

# Does Visual Rendering Bypass Tokenization? Investigating Script-Tokenizer Misalignment in Pixel-Based Language Models

Anonymous ACL submission

## Abstract

While pixel-based language modeling aims to bypass the sub-word tokenization bottleneck by rendering text as images, recent multimodal variants such as DualGPT reintroduce text tokenizers to improve autoregressive performance. We investigate a fundamental question, **does visual rendering truly decouple a model from tokenization constraints?** Focusing on four Indonesian low-resource local languages that have their own non-Latin scripts (i.e., Javanese, Balinese, Sundanese, and Lampungnese), we evaluate the impact of script-tokenizer alignment within the DualGPT architecture. Our results show that, despite visual rendering, reintegrating a text tokenizer into the architecture reintroduces the same issue that pixel-based language modeling aims to resolve, which is the tokenizer misalignment problem. Despite having lower OOV and fertility rates, we show that the Llama 2 tokenizer performs significantly worse than a custom tokenizer, with improvements of up to 30.15 chrF++. Our findings serve as a warning for future multimodal variants, as text tokenizers remain a significant barrier to equitable models.

## 1 Introduction

While many Visual Language Models (VLMs) claim broad multilingual support (see Figure 1; Bercovich et al., 2025; Gemma Team, 2025; Wang et al., 2025; Bai et al., 2025; OpenAI, 2025), their performance on low-resource languages with non-Latin scripts remains poor (Han et al., 2025). A key challenge is script-tokenizer misalignment, where tokenizers trained predominantly on English data exhibit orthographic bias, causing excessive fragmentation of non-Latin scripts where single semantic units are split into multiple sub-word pieces.

Recent Pixel-based approaches introduced by Rust et al. (2022) avoid tokenization by processing rendered text as images. PixelGPT (Chai et al., 2024) extends this with autoregressive variants,

Five-shot Transliteration Performance of VLMs Across Four Languages




Figure 1: General VLMs image transliteration performance averaged across Javanese, Balinese, Sundanese, and Lampung on the NusaAksara evaluation dataset. Llama-3.1-NN-VL-8B-V1 is the Nemotron-Nano variant.

reintroducing text tokenizers in multimodal variants such as DualGPT. **If the visual modality truly decoupled the model from tokenization constraints, any tokenizer with high vocabulary coverage should yield similar results.** In contrast, our DualGPT results on four Indonesian low-resource local languages and scripts show that script-tokenizer misalignment remains a key bottleneck once tokenizers are reintroduced, even in pixel-based models.

## 2 Background and Motivation

**Tokenization and Cross-lingual Disparities.** Current LLMs predominantly use BPE-based tokenization (Sennrich et al., 2016; Abagyan et al., 2025), including Llama 2 (Touvron et al., 2023), which PixelGPT adopts (Chai et al., 2024). However, BPE tokenizers trained on English-dominant corpora create inequities for underrepresented languages through excessive subword fragmentation (Foroutan et al., 2025; Petrov et al., 2023). Research on tokenization-free models (Clark et al., 2022; Xue et al., 2022) also offers an alternative approach for handling multilingual data. However, its application on generation tasks remains under-explored (Sun et al., 2023).

**Pixel-based Language Modeling.** PIXEL (Rust et al., 2022) addresses tokenization disparities by rendering text as images, theoretically enabling unlimited multilingual vocabulary. Kesen et al. (2025) demonstrate strong low-resource performance with this approach. However, rendering strategies affect patch diversity and learned representations (Lotz et al., 2023; Tatariya et al., 2024), with Tatariya et al. (2024) noting that structured rendering may reintroduce tokenization-like constraints.

**Multimodal Variants.** Recent work explores combining pixel-based and tokenization-based approaches. PixelGPT (Chai et al., 2024) extends PIXEL to autoregressive modeling, with its multimodal variant DualGPT integrating both visual rendering and text tokenization to leverage complementary strengths of each modality. While this hybrid approach demonstrates improved downstream performance, it raises an important question, **does reintroducing text tokenizers reintroduce the cross-lingual disparities that motivated pixel-based approaches?** We investigate this question by examining tokenizer alignment effects in DualGPT across low-resource Indonesian scripts and languages.

### 3 Methodology

#### 3.1 Datasets

In this work, we analyze four Indonesian low-resource local scripts: Javanese (jav), Sundanese (sun), Balinese (ban), and Lampung (ljp), representing a range of resource levels and linguistic diversity.

Language	# Train	# Eval
Javanese	400,726	816
Sundanese	293,933	823
Balinese	54,017	450
Lampung	945	84

Table 1: Dataset Statistics.

Training data comes from Wikidumps (July 2025) and digitalized Indonesian folklore<sup>1</sup> for Javanese, Sundanese, and Balinese, preprocessed as described in Section 3.3. We use NusaAksara (Adilazuarda et al., 2025) as the evaluation dataset and preprocess it in the same manner as the training data. Moreover, we split the NusaAksara’s Lampung data for training and evaluation. Dataset statistics are presented in Table 1<sup>2</sup>.

<sup>1</sup>extracted using pdfplumber.

<sup>2</sup>Full dataset will be released after the double-blind review.

#### 3.2 PixelGPT

We use DualGPT (Chai et al., 2024)<sup>3</sup>, pretraining the model on text, images, and paired text-image data, then finetuned for image-to-text transliteration. Each model went through a single training and finetuning run, with all reported metrics based on that execution.

#### 3.3 Tokenizer and Renderer

Both a tokenizer and a renderer are required as the DualGPT model relies on both text and image modalities during pre-training. The visualization for the data preprocessing pipeline can be seen in Appendix B.

**Tokenizer.** We train our models using two different tokenizers: the default tokenizer used in the PixelGPT architecture (Llama 2 tokenizer) and a custom-built tokenizer. Our tokenizer is inspired by a grapheme-based approach (Basher et al., 2023), rather than the default BPE-based method. First, we back-transliterate our training dataset from Latin script text into *aksara* (the Indonesian local scripts) using community-built tools<sup>4</sup>. We then tokenize the *aksara* at the grapheme level and transliterate it back into Latin script text. Finally, we apply a word-level tokenization scheme to map the text into token IDs.

**Renderer.** With the findings of Tatariya et al. (2024) and the nature of Javanese and Balinese, which do not mark word boundaries with whitespace, we adopt a continuous rendering strategy following Lotz et al. (2023). We use a custom Pillow-based renderer instead of the default Pygame renderer on our back-transliterated datasets, as the default renderer is unable to correctly render additional characters and diacritics. The renderer parameters can be seen in Appendix C.

#### 3.4 Evaluation

We use image transliteration as the evaluation task, which directly assesses script-tokenizer alignment by measuring surface-form preservation with minimal confounding variables. We report ChrF++ (Popović, 2017) in the main works, alongside BLEU (Papineni et al., 2002) and Word Error Rate in percentage (WER) in the Appendix (see Appendix E).

<sup>3</sup>Weights taken from ernie-research/DualGPT

<sup>4</sup>Heuristic transliteration for Javanese, Balinese, and Sundanese; and a custom font for Lampung.

### 3.5 Model Setup

Models are trained on 4x A40 GPUs with a maximum training time of 24 hours and early stopping (patience = 5) enabled. Training and finetuning hyperparameters are detailed in Appendix A.

### 3.6 Tokenizer Statistics

Language	Llama 2	Custom	Inflation
Javanese <sup>†</sup>	45.97	65.11	+41.6%
Sundanese <sup>‡</sup>	39.41	61.42	+55.8%
Balinese <sup>†</sup>	130.04	130.86	+0.6%
Lampung <sup>‡</sup>	3.73	4.21	+12.8%

Table 2: Average text token length on train data. <sup>†</sup>Processed using the Javanese tokenizer. <sup>‡</sup>Processed using the Sundanese tokenizer.

Table 2 and Figure 2 show that the Llama 2 tokenizer produces shorter sequences with lower fertility rates than custom tokenizers. The custom Javanese and Sundanese tokenizers achieve near-0% Out-of-Vocabulary (OOV) rates for their target languages but exhibit 22% and 10% OOV rates on Balinese and Lampung, respectively. The Llama 2 tokenizer maintains 0% OOV across all languages due to its broader vocabulary coverage.

By conventional metrics, sequence length, fertility rate, and OOV rate, the Llama 2 tokenizer appears superior. However, as we demonstrate in Section 4, these efficiency metrics do not translate into real model performance on low-resource scripts.

## 4 Results

### 4.1 VLM Evaluation

Five-shot in-context prompting reveals that current VLMs fail to transliterate rendered Indonesian




Figure 2: Fertility rate on the evaluation dataset per language. Bali<sup>†</sup> and Lampung<sup>‡</sup> uses Javanese’s and Sundanese’s tokenizer, respectively.

Effect of Tokenizer Choice in Monolingual Finetuning (DualGPT)

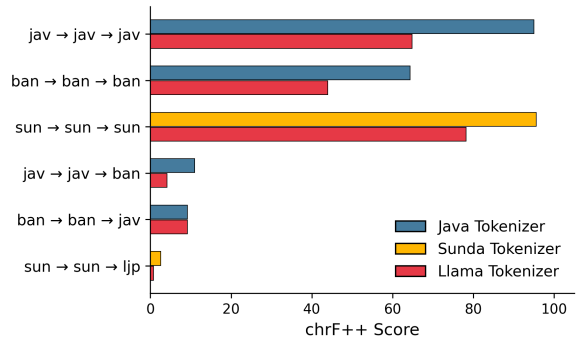


Figure 3: Impact of tokenizer choice on **Monolingual** DualGPT pretraining and finetuning: Pretraining → Finetuning → Evaluation.

scripts (Figure 1). While Llama-3.1-Nemotron-Nano-VL-8B-V1 achieves the highest chrF++ among tested models, its performance is negligible, with an average BLEU of 1.64 and a WER of 1,145.55 (many more errors than words in the reference). We report the full generation setup and results in Appendix D.

### 4.2 DualGPT Architecture

Custom tokenizers significantly outperform Llama 2 in monolingual settings (Figure 3), with chrF++ improvements of +30.15 for Javanese, +20.45 for Balinese, +17.40 for Sundanese, and +1.8 for Lampung (see Appendix E for complete metrics). However, zero-shot cross-lingual transfer fails for both tokenizers (e.g., Java→Bali only achieves 10.94 chrF++ with custom vs. 4.15 with Llama 2), indicating that improved in-language alignment does not enable cross-lingual generalization.

In multilingual training (Table 3), custom tokenizers maintain their advantage. Joint Java+Bali training with the Javanese tokenizer achieves 69.2 chrF++, more than 6 times Java→Bali monolingual training performance, suggesting that multilingual exposure improves cross-lingual transfer, though high WER (108.73%) indicates continued word-level challenges.

Crucially, the Llama 2 tokenizer’s 0% OOV rate did not translate to superior performance. Despite having significantly higher OOV rates in cross-lingual settings (Section 3.6), our custom tokenizers maintained a substantial lead, suggesting that visual features do not compensate for poorly aligned text embeddings.

This script-tokenizer alignment problem is also evidenced by the Javanese and Balinese performance mismatch, although to a lesser extent than

Pretrain Data	Finetune Data	Eval Data	Tokenizer	ChrF++ $\uparrow$	WER $\downarrow$	BLEU $\uparrow$
Javanese	Javanese + Balinese	Javanese	Llama Java	62.52 <b>94.27</b>	125.42 <b>21.79</b>	16.18 <b>78.05</b>
		Balinese	Llama Java	42.29 <b>64.11</b>	<b>95.65</b> 135.62	6.51 <b>7.92</b>
Balinese	Javanese + Balinese	Javanese	Llama Java	62.48 <b>87.93</b>	<b>74.84</b> 90.60	23.88 <b>42.30</b>
		Balinese	Llama Java	42.64 <b>62.40</b>	<b>105.35</b> 285.08	<b>5.98</b> 4.40
Java + Balinese	Javanese + Balinese	Javanese	Llama Java	64.06 <b>94.96</b>	127.51 <b>16.85</b>	17.33 <b>81.30</b>
		Balinese	Llama Java	44.89 <b>69.20</b>	159.35 <b>108.73</b>	5.11 <b>10.18</b>
Sundanese	Sunda + Lampung	Sundanese	Llama Sunda	78.35 <b>96.02</b>	40.16 <b>5.33</b>	59.33 <b>92.45</b>
		Lampung	Llama Sunda	2.95 <b>9.67</b>	194.12 <b>159.85</b>	0.05 <b>0.13</b>

Table 3: Image transliteration performance comparing Llama and Custom tokenizers under a multilingual training setup. **Bolded** values indicate the best performance across tokenizers for each metric within each evaluation language. We use chrF++ as the main metrics for analysis (see Appendix F).

the Llama 2 tokenizer. Despite their shared Austronesian roots, Romanization conventions, and similar orthographic structures, the Javanese and Balinese models exhibit a significant chrF++ performance mismatch. Analyzing the vocabulary overlap, we find that a grapheme-based tokenizer trained on the Balinese data only has an overlap of 16.4% with the used Javanese tokenizer (4,994 shared tokens out of 30,346 unique tokens). This indicates that the Balinese’s performance mismatch stems from embedding space mismatch caused by the script-tokenizer misalignment rather than mere data scarcity.

These results demonstrate that tokenizer alignment critically affects performance, contradicting traditional efficiency metrics (Section 3.6): the Llama 2 tokenizer’s superior fertility rate and zero OOV did not translate to better transliteration.

## 5 Discussion and Future Works

**Tokenizer metrics are not indicative of tokenizer fit** While standard metrics (Section 3.6) suggest the Llama 2 tokenizer is superior for all four languages, each model performs significantly better with the custom tokenizer. We thus urge the creation of evaluation frameworks that better align tokenizer quality with model performance for both monolingual and multilingual cases (Chelombitko et al., 2024).

**Text tokenizer remains a bottleneck despite pixel-based modeling** One might expect that pixel-based architectures would not require a robust text tokenizer. However, we show that the text tokenizer remains a crucial bottleneck. While this work focuses on pixel-based models, current LLMs likely face similar issues with low-resource scripts, even when Romanized. We require further analysis on the impact of tokenizer choice on multilingual LLM performance (Limisiewicz et al., 2023).

**Cross-lingual transfers are still limited by the tokenizer** Our analysis indicates that poor performance stems not only from data scarcity but from tokenizer mismatch. Despite Javanese and Balinese being linguistically highly related, cross-lingual transfer remains ineffective. This suggests that as long as the tokenizer issue persists, low-resource languages may benefit more from specialized models than from cross-lingual transfer. Advancements in architectures that balance specialization with generalization are vital for equitable performance.

**Leveraging Rendered Representations** Despite our criticism, we believe pixel-based language modeling is a promising path toward more equitable technology. This work acts as a warning, reintroducing text tokenizers into such architectures may bring unintended consequences and must be done with proper justification and extreme care.

## 271 Limitations

272 **Scope of Architectural Analysis.** This work fo-  
273 cuses specifically on the impact of tokenizer align-  
274 ment within hybrid vision-text architectures (e.g.,  
275 DualGPT). While a comparison with pure vision-  
276 based models (without text-heads) would provide  
277 additional context, our study is designed to iso-  
278 late the behavior of the text-modality when reintro-  
279 duced into pixel-based frameworks. Future work  
280 will expand this to a wider range of architectural  
281 configurations.

282 **Resource Disparity.** The available training data  
283 for Javanese, Balinese, Sundanese, and Lampung  
284 varies significantly, which is reflected in the perfor-  
285 mance gap observed in the Lampung experiments.  
286 However, the consistent improvement across all  
287 languages when using aligned tokenizers suggests  
288 that our findings are robust to variations in data  
289 scale.

290 **Task Specificity.** We utilize image transliteration  
291 as a primary probe because it directly measures  
292 surface-form preservation and script-alignment  
293 with minimal confounding variables. While this  
294 task effectively exposes the tokenizer bottleneck,  
295 further research is needed to determine how these  
296 alignment issues affect higher-level semantic tasks  
297 like NLI or abstractive summarization in pixel-  
298 based models.

## 299 Acknowledgment

300 Anonymized due to double blind.

## 301 Ethical Consideration

302 Despite rendering native scripts as images, we still  
303 use the Romanized form to simplify the analysis in  
304 this work. This approach is not meant to supplant  
305 the use of native scripts with Latin texts, but is out  
306 of necessity, as many local Indonesian languages’  
307 scripts are not supported by Unicode. Further ad-  
308 vancements in this line of research are pivotal to  
309 overcoming this digital infrastructure gap along-  
310 side improving language and script equity in the  
311 field of AI.

## 312 References

313 Diana Abagyan, Alejandro R. Salamanca, Andres Felipe  
314 Cruz-Salinas, Kris Cao, Hangyu Lin, Acyr Locatelli,  
315 Marzieh Fadaee, Ahmet Üstün, and Sara Hooker.

2025. [One tokenizer to rule them all: Emergent lan-  
guage plasticity via multilingual tokenizers.](#) *Preprint*,  
arXiv:2506.10766. 316  
317  
318

Muhammad Farid Adilazuarda, Musa Izzanardi 319  
Wijanarko, Lucky Susanto, Khumaisa Nur’aini, 320  
Derry Tanti Wijaya, and Alham Fikri Aji. 2025. 321  
[NusaAksara: A multimodal and multilingual bench-  
mark for preserving Indonesian indigenous scripts.](#) 322  
In *Proceedings of the 63rd Annual Meeting of the  
Association for Computational Linguistics (Volume 1:  
Long Papers)*, pages 28371–28401, Vienna, Austria. 323  
Association for Computational Linguistics. 324  
325  
326  
327

Shuai Bai and 1 others. 2025. [Qwen3-vl technical re-  
port.](#) *Preprint*, arXiv:2511.21631. 328  
329

Mohammad Jahid Ibna Basher, Mohammad Raghif 330  
Noor, Sadia Afroze, Ikbal Ahmed, and Mo- 331  
hammed Moshul Hoque. 2023. [Bngraphemizer:  
A grapheme-based tokenizer for bengali handwrit-  
ten text recognition.](#) *2023 IEEE 9th International  
Women in Engineering (WIE) Conference on Elec-  
trical and Computer Engineering (WIECON-ECE)*,  
pages 183–188. 332  
333  
334  
335  
336  
337

Akhiaid Bercovich and 1 others. 2025. [Llama-  
nemotron: Efficient reasoning models.](#) *Preprint*,  
arXiv:2505.00949. 338  
339  
340

Yekun Chai, Qingyi Liu, Jingwu Xiao, Shuohuan Wang, 341  
Yu Sun, and Hua Wu. 2024. [Autoregressive pre-  
training on pixels and texts.](#) In *Proceedings of the  
2024 Conference on Empirical Methods in Natu-  
ral Language Processing*, pages 3106–3125, Miami,  
Florida, USA. Association for Computational Lin- 342  
guistics. 343  
344  
345  
346  
347

Iaroslav Chelombitko, Egor Safronov, and Aleksey 348  
Komissarov. 2024. [Qtok: A comprehensive frame-  
work for evaluating multilingual tokenizer quality in  
large language models.](#) *Preprint*, arXiv:2410.12989. 349  
350  
351

Jonathan H. Clark, Dan Garrette, Iulia Turc, and John 352  
Wieting. 2022. [Canine: Pre-training an efficient  
tokenization-free encoder for language representa-  
tion.](#) *Transactions of the Association for Computa-  
tional Linguistics*, 10:73–91. 353  
354  
355  
356

Negar Foroutan, Clara Meister, Debjit Paul, Joel 357  
Niklaus, Sina Ahmadi, Antoine Bosselut, and Rico 358  
Sennrich. 2025. [Parity-aware byte-pair encoding:  
Improving cross-lingual fairness in tokenization.](#)  
*Preprint*, arXiv:2508.04796. 359  
360  
361

Gemma Team. 2025. [Gemma 3 technical report.](#)  
*Preprint*, arXiv:2503.19786. 362  
363

Wenhan Han, Yifan Zhang, Zhixun Chen, Binbin Liu, 364  
Haobin Lin, Bingni Zhang, Taifeng Wang, Mykola 365  
Pechenizkiy, Meng Fang, and Yin Zheng. 2025. 366  
[Mubench: Assessment of multilingual capabilities of  
large language models across 61 languages.](#) *Preprint*,  
arXiv:2506.19468. 367  
368  
369



Category	Hyperparameter	Value
Model Setup	Pretrained Weight	DualGPT-pt
	Dropout	0.1
	Frozen Layers	lm_pixel_head
Optimization	Learning Rate	$2 \times 10^{-5}$
	LR Scheduler	Cosine
	Warmup Ratio	0.03
	Weight Decay	0.1
	Numerical Precision	BF16
Training	Per-device Batch	2
	Grad. Accumulation	4
	Num. GPUs (A40)	4
	Effective Batch Size	32
	Max Epochs	10
	Early Stop Patience	5
	Eval / Save Steps	1000

Table 5: Hyperparameters for image-to-text transliteration fine-tuning. We swap **Warmup steps** to **Warmup ratio** to account for Lampung finetuning data (945 entries). DualGPT-pt indicates the pretrained DualGPT model.

## B Dataset Construction

Figure 4 illustrates the dataset construction pipeline used for DualGPT training. Starting from romanized text, each word is back-transliterated into Indonesian scripts (Aksara), which is then tokenized using a grapheme-based scheme to preserve script-specific character structures. The Aksara text is rendered into an image using a custom renderer, forming the visual modality, while the tokenized Aksara is transliterated back into Latin and converted into token IDs using a custom tokenizer. The resulting paired image–token representation constitutes the final training dataset.

## C Renderer Parameters

We follow the original PixelGPT renderer parameters, modifying only the font and font size to prevent text from overflowing beyond the image boundaries as shown in Table 6. Lampung language used **custom font** since it is not inherently supported by Unicode.

## D VLM Evaluation Setups and Results

We use open small-sized to medium-sized VLMs, including Gemma-3-12B-it (Gemma Team, 2025), Qwen3-VL-8B-Instruct (Bai et al., 2025), InternVL3.5-8B (Wang et al., 2025), and Llama-3.1-Nemotron-Nano-8B-V1 (Bercovich et al., 2025), as well as a closed-source model, GPT-5-nano (OpenAI, 2025). Few-shot examples

Parameter	jav/sun/ban	ljp
Background Color	White	White
DPI	120	120
Font Color	Black	Black
Font Type	Noto Sans	Had Lampung Yuzu Rounded
Font Size	7	10
Max sequence length	1024	1024
Padding size	3	3
Pixels per patch	$16 \times 16$	$16 \times 16$

Table 6: Configuration of text rendering for Javanese, Sundanese, Balinese, and Lampung language.

are randomly selected from pretrain data to avoid leakage. For the open models, we deploy them using vLLM (Kwon et al., 2023). For the generation hyperparameters, we constrain temperature = 0.4 and max\_output\_tokens = 2048. Full VLMs evaluation performance is available at Table 8.

## E Full Monolingual Model Performance

Full model monolingual performance is available in Table 9.

## F chrF++ as the Main Metric

Due to the nature of languages we are working with (i.e., low resource, native scripts), finding a suitable dataset for the evaluation remains a significant challenge. NusaAksara exist as one of the only benchmark available to support our transliteration task. However, NusaAksara has limitations. Shown in Table 7, NusaAksara’s Balinese entry consist mainly of short word entries, disadvantaging metrics such as BLEU and WER.

ID	Reference	Prediction
7	prathiwitala	prathiwitala
9	mandaaakranta	mantaaakranba
11	134 ,	13 ,
13	dumala	meralumala

Table 7: Transliteration sample for Balinese language. ID refers to NusaAksara’s dataset entry ID.

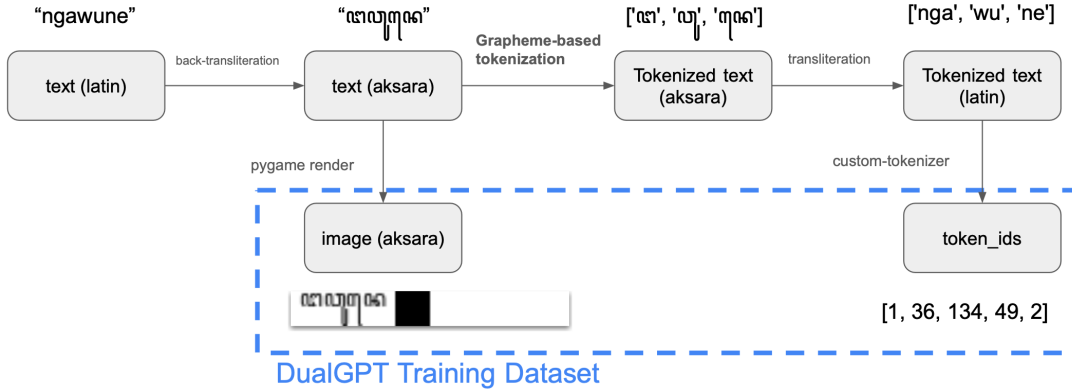


Figure 4: Dataset building process using custom tokenization rules.

Model	Balinese			Javanese			Lampung			Sundanese		
	BLEU $\uparrow$	chrF++ $\uparrow$	WER $\downarrow$	BLEU $\uparrow$	chrF++ $\uparrow$	WER $\downarrow$	BLEU $\uparrow$	chrF++ $\uparrow$	WER $\downarrow$	BLEU $\uparrow$	chrF++ $\uparrow$	WER $\downarrow$
InternVL3.5-8B	0.00	9.18	564.94	0.00	17.10	1947.98	0.00	7.92	625.11	6.57	19.08	1444.19
Llama-3.1-Nemotron-Nano-VL-8B-V1	4.99	<b>19.44</b>	794.22	0.52	8.82	907.54	<b>1.09</b>	<b>11.83</b>	2086.88	7.81	<b>30.91</b>	1200.58
Qwen3-VL-8B	1.62	14.21	238.35	2.91	16.50	316.89	0.00	0.00	1447.06	<b>8.12</b>	15.34	710.77
Gemma-3-12B-it	<b>6.57</b>	15.26	295.63	<b>3.83</b>	<b>24.15</b>	625.47	0.00	6.01	225.57	5.52	13.20	519.27
GPT-5-Nano	0.00	0.00	<b>205.42</b>	0.00	0.00	<b>161.94</b>	0.00	9.36	<b>101.81</b>	0.00	0.00	<b>123.71</b>

Table 8: VLM Zero-shot image transliteration performance across four Indonesian languages. Higher is better for BLEU and chrF++; lower is better for WER. Note that extremely low WER in models with zero BLEU (e.g., GPT-5-Nano) typically indicates empty or extremely short output rather than accuracy.

Pretrain & Finetune	Eval	Tokenizer	ChrF++ $\uparrow$	WER $\downarrow$	BLEU $\uparrow$
Javanese	Javanese	Llama	64.83	79.01	24.99
		Java	<b>94.98</b>	<b>19.95</b>	<b>80.28</b>
Balinese	Balinese	Llama	4.10	<b>506.99</b>	0.01
		Java	<b>10.94</b>	550.00	0.01
Balinese	Javanese	Llama	<b>9.20</b>	250.00	0.03
		Java	9.16	<b>102.30</b>	0.03
Balinese	Balinese	Llama	43.91	<b>86.16</b>	<b>7.53</b>
		Java	<b>64.36</b>	131.91	7.00
Sundanese	Sundanese	Llama	78.19	40.54	58.93
		Sunda	<b>95.59</b>	<b>7.87</b>	<b>90.55</b>
Sundanese	Lampung	Llama	0.77	<b>325.79</b>	0.00
		Sunda	<b>2.57</b>	585.64	0.00

Table 9: Image transliteration performance comparing Llama and Custom tokenizers under a monolingual training setup. Bold values indicate which tokenizer performs better.