table 1a bottom), we observe the same trends: steady improvements from in-language training, with negligible (if any) performance drops for other languages. A subset of experiments with Llama 3 (setups: English, T5, and L100) in Table 13 in the Appendix confirms these trends observed with Phi

⁶More specifically, 50% of the 80% of non-text-in-image data, which is excluded from translation.

Train Lang.	T1	T2	T3	T4	T5	en
All tasks						
English	14.4	30.4	24.4	23.6	28.5	53.6
T5	16.5	31.0	26.3	26.7	34.0	53.7
T5-4	17.4	30.6	27.9	29.6	33.5	51.5
T5-3	<u>17.7</u>	31.4	32.1	<u>29.0</u>	<u>34.1</u>	52.7
T5-2	17.0	34.5	30.0	28.2	33.4	54.1
L100	19.3	32.6	30.7	28.9	34.4	52.6
Tasks unaffected by language fidelity						
English	33.0	32.5	36.3	38.5	42.9	55.7
T5	35.3	33.2	36.4	38.7	42.4	56.0
T5-T4	35.8	32.6	37.8	40.1	42.2	55.7
T5-T3	<u>35.9</u>	33.6	40.5	39.7	42.6	56.3
T5-T2	35.2	36.5	38.5	39.5	42.8	55.5
L100	36.1	<u>34.3</u>	<u>39.1</u>	<u>39.8</u>	42.7	54.6

(a) Scores are averaged over results from all tasks grouped by language tier. The performance on the following tasks is affected by language fidelity: XM3600, MaXM, MTVQA.

Train Lang.	T1	T2	T3	T4	T5	en
English	0.2	0.2	0.1	2.4	6.2	100.0
T5	39.1	36.1	82.2	83.9	99.1	100.0
T5-T4	61.8	84.6	87.5	99.2	<u>98.4</u>	100.0
T5-T3	<u>72.9</u>	84.4	98.2	95.2	97.9	100.0
T5-T2	<u>68.5</u>	99.0	97.9	98.4	98.1	100.0
L100	72.9	<u>98.2</u>	95.4	97.8	98.2	100.0

(b) Average language fidelity on XM3600 in %.

Table 1: **RQ1** (§2.2) results for models trained with different sets of languages. We emphasize the **best** and second-best result in each column.

3.5: in fact, we see even larger gains over all tasks when training with more languages.

2.3 RQ2: Language Distribution in Instruction-Tuning

RQ1 experiments show that massively multilingual instruction-tuning data is beneficial across the board. We now analyze *how much* of the training data should be multilingual. On the one hand, intuitively, increasing the non-English portion of the training data budget could then lead to further gains. On the other hand, the gains from more multilingual training are, at some point, likely to be offset by the fact that we are adding noisy (MT-obtained) data at the expense of clean (English) data.

Setup. We opt for the full set of 100 languages (*L100*) in this experiment due to their robust multilingual performance. However, we adjust the language distribution by keeping $E\%$ of the data budget in English and splitting the remaining $100 - E\%$ equally across the other 99 languages⁷. We consider the following six setups: $E \in \{1, 10, 25, 50, 75, 90\}$.

Results. We present the results in Table 2. We ob-

⁷We observed no benefits from an unequal allocation that up-samples low(er)-resource languages; see §D.2)

English %	T1	T2	T3	T4	T5	en
1	19.1	30.3	28.8	27.1	31.7	48.9
10	18.1	32.4	29.4	27.4	32.5	50.1
25	19.7	35.5	29.9	27.9	33.0	50.3
50	<u>19.3</u>	<u>32.6</u>	<u>30.7</u>	28.9	<u>34.4</u>	52.6
75	18.5	31.5	30.7	<u>28.4</u>	34.6	<u>54.1</u>
90	15.9	31.2	27.6	26.9	34.1	54.8

Table 2: **RQ2** (§2.3) results for models trained with different ratios of English to multilingual data in the instruction-tuning phase. Scores are averaged over results from all tasks grouped by language tier.

serve peak performance for all language tiers when between 25% and 75% of the training data is in English. For some tasks (e.g., XM3600, MaXM, BIN-MC), we see weaker performance with more English data, while for others (e.g., MTVQA, xGQA, MaRVL) more multilingual data leads to slight performance drops (see per-task results in F.1). Overall, lower-resource languages benefit from more multilingual data and, conversely, higher-resource languages benefit from more English data. However, this is in part also a consequence of language coverage across tasks: XM3600 and BIN-MC profit from a more multilingual training mix; at the same time, they are the tasks that encompass the most low(er)-resource languages.

Results obtained with the Llama 3 backbone (see Table 14 in the Appendix) follow the same pattern: we observe the best performance in T1 and T2 with $E = 10$; and for T5 and English with $E = 90$; $E = 50$ yields the best results overall, considering all tiers. Our findings align with concurrent work by Yue et al. (2025), who found that anywhere between 20 and 80% of English data yields good global performance. Following these results, we choose $E = 50$ as a robust value for the training.

2.4 RQ3: Language Distribution in Pre-Training

As hinted by (Liu et al., 2023, 2024b) and explicitly demonstrated by Tong et al. (2024), pre-training on image-caption pairs improves the LVLM’s performance. We therefore, after identifying an effective distribution of instruction-tuning data, next explore the effect of different distributions of pre-training data across languages. Specifically, we test if balancing out the English and multilingual portions delivers better performance than unbalanced distributions, that assign more training budget to English or the multilingual mix, respectively.

Setup. In these experiments, we fix the instruction-

English %	T1	T2	T3	T4	T5	en
No pre-training	19.3	32.6	30.7	28.9	34.4	52.6
100	19.3	33.3	32.1	29.4	34.5	55.2
50	22.8	39.5	33.8	<u>30.8</u>	35.7	54.9
1	<u>22.7</u>	<u>38.9</u>	<u>33.7</u>	31.2	<u>35.4</u>	<u>55.1</u>

Table 3: **RQ3** (§2.4) results with different English-to-multilingual ratios ($L100$) for pre-training. All variants are identically instruction-tuned ($L100, 50\% En.$).

tuning mix to $L100$ with $E_{IT} = 50\%$ of data in English, which we found in the previous section to produce overall most balanced results. For the pre-training data mix, we select the same 100 languages, varying the portion of English image-caption pairs, $E_{PT} \in \{100\%, 50\%, 1\%\}$; as in instruction-tuning, the non-English data budget is equally distributed across the other 99 languages.

Results. Scores in Table 3 reveal that while English-only pre-training yields downstream benefits on English tasks, it has a largely negligible effect on other languages. The multilingual mixes substantially improve the performance for virtually all language tiers, with gains being the most prominent for lowest-resource languages from T2 and T1. In contrast to instruction-tuning, a very low proportion of clean English data does not result in tangible performance degradation, but it generally does not improve the multilingual performance either. We thus select $E_{PT} = 50\%$ as the “optimal” choice for subsequent experiments. Experiments with Llama 3, with 1% and 100% of English data (see Table 15 in the Appendix) support this finding that having a highly multilingual pre-training benefits multilingual downstream performance.

2.5 RQ4: Improving on Multilingual Text-in-Image Tasks

Finally, we focus on the models’ multilingual understanding of text in images and how to improve it. Unlike tasks based on natural images, text-in-image tasks cannot be translated trivially from English: even if the prompt and output text are translated, the text in the image remains in English. Because of this, we test how *synthetic* multilingual OCR data, which can be generated at scale in any number of languages, can help improve performance.

Novel Evaluation Dataset. To this end, we introduce **SMPQA** (Synthetic Multilingual Plot QA) a new multilingual evaluation dataset, which focuses on two fundamental skills required in text-in-image tasks: 1) **reading** (and outputting) the text from an

Setup	SMPQA Ground			SMPQA Read		
	en	Latin	other	en	Latin	other
No pre-training	69.6	67.2	51.9	33.4	12.8	0.1
No OCR	76.1	73.0	55.3	41.8	23.1	0.2
100% Eng.	78.4	74.7	57.9	55.8	39.9	3.9
50% Eng.	81.2	76.7	60.0	53.8	41.8	7.1
50% (frozen)	76.1	70.8	56.3	47.2	34.1	3.5
1% Eng.	81.0	78.3	64.1	54.8	43.5	8.0
Latin down	78.9	74.2	59.5	54.6	41.0	9.9

Table 4: **RQ4** (§2.5) results of models trained with additional synthetic OCR data on SMPQA for English, Latin-script languages, and languages with other scripts. **No pre-training:** from Table 2; **No OCR:** from Table 3; **frozen:** image encoder frozen; **N% Eng.:** N% of OCR data is English, rest uniform distributed over L100 languages; **Latin down:** 2.5k samples for all Latin-script languages, 10k samples for others.

image and 2) **grounding** the input text (given as part of the prompt) to the corresponding text in the image (via balanced ‘yes/no’ questions, e.g., “Is the bar with label *\$Label* the largest?”). We provide further details on the construction and examples in Appendix C.5.⁸ We construct SMPQA to cover (i) 5 Latin-script languages, one from each tier, and (ii) 6 major languages with different non-Latin scripts.

Setup. We generate multilingual synthetic text-in-image data for training following the Synthdog approach (Kim et al., 2022) (see B.3 for details). We again adopt the training setup *L100* with 50% English data, both in pre-training and fine-tuning, now adding 500k Synthdog samples to pre-training and a subset of 50k instances to the instruction-tuning mix. As before, we select $E \in \{100, 50, 1\}\%$ English samples, distributing the rest of the budget equally over the other 99 languages. We test an additional *Latin-down* distribution: we double the budget allocated to 32 non-Latin-script languages and cut the training budget for Latin-script languages (other than English) in half. Importantly, in these experiments we *unfreeze* the image encoder and fine-tune its parameters as well.

Results. Table 4 summarizes the results. The models from prior experiments, *No pre-training* and *No OCR*, succeed for English and other Latin-script languages but utterly fail on non-Latin scripts with near-random performance. We note that the model with the pre-training step (without the additional OCR data) already performs better than the model trained just via instruction-tuning; this is likely due

⁸While MTVQA and M3Exam also require OCR capabilities, they also require input image resolution that is far greater than what we use in our experiments (384px); SMPQA uses bigger letters, making performance effects from multilingual training on text-in-image understanding easier to measure.

to the presence of images with text coupled with captions that explicitly mention this text. Training with synthetic data greatly improves the performance across all languages even if all of the OCR data is in English (100% Eng.). Nonetheless, using multilingual synthetic OCR data is very effective and, importantly, does not degrade English SMPQA performance even if English constitutes only 1% of the training data. We note that *unfreezing* and training the image encoder is critical for optimal performance in *all scripts*. Despite all this, we still observe a large performance gap between Latin-and non-Latin-script languages, even if we skew the training budget towards the non-Latin scripts (*Latin-down*). We hypothesize that orders of magnitude more text-in-image training data for other scripts are required for adequate performance.⁹

3 Centurio: Applying Lessons Learned

Our answers to **RQ1–RQ4** (see §2) point to the feasibility of training massively multilingual LVLMs supporting 100 languages with a “sweet spot” of roughly 50% of the English data being MT-translated to the languages covered. For improving multilingual OCR capabilities, training on large-scale synthetic data with an unfrozen image encoder has proven effective. Demonstrating the practicability of our findings, we now train state-of-the-art multilingual LVLMs applying our lessons learned, which we call **Centurio**. We briefly describe further design choices below.

3.1 Design Choices

Text Encoder. The choice of the LLM greatly impacts multilingual performance. We benchmark several LLMs (with 7–9B parameters) following the evaluation setup described in §2 for L100 languages and translations for 50% of the English instruct data to find candidates for **Centurio** (details in Appendix D.3). The best performances were obtained with Aya-Expanse (Dang et al., 2024) and Qwen 2.5 (Yang et al., 2024) as backbones.

Image Tiling and Projection. Image tiling methods (Lin et al., 2024; Liu et al., 2024b) increase the image resolution by concatenating encodings of n non-overlapping tiles of an input image together, which significantly helps with ‘reading’ small text

⁹Concurrently, in a preliminary exploration of text-in-image capabilities, Yue et al. (2025) noted steady gains with 50k samples per language but also observed worse performance for non-Latin-script languages.

	AVG.	XM3600			MT-VQA	SMPQA G.		SMPQA N.		M3Exam		xMMMU		C-VQA
		en	mul	fid.		en	mul	en	mul	en	mul	en	mul	
Parrot	25.8	5.6	0.4	25.0	2.0	51.0	49.9	0.0	0.0	46.6	36.2	35.3	32.4	41.1
PALO 7B	28.7	65.9	13.5	72.0	5.8	55.5	52.8	22.4	2.7	41.0	29.1	31.8	30.9	37.1
PALO 13B	29.9	67.3	17.0	60.1	6.3	54.0	51.5	25.6	4.0	45.2	28.3	32.4	28.9	39.6
Llama-Vision 3.2 11B	*32.3	35.9	7.2	33.3	15.2	91.1	84.8	58.4	22.8	—	—	—	—	38.8
Maya	33.4	55.9	14.6	65.7	5.3	51.4	50.9	14.6	1.8	49.2	36.3	37.9	33.3	39.8
Pixtral 12B	38.1	26.5	22.1	96.8	14.1	91.1	71.0	85.0	35.9	49.4	33.7	30.3	26.2	33.5
Phi 3.5 Vision	39.5	32.3	6.3	40.8	11.1	92.2	79.4	84.8	35.9	56.3	40.7	41.7	37.4	40.9
Qwen2VL 2B	41.2	68.8	5.2	13.2	19.0	85.0	83.5	68.8	47.4	47.9	40.5	36.8	35.5	33.6
MiniCPM 2.6	41.7	87.5	14.2	92.3	16.1	89.0	74.3	80.8	39.3	55.0	48.2	39.1	36.5	34.1
InternVL 2.5 4B	45.3	38.9	17.5	91.0	25.1	87.0	78.3	77.8	47.5	63.2	50.3	49.2	42.7	48.1
InternVL 2.5 8B	47.4	38.3	15.7	91.1	25.0	91.0	79.2	80.6	48.2	67.0	53.3	50.7	45.2	48.6
Qwen2VL 7B	47.7	50.3	24.6	90.0	23.2	91.2	90.9	85.0	64.9	56.1	49.7	43.0	40.7	37.6
Pangea	48.2	70.1	34.6	87.9	19.3	87.2	72.2	72.0	23.8	58.0	45.5	43.1	42.0	55.2
Centurio Aya	48.5	78.4	39.2	95.7	11.1	83.1	74.2	60.0	30.1	53.0	41.2	37.6	37.2	49.4
Centurio Qwen	51.6	79.1	34.4	95.2	11.9	84.8	76.1	65.2	31.7	61.2	46.9	46.4	43.0	52.9
	MAXM	xGQA	BIN-MC	XVNLI	MaRVL	VGR	VLOD							
	en	mul												
Parrot	28.2	3.6	37.7	21.2	30.5	25.7	28.7	31.4	63.5	55.1	59.2	52.9	0.0	0.0
PALO 7B	54.0	22.5	59.1	36.6	58.7	38.6	58.0	53.4	62.7	24.1	48.3	25.6	5.8	6.8
PALO 13B	51.7	33.1	58.0	27.8	61.4	41.1	56.6	53.6	63.8	33.1	63.3	26.2	2.5	4.9
Llama-Vision 3.2 11B	0.0	4.7	39.3	27.6	75.6	50.8	—	—	—	—	—	—	—	—
Maya	55.4	17.3	58.2	49.1	54.0	43.2	50.1	43.9	60.3	56.3	46.7	42.3	20.0	20.1
Pixtral 12B	59.4	43.4	59.9	3.8	71.0	54.2	60.9	52.7	67.7	60.7	55.8	47.7	9.2	12.4
Phi 3.5 Vision	43.6	17.9	65.2	38.0	63.1	36.8	58.9	53.3	73.4	46.4	81.7	50.3	45.8	31.5
Qwen2VL 2B	53.7	26.5	60.5	38.2	78.2	47.2	61.9	56.2	67.9	55.9	61.7	50.5	22.5	20.4
MiniCPM 2.6	53.4	22.3	57.9	45.7	72.6	47.4	71.9	65.4	70.2	57.9	52.5	49.1	9.2	14.6
InternVL 2.5 4B	46.0	42.5	63.6	28.0	68.4	45.4	69.0	58.7	74.9	59.0	72.5	49.7	24.2	21.0
InternVL 2.5 8B	45.6	38.2	63.4	32.0	70.3	44.2	73.5	66.4	83.0	63.3	87.5	51.6	57.5	29.0
Qwen2VL 7B	54.7	31.2	62.5	49.3	80.7	57.5	62.1	59.6	69.8	60.2	60.0	52.9	5.8	13.2
Pangea	61.4	55.0	64.6	60.4	70.3	52.1	69.0	65.2	75.8	70.5	69.2	58.9	0.0	6.7
Centurio Aya	55.7	49.3	59.1	53.2	69.7	54.7	65.0	62.4	85.0	77.9	82.5	66.8	12.5	20.7
Centurio Qwen	60.1	47.7	60.6	54.8	72.7	56.2	75.4	70.2	89.6	81.7	87.5	73.1	28.3	27.0

Table 5: Comparison of Centurio and 13 other LVLMs across 14 tasks. We highlight the **best** and second-best results. Scores are accuracy (CIDEr for XM3600). **en** & **mul** are the English and averaged multilingual results. **XM3600 fid.** is the language fidelity over all languages; **SMPQA G. & N** are Grounding and Naming. *: supports only single-image input. **AVG.:** average over all tasks. Details on the setup and models are provided in Appendix C.

Model	T1	T2	T3	T4	T5	en
Centurio Aya	35.1	46.4	47.0	46.7	48.3	60.6
Centurio Qwen	38.1	51.0	48.3	47.0	50.9	66.6
InternVL 2.5 8B	29.9	37.0	37.4	41.0	50.5	64.4
Qwen2VL 7B	30.6	36.8	40.5	46.2	48.0	56.8
Pangea	38.5	38.6	46.9	44.2	49.9	59.8
Without multi-image tasks (MaRVL, VGR, VLOD):						
Centurio Aya	35.1	44.5	45.7	46.2	47.7	60.7
Centurio Qwen	38.1	49.5	45.6	45.8	49.6	66.0
InternVL 2.5 8B	29.9	40.4	35.2	39.4	49.7	62.3
Qwen2VL 7B	30.6	38.7	40.8	46.8	48.3	61.7
Pangea	38.5	46.5	47.7	44.4	49.9	64.9

Table 6: Comparison between Centurio and the top-3 models of Table 5. Scores are averages over results from all 14 tasks grouped by language tier.

in images. However, they also greatly increase the input length: a 2×2 tiling would yield 3,645 tokens per image with our model.¹⁰ Instead, we adopt the method by Shi et al. (2024), which concatenates the tokens of the whole image and the tiles along the *feature dimension* before projection by the MLP. This gives an efficient trade-off between computing cost—the number of tokens stays constant—and performance gains for fine-grained content.

¹⁰The whole image plus four tiles, each with 729 tokens.

Training Data. We increase the amount of the pre-training and instruct tuning data to further improve performance beyond our analysis setup. For pre-training, we add the 0.7M ALLaVA captions (Chen et al., 2024a) to the ShareGPT-4V captions and we use all synthetic OCR data generated in §2.5 (1.16M total: 500k English, 5k for Latin-script language, 10k for other scripts). For instruction-tuning, we incorporate additional datasets from the Cambrian collection (Tong et al., 2024) along with several text-only instruction-tuning datasets (see Appendix B.2 for a list). We translate the data to the *L100 50% En.* setup, excluding text-heavy datasets and others that are problematic for MT.

3.2 Results

We compare our Centurio models against 13 other multilingual LVLMs across the 13 tasks used in §2, and additionally evaluate them on CVQA¹¹, testing the models’ capabilities across 56 languages. We provide details for all models in Appendix C.6.

On average, Centurio achieves the best results

¹¹CVQA has a private test set and only allows limited submissions hence we left it out for our analysis experiments.

across all 14 tasks (cf. Table 5). Focusing on multilingual single-task performance, our model sometimes lags behind the best-performing models but still achieves best or second-best results in 9 of 14 tasks. While Centurio performs strongly on English, it is more often surpassed by other models. These results prove the effectiveness of our training composition: we are able to retain high English performances while maximizing the models’ multilingual capabilities. When analyzing these results grouped by language tier (Table 6), we find that our models shine in the low-resource tiers T1¹² and T2, with competitive results for higher-resource languages—even when excluding multi-image tasks (VGR, MaRVL, VLOD), where our models greatly outperform most others.

For text-heavy tasks (primarily MTVQA and SMPQA), we find that Centurio falls behind. While we show the importance of multilingual OCR training—Centurio succeeds at the SMPQA reading task in more languages than, for example, Pangea—the limited input resolution and magnitudes less OCR data compared to Qwen2-VL and others result in comparably poor performance.

4 Related Work

Multilingual LVLMs. Building on the success of monolingual LVLMs like BLIP-2 (Li et al., 2023) and LLaVA (Liu et al., 2023, 2024a), researchers extended the English training protocols to include multilingual data for obtaining massively multilingual LVLMs (e.g., Maaz et al., 2025; Geigle et al., 2023). As such, Google’s PaLI models (Chen et al., 2022, 2023) were the first closed-weight models trained on multilingual captions and VQA data with the recent open-weight PaliGemma (Beyer et al., 2024) following a similar training strategy. Geigle et al. (2023) presented with mBLIP the first open model, trained with image captions and a limited mix of instruct data translated to 98 languages. Subsequent models similarly followed an established procedure by directly translating parts of the English training data (Maaz et al., 2025; Hu et al., 2024; Alam et al., 2024). For the concurrent Pangea, Yue et al. (2025) optimized for multicultural aspects and used a mix of machine-translated data, existing multilingual data, and synthetically generated data. While they analyze the ratio between English and multilingual data, they do not

¹²Despite 4 of 7 T1 CVQA languages **not** in our training data.

vary the number of languages, fixing it at 39. Interestingly, most researchers either (i) did not properly motivate their multilingual data mix (e.g., Geigle et al., 2023; Alam et al., 2024; Beyer et al., 2024), or (ii) did not provide any details on the training data composition (e.g., Wang et al., 2024b; Yao et al., 2024; Chen et al., 2024d))

Multilingual OCR with LVLMs. While OCR recently gained popularity for English LVLMs (Lu et al., 2024; Tong et al., 2024), multilingual OCR has largely been neglected in prior work. As an exception, Qwen2-VL (Wang et al., 2024b) and InternVL 2.5 (Chen et al., 2024d) exhibit excellent multilingual OCR capabilities, but no training details are known. Towards open knowledge on improving multilingual OCR, Yue et al. (2025) performed preliminary experiments leveraging data in 10 languages. However, such efforts are still hindered by the lack of evaluation resources: MTVQA (Tang et al., 2024) and M3Exam (Zhang et al., 2023a) only cover up to 9 languages and conflate language understanding (in the text input) with understanding text on images. In this work, we push multilingual OCR research by presenting the novel SMPQA dataset dedicated to evaluation of multilingual OCR. We further explore how synthetic training data can improve models’ capabilities.

Multilingual Instruction Tuning of LLMs.. While older LLMs struggled in multilingual tasks (Ahuja et al., 2024), more recent ones like Qwen 2.5 (Yang et al., 2024), Llama 3 (Llama Team, 2024), Gemma 2 (Team et al., 2024), or Aya (Aryabumi et al., 2024) have greatly improved in that respect, making them usable in many languages besides English. Still, current LLMs often fail to respond faithfully to the prompting language if that language is not English, especially for low-resource languages (Holtermann et al., 2024; Kew et al., 2024; Marchisio et al., 2024). To mitigate this issue, several works have analyzed the importance of multilingual instruction tuning. Weber et al. (2024) demonstrated that multilingual training is crucial for downstream performance even if the base models are pre-trained on multilingual data mixtures. Others showed that just a small set of languages is sufficient to improve cross-lingual transfer for multilingual downstream tasks significantly (Shaham et al., 2024; Chen et al., 2024c; Kew et al., 2024). However, they focus on a small set of primarily higher-resource languages, while we consider the problem in the vision-language

context for a wider language selection.

In (Soykan and Sahin, 2024), the authors propose methods to select the optimal mix of languages for instruction tuning in a "linguistically-informed manner". However, they find no general best selection, and instead a task- and model-dependent selection is necessary. Therefore, in our work, we do not apply these techniques and instead choose languages based on the taxonomy introduced by Joshi et al. (2020).

5 Conclusion

In this study, we systematically investigated the optimal data composition for training a multilingual LVLM through four progressively refined analysis setups. Our findings reveal that massively multilingual training with 100 languages is highly effective, achieving comparable results to configurations with fewer languages. Moreover, only 25–50% of the training data needs to be non-English, keeping the cost of multilingual data production low. To enhance multilingual text understanding in images, we introduced a novel evaluation benchmark and demonstrated the importance and effectiveness of including multilingual synthetic OCR data in the training mix. Finally, we apply our findings to train Centurio, massively multilingual LVLMs trained with 100 languages, and achieve state-of-the-art results on our evaluation suite covering 14 tasks and 56 language tasks against 13 other LVLMs.

6 Limitations

Lack of Explicit Multicultural Training The focus of this work is on *language understanding* in a massively multilingual setup, that is, how to train the model to maximize its ability to understand and generate text in various languages. We do not consider the multicultural aspect, that is, training a model so that it is also more knowledgeable about concepts from the countries whose languages it can understand as measured by benchmarks like CVQA or CulturalVQA (Nayak et al., 2024). While the two aspects — multilingual and multicultural knowledge — can be intermingled in practice, they require distinct approaches in training: Multilingual data is necessary for multilingual language understanding, as we have shown. However, multicultural knowledge can be learned from multilingual resources as created by, for example, by Yue et al. (2025), but also from fully English resources like Wikipedia (Srinivasan et al., 2021).

Using Machine-Translated Training Data We train our model using machine-translated (MT) data derived from high-quality English datasets. This is advantageous because it allows us to create comparable setups for our analyses with full control over the languages and their proportions. While the data proves effective in increasing multilingual performance, MT data, especially for low-resource languages, can be of low quality and, even in higher-resource languages, might exhibit unwanted "translationese" artifacts. This can negatively impact the quality of generated text in a way that the metrics employed in our evaluation suite do not adequately measure. While native multilingual training data is available, it is not available for all tasks or languages equally, or, for most languages, not at all. Future work should consider evaluation setups to quantify the effect the MT data has on the final model, work on better MT pipelines, or create more data through native speakers.

Using Synthetically Generated OCR Data The text-heavy, "real-world" tasks in some datasets of our instruction tuning mix, which cover diverse image types such as plots, scans, application screenshots, or screenshots of webpages, are still entirely in English. Due to the issues that arise when translating such samples, we do not translate them. Hence, our methods to improve the understanding of multilingual texts in images are limited to only using synthetically generated images. While we have seen that our synthetic data positively impacts the performance of models on the respective tasks, future work should explore methods for collecting or generating more diverse data in different languages beyond our synthetic OCR data.

Another limitation regarding OCR capabilities is our relatively small image input resolution compared to models like Qwen2-VL or InternVL 2.5 — both of which support image inputs in native resolution at the cost of thousands of tokens per image —, which limits the performance of Centurio for images with small text.

Acknowledgments

Simulations were performed with computing resources granted by WestAI under project 9148.

Simulations were performed with computing resources from Julia 2. Julia 2 was funded as DFG project as "Forschungsgroßgerät nach Art 91b GG" under INST 93/1145-1 FUGG

The work of Gregor Geigle was in part supported by the Alexander von Humboldt Foundation.

The work of Carolin Holtermann and Anne Lauscher is funded under the Excellence Strategy of the German Federal Government and the States.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024a. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. *arXiv preprint arXiv:2404.14219*.
- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Parul Chopra, Allie Del Giorno, Gustavo de Rosa, Matthew Dixon, Ronen Eldan, Dan Iter, Amit Garg, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Jamie Huynh, Mojgan Javaheripi, Xin Jin, Piero Kauffmann, Nikos Karampatziakis, Dongwoo Kim, Ma-houd Khademi, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Chen Liang, Weishung Liu, Eric Lin, Zeqi Lin, Piyush Madan, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Corby Rosset, Sam-budha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Xia Song, Masahiro Tanaka, Xin Wang, Rachel Ward, Guanhua Wang, Philipp Witte, Michael Wyatt, Can Xu, Jiahang Xu, Sonali Yadav, Fan Yang, Ziyi Yang, Donghan Yu, Chengrudong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyra Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024b. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. *arXiv preprint arXiv:2404.14219*.
- Manoj Acharya, Kushal Kafle, and Christopher Kanan. 2019. [TallyQA: Answering Complex Counting Questions](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8076–8084.
- Željko Agić and Natalie Schluter. 2018. Baselines and Test Data for Cross-Lingual Inference. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 3890–3894, Miyazaki, Japan.
- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, et al. 2024. Pixtral 12B. *arXiv preprint arXiv:2410.07073*.
- Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, Ishaan Watts, Ashutosh Sathe, Millicent Ochieng, Rishav Hada, Prachi Jain, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2024. [MEGAVERSE: Benchmarking Large Language Models Across Languages, Modalities, Models and Tasks](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2598–2637, Mexico City, Mexico.
- AI@Meta. 2024. [Llama 3 Model Card](#).
- Nahid Alam, Karthik Reddy Kanjula, Surya Guthikonda, Timothy Chung, Bala Krishna S Vegevna, Abhipsha Das, Anthony Susevski, Ryan Sze-Yin Chan, SM Uddin, Shayekh Bin Islam, et al. 2024. Maya: An Instruction Finetuned Multilingual Multimodal Model. *arXiv preprint arXiv:2412.07112*.
- Viraat Aryabumi, John Dang, Dwarak Talupuru, Saurabh Dash, David Cairuz, Hangyu Lin, Bharat Venkitesh, Madeline Smith, Jon Ander Campos, Yi Chern Tan, Kelly Marchisio, Max Bartolo, Sebastian Ruder, Acyr Locatelli, Julia Kreutzer, Nick Frosst, Aidan N. Gomez, Phil Blunsom, Marzieh Fadaee, Ahmet Üstün, and Sara Hooker. 2024. Aya 23: Open Weight Releases to Further Multilingual Progress. *arXiv preprint arXiv:2405.15032*.
- Jonas Belouadi, Anne Lauscher, and Steffen Eger. 2024. [AutomaTikZ: Text-Guided Synthesis of Scientific Vector Graphics with TikZ](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024*, Vienna, Austria.
- Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, Thomas Unterthiner, Daniel Keysers, Skanda Koppula, Fangyu Liu, Adam Grycner, Alexey A. Gritsenko, Neil Houlsby, Manoj Kumar, Keran Rong, Julian Eisen-schlos, Rishabh Kabra, Matthias Bauer, Matko Bosnjak, Xi Chen, Matthias Minderer, Paul Voigtlaender, Ioana Bica, Ivana Balazevic, Joan Puigcerver, Pinelopi Papalampidi, Olivier J. Hénaff, Xi Xiong, Radu Soriciu, Jeremiah Harmsen, and Xiaohua Zhai. 2024. [PaliGemma: A versatile 3B VLM for transfer](#). *arXiv preprint arXiv:2407.07726*.
- Ali Furkan Biten, Rubèn Tito, Andrés Mafla, Lluís Gómez i Bigorda, Marçal Rusiñol, C. V. Jawahar, Ernest Valveny, and Dimosthenis Karatzas. 2019. [Scene Text Visual Question Answering](#). In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019*, pages 4290–4300, Seoul, Korea (South).
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, Virtual.
- Emanuele Bugliarello, Fangyu Liu, Jonas Pfeiffer, Siva Reddy, Desmond Elliott, Edoardo Maria Ponti, and Ivan Vulic. 2022. IGLUE: A Benchmark for Transfer Learning across Modalities, Tasks, and Languages. In *International Conference on Machine Learning*, pages 2370–2392, Baltimore, MD, USA. PMLR.
- Soravit Changpinyo, Linting Xue, Michal Yarom, Ashish Thapliyal, Idan Szpektor, Julien Amelot, Xi Chen, and Radu Soricut. 2023. **MaXM: Towards Multilingual Visual Question Answering**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2667–2682, Singapore.
- Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. 2024a. **ALLaVA: Harnessing GPT4V-synthesized Data for A Lite Vision-Language Model**. *arXiv preprint arXiv:2402.11684*.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. 2024b. ShareGPT4V: Improving Large Multi-modal Models with Better Captions. In *European Conference on Computer Vision*, pages 370–387, Milan, Italy. Springer.
- Pinzhen Chen, Shaoxiong Ji, Nikolay Bogoychev, Andrey Kutuzov, Barry Haddow, and Kenneth Heafield. 2024c. **Monolingual or Multilingual Instruction Tuning: Which Makes a Better Alpaca**. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1347–1356, St. Julian’s, Malta.
- Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, Siamak Shakeri, Mostafa Dehghani, Daniel Salz, Mario Lucic, Michael Tschannen, Arsha Nagrani, Hexiang Hu, Mandar Joshi, Bo Pang, Ceslee Montgomery, Paulina Pietrzyk, Marvin Ritter, A. J. Piergiovanni, Matthias Minderer, Filip Pavetic, Austin Waters, Gang Li, Ibrahim Alabdulmohsin, Lucas Beyer, Julien Amelot, Kenton Lee, Andreas Peter Steiner, Yang Li, Daniel Keysers, Anurag Arnab, Yuanzhong Xu, Keran Rong, Alexander Kolesnikov, Mojtaba Seyedhosseini, Anelia Angelova, Xiaohua Zhai, Neil Houlsby, and Radu Soricut. 2023. **PaLI-X: On Scaling up a Multilingual Vision and Language Model**. *arXiv preprint arXiv:2305.18565*.
- 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 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. In *The Eleventh International Conference on Learning Representations*, Kigali, Rwanda.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yimin Ren, Zixuan Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, Conghui He, Botian Shi, Xingcheng Zhang, Han Lv, Yi Wang, Wenqi Shao, Pei Chu, Zhongying Tu, Tong He, Zhiyong Wu, Huipeng Deng, Jiaye Ge, Kai Chen, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhui Wang. 2024d. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. *arXiv preprint arXiv:2412.05271*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024e. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. *arXiv preprint arXiv:2412.05271*.
- Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan R. Routledge, and William Yang Wang. 2021. **FinQA: A Dataset of Numerical Reasoning over Financial Data**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021*, pages 3697–3711, Virtual Event / Punta Cana, Dominican Republic.
- Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, and Dongmei Zhang. 2022. **HiTab: A Hierarchical Table Dataset for Question Answering and Natural Language Generation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, pages 1094–1110, Dublin, Ireland.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. **Unsupervised Cross-lingual Representation Learning at Scale**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pages 8440–8451, Online.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barraud, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram

- Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. **No Language Left Behind: Scaling Human-Centered Machine Translation**. *arXiv preprint arXiv:2207.04672*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. 2023. **InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning**. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, volume abs/2305.06500, pages 49250–49267, New Orleans, LA, USA.
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, Acyr Locatelli, Hangyu Lin, Dwarak Talupuru, Bharat Venkitesh, David Cairuz, Bowen Yang, Tim Chung, Wei-Yin Ko, Sylvie Shang Shi, Amir Shukayev, Sammie Bae, Aleksandra Piktus, Roman Castagné, Felipe Cruz-Salinas, Eddie Kim, Lucas Crawhall-Stein, Adrien Morisot, Sudip Roy, Phil Blunsom, Ivan Zhang, Aidan Gomez, Nick Frosst, Marzieh Fadaee, Beyza Ermis, Ahmet Üstün, and Sara Hooker. 2024. **Aya Expanse: Combining Research Breakthroughs for a New Multilingual Frontier**. *arXiv preprint arXiv:2412.04261*.
- J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. 2009. **ImageNet: A large-scale hierarchical image database**. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, Miami, FL, USA.
- Peter Devine. 2024. **Tagengo: A Multilingual Chat Dataset**. In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*, pages 106–113, Miami, FL, USA.
- Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wan-jun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, and Lingpeng Kong. 2023. **G-LLaVA: Solving Geometric Problem with Multi-Modal Large Language Model**. In *The Thirteenth International Conference on Learning Representations*, Singapore.
- William Gaviria Rojas, Sudnya Diamos, Keertan Kini, David Kanter, Vijay Janapa Reddi, and Cody Coleman. 2022. The Dollar Street Dataset: Images Representing the Geographic and Socioeconomic Diversity of the World. In *Advances in Neural Information Processing Systems*, volume 35, pages 12979–12990.
- Gregor Geigle, Abhay Jain, Radu Timofte, and Goran Glavas. 2023. **mBLIP: Efficient Bootstrapping of Multilingual Vision-LLMs**. *arXiv preprint arXiv:2307.06930*.
- Gregor Geigle, Radu Timofte, and Goran Glavas. 2024a. **African or European Swallow? Benchmarking Large Vision-Language Models for Fine-Grained Object Classification**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024*, pages 2653–2669, Miami, FL, USA.
- Gregor Geigle, Radu Timofte, and Goran Glavaš. 2024b. **Babel-ImageNet: Massively Multilingual Evaluation of Vision-and-Language Representations**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5064–5084, Bangkok, Thailand.
- Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. 2017. **Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6325–6334, Honolulu, HI, USA.
- 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 *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, pages 3608–3617, Salt Lake City, UT, USA.
- Musashi Hinck, Carolin Holtermann, Matthew L. Olson, Florian Schneider, Sungduk Yu, Anahita Bhiwandiwala, Anne Lauscher, Shao-Yen Tseng, and Vasudev Lal. 2024. **Why do LLaVA Vision-Language Models Reply to Images in English?** In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13402–13421, Miami, FL, USA.
- Carolin Holtermann, Paul Röttger, Timm Dill, and Anne Lauscher. 2024. **Evaluating the Elementary Multilingual Capabilities of Large Language Models with MultiQ**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4476–4494, Bangkok, Thailand.
- Yu-Chung Hsiao, Fedir Zubach, Gilles Baechler, Srinivas Sunkara, Victor Carbune, Jason Lin, Maria Wang, Yun Zhu, and Jindong Chen. 2025. **ScreenQA: Large-Scale Question-Answer Pairs Over Mobile App Screenshots**. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9427–9452, Albuquerque, New Mexico, USA.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. **LoRA: Low-Rank Adaptation of Large Language Models**. In *The Tenth International Conference on Learning Representations, ICLR 2022*, Virtual.

- Jinyi Hu, Yuan Yao, Chongyi Wang, Shan Wang, Yinxu Pan, Qianyu Chen, Tianyu Yu, Hanghao Wu, Yue Zhao, Haoye Zhang, Xu Han, Yankai Lin, Jiao Xue, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. **Large Multilingual Models Pivot Zero-Shot Multi-modal Learning across Languages**. In *The Twelfth International Conference on Learning Representations, ICLR 2024*, Vienna, Austria.
- Drew A. Hudson and Christopher D. Manning. 2019. **GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019*, pages 6700–6709, Long Beach, CA, USA.
- Harsh Jhamtani and Taylor Berg-Kirkpatrick. 2018. **Learning to Describe Differences Between Pairs of Similar Images**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4024–4034, Brussels, Belgium.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. 2017. **CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, pages 1988–1997, Honolulu, HI, USA.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. **The State and Fate of Linguistic Diversity and Inclusion in the NLP World**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, pages 6282–6293, Online.
- Kushal Kafle, Brian L. Price, Scott Cohen, and Christopher Kanan. 2018. **DVQA: Understanding Data Visualizations via Question Answering**. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018*, pages 5648–5656, Salt Lake City, UT, USA.
- Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa Sadigh. 2024. **Prismatic VLMs: investigating the design space of visually-conditioned language models**. In *Proceedings of the 41st International Conference on Machine Learning*, pages 23123–23144, Vienna, Austria.
- Mehran Kazemi, Hamidreza Alvari, Ankit Anand, Jialin Wu, Xi Chen, and Radu Soricut. 2023. **GeomVerse: A Systematic Evaluation of Large Models for Geometric Reasoning**. In *AI for Math Workshop@ ICML 2024*.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. **ReferItGame: Referring to Objects in Photographs of Natural Scenes**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 787–798, Doha, Qatar.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. **A Diagram is Worth a Dozen Images**. In *Computer Vision – ECCV 2016*, pages 235–251, Amsterdam, The Netherlands.
- Aniruddha Kembhavi, Min Joon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. **Are You Smarter Than a Sixth Grader? Textbook Question Answering for Multimodal Machine Comprehension**. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, pages 5376–5384, Honolulu, HI, USA.
- Tannon Kew, Florian Schottmann, and Rico Sennrich. 2024. **Turning English-centric LLMs Into Polyglots: How Much Multilinguality Is Needed?** In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13097–13124, Miami, Florida, USA.
- Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. 2022. **OCR-Free Document Understanding Transformer**. In *Computer Vision - ECCV 2022 - 17th European Conference*, volume 13688, pages 498–517, Tel Aviv, Israel.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. **Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations**. *Int. J. Comput. Vision*, 123(1):32–73.
- Hugo Laurençon, Léo Tronchon, and Victor Sanh. 2024. **Unlocking the conversion of Web Screenshots into HTML Code with the WebSight Dataset**. *arXiv preprint arXiv:2403.09029*.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2025. **What matters when building vision-language models?** In *Proceedings of the Winter Conference on Applications of Computer Vision*, pages 1372–1381.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. **BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models**. In *International Conference on Machine Learning*, pages 19730–19742, Honolulu, HI, USA. PMLR.
- Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. 2024. **Multi-modal ArXiv: A Dataset for Improving Scientific Comprehension of Large Vision-Language Models**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024*, pages 14369–14387, Bangkok, Thailand.
- Tsung-Yi Lin, Michael Maire, Serge J. 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*, volume 8693, pages 740–755, Zurich, Switzerland.
- Ziyi Lin, Dongyang Liu, Renrui Zhang, Peng Gao, Longtian Qiu, Han Xiao, Han Qiu, Wenqi Shao, Keqin Chen, Jiaming Han, Siyuan Huang, Yichi Zhang, Xuming He, Yu Qiao, and Hongsheng Li. 2024. **SPHINX: A Mixer of Weights, Visual Embeddings and Image Scales for Multi-modal Large Language Models**. In *Computer Vision - ECCV 2024 - 18th European Conference*, volume 15120, pages 36–55, Milan, Italy.
- Adam Dahlgren Lindström and Savitha Sam Abraham. 2022. **CLEVR-Math: A Dataset for Compositional Language, Visual and Mathematical Reasoning**. In *Proceedings of the 16th International Workshop on Neural-Symbolic Learning and Reasoning as part of the 2nd International Joint Conference on Learning & Reasoning (IJCLR 2022)*, volume 3212 of *CEUR Workshop Proceedings*, pages 155–170, Windsor Great Park, UK.
- 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, EMNLP 2021*, pages 10467–10485, Virtual Event / Punta Cana, Dominican Republic.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved Baselines with Visual Instruction Tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020*, pages 26296–26306, Seattle, WA, USA.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024b. **LLaVA-NeXT: Improved reasoning, OCR, and world knowledge**.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. **Visual Instruction Tuning**. In *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916.
- Llama Team. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled Weight Decay Regularization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, Yaofeng Sun, Chengqi Deng, Hanwei Xu, Zhenda Xie, and Chong Ruan. 2024. **DeepSeek-VL: Towards Real-World Vision-Language Understanding**. *arXiv preprint arXiv:2403.05525*.
- Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, and Song-Chun Zhu. 2021a. **Inter-GPS: Interpretable Geometry Problem Solving with Formal Language and Symbolic Reasoning**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers)*, pages 6774–6786, Virtual Event.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. **Learn to Explain: Multimodal Reasoning via Thought Chains for Science Question Answering**. In *Advances in Neural Information Processing Systems 35, NeurIPS 2022*, New Orleans, LA, USA.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. 2023. **Dynamic Prompt Learning via Policy Gradient for Semi-structured Mathematical Reasoning**. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda*.
- Pan Lu, Liang Qiu, Jiaqi Chen, Tanglin Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. 2021b. **IconQA: A New Benchmark for Abstract Diagram Understanding and Visual Language Reasoning**. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021*, Virtual.
- Muhammad Maaz, Hanoona Rasheed, Abdelrahman Shaker, Salman Khan, Hisham Cholakal, Rao M Anwer, Tim Baldwin, Michael Felsberg, and Fahad S Khan. 2025. **Palo: A Polyglot Large Multimodal Model for 5B People**. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1745–1754, Tucson, Arizona, USA.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. 2016. **Generation and Comprehension of Unambiguous Object Descriptions**. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, pages 11–20, Las Vegas, NV, USA.
- Kelly Marchisio, Wei-Yin Ko, Alexandre Berard, Théo Dehaze, and Sebastian Ruder. 2024. **Understanding and Mitigating Language Confusion in LLMs**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024*, pages 6653–6677, Miami, FL, USA.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. **OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge**. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3190–3199, Long Beach, CA, USA.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq R. Joty, and Enamul Hoque. 2022. **ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning**. In *Findings of the As-*

- sociation for Computational Linguistics: ACL 2022*, pages 2263–2279, Dublin, Ireland.
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and C. V. Jawahar. 2022. **InfographicVQA**. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022*, pages 2582–2591, Waikoloa, HI, USA.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. **DocVQA: A Dataset for VQA on Document Images**. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2021*, pages 2199–2208, Waikoloa, HI, USA.
- Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. 2019. **OCR-VQA: Visual Question Answering by Reading Text in Images**. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019*, pages 947–952, Sydney, Australia.
- 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, EMNLP 2024*, pages 5769–5790, Miami, FL, USA.
- Jason Obeid and Enamul Hoque. 2020. **Chart-to-Text: Generating Natural Language Descriptions for Charts by Adapting the Transformer Model**. In *Proceedings of the 13th International Conference on Natural Language Generation, INLG 2020*, pages 138–147, Dublin, Ireland.
- Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin O. Steitz, Stefan Roth, Ivan Vulic, and Iryna Gurevych. 2022a. **xGQA: Cross-Lingual Visual Question Answering**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2497–2511, Dublin, Ireland.
- Jonas Pfeiffer, Naman Goyal, Xi Lin, Xian Li, James Cross, Sebastian Riedel, and Mikel Artetxe. 2022b. **Lifting the Curse of Multilinguality by Pre-training Modular Transformers**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3479–3495, Seattle, WA, USA.
- 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, Joelyn Dunstan, Laura Alonso Alemany, Kumaranage Ravindu Yasas Nagasinghe, Luciana Benotti, Luis Fernando D' Haro, Marcelo Viridiano, Marcos Estecha-Garitagoitia, Maria Camila Buitrago Cabrera, Mario Rodríguez-Cantelar, Mélanie Jouitteau, Mihail Mihaylov, Naome Etori, Mohamed Faizli 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 Kurabayashi, Teresa Clifford, Thanmay Jayakumar, Tiago Timponi Torrent, Toqueer Ehsan, Vladimir Araujo, Yova Kementchedjhieva, Zara Burzo, Zheng Wei Lim, Zheng Xin Yong, Oana Ignat, Joan Nwatu, Rada Mihalcea, Thamar 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, Vancouver, Canada.
- Florian Schneider and Sunayana Sitaram. 2024. **M5 – A Diverse Benchmark to Assess the Performance of Large Multimodal Models Across Multilingual and Multicultural Vision-Language Tasks**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4309–4345, Miami, FL, USA.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. **LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models**. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, pages 25278–25294, New Orleans, LA, USA.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. **A-OKVQA: A Benchmark for Visual Question Answering Using World Knowledge**. In *Computer Vision - ECCV 2022 - 17th European Conference*, volume 13668, pages 146–162, Tel Aviv, Israel.
- Uri Shaham, Jonathan Herzig, Roee Aharoni, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2024. **Multilingual Instruction Tuning With Just a Pinch of Multilinguality**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2304–2317, Bangkok, Thailand.
- Baifeng Shi, Ziyang Wu, Maolin Mao, Xin Wang, and Trevor Darrell. 2024. **When Do We Not Need Larger Vision Models?** In *Computer Vision - ECCV 2024 - 18th European Conference*, volume 15066, pages 444–462, Milan, Italy.
- Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. 2024. **Design2Code: How Far**

Are We From Automating Front-End Engineering?
arXiv preprint arXiv:2403.03163.

Shivalika Singh, Freddie Vargus, Daniel D’souza, Börje Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura O’Mahony, Mike Zhang, Ramith Het-tiarachchi, Joseph Wilson, Marina Machado, Luisa Moura, Dominik Krzeminski, Hakimeh Fadaei, Irem Ergun, Ifeoma Okoh, Aisha Alaagib, Oshan Mudannayake, Zaid Alyafeai, Vu Chien, Sebastian Ruder, Surya Guthikonda, Emad Alghamdi, Sebastian Gehrmann, Niklas Muennighoff, Max Bartolo, Julia Kreutzer, Ahmet Üstün, Marzieh Fadaee, and Sara Hooker. 2024. **Aya Dataset: An Open-Access Collection for Multilingual Instruction Tuning**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11521–11567, Bangkok, Thailand.

Gurkan Soykan and Gözde Güll Sahin. 2024. Linguistically-Informed Multilingual Instruction Tuning: Is There an Optimal Set of Languages to Tune? *arXiv preprint arXiv:2410.07809*.

Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. **WIT: Wikipedia-based Image Text Dataset for Multimodal Multilingual Machine Learning**. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’21*, pages 2443–2449, New York, NY, USA.

Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. **A Corpus for Reasoning about Natural Language Grounded in Photographs**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428, Florence, Italy.

Hai-Long Sun, Da-Wei Zhou, Yang Li, Shiyin Lu, Chao Yi, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, De-Chuan Zhan, et al. 2024. Parrot: Multilingual Visual Instruction Tuning. *arXiv preprint arXiv:2406.02539*.

Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. 2021. **VisualMRC: Machine Reading Comprehension on Document Images**. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, pages 13878–13888, Virtual.

Benny J. Tang, Angie W. Boggust, and Arvind Satyanarayan. 2023. **VisText: A Benchmark for Semantically Rich Chart Captioning**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023*, pages 7268–7298, Toronto, Canada.

Jingqun Tang, Qi Liu, Yongjie Ye, Jinghui Lu, Shu Wei, Chunhui Lin, Wanqing Li, Mohamad Fitri Faiz Bin Mahmood, Hao Feng, Zhen Zhao, Yanjie Wang, Yuliang Liu, Hao Liu, Xiang Bai, and Can Huang.

2024. **MTVQA: Benchmarking Multilingual Text-Centric Visual Question Answering**. *arXiv preprint arXiv:2405.11985*.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. **Gemma 2: Improving Open Language Models at a Practical Size**. *arXiv preprint arXiv:2408.00118*.

Ashish V. Thapliyal, Jordi Pont Tuset, Xi Chen, and Radu Soricut. 2022. Crossmodal-3600: A Massively Multilingual Multimodal Evaluation Dataset. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 715–729, Abu Dhabi, United Arab Emirates.

Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. 2024. Cambrian-1: A Fully Open, Vision-Centric Exploration of Multimodal LLMs. In *Advances in Neural Information Processing Systems*, volume 37, pages 87310–87356, Vancouver, Canada.

Haoqin Tu, Chenhang Cui, Zijun Wang, Yiyang Zhou, Bingchen Zhao, Junlin Han, Wangchunshu Zhou, Huaxiu Yao, and Cihang Xie. 2024. **How Many Are in This Image A Safety Evaluation Benchmark for Vision LLMs**. In *Computer Vision – ECCV 2024: 18th European Conference*, page 37–55, Milan, Italy.

Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. **CIDEr: Consensus-based image description evaluation**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015*, pages 4566–4575, Boston, MA, USA.

Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. 2023. **To See is to Believe: Prompting GPT-4V for Better Visual Instruction Tuning**. *arXiv preprint arXiv:2311.07574*.

Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. 2024a. **Measuring Multimodal Mathematical Reasoning with MATH-Vision Dataset**. In *Advances in Neural Information Processing Systems*, volume 37, pages 95095–95169, Vancouver, Canada.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhi-hao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024b. **Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution**. *arXiv preprint arXiv:2409.12191*.

Alexander Arno Weber, Klaudia Thellmann, Jan Ebert, Nicolas Flores-Herr, Jens Lehmann, Michael Fromm, and Mehdi Ali. 2024. **Investigating Multilingual Instruction-Tuning: Do Polyglot Models Demand for Multilingual Instructions?** In *Proceedings of the*

2024 Conference on Empirical Methods in Natural Language Processing, pages 20829–20855, Miami, Florida, USA.

Ning Xie, Farley Lai, Derek Doran, and Asim Kadav. 2019. Visual Entailment: A Novel Task for Fine-grained Image Understanding. *arXiv preprint arXiv:1901.06706*.

Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. 2024. **Magpie: Alignment Data Synthesis from Scratch by Prompting Aligned LLMs with Nothing**. *CoRR*, abs/2406.08464.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. **Qwen2 Technical Report**. *CoRR*, abs/2407.10671.

Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianshi Cai, Haoyu Li, Weilin Zhao, Zhihui He, Qianyu Chen, Huarong Zhou, Zhensheng Zou, Haoye Zhang, Shengding Hu, Zhi Zheng, Jie Zhou, Jie Cai, Xu Han, Guoyang Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. **MiniCPM-V: A GPT-4V Level MLLM on Your Phone**. *CoRR*, abs/2408.01800.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Bota Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2024. **MMMU: A Massive Multi-discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI**. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, volume abs/2311.16502, pages 9556–9567, Seattle, WA, USA.

Xiang Yue, Yueqi Song, Akari Asai, Seungone Kim, Jean de Dieu Nyandwi, Simran Khanuja, Anjali Kantharuban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, and Graham Neubig. 2025. **Pangea: A Fully Open Multilingual Multimodal LLM for 39 Languages**. In *The Thirteenth International Conference on Learning Representations*, Singapore, Republic of Singapore.

Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. **Sigmoid Loss for Language Image Pre-Training**. *CoRR*, abs/2303.15343.

Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. 2019. **RAVEN: A Dataset for Relational and Analogical Visual REasoNing**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019*, pages 5317–5327, Long Beach, CA, USA.

Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. 2024. **MM-LLMs: Recent Advances in MultiModal Large Language Models**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12401–12430, Bangkok, Thailand.

Wenxuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2023a. **M3Exam: A Multilingual, Multimodal, Multilevel Benchmark for Examining Large Language Models**. In *Advances in Neural Information Processing Systems*, volume 36, pages 5484–5505, New Orleans, LA, USA.

Yanzhe Zhang, Ruiyi Zhang, Jiuxiang Gu, Yufan Zhou, Nedim Lipka, Diyi Yang, and Tong Sun. 2023b. **LLaVAR: Enhanced Visual Instruction Tuning for Text-Rich Image Understanding**. *CoRR*, abs/2306.17107.

Yilun Zhao, Chen Zhao, Linyong Nan, Zhenting Qi, Wenlin Zhang, Xiangru Tang, Boyu Mi, and Dragomir Radev. 2023. **RobuT: A Systematic Study of Table QA Robustness Against Human-Annotated Adversarial Perturbations**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, pages 6064–6081, Toronto, Canada.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. **TAT-QA: A Question Answering Benchmark on a Hybrid of Tabular and Textual Content in Finance**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers)*, pages 3277–3287, Virtual Event.

Yuke Zhu, Oliver Groth, Michael S. Bernstein, and Li Fei-Fei. 2016. **Visual7W: Grounded Question Answering in Images**. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, pages 4995–5004, Las Vegas, NV, USA.

Appendix Overview

Due to the number of experiments, the general density of our work, and our aim to be as transparent as possible in the sense of open science, the following appendix is extensive. Hence, we provide a brief outline of its content to ease navigation and to get an overview quickly.

A Training Setup

Hyperparameters and Languages

B Training Data

Training Datasets

C Evaluation Setup

Hyperparameters, Metrics, Datasets

D Additional Experiments

Ablation Studies on Architecture and Data

E Qualitative Examples

Showcasing Centurio

B Full Results

All Models, Datasets, and Languages

A Training Setup

All models are trained with the following hyperparameters: AdamW optimizer (Loshchilov and Hutter, 2019) with cosine learning rate schedule and 3% linear warmup; LORA (Hu et al., 2022) is used with rank 256 and α 512 and applied to all matrices in the LLM – the LLM is otherwise frozen; the image encoder is frozen in the first three experiments and jointly trained with the model otherwise; weight decay is 0; batch size is 32 using gradient accumulation; learning rate is $1e - 6$ for the image encoder, $1e - 4$ for LORA and the MLP in general except when training Centurio, we use $5e - 5$ during pretraining and $3e - 5$ during instruct tuning. Models are always trained for one epoch on the entire data. The training loss is causal language modeling and we mask both image and input prompt tokens for calculating the loss.

For going from the pretraining to the instruct tuning phase, we found it best to continue training the same LORA adapter; merging the LORA weights after pretraining and initializing a new adapter gave worse results.

Hyperparameter (LORA rank & α , learning rates, weight decay) were tuned for Phi 3.5 and transferred to the other LLMs.

All models were trained with 4 H100 GPUs. Training Centurio took \approx 6 days (half for pre-training, half for instruct tuning).

Training of one Phi 3.5 model for §2 takes 8–10h for instruct tuning, and for pre-training 12 to

20h (with synthetic OCR data and unfrozen image encoder).

A.1 Training Languages

We list the 100 languages used in training in Table 7.

B Training Data

B.1 For Analysis Experiments

The collections of datasets used in the instruct tuning phase for the analysis experiments (§2) is adapted from LLaVA-Next (Liu et al., 2024b). As multiple evaluation datasets contain multiple images in the input (MaRVL, VGR, VLLOD, M3Exam, xMMMU), we include additional datasets to improve capabilities for this situation. See Table 8 for the full list.

B.2 For Training Centurio

For training Centurio, we combine the datasets from Table 8 with additional datasets listed in Table 9.

B.3 Synthetic OCR Data

We use the official Synthdog code¹³ to generate the samples using the Google Noto font and with images from the ImageNet train split as background. Text is sampled from Wikipedias of the respective languages.

We consider the following 32 languages as not using the Latin script: *am, ar, as, azb, be, bg, bn, bo, el, fa, he, hi, ja, ka, kk, km, ko, lo, mr, my, pa, ru, sa, sd, sr, ta, te, th, ti, uk, ur, zh*.

C Evaluation Setup

This section describes the details of our evaluation setup.

C.1 Generation Parameters

In all experiments of our test suite, we use greedy decoding (`temperature=0.0; do_sample=False;`).

C.2 Metrics

Depending on the dataset and task, we employ either CIDEr (Vedantam et al., 2015), the exact match accuracy, or a relaxed match accuracy (see Table C.4).

¹³<https://github.com/clovaai/donut/tree/master/synthdog>

Name	Script	ISO-639	Flores-200	Tier	Name	Script	ISO-639	Flores-200	Tier
Arabic	Arabic	ar	arb_Arab	5	Urdu	Arabic	ur	urd_Arab	3
Chinese	Trad. Han	zh	zho_Hant	5	Uzbek	Latin	uz	uzn_Latn	3
English	Latin	en	eng_Latn	5	Hebrew	Hebrew	iwhe	heb_Hebr	3
French	Latin	fr	fra_Latn	5	Amharic	Ethiopic	am	amb_Ethi	2
German	Latin	de	deu_Latn	5	Haitian	Latin	ht	hat_Latn	2
Japanese	Japanese	ja	jpn_Jpan	5	Hausa	Latin	ha	hau_Latn	2
Spanish	Latin	es	spa_Latn	5	Icelandic	Latin	is	isl_Latn	2
Basque	Latin	eu	eus_Latn	4	Irish	Latin	ga	gle_Latn	2
Catalan	Latin	ca	cat_Latn	4	Lao	Latin	lo	lao_Laoo	2
Croatian	Latin	hr	hrv_Latn	4	Maltese	Latin	mt	mlt_Latn	2
Czech	Latin	cs	ces_Latn	4	Marathi	Devanagari	mr	mar_Deva	2
Dutch	Latin	nl	nld_Latn	4	Punjabi	Gurmukhi	pa	pan_Guru	2
Finnish	Latin	fi	fin_Latn	4	Sanskrit	Devanagari	sa	san_Deva	2
Hindi	Devanagari	hi	hin_Deva	4	Swahili	Latin	sw	swh_Latn	2
Hungarian	Latin	hu	hun_Latn	4	Tigrinya	Ethiopic	ti	tir_Ethi	2
Italian	Latin	it	ita_Latn	4	Tswana	Latin	tn	tsn_Latn	2
Korean	Hangul	ko	kor_Hang	4	Wolof	Latin	wo	wol_Latn	2
Persian	Arabic	fa	pes_Arab	4	Xhosa	Latin	xh	xho_Latn	2
Polish	Latin	pl	pol_Latn	4	Yoruba	Latin	yo	yor_Latn	2
Portuguese	Latin	pt	por_Latn	4	Zulu	Latin	zu	zul_Latn	2
Russian	Cyrillic	ru	rus_Cyrl	4	Albanian	Latin	sq	als_Latn	1
Serbian	Cyrillic	sr	srp_Cyrl	4	Assamese	Bengali	as	asm_Beng	1
Swedish	Latin	sv	swe_Latn	4	Azerbaijani	Arabic	azb	azb_Arab	1
Turkish	Latin	tr	tur_Latn	4	Bambara	Latin	bm	bam_Latn	1
Vietnamese	Latin	vi	vie_Latn	4	Burmese	Myanmar	my	mya_Mymr	1
Afrikaans	Latin	af	afr_Latn	3	Esperanto	Latin	eo	epo_Latn	1
Bangla	Bengali	bn	ben_Beng	3	Igbo	Latin	ig	ibo_Latn	1
Belarusian	Cyrillic	be	bel_Cyrl	3	Javanese	Latin	jv	jav_Latn	1
Bosnian	Latin	bs	bos_Latn	3	Khmer	Khmer	km	khm_Khmr	1
Bulgarian	Cyrillic	bg	bul_Cyrl	3	Kikuyu	Latin	ki	kik_Latn	1
Cebuano	Latin	ceb	ceb_Latn	3	Lingala	Latin	ln	lin_Latn	1
Danish	Latin	da	dan_Latn	3	Luxembourgish	Latin	lb	ltz_Latn	1
Egyptian Arabic	Arabic	ar-eg	arz_Arab	3	Maori	Latin	mi	mri_Latn	1
Estonian	Latin	et	est_Latn	3	Norwegian	Latin	no	nob_Latn	1
Galician	Latin	gl	glg_Latn	3	Occitan	Latin	oc	oci_Latn	1
Georgian	Georgian	ka	kat_Geor	3	Quechua	Latin	qu	quy_Latn	1
Greek	Greek	el	ell_Grek	3	Samoan	Latin	sm	smo_Latn	1
Indonesian	Latin	id	ind_Latn	3	Sango	Latin	sg	sag_Latn	1
Kazakh	Cyrillic	kk	kaz_Cyrl	3	Sardinian	Latin	sc	srd_Latn	1
Latin	Latin	la	NO	3	Scottish Gaelic	Latin	gd	gla_Latn	1
Latvian	Latin	lv	lvs_Latn	3	Sindhi	Arabic	sd	snd_Arab	1
Lithuanian	Latin	lt	lit_Latn	3	Somali	Latin	so	som_Latn	1
Malay	Latin	ms	zsm_Latn	3	Swati	Latin	ss	ssw_Latn	1
Romanian	Latin	ro	ron_Latn	3	Telugu	Telugu	te	tel_Telu	1
Slovak	Latin	sk	slk_Latn	3	Tibetan	Tibetan	bo	bod_Tibt	1
Slovenian	Latin	sl	slv_Latn	3	Tok Pisin	Latin	tpi	tpi_Latn	1
Tagalog	Latin	tl	tgl_Latn	3	Tsonga	Latin	ts	tso_Latn	1
Tamil	Tamil	ta	tam_Taml	3	Twi	Latin	tw	twi_Latn	1
Thai	Thai	th	tha_Thai	3	Waray	Latin	war	war_Latn	1
Ukrainian	Cyrillic	uk	ukr_Cyrl	3	Welsh	Latin	cy	cym_Latn	1

Table 7: The list of 100 languages used in our training experiments. The "Tier" column represents the tier in the taxonomy proposed by Joshi et al. (2020), where a higher tier indicates more available resources, i.e., data, in the respective language.

Dataset	Size (Images)	Translated?
Natural Image:		
LLaVA Instruct (Liu et al., 2023)	160k	yes
VQAv2 (Goyal et al., 2017)	83k	yes
GQA (Hudson and Manning, 2019)	72k	yes
OKVQA (Marino et al., 2019)	9k	yes
A-OKVQA (Schwenk et al., 2022)	30k	yes
RefCOCO (Kazemzadeh et al., 2014; Mao et al., 2016)	48k	yes
VG (Krishna et al., 2017)	86k	yes
MSCOCO (Lin et al., 2014)	50k (subset)	yes
Multiple Images:		
NLVR (Suh et al., 2019)	86k	yes
Spot-the-difference (Jhamtani and Berg-Kirkpatrick, 2018)	8k	yes
OCR:		
OCRVQA (Mishra et al., 2019)	50k (subset)	no
DocVQA (Mathew et al., 2021)	10k	no
AI2D (Kembhavi et al., 2016)	3k	no
ChartQA (Masry et al., 2022)	18k	no
DVQA (Kafe et al., 2018)	50k (subset)	no
ScienceQA (Lu et al., 2022)	6k	no
Total	766k	

Table 8: List of datasets included in the *instruct tuning phase* in our analysis experiments. All sizes are based on unique images; examples about the same image are packed into one sequence.

For the relaxed match accuracy, we consider an answer correct if it starts with the correct choice letter. For example, answers like "A." are also counted as correct if the gold label is "A".

C.3 Prompts

We list the prompts for each dataset in our test suite used for all models in Figure 2.

C.4 Datasets

In the following, datasets included in our test suite are briefly introduced. An overview is provided in Table 10. Details about the languages covered by the datasets are listed in Table 11.

xGQA The xGQA dataset (Pfeiffer et al., 2022a) is a cross-lingual visual question-answering dataset. It extends the well-known English-only GQA dataset (Hudson and Manning, 2019) by manually translating the questions in the balanced *test-dev* set. Each of the 9666 questions is available in eight languages covering five scripts, while the answers are in English only. The dataset holds 300 unique images from Visual Genome (Krishna et al., 2017).

MaXM The MaXM dataset was introduced by Changpinyo et al. (2023) and is a VQA dataset comprising seven languages in five scripts. In MaXM, the questions and their respective answers are in the same language. The images are a subset of the XM3600 (Thapliyal et al., 2022) dataset and are chosen to match a region where the language of the question-answer pair is spoken. This ensures cultural diversity in the images in addition to the language diversity in the question-answer texts.

Dataset	Size (Images)	Translated?
Natural Image:		
ALLaVA Instruct ¹ (Chen et al., 2024a)	760k	yes
LVIS Instruct4V (Wang et al., 2023)	223k	yes
Visual7W (Zhu et al., 2016)	14k	no
VizWiz QA (Gurari et al., 2018)	21k	no
TallyQA (Acharya et al., 2019)	133k	yes
SketchyVQA (Tu et al., 2024)	4k	yes
OODVQA (Tu et al., 2024)	3k	no
OCR:		
ScienceQA (Cambrian version)	6k	no
AI2D (Cambrian version)	4k	no
Rendered Text ²	10k	no
ScreenQA (Hsiao et al., 2025)	33k	no
LLaVAR (Zhang et al., 2023b)	20k	no
ArxivQA (Li et al., 2024)	54k	no
Chart2Text (Obeid and Hoque, 2020)	25k	no
InfographicVQA (Mathew et al., 2022)	2k	no
VisText (Tang et al., 2023)	10k	no
TQA (Kembhavi et al., 2017)	1k	no
STVQA (Biten et al., 2019)	17k	no
TAT-QA (Zhu et al., 2021)	2k	no
TabMWP (Lu et al., 2023)	23k	no
HiTab (Cheng et al., 2022)	2k	no
IconQA (Lu et al., 2021b)	27k	no
VisualMRC (Tanaka et al., 2021)	3k	no
RobuT (Zhao et al., 2023)	113k	no
FinQA (Chen et al., 2021)	5k	no
Math & Code:		
WebSight (Laurençon et al., 2024)	10k	yes
Design2Code (Si et al., 2024)	0k	yes
DaTikz (Belouadi et al., 2024)	48k	no
CLEVR (Johnson et al., 2017)	70k	yes
CLEVR-Math (Lindström and Abraham, 2022)	70k	yes
Geo170k (Gao et al., 2023)	9k	no
GeomVers (Kazemi et al., 2023)	9k	no
Inter-GPS (Lu et al., 2021a)	1k	no
MathVision (Wang et al., 2024a)	3k	no
Raven (Zhang et al., 2019)	42k	no
Text (no images):		
Aya Dataset (Singh et al., 2024)	202k	–
Tagengo-GPT4 (Devine, 2024)	70k	–
Magpie ² (Xu et al., 2024)	400k	–
Total	2.47M	

Table 9: Datasets used on top of the datasets from Table 8 for the instruct tuning phase of Centurio.

¹: also contains web-scraped images from LAION (Schuhmann et al., 2022) which contain textual elements. ²: <https://huggingface.co/datasets/wendler/RenderedText>. ²: Combining magpie-ultra-v0.1 (50k), Magpie-Qwen2-Pro-200K-English (200k), Magpie-Llama-3.1-Pro-MT-300K-Filtered (150k subset).

XVNLI The XVNLI dataset (Bugliarello et al., 2022) introduces the task of Cross-lingual Visual Natural Language Inference where a model needs to predict whether a textual hypothesis *entails*, *contradicts*, or is *neutral* concerning a visual premise. XVNLI comprises five languages covering three scripts and 357 unique images from Visual Genome. It is based on a combination of the text-only SNLI (Bowman et al., 2015) dataset and its cross-lingual (Agić and Schluter, 2018) and cross-modal (Xie et al., 2019) equivalents.

MaRVL The MaRVL dataset (Liu et al., 2021) aims to benchmark models on Multicultural Reasoning over Vision and Language. A task sample

SMPQA

{QUESTION}\nAnswer the question using a single word or phrase.

CVQA

{QUESTION}\nThere are several options:\nA. {OPTION A}\nB. {OPTION B}\nC. {OPTION C}\nD. {OPTION D}\nAnswer with the option's letter from the given choices directly.

xMMMU

{QUESTION}\nThere are several options:\nA. {OPTION A}\nB. {OPTION B}\nC. {OPTION C}\nD. {OPTION D}\nAnswer with the option's letter from the given choices directly.

MTVQA

{QUESTION}\nAnswer the question using a single word or phrase.\nAnswer in {LANGUAGE}.

M3Exam

{QUESTION}\nOptions:\nA. {OPTION A}\nB. {OPTION B}\nC. {OPTION C}\nD. {OPTION D}\nAnswer with the option's letter from the given choices directly.

BIN-MC

Which of these choices (in English) is shown in the image?\nChoices:\nA. {CHOICE A}\nB. {CHOICE B}\nC. {CHOICE C}\nD. {CHOICE D}\nAnswer with the letter from the given choices directly.

xGQA

{QUESTION}\nAnswer the question using a single word or phrase.\nAnswer in English.

MaXM

{QUESTION}\nAnswer the question using a single word or phrase.\nAnswer in {LANGUAGE}.

MaRVL

Given the two images , is it correct to say “[HYPOTHESIS]”? Answer yes or no.’

XVNLI

Is it guaranteed true that “[HYPOTHESIS]”? Yes, no, or maybe? Answer in English.

M5-VGR

Given the two images , is it correct to say “[HYPOTHESIS]”? Answer yes or no.’

M5-VLOD

Based on the 5 images ordered from top-left to bottom-right, which image does not match the hypothesis “[HYPOTHESIS]”? Choose one from [A, B, C, D, E] and only output a single letter:

XM3600

Briefly describe the image in {LANGUAGE} in one sentence.

Figure 2: Prompts used for the different datasets of our test suite. For M3Exam and xMMMU, the questions contain images at individual positions, and also the options can consist of images. In total, a sample of M3Exam can contain up to 8 images and 8 options, and a sample of xMMMU can contain up to 4 images and 4 options.

Dataset	Task	Visual Input	Textual Input	Target Output	Metric	#Lang.
MaXM	VQA	Single-Image	Question (TL)	WoP (TL)	E. Acc.	6
xGQA	VQA	Single-Image	Question (TL)	WoP (EN)	E. Acc.	8
XVNLI	VNLI	Single-Image	Hypothesis (TL)	'yes' / 'no' / 'maybe'	E. Acc.	5
M5B-VLOD	VLOD	Multi-Image	Hypothesis (TL)	LoC	R. Acc.	12
M5B-VGR	VGR	Multi-Image	Hypothesis (TL)	'yes' / 'no'	E. Acc.	12
MaRVL	VGR	Multi-Image	Hypothesis (TL)	'yes' / 'no'	E. Acc.	6
MTVQA	TH VQA	Single-Image	Question (TL)	WoP (TL)	E. Acc.	9
SMPQA - Name	TH VQA	Single-Image	Question (TL)	WoP (TL)	E. Acc.	11
SMPQA - Ground	TH VGR	Single-Image	Question (TL)	'yes' / 'no'	E. Acc.	11
M3Exam	TH MC VQA	Single or Multi-Image	Question (TL)	LoC	R. Acc.	7
MMMU	TH MC VQA	Single or Multi-Image	Question (EN)	LoC	R. Acc.	1
xMMMU	TH MC VQA	Single or Multi-Image	Question (TL)	LoC	R. Acc.	7
BabelImageNet-MC	MC VQA	Single-Image	Question (TL)	LoC	R. Acc.	20
CVQA	MC VQA	Single-Image	Question (TL)	LoC	R. Acc.	39
XM3600	Captioning	Single-Image	Prompt (EN)	Caption (TL)	CIDEr	36

Table 10: List of datasets contained in our test suite. In the Task column, "VQA", "VNLI", "VLOD", "VGR", "TH", and "MC" are acronyms for "Visual Question Answering", "Visual Natural Language Inference", "Visio-Linguistic Outlier Detection", "Visually Grounded Reasoning", "Text-Heavy", and "Multiple-Choice", respectively. In the "Textual Input" and "Target Output" columns, the acronyms "WoP", "LoC", "TL", and "EN" stand for "(Single) Word or Phrase", "Letter of the correct Choice", "Target Language", and "English", respectively. Further, "E. Acc." is "Exact Accuracy" and "R. Acc." is "Relaxed Accuracy". CVQA is not used in §2 due to its hidden test set with limited submissions.

comprises two images, a textual statement, and a binary true or false answer grounded in the images. MaRVL comprises five languages covering three scripts and 4914 culturally diverse images that match the respective languages. The images in a sample are chosen to match the culture of the annotator who has written the textual statement in his or her native language.

XM3600 The XM3600 dataset (Thapliyal et al., 2022) is a large multilingual image captioning dataset comprising 36 languages with 261375 captions covering 13 different scripts for 100 unique images per language. The images are selected to match the language’s cultural background, ensuring cultural and linguistic diversity. The captions were not automatically translated but manually created by professional annotators who are native speakers of the respective language.

We only use a subset of 500/3600 images (selected randomly) per language when evaluating XM3600 due to its size.

Babel-ImageNet (multiple-choice) (BIN-MC) Babel-ImageNet (Geigle et al., 2024b) translates the labels of ImageNet (Deng et al., 2009) to nearly 300 languages, which allows us to test if models are capable of recognizing and linking the diverse objects of ImageNet to their correct label in the tested language. Testing all 300 languages would be too expensive, instead we use it to deepen our evaluation in languages appearing in only 1 or 2 other datasets, plus English and select few high-resource

languages. Also, we only use 10 images per class instead of 50, again, to keep computational cost reasonable.

We formulate the task as a multiple-choice problem, following the approach by Geigle et al. (2024a) to mine hard negative options from the total label pool. This avoids problems of unclear or underspecified answers that appear in a traditional open-ended VQA formulation. We mine negatives with the English labels, filtering out all candidates not translated by Babel-ImageNet in the target language, that is, in the end, we select the three most similar negative labels that appear in the Babel-ImageNet labels of a given language.

SMPQA We propose **SMPQA** (Synthetic Multilingual Plot QA) as a novel test dataset for evaluating multilingual OCR capabilities in images – bar plots and pie charts to be specific – in 11 languages, covering different scripts and resource levels. See §C.5 for details.

M5B-VGR The M5B-VGR dataset is a Visually Grounded Reasoning dataset similar to MaRVL and was introduced by (Schneider and Sitaram, 2024). A sample comprises two images, a textual statement, and a binary true or false answer grounded in the images. It comprises 12 languages covering 7 scripts and culturally diverse photos taken in regions where the respective language is spoken. The images are sampled from the Dollar Street (Gaviria Rojas et al., 2022) dataset. For each language, there are 120 samples.

Name	Tier	ISO-639-3	ISO-639-1	Datasets
Afrikaans	3	afr	af	BabelImageNet-MC, M3Exam
Amharic	2	amh	am	BabelImageNet-MC, CVQA, M5B-VGR, M5B-VLOD
Arabic	5	ara	ar	MTVQA, SMPQA, XM3600, xMMMU, XVNLI
Bengali	3	ben	bn	CVQA, M5B-VGR, M5B-VLOD, xGQA, XM3600
Berber (macrolanguage)	0	ber	-	M5B-VGR, M5B-VLOD
Breton	1	bre	br	CVQA
Bulgarian	3	bul	bg	CVQA
Chinese	5	zho	zh	CVQA, M3Exam, MaRVL, MaXM, SMPQA, xGQA, XM3600
Croatian	4	hrv	hr	BabelImageNet-MC, XM3600
Cusco Quechua	1	quz	-	XM3600
Czech	4	ces	cs	BabelImageNet-MC, XM3600
Danish	3	dan	da	XM3600
Dutch	4	nld	nl	BabelImageNet-MC, XM3600
Egyptian Arabic	3	arz	-	CVQA
English	5	eng	en	BabelImageNet-MC, M3Exam, M5B-VGR, M5B-VLOD, MaRVL, MaXM, MME, MMMU, SMPQA, xGQA, XM3600, xMMMU, XVNLI
Filipino	3	fil	-	CVQA, M5B-VGR, M5B-VLOD, XM3600
Finnish	4	fin	fi	BabelImageNet-MC, XM3600
French	5	fra	fr	MaXM, MTVQA, XM3600, xMMMU, XVNLI
German	5	deu	de	M5B-VGR, M5B-VLOD, MTVQA, SMPQA, xGQA, XM3600
Hausa	2	hau	ha	BabelImageNet-MC, M5B-VGR, M5B-VLOD
Hebrew	3	heb	he	XM3600
Hindi	4	hin	hi	M5B-VGR, M5B-VLOD, MaXM, SMPQA, XM3600, xMMMU
Hungarian	4	hun	hu	BabelImageNet-MC, XM3600
Igbo	1	ibo	ig	CVQA
Indonesian	3	ind	id	CVQA, MaRVL, SMPQA, xGQA, XM3600, xMMMU
Irish	2	gle	ga	CVQA
Italian	4	ita	it	M3Exam, MTVQA, SMPQA, XM3600
Japanese	5	jpn	ja	BabelImageNet-MC, CVQA, MTVQA, XM3600, xMMMU
Javanese	1	jav	jv	CVQA
Kanuri	0	kau	kr	CVQA
Kinyarwanda	1	kin	rw	CVQA
Korean	4	kor	ko	CVQA, SMPQA, xGQA, XM3600
Malay (macrolanguage)	3	msa	ms	CVQA
Maori	1	mri	mi	BabelImageNet-MC, XM3600
Mi-gkabau	1	min	-	CVQA
Modern Greek	3	ell	el	BabelImageNet-MC, XM3600
Mongolian	1	mon	mn	CVQA
Norwegian	1	nor	no	BabelImageNet-MC, CVQA, XM3600
Oromo	1	orm	om	CVQA
Persian	4	fas	fa	BabelImageNet-MC, XM3600
Polish	4	pol	pl	BabelImageNet-MC, XM3600
Portuguese	4	por	pt	CVQA, M3Exam, xGQA, XM3600, xMMMU
Romanian	3	ron	ro	BabelImageNet-MC, CVQA, MaXM, XM3600
Russian	4	rus	ru	CVQA, M5B-VGR, M5B-VLOD, MTVQA, SMPQA, xGQA, XM3600, XVNLI
Sinhala	0	sin	si	CVQA
Spanish	5	spa	es	BabelImageNet-MC, CVQA, XM3600, XVNLI
Sundanese	1	sun	su	CVQA
Swahili (macrolanguage)	2	swa	sw	CVQA, M5B-VGR, M5B-VLOD, MaRVL, XM3600
Swedish	4	swe	sv	XM3600
Tamil	3	tam	ta	BabelImageNet-MC, CVQA, MaRVL
Telugu	1	tel	te	BabelImageNet-MC, CVQA, XM3600
Thai	3	tha	th	M3Exam, M5B-VGR, M5B-VLOD, MaXM, MTVQA, SMPQA, XM3600
Turkish	4	tur	tr	MaRVL, XM3600
Ukrainian	3	ukr	uk	XM3600
Urdu	3	urd	ur	CVQA
Vietnamese	4	vie	vi	M3Exam, MTVQA, XM3600
Zulu	2	zul	zu	BabelImageNet-MC, M5B-VGR, M5B-VLOD, SMPQA
Unique Languages	56 (43 without CVQA)			

Table 11: List of languages covered in the datasets of our test suite. The "Tier" column represents the tier in the taxonomy proposed by Joshi et al. (2020), where a higher tier indicates more available resources, i.e., data, in the respective language. CVQA is not used in §2 due to its hidden test set with limited submissions.

M5B-VLOD The M5B-VLOD (Visio-Linguistic Outlier Detection) dataset was introduced by (Schneider and Sitaram, 2024). A sample comprises five images and a textual statement that is true for all but one of the images. The task is to find the outlier image, that is, the image that does not match the statement. It comprises the same 12

languages as M5B-VGR and images sampled with a similar strategy from the same dataset. For each language, there are 120 samples.

MTVQA The MTVQA dataset was introduced by (Tang et al., 2024) and comprises text-heavy Visual Question Answering (VQA) tasks. It features human expert annotations across 9 diverse

languages, consisting of a total of 6778 question-answer pairs across 2116 images. The images primarily contain text in the respective language and the question (and answer) related to that text. The images are sampled from different publicly available datasets.

CVQA The CVQA dataset was introduced by (Romero et al., 2024) and is a multilingual, culturally nuanced VQA benchmark that includes a diverse set of languages, many of them underrepresented and understudied in NLP. It consists of 10000 questions across 30 countries, covering 31 languages, and in 39 distinct country-language pairs (e.g., the dataset includes 7 different splits for Spanish because it contains 7 countries where Spanish is spoken). The images in the dataset were manually gathered by human annotators to match and depict the culture of the respective country-language pair.

A sample consists of one image and a question related to the image in the respective language. The authors did not release the test set publicly but allowed up to 5 daily submissions to their leaderboard to obtain evaluation results.

M3Exam The M3Exam dataset was introduced by (Zhang et al., 2023a). It contains real-world exam questions in 9 languages, which are either text-only or multi-modal. In our test suite, we only consider samples that require at least one image. Further, due to the low number of resulting samples for Swahili and Javanese, we only include the remaining 7 languages. The remaining samples consist of multiple-choice questions in the target language and up to 8 images that can appear both in the question and the answer options. Further, the number of options ranges from 4 to 8 depending on the individual sample.

xMMMU The xMMMU was introduced by (Yue et al., 2025) and consists of college-level multiple-choice VQA samples across seven languages. It was automatically translated using GPT4o from a subset of 300 randomly selected questions from the MMMU (Yue et al., 2024) validation split.

C.5 Details for SMPQA

We propose **SMPQA** (Synthetic Multilingual Plot QA) as a test dataset for evaluating multilingual OCR capabilities, that is capabilities to identify and read text in various languages in images, specifically bar plots and pie charts.

We test the capabilities in two directions: i) *grounding* requires the model to ground a given label in the user prompt to the corresponding part in the plot to answer a yes/no question (“Is the bar with label \$X the biggest?”); ii) *reading* requires the model to output the label of a specified part of the plot (“What is the label of the biggest slice?”). The questions are simple by design, requiring minimal reasoning, math, or multi-hop capabilities, as to isolate solely the OCR capabilities in the tested language. We show example plots and questions in Figure 3.

We use exact match accuracy for both tasks. For reading, edit distance to the correct word would be a fine-grained alternative but since word lengths differ between languages – Chinese can be 1-2 characters while Indonesian can be >10 – we opt against it to more easily compare results between languages. To have a fair comparison between languages, we construct the dataset in a way that plots and questions about them are identical between languages (except for labels in the respective languages, obviously).

Construction: SMPQA is constructed with a deterministic pipeline yielding identical results for each language.

1. We define a list of diverse pie charts and bar plots by randomly sampling the number of bars/slices, the size of each, their colors, the plot size and aspect ratio, and vertical/horizontal orientation for bar plots, and exploding some slices in pie charts. For each plot type, we define 50 configurations, so we have 100 plots/images in total per language.
2. Using word lists of common words in the languages, we sample words for use as labels for the bars and pie slices to fill and ultimately render the pre-defined plots. This means the plots are identical between languages except for the labels and some size adjustments caused by different word lengths.
3. For each plot, we use templates to generate 5 questions for reading and 8 questions for grounding (with balanced ‘yes’ and ‘no’ as answers). The questions are always the same for a plot, so each language has the same questions, just with different labels.

Language Selection: We selected the languages as follows: For Latin-script languages, we chose English and one language from Tier 5 to 2 to have both high- and low-resource languages: German, Italian, Indonesian, and Zulu. For non-Latin scripts,

HuggingFace Model ID	Params
Qwen/Qwen2-VL-2B-Instruct (Wang et al., 2024b)	2B
Qwen/Qwen2-VL-7B-Instruct (Wang et al., 2024b)	7B
microsoft/Phi-3.5-vision-instruct (Abdin et al., 2024a)	4B
neulab/Pangea-7B-hf (Yue et al., 2025)	7B
openbmb/MiniCPM-V2_6 (Yao et al., 2024)	8B
meta-llama/Llama-3.2-11B-Vision-Instruct (AI@Meta, 2024)	11B
mistralai/Pixtral-12B-2409 (Agrawal et al., 2024)	12B
AIDC-AI/Parrot-7B (Sun et al., 2024)	7B
MBZUAI/PALO-7B (Maaz et al., 2025)	7B
MBZUAI/PALO-13B (Maaz et al., 2025)	13B
OpenGVLab/InternVL2_5-4B (Chen et al., 2024e)	4B
OpenGVLab/InternVL2_5-8B (Chen et al., 2024e)	8B
maya-multimodal/maya (Alam et al., 2024)	8B

Table 12: List of models considered in our evaluation experiments.

we select 6 languages to represent scripts with high usage in the world: Russian (Cyrillic), Chinese, Korean (Hangul), Hindi (Devanagari), Arabic, and Thai.

We note that our dataset construction can easily be extended to other languages if needed (as long as word lists are available) to test, for example, more scripts (Telugu, Greek, Hebrew, ...) or languages using the Latin script with heavy use of diacritics (Vietnamese, Turkish, ...). This makes SMPQA an ideal starting point for probing OCR capabilities in diverse languages.

C.6 Baseline Models

We list the evaluated baseline models in Table 12. In all baseline evaluation experiments, we use greedy decoding (temperature=0.0; do_sample=False;). Further, we do not preprocess the images in any way and use the provided code for inference with the respective model.

We use relaxed match accuracy for all tasks even if Centurio uses exact match for a fairer comparison because some models struggled with replying just ‘yes’/‘no’ and similar issues.

D Additional Experiments

D.1 Analysis Results with Llama 3

We report the results with Llama 3 when repeating the experiments of §2.2, 2.3, 2.4 in Table 13, 14, 15.

D.2 Non-Uniform Language Allocation

In our experiments in §2, we distribute the non-English portion of the data uniformly over all languages. We now consider two *stratified* distributions that upsample low-resource languages. A language with taxonomy i will get allocated the

Train Lang.	T1	T2	T3	T4	T5	en
English	16.1	34.7	26.3	24.3	26.2	56.4
T5	19.1	32.5	29.3	27.2	35.5	54.3
L100	31.1	43.0	39.4	35.9	36.4	56.6
Without tasks affected by language fidelity:						
English	36.6	37.1	39.0	39.6	40.0	54.6
T5	38.8	34.8	40.1	40.2	40.4	53.5
L100	46.3	44.0	45.0	42.8	42.9	55.3

Table 13: Experimental setup of Table 1 repeated with Llama 3 and the setups: just English, T5 languages, and L100 languages.

English %	T1	T2	T3	T4	T5	en
10	32.9	43.1	38.7	35.4	35.4	54.2
50	31.1	43.0	39.4	35.9	36.4	56.6
90	26.9	38.7	36.9	34.2	35.8	56.6

Table 14: Experimental setup of Table 2 repeated with Llama 3 and the setups: 10, 50, and 90% English instruct tune data.

following portion of the non-English data:

$$p(i) = \frac{f(i)}{\sum_{j \in \text{TrainLanguages}} f(j)} \quad (1)$$

with $f(i) = \frac{1}{i}$ for **Stratified-1** and $f(i) = \frac{1}{\exp i}$ for **Stratified-2**. This effectively doubles the allocated data for T1 languages, and divides the data for T5 languages by a factor 3 or 20 (depending on Stratified-1 or -2).

Results are reported in Table 16. We do observe a small decrease for T5 and T4 languages, and also for T3 with Stratified-2, but results for T1 and T2 languages stay relatively constant despite more data. This suggests that higher-resource languages can be quite sample efficient even with what amounts to a few hundred samples (at least in the instruct tuning phase) but as the stratified distributions fail to improve lower-resource languages, there is little reason in practice to not use the uniform distribution which makes no assumptions about the resource-level of a language.

D.3 LLM Comparison

We train several recent 7-9B parameter LLMs on the instruct tuning data mix used in our analysis with L100 languages and 50% English. All models are trained with the same hyperparameters. We compare Llama 3 (Llama Team, 2024)¹⁴, Gemma 2

¹⁴While Llama 3.1 has officially better multilingual capabilities, we found Llama 3 to work better.

English %	T1	T2	T3	T4	T5	en
No pretrain	31.1	43.0	39.4	35.9	36.4	56.6
100	33.9	44.7	43.3	39.9	39.9	60.8
1	37.8	47.4	45.0	41.1	40.7	61.4

Table 15: Results of Table 3 repeated with Llama 3 and the setups: 1 and 100% English pre-train data.

Distribution	T1	T2	T3	T4	T5	en
Uniform	<u>18.9</u>	<u>32.6</u>	<u>30.7</u>	28.8	34.4	<u>52.6</u>
Stratified-1	18.6	32.5	30.7	<u>28.0</u>	33.8	53.0
Stratified-2	19.2	32.6	29.5	27.4	<u>33.9</u>	52.0

Table 16: Comparison between our uniform allocation of data compared to two stratified allocations that upsample low-resource languages.

(Team et al., 2024), Aya-Expanse (Aryabumi et al., 2024; Dang et al., 2024), and Qwen 2.5 (Yang et al., 2024).

Table 17 shows that Qwen and Aya yield the overall best results with Aya slightly ahead in T3-T5 and Qwen (with Llama 3) better in T1 and T2. Qwen achieves by far the best English results. Qwen is also notably strong in the exam tasks (M3Exam, xMMMU) and is the only model with better-than-guessing results on average for VLOD.

D.4 Qualitative Analysis of the Effect of Synthetic OCR-Training Data on Non-Latin Scripts

We performed a qualitative analysis of the predictions in the SMPQA Naming subtask to identify potential issues with non-Latin scripts when synthetic OCR training data is used. However, the results were not conclusive. For example, without the OCR training data, some models answered with colors or values (of the chart elements) but only in non-Latin scripts, while they tried to give a proper answer in other languages. We also saw (for models trained with OCR data) that in some scripts (e.g., Thai, Hindi), answers were sometimes nearly correct with only 1 or 2 characters wrong, but e.g., in Chinese or Korean, predicted answers were often completely wrong, i.e., did not just differ by 1 or 2 incorrect radicals.

E Qualitative Examples

We provide some qualitative examples of our Centurio models. Figure 4, 5, and 6 show results with (non-English) text in images with English prompts. Figure 7 and 8 show examples for multilingual prompts (and responses).

LLM	T1	T2	T3	T4	T5	en
Phi-3.5-mini-instruct	18.9	32.6	30.7	28.8	34.4	52.6
gemma-2-9b-it	29.2	40.9	36.4	33.5	35.3	52.8
Meta-Llama-3-8B-Instruct	31.1	<u>43.0</u>	39.4	35.9	36.4	56.6
Qwen2.5-7B-Instruct	<u>30.7</u>	43.7	<u>42.0</u>	<u>38.1</u>	<u>40.5</u>	62.7
aya-expansse-8b	28.3	42.5	43.0	39.8	40.9	<u>59.9</u>

Table 17: Comparison between different LLM backbones all trained with the instruct tuning data with L100 languages and 50% English (as in §2.3).

F Full Results

We report the full results of all tasks with all language-specific results. **Avg.** refers to the average without English. Metric are the same as in §C.

F.1 Analysis Experiments

We report the full results for all models trained for §2 (and also the LLMs tested for §D.3).

The following holds for all Tables: Models of the form ‘**Phi 3.5 - T5 50**’ are to be interpreted as using the LLM Phi 3.5 with the T5 languages and 50% English with analog interpretation for other rows.

‘**Phi 3.5 - PT 1**’ means the model was pretrained with 1% English and then instruct-tuned with the L100 50% English mix (see §3).

‘**Phi 3.5 - OCR 1**’ means the model was pretrained with 50% English for the captions and 1% English for the OCR data and then instruct-tuned with the L100 50% English mix (see §4).

BIN-MC Table 18

M3Exam Table 19

VGR Table 20

VLOD Table 21

MaRVL Table 22

MaXM Table 23

MTVQA Table 24

xGQA Table 25

XM3600 Table 26. Language fidelity for §2.2 in Table 27.

XVNLI Table 28

xMMMU Table 29

SMPQA - Ground Table 30

SMPQA - Name Table 31

	en	avg.	af	am	cs	el	es	fa	fi	ha	hr	hu	ja	mi	nl	no	pl	ro	ta	te	zu
Phi 3.5 - English	64.7	38.1	43.3	29.7	41.5	35.5	55.9	33.6	36.4	24.5	43.3	39.0	49.8	27.8	47.3	44.2	41.4	42.8	30.0	27.1	31.1
Phi 3.5 - T5 50	66.0	39.6	46.0	30.3	43.1	36.3	56.3	33.4	36.5	35.1	45.1	40.7	50.9	30.1	48.4	46.2	41.1	43.1	31.0	29.6	29.1
Phi 3.5 - T5-4 50	65.2	40.6	46.8	29.6	44.6	37.9	59.1	36.7	37.5	29.0	46.4	42.5	52.0	31.1	50.7	47.4	43.0	43.5	31.5	29.0	32.4
Phi 3.5 - T5-3 50	65.5	40.6	50.0	28.8	43.3	37.4	58.6	34.4	38.4	33.0	46.1	41.4	50.9	31.2	49.7	47.0	41.9	43.8	32.5	29.6	32.7
Phi 3.5 - T5-2 50	64.8	39.1	47.2	25.6	41.9	35.8	57.9	34.0	36.0	29.8	44.8	39.5	50.0	30.5	47.7	45.8	41.2	42.4	30.4	29.2	33.9
Phi 3.5 - L100 50	64.7	39.9	48.1	28.2	42.8	36.8	57.2	34.7	37.0	28.2	44.7	40.4	51.2	31.6	47.8	46.4	40.9	43.8	30.6	30.1	37.5
Llama 3 - English	65.4	40.9	44.0	28.2	46.9	42.2	53.0	42.4	38.7	31.1	47.6	46.3	48.6	30.1	48.2	47.4	44.0	44.9	31.6	32.5	29.3
Llama 3 - T5 50	63.9	43.7	50.6	28.7	49.2	46.4	54.6	46.6	41.7	35.4	50.7	50.8	51.9	30.0	51.2	50.9	47.0	48.4	31.1	35.6	30.1
Llama 3 - L100 50	66.2	48.8	55.3	35.1	54.2	51.2	56.2	47.6	46.2	37.2	56.1	54.1	53.3	33.7	54.6	54.3	50.8	51.9	43.6	50.9	40.8
Phi 3.5 - L100 1	63.1	39.7	47.4	26.8	42.9	36.7	56.2	34.3	35.9	33.5	46.8	40.5	49.0	32.7	48.4	46.7	41.3	43.1	29.0	29.9	34.2
Phi 3.5 - L100 10	62.7	39.4	47.1	27.1	43.1	36.8	56.5	34.4	36.5	29.3	43.8	40.9	49.8	29.8	47.2	48.2	41.4	43.6	30.2	28.0	34.4
Phi 3.5 - L100 24	63.3	40.4	48.0	29.0	43.3	37.7	56.5	35.2	36.7	32.4	46.9	40.7	50.7	33.0	49.3	47.2	41.8	44.4	31.9	31.2	31.1
Phi 3.5 - L100 50	64.7	39.9	48.1	28.2	42.8	36.8	57.2	34.7	37.0	28.2	44.7	40.4	51.2	31.6	47.8	46.4	40.9	43.8	30.6	30.1	37.5
Phi 3.5 - L100 75	65.4	39.8	47.1	26.0	42.0	37.1	57.4	34.7	36.9	32.2	44.2	40.4	51.3	31.5	49.4	46.8	41.6	42.9	31.1	28.5	34.2
Phi 3.5 - L100 90	64.7	37.5	43.8	24.1	40.3	35.8	57.1	31.8	35.7	25.0	43.1	39.2	49.1	24.9	47.9	44.4	39.3	42.8	28.7	27.7	31.9
Llama 3 - L100 10	65.9	49.8	58.4	38.1	55.0	50.9	58.5	49.3	45.7	40.7	59.4	56.3	54.1	34.9	53.7	56.8	51.8	51.3	42.6	51.9	36.0
Llama 3 - L100 50	66.2	48.8	55.3	35.1	54.2	51.2	56.2	47.6	46.2	37.2	56.1	54.1	53.3	33.7	54.6	54.3	50.8	51.9	43.6	50.9	40.8
Llama 3 - L100 90	64.4	45.3	52.5	26.8	51.0	47.2	54.8	45.9	44.0	29.5	54.1	50.0	51.2	31.3	52.1	51.8	48.6	49.8	36.5	48.3	34.7
Phi 3.5 - L100 50	64.7	39.9	48.1	28.2	42.8	36.8	57.2	34.7	37.0	28.2	44.7	40.4	51.2	31.6	47.8	46.4	40.9	43.8	30.6	30.1	37.5
Phi 3.5 - PT 100	66.3	38.9	48.4	25.0	42.7	36.0	57.3	33.2	36.5	22.3	44.8	39.8	49.9	31.3	48.6	46.4	41.3	43.2	30.5	30.8	30.6
Phi 3.5 - PT 50	65.7	42.2	50.0	37.8	44.2	40.0	57.8	36.0	36.5	33.0	45.2	41.8	49.3	35.0	49.0	48.1	42.0	44.1	33.7	37.7	40.6
Phi 3.5 - PT 1	65.8	42.8	50.1	35.1	44.8	38.9	56.9	37.9	37.5	41.2	49.1	42.1	49.6	33.4	49.6	48.2	43.6	45.9	34.9	36.1	38.5
Llama 3 - L100 50	66.2	48.8	55.3	35.1	54.2	51.2	56.2	47.6	46.2	37.2	56.1	54.1	53.3	33.7	54.6	54.3	50.8	51.9	43.6	50.9	40.8
Llama 3 - PT 1	69.6	55.5	62.4	44.0	60.4	60.4	62.9	55.3	51.7	40.4	63.0	62.1	59.9	36.6	59.4	62.1	58.0	58.6	50.6	60.6	45.7
Llama 3 - PT 100	68.7	53.6	63.4	36.8	59.6	58.1	62.5	54.1	50.8	37.5	63.1	61.6	60.7	36.9	59.9	61.0	58.0	58.0	46.5	54.0	34.9
Gemma 2 - L100 50	60.5	44.8	49.1	42.5	47.5	45.3	52.0	44.8	41.6	30.9	50.7	47.6	51.4	32.8	49.8	51.1	47.2	47.5	41.8	45.1	32.1
Llama 3 - L100 50	66.2	48.8	55.3	35.1	54.2	51.2	56.2	47.6	46.2	37.2	56.1	54.1	53.3	33.7	54.6	54.3	50.8	51.9	43.6	50.9	40.8
Qwen 2.5 - L100 50	68.2	50.6	62.4	37.1	57.9	50.8	63.4	49.6	42.6	28.7	61.0	48.3	63.1	33.5	58.8	58.2	57.2	55.4	36.8	55.6	40.6
Aya-Expanse - L100 50	67.6	52.0	62.2	31.0	65.3	65.5	63.2	58.9	39.8	33.2	60.8	46.3	65.1	33.1	61.3	55.5	60.2	61.9	43.5	43.2	37.2
Centurio Aya	69.7	54.7	63.6	29.4	66.2	67.8	65.1	60.0	43.3	37.5	63.6	49.8	66.7	37.0	62.4	59.1	62.6	64.0	46.9	50.9	42.6
Centurio Qwen	72.7	56.2	65.3	47.4	62.2	56.7	67.0	53.6	48.8	36.7	65.4	54.1	67.6	39.1	63.7	63.6	60.4	58.5	45.2	63.4	49.5

Table 18: BIN-MC

F.2 Comparison with Centurio

BIN-MC Table 32

M3Exam Table 33

VGR Table 34

VLOD Table 35

MaRVL Table 36

MaXM Table 37

MTVQA Table 38

xGQA Table 39

XM3600 Table 40. Language fidelity Table 41.

XVNLI Table 42

xMMMU Table 43

CVQA Table 44

SMPQA - Ground Table 45

SMPQA - Name Table 46

	en	avg.	af	zh	it	pt	th	vi
Phi 3.5 - English	52.9	32.7	32.5	37.0	49.6	39.7	25.4	12.2
Phi 3.5 - T5 50	51.2	35.3	39.9	35.9	46.4	39.7	28.2	21.7
Phi 3.5 - T5-4 50	52.2	34.2	40.5	32.4	49.1	38.6	25.2	19.1
Phi 3.5 - T5-3 50	51.3	35.3	43.6	34.0	47.4	37.3	27.9	21.7
Phi 3.5 - T5-2 50	49.2	33.7	39.3	32.9	45.1	38.4	22.2	24.3
Phi 3.5 - L100 50	50.8	36.0	39.3	36.1	50.9	40.1	26.2	23.5
Llama 3 - English	46.1	32.5	38.6	32.6	41.6	35.0	25.9	20.9
Llama 3 - T5 50	45.0	33.8	40.5	34.3	41.9	34.1	25.7	26.1
Llama 3 - L100 50	46.6	34.2	44.2	31.0	42.4	34.6	27.2	26.1
Phi 3.5 - L100 1	50.3	35.1	39.9	35.4	46.6	39.2	23.9	25.2
Phi 3.5 - L100 10	48.8	33.9	35.0	33.6	48.1	36.1	24.7	26.1
Phi 3.5 - L100 24	50.8	36.5	41.7	37.0	51.6	35.9	27.7	25.2
Phi 3.5 - L100 50	50.8	36.0	39.3	36.1	50.9	40.1	26.2	23.5
Phi 3.5 - L100 75	48.0	36.1	44.2	35.9	47.1	38.4	26.7	24.3
Phi 3.5 - L100 90	51.7	35.1	36.8	38.0	48.1	36.8	26.4	24.3
Llama 3 - L100 10	43.7	33.6	41.7	29.4	44.9	35.3	23.7	27.0
Llama 3 - L100 50	46.6	34.2	44.2	31.0	42.4	34.6	27.2	26.1
Llama 3 - L100 90	43.3	34.6	37.4	32.2	44.9	35.3	30.2	27.8
Phi 3.5 - L100 50	50.8	36.0	39.3	36.1	50.9	40.1	26.2	23.5
Phi 3.5 - PT 100	50.3	35.8	41.7	37.5	49.4	36.6	24.2	25.2
Phi 3.5 - PT 50	49.7	33.1	41.1	36.1	44.4	35.0	21.7	20.0
Phi 3.5 - PT 1	48.4	33.8	41.7	35.9	46.4	34.8	23.2	20.9
Llama 3 - L100 50	46.6	34.2	44.2	31.0	42.4	34.6	27.2	26.1
Llama 3 - PT 1	50.2	37.9	44.8	34.7	48.1	40.6	31.4	27.8
Llama 3 - PT 100	52.9	37.1	50.3	33.8	46.6	37.5	30.2	24.3
Gemma 2 - L100 50	42.5	33.4	43.6	33.6	41.6	30.4	27.7	23.5
Llama 3 - L100 50	46.6	34.2	44.2	31.0	42.4	34.6	27.2	26.1
Qwen 2.5 - L100 50	53.6	39.6	46.0	44.7	50.6	42.4	29.7	24.3
Aya-Expanse - L100 50	49.3	36.5	46.6	36.8	51.9	39.0	26.2	18.3
Centurio Aya	53.0	41.2	52.8	40.3	51.4	47.7	27.4	27.8
Centurio Qwen	61.2	46.9	50.9	64.1	55.6	49.0	31.9	29.6

Table 19: M3Exam

	en	avg.	am	ber	bn	de	fil	ha	hi	ru	sw	th	zu
Phi 3.5 - English	80.8	54.1	45.0	50.8	41.5	71.7	55.8	41.7	62.7	85.0	35.8	68.3	36.2
Phi 3.5 - T5 50	75.8	50.9	49.2	49.2	40.7	72.5	55.0	42.5	54.2	60.8	37.5	60.8	37.9
Phi 3.5 - T5-4 50	83.3	55.1	51.7	43.3	49.2	70.8	65.8	42.5	61.9	70.8	38.3	75.0	36.2
Phi 3.5 - T5-3 50	83.3	56.6	43.3	50.8	50.8	74.2	69.2	42.5	57.6	76.7	43.3	71.7	42.2
Phi 3.5 - T5-2 50	81.7	57.5	45.8	52.5	44.1	73.3	64.2	39.2	59.3	73.3	60.0	60.8	59.5
Phi 3.5 - L100 50	76.7	56.4	46.7	46.7	54.2	71.7	60.0	45.0	57.6	70.8	57.5	65.8	44.0
Llama 3 - English	82.5	56.3	66.7	30.8	49.2	77.5	50.8	48.3	63.6	75.8	46.7	70.0	39.7
Llama 3 - T5 50	77.5	55.9	47.5	49.2	49.2	71.7	63.3	42.5	62.7	73.3	45.8	70.8	38.8
Llama 3 - L100 50	80.0	64.8	58.3	47.5	64.4	75.8	61.7	67.5	64.4	73.3	59.2	67.5	73.3
Phi 3.5 - L100 1	65.0	47.5	42.5	50.0	38.1	65.0	58.3	40.0	45.8	58.3	39.2	42.5	42.2
Phi 3.5 - L100 10	73.3	54.5	43.3	50.0	51.7	67.5	60.0	45.0	51.7	63.3	53.3	63.3	50.0
Phi 3.5 - L100 24	73.3	60.3	54.2	47.5	58.5	72.5	55.0	58.3	60.2	72.5	64.2	59.2	61.2
Phi 3.5 - L100 50	76.7	56.4	46.7	46.7	54.2	71.7	60.0	45.0	57.6	70.8	57.5	65.8	44.0
Phi 3.5 - L100 75	80.0	56.7	51.7	53.3	55.1	70.8	67.5	41.7	63.6	75.8	38.3	69.2	36.2
Phi 3.5 - L100 90	79.2	54.6	43.3	50.0	44.9	80.8	60.0	42.5	55.9	77.5	45.0	55.8	44.8
Llama 3 - L100 10	77.5	65.4	65.0	45.0	63.6	76.7	58.3	70.8	64.4	74.2	63.3	69.2	69.0
Llama 3 - L100 50	80.0	64.8	58.3	47.5	64.4	75.8	61.7	67.5	64.4	73.3	59.2	67.5	73.3
Llama 3 - L100 90	82.5	63.0	45.8	39.2	66.1	80.8	58.3	68.3	61.9	75.0	63.3	75.0	59.5
Phi 3.5 - L100 50	76.7	56.4	46.7	46.7	54.2	71.7	60.0	45.0	57.6	70.8	57.5	65.8	44.0
Phi 3.5 - PT 100	80.8	58.6	44.2	49.2	56.8	78.3	56.7	47.5	65.3	75.0	47.5	73.3	50.9
Phi 3.5 - PT 50	80.0	63.2	58.3	50.0	55.1	78.3	63.3	60.0	61.9	76.7	55.0	75.0	61.2
Phi 3.5 - PT 1	80.0	62.0	55.8	50.0	51.7	81.7	62.5	60.0	66.1	75.0	50.0	66.7	62.1
Llama 3 - L100 50	80.0	64.8	58.3	47.5	64.4	75.8	61.7	67.5	64.4	73.3	59.2	67.5	73.3
Llama 3 - PT 1	87.5	71.2	70.0	50.8	65.3	79.2	63.3	83.3	68.6	82.5	66.7	85.8	68.1
Llama 3 - PT 100	85.0	68.8	65.8	49.2	67.8	80.8	61.7	70.0	66.9	85.0	70.0	74.2	65.5
Gemma 2 - L100 50	77.5	61.8	64.2	52.5	48.3	70.8	51.7	64.2	58.5	71.7	54.2	70.8	73.3
Llama 3 - L100 50	80.0	64.8	58.3	47.5	64.4	75.8	61.7	67.5	64.4	73.3	59.2	67.5	73.3
Qwen 2.5 - L100 50	91.7	71.2	76.7	50.0	69.5	81.7	77.5	57.5	72.9	83.3	71.7	80.8	62.1
Aya-Expanse - L100 50	92.5	69.9	52.5	54.2	55.9	80.8	85.0	72.5	79.7	83.3	63.3	78.3	63.8
Centurio Aya	82.5	66.8	71.7	54.2	59.3	73.3	59.2	65.0	71.2	75.8	67.5	72.5	65.5
Centurio Qwen	87.5	73.1	77.5	49.2	62.7	80.8	78.3	76.7	72.9	85.0	70.0	81.7	69.0

Table 20: VGR

	en	avg.	am	ber	bn	de	fil	ha	hi	ru	sw	th	zu
Phi 3.5 - English	16.7	21.3	20.8	20.8	19.2	16.7	25.8	28.3	17.0	12.5	25.0	26.7	22.0
Phi 3.5 - T5 50	23.3	20.0	15.0	18.3	20.8	21.7	16.7	20.0	23.2	27.5	22.3	15.8	18.6
Phi 3.5 - T5-4 50	17.5	18.2	19.2	20.8	13.3	20.8	17.5	16.7	21.4	26.7	16.1	10.0	17.8
Phi 3.5 - T5-3 50	25.8	19.8	16.7	17.5	21.7	21.7	20.0	21.7	23.2	20.8	18.8	16.7	18.6
Phi 3.5 - T5-2 50	21.7	20.5	21.7	18.3	16.7	22.5	27.5	27.5	17.9	21.7	17.0	13.3	21.2
Phi 3.5 - L100 50	18.3	19.5	16.7	20.8	19.2	25.8	20.0	16.7	25.0	20.8	13.4	17.5	18.6
Llama 3 - English	12.5	20.8	18.3	21.7	20.0	10.8	24.2	29.2	15.2	12.5	28.6	29.2	19.5
Llama 3 - T5 50	20.8	20.1	18.3	19.2	17.5	16.7	25.0	21.7	24.1	15.0	19.6	23.3	20.3
Llama 3 - L100 50	12.5	20.6	19.2	20.8	20.0	10.8	24.2	30.0	15.2	10.8	28.6	27.5	19.5
Phi 3.5 - L100 1	24.2	19.3	15.0	21.7	17.5	20.0	29.2	22.5	17.9	14.2	16.1	22.5	16.1
Phi 3.5 - L100 10	23.3	19.2	23.3	15.0	16.7	21.7	20.8	20.8	20.5	24.2	10.7	15.8	22.0
Phi 3.5 - L100 24	25.0	18.3	20.8	18.3	16.7	20.8	16.7	20.8	17.9	21.7	14.3	16.7	16.9
Phi 3.5 - L100 50	18.3	19.5	16.7	20.8	19.2	25.8	20.0	16.7	25.0	20.8	13.4	17.5	18.6
Phi 3.5 - L100 75	16.7	18.0	15.0	20.0	19.2	19.2	16.7	23.3	17.0	13.3	17.9	15.8	20.3
Phi 3.5 - L100 90	22.5	19.0	20.0	16.7	15.8	20.0	16.7	23.3	21.4	23.3	16.1	15.8	19.5
Llama 3 - L100 10	13.3	20.4	18.3	21.7	19.2	10.8	23.3	26.7	17.9	10.0	28.6	28.3	19.5
Llama 3 - L100 50	12.5	20.6	19.2	20.8	20.0	10.8	24.2	30.0	15.2	10.8	28.6	27.5	19.5
Llama 3 - L100 90	12.5	19.9	18.3	21.7	15.0	10.8	22.5	28.3	15.2	10.8	28.6	28.3	19.5
Phi 3.5 - L100 50	18.3	19.5	16.7	20.8	19.2	25.8	20.0	16.7	25.0	20.8	13.4	17.5	18.6
Phi 3.5 - PT 100	23.3	20.0	16.7	16.7	24.2	20.0	25.0	21.7	19.6	15.0	20.5	20.0	20.3
Phi 3.5 - PT 50	20.0	18.6	18.3	17.5	15.0	15.8	14.2	21.7	17.9	23.3	20.5	20.8	19.5
Phi 3.5 - PT 1	25.0	19.4	21.7	22.5	19.2	22.5	16.7	15.8	20.5	21.7	16.1	15.0	22.0
Llama 3 - L100 50	12.5	20.6	19.2	20.8	20.0	10.8	24.2	30.0	15.2	10.8	28.6	27.5	19.5
Llama 3 - PT 1	19.2	20.5	15.8	19.2	22.5	15.0	23.3	23.3	17.9	13.3	25.9	28.3	21.2
Llama 3 - PT 100	13.3	20.8	18.3	21.7	20.0	12.5	23.3	29.2	17.0	10.8	28.6	28.3	19.5
Gemma 2 - L100 50	14.2	21.1	18.3	22.5	20.8	10.8	25.0	28.3	16.1	11.7	27.7	30.0	20.3
Llama 3 - L100 50	12.5	20.6	19.2	20.8	20.0	10.8	24.2	30.0	15.2	10.8	28.6	27.5	19.5
Qwen 2.5 - L100 50	26.7	27.3	25.0	21.7	26.7	27.5	27.5	25.0	29.5	25.0	29.5	40.0	22.9
Aya-Expanse - L100 50	12.5	20.7	18.3	21.7	20.0	10.8	24.2	29.2	15.2	10.8	28.6	29.2	19.5
Centurio Aya	12.5	20.7	18.3	21.7	20.0	11.7	24.2	29.2	15.2	10.8	28.6	29.2	19.5
Centurio Qwen	28.3	27.0	18.3	20.0	33.3	32.5	29.2	22.5	25.0	22.5	30.4	30.0	33.1

Table 21: VLOD

	en	avg.	id	sw	ta	tr	zh
Phi 3.5 - English	82.1	61.4	65.6	50.8	53.3	63.8	73.2
Phi 3.5 - T5 50	81.5	61.8	66.4	53.4	53.7	61.6	73.8
Phi 3.5 - T5-4 50	81.2	64.3	68.7	52.3	54.3	70.2	76.2
Phi 3.5 - T5-3 50	81.5	65.9	70.8	56.4	56.7	68.9	76.7
Phi 3.5 - T5-2 50	79.7	66.4	70.2	62.2	57.5	66.7	75.4
Phi 3.5 - L100 50	79.6	64.4	69.0	59.0	53.6	67.5	73.0
Llama 3 - English	85.2	65.0	68.8	52.5	54.3	69.7	79.8
Llama 3 - T5 50	84.5	67.1	73.8	55.7	53.6	72.7	79.6
Llama 3 - L100 50	83.7	74.2	75.3	71.4	68.4	79.8	76.0
Phi 3.5 - L100 1	71.9	61.4	65.1	56.1	54.3	65.2	66.1
Phi 3.5 - L100 10	74.1	63.4	66.8	58.1	57.2	65.1	70.0
Phi 3.5 - L100 24	76.0	61.6	63.4	57.6	56.9	64.0	66.3
Phi 3.5 - L100 50	79.6	64.4	69.0	59.0	53.6	67.5	73.0
Phi 3.5 - L100 75	81.7	64.7	71.3	54.4	56.1	64.8	77.0
Phi 3.5 - L100 90	83.1	64.3	70.7	56.3	53.8	62.8	77.8
Llama 3 - L100 10	80.0	72.9	71.9	70.8	71.7	75.7	74.2
Llama 3 - L100 50	83.7	74.2	75.3	71.4	68.4	79.8	76.0
Llama 3 - L100 90	85.1	71.1	73.4	63.7	65.1	75.7	77.6
Phi 3.5 - L100 50	79.6	64.4	69.0	59.0	53.6	67.5	73.0
Phi 3.5 - PT 100	82.0	65.6	68.6	59.4	57.9	67.6	74.5
Phi 3.5 - PT 50	82.5	69.9	75.2	64.0	64.1	71.1	74.9
Phi 3.5 - PT 1	81.9	67.9	74.0	64.0	60.2	68.0	73.4
Llama 3 - L100 50	83.7	74.2	75.3	71.4	68.4	79.8	76.0
Llama 3 - PT 1	87.5	80.4	82.5	75.5	77.1	84.5	82.3
Llama 3 - PT 100	86.5	78.9	81.3	73.0	75.1	83.4	81.5
Gemma 2 - L100 50	82.5	73.0	72.6	71.4	68.3	76.4	76.2
Llama 3 - L100 50	83.7	74.2	75.3	71.4	68.4	79.8	76.0
Qwen 2.5 - L100 50	89.6	79.4	84.8	73.9	65.2	86.6	86.6
Aya-Expanse - L100 50	87.0	80.2	83.9	75.6	71.7	86.9	83.0
Centurio Aya	85.0	77.9	79.5	70.9	73.4	83.4	82.4
Centurio Qwen	89.6	81.7	85.0	76.8	76.0	84.2	86.7

Table 22: MaRVL

	en	avg.	fr	hi	he	ro	th	zh
Phi 3.5 - English	53.0	9.2	14.3	11.9	7.9	7.2	7.0	7.2
Phi 3.5 - T5 50	51.3	25.6	41.0	30.6	17.5	15.6	27.5	21.5
Phi 3.5 - T5-4 50	51.0	33.1	45.4	50.7	27.0	23.7	32.5	19.5
Phi 3.5 - T5-3 50	53.7	36.7	41.0	45.9	33.0	36.6	40.4	23.5
Phi 3.5 - T5-2 50	53.4	35.9	42.3	48.0	33.3	35.1	32.8	23.8
Phi 3.5 - L100 50	54.4	36.6	43.0	48.0	30.8	35.1	39.1	23.5
Llama 3 - English	55.4	7.7	9.2	10.9	6.7	4.5	8.3	6.8
Llama 3 - T5 50	41.3	20.2	45.1	12.6	2.9	24.3	14.6	21.8
Llama 3 - L100 50	52.7	42.3	42.3	54.4	40.6	40.5	52.6	23.1
Phi 3.5 - L100 1	48.0	33.8	39.9	45.2	32.4	32.4	32.8	19.9
Phi 3.5 - L100 10	52.0	35.4	44.7	45.6	34.6	36.0	29.5	22.1
Phi 3.5 - L100 24	50.7	35.1	44.0	44.6	29.8	33.0	38.1	21.2
Phi 3.5 - L100 50	54.4	36.6	43.0	48.0	30.8	35.1	39.1	23.5
Phi 3.5 - L100 75	51.0	32.5	42.0	36.4	29.8	33.3	31.8	21.8
Phi 3.5 - L100 90	54.7	29.7	41.6	28.2	27.3	28.5	30.5	21.8
Llama 3 - L100 10	49.0	41.9	37.9	53.4	45.7	41.4	51.0	21.8
Llama 3 - L100 50	52.7	42.3	42.3	54.4	40.6	40.5	52.6	23.1
Llama 3 - L100 90	52.7	40.6	43.3	52.7	36.2	40.2	49.0	22.1
Phi 3.5 - L100 50	54.4	36.6	43.0	48.0	30.8	35.1	39.1	23.5
Phi 3.5 - PT 100	54.0	36.2	44.0	48.6	32.4	33.9	36.8	21.5
Phi 3.5 - PT 50	53.4	39.0	45.7	49.3	39.4	36.6	40.7	22.1
Phi 3.5 - PT 1	55.7	39.7	44.7	52.0	41.0	40.8	40.1	19.9
Llama 3 - L100 50	52.7	42.3	42.3	54.4	40.6	40.5	52.6	23.1
Llama 3 - PT 1	55.0	48.5	47.4	57.1	56.2	47.4	57.3	25.7
Llama 3 - PT 100	58.1	47.4	44.7	54.8	54.0	47.1	57.3	26.4
Gemma 2 - L100 50	51.7	41.5	39.6	52.4	44.1	39.3	48.7	24.8
Llama 3 - L100 50	52.7	42.3	42.3	54.4	40.6	40.5	52.6	23.1
Qwen 2.5 - L100 50	58.7	45.8	46.4	51.4	50.2	41.7	57.9	27.0
Aya-Expanse - L100 50	53.4	47.2	46.4	58.8	59.4	49.9	41.4	27.4
Centurio Aya	55.7	49.3	45.1	62.9	58.7	51.1	46.7	31.6
Centurio Qwen	60.1	47.7	47.1	56.8	45.1	47.7	57.0	32.2

Table 23: MaXM

	avg.	ar	de	fr	it	ja	ko	ru	th	vi
Phi 3.5 - English	3.2	0.9	6.5	9.3	8.1	0.8	0.7	1.6	0.0	1.1
Phi 3.5 - T5 50	5.7	1.7	12.0	15.9	10.1	2.4	3.8	2.6	0.9	1.8
Phi 3.5 - T5-4 50	5.9	2.7	14.0	15.1	9.6	3.5	3.8	1.9	0.9	1.6
Phi 3.5 - T5-3 50	5.8	2.0	13.5	14.6	9.4	3.9	3.8	2.4	0.9	2.0
Phi 3.5 - T5-2 50	6.6	5.3	15.9	15.1	9.4	4.1	3.8	2.5	0.4	2.7
Phi 3.5 - L100 50	6.3	2.8	15.8	16.8	8.9	3.9	2.7	2.8	0.4	2.9
Llama 3 - English	3.2	0.3	6.9	8.0	8.7	0.7	0.5	0.7	0.4	2.7
Llama 3 - T5 50	5.6	2.0	14.2	15.0	9.1	1.9	1.4	2.6	1.3	2.8
Llama 3 - L100 50	6.0	2.1	11.9	15.8	7.2	2.1	3.2	2.4	4.8	4.1
Phi 3.5 - L100 1	4.7	2.0	12.0	9.4	7.5	3.4	3.4	1.9	0.9	2.3
Phi 3.5 - L100 10	5.7	3.0	12.1	14.2	8.6	4.6	4.1	2.1	0.9	1.5
Phi 3.5 - L100 24	6.2	3.6	14.0	15.8	8.7	3.1	3.8	3.3	0.9	2.5
Phi 3.5 - L100 50	6.3	2.8	15.8	16.8	8.9	3.9	2.7	2.8	0.4	2.9
Phi 3.5 - L100 75	6.3	2.6	13.8	18.3	8.7	4.3	2.9	2.8	0.9	2.8
Phi 3.5 - L100 90	7.0	2.6	14.7	19.3	10.4	3.6	4.1	3.2	3.5	1.5
Llama 3 - L100 10	5.3	1.6	11.3	13.8	7.5	2.9	3.4	2.6	0.9	3.5
Llama 3 - L100 50	6.0	2.1	11.9	15.8	7.2	2.1	3.2	2.4	4.8	4.1
Llama 3 - L100 90	6.5	2.1	14.0	17.8	9.7	2.5	3.8	2.8	2.2	3.5
Phi 3.5 - L100 50	6.3	2.8	15.8	16.8	8.9	3.9	2.7	2.8	0.4	2.9
Phi 3.5 - PT 100	6.9	3.7	16.0	15.9	11.3	3.4	3.2	2.9	2.2	3.5
Phi 3.5 - PT 50	6.1	1.8	14.8	15.8	10.5	3.5	2.9	2.6	0.9	2.1
Phi 3.5 - PT 1	6.2	1.6	14.9	15.9	11.1	3.7	3.0	1.7	0.9	2.7
Llama 3 - L100 50	6.0	2.1	11.9	15.8	7.2	2.1	3.2	2.4	4.8	4.1
Llama 3 - PT 1	6.9	2.4	17.1	16.6	9.1	3.4	4.5	2.5	1.7	5.2
Llama 3 - PT 100	8.3	2.6	18.7	19.6	11.4	4.0	4.3	4.0	4.8	5.3
Gemma 2 - L100 50	4.3	1.7	11.1	8.1	7.1	3.0	2.3	2.1	1.7	1.7
Llama 3 - L100 50	6.0	2.1	11.9	15.8	7.2	2.1	3.2	2.4	4.8	4.1
Qwen 2.5 - L100 50	6.4	5.5	12.0	13.0	10.3	3.0	3.2	2.9	2.2	5.2
Aya-Expanse - L100 50	6.2	3.7	13.2	13.9	9.5	3.0	3.4	3.4	1.7	3.6
Centurio Aya	11.1	6.7	19.9	22.5	16.7	5.0	9.0	5.2	5.2	9.7
Centurio Qwen	11.9	4.6	22.7	26.5	18.6	5.9	9.9	5.0	5.2	8.9

Table 24: MTVQA

	en	avg.	bn	de	id	ko	pt	ru	zh
Phi 3.5 - English	59.7	37.2	4.9	47.8	33.2	38.2	47.1	42.1	47.2
Phi 3.5 - T5 50	54.1	34.1	2.6	44.6	34.3	36.3	43.8	36.4	41.0
Phi 3.5 - T5-4 50	52.0	37.4	5.7	45.6	38.7	40.4	45.2	43.4	42.7
Phi 3.5 - T5-3 50	54.8	40.6	22.7	46.5	42.1	39.8	46.0	43.6	43.6
Phi 3.5 - T5-2 50	57.8	45.3	27.4	50.3	46.0	46.4	48.6	49.5	48.9
Phi 3.5 - L100 50	56.6	45.1	27.0	51.4	44.8	44.9	50.8	48.2	48.7
Llama 3 - English	61.9	39.2	13.2	49.0	35.6	39.1	44.9	44.1	48.4
Llama 3 - T5 50	49.3	33.8	5.9	43.8	38.0	32.4	41.7	37.3	37.4
Llama 3 - L100 50	60.6	51.0	46.7	54.1	51.2	49.4	53.4	51.2	51.3
Phi 3.5 - L100 1	48.4	40.3	28.2	43.9	41.1	40.6	43.0	42.8	42.8
Phi 3.5 - L100 10	51.8	42.2	27.6	46.3	43.0	42.2	45.7	44.8	45.6
Phi 3.5 - L100 24	53.8	42.9	29.1	47.6	43.4	42.2	46.6	45.9	45.4
Phi 3.5 - L100 50	56.6	45.1	27.0	51.4	44.8	44.9	50.8	48.2	48.7
Phi 3.5 - L100 75	58.6	45.8	26.4	52.4	44.4	45.4	51.9	49.9	50.0
Phi 3.5 - L100 90	58.5	42.1	14.2	53.0	39.8	43.3	51.4	45.7	47.7
Llama 3 - L100 10	54.9	45.0	40.5	46.4	45.7	42.5	46.5	46.2	47.4
Llama 3 - L100 50	60.6	51.0	46.7	54.1	51.2	49.4	53.4	51.2	51.3
Llama 3 - L100 90	61.9	51.4	42.5	56.2	51.7	50.1	54.6	52.5	52.1
Phi 3.5 - L100 50	56.6	45.1	27.0	51.4	44.8	44.9	50.8	48.2	48.7
Phi 3.5 - PT 100	58.0	46.1	29.5	52.8	46.1	44.5	51.7	49.5	48.3
Phi 3.5 - PT 50	58.3	47.6	35.4	52.8	48.7	45.5	52.5	49.6	48.6
Phi 3.5 - PT 1	58.3	47.0	37.6	52.6	46.8	44.1	51.5	48.1	48.1
Llama 3 - L100 50	60.6	51.0	46.7	54.1	51.2	49.4	53.4	51.2	51.3
Llama 3 - PT 1	61.1	55.1	52.8	56.6	56.0	53.9	56.0	55.4	55.0
Llama 3 - PT 100	61.6	53.0	50.4	54.9	53.6	52.4	53.0	53.1	53.4
Gemma 2 - L100 50	56.5	47.5	43.9	51.6	47.6	44.2	50.1	47.5	47.5
Llama 3 - L100 50	60.6	51.0	46.7	54.1	51.2	49.4	53.4	51.2	51.3
Qwen 2.5 - L100 50	60.3	51.9	44.2	54.8	53.1	51.3	54.3	53.2	52.8
Aya-Expanse - L100 50	60.5	52.5	45.2	54.6	53.8	51.7	54.7	53.9	53.4
Centurio Aya	59.1	53.2	43.4	56.9	54.4	53.6	56.2	54.0	54.3
Centurio Qwen	60.6	54.8	49.9	57.0	54.9	53.5	57.2	55.8	55.6

Table 25: xGQA

	en	avg.	ar	bn	cs	da	de	el	es	fa	fi	fil	fr	he	hi	hr	hu	id	
Phi 3.5 - English	33.6	1.2	0.0	0.0	0.7	1.1	1.7	0.0	10.5	0.0	0.4	1.6	4.4	0.0	0.0	0.5	0.6	1.5	
Phi 3.5 - T5 50	33.0	9.5	7.8	0.6	3.9	8.2	24.7	0.6	34.4	0.4	1.8	1.9	39.1	3.4	3.5	2.6	4.0	7.9	
Phi 3.5 - T5-4 50	25.2	11.8	6.9	1.0	13.8	10.8	24.4	1.4	27.3	8.6	7.0	3.6	31.3	3.5	9.7	9.2	10.5	8.4	
Phi 3.5 - T5-3 50	32.7	13.6	6.5	5.9	13.8	18.2	25.4	7.0	31.0	7.1	5.7	11.9	30.7	7.6	7.2	9.4	8.2	22.9	
Phi 3.5 - T5-2 50	29.9	11.1	4.5	5.5	10.7	12.4	20.4	6.4	20.4	5.7	5.5	10.1	24.8	7.4	6.7	8.0	7.1	17.5	
Phi 3.5 - L100 50	31.0	13.2	5.6	3.3	10.9	18.0	26.4	4.5	30.9	4.1	4.4	11.1	38.5	6.7	6.3	7.3	7.8	22.4	
Llama 3 - English	75.6	1.1	0.1	0.0	1.3	2.3	1.6	0.1	2.1	0.1	0.8	2.9	3.4	0.0	0.0	0.7	1.3	2.1	
Llama 3 - T5 50	76.1	12.6	27.9	0.5	1.6	20.7	29.2	0.6	61.6	0.4	1.1	2.9	58.2	0.0	0.4	0.9	1.6	19.6	
Llama 3 - L100 50	72.6	28.5	25.6	14.0	30.5	38.1	27.6	23.3	56.0	24.3	13.8	29.0	50.9	15.5	22.5	19.7	18.0	39.9	
Phi 3.5 - L100 1	43.3	13.3	5.2	4.3	11.1	15.3	24.5	5.3	25.4	5.4	6.1	13.3	37.1	7.0	7.2	7.9	6.7	22.6	
Phi 3.5 - L100 10	38.9	12.7	4.7	4.1	11.4	12.5	23.6	5.9	28.1	4.8	5.3	14.2	33.3	8.5	8.6	8.2	6.3	22.6	
Phi 3.5 - L100 24	31.5	13.2	5.2	5.0	11.6	15.1	22.6	6.9	30.9	4.9	6.1	12.0	39.2	8.2	7.7	8.2	5.1	22.4	
Phi 3.5 - L100 50	31.0	13.2	5.6	3.3	10.9	18.0	26.4	4.5	30.9	4.1	4.4	11.1	38.5	6.7	6.3	7.3	7.8	22.4	
Phi 3.5 - L100 75	36.5	12.0	4.4	2.5	9.1	13.6	25.0	3.0	25.7	3.4	3.8	7.1	33.6	6.3	5.2	5.9	7.0	20.2	
Phi 3.5 - L100 90	34.2	9.4	4.0	1.9	6.7	9.2	21.8	2.8	23.4	2.0	3.8	4.1	26.2	4.4	3.7	4.7	4.9	12.5	
Llama 3 - L100 10	74.8	28.9	23.0	11.9	25.8	43.6	26.0	24.6	53.7	24.9	16.0	30.2	52.6	17.1	20.1	20.5	18.5	43.3	
Llama 3 - L100 50	72.6	28.5	25.6	14.0	30.5	38.1	27.6	23.3	56.0	24.3	13.8	29.0	50.9	15.5	22.5	19.7	18.0	39.9	
Llama 3 - L100 90	73.6	23.0	18.2	7.8	24.1	36.7	23.0	19.5	54.2	17.6	10.5	24.0	51.9	9.7	20.2	15.4	15.3	33.0	
Phi 3.5 - L100 50	31.0	13.2	5.6	3.3	10.9	18.0	26.4	4.5	30.9	4.1	4.4	11.1	38.5	6.7	6.3	7.3	7.8	22.4	
Phi 3.5 - PT 100	35.9	13.5	5.3	5.0	13.9	15.7	26.5	5.9	29.6	5.4	4.1	9.1	33.5	8.3	6.9	8.8	7.1	22.3	
Phi 3.5 - PT 50	37.1	17.3	7.7	9.0	16.5	21.2	27.8	9.3	38.0	8.2	7.0	15.2	42.4	10.9	10.4	11.8	9.7	28.5	
Phi 3.5 - PT 1	33.1	17.4	6.3	9.3	17.2	22.1	26.9	8.2	37.5	9.1	7.2	13.9	40.6	12.2	9.1	11.5	11.1	28.9	
Llama 3 - L100 50	72.6	28.5	25.6	14.0	30.5	38.1	27.6	23.3	56.0	24.3	13.8	29.0	50.9	15.5	22.5	19.7	18.0	39.9	
Llama 3 - PT 1	80.8	35.3	30.6	15.4	35.5	51.3	34.0	28.2	65.4	32.3	17.9	36.3	62.5	24.6	27.4	26.9	24.7	49.2	
Llama 3 - PT 100	77.9	31.8	26.1	14.4	35.4	43.5	33.4	27.0	60.7	23.0	14.6	31.7	58.9	18.2	24.6	22.2	22.6	45.2	
Gemma 2 - L100 50	66.6	27.5	24.5	17.9	28.6	35.1	26.0	18.2	54.7	29.4	13.7	26.8	54.3	22.8	21.6	17.7	20.1	43.8	
Qwen 2.5 - L100 50	74.8	27.8	28.6	13.9	26.1	35.4	29.6	11.6	58.7	17.1	10.2	26.0	55.6	22.2	16.3	19.8	11.7	40.3	
Aya-Expanse - L100 50	75.6	33.4	40.4	12.2	39.4	37.7	32.7	31.9	69.2	41.6	8.4	26.2	67.6	42.5	24.5	19.1	12.5	50.3	
Centurio Aya	78.4	39.2	40.4	18.5	33.9	40.0	38.6	35.3	69.7	55.8	11.0	34.0	71.3	47.1	26.3	24.9	19.6	58.3	
Centurio Qwen	79.1	34.4	36.6	17.1	29.7	43.1	32.0	19.2	69.2	31.2	12.0	33.6	67.6	27.6	20.3	22.0	18.7	50.4	
	it	ja	ko	mi	nl	no	pl	pt	quz	ro	ru	sv	sw	te	th	tr	uk	vi	zh
Phi 3.5 - English	9.2	0.1	0.0	0.0	1.6	1.5	0.7	1.9	0.1	0.7	0.4	1.1	0.6	0.0	0.2	0.3	0.1	0.6	0.0
Phi 3.5 - T5 50	28.5	27.6	1.6	0.1	20.4	8.8	4.8	30.2	0.7	5.4	17.5	11.8	1.3	0.0	4.0	3.2	5.7	2.6	12.8
Phi 3.5 - T5-4 50	24.9	27.9	3.1	2.1	21.7	12.0	14.1	24.0	0.5	6.0	22.5	24.1	2.1	0.0	5.8	9.4	4.6	18.8	10.6
Phi 3.5 - T5-3 50	26.0	27.3	3.0	1.8	28.5	14.1	13.6	22.1	0.4	11.0	16.7	23.8	1.5	0.0	12.9	7.6	10.0	15.9	20.2
Phi 3.5 - T5-2 50	18.5	27.0	2.5	2.0	18.9	11.2	10.0	18.1	0.4	7.6	21.1	17.8	7.0	0.0	12.6	7.9	10.3	12.7	9.5
Phi 3.5 - L100 50	30.6	23.7	2.5	2.6	24.2	18.8	10.8	29.0	1.1	8.1	18.6	17.8	8.1	1.5	10.2	7.2	8.3	14.7	15.6
Llama 3 - English	1.5	0.2	0.0	0.2	3.0	1.8	1.1	2.6	0.8	1.2	0.7	2.0	0.8	0.0	0.4	0.5	0.3	0.3	0.3
Llama 3 - T5 50	26.8	35.8	0.1	0.2	34.9	15.7	11.3	12.5	1.3	5.4	12.8	18.3	0.6	0.0	4.4	6.1	0.2	10.1	17.2
Llama 3 - L100 50	44.1	33.8	13.4	24.9	50.1	41.5	26.9	45.1	1.3	22.6	23.5	42.7	28.9	11.0	27.7	21.8	21.9	49.7	16.9
Phi 3.5 - L100 1	28.1	25.4	2.1	4.0	22.3	17.7	12.0	24.9	1.4	10.7	21.4	17.8	11.1	1.4	12.6	7.6	8.2	17.2	16.9
Phi 3.5 - L100 10	23.3	25.2	2.4	3.0	23.0	15.6	12.0	25.8	0.7	8.2	20.0	15.9	9.9	1.8	11.8	6.7	11.9	14.7	11.8
Phi 3.5 - L100 24	24.7	25.8	3.3	4.8	24.8	18.5	13.6	17.7	0.7	9.1	17.8	17.5	9.3	2.3	13.3	6.5	9.6	15.4	14.4
Phi 3.5 - L100 50	30.6	23.7	2.5	2.6	24.2	18.8	10.8	29.0	1.1	8.1	18.6	17.8	8.1	1.5	10.2	7.2	8.3	14.7	15.6
Phi 3.5 - L100 75	29.7	24.8	2.0	3.3	23.0	17.1	9.8	27.8	0.8	6.2	19.6	15.2	5.0	1.4	10.9	5.7	8.9	13.0	18.5
Phi 3.5 - L100 90	21.3	21.0	2.0	1.2	12.7	11.8	7.3	22.5	0.8	5.9	16.3	16.1	5.6	0.6	7.5	3.9	7.8	8.8	20.3
Llama 3 - L100 10	40.3	35.0	13.9	29.4	53.4	41.9	25.6	24.8	1.6	19.8	25.3	44.0	30.3	13.8	28.8	22.1	21.1	47.4	20.2
Llama 3 - L100 50	44.1	33.8	13.4	24.9	50.1	41.5	26.9	45.1	1.3	22.6	23.5	42.7	28.9	11.0	27.7	21.8	21.9	49.7	16.9
Llama 3 - L100 90	38.0	25.6	10.4	17.5	46.1	33.1	19.8	41.2	0.2	17.1	20.6	38.1	14.6	5.8	23.2	16.3	0.3	43.9	13.3
Phi 3.5 - L100 50	30.6	23.7	2.5	2.6	24.2	18.8	10.8	29.0	1.1	8.1	18.6	17.8	8.1	1.5	10.2	7.2	8.3	14.7	15.6
Phi 3.5 - PT 100	30.3	25.7	2.7	3.9	21.6	20.1	12.0	21.8	0.9	9.5	19.5	18.9	8.5	1.4	13.6	7.5	8.5	14.9	23.9
Phi 3.5 - PT 50	33.9	26.2	3.2	7.2	30.0	24.7	16.1	29.1	2.5	14.7	21.3	24.1	15.3	4.3	18.6	8.0	9.8	19.4	22.3
Phi 3.5 - PT 1	34.8	30.5	2.9	7.9	27.7	26.4	14.9	31.2	2.3	14.4	22.4	23.8	14.7	4.4	18.4	10.8	10.5	18.8	21.1
Llama 3 - L100 50	44.1	33.8	13.4	24.9	50.1	41.5	26.9	45.1	1.3	22.6	23.5	42.7	28.9	11.0	27.7	21.8	21.9	49.7	16.9
Llama 3 - PT 1	51.6	38.7	15.2	35.9	59.1	49.2	32.5	51.1	2.9	30.7	32.3	51.8	38.0	17.4	36.2	29.3	26.4	58.1	16.5
Llama 3 - PT 100	49.2	35.9	14.0	26.5	55.3	45.8	32.3	51.6	0.9	25.2	30.7	47.7	29.3	12.4	32.0	27.0	25.2	56.0	15.2
Gemma 2 - L100 50	39.7	36.3	11.5	21.5	46.2	38.7	25.1	45.3	1.8	21.7	24.0	37.9	27.2	13.1	28.5	19.0	0.2	50.1	18.5
Qwen 2.5 - L100 50	43.8	39.0	13.4	21.8	48.9	36.6	25.0	52.3	0.9	16.9	33.6	36.7	18.9	8.0	38.1	17.4	18.1	59.5	19.9
Aya-Expanse - L100 50	53.6	39.5	18.4	20.1	59.6	36.3	35.5	50.6	0.7	31.3	31.9	34.9	21.3	9.2	19.0	29.8	27.7	68.2	26.1
Centurio Aya	60.4	49.1	21.3	33.7	61.7	42.5	37.9	59.3	2.3	34.6	38.0	45.9	29.9	15.1	26.0	30.6	30.6</		

	en	avg.	ar	es	fr	ru
Phi 3.5 - English	59.6	55.0	52.3	54.9	57.6	55.2
Phi 3.5 - T5 50	59.9	51.8	49.7	51.3	55.2	51.0
Phi 3.5 - T5-4 50	58.9	48.3	47.7	47.1	51.2	47.3
Phi 3.5 - T5-3 50	58.6	50.5	46.6	51.3	52.6	51.3
Phi 3.5 - T5-2 50	58.5	53.6	50.7	54.7	55.1	54.0
Phi 3.5 - L100 50	59.6	53.3	49.9	53.7	56.8	52.6
Llama 3 - English	46.1	36.3	33.4	37.0	36.5	38.2
Llama 3 - T5 50	45.4	37.5	36.6	36.1	38.7	38.7
Llama 3 - L100 50	59.7	54.8	53.0	54.7	56.3	55.4
Phi 3.5 - L100 1	55.2	48.2	42.2	50.6	51.1	48.8
Phi 3.5 - L100 10	58.3	53.4	50.5	54.3	55.7	53.1
Phi 3.5 - L100 24	58.2	48.4	43.5	50.3	52.7	47.3
Phi 3.5 - L100 50	59.6	53.3	49.9	53.7	56.8	52.6
Phi 3.5 - L100 75	61.9	54.5	50.3	56.2	57.5	54.2
Phi 3.5 - L100 90	60.0	50.5	43.6	53.8	55.1	49.5
Llama 3 - L100 10	60.8	55.3	56.3	52.6	55.0	57.1
Llama 3 - L100 50	59.7	54.8	53.0	54.7	56.3	55.4
Llama 3 - L100 90	57.8	51.1	48.9	52.3	52.1	51.0
Phi 3.5 - L100 50	59.6	53.3	49.9	53.7	56.8	52.6
Phi 3.5 - PT 100	54.3	45.4	40.1	47.8	49.4	44.4
Phi 3.5 - PT 50	58.9	52.5	49.0	53.8	54.6	52.5
Phi 3.5 - PT 1	56.8	49.7	46.8	49.6	53.9	48.6
Llama 3 - L100 50	59.7	54.8	53.0	54.7	56.3	55.4
Llama 3 - PT 1	61.7	59.4	58.8	59.0	60.0	59.7
Llama 3 - PT 100	60.3	57.3	56.5	56.5	58.3	57.8
Gemma 2 - L100 50	59.9	55.0	53.1	54.6	57.1	55.1
Llama 3 - L100 50	59.7	54.8	53.0	54.7	56.3	55.4
Qwen 2.5 - L100 50	57.8	52.6	55.7	47.5	52.5	54.8
Aya-Expanse - L100 50	58.2	54.7	54.7	54.0	56.4	53.5
Centurio Aya	65.0	62.4	61.7	61.0	64.3	62.7
Centurio Qwen	75.4	70.2	68.8	70.9	70.5	70.8

Table 28: XVNLI

	en	avg.	ar	fr	hi	id	ja	pt
Phi 3.5 - English	38.4	36.2	36.2	41.9	29.9	35.4	34.2	39.7
Phi 3.5 - T5 50	36.7	36.2	31.5	38.9	31.6	37.0	34.9	43.4
Phi 3.5 - T5-4 50	37.0	33.9	33.2	39.9	29.2	32.3	31.2	37.7
Phi 3.5 - T5-3 50	37.3	35.8	32.5	39.3	32.3	37.0	36.8	37.0
Phi 3.5 - T5-2 50	37.6	35.1	32.2	40.3	32.6	34.7	32.3	38.7
Phi 3.5 - L100 50	36.6	32.0	28.5	35.9	27.8	32.0	31.2	36.7
Llama 3 - English	33.2	32.4	30.9	34.2	30.6	32.7	30.5	35.7
Llama 3 - T5 50	33.4	32.4	34.9	36.6	28.9	31.3	30.9	31.6
Llama 3 - L100 50	33.0	31.7	31.5	34.6	34.0	31.6	27.9	30.6
Phi 3.5 - L100 1	37.3	34.1	32.5	40.3	30.9	31.3	31.6	37.7
Phi 3.5 - L100 10	36.1	30.9	27.5	33.9	28.2	28.6	32.7	34.7
Phi 3.5 - L100 24	34.4	31.9	28.5	35.9	29.2	30.3	33.5	34.0
Phi 3.5 - L100 50	36.6	32.0	28.5	35.9	27.8	32.0	31.2	36.7
Phi 3.5 - L100 75	36.2	33.2	31.9	38.9	29.2	32.7	29.0	37.4
Phi 3.5 - L100 90	37.1	31.9	30.5	35.6	25.8	31.0	33.8	34.7
Llama 3 - L100 10	32.6	30.0	26.8	31.5	26.8	31.6	32.0	31.3
Llama 3 - L100 50	33.0	31.7	31.5	34.6	34.0	31.6	27.9	30.6
Llama 3 - L100 90	32.7	33.5	30.5	35.9	30.9	35.4	31.2	37.0
Phi 3.5 - L100 50	36.6	32.0	28.5	35.9	27.8	32.0	31.2	36.7
Phi 3.5 - PT 100	33.4	30.2	28.5	32.9	28.9	30.0	27.5	33.3
Phi 3.5 - PT 50	35.0	33.4	30.9	39.3	33.7	31.0	30.5	35.0
Phi 3.5 - PT 1	36.0	31.3	26.5	35.9	29.2	32.0	28.3	36.0
Llama 3 - L100 50	33.0	31.7	31.5	34.6	34.0	31.6	27.9	30.6
Llama 3 - PT 1	38.6	35.2	33.9	34.2	34.0	35.0	36.1	38.0
Llama 3 - PT 100	36.9	36.1	34.6	36.2	36.8	36.7	36.1	36.0
Gemma 2 - L100 50	32.8	32.0	32.5	30.9	33.0	30.6	32.7	32.0
Llama 3 - L100 50	33.0	31.7	31.5	34.6	34.0	31.6	27.9	30.6
Qwen 2.5 - L100 50	39.8	39.7	38.6	40.3	34.4	40.7	38.7	45.5
Aya-Expanse - L100 50	36.8	35.4	34.9	35.2	37.5	36.4	34.6	33.7
Centurio Aya	37.6	37.2	36.2	38.9	38.8	39.7	34.2	35.4
Centurio Qwen	46.4	43.0	39.6	45.0	41.6	44.1	43.5	44.1

Table 29: xMMMU

	en	avg.	avg. Latin	avg. other	ar	de	hi	id	it	ko	ru	th	zh	zu
Phi 3.5 - English	65.8	55.8	62.3	51.5	50.2	63.5	58.5	61.4	64.0	49.0	52.1	49.1	49.8	60.2
Phi 3.5 - T5 50	75.2	60.2	70.9	53.1	50.2	70.8	65.4	71.8	71.6	49.8	54.1	51.0	48.0	69.4
Phi 3.5 - T5-4 50	74.2	60.8	71.4	53.7	52.2	71.5	65.5	72.8	73.1	51.1	53.9	49.6	49.6	68.4
Phi 3.5 - T5-3 50	70.4	58.7	67.7	52.8	51.6	66.9	61.0	69.6	67.2	50.0	53.6	48.9	51.4	67.0
Phi 3.5 - T5-2 50	68.4	56.2	64.2	50.8	49.5	64.5	58.4	65.4	64.9	50.0	50.5	48.8	47.9	62.0
Phi 3.5 - L100 50	69.6	58.0	67.2	51.9	49.9	68.0	62.4	69.0	67.9	48.6	52.5	49.6	48.4	64.1
Llama 3 - English	72.0	60.5	69.6	54.4	53.5	69.9	67.2	71.1	70.9	48.9	57.5	50.1	49.4	66.5
Llama 3 - T5 50	73.4	62.2	72.5	55.4	54.5	72.2	67.1	74.4	71.5	50.5	56.6	51.6	51.9	72.0
Llama 3 - L100 50	72.0	58.4	67.9	52.1	51.6	69.6	62.0	65.9	70.4	49.9	52.0	48.8	48.4	65.6
Phi 3.5 - L100 1	58.4	52.6	55.7	50.5	50.2	55.2	53.5	57.5	56.4	49.6	50.9	48.2	50.5	53.5
Phi 3.5 - L100 10	56.9	51.6	54.9	49.4	48.5	54.8	49.6	55.1	56.8	50.5	48.2	49.6	50.0	52.9
Phi 3.5 - L100 24	60.4	54.0	58.8	50.8	51.8	58.9	54.5	58.1	60.0	50.0	51.1	48.4	49.2	58.0
Phi 3.5 - L100 50	69.6	58.0	67.2	51.9	49.9	68.0	62.4	69.0	67.9	48.6	52.5	49.6	48.4	64.1
Phi 3.5 - L100 75	74.5	61.2	71.6	54.2	53.2	71.2	63.8	74.0	70.5	50.5	54.2	51.9	51.8	70.6
Phi 3.5 - L100 90	71.6	59.4	69.2	52.9	51.0	70.2	60.5	69.4	71.2	49.6	54.1	50.4	51.8	66.1
Llama 3 - L100 10	65.9	56.6	62.6	52.6	51.5	62.1	59.5	62.6	65.8	50.8	54.5	50.4	48.8	59.9
Llama 3 - L100 50	72.0	58.4	67.9	52.1	51.6	69.6	62.0	65.9	70.4	49.9	52.0	48.8	48.4	65.6
Llama 3 - L100 90	73.1	59.4	68.4	53.3	51.0	67.4	65.8	71.0	69.0	50.6	52.6	49.9	50.1	66.2
Phi 3.5 - L100 50	69.6	58.0	67.2	51.9	49.9	68.0	62.4	69.0	67.9	48.6	52.5	49.6	48.4	64.1
Phi 3.5 - PT 100	79.5	63.3	74.8	55.6	52.8	75.8	68.5	76.2	76.5	50.8	59.6	50.9	51.0	70.8
Phi 3.5 - PT 50	76.1	62.4	73.0	55.3	52.4	72.2	69.6	73.6	73.8	49.2	59.9	50.0	50.6	72.2
Phi 3.5 - PT 1	78.1	64.5	74.5	57.7	57.0	74.0	72.8	76.8	75.0	52.8	62.2	51.4	50.2	72.4
Llama 3 - L100 50	72.0	58.4	67.9	52.1	51.6	69.6	62.0	65.9	70.4	49.9	52.0	48.8	48.4	65.6
Llama 3 - PT 1	76.9	65.1	74.4	58.9	55.0	74.8	73.0	75.5	74.4	53.4	65.9	52.5	53.8	72.9
Llama 3 - PT 100	79.9	65.2	77.4	57.0	52.6	77.6	73.4	78.1	78.2	51.0	64.0	49.1	51.8	75.8
Phi 3.5 - OCR English	78.4	64.6	74.7	57.9	59.1	77.1	70.9	73.6	74.5	50.6	66.5	51.1	49.0	73.6
Phi 3.5 - OCR 50	81.2	66.7	76.7	60.0	61.4	78.6	72.1	76.0	77.1	51.5	71.5	52.1	51.6	75.0
Phi 3.5 - OCR 1	81.0	69.8	78.3	64.1	66.8	78.0	76.8	78.5	79.1	56.9	73.2	58.6	52.4	77.6
Phi 3.5 - OCR Latin-down	78.9	65.4	74.2	59.5	57.8	75.5	67.6	75.0	75.0	56.4	67.8	55.0	52.5	71.1
Phi 3.5 - OCR 50 (frozen)	76.1	62.1	70.8	56.3	59.2	73.2	63.2	66.2	76.1	50.0	68.0	47.8	49.8	67.8
Gemma 2 - L100 50	59.9	53.5	57.1	51.1	49.6	59.1	56.5	56.8	58.9	49.9	51.0	50.6	49.2	53.6
Llama 3 - L100 50	72.0	58.4	67.9	52.1	51.6	69.6	62.0	65.9	70.4	49.9	52.0	48.8	48.4	65.6
Qwen 2.5 - L100 50	82.8	62.5	75.1	54.0	51.5	76.4	66.5	76.5	76.5	50.1	55.2	51.0	49.8	71.1
Aya-Expanse - L100 50	79.1	63.5	75.2	55.7	53.9	77.2	71.4	75.6	75.0	50.6	56.0	51.1	51.0	73.1
Centurio Aya	83.1	74.2	80.9	69.7	75.9	82.1	80.1	81.4	80.6	68.8	73.5	66.5	53.4	79.5
Centurio Qwen	84.8	76.1	82.7	71.8	76.9	83.5	82.4	83.8	83.1	72.4	75.6	64.4	58.9	80.2

Table 30: SMPQA Ground

	en	avg.	avg. Latin	avg. other	ar	de	hi	id	it	ko	ru	th	zh	zu
Phi 3.5 - English	36.2	5.0	12.4	0.0	0.0	17.4	0.0	12.6	15.2	0.0	0.0	0.0	0.0	4.4
Phi 3.5 - T5 50	36.4	5.4	13.6	0.0	0.0	21.2	0.0	13.2	16.0	0.0	0.0	0.0	0.0	3.8
Phi 3.5 - T5-4 50	35.0	5.8	14.4	0.0	0.0	20.0	0.0	14.6	16.6	0.0	0.0	0.0	0.0	6.4
Phi 3.5 - T5-3 50	34.6	5.8	14.4	0.0	0.0	16.0	0.0	16.6	20.4	0.0	0.0	0.0	0.0	4.8
Phi 3.5 - T5-2 50	35.8	5.8	14.5	0.0	0.0	18.0	0.0	14.8	19.6	0.0	0.0	0.0	0.0	5.6
Phi 3.5 - L100 50	33.4	5.2	12.8	0.1	0.0	17.4	0.0	14.0	14.6	0.0	0.2	0.2	0.0	5.2
Llama 3 - English	41.0	8.5	21.1	0.0	0.0	24.4	0.0	21.6	23.8	0.0	0.0	0.2	0.0	14.8
Llama 3 - T5 50	41.4	8.2	20.4	0.0	0.0	25.2	0.0	21.8	23.4	0.0	0.0	0.2	0.0	11.2
Llama 3 - L100 50	39.2	7.3	18.2	0.0	0.0	21.6	0.0	18.8	21.6	0.0	0.0	0.2	0.0	10.8
Phi 3.5 - L100 1	22.0	4.0	10.1	0.0	0.0	12.0	0.0	9.0	14.0	0.0	0.0	0.0	0.0	5.2
Phi 3.5 - L100 10	24.6	4.1	10.3	0.0	0.0	11.6	0.0	10.0	14.2	0.0	0.0	0.0	0.0	5.4
Phi 3.5 - L100 24	26.0	3.8	9.5	0.1	0.0	12.2	0.0	8.4	12.6	0.0	0.0	0.4	0.0	4.8
Phi 3.5 - L100 50	33.4	5.2	12.8	0.1	0.0	17.4	0.0	14.0	14.6	0.0	0.2	0.2	0.0	5.2
Phi 3.5 - L100 75	38.4	6.0	15.1	0.0	0.0	21.0	0.0	14.8	18.6	0.0	0.2	0.0	0.0	5.8
Phi 3.5 - L100 90	39.8	6.5	16.1	0.0	0.0	21.0	0.0	17.0	21.8	0.0	0.0	0.0	0.0	4.8
Llama 3 - L100 10	32.0	6.3	15.6	0.1	0.0	17.8	0.0	15.8	19.2	0.0	0.0	0.4	0.0	9.6
Llama 3 - L100 50	39.2	7.3	18.2	0.0	0.0	21.6	0.0	18.8	21.6	0.0	0.0	0.2	0.0	10.8
Llama 3 - L100 90	40.0	7.5	18.8	0.0	0.0	21.2	0.0	21.0	20.4	0.0	0.0	0.2	0.0	12.6
Phi 3.5 - L100 50	33.4	5.2	12.8	0.1	0.0	17.4	0.0	14.0	14.6	0.0	0.2	0.2	0.0	5.2
Phi 3.5 - PT 100	44.0	9.9	24.5	0.2	0.0	31.4	0.0	25.6	26.8	0.0	1.2	0.2	0.0	14.0
Phi 3.5 - PT 50	41.8	9.4	23.1	0.2	0.0	27.8	0.0	24.4	25.0	0.0	1.2	0.2	0.0	15.0
Phi 3.5 - PT 1	42.2	9.5	23.7	0.1	0.0	27.2	0.0	24.4	29.0	0.0	0.4	0.0	0.0	14.0
Llama 3 - L100 50	39.2	7.3	18.2	0.0	0.0	21.6	0.0	18.8	21.6	0.0	0.0	0.2	0.0	10.8
Llama 3 - PT 1	48.4	11.4	27.9	0.4	0.0	29.6	0.2	30.6	30.6	0.0	1.6	0.4	0.0	20.6
Llama 3 - PT 100	48.8	10.5	25.0	0.8	0.0	28.8	2.6	26.2	28.4	0.2	1.8	0.4	0.0	16.6
Phi 3.5 - OCR English	55.8	18.3	39.9	3.9	5.2	38.6	2.4	43.2	41.6	0.0	15.2	0.4	0.0	36.4
Phi 3.5 - OCR 50	53.8	21.0	41.8	7.1	14.4	42.2	6.4	45.8	42.6	0.2	21.2	0.6	0.0	36.4
Phi 3.5 - OCR 1	54.8	22.2	43.5	8.0	17.2	43.8	6.2	46.4	42.8	1.2	21.4	1.8	0.0	40.8
Phi 3.5 - OCR Latin-down	54.6	22.4	41.0	9.9	20.2	41.6	7.0	42.6	43.0	2.8	25.6	3.4	0.6	36.8
Phi 3.5 - OCR 50 (frozen)	47.2	15.7	34.1	3.5	5.2	36.4	3.8	37.2	33.0	0.0	11.8	0.2	0.0	29.6
Gemma 2 - L100 50	28.6	3.8	9.4	0.1	0.0	13.8	0.0	10.4	8.4	0.0	0.0	0.4	0.0	5.0
Llama 3 - L100 50	39.2	7.3	18.2	0.0	0.0	21.6	0.0	18.8	21.6	0.0	0.0	0.2	0.0	10.8
Qwen 2.5 - L100 50	48.8	10.1	25.1	0.1	0.0	32.0	0.0	23.8	29.0	0.0	0.2	0.2	0.0	15.6
Aya-Expanse - L100 50	46.6	10.2	25.4	0.1	0.0	27.4	0.0	28.8	27.4	0.0	0.0	0.4	0.0	18.0
Centurio Aya	60.0	30.1	49.8	17.0	29.2	50.2	17.6	52.6	51.2	11.2	38.2	4.8	0.8	45.2
Centurio Qwen	65.2	31.7	54.3	16.6	21.4	53.2	21.4	55.4	56.6	16.2	34.8	5.2	0.6	52.2

Table 31: SMPQA Name

	en	avg.	af	am	cs	el	es	fa	fi	ha	hr	hu	ja	mi	nl	no	pl	ro	ta	te	zu
Centurio Aya	69.7	54.7	63.6	29.4	66.2	67.8	65.1	60.0	43.3	37.5	63.6	49.8	66.7	37.0	62.4	59.1	62.6	64.0	46.9	50.9	42.6
Centurio Qwen	72.7	56.2	65.3	47.4	62.2	56.7	67.0	53.6	48.8	36.7	65.4	54.1	67.6	39.1	63.7	63.6	60.4	58.5	45.2	63.4	49.5
Parrot	30.5	25.7	26.0	22.8	26.1	25.5	27.3	25.9	26.4	23.7	25.3	25.6	26.7	25.4	28.0	26.6	26.5	26.8	25.5	23.9	24.0
PALO 13B	61.4	41.1	48.4	25.9	47.9	35.8	53.2	37.5	42.7	26.1	52.3	47.9	49.1	31.0	48.9	51.2	46.1	46.5	28.9	32.2	28.3
PALO 7B	58.7	38.6	44.2	28.4	43.6	33.5	49.9	36.9	39.1	24.5	49.6	45.4	48.8	27.8	45.1	45.8	42.0	44.0	26.7	30.1	28.3
InternVL 2.5 4B	68.4	45.4	53.2	31.3	53.2	42.3	60.8	45.4	38.3	26.3	55.2	42.1	60.5	29.5	56.6	53.7	53.1	49.7	35.3	50.1	26.5
InternVL 2.5 8B	70.3	44.2	54.4	29.1	52.8	43.3	57.8	40.5	41.3	25.8	55.6	44.9	57.3	30.0	51.8	54.8	50.3	48.9	33.2	41.2	27.3
Qwen2-VL 2B	78.2	47.2	56.6	30.3	56.7	47.2	64.0	48.7	41.7	26.1	57.1	48.0	62.2	30.0	59.2	57.8	54.6	54.5	31.9	43.4	27.6
Qwen2-VL 7B	80.7	57.5	68.9	37.2	68.5	62.2	72.6	59.8	55.1	27.1	72.2	61.8	71.8	29.5	69.5	69.6	67.5	65.6	42.7	62.3	29.3
Maya	54.0	43.2	50.6	27.1	53.3	53.6	52.7	48.7	35.3	23.7	50.5	39.3	55.2	28.6	51.4	46.4	50.0	51.3	31.9	36.9	33.4
Llama-Vision	75.6	50.8	65.1	30.6	61.3	42.9	65.1	49.9	51.5	31.1	60.9	65.0	46.3	32.8	61.5	61.8	55.7	57.3	42.0	51.6	31.9
Phi 3.5 Vision	63.1	36.8	40.9	28.7	41.0	34.7	52.7	33.5	34.9	27.1	40.5	36.8	45.9	28.2	43.6	44.4	38.5	39.8	30.9	28.1	28.1
Pixtral 12B	71.0	54.2	62.3	34.3	61.6	58.3	66.1	57.3	52.0	27.7	67.1	60.4	64.8	31.9	58.6	62.1	59.8	59.0	56.7	64.5	25.0
Pangea	70.3	52.1	61.4	34.3	59.6	54.2	64.4	54.9	45.4	27.9	63.0	49.8	65.5	29.6	61.0	64.1	59.5	60.6	42.4	62.7	29.3
MiniCPM 2.6	72.6	47.4	56.0	29.9	55.1	46.6	62.1	48.5	41.8	22.9	59.5	44.9	62.9	29.0	57.8	55.2	54.7	52.7	34.5	53.9	33.4

Table 32: BIN-MC

	en	avg.	af	zh	it	pt	th	vi
Centurio Aya	53.0	41.2	52.8	51.4	47.7	27.4	27.8	40.3
Centurio Qwen	61.2	46.9	50.9	55.6	49.0	31.9	29.6	64.1
Parrot	46.6	36.2	38.0	37.8	36.8	25.9	23.5	55.1
PALO 13B	45.2	28.3	33.1	31.3	36.5	19.3	20.2	29.2
PALO 7B	41.0	29.1	34.4	31.5	32.7	21.8	21.1	33.4
InternVL 2.5 4B	63.2	50.3	46.0	60.9	50.3	34.9	39.1	70.4
InternVL 2.5 8B	67.0	53.3	57.7	61.7	53.2	33.0	39.1	75.2
Qwen2-VL 2B	47.9	40.5	38.0	51.6	36.4	36.2	26.1	54.9
Qwen2-VL 7B	56.1	49.7	50.9	58.6	46.8	34.7	38.3	69.0
Maya	49.2	36.3	48.5	46.4	36.6	25.9	20.0	40.3
Phi 3.5 Vision	56.3	40.7	51.5	54.4	44.1	25.2	24.3	44.4
Pixtral 12B	49.4	33.7	39.9	53.6	34.4	19.5	7.0	47.7
Pangea	58.0	45.5	50.3	58.6	49.0	32.2	27.8	55.3
MiniCPM 2.6	55.0	48.2	44.2	54.6	44.3	36.9	38.3	70.8

Table 33: M3Exam

	en	avg.	am	ber	bn	de	fil	ha	hi	ru	sw	th	zu
Centurio Aya	82.5	66.8	71.7	54.2	59.3	73.3	59.2	65.0	71.2	75.8	67.5	72.5	65.5
Centurio Qwen	87.5	73.1	77.5	49.2	62.7	80.8	78.3	76.7	72.9	85.0	70.0	81.7	69.0
Parrot	59.2	52.9	45.0	64.2	53.4	63.3	49.2	41.7	62.7	62.5	35.8	67.5	36.2
PALO 13B	63.3	26.2	25.0	55.0	0.8	44.2	47.5	40.0	0.0	5.8	32.5	0.0	37.1
PALO 7B	48.3	25.6	40.8	75.0	0.0	0.0	49.2	40.0	0.0	0.0	39.2	0.0	37.9
InternVL 2.5 4B	72.5	49.7	43.3	50.0	40.7	62.5	56.7	41.7	42.4	63.3	35.8	74.2	36.2
InternVL 2.5 8B	87.5	51.6	43.3	50.0	41.5	64.2	49.2	41.7	59.3	75.8	36.7	68.3	37.1
Qwen2-VL 2B	61.7	50.5	44.2	50.0	43.2	65.0	53.3	41.7	61.0	52.5	38.3	67.5	38.8
Qwen2-VL 7B	60.0	52.9	48.3	50.0	46.6	60.0	50.0	46.7	48.3	63.3	58.3	60.8	49.1
Maya	46.7	42.3	43.3	48.3	33.9	50.8	51.7	40.8	42.4	45.8	34.2	38.3	35.3
Phi 3.5 Vision	81.7	50.3	45.8	49.2	56.8	73.3	54.2	41.7	56.8	85.8	38.3	15.0	36.2
Pixtral 12B	55.8	47.7	51.7	32.5	47.5	63.3	51.7	44.2	16.1	54.2	65.8	53.3	44.0
Pangea	69.2	58.9	45.8	90.0	53.4	61.7	55.0	41.7	60.2	74.2	54.2	75.8	36.2
MiniCPM 2.6	52.5	49.1	45.0	55.8	49.2	45.8	48.3	40.8	44.1	59.2	48.3	65.8	37.9

Table 34: VGR

	en	avg.	am	ber	bn	de	fil	ha	hi	ru	sw	th	zu
Centurio Aya	12.5	20.7	18.3	21.7	20.0	11.7	24.2	29.2	15.2	10.8	28.6	29.2	19.5
Centurio Qwen	28.3	27.0	18.3	20.0	33.3	32.5	29.2	22.5	25.0	22.5	30.4	30.0	33.1
Parrot	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
PALO 13B	2.5	4.9	6.7	5.0	6.7	5.0	5.8	2.5	3.6	4.2	5.4	5.0	4.2
PALO 7B	5.8	6.8	8.3	9.2	10.0	5.8	6.7	4.2	9.8	5.0	4.5	5.8	5.1
InternVL 2.5 4B	24.2	21.0	18.3	26.7	17.5	20.8	20.0	23.3	22.3	20.0	23.2	20.0	18.6
InternVL 2.5 8B	57.5	29.0	25.0	22.5	25.8	38.3	36.7	25.8	41.1	35.8	15.2	30.0	22.9
Qwen2-VL 2B	22.5	20.4	17.5	20.0	13.3	26.7	25.0	24.2	20.5	16.7	21.4	15.8	23.7
Qwen2-VL 7B	5.8	13.2	14.2	15.8	13.3	11.7	10.0	15.0	12.5	12.5	13.4	13.3	13.6
Maya	20.0	20.1	20.0	25.8	19.2	20.8	15.0	25.8	17.9	23.3	21.4	15.8	16.1
Phi 3.5 Vision	45.8	31.5	27.5	29.2	23.3	36.7	30.0	31.7	33.9	29.2	37.5	35.8	31.4
Pixtral 12B	9.2	12.4	17.5	13.3	10.0	16.7	10.0	16.7	3.6	14.2	8.9	12.5	13.6
Pangea	0.0	6.7	0.0	0.8	0.0	20.8	24.2	15.8	6.2	0.8	3.6	0.8	0.8
MiniCPM 2.6	9.2	14.6	11.7	19.2	12.5	10.8	10.0	22.5	10.7	12.5	19.6	11.7	19.5

Table 35: VLOD

	en	avg.	id	sw	ta	tr	zh
Centurio Aya	85.0	77.9	79.5	70.9	73.4	83.4	82.4
Centurio Qwen	89.6	81.7	85.0	76.8	76.0	84.2	86.7
Parrot	63.5	55.1	56.6	51.2	50.7	58.6	58.2
PALO 13B	63.8	33.1	58.7	50.9	2.6	53.1	0.2
PALO 7B	62.7	24.1	33.6	47.8	0.4	38.5	0.0
InternVL 2.5 4B	74.9	59.0	65.7	50.7	50.9	64.2	63.5
InternVL 2.5 8B	83.0	63.3	63.2	51.4	54.6	67.2	79.9
Qwen2-VL 2B	67.9	55.9	60.9	51.8	52.2	59.0	55.8
Qwen2-VL 7B	69.8	60.2	61.1	53.1	60.9	65.3	60.7
Maya	60.3	56.3	60.3	50.7	50.6	58.9	61.2
Phi 3.5 Vision	73.4	46.4	56.4	51.3	50.8	58.0	15.7
Pixtral 12B	67.7	60.7	62.5	54.4	61.8	65.5	59.1
Pangea	75.8	70.5	74.3	70.9	66.6	71.1	69.6
MiniCPM 2.6	70.2	57.9	57.8	54.2	57.2	63.3	57.2

Table 36: MaRVL

	en	avg.	fr	hi	he	ro	th	zh
Centurio Aya	55.7	49.3	45.1	58.7	62.9	51.1	46.7	31.6
Centurio Qwen	60.1	47.7	47.1	45.1	56.8	47.7	57.0	32.2
Parrot	28.2	3.6	2.7	2.9	1.4	1.2	3.0	10.7
PALO 13B	51.7	33.1	42.0	17.5	53.4	34.2	20.9	30.6
PALO 7B	54.0	22.5	39.9	9.2	30.6	16.8	12.3	26.4
InternVL 2.5 4B	46.0	42.5	45.7	37.1	38.8	31.5	51.0	50.8
InternVL 2.5 8B	45.6	38.2	51.2	27.9	24.5	35.7	36.4	53.4
Qwen2-VL 2B	53.7	26.5	40.3	10.8	9.5	15.6	38.1	44.6
Qwen2-VL 7B	54.7	31.2	38.6	18.7	13.9	37.2	42.1	36.8
Maya	55.4	17.3	19.1	13.0	21.1	18.0	11.6	20.8
Llama-Vision	0.0	4.7	0.0	0.6	2.4	0.3	24.8	0.0
Phi 3.5 Vision	43.6	17.9	23.5	12.1	16.3	7.8	20.9	27.0
Pixtral 12B	59.4	43.4	46.8	31.7	54.4	44.1	44.4	39.1
Pangea	61.4	55.0	47.4	61.0	53.7	52.9	67.2	47.9
MiniCPM 2.6	53.4	22.3	14.3	12.1	5.1	19.5	53.6	29.3

Table 37: MaXM

	avg.	ar	de	fr	it	ja	ko	ru	th	vi
Centurio Aya	11.1	6.7	19.9	22.5	16.7	5.0	9.0	5.2	5.2	9.7
Centurio Qwen	11.9	4.6	22.7	26.5	18.6	5.9	9.9	5.0	5.2	8.9
Parrot	2.0	1.4	1.9	0.9	1.6	1.6	2.7	2.0	5.2	0.9
PALO 13B	6.3	2.6	15.6	12.1	10.4	4.0	4.3	4.0	0.0	4.2
PALO 7B	5.8	1.8	14.3	13.3	8.3	3.4	3.2	3.6	0.4	4.1
InternVL 2.5 4B	25.1	11.2	34.4	38.4	33.5	18.4	29.0	9.8	16.5	34.6
InternVL 2.5 8B	25.0	11.5	33.8	37.4	35.3	19.7	30.3	10.4	16.5	30.4
Qwen2-VL 2B	19.0	6.1	26.8	30.9	30.7	13.5	21.1	9.3	10.0	22.4
Qwen2-VL 7B	23.2	16.9	27.3	31.7	35.2	16.1	24.6	10.8	15.6	30.7
Maya	5.3	2.8	13.1	12.2	6.6	2.8	4.8	2.9	0.4	2.3
Llama-Vision	15.2	7.4	24.0	18.7	25.3	9.4	14.5	6.1	15.2	15.8
Phi 3.5 Vision	11.1	3.3	18.2	20.2	25.2	5.6	8.8	5.4	3.0	10.5
Pixtral 12B	14.1	4.3	25.7	27.3	25.2	5.9	9.1	7.5	5.2	16.6
Pangea	19.3	8.3	29.5	35.2	29.2	9.3	14.5	7.4	10.8	29.2
MiniCPM 2.6	16.1	2.3	23.9	27.5	32.7	11.7	12.7	7.3	10.0	16.5

Table 38: MTVQA

	en	avg.	bn	de	id	ko	pt	ru	zh
Parrot	37.7	21.2	20.2	23.2	19.8	22.8	21.7	19.7	21.2
PALO 13B	58.0	27.8	26.3	14.7	29.6	30.9	17.8	30.9	44.1
PALO 7B	59.1	36.6	42.8	34.5	30.0	40.8	27.7	32.2	47.9
InternVL 2.5 4B	63.6	28.0	28.1	29.2	15.4	38.3	27.2	31.5	25.9
InternVL 2.5 8B	63.4	32.0	17.4	23.8	25.0	38.2	27.6	36.4	55.2
Qwen2-VL 2B	60.5	38.2	18.6	43.2	32.6	39.0	39.9	44.1	50.3
Qwen2-VL 7B	62.5	49.3	37.4	51.1	48.4	50.3	51.8	52.1	54.1
Maya	58.2	49.1	40.1	53.2	49.7	47.2	52.5	50.6	50.1
Llama-Vision	39.3	27.6	26.0	29.2	26.8	24.9	27.9	30.7	27.9
Phi 3.5 Vision	65.2	38.0	5.0	51.9	37.3	35.6	50.6	45.9	39.5
Pixtral 12B	59.9	3.8	0.7	5.4	14.0	0.3	3.6	0.4	1.9
Pangea	64.6	60.4	59.1	61.6	60.7	58.8	62.1	60.7	59.6
MiniCPM 2.6	57.9	45.7	33.9	49.0	46.3	42.1	51.0	48.7	48.6

Table 39: xGQA

	en	avg.	ar	bn	cs	da	de	el	es	fa	fi	fil	fr	he	hi	hr	hu	id
Centurio Aya	78.4	39.2	40.4	18.5	33.9	40.0	38.6	35.3	69.7	55.8	11.0	34.0	71.3	47.1	26.3	24.9	19.6	58.3
Centurio Qwen	79.1	34.4	36.6	17.1	29.7	43.1	32.0	19.2	69.2	31.2	12.0	33.6	67.6	27.6	20.3	22.0	18.7	50.4
Parrot	5.6	0.4	0.6	0.0	0.2	0.0	0.0	0.0	3.3	2.3	0.0	0.0	0.3	0.0	0.0	0.0	0.0	0.0
PALO 13B	67.3	17.0	23.5	22.9	7.9	30.3	32.4	0.2	57.0	1.5	6.6	8.4	66.2	0.6	25.0	9.9	2.7	22.7
PALO 7B	65.9	13.5	17.3	18.8	5.8	18.5	23.3	0.1	48.3	1.5	4.0	2.7	59.1	0.2	21.2	2.8	6.3	20.2
InternVL 2.5 4B	38.9	17.5	12.1	3.7	9.4	13.6	28.0	2.0	39.7	10.1	2.6	6.2	49.2	8.5	5.4	5.6	3.8	39.9
InternVL 2.5 8B	38.3	15.7	7.9	4.0	10.7	19.2	27.8	2.9	35.0	8.7	5.0	10.9	47.0	8.2	8.3	7.9	5.3	24.7
Qwen2-VL 2B	68.8	5.2	0.8	0.0	1.7	7.2	7.0	0.2	5.1	9.0	1.2	2.9	9.4	0.4	0.0	1.4	2.1	8.9
Qwen2-VL 7B	50.3	24.6	17.9	11.5	23.8	32.3	36.1	13.5	38.9	23.6	8.0	8.3	50.6	13.7	6.7	11.6	15.5	45.4
Maya	55.9	14.6	20.6	18.4	11.4	10.6	23.6	10.7	38.2	1.5	0.5	2.1	47.3	18.9	15.0	2.0	0.9	19.4
Llama-Vision	35.9	7.2	0.0	0.0	0.9	15.5	22.4	0.5	14.7	0.0	4.0	13.1	32.1	0.0	0.0	2.9	13.2	2.2
Phi 3.5 Vision	32.3	6.3	2.8	0.0	0.6	10.5	21.3	0.1	21.9	0.1	0.9	2.5	32.5	1.0	0.1	1.5	2.6	4.2
Pixtral 12B	26.5	22.1	18.6	9.6	16.8	24.4	33.2	8.9	36.5	20.5	10.4	15.3	47.8	18.0	6.3	18.7	15.6	44.6
Pangea	70.1	34.6	33.3	30.8	19.4	25.2	39.4	13.0	61.4	25.4	4.2	6.7	69.7	42.7	21.5	9.5	3.6	70.9
MiniCPM 2.6	87.5	14.2	6.7	3.3	8.5	8.7	27.5	1.7	44.0	5.8	3.2	5.0	52.1	1.5	3.0	6.1	5.8	24.6

	it	ja	ko	mi	nl	no	pl	pt	quz	ro	ru	sv	sw	te	th	tr	uk	vi	zh
Centurio Aya	60.4	49.1	21.3	33.7	61.7	42.5	37.9	59.3	1.7	34.6	38.0	45.9	29.9	15.1	26.0	30.6	30.6	72.7	56.9
Centurio Qwen	53.7	43.5	13.4	34.9	56.2	41.4	30.0	59.9	2.1	23.4	39.2	42.7	30.2	13.5	42.3	23.3	20.3	69.4	33.8
Parrot	0.2	0.0	0.0	1.6	0.0	0.0	4.0	0.2	0.4	0.4	0.0	0.0	0.8	0.0	1.0	0.0	0.0	0.0	
PALO 13B	40.4	19.7	0.2	0.3	36.5	31.0	9.1	13.8	0.8	14.5	21.3	33.9	0.8	0.0	0.5	0.6	2.6	15.6	37.0
PALO 7B	31.0	29.8	2.4	0.3	29.8	16.5	8.4	8.9	0.5	2.6	19.7	23.3	0.0	0.0	0.5	0.1	17.4	29.7	
InternVL 2.5 4B	33.1	33.1	8.9	0.8	29.3	14.2	12.9	39.0	0.2	9.9	23.4	17.1	0.6	1.1	27.9	7.8	7.1	61.3	44.1
InternVL 2.5 8B	27.5	22.0	6.7	0.9	26.8	16.6	12.0	35.0	0.8	12.0	22.6	20.5	1.0	2.6	7.2	9.3	4.7	46.2	40.1
Qwen2-VL 2B	5.9	8.4	1.0	0.3	2.9	5.2	1.5	21.0	1.0	3.7	1.1	13.0	1.1	0.0	1.3	0.9	0.6	7.9	49.5
Qwen2-VL 7B	38.7	32.0	9.1	0.9	39.1	35.7	30.1	48.8	0.9	19.0	37.9	43.1	2.4	3.9	31.2	15.8	16.6	55.6	41.8
Maya	34.4	26.3	8.9	0.3	28.8	9.4	15.8	16.4	0.6	22.0	19.9	11.4	0.5	0.0	0.2	13.5	1.5	31.8	26.9
Llama-Vision	33.5	0.2	0.1	0.8	30.1	2.8	2.4	15.7	0.2	23.4	0.3	11.2	6.8	0.0	1.2	0.6	0.1	0.8	0.0
Phi 3.5 Vision	23.6	8.0	0.3	0.2	19.8	10.7	1.7	25.8	0.4	3.0	0.5	10.2	0.5	0.0	1.0	1.7	0.1	2.6	8.1
Pixtral 12B	32.8	21.8	12.0	5.9	29.7	26.0	19.6	42.4	1.0	20.2	33.8	30.0	10.4	6.2	23.8	14.9	18.4	51.7	28.1
Pangea	53.5	63.3	20.3	0.3	44.9	48.5	24.1	64.6	1.7	38.7	47.3	20.1	40.7	21.8	61.4	30.2	20.7	81.3	50.7
MiniCPM 2.6	24.6	18.6	4.4	2.2	27.8	12.0	36.0	0.2	10.0	20.0	17.0	1.5	0.5	20.9	8.0	7.5	25.8	39.4	

Table 40: XM3600

	en	avg.	ar	bn	cs	da	de	el	es	fa	fi	fil	fr	he	hi	hr	hu	id
Centurio Aya	100.0	95.7	93.6	100.0	97.7	96.7	100.0	100.0	99.8	100.0	99.8	100.0	100.0	99.8	99.6	84.6	99.8	99.2
Centurio Qwen	99.8	95.2	95.1	100.0	98.6	93.9	100.0	100.0	99.4	100.0	100.0	100.0	100.0	98.8	99.0	80.9	100.0	96.7
Parrot	100.0	25.0	100.0	0.0	0.0	0.0	0.0	0.0	100.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
PALO 13B	100.0	60.1	98.6	93.9	47.1	87.5	100.0	60.7	99.6	0.0	74.0	71.5	99.8	35.4	98.2	70.1	9.0	66.4
PALO 7B	100.0	72.0	99.6	98.8	47.5	93.4	100.0	58.2	99.8	0.0	91.8	52.7	100.0	30.7	98.8	27.0	90.8	96.9
InternVL 2.5 4B	100.0	91.0	96.7	93.9	97.1	82.8	100.0	99.0	99.8	98.8	96.1	95.3	100.0	96.7	88.9	91.4	96.1	96.9
InternVL 2.5 8B	100.0	91.1	99.4	95.3	97.7	82.8	100.0	100.0	99.4	97.9	98.2	96.3	100.0	98.4	95.1	83.2	98.2	96.7
Qwen2-VL 2B	100.0	13.2	8.2	0.0	0.0	9.6	12.9	0.2	5.9	58.4	0.2	10.9	10.0	4.5	0.0	3.1	0.2	12.7
Qwen2-VL 7B	100.0	90.0	96.5	98.2	93.9	86.1	99.8	99.4	99.2	99.2	95.7	96.3	98.2	98.2	60.2	79.1	75.4	86.9
Maya	100.0	65.7	99.0	96.1	67.6	85.5	98.6	92.0	99.8	0.2	12.1	1.0	100.0	77.0	98.4	20.7	60.7	40.6
Llama-Vision	100.0	33.3	0.0	0.0	4.9	68.8	95.5	7.0	52.7	0.0	35.0	80.7	88.3	0.0	0.0	17.0	91.0	1.6
Phi 3.5 Vision	100.0	40.8	58.4	0.6	1.4	85.4	99.2	16.2	99.4	0.0	15.2	30.1	99.8	14.8	4.7	25.2	56.4	31.6
Pixtral 12B	100.0	96.8	99.8	99.6	98.8	95.9	100.0	99.4	100.0	99.8	99.8	100.0	99.8	100.0	93.4	100.0	99.6	
Pangea	99.8	87.9	98.8	99.0	97.9	19.1	99.6	99.8	99.2	98.4	91.6	68.9	100.0	100.0	98.2	67.8	93.6	97.9
MiniCPM 2.6	99.8	92.3	94.7	96.5	95.5	96.3	100.0	99.8	99.8	99.0	98.4	97.9	100.0	62.9	92.6	77.3	94.5	93.6

	it	ja	ko	mi	nl	no	pl	pt	quz	ro	ru	sv	sw	te	th	tr	uk	vi	zh
Centurio Aya	99.6	98.8	100.0	98.8	100.0	97.3	99.8	100.0	1.8	100.0	99.6	98.8	90.6	100.0	99.6	100.0	99.8	100.0	89.6
Centurio Qwen	99.8	98.8	100.0	100.0	100.0	95.7	99.6	99.4	3.7	100.0	99.4	98.2	86.5	100.0	99.8	99.6	99.2	100.0	86.5
Parrot	0.0	0.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0	100.0	0.0	0.0
PALO 13B	99.2	61.5	9.8	0.0	99.8	92.2	41.2	27.5	0.0	95.5	68.2	94.9	1.6	26.4	68.0	9.4	11.3	54.7	88.9
PALO 7B	99.4	99.2	91.6	0.0	99.4	95.1	95.5	27.0	0.0	69.1	100.0	96.9	0.0	56.8	91.6	84.4	0.0	99.6	99.8
InternVL 2.5 4B	99.6	100.0	99.2	48.2	99.6	83.0	99.6	100.0	7.0	98.8	99.2	97.7	34.6	90.8	98.2	97.7	95.7	96.1	99.8
InternVL 2.5 8B	100.0	99.8	99.2	66.8	99.2	86.5	99.6	99.8	1.2	99.8	99.8	98.2	54.7	99.4	99.2	98.4	40.6	100.0	99.8
Qwen2-VL 2B	3.9	19.3	18.0	0.2	0.0	5.3	0.0	34.0	0.0	15.4	0.0	23.6	4.5	0.0	1.0	1.0	0.8	12.7	98.8
Qwen2-VL 7B	99.0	98.8	99.0	64.5	99.2	94.1	95.7	95.3	0.2	98.2	99.4	97.9	72.1	95.1	98.8	89.8	83.2	98.2	99.0
Maya	99.6	99.8	91.4	0.0	99.8	80.1	92.0	43.9	0.0	95.7	100.0	96.7	1.6	0.0	47.9	96.3	7.2	62.5	99.8
Llama-Vision	94.7	0.0	0.0	93.0	99.6	9.0	9.2	48.2	0.8	92.6	0.2	33.8	73.6	0.0	2.3	0.2	0.0	0.2	0.0
Phi 3.5 Vision	99.0	58.6	9.0	1.0	93.6	89.8	27.3	63.9	0.0	56.8	0.0	85.0	2.9	13.7	3.5	38.1	0.0	51.8	37.9
Pixtral 12B	100.0	100.0	100.0	99.6	100.0	95.5	100.0	100.0	9.4	99.6	100.0	95.9	99.4	100.0	99.8	100.0	100.0	99.8	
Pangea	100.0	99.4	100.0	0.8	99.6	95.7	99.6	100.0	0.0	100.0	99.6	67.0	82.4	99.8	100.0	99.8	91.0	99.8	99.0
MiniCPM 2.6	100.0	98.4	99.2	99.2	99.6	95.5	96.5	99.8	10.0	98.2	97.3	99.4	66.4	85.5	96.3	95.1	99.2	90.6	97.1

Table 41: XM3600 Language Fidelity

	en	avg.	ar	es	fr	ru
Centurio Aya	65.0	62.4	61.7	61.0	64.3	62.7
Centurio Qwen	75.4	70.2	68.8	70.9	70.5	70.8
Parrot	28.7	31.4	34.0	24.3	30.0	37.4
PALO 13B	56.6	53.6	51.8	52.7	54.9	55.0
PALO 7B	58.0	53.4	52.5	52.3	53.7	55.1
InternVL 2.5 4B	69.0	58.7	55.7	58.8	61.4	59.0
InternVL 2.5 8B	73.5	66.4	61.8	68.0	68.4	67.3
Qwen2-VL 2B	61.9	56.2	52.9	55.3	58.6	57.9
Qwen2-VL 7B	62.1	59.6	59.2	58.9	60.0	60.3
Maya	50.1	43.9	45.3	42.7	45.8	41.8
Phi 3.5 Vision	58.9	53.3	49.7	52.7	56.4	54.3
Pixtral 12B	60.9	52.7	36.0	57.9	59.0	58.1
Pangea	69.0	65.2	64.5	64.3	66.3	65.7
MiniCPM 2.6	71.9	65.4	61.1	67.5	67.0	66.1

Table 42: XVNLI

	en	avg.	ar	fr	hi	id	ja	pt
Centurio Aya	37.6	37.2	36.2	38.9	38.8	39.7	34.2	35.4
Centurio Qwen	46.4	43.0	39.6	45.0	41.6	44.1	43.5	44.1
Parrot	35.3	32.4	31.9	34.9	26.1	31.3	34.9	35.4
PALO 13B	32.4	28.9	24.2	34.9	24.2	31.6	26.4	32.3
PALO 7B	31.8	30.9	28.2	33.6	27.3	30.6	32.3	33.3
InternVL 2.5 4B	49.2	42.7	41.6	45.6	33.7	43.4	44.2	47.8
InternVL 2.5 8B	50.7	45.2	40.3	48.7	41.2	43.1	47.6	50.2
Qwen2-VL 2B	36.8	35.5	31.5	41.3	30.2	36.7	36.1	37.0
Qwen2-VL 7B	43.0	40.7	36.9	42.6	38.5	41.1	41.3	43.8
Maya	37.9	33.3	32.6	36.6	31.3	31.6	32.0	36.0
Phi 3.5 Vision	41.7	37.4	34.9	44.3	29.2	37.7	35.7	42.4
Pixtral 12B	30.3	26.2	19.1	28.5	19.2	27.3	28.6	34.7
Pangea	43.1	42.0	37.6	43.0	38.5	46.8	41.6	44.8
MiniCPM 2.6	39.1	36.5	30.5	38.9	33.7	37.7	37.2	40.7

Table 43: xMMMU

	avg.	amh-ethiopia	arz-egypt	ben-india	bre-france	bul-bulgaria	fil-philippines	gle-ireland	hin-india	ibo-nigeria
Centurio Aya	49.4	32.1	<u>52.7</u>	45.8	30.4	<u>50.1</u>	48.8	41.7	67.2	31.0
Centurio Qwen	<u>52.9</u>	38.0	<u>52.7</u>	<u>54.9</u>	30.6	<u>49.6</u>	<u>51.7</u>	<u>48.5</u>	65.7	30.5
Parrot	41.1	31.6	35.5	33.9	31.1	38.8	45.3	45.1	41.8	35.5
PALO 13B	39.6	26.1	31.5	33.9	31.6	35.3	45.3	41.1	40.3	32.5
PALO 7B	37.1	20.5	25.1	30.1	27.7	32.6	43.3	37.1	43.3	31.0
InternVL 2.5 4B	48.1	36.3	38.9	42.7	<u>33.1</u>	42.0	47.8	45.7	51.7	29.0
InternVL 2.5 8B	48.6	<u>29.5</u>	41.4	42.3	29.4	47.4	46.8	47.5	50.2	33.5
Qwen2-VL 2B	33.6	27.4	31.0	33.6	25.9	32.9	32.0	31.3	34.8	35.5
Qwen2-VL 7B	37.6	31.2	35.5	31.5	31.1	35.6	40.9	37.1	39.3	31.0
Maya	39.8	30.3	41.9	38.8	30.6	36.7	35.0	34.4	46.8	31.0
Llama-Vision	38.8	32.1	5.4	60.1	13.6	22.4	43.8	35.9	46.3	28.5
Phi 3.5 Vision	40.9	28.6	38.4	28.7	28.9	33.7	45.3	40.2	42.8	<u>38.0</u>
Pixtral 12B	33.5	22.6	27.1	21.7	24.9	30.5	35.5	38.7	41.3	26.5
Pangea	55.2	35.5	49.3	53.5	<u>33.1</u>	52.3	56.7	53.1	<u>66.2</u>	40.0
MiniCPM 2.6	34.1	26.9	31.5	27.6	25.9	32.6	31.5	37.1	32.3	36.0
	ind-indonesia	jav-indonesia	jpn-japan	kin-rwanda	kor-south korea	mar-india	min-indonesia	mon-mongolia	msa-malaysia	nor-norway
Centurio Aya	53.6	41.1	<u>44.8</u>	32.8	61.7	<u>56.9</u>	42.6	29.2	52.7	55.5
Centurio Qwen	<u>54.9</u>	<u>46.1</u>	<u>44.8</u>	42.1	<u>66.2</u>	55.9	41.8	33.3	<u>55.9</u>	<u>58.2</u>
Parrot	43.4	37.7	34.5	32.8	47.6	32.2	36.3	34.3	42.9	49.5
PALO 13B	41.0	35.4	33.0	28.9	42.8	37.6	37.8	26.3	44.1	54.5
PALO 7B	37.9	33.7	29.1	29.4	42.4	33.7	31.1	25.6	36.8	47.5
InternVL 2.5 4B	54.6	44.4	39.4	34.9	65.9	48.0	41.0	27.6	53.0	52.8
InternVL 2.5 8B	<u>54.9</u>	44.8	41.4	32.8	56.9	<u>43.6</u>	<u>44.6</u>	33.0	54.3	55.2
Qwen2-VL 2B	36.9	31.6	25.1	31.5	37.6	24.3	27.9	31.1	32.1	40.5
Qwen2-VL 7B	40.8	32.3	36.0	30.2	43.4	31.2	33.5	34.0	43.5	42.8
Maya	36.2	34.7	29.1	31.5	50.0	42.6	33.9	31.1	44.4	47.5
Llama-Vision	42.0	34.0	25.1	18.3	45.9	38.1	34.3	27.9	40.6	48.5
Phi 3.5 Vision	40.8	35.4	36.9	<u>36.2</u>	44.1	33.7	39.0	<u>35.3</u>	41.3	48.2
Pixtral 12B	36.9	32.3	27.6	25.1	39.7	24.3	28.3	24.7	36.8	32.1
Pangea	60.0	50.5	42.9	33.6	68.3	57.9	48.2	40.7	60.3	<u>58.2</u>
MiniCPM 2.6	36.2	31.6	33.5	30.6	31.0	32.2	29.9	29.2	34.6	40.5
	orm-ethiopia	por-brazil	ron-romania	rus-russia	sin-sri lanka	spa-argentina	spa-chile	spa-colombia	spa-ecuador	spa-mexico
Centurio Aya	36.4	65.1	<u>61.6</u>	65.5	28.9	60.0	58.5	56.0	55.2	52.6
Centurio Qwen	35.0	70.8	57.0	<u>67.5</u>	<u>42.7</u>	<u>63.8</u>	<u>66.2</u>	<u>59.8</u>	58.8	61.9
Parrot	34.6	60.9	44.0	45.0	28.0	47.9	52.1	48.5	43.9	44.3
PALO 13B	29.4	53.9	52.6	44.0	24.9	47.5	50.0	47.3	47.8	44.9
PALO 7B	51.1	49.3	45.5	24.0	47.5	49.1	45.2	45.9	42.4	57.5
InternVL 2.5 4B	34.6	66.5	49.7	65.5	34.7	60.4	59.8	54.4	56.6	53.9
InternVL 2.5 8B	32.2	64.4	60.3	62.5	29.8	60.4	65.4	56.4	59.7	57.0
Qwen2-VL 2B	33.2	39.8	32.5	33.0	24.9	40.0	40.2	40.7	39.0	35.0
Qwen2-VL 7B	37.4	47.9	34.8	47.5	30.7	44.5	47.9	40.7	42.3	40.2
Maya	29.9	53.5	51.3	42.0	30.2	44.9	47.4	45.2	45.6	39.3
Llama-Vision	23.4	38.7	47.7	52.0	48.4	47.9	55.1	51.0	48.1	48.3
Phi 3.5 Vision	33.6	62.0	43.7	45.0	29.3	55.1	59.0	51.9	52.8	48.0
Pixtral 12B	21.5	41.2	28.8	40.5	23.6	49.1	54.3	43.6	48.3	40.6
Pangea	<u>36.4</u>	<u>69.7</u>	62.9	73.5	36.0	<u>63.8</u>	67.1	61.4	62.4	60.7
MiniCPM 2.6	36.0	45.1	33.4	37.0	26.7	37.7	44.4	40.7	38.7	35.9
	spa-spain	spa-uruguay	sun-indonesia	swa-kenya	tam-india	tel-india	urd-india	urd-pakistan	zho-china	zho-singapore
Centurio Aya	68.2	40.3	42.0	50.2	<u>49.5</u>	37.0	47.3	50.5	64.6	65.6
Centurio Qwen	70.4	42.5	41.5	<u>56.4</u>	43.0	50.5	<u>52.7</u>	<u>56.9</u>	71.1	73.1
Parrot	65.7	36.2	31.5	40.7	36.9	29.5	31.8	32.9	64.3	55.2
PALO 13B	63.2	39.4	37.5	39.6	31.3	29.5	35.9	40.3	46.6	39.6
PALO 7B	57.5	36.2	31.5	31.1	29.0	28.5	34.1	40.3	46.3	42.0
InternVL 2.5 4B	68.2	<u>44.8</u>	<u>44.0</u>	45.8	35.5	41.0	47.7	41.7	<u>74.6</u>	66.5
InternVL 2.5 8B	72.3	<u>44.8</u>	41.5	50.9	35.5	39.5	44.1	39.8	78.5	72.6
Qwen2-VL 2B	42.1	34.6	33.5	37.7	25.2	27.0	30.5	31.9	44.1	41.5
Qwen2-VL 7B	47.8	41.6	31.5	34.1	26.6	26.5	37.3	31.0	51.4	43.4
Maya	55.7	34.3	33.5	38.8	32.2	29.0	43.2	48.1	50.5	50.9
Llama-Vision	70.1	37.1	29.5	52.0	60.7	62.5	24.5	15.3	36.3	23.1
Phi 3.5 Vision	64.2	45.1	32.0	46.5	29.4	32.5	29.5	26.9	51.1	40.6
Pixtral 12B	48.4	40.3	27.0	42.1	22.9	19.0	27.7	23.6	47.3	39.6
Pangea	73.0	<u>44.8</u>	49.0	65.6	46.7	<u>55.0</u>	57.7	65.7	71.7	68.4
MiniCPM 2.6	42.5	36.2	29.0	34.1	24.3	28.5	32.7	22.7	48.2	46.2

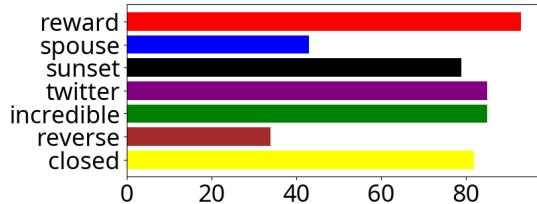
Table 44: CVQA

	en	avg.	avg.	Latin	avg.	other	ar	de	hi	id	it	ko	ru	th	zh	zu
Centurio Aya	83.1	74.2		80.9		69.7	75.9	82.1	80.1	81.4	80.6	68.8	73.5	66.5	53.4	79.5
Centurio Qwen	84.8	76.1		82.7		71.8	<u>76.9</u>	83.5	82.4	83.8	83.1	72.4	75.6	64.4	58.9	80.2
Parrot	51.0	49.9		50.5		49.5	<u>50.4</u>	51.6	49.6	51.0	49.8	50.4	50.5	48.2	47.8	49.5
PALO 13B	54.0	51.5		52.7		50.7	50.9	53.2	51.2	52.5	52.8	51.0	49.5	51.0	50.7	52.1
PALO 7B	55.5	52.8		55.4		51.0	50.4	56.9	51.0	55.0	54.1	51.6	51.1	51.4	50.2	55.8
InternVL 2.5 4B	87.0	78.3		86.9		72.6	54.9	87.6	59.8	87.0	88.2	89.4	86.4	55.1	90.4	84.8
InternVL 2.5 8B	91.0	79.2		88.7		72.8	55.8	89.8	54.9	89.1	89.1	<u>92.5</u>	86.9	53.1	93.6	86.9
Qwen2-VL 2B	85.0	83.5		83.4		<u>83.5</u>	70.6	84.4	86.5	84.1	83.5	88.1	78.8	<u>86.4</u>	90.4	81.8
Qwen2-VL 7B	<u>91.2</u>	90.9		90.1		91.4	83.4	90.5	94.8	<u>91.0</u>	90.8	93.8	<u>87.5</u>	94.1	94.9	88.2
Maya	51.4	50.9		51.6		50.4	50.4	53.4	50.1	51.5	50.0	49.9	<u>49.5</u>	51.1	51.6	51.6
Llama-Vision	91.1	<u>84.8</u>		89.9		81.5	63.2	90.1	<u>91.1</u>	89.5	91.9	87.4	83.0	84.8	79.5	88.0
Phi 3.5 Vision	92.2	79.4		<u>90.2</u>		72.2	53.1	91.9	83.8	89.2	<u>90.9</u>	87.1	77.9	86.6	55.5	76.5
Pixtral 12B	91.1	71.0		90.5		58.0	50.4	<u>91.5</u>	53.6	91.1	90.9	49.5	88.2	52.9	53.4	<u>88.4</u>
Pangea	87.2	72.2		85.7		63.1	51.5	86.6	69.4	86.2	87.1	71.4	79.2	54.4	52.9	82.9
MiniCPM 2.6	89.0	74.3		88.0		65.2	52.0	89.0	53.1	87.9	89.0	54.8	84.0	53.1	<u>94.5</u>	86.0

Table 45: SMPQA Ground

	en	avg.	avg. Latin	avg. other	ar	de	hi	id	it	ko	ru	th	zh	zu
Centurio Aya	60.0	30.1	49.8	17.0	29.2	50.2	17.6	52.6	51.2	11.2	38.2	4.8	0.8	45.2
Centurio Qwen	65.2	31.7	54.3	16.6	<u>21.4</u>	53.2	21.4	55.4	56.6	16.2	34.8	5.2	0.6	52.2
Parrot	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.4	0.0	0.0	0.0	0.0
PALO 13B	25.6	4.0	9.9	0.1	0.0	12.0	0.0	10.2	12.4	0.4	0.0	0.0	0.0	5.0
PALO 7B	22.4	2.7	6.7	0.1	0.0	8.4	0.0	7.0	7.0	0.4	0.0	0.0	0.0	4.4
InternVL 2.5 4B	77.8	47.5	67.7	34.0	0.0	71.0	0.0	69.8	69.6	<u>69.0</u>	54.4	0.2	80.2	60.4
InternVL 2.5 8B	80.6	<u>48.2</u>	68.1	34.9	0.0	69.2	0.0	70.4	70.8	67.2	61.2	0.2	80.8	62.2
Qwen2-VL 2B	68.8	47.4	60.0	39.0	0.2	61.2	24.8	59.4	61.2	66.0	46.8	24.0	72.0	58.2
Qwen2-VL 7B	<u>85.0</u>	64.9	76.2	57.4	1.8	80.6	58.6	75.8	<u>79.2</u>	77.6	70.6	43.8	92.0	69.2
Maya	14.6	1.8	4.3	0.1	0.0	8.2	0.0	3.6	4.6	0.4	0.0	0.0	0.0	0.8
Llama-Vision	58.4	22.8	46.6	6.9	0.0	55.4	2.4	38.4	37.2	8.4	13.0	6.0	11.8	55.4
Phi 3.5 Vision	84.8	35.9	69.4	13.5	0.2	70.8	12.0	69.4	76.6	15.4	40.4	0.2	12.8	61.0
Pixtral 12B	<u>85.0</u>	35.9	<u>73.3</u>	10.9	0.0	<u>71.8</u>	0.0	<u>75.4</u>	81.6	0.4	<u>64.6</u>	0.4	0.0	<u>64.6</u>
Pangea	<u>72.0</u>	23.8	<u>54.4</u>	3.4	0.0	<u>58.6</u>	0.2	<u>57.2</u>	64.4	0.4	<u>19.2</u>	0.4	0.0	<u>37.4</u>
MiniCPM 2.6	80.8	39.3	67.5	20.6	0.0	67.2	0.0	69.8	71.4	1.0	38.4	0.4	<u>83.6</u>	61.6

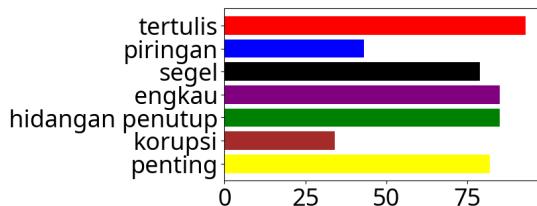
Table 46: SMPQA Name



(a) Example of a bar plot in SMPQA for English.

Questions for Grounding: "Is the bar with label 'reward' the biggest?", "Is the bar with label 'incredible' the biggest?", "Is the bar with label 'reverse' the smallest?", "Is the bar with label 'sunset' the smallest?", "Is the bar with label 'closed' colored in yellow?", "Is the bar with label 'closed' colored in purple?", "Is the bar with label 'twitter' colored in purple?", "Is the bar with label 'twitter' colored in red?"

Questions for Reading: "What is the label of the biggest bar?", "What is the label of the smallest bar?", "What is the label of the yellow bar?", "What is the label of the red bar?", "What is the label of the purple bar?"



(b) The same plot in Indonesian. Note that all questions refer to the same parts of the plot as the English version just with different words for labels.

Questions for Grounding: "Is the bar with label 'tertulis' the biggest?", "Is the bar with label 'hidangan penutup' the biggest?", "Is the bar with label 'korupsi' the smallest?", "Is the bar with label 'segel' the smallest?", "Is the bar with label 'penting' colored in yellow?", "Is the bar with label 'penting' colored in purple?", "Is the bar with label 'engkau' colored in purple?", "Is the bar with label 'engkau' colored in red?"

Questions for Reading: "What is the label of the biggest bar?", "What is the label of the smallest bar?", "What is the label of the yellow bar?", "What is the label of the red bar?", "What is the label of the purple bar?"

Figure 3: Examples of one plot configuration in SMPQA for English and Indonesian.

Jesus und seine Anhänger



Can you explain the joke of this German meme?

Centurio Aya: The humor in this image comes from juxtaposing a biblical figure, Jesus, with an everyday scene that is both ordinary and unexpected. The man's appearance as Jesus, complete with long hair and a beard, contrasts with the mundane setting of a caravan park populated by caravans (trailers). This creates a humorous effect because it takes something familiar (Jesus) and places it into a context that is not only unfamiliar but also somewhat banal or trivial. It plays on the idea of finding divinity in the most unexpected places, which can be seen as a commentary on how we perceive holiness or significance. Additionally, the caption "Jesus und seine Anhänger" further emphasizes the absurdity by referencing followers, implying that even in such an unremarkable location, one could find disciples or supporters, thus adding another layer to the comedic effect.

Centurio Qwen: The joke in this German meme plays on a pun. The phrase "Jesus und seine Anhänger" translates to "Jesus and his followers." However, the word "Anhänger," which means "followers," is also slang for caravans or RVs (recreational vehicles) in some contexts. Therefore, the image humorously suggests that Jesus has modern-day followers who travel together in caravans instead of following him spiritually as he did with his original disciples. This juxtaposition creates a humorous effect by combining religious iconography with contemporary culture.

Figure 4: Centurio explaining a German meme. Both models correctly read the caption but only Qwen ‘understands’ the pun and gets the explanation (mostly) correct.

	AVG.	XM3600			M _T -VQA	SMPQA G.	SMPQA N.	M3Exam	sMAMMU	C-QQA	
		en	mul	fid.	en	mul	en	mul	en	mul	
Parrot	25.8	5.6	0.4	25.0	2.0	51.0	49.9	0.0	0.0	46.6	36.2
PALO 7B	28.7	65.9	13.5	72.0	5.8	55.5	52.8	22.4	2.7	41.0	29.1
PALO 13B	29.9	72.9	17.0	79.0	11.0	60.0	57.8	23.8	2.7	45.2	28.3
Lama Vision 3.2 11B	*32.3	35.9	7.2	33.3	15.2	91.1	84.8	58.4	22.8	42.9	28.9
Maya	33.4	55.9	14.6	65.7	5.3	51.4	49.4	14.6	49.2	36.3	37.9
Pixel 12B	38.1	26.5	22.1	96.8	14.1	91.1	71.0	85.0	35.9	49.4	33.7
Phi 3.5 Vision	39.1	35.9	6.3	50.1	11.0	95.0	79.0	84.8	35.7	56.1	40.3
Qwen-VL 2B	41.2	68.8	5.2	13.2	9.0	65.0	63.5	8.8	14.4	40.5	36.5
MiniCPM 2.6	41.7	87.5	14.2	92.3	16.1	89.0	74.3	80.8	39.3	55.0	48.2
InterVL 2.5 4B	45.3	38.9	17.5	91.0	25.1	87.0	78.3	77.8	47.5	63.2	50.3
InterVL 2.5 SB	47.1	37.0	15.7	79.1	20.0	91.0	87.0	79.0	28.5	58.1	52.2
Qwen-VL 7B	45.7	50.3	25.6	90.0	11.2	99.9	85.0	85.0	56.1	97.7	43.0
Pangue	48.2	70.1	34.6	87.9	19.3	87.2	72.2	72.0	23.9	58.0	45.5
Centurio Aya	48.5	78.4	39.2	95.7	11.1	83.1	74.2	60.0	30.1	53.0	41.2
Centurio Qwen	51.6	79.1	34.4	95.2	11.9	84.8	76.1	65.2	31.7	61.2	46.9
<hr/>											
	MAX	xGQA	BIN-MC	XNLI	MARVL	VGR	VLOD				
	en	mul	en	mul	en	mul	en	mul	en	mul	
Parrot	28.2	3.6	37.7	21.2	30.5	25.7	28.7	31.4	63.5	55.1	52.9
PALO 7B	54.0	22.5	59.1	36.6	58.7	38.6	58.0	53.4	62.7	24.1	48.3
PALO 13B	51.7	33.2	58.0	27.8	61.4	41.7	56.6	53.6	63.8	33.1	63.3
Lama-Vision 3.2 11B	51.0	4.7	57.7	27.6	59.0	40.8	59.8	50.8	62.3	26.2	2.5
Maya	55.4	17.3	58.2	49.1	54.0	43.2	50.1	43.9	60.3	56.3	46.7
Pixel 12B	59.4	43.4	59.9	3.8	71.0	54.2	60.9	52.7	67.7	60.7	55.8
Phi 3.5 Vision	43.6	17.9	65.2	38.0	63.1	36.8	58.9	53.3	73.4	46.4	81.7
Qwen-VL 2B	53.7	22.3	58.2	38.2	73.2	47.4	69.9	54.2	73.7	53.3	54.4
MiniCPM 2.6	53.4	22.3	57.9	45.7	72.6	47.4	71.9	65.4	70.2	57.9	52.5
InterVL 2.5 4B	46.0	42.5	63.6	28.0	68.4	45.4	69.0	58.7	74.9	59.9	72.5
InterVL 2.5 SB	45.6	38.2	63.4	32.0	70.3	44.2	73.4	66.4	83.0	63.3	87.5
Qwen-VL 7B	54.7	45.7	64.9	45.3	74.1	47.1	74.1	61.7	80.3	64.1	52.2
Pangue	61.4	55.0	64.6	60.4	70.3	52.1	69.0	65.2	75.8	70.3	69.2
Centurio Aya	55.7	49.3	59.1	53.2	69.7	54.7	65.0	62.4	85.0	72.9	82.5
Centurio Qwen	60.1	47.7	60.6	54.8	72.7	56.2	75.4	70.2	89.0	81.7	87.5
<hr/>											

Table 6: Comparison of Centurio and 13 other LLMs over 14 tasks. We mark best and second-best results. Scores are accuracy (CIDEr for XM3600). en & mul are the English and averaged multilingual results. XM3600 fid. is the language fidelity over all languages; SMPQA G. & N are Grounding and Naming. *: only single-image input was supported. AVG. is the average over all task. Details on the evaluation setup and models in Appendix B.

Summarize the table for me.

Centurio Aya: This is a detailed summary of the table:

The table compares Centurio and 13 other LLM models on 14 different tasks, using CIDEr as an evaluation metric. It presents accuracy scores (AVG., MTXMUL) and their standard deviations (SD.) across various datasets like SQuAD G & N, Image captioning, etc. Models are listed in descending order based on AVG score, with Paratrol being at the top followed by others like PALO-TB, Lama Vision, etc. There's also information about model names, task descriptions, and additional metrics like BIN_MU, XNLI, MARVL, VGRU, VLOD, and C.Q.A. A note explains that only one image input was used for evaluating language fidelity over all languages, and it mentions "CENTURIO QEN" which might be a specific version or variant of the Centurio model. The bottom part of the table provides further details about the evaluation setup and models' performance on these tasks. Overall, this table seems to evaluate the Centurio model against several others in terms of its ability to perform well on multiple NLP tasks.

Centurio Qwen: Here a summary:

The best model is Centurio Aya.

The worst model is Palo 7B.

The average score of all models is around 50%.

The most accurate model (Centurio Aya) has an accuracy over 63% in all tasks, while the least accurate one (Palo 7B) has an accuracy under 28%.

In multilingual results, the best model is also the most accurate one and vice versa.

The most accurate model has an overall language fidelity over 94%, whereas the least accurate one has an overall language fidelity below 13%.

All models have an input supported AVG. above 40%.

The highest number of errors made by any model was 128, which occurred on the Inten/VL 2.5 SB task.

The lowest number of errors made by any model was 1, which occurred on the MiniCPM 2.6 task.

What model is best? Explain why.

Centurio Aya: Centurio! QEN, with an AVG., MAMM/U 32.0 and C.QAL 41.8. It has the highest average over all tasks compared to other models listed in table 6, indicating its superior performance across multiple linguistic tests.

Centurio Qwen: Centurio Aya Centurio'Queen. It has the highest accuracy score (65.7) and lowest error rate (20.3).

Figure 5: Centurio struggle with images containing small text.



What is written on this bottle?

Centurio Aya: The text on the bottle, in German, reads "Saxia Natürliche Mineralwasser" which translates to "Saxia Natural Mineral Water." Below that, it says "Wort am Rhein," likely referring to a specific source or brand name associated with the Rhine River. The phrase "ohne Kohlensäure" indicates that the water contains no carbonation. There's also some additional information such as size ("0,5L") and possibly an expiration date or batch number at the bottom of the label, but these are not clearly visible in the image provided.

Centurio Qwen: The text on the label of the bottle reads "Saskia," followed by a description in German that translates to "Natural Mineral Water" and "Mineral water from Würtemberg am Rhein." Below, it states "still," indicating the type of mineral water. The bottom part mentions "ohne Kohlenstoffe," which means "without carbonation," confirming that the water is still (non-carbonated). Additionally, there's a logo or emblem above the brand name "Saskia," but due to the angle and focus of the image, specific details about the logo cannot be discerned. There are also some numbers at the very top right corner of the label, likely related to product information such as batch number or size, but they are not clearly visible. The overall design suggests a natural, possibly organic origin for the mineral water. However, without clearer visibility of all elements, more detailed information would require further examination or context. If you're looking for precise nutritional facts or chemical composition, those specifics might not be available from the visual content alone. For accurate information regarding the contents or any health-related aspects of the beverage, consulting the official product packaging or manufacturer's website would be advisable.

Please note: This analysis focuses solely on what can be seen within the provided image and does not include speculative or unverifiable claims beyond what is directly observable. Any additional interpretation goes beyond the scope of the given visual data.

If you need assistance with something else or have another question, feel free to ask!

In summary:

Brand Name: Saskia Type: Natural Mineral Water Origin: Würtemberg am Rhein Carbonation Status: Still (Non-Carbonated) Additional Information: Without Carbon Dioxide (Kohlensäure) Label Design Suggests Organic/Healthy Origin Specific Nutritional Facts or Chemical Composition Not Clearly Visible in Image

For comprehensive details, refer to the actual product packaging or manufacturer's resources.

Figure 6: Centurio summarizing the label of a water bottle. While mostly correct, there are some OCR mistakes, e.g., with the city name ('Wörth')



Opisite sliku što detaljnije. (Describe the image in as much detail as possible.)

Centurio Aya: Slika prikazuje mirno i živopisno prizoru vodopada koji se nalazi u bujnim zelenim krajoliku. Vodopad, s više padina, stvara dinamičnu scenu sa čistim plavom vodi koja kontrastira sa okolnom zelenilom. Okolište je gusto naseljen drvećem, a nekoliko oblaka može se vidjeti na nebu, dodajući dubinu sceni. (*The image depicts a peaceful and picturesque scene of a waterfall set in a lush green landscape. The waterfall, with multiple slopes, creates a dynamic scene with its clear blue water contrasting with the surrounding greenery. The surroundings are densely populated with trees, and a few clouds can be seen in the sky, adding depth to the scene.*)

Centurio Qwen: Slika prikazuje mirnu vodopadsku scenu, karakterizirana svojim više padavina koji se spuštaju u čisto plavo-zeleno jezero ispod njih. Voda izgleda glatko i neometano, sugerirajući dubinu i čistost. Okolišna vegetacija je bogata i živahna, ukazujući na dobro održavan ekosistem sa obiljem svježe vode. Nebo je jasno i plavo, s minimalnim oblaci, što sugerire sunčan dan bez vidljive atmosferske promene ili udaljenog događaja. Sama vodopada nema ljudskih elemenata kao što su mostovi, klupi ili barijeri, što bi moglo značiti da je lokacija priloživa i možda nepovremena za putnike. (*The image depicts a tranquil waterfall scene, characterized by its multiple falls cascading into a clear blue-green lake below. The water appears smooth and undisturbed, suggesting depth and purity. The surrounding vegetation is rich and vibrant, indicating a well-maintained ecosystem with an abundance of fresh water. The sky is clear and blue, with minimal clouds, suggesting a sunny day with no visible atmospheric change or distant event. The waterfall itself has no human elements such as bridges, benches, or barriers, which could indicate that the location is approachable and perhaps uninhabited by travelers.*)

Figure 7: Centurio describing an image in Croatian. Translations by Google Translate are provided in brackets for the reader.



Cine este prezentat în acest portret? (Who is shown in this portrait?)

Centurio Aya: Portretul prezintă un om care pare a fi o figură istorică sau fictivă, probabil de rang înalt datorită îmbrăcăminte sale regale și a bijuteriilor. Stilul picturii sugerează că ar putea fi din perioada Renașterii sau barocă, având în vedere detaliile meticuloase ale costumului și tehnica picturilor cu ulei. Fără context suplimentar, nu este posibil să se identifice persoana exactă reprezentată, dar hainele și accesorile sunt caracteristice celor care au detinut poziții semnificative în istoria europeană. (*The portrait depicts a man who appears to be a historical or fictional figure, likely of high rank due to his regal attire and jewelry. The style of the painting suggests that it may be from the Renaissance or Baroque period, given the meticulous detail of the costume and the oil painting technique. Without further context, it is not possible to identify the exact person depicted, but the clothing and accessories are characteristic of those who held significant positions in European history.*)

Centurio Qwen: Persoana din portret pare a fi un om de natură, probabil o figură istorică sau nobilă, având în vedere îmbrăcăminta și simbolismul capului. În mod specific, el poartă o coroană cu modele complicate, care ar putea indica putere sau statut. Stila îmbrăcăminte sugerează că ar putea fi din secolul al XVIII-lea, ceea ce indică faptul că figura poate fi asociată cu această epocă. Fără context suplimentar, nu este posibil să se determine identitatea exactă, dar caracteristicile sale distincte sugerează că ar putea fi un lider politic, un regiune sau un alt individ important dintr-o perioadă istorică. Portretul servește pentru a transmite autoritatea și statusul sa prin utilizarea elementelor vizuale ale îmbrăcăminte și accesorilor. Este important de menționat că, deși detaliul exact al persoanei este neclar, natura portretului sugerează că ea are un rol semnificativ în istoria respectivă. Portretul servește pentru a transmite autoritatea și statusul sa prin utilizarea elementelor vizuale ale îmbrăcăminte și accesorilor. (*The person in the portrait appears to be a man of character, perhaps a historical or noble figure, given the clothing and symbolism of the head. Specifically, he is wearing an intricately patterned crown, which could indicate power or status. The style of the clothing suggests that it may be from the 18th century, indicating that the figure can be associated with this era. Without additional context, it is not possible to determine the exact identity, but his distinctive features suggest that he may be a political leader, a regional leader, or another important individual from a historical period. The portrait serves to convey his authority and status through the use of visual elements of his clothing and accessories. It is important to note that while the exact detail of the person is unclear, the nature of the portrait suggests that he has a significant role in the history in question. The portrait serves to convey his authority and status through the use of visual elements of his clothing and accessories.*)

Figure 8: Centurio answering a question in Romanian at length. Still, neither model correctly identifies the famous portrait of Vlad III Dracula and both models are wrong with their guess of creation time (16th century). Translations by Google Translate are provided in brackets for the reader.