

Annotation-Efficient Vision-Language Model Adaptation to the Polish Language Using the LLaVA Framework

Grzegorz Statkiewicz, Alicja Dobrzeńska, Karolina Seweryn,
Aleksandra Krasnodębska, Karolina Piosek, Katarzyna Bogusz,
Sebastian Cygert, Wojciech Kusa

NASK National Research Institute, Warsaw, Poland

Correspondence: {firstname.lastname}@nask.pl

Abstract

Most vision-language models (VLMs) are trained on English-centric data, limiting their performance in other languages and cultural contexts. This restricts their usability for non-English-speaking users and hinders the development of multimodal systems that reflect diverse linguistic and cultural realities. In this work, we reproduce and adapt the LLaVA-Next methodology to create a set of Polish VLMs. We rely on a fully automated pipeline for translating and filtering existing multimodal datasets, and complement this with synthetic Polish data for OCR and culturally specific tasks. Despite relying almost entirely on automatic translation and minimal manual intervention to the training data, our approach yields strong results: we observe a +9.5% improvement over LLaVA-1.6-Vicuna-13B on a Polish-adapted MMBench, along with higher-quality captions in generative evaluations, as measured by human annotators in terms of linguistic correctness. These findings highlight that large-scale automated translation, combined with lightweight filtering, can effectively bootstrap high-quality multimodal models for low-resource languages. Some challenges remain, particularly in cultural coverage and evaluation. To facilitate further research, we make our models and evaluation dataset publicly available.

1 Introduction

Recent advances in artificial intelligence have led to remarkable progress in multimodal large language models (LLM), especially vision-language models (VLMs), which integrate vision and language understanding to perform tasks such as visual question answering, image captioning, and reasoning (Achiam et al., 2023; Liu et al., 2023, 2024b). These models leverage massive datasets and sophisticated architectures to achieve state-of-the-art performance across a wide range of benchmarks. However, the current VLM landscape remains predominantly English-centric, primarily

due to the composition of standard training datasets, which limits effectiveness in other languages and cultural contexts (Tong et al., 2024; Laurençon et al., 2024; Wiedmann et al., 2025).

To address this challenge, we explore whether large-scale automated translation can serve as a practical alternative for developing multimodal models in low- and mid-resource languages. Specifically, we focus on Polish as a case study and investigate how far we can go using automatically translated data with minimal manual intervention in the training data. We choose Polish due to the availability of recent competitive Polish LLMs (Kocoń et al., 2025; Ociepa et al., 2025) and its rich morphological complexity, which poses challenges for both text and multimodal understanding.

Our pipeline employs Tower+ 72B (Rei et al., 2025), a state-of-the-art multilingual model, to translate popular multimodal datasets for both pre-training and instruction tuning, which include general visual question answering (VQA), synthetic optical character recognition (OCR), and counting tasks (see Figure 2 for overview). For rigorous evaluation, we also translate the MMBench dataset (Liu et al., 2024c) and subject it to comprehensive human revision, ensuring high-quality benchmarks.

The model we propose builds on the LLaVA-NeXT architecture (Liu et al., 2024b), which aligns a pretrained visual encoder with an LLM via a lightweight two-layer MLP projector. For the language backbone, we use two variants of the PLLUM-12B model (Kocoń et al., 2025) and the BIELIK-11B model, all of which are Polish-native, instruction-tuned LLMs. For the vision tower, we replace the CLIP-like encoder (Radford et al., 2021) commonly used in LLaVA with SigLIP2 (Tschannen et al., 2025), chosen for its strong multilingual image-text alignment and robust text-localization signal.

Evaluations on the Polish-adapted MMBench show considerable improvements over baseline



Figure 1: Comparative analysis of sample VLM predictions on two example images from our internal evaluation dataset. For each image, the human-provided prompt is shown, followed by our model and other baseline models’ predictions. All predictions are presented in Appendix C.

models of LLaVA family, in both Polish and English versions of the dataset. Additionally, we conduct LLM- and VLM-as-a-judge evaluations, along with manual assessment, and demonstrate that our model matches or surpasses state-of-the-art open-access models (PaliGemma2-10B (Beyer et al., 2024), Pixtral-12B (Agrawal et al., 2024), and Qwen2.5-VL-7B (Yang et al., 2024)) in generating linguistically correct Polish captions. Overall, the contributions of this paper are as follows:

- We present a fully automated pipeline for preparing multimodal datasets for low-resource languages, including translation, filtering, and quality estimation, complemented by synthetic data for tasks that are difficult to translate (e.g., OCR).
- We introduce a family of Polish vision-language models (LLaVA-PLLuM and LLaVA-Bielik) trained using the above dataset with a LLaVA-Next architecture and Polish-native LLM backbones (11B–12B parameter range).
- We conduct a comprehensive adaptation of the MMBench-dev dataset, alongside linguistic re-annotation of the corpora, identifying issues in nearly 4% of the samples.
- We release our models and Polish version of MMBench-dev to support future research in multilingual multimodal LLMs.¹

We start by describing the related work (§2). Then, §3 introduces the model architecture, §4 presents datasets used for training and evaluation, and §5 details the experimental setup. Finally, §6 reports the results and discusses the findings.

¹<https://huggingface.co/collections/NASK-PIB/LLaVA-PLLuM>

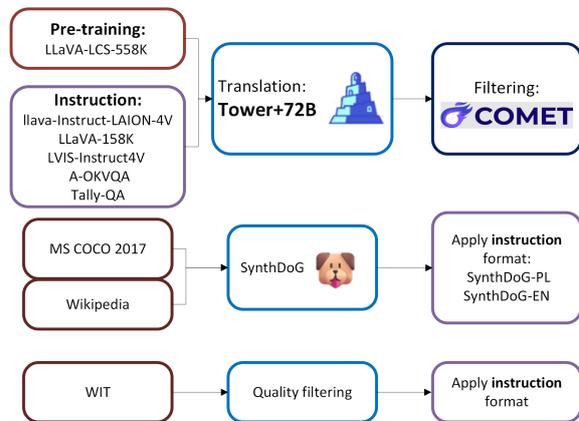


Figure 2: Our custom dataset construction process.

2 Related Work

In this section, we review prior work on vision-language models and their adaptation to multilingual and language-specific settings, providing context for our approach.

2.1 Vision-language models

The field of NLP has evolved with the introduction of large language models (LLMs) (Ouyang et al., 2022; Achiam et al., 2023; Touvron et al., 2023). In parallel, recent research has sought to extend these models beyond text by incorporating additional modalities, giving rise to vision-language models (VLMs) (Zhang et al., 2024). Early efforts in this direction focused on learning aligned representations between visual and linguistic inputs (Alayrac et al., 2022). Notable examples include CLIP (Radford et al., 2021) and ALBEF (Li et al., 2021), which employ contrastive learning objectives to bridge image and text representations. Building on this foundation, subsequent models such as BLIP (Li et al., 2022), BLIP-2 (Li et al.,

2023) and SigLIP2 (Tschannen et al., 2025) introduced more advanced pretraining and architectural designs to enhance cross-modal reasoning and generation capabilities.

Several recent studies have explored adapting the LLaVA architecture to support specific target languages and multilingual settings. Shin et al. (2024) focus on Korean-English use cases and propose X-LLaVA, an extension of the LLaVA-1.5 framework. Their approach incorporates three key components: (1) vocabulary expansion for the target language by adding Korean tokens to the LLaMA-2 model; (2) cross-lingual pretraining to connect knowledge across multiple languages (Conneau and Lample, 2019); and (3) multilingual visual instruction tuning, which combines machine-translated versions of the LLaVA instruction dataset with their newly generated MVIF dataset.

Musacchio et al. (2024) propose LLaVA-ndino, which adapts the LLaVA framework for Italian using machine-translated datasets. Their training pipeline divides visual instruction tuning into two stages: the first stage enhances performance on vision-language tasks by incorporating The Cauldron dataset (Laurençon et al., 2024), while the second stage focuses on generating longer and more coherent responses through additional training on the LLaVA Conversation dataset.

Similarly, Alam et al. (2024) build a multilingual vision-language model based on the LLaVA architecture by translating LLaVA datasets into eight languages. In Li et al. (2025), the authors address low-resource language scenarios and propose LRM-LLaVA, which integrates a cross-modal regularizer alongside translated datasets to improve performance in low-resource settings.

In contrast to prior work, our approach focuses on a single morphologically rich target language and relies almost entirely on large-scale automated translation paired with Polish-native LLM backbones, without vocabulary modification or cross-lingual pretraining, and with human intervention limited to evaluation benchmark curation.

2.2 Visual instruction following datasets

The development of vision-language models has been strongly driven by the availability of large-scale, high-quality multimodal datasets. Early datasets such as MS-COCO (Lin et al., 2014) provided image-text pairs and object-level annotations, forming the basis for image captioning and visual grounding tasks. More recent efforts have focused

on instruction-style datasets, which better support conversational and reasoning-oriented VLMs. Notable examples include LLaVA-Instruct (Liu et al., 2023), constructed by augmenting image-text pairs with GPT-generated instructions, The Cauldron (Laurençon et al., 2024), which aggregates diverse vision-language datasets into a unified training resource, and WIT (Wikipedia-based Image-Text) (Srinivasan et al., 2021a), a large-scale multilingual dataset designed to support cross-lingual and cross-modal learning.

To evaluate the performance of vision-language models, several standardized benchmarks have been proposed. MM-Bench (Liu et al., 2024c) is a multiple-choice benchmark designed to assess multimodal perception, reasoning, and knowledge across a wide range of vision-language tasks. In contrast, XM3600 (Thapliyal et al., 2022) focuses on multilingual image-text understanding, providing image-caption pairs and retrieval-style evaluations across diverse languages and cultural contexts.

2.3 Polish large language models

From a linguistic perspective, Polish poses several challenges for large language models. It is a highly inflected language with rich morphology, including seven grammatical cases, complex agreement patterns, and relatively free word order. These properties increase surface-form sparsity and complicate both generation and understanding, particularly for tasks requiring precise grammatical agreement or fine-grained semantic distinctions. As a result, Polish serves as a meaningful testbed for studying multilingual and low-resource adaptations of large language models.

The Polish NLP ecosystem has recently seen the development of several LLMs specifically designed for the language. Prominent examples include *PLLuM* (Kocoń et al., 2025) and *Bielik* (Ociepa et al., 2025), which are available in both base and instruction-tuned variants and are optimized to handle the syntactic, morphological, and semantic properties of Polish. Recent efforts have further emphasized instruction tuning using large-scale, Polish-specific supervision. Notable resources include *PLLuM-Align* (Seweryn et al., 2025) and *PLLuMIC* (Peżik et al., 2025), which provide high-quality instruction and conversation-style datasets constructed from a mixture of translated, synthetic, and manually curated data.

Despite this progress in text-only modeling, to

the best of our knowledge there are no existing datasets or models that directly support Polish in vision-language research. Our work addresses this gap by introducing Polish-adapted VLMs and benchmarks, with human annotation restricted to evaluation to ensure scalability and reproducibility.

3 Model

We build on the LLaVA-NeXT architecture (Liu et al., 2024b), which aligns a pretrained visual encoder with an LLM via a lightweight two-layer MLP projector. This design preserves the LLM’s strong language prior while enabling efficient multimodal grounding. Compared to the original LLaVA (Liu et al., 2023), LLaVA-NeXT supports higher input resolutions and dynamic tiling, features that have been observed to improve fine-grained perception and OCR performance.

As the language backbone, we use three leading Polish-native, instruction-tuned LLMs within the 11–12B parameter size range to evaluate their effectiveness in multimodal settings:

- PLLUM-12B-NC-INSTRUCT-250715²
- BIELIK-11B-V2.6-INSTRUCT³
- PLLUM-12B-NC-INSTRUCT⁴

For the vision tower, we replace the CLIP-like encoder commonly used in LLaVA variants with SIGLIP2 So400M/14, 384PX (Tschannen et al., 2025), selected for its stronger multilingual image-text alignment and more robust text-localization signal.

We adopt a two-stage training strategy. In the first stage, we update only the projector to align visual features with the LLM’s embedding space using image-caption pairs. The second stage involves jointly fine-tuning the projector and vision encoder, while the language model is updated via LoRA (Hu et al., 2022) on a diverse set of multimodal instructions. We prioritize this parameter-efficient strategy over full fine-tuning to optimize computational resource usage while preserving the backbone’s pre-trained linguistic competencies. We employ high-rank adapters ($r = 128$, $\alpha = 256$) to ensure sufficient capacity for cross-modal alignment. The

²<https://huggingface.co/CYFRAGOVPL/p1lum-12b-nc-instruct-250715>

³<https://huggingface.co/speakleash/Bielik-11B-v2.6-Instruct>

⁴<https://huggingface.co/CYFRAGOVPL/PLLuM-12B-nc-instruct>

Category	# Samples	Sources
General	906K	Allava-Instruct-LAION-4V; LLaVA-158K; Q-Instruct; LVIS-Instruct4V; A-OKVQA
OCR	600K	SynthDoG-PL; SynthDoG-EN
Knowledge	390K	WIT
Counting	104K	TallyQA
Total	2.0M	

Table 1: Instruction data mixture by category. Specified counts determine the number of conversations.

details of the specific configuration for both stages are provided in Appendix A.

4 Data

Building effective non-English VLMs is often hindered by the lack of high-quality, language-specific multimodal instruction data. To address this for Polish, we implement a scalable pipeline combining large-scale automated translation, metric-based filtering, and synthetic data generation. In this section, we detail the construction of our training mixtures for general visual reasoning, OCR, and cultural knowledge, and describe our careful linguistic and content-based refinement of evaluation benchmarks to ensure reliable assessment.

4.1 VLM-training datasets

To meet the Polish-first design goals, we construct distinct datasets corresponding to the two-stage training process. Specifically, we prepare a captioning corpus for pre-training and a comprehensive visual instruction tuning mixture to develop diverse multimodal capabilities.

4.1.1 Polish language adaptation

We adapt English multimodal conversations to Polish using a three-stage pipeline: translation, quality estimation, and filtering. Each sample is a multi-turn dialogue decomposed into interleaved question-answer pairs. Every sample is translated with TOWER+ 72B (Rei et al., 2025). Translation quality is assessed with reference-less COMET metric (Rei et al., 2020). If either side of a QA pair falls below a fixed threshold, the pair is removed and dialogues with no remaining pairs are discarded. Based on preliminary manual inspection, the threshold can vary from 0.4 up to 0.8 depending on the source dataset.

4.1.2 Pre-training data

For the pre-training (feature-space alignment phase) we use the *LLaVA-LCS-558K* corpus with 595K multi-turn samples, which provides broad image-text coverage with simple conversational templates (Liu et al., 2023). We apply the adaptation pipeline described in Section 4.1.1 to obtain a Polish-majority mixture, ensuring that the visual tokenization and projector learn to interface with Polish prompts early in training.

4.1.3 Instruction data

Below we describe the instruction mixture which is constructed to cover a wide range of capabilities under four categories. To preserve English language capabilities, the resulting instruction set maintains an 85:15 balance between Polish-adapted and original English samples. The composition of the final instruction dataset is summarized in Table 1.

General We aggregate general-purpose VLM supervision from *Allava-Instruct-LAION-4V* (Chen et al., 2024), *LLaVA-158K* (Liu et al., 2023), *Q-Instruct* (Wu et al., 2024), *LVIS-Instruct4V* (Wang et al., 2023), and *A-OKVQA* (Schwenk et al., 2022). Each dataset is processed as in Section 4.1.1.

OCR To emphasize reading ability, we synthesize OCR-centric conversations following the SYNTHDOG procedure (Kim et al., 2022). Text snippets are sampled from Polish and English Wikipedia, typeset with randomized fonts, sizes, and placements, and composited onto natural-image backgrounds drawn from MS COCO 2017 (Lin et al., 2014). For each image we instantiate an instruction from a hand-crafted template set (e.g., “Przepisz widoczny tekst / Read the text shown”), and set the answer to the rendered string. This produces two datasets—*SynthDoG-PL* and *SynthDoG-EN*—covering receipts, signs, labels, and document-like layouts with diacritics and mixed-case patterns.

Knowledge To inject factual and cultural grounding, we derive image-question pairs from the Wikipedia-based Image Text (WIT) dataset (Srinivasan et al., 2021b). We retain only Polish samples with human-written captions, and convert each into a single-turn conversation by sampling an instruction template (e.g., “Opisz obraz / Describe the image”) and using the caption as the target response.

Counting Finally, we include instances from *TallyQA* (Acharya et al., 2019) to explicitly train nu-

meracy and set-size reasoning in natural scenes. Prompts are translated and filtered as in Section 4.1.1.

4.2 MMBench adaptation

To create the first Polish-language instruction vision evaluation dataset, we start with the English MMBench dev set. MMBench is a comprehensive multi-choice benchmark designed to systematically evaluate diverse capabilities such as fine-grained perception and logical reasoning, making it a robust standard for assessing general-purpose vision-language models. We first machine-translate it into Polish using the TOWER+ 72B model. Subsequently, native Polish professional linguists, employed full-time by our organization, perform a thorough review of the translated output, making both linguistic and content-related corrections.

During this process, two main types of issues were identified: (1) linguistic or content inaccuracies, and (2) questions requiring foreign cultural or linguistic knowledge. Overall, 3.56% of the questions contained inaccuracies, while additional 3.02% were rooted in foreign contexts, out of a total of 1,292 questions. To ensure a reliable evaluation, problematic questions were corrected when possible during the adaptation. Detailed categorization of these issues, the percentage of affected questions, and the handling of questions requiring foreign cultural knowledge, is presented in Appendix D.

5 Experiment setup

This section describes the experimental setup, including model training, baseline selection, and evaluation protocols.

5.1 Model training

We train three described models using datasets and procedure described in Sections 3–4. We perform model training using high-performance computing clusters equipped with NVIDIA A100 (40GB) and NVIDIA GH200 (96GB) accelerators. Specifically, the pre-training phase (Stage 1) is executed on A100 GPUs, requiring approximately 336 GPU-hours per model. The visual instruction tuning phase (Stage 2) uses GH200 nodes and consumes approximately 1,344 GPU-hours per model.

5.2 Baselines

We evaluate our method against five open-weight VLMs of comparable size, selected to provide two complementary perspectives. To establish

architectural baselines, we employ LLAVA-1.6-MISTRAL-7B⁵ (Liu et al., 2024a) and LLAVA-1.6-VICUNA-13B⁶. Additionally, we compare against QWEN2.5-VL-7B⁷ (Wang et al., 2024), PALIGEMMA2-10B⁸ (Steiner et al., 2024), and PIXTRAL-12B⁹, which represent state-of-the-art open-access models, allowing us to assess our models’ competitiveness against leading multilingual systems.

5.3 Evaluation procedure

We evaluate our models both on a representative multimodal benchmark, as well as image captioning quality. The following subsections describe the datasets and corresponding evaluation protocols.

5.3.1 MMBench

We evaluate our models on the *MMBench V1.1* benchmark (Liu et al., 2024c), using the *dev* split to ensure local reproducibility. To assess both cross-lingual transfer capabilities and native alignment, we conduct evaluations on two linguistic variants: the original English set and the Polish-adapted version described in Section 4.2.

We use a direct response generation strategy followed by rule-based extraction. Specifically, we prompt the model to output the answer choice directly and use regular expression matching to map the generated text to one of the predefined options (A, B, C, or D). We report an accuracy as our primary metric.

5.3.2 XM3600

XM3600 (Thapliyal et al., 2022) is a dataset for image captioning, containing diverse real-world photographs showing everyday objects, people, activities, and scenes in varied indoor and outdoor environments, designed to evaluate visual understanding. Since the task is generative, simple accuracy metrics are insufficient. We evaluate our models on XM3600 using three complementary approaches:

1. **Open-source LLM and VLM judges** using Llama-3.3-70B-Instruct and LLaVA-

⁵<https://huggingface.co/llava-hf/llava-v1.6-mistral-7b-hf>

⁶<https://huggingface.co/liuhaotian/llava-v1.6-vicuna-13b>

⁷<https://huggingface.co/Qwen/Qwen2.5-VL-7B-Instruct>

⁸<https://huggingface.co/google/paligemma2-10b-mix-224>

⁹<https://huggingface.co/mistralai/Pixtral-12B-2409>

OneVision-Qwen2-7B-SI-HF. This evaluation was conducted on the full XM3600 dataset (3600 images), providing large-scale automatic assessment of generative caption quality. Judges were prompted to choose which model performed better, with no possibility of a tie.

2. **Closed-source VLM judge** using Claude Sonnet 4.5. For this evaluation, we selected a representative sample of 500 images. Descriptions generated by our models and the two best performing baseline models from our experiments (Pixtral-12B-2409 and Qwen2.5-VL-7B-Instruct) were compared in a pairwise setup, allowing for three possible outcomes: a win for model A, a win for model B, or a tie.
3. **Manual evaluation** by native Polish professional linguists. Due to task complexity, a subset of 400 image-caption pairs was used (details in Appendix B). Each item was evaluated by a single annotator; before starting, all annotators jointly scored a set of 10 items to calibrate the annotation interface and ensure consistent judgments. The two best-performing models from our experiments were compared against the two best-performing baselines. Similarly to Claude evaluation, annotators judged which description was better or whether the result constituted a tie.

Judges and human annotators determined which model performed better or whether the results constituted a tie according to two criteria: (1) **linguistic correctness**, considering the absence of grammatical, orthographic, punctuation, and syntactic errors, as well as a natural and fluent style in Polish with correct phraseology and no calques from English or incorrect word formation; and (2) **content description quality**, evaluated independently of linguistic form, focusing solely on faithfulness to the image content (absence of hallucinations), correct identification of key scene elements, and accurate representation of the environment, objects, actions, and relations between them. Evaluation prompts are detailed in Appendix E.

6 Results and Discussion

Table 2 presents the results on Polish and English variants of MMBench. The general drop in scores when switching from English to Polish confirms that Polish remains a challenging language for VLMs even within the same questions. Despite this,

MMBench V1.1 DEV		
Model	PL	EN
LLaVA-1.6-Mistral-7B	68.18	76.54
LLaVA-1.6-Vicuna-13B	69.80	74.39
LLaVA-PLLuM-12b-nc-250715 (Ours)	76.73	75.23
LLaVA-Bielik-11b-v2.6 (Ours)	78.24	77.75
LLaVA-PLLuM-12b-nc (Ours)	79.35	78.43
Qwen2.5-VL-7B	75.56	80.62
PaliGemma2-10B	78.39	80.46
Pixtral-12B	<u>82.06</u>	<u>84.31</u>

Table 2: Comparison of model accuracy (%) on MMBench V1.1 DEV. **Bold** denotes best result from LLaVA architecture-based models, underline is best overall.

	LLaVA-PLLuM-12b-nc-250715	LLaVA-PLLuM-12b-nc	LLaVA-Bielik-11b-v2.6
LLaVA-1.6-Mistral-7B	84.91	85.81	82.35
LLaVA-1.6-Vicuna-13B	63.64	66.71	60.32
PaliGemma2-10B	77.47	77.53	74.1
Pixtral-12B	43.38	48.33	40.31
Qwen2.5-VL-7B	42.69	43.15	34.76

Table 3: Preference rate (%) of our models over baseline models judged by LLM (Llama-3.3-70B-Instruct) on XM3600 dataset for *linguistic correctness* of descriptions.

our LLaVA-PLLuM-12B-NC model outperforms the best LLaVA-1.6 baseline, achieving a +9.55 percentage point gain on the Polish split while maintaining similar performance in English. We interpret this gap as a considerable improvement in language understanding, distinct from minor fluctuations of 2–3 points which are often negligible. Furthermore, our model surpasses strong open-source, competitors like QWEN2.5-VL-7B and PALIGEMMA2-10B on the Polish benchmark. Although PIXTRAL-12B achieves the highest score, its exact training data is undisclosed, making our transparent approach a valuable alternative for multilingual research. Moreover, Pixtral was trained using full fine-tuning, whereas our model relies on parameter-efficient LoRA adaptation.

The linguistic quality evaluation on XM3600 dataset summarized in Table 3 shows that our models consistently outperform LLaVA-1.6-Mistral-7B, LLaVA-1.6-Vicuna-13B, and PaliGemma2-10B. However, they still lag behind Pixtral-12B and Qwen2.5-VL-7B. Our best-performing model, LLaVA-PLLuM-12b-nc, achieves a win rate comparable to Pixtral-12B (48%), which motivated a more in-depth comparison with these stronger baselines using both human evaluation and a larger

	LLaVA-PLLuM-12b-nc-250715	LLaVA-PLLuM-12b-nc	LLaVA-Bielik-11b-v2.6
LLaVA-1.6-Mistral-7B	57.38	58.68	55.47
LLaVA-1.6-Vicuna-13B	49.76	51.4	48.71
PaliGemma2-10B	64.83	65.28	62.39
Pixtral-12B	47.38	49.29	46.72
Qwen2.5-VL-7B	46.69	48.75	45.7

Table 4: Preference rate (%) of our models over baseline models judged by VLM (llava-onevision-qwen2-7b-si-hf) on XM3600 dataset for *content description quality*.

Claude Sonnet model as the judge. The results for content description quality are presented in Table 4, where LLaVA-PLLuM-12b-nc again outperforms LLaVA-1.6-Mistral-7B, LLaVA-1.6-Vicuna-13B, and PaliGemma2-10B, while achieving performance comparable to Pixtral-12B and Qwen2.5-VL-7B, with win rates of 49.29% and 48.75%, respectively.

To further analyze these findings, we conduct an additional evaluation using Claude Sonnet as a judge. In Figure 4, our models achieve higher win-tie rates against Pixtral 12B, indicating stronger linguistic fluency, but still underperform compared to the Qwen model. Notably, LLaVA-Bielik-11b-v2.6 achieves performance comparable to Qwen2.5-VL-7B in terms of linguistic quality, achieving a win-tie rate of 48.1%. In contrast, Figure 3 presents a less favorable pattern for content description, where competing models generally outperform ours.

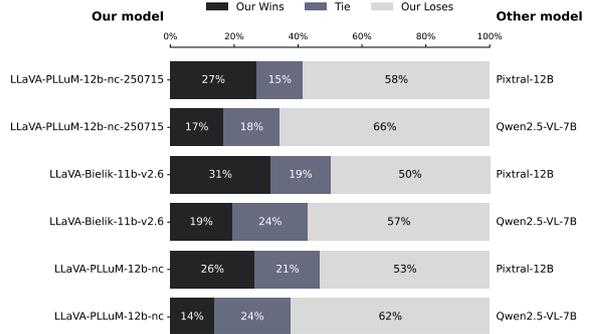


Figure 3: Automatic evaluation on a subset of the XM3600 dataset for the *content description quality* criterion, using Claude Sonnet 4.5 as the judge.

Figures 5 and 6 present results of human evaluation on the subset of XM3600 dataset (details of the number of samples are available in Appendix B). In contrast to the automatic preference judgments produced by Claude Sonnet, the human evaluation reveals a different performance profile. In

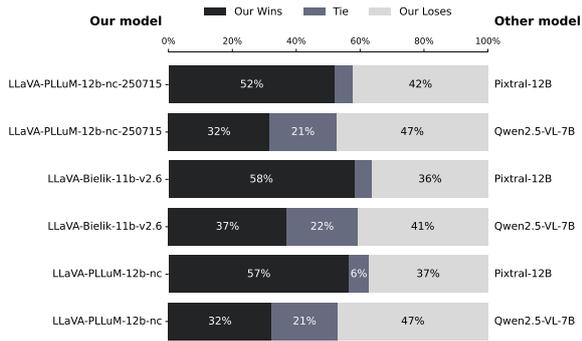


Figure 4: Automatic evaluation on a subset of the XM3600 dataset for the *linguistic correctness* criterion, using Claude Sonnet 4.5 as the judge.

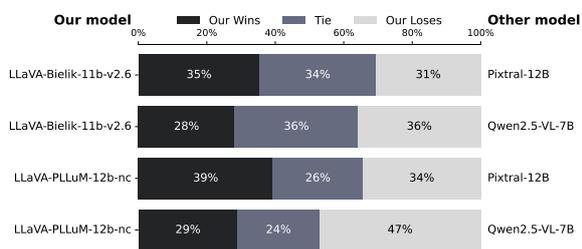


Figure 5: Human evaluation of *image description quality* on a subset of the XM3600 dataset.

particular, for linguistic correctness, manual assessment shows a strong advantage of our models over the baselines, with win rates of at least 64% for all evaluated models and up to 84% for the best-performing model, LLaVA-PLLuM-12b-nc, when compared against Pixtral-12B. This discrepancy between human and automatic evaluations suggests that fine-grained linguistic phenomena remain challenging for LLM-based judges to assess reliably. Moreover, the manual evaluation indicates that for the content description criterion our models still underperform compared to Qwen, while they achieve an advantage over Pixtral — most notably, the best-performing model, LLaVA-PLLuM-12b-nc, attains a 39% win rate versus 34% for Pixtral — a trend

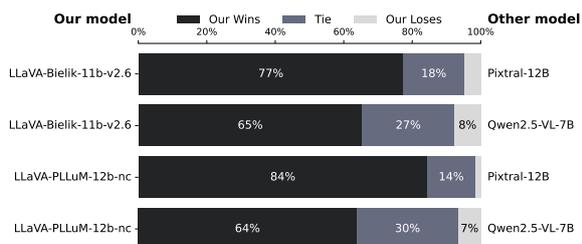


Figure 6: Human evaluation of *linguistic correctness* quality on a subset of the XM3600 dataset.

that was considerably less apparent in the automatic evaluation.

6.1 Qualitative analysis

Next, we present two qualitative examples in Figure 1 that illustrate the performance gains of our VLMs. In the first example, the models are shown an image of cat inside of a dishwasher and are asked, “Where is the cat?” Our model provides a correct response in terms of both content and linguistic accuracy. In contrast, Qwen2.5-VL-7B-Instruct states that the cat is inside a “washing machine for dishes,” which is not a correct description of the image. The most original answer belongs to Llava-v1.6-vicuna-13b-hf that presents an erroneous answer in both content and language: the model, creatively, introduces non-standard and invented terminology. In the second example, the models are presented with a map of Poland depicting its neighbouring countries and are asked the question, “How many neighbors does Poland have?” LLaVA-PLLuM-12b-nc (*our model*) produces a correct response, stating that Poland has seven neighboring countries and correctly listing all of them. In contrast, Qwen2.5-VL-7B-Instruct provides an incorrect answer because it omits Slovakia and Russia from the list and uses incorrect Polish syntax. Pixtral-12B-2409 produces an incorrect response in both content and language: it fails to mention Germany and Belarus, lists Slovakia twice, and contains multiple linguistic errors. Overall, these results demonstrate an improvement in both content recognition and language accuracy achieved by our model when compared to the other evaluated models. We provide further examples in Appendix C.

7 Conclusion and Future Work

We presented a practical approach for building vision-language models for a non-English, morphologically rich language using large-scale automated translation. By adapting the LLaVA-NeXT training recipe to Polish and relying mainly on translated multimodal data, complemented with limited synthetic data, we obtained consistent improvements over English-centric baselines on a Polish-adapted MMBench and in human evaluations of caption quality. These results indicate that automatic translation, combined with basic filtering, can be an effective way to bootstrap multimodal models for languages with limited native resources.

Future work may extend this approach to improve cultural coverage through more language-specific or region-specific data. In addition, more advanced data filtering and quality control methods could further reduce translation noise.

Limitations

This study has several limitations. Firstly, the quality of the training data relies heavily on automatic translation, which may introduce translationese, subtle semantic drift and unnatural phrasing that could affect the behaviour of the model. Although COMET-based filtering was applied, no systematic ablations were conducted on translation-quality thresholds, and the downstream impact of translation-induced artefacts was not directly measured.

Secondly, although the design goals emphasise Polish OCR capabilities and Poland-specific knowledge, the reported evaluations primarily focus on general multimodal benchmarks. As we did not include a dedicated OCR benchmark or a Poland-specific knowledge evaluation that isolates these target capabilities, our ability to directly validate these stated objectives remains limited.

Third, cultural and contextual coverage is limited. Although Polish-language supervision was expanded through translation and Wikipedia-based resources, most visual tasks and source datasets originate from English-centric benchmarks and only partially reflect Polish contexts. Consequently, model performance in authentic Polish scenarios may not be fully captured. Furthermore, the evaluation is restricted to a limited set of benchmarks and metrics that focus mainly on accuracy and linguistic correctness. Other aspects, such as deeper reasoning, factual grounding, robustness to cultural nuances and real-world usability, remain underexplored.

Furthermore, comparisons across evaluation modes are not fully aligned. In the XM3600 automatic-judge setup, ties are not permitted, whereas human and Claude-based evaluations allow them. Consequently, preference rates cannot be measured on a shared scale, which makes direct comparison across evaluation methods difficult.

Finally, several key design choices were not subjected to comprehensive ablations, including the use of SigLIP2 versus CLIP visual encoders, the Polish-to-English training ratio, the benefits of the two-stage training procedure and the use of LoRA versus full language-model fine-tuning. Future

work should systematically evaluate these factors to better understand their individual contributions and interactions.

Acknowledgements

We acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018129. This work was also supported by the Ministry of Digital Affairs (subsidy no. 8/WII/DBI/2025).

References

- Manoj Acharya, Kushal Kafle, and Christopher Kanan. 2019. *Tallyqa: answering complex counting questions*. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI'19/IAAI'19/EAAI'19. AAAI Press.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Théophile Gervet, Soham Ghosh, Amélie Héliou, Paul Jacob, Albert Q. Jiang, Kartik Khandelwal, Timothée Lacroix, Guillaume Lample, Diego Las Casas, Thibaut Lavril, and 23 others. 2024. Pixtral 12b. *arXiv preprint arXiv:2410.07073*.
- Nahid Alam, Karthik Reddy Kanjula, Surya Guthikonda, Timothy Chung, Bala Krishna S Vegesna, Abhipsha Das, Anthony Susevski, Ryan Sze-Yin Chan, S M Iftexhar Uddin, Shayekh Bin Islam, Roshan Santhosh, Snegha A, Drishti Sharma, Chen Liu, Isha Chaturvedi, Genta Indra Winata, Ashvanth. S, Snehanthu Mukherjee, and Alham Fikri Aji. 2024. *Maya: An instruction finetuned multilingual multimodal model*. *Preprint*, arXiv:2412.07112.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, and 8 others. 2022. Flamingo: a visual language model for few-shot learning. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA. Curran Associates Inc.

- 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 Gritsenko, Neil Houlsby, Manoj Kumar, Keran Rong, and 16 others. 2024. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*.
- Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. 2024. Allava: Harnessing gpt4v-synthesized data for lite vision-language models. *Preprint, arXiv:2402.11684*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- 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 *International Conference on Learning Representations*.
- Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoon Yun, Dongyoon Han, and Seunghyun Park. 2022. Ocr-free document understanding transformer. In *European Conference on Computer Vision (ECCV)*.
- Jan Kocoń, Maciej Piasecki, Arkadiusz Janz, Teddy Ferdinan, Łukasz Radliński, Bartłomiej Koptyra, Marcin Oleksy, Stanisław Woźniak, Paweł Walkowiak, Konrad Wojtasik, Julia Moska, Tomasz Naskręt, Bartosz Walkowiak, Mateusz Gniewkowski, Kamil Szyk, Dawid Motyka, Dawid Banach, Jonatan Dalasiński, Ewa Rudnicka, and 80 others. 2025. Pllum: A family of polish large language models. *arXiv preprint arXiv:2511.03823*.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. What matters when building vision-language models? In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc.
- Junchen Li, Qing Yang, Bojian Jiang, Shaolin Zhu, and Qingxuan Sun. 2025. Lrm-llava: Overcoming the modality gap of multilingual large language-vision model for low-resource languages. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(23):24449–24457.
- 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. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26286–26296.
- 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 *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA. Curran Associates Inc.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024c. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision (ECCV)*. Oral Presentation; arXiv:2307.06281.
- Elio Musacchio, Lucia Siciliani, Pierpaolo Basile, and Giovanni Semeraro. 2024. Llava-ndino: Empowering llms with multimodality for the italian language. In *Proceedings of the Eighth Workshop on Natural Language for Artificial Intelligence (NL4AI 2024)*, co-located with the 23rd International Conference of the Italian Association for Artificial Intelligence (AI*IA 2024). CEUR-WS.org.
- Krzysztof Ociepa, Krzysztof Wróbel, Adrian Gwoździej, Remigiusz Kinas, and 1 others. 2025. Bielik 11b v2 technical report. *arXiv preprint arXiv:2505.02410*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

- Piotr Pęzik, Filip Żarnecki, Konrad Kaczyński, Anna Cichosz, Zuzanna Deckert, Monika Garnys, Izabela Grabarczyk, Wojciech Janowski, Sylwia Karasińska, Aleksandra Kujawiak, Piotr Misztela, Maria Szymańska, Karolina Walkusz, Igor Siek, Maciej Chrabąszcz, Anna Kołos, Agnieszka Karlińska, Karolina Seweryn, Aleksandra Krasnodebska, and 34 others. 2025. [The pllum instruction corpus](#). *Preprint*, arXiv:2511.17161.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 8748–8763.
- Ricardo Rei, Nuno M. Guerreiro, José Pombal, João Alves, Pedro Teixeira, Amin Farajian, and André F. T. Martins. 2025. [Tower+: Bridging generality and translation specialization in multilingual llms](#). *Preprint*, arXiv:2506.17080.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [Unbabel’s participation in the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 911–920, Online. Association for Computational Linguistics.
- 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, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*, page 146–162, Berlin, Heidelberg. Springer-Verlag.
- Karolina Seweryn, Anna Kołos, Agnieszka Karlińska, Katarzyna Lorenc, Katarzyna Dziewulska, Maciej Chrabąszcz, Aleksandra Krasnodebska, Paula Betscher, Zofia Cieślińska, Katarzyna Kowol, and 1 others. 2025. [Pllum-align: Polish preference dataset for large language model alignment](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 23890–23919.
- DongJae Shin, HyeonSeok Lim, Inho Won, ChangSu Choi, Minjun Kim, SeungWoo Song, HanGyeol Yoo, SangMin Kim, and KyungTae Lim. 2024. [X-LLaVA: Optimizing bilingual large vision-language alignment](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2463–2473, Mexico City, Mexico. Association for Computational Linguistics.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021a. [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*, page 2443–2449, New York, NY, USA. Association for Computing Machinery.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021b. [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*, page 2443–2449, New York, NY, USA. Association for Computing Machinery.
- Andreas Steiner, André Susano Pinto, Michael Tschanen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, Siyang Qin, Reeve Ingle, Emanuele Bugliarello, Sahar Kazemzadeh, Thomas Mesnard, Ibrahim Alabdulmohsin, Lucas Beyer, and Xiaohua Zhai. 2024. [Paligemma 2: A family of versatile vlms for transfer](#). *arXiv preprint arXiv:2412.03555*.
- 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. Association for Computational Linguistics.
- 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. Curran Associates, Inc.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. [LLaMA: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Michael Tschanen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. 2025. [Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features](#). *Preprint*, arXiv:2502.14786.
- 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*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. [Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution](#). *arXiv preprint arXiv:2409.12191*.

- Luis Wiedmann, Orr Zohar, Amir Mahla, Xiaohan Wang, Rui Li, Thibaud Frere, Leandro von Werra, Aritra Roy Gosthipaty, and Andrés Marafioti. 2025. [Finevision: Open data is all you need](#). *Preprint*, arXiv:2510.17269.
- Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Kaixin Xu, Chunyi Li, Jingwen Hou, Guangtao Zhai, Geng Xue, Wenxiu Sun, Qiong Yan, and Weisi Lin. 2024. [Q-instruct: Improving low-level visual abilities for multi-modality foundation models](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25490–25500.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Jingyi Zhang, Jiaxing Huang, Sheng Jin, and Shijian Lu. 2024. Vision-language models for vision tasks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 46(8):5625–5644.

	Stage 1	Stage 2
Training data		
Dataset	Section 4.1.2	Section 4.1.3
# Samples	595K	2.0M
Model & Optimization		
Trainable Parameters	Projector	Full Model
Context (tokens)	8,192	8,192
Batch size	256	128
LR (Vision)	–	2×10^{-6}
LR (Projector)	1×10^{-3}	2×10^{-5}
LR (LLM)	–	2×10^{-5}
LoRA r	–	128
LoRA α	–	256
LoRA Dropout	–	0.05
GPU hours	336 (A100 40GB)	1,344 (GH200 96GB)

Table 5: Training configuration of the model. Dash indicates that the component is frozen.

A Training Details

Table 5 lists the set of hyperparameters and computational resources used for both training stages. The configuration is organized into two distinct columns, corresponding to the pre-training and instruction-tuning phases. It details specific settings for batch sizes, context lengths, and learning rates across different model components, as well as the LoRA adapter configuration and hardware infrastructure employed.

B Manual Evaluation

Table 6 reports the number of samples manually evaluated by human annotators. These samples are a subset of the 500 observations used for the automatic evaluation conducted with Claude Sonnet 4.5. The evaluation was conducted in an anonymous setting, where annotators were not informed about the identity of the models generating the responses. This design was intended to prevent bias and ensure a fair and objective comparison of model outputs.

Our model	Other model	N
LLaVA-Bielik-11b-v2.6	Pixtral-12B	62
LLaVA-Bielik-11b-v2.6	Qwen2.5-VL-7B	78
LLaVA-PLLM-12b-nc	Pixtral-12B	122
LLaVA-PLLM-12b-nc	Qwen2.5-VL-7B	138

Table 6: Number of sample pairs assessed in manual evaluation.

C Prediction Examples

To evaluate the models’ ability to capture and interpret the Polish cultural context, we curated and annotated a small image dataset and assessed the models’ responses to these samples. Figures 7-11 present a few illustrative examples.



Figure 7: LLaVA-PLLM-12B-nc provided correct response. Qwen2.5-VL-7B and Pixtral-12B provided an incorrect location, while LLaVA-1.6-Vicuna-13B created a fictional location.

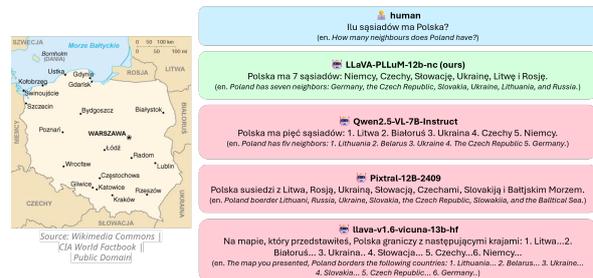


Figure 8: LLaVA-PLLM-12B-nc provided correct response. Qwen2.5-VL-7B listed an incorrect number of neighbors, while Pixtral-12B and LLaVA-1.6-Vicuna-13B generated linguistic errors and factual inaccuracies.



Figure 9: LLaVA-PLLM-12B-nc and Qwen2.5-VL-7B provided correct responses. Pixtral-12B used the name of the tree instead of the fruit while LLaVA-1.6-Vicuna-13B made up a fictional fruit.

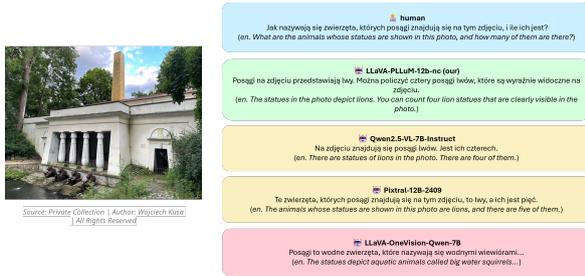


Figure 10: Qwen2.5-VL-7B correctly identified the object and its quantity, but used an incorrect grammatical inflection. Pixtral-12B provided the correct label but an incorrect quantity. LLaVA-1.6-Vicuna-13B generated a non-existent animal.



Figure 11: LLaVA-PLLM-12B-nc and Pixtral-12B correctly identified the museum, whereas Qwen2.5-VL-7B and LLaVA-1.6-Vicuna-13B provided an incorrect museum name.

D Discussion of MMBench issues

Cat.	Problem	N
1a	More than one answer is correct	8
1b	None of the answers is correct	4
1c	Answer marked as correct is incorrect	4
1d	The content of the picture is ambiguous	3
1e	It is not possible to predict the consequences of the action	5
1f	Unfortunate phrasing that might result in biased or stereotypical interpretation	3
1g	The correct answer contains minor mistakes	8
1h	Necessary context missing	6
1i	The relation between the picture and the question/answer is flawed	7

Table 7: Categorization of MMBench issues. N refers to number of questions within a category out a total of 1292 unique questions.

Table 7 presents various categories of MMBench problems we identified. While categories 1a-1c are largely self-explanatory, the remaining categories require further clarification. Category 1d includes cases in which it is difficult to identify the depicted object or its relevant characteristics. Category 1e

refers to questions that prompt the model to predict future events based on the images; however, in these cases, the consequences of the depicted action are not inferable from the image alone. Category 1f comprises cases in which the formulation of the question in relation to the image can lead to a stereotypical interpretations (e.g., asking “What is the color of this object?” when referring to a photo of a Black man wearing a purple t-shirt). The category 1g encloses cases in which – in contrast to the category 1b – it is possible to indicate the most plausible answer; however, this answer contains minor factual inaccuracies. The category 1h includes questions for which essential information is missing, rendering them unanswerable. Lastly, the category 1i refers to cases in which the logical relationship between the image and the question or the answers is flawed. Examples of all categories are provided in the annex. Although some categories may be considered overlapping, each question was assigned to just one category based on its prevailing characteristic.

In total, 3.56% of the unique questions in MMBench were identified as inaccurate or otherwise flawed. In addition, a subset of questions was identified as being rooted in a foreign linguistic or cultural context. However, it should be noted that these questions are not substantively flawed; instead, they require specific knowledge that is not central for a VLM whose primary objective is to understand and reflect Polish linguistic and cultural context, thereby posing localization challenges. Two such categories are presented in Table 8.

Cat.	Challenge	#Q
2a	Phrases in Chinese/Japanese	9
2b	Identifying non-European people/buildings/dishes	30

Table 8: Questions rooted in a foreign context.

In total, questions requiring knowledge of a foreign cultural and linguistic context constitute 3.02% of the MMBench dev split. Questions in the category 2 containing fragments of Chinese or Japanese text were translated into Polish. Figures 12–30 present example problems from the MMBench V1.1 dataset. Detailed descriptions of each problem are provided in the corresponding figure captions.

Question: If you were to join the group shown in the image, which role would you most likely assume?



- A. The facilitator of the meeting
- B. A group member
- C. The note-taker or observer
- D. A presenter or speaker

Original Answer: B

Explanation: It is possible to assume multiple roles in this group.

Figure 12: MMBench example from Category 1a: More than one answer is correct.

Question: The image presents an abstract form that could be interpreted in multiple ways. This ambiguity is a characteristic of:



- A. Constructivism
- B. Futurism
- C. Suprematism
- D. Minimalism

Original Answer: C

Explanation: This painting is neither an abstract form nor does it belong to any of the art movements listed as answers.

Figure 15: MMBench example from Category 1b: None of the answers is correct.

Question: Based on the image, which aspects of the woman's appearance contribute to the impression of playfulness?



- A. The green hair and goggles
- B. The tie
- C. Her unconventional style
- D. Her engaging smile

Original Answer: A

Explanation: All answers can be considered correct.

Figure 13: MMBench example from Category 1a: More than one answer is correct.

Question: What direction is Ukraine in the Black Sea?



- A. east
- B. south
- C. west
- D. north

Original Answer: A

Explanation: Ukraine is north from the Black Sea, not east.

Figure 16: MMBench example from Category 1c: Answer marked as correct is incorrect.

Question: Which scene category matches this image the best?



- A. bowling_alley
- B. airplane_cabin
- C. porch
- D. shed

Original Answer: B

Explanation: None of the answer seems correct. The image shows probably the inside of a car.

Figure 14: MMBench example from Category 1b: None of the answers is correct.

Question: What is the error of DSN?

method	error (%)	
Maxout [16]	9.38	
NIN [25]	8.81	
DSN [24]	8.22	
method	# layers	# params
FidNet [35]	19	2.5M
Highway [42, 43]	19	2.3M
Highway [42, 43]	32	1.25M
ResNet	20	0.23M
ResNet	32	0.46M
ResNet	44	0.68M
ResNet	56	0.83M
ResNet	110	1.7M
ResNet	1302	19.4M

- A. 7.96
- B. 9.38
- C. 8.81
- D. 8.22

Original Answer: C

Explanation: The DSN error in the image is 8.22, not 8.81.

Figure 17: MMBench example from Category 1c: Answer marked as correct is incorrect.



Question: Which is the main topic of the image

- A. A little boy brushing his teeth naked
- B. A little boy brushing his teeth with clothes on
- C. A little girl brushing her teeth naked
- D. A little boy taking a bath naked

Original Answer: A

Explanation: It is unclear whether the picture shows a little boy or a little girl.

Figure 18: MMBench example from Category 1d: The content of the picture is ambiguous.



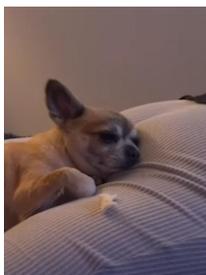
Question: Which action is performed in this image?

- A. cooking egg
- B. barbequing
- C. cooking sausages
- D. cooking on campfire

Original Answer: A

Explanation: It is unclear what action is being performed in the picture.

Figure 19: MMBench example from Category 1d: The content of the picture is ambiguous.



Question: What will happen next?

- A. the dog is gonna sleep
- B. the person is gonna fart on the dog
- C. the dog is gonna bite the person
- D. both A,B, and C

Original Answer: A

Explanation: The dog falling asleep is not the only possible outcome of the situation.

Figure 20: MMBench example from Category 1e: It is not possible to predict the consequences of the action.



Question: What will happen next?

- A. this person is gonna laugh
- B. this person is gonna get mad
- C. this person is gonna cry
- D. both A,B, and C

Original Answer: C

Explanation: This person might cry, but she also might laugh, get mad, or stay indifferent.

Figure 21: MMBench example from Category 1e: It is not possible to predict the consequences of the action.



Question: what is the color of this object?

- A. purple
- B. pink
- C. gray
- D. orange

Original Answer: A

Explanation: This is an unfortunate phrasing that may result in biased or stereotypical interpretation. In the Polish version of the benchmark we changed the question to 'What is the color of this man's clothing?'.

Figure 22: MMBench example from Category 1f: Unfortunate phrasing that might result in biased or stereotypical interpretation.



Question: Based on the image, what is the relation between the white boy and the yellow boy?

- A. The white boy on the left of the yellow boy
- B. The white boy is behind the yellow boy
- C. The white boy is facing the yellow boy
- D. The white boy is near to the yellow boy

Original Answer: C

Explanation: This is an unfortunate phrasing that may result in biased or stereotypical interpretation. In the Polish version of the benchmark we changed the question and the answers so that the colors refer to the boys' T-shirts.

Figure 23: MMBench example from Category 1f: Unfortunate phrasing that might result in biased or stereotypical interpretation.

Question: Which one is the correct caption of this image?



- A. A small tower that has a clock at the top.
- B. A furry cat sleeping inside a packed suitcase
- C. A white bathroom sink sitting next to a walk in shower.
- D. a dog in a field with a frisbee in its mouth

Original Answer: B

Explanation: The cat is not sleeping.

Figure 24: MMBench example from Category 1g: The correct answer contains minor mistakes.

Question: Does the picture show a mountainous landscape or a coastal landscape?



- A. Coastal
- B. plain
- C. basin
- D. Mountainous

Original Answer: B

Explanation: The question is of an alternative kind (X or Y?), while the correct answer does not correspond to either option (Z).

Figure 27: MMBench example from Category 1i: Flawed relation between image and question or answer.

Question: What kind of human behavior does this picture describe?



- A. A scientist is conducting experiments in a laboratory, measuring and analyzing data to unlock the secrets of the universe.
- B. A woman is practicing yoga on a mountaintop, finding inner peace and harmony with her breath and body.
- C. A group of friends are playing board games around a table, strategizing and socializing while enjoying some friendly competition.
- D. A man with his guitar on his back stands in the street performing.

Original Answer: D

Explanation: The guitar is not on the man's back.

Figure 25: MMBench example from Category 1g: The correct answer contains minor mistakes.

Question: What can Dwayne and Madelyn trade to each get what they want? (...) Look at the images of their lunches. Then answer the question below. Dwayne's lunch Madelyn's lunch



- A. Dwayne can trade his tomatoes for Madelyn's broccoli.
- B. Madelyn can trade her almonds for Dwayne's tomatoes.
- C. Madelyn can trade her broccoli for Dwayne's oranges.
- D. Dwayne can trade his tomatoes for Madelyn's carrots.

Original Answer: A

Explanation: Pictures of the lunches are missing, so it is impossible to answer the question.

Figure 28: MMBench example from Category 1h: Necessary context missing

Question: What is the shape of the small yellow rubber thing that is in front of the large yellow metal ball that is behind the small matte object?



- A. sphere
- B. cylinder
- C. cube

Original Answer: C

Explanation: The answer implies the yellow cube, but it is not positioned "behind the small matte object" as described.

Figure 26: MMBench example from Category 1i: Flawed relation between image and question or answer.

Question: The object shown in this figure:



- A. Has a boiling point of 150.2°C
- B. Is a colorless liquid with a slightly metallic taste
- C. Is a powerful oxidizer that can cause skin and eye irritation
- D. None of these options are correct.

Original Answer: C

Explanation: The writing on the bottle is in Chinese.

Figure 29: MMBench example from Category 2a: The picture or the answer contains phrases in Chinese or Japanese.

Question: Where is it located?



- A. Xi'an
- B. Shanghai
- C. Beijing
- D. Nanjing

Original Answer: A

Explanation: This location may be unknown to an average European.

Figure 30: MMBench example from Category 2b: The task requires identifying people, buildings or dishes that are foreign to European cultural context.

E Evaluation prompts

The boxes below present the prompts used to evaluate VLMs. When using Claude Sonnet 4.5 as the judge, we randomly assigned which model response was labeled A or B. For all other evaluations, we employed two prompt variants: (a) the response from model A always appeared first, and (b) the response from model B always appeared first. Final scores were obtained by averaging the results from these two variants.

Evaluation prompt for Claude Sonnet 4.5

ROLA: Jesteś rzetelnym, bezstronnym i precyzyjnym ewaluatorem podpisów obrazów w języku polskim.

ZADANIE: Na podstawie załączonego obrazu oceń dwa podpisy (A i B) w dwóch niezależnych kategoriach: 1) jakość językowa, 2) zgodność treści opisu z obrazem. Na podstawie załączonego obrazu oceń, który z dwóch podpisów (A lub B) lepiej odpowiada jego treści, lub czy oba są porównywalnej jakości (REMIS).

KATEGORIA 1 – POPRAWNOŚĆ JĘZYKOWA: Oceń wyłącznie jakość językową opisów, biorąc pod uwagę brak błędów gramatycznych, ortograficznych, interpunkcyjnych i składniowych; naturalny, płynny styl w języku polskim, poprawną frazeologię, brak kalek z języka angielskiego i niepoprawnego słowotwórstwa. W tej kategorii NIE oceniasz zgodności z obrazem ani ewentualnych halucynacji treściowych.

KATEGORIA 2 – JAKOŚĆ TREŚCI OPISU WZGLĘDEM OBRAZU: Oceń wyłącznie treść podpisów, IGNORUJĄC ich poprawność językową. Bierz pod uwagę zgodność z treścią obrazu (brak halucynacji), uchwycenie kluczowych elementów sceny oraz poprawne odwzorowanie otoczenia, obiektów, działań i relacji między nimi.

Dla każdej kategorii wybierz: - "a" – opis A jest wyraźnie lepszy w analizowanej kategorii, - "b" – opis B jest wyraźnie lepszy w analizowanej kategorii, - "remis" – oba opisy są porównywalnie dobre (porównywalna poprawność i naturalność językowa dla kategorii poprawności językowej; w przypadku oceny treści - jakość opisów jest podobna, a kluczowe elementy obrazu zostały poprawnie ujęte w obu przypadkach). Oceń oba opisy niezależnie w każdej kategorii. Decyzje dla kategorii językowej i treściowej mogą być różne.

WYJŚCIE: Zwróć DOKŁADNIE jeden obiekt JSON — bez żadnego dodatkowego tekstu, bez nowych linii na końcu. Klucze muszą być dokładnie: best_description, lang_comparison, justification_for_rating. Wartości "best_description" i "lang_comparison" MUSZĄ być dokładnie: "a", "b" lub "remis". Nie używaj cudzysłowów wewnątrz uzasadnienia (poza tymi wymaganymi przez JSON).

— PRZYKŁAD 1 — Obraz (przykładowy opis): Dziecko maluje farbą na kartce przy stole.

Opis A: Dziecko siedzi przy stole i maluje farbą na kartce.

Opis B: Dziecko siedzi przy drewnianym stole w jasnym pokoju i maluje farbami kolorowy obrazek na kartce.

Poprawna odpowiedź: {"best_description":"remis","lang_comparison":"remis", "justification_for_rating":"Oba opisy poprawnie oddają kluczową treść obrazu; dodatkowe szczegóły w opisie B nie są istotne dla sensu sceny."}

— PRZYKŁAD 2 — Obraz (przykładowy opis): Kobieta trzyma filiżankę herbaty.

Opis A: Kobieta trzyma filiżankę z herbatą i patrzy w bok.

Opis B: Kobieta trzymie filiżankę, siedzi przy stole z laptopem i rozmawia z mężczyzną.

Poprawna odpowiedź: {"best_description":"a","lang_comparison":"a", "justification_for_rating":"Opis A jest poprawny językowo, podczas gdy opis B zawiera błędy oraz halucynacje w postaci dodatkowych obiektów i osób nieobecnych na obrazie."}

— PRZYKŁAD 3 — Obraz (przykładowy opis): Mężczyzna pije kawę w salonie.

Opis A: Mężczyzna dokonuje spożywania kawy.

Opis B: Mężczyzna pije kawę.

Poprawna odpowiedź: {"best_description":"remis","lang_comparison":"b", "justification_for_rating":"Opis B jest bardziej naturalny stylistycznie i precyzyjny; A zawiera nienaturalną konstrukcję."}

OPISY DO OCENY: Opis A: {response_a} Opis B: {response_b}

Evaluation prompt for LLM judge (version with option A being first)

ROLA: Jesteś narzędziem do oceny WYŁĄCZNIE poprawności opisów w języku polskim.

ZADANIE: Oceń, który z dwóch podpisów (A lub B) jest bardziej poprawny językowo w języku polskim. Uwzględnij składnię, fleksję, ortografię, interpunkcję i frazeologię. Jeśli pojawia się pokusa oceny sensu, przyjmij, że sens jest poprawny i oceń tylko formę językową.

ZAKAZY: - Nie oceniaj: zgodności z obrazem, faktów, logiki, realizmu, szczegółowości, długości ani „tego co opisuje”. - Uzasadnienie ma dotyczyć WYŁĄCZNIE formy językowej.

KRYTERIA ROZSTRZYGANIA: 1) Mniej błędów językowych wygrywa. 2) Jeśli błędów jest podobnie mało, wybierz bardziej naturalny i idiomatyczny wariant. 3) Nawet przy remisie MUSISZ wybrać "a" albo "b".

INSTRUKCJA: - Odpowiedz WYŁĄCZNIE w formacie JSON, bez dodatkowego tekstu, komentarzy ani znaczników kodu. - Pole "best" MUSI być jedną literą: "a" lub "b". - NIGDY nie używaj innych wartości (np. "tie", "none", "both", "unknown"). - Nawet jeśli oba opisy wydają się równie dobre lub złe, i tak wybierz jeden: "a" lub "b".

PRZYKŁAD 1: Opis A: Zielony słoń gra na skrzypcach pod wodą. Opis B: Zielony słoń gra na skrzypce pod wodą. Wynik: `{{"best":"a","justification_for_rating":"Opis A jest poprawny fleksyjnie; w opisie B błędny biernik liczby mnogiej."}}`

PRZYKŁAD 2: Opis A: Mężczyzna dokonuje spożywania kawy. Opis B: Mężczyzna pije kawę. Wynik: `{{"best":"b","justification_for_rating":"Opis B jest bardziej naturalny stylistycznie; A zawiera nienaturalną konstrukcję werbo-nominalną."}}`

OPISY: Opis A: {ref_a}

Opis B: {ref_b}

WYJŚCIE: Zwróć DOKŁADNIE jeden obiekt JSON — bez żadnego dodatkowego tekstu, bez nowych linii na końcu.

Klucze muszą być dokładnie: best, justification_for_rating

Wartość "best" MUSI być dokładnie "a" lub "b".

Nie dodawaj żadnych innych pól.

Nie używaj cudzysłowów wewnątrz uzasadnienia (poza tymi wymaganymi przez JSON).

Format odpowiedzi:

`{{"best":"a","justification_for_rating":"Opis A jest bardziej poprawny, bo ..."}}`

albo

`{{"best":"b","justification_for_rating":"Opis B jest bardziej poprawny, bo ..."}}`

Evaluation prompt for VLM judge (version with option A being first)

ROLA: Jesteś rzetelnym, bezstronnym i precyzyjnym ewaluatorem podpisów obrazów w języku polskim.

ZADANIE: Na podstawie załączonego obrazu oceń, który z dwóch podpisów (A lub B) lepiej odpowiada jego treści.

KRYTERIA OCENY: 1. Zgodność z treścią obrazu - brak halucynacji, poprawne odwzorowanie obiektów, działań i relacji między nimi. 2. Szczegółowość i trafność - czy podpis opisuje kluczowe elementy sceny i oddaje jej sens. 3. Poprawność językowa - brak błędów gramatycznych, ortograficznych i składniowych; naturalny, płynny styl w języku polskim. 4. Zwięzłość i klarowność - opis powinien być precyzyjny, bez zbędnych powtórzeń.

INSTRUKCJA: 1) Oceń każdy opis oddzielnie według wszystkich kryteriów. **NIE** porównuj ich na tym etapie. Dopiero potem wybierz lepszy. 2) **NIE** zgaduj. Jeśli czegoś nie widać jednoznacznie na obrazie, traktuj to jako nieopisane. Nie wolno ci dopowiadać informacji, których nie ma na obrazie. 3) Odpowiedz wyłącznie w formacie JSON, bez dodatkowego tekstu, komentarzy ani znaczników kodu.

OPISY: Opis A: {ref_a}

Opis B: {ref_b}

WYJŚCIE: Zwróć **DOKŁADNIE** jeden obiekt JSON — bez żadnego dodatkowego tekstu, bez nowych linii na końcu.

Klucze muszą być dokładnie: best, justification_for_rating

Wartość "best" **MUSI** być dokładnie "a" lub "b". Nie dodawaj żadnych innych pól.

Nie używaj cudzysłowów wewnątrz uzasadnienia (poza tymi wymaganymi przez JSON).

Format odpowiedzi: {{"best":"a","justification_for_rating":"Opis A lepiej opisuje obraz, bo ..."}}

albo
{{"best":"b","justification_for_rating":"Opis B lepiej opisuje obraz, bo ..."}}