

Baseer: A Vision-Language Model for Arabic Document-to-Markdown OCR

Khalil Hennara
Misraj AI
Khobar, Saudi Arabia
hennara@misraj.ai

Muhammad Hreden
Misraj AI
Khobar, Saudi Arabia
hreden@misraj.ai

Mohamed Motasim Hamed
Misraj AI
Khobar, Saudi Arabia
hamed@misraj.ai

Ahmad Bastati
Misraj AI
Khobar, Saudi Arabia
bastati@misraj.ai

Zeina Aldallal
Misraj AI
Khobar, Saudi Arabia
aldallal@misraj.ai

Sara Chrouf
Misraj AI
Khobar, Saudi Arabia
chrouf@misraj.ai

Safwan AlModhayan
Misraj AI
Khobar, Saudi Arabia
safwan@misraj.ai

Abstract

Arabic Optical Character Recognition (OCR) remains a challenging task due to the language’s cursive script, diverse fonts, diacritics, and right-to-left orientation. While modern Multimodal Large Language Models (MLLMs) have advanced document understanding for high-resource languages, their performance on Arabic remains limited. In this work, we introduce **Baseer**, a vision-language model fine-tuned specifically for Arabic document OCR. Leveraging a large-scale dataset combining synthetic and real-world documents, Baseer is trained using a decoder-only fine-tuning strategy to adapt a pre-trained MLLM while preserving general visual features. We also present Misraj-DocOCR, a high-quality, expert-verified benchmark designed for rigorous evaluation of Arabic OCR systems. Our experiments show that Baseer consistently outperforms existing open-source and commercial solutions on the benchmarks evaluated in this work, achieving state-of-the-art results for Arabic document OCR under these settings. Our results highlight the benefits of domain-specific adaptation of general-purpose MLLMs and establish a strong baseline for high-accuracy OCR on morphologically rich languages like Arabic.

1 Introduction

The recent and rapid advancements in Multimodal Large Language Models (MLLMs) have fundamen-

tally reshaped the landscape of how machines perceive and process complex visual and textual data (Hurst et al., 2024; Comanici et al., 2025; Zhu et al., 2025; Bai et al., 2025). Among the myriad applications of these models, OCR and comprehensive document understanding continue to present significant challenges. This is particularly true for languages that are morphologically rich and structurally complex, such as Arabic. While contemporary OCR solutions have achieved remarkable performance for English and other high-resource languages (Comanici et al., 2025; Hurst et al., 2024), their efficacy does not readily generalize to Arabic documents. The inherent complexities of Arabic script, including its cursive nature, extensive ligature formation, the wide variety of fonts and styles, the critical role of diacritics, and the right-to-left text orientation, render Arabic OCR a task of considerable difficulty.

In parallel, progress in multimodal architectures has paved the way for unified vision-language reasoning, which enables models to concurrently extract both textual content and structural information from documents (Li et al., 2025; Mandal et al., 2025). Despite these technological strides, modern multimodal frameworks have seldom been specialized for the distinct demands of Arabic OCR and document parsing. This significant gap in research and development leaves academics, practitioners, and industries without robust, dedicated tools for processing real-world Arabic documents, which are prevalent across academic, commercial, and cultural heritage domains.

***Baseer**: meaning “one who sees clearly” and “insightful.” The name reflects the model’s ability to “see” and interpret documents with clarity.

In this work, we introduce **Baseer**[†], a vision-language model fine-tuned for Arabic document OCR. Our approach adapts a powerful general-purpose MLLM to the unique challenges of Arabic document analysis. To facilitate this specialization, Baseer was trained on a large-scale, diverse dataset composed of both synthetically generated and authentic real-world Arabic documents. This dataset was curated to encompass the extensive variety of formats, fonts, and layouts encountered in practical applications. Furthermore, we present **Misraj-DocOCR**, a novel benchmark specifically engineered for the evaluation of Arabic OCR systems, featuring high-quality, expert-verified annotations to ensure reliability.

Our primary contributions are threefold:

1. We present the development and fine-tuning of **Baseer**, demonstrating that an efficient, decoder-only fine-tuning strategy can achieve state-of-the-art performance on the Arabic document OCR benchmarks evaluated in this work.
2. We introduce **Misraj-DocOCR***, a new, reliable, and openly available benchmark designed to provide a standardized and rigorous evaluation framework for Arabic OCR systems.
3. We conduct a thorough analysis of the **KITAB-pdf-to-markdown**[‡] benchmark, providing a revised and improved version that addresses significant inaccuracies to enhance its accuracy and utility for the research community.

The remainder of this paper is organized as follows: Section 2 reviews related work, while Section 3 describes the dataset. Section 4 introduces our new benchmark, followed by Section 5, which details the proposed methodology. Finally, Section 6 outlines the experimental setup, and Section 7 presents the results and analysis.

2 Related Work

We situate our work at the intersection of two major research domains. First, we review the rapid advancements in Multimodal Large Language Models

(MLLMs), which provide the architectural foundation for our approach. Second, we delve into the field of Optical Character Recognition (OCR) and Document Understanding, examining its evolution and highlighting the persistent challenges that motivate our research, particularly for morphologically complex languages like Arabic.

2.1 Multimodal Large Language Models

The paradigm of Large Language Models (LLMs) has recently been extended to handle multimodal input, leading to the development of powerful models capable of joint vision-language reasoning. Research in this area has generally progressed along two main architectural paths. One approach involves the modular integration of pre-trained components, where a specialized frozen vision encoder is connected to a large language decoder via a lightweight adapter. This design is seen in influential models like LLaVA (Liu et al., 2023, 2024; Li et al., 2024), Aya-Vision (Dash et al., 2025), Idefics (Laurençon et al., 2024b,a), and more compact architectures such as SmolVLM (Marafioti et al., 2025). These architectures achieve impressive zero-shot and few-shot performance on a diverse range of multimodal tasks with high parameter efficiency. The second approach focuses on training massive, end-to-end vision-language models. This category includes state-of-the-art systems such as InternVL (Chen et al., 2024; Zhu et al., 2025), Gemma (Team et al., 2025), PaliGemma (Steiner et al., 2024), and Qwen-VL (Bai et al., 2025). These models, with parameter counts scaling up to 70 billion in the case of Qwen2.5-VL, have demonstrated remarkable general-purpose capabilities. However, their broad, generalist training often leaves them unspecialized for precision-critical, niche domains. As we will discuss, high-fidelity document OCR represents one such domain where these powerful models still exhibit significant limitations.

2.2 OCR and Document Understanding

Leading efforts in this domain, such as Idefics3 (Laurençon et al., 2024a), Monkey-OCR (Li et al., 2025), SmolDocling (Nassar et al., 2025), and commercial systems such as Nanonets-OCR-s (Mandal et al., 2025), have established high benchmarks for performance on standard document types. More recent attempts, such as Qari (Wasfy et al., 2025), have tackled Arabic OCR directly, but their scope remains

[†]<https://baseerocr.com/>

^{*}<https://huggingface.co/datasets/Misraj/Misraj-DocOCR>

[‡]https://huggingface.co/datasets/Misraj/KITAB_pdf_to_markdown_reviewed

limited compared to comprehensive document understanding systems.

A small but growing body of work targets Arabic document OCR and the resources needed to train it. Arabic-Nougat (Rashad, 2024) adapts the Nougat architecture to convert Arabic book pages into Markdown and releases *arabic-img2md*, a synthetic image–Markdown corpus of roughly 13.7K pages; its focus, however, is on running text rather than the full range of structured elements (e.g., HTML tables) targeted here. SARD (RIOTU Lab, 2025) contributes a large-scale (~843K pages) synthetic Arabic corpus that systematically varies fonts and book-style layouts, but by construction it is noise-free and omits real-world scans, embedded tables, and multi-column layouts. QARI-OCR (Wasfy et al., 2025) similarly releases synthetic Arabic OCR resources and a fine-tuned model. Baseer is complementary to and broader than these efforts along two axes: (i) its training mixture pairs synthetic data with a large set of *expert-verified real-world* pages and dedicated table-centric and non-text data, and (ii) it explicitly models structured output. We position *Misraj-DocOCR* against these resources accordingly: whereas KITAB-Bench (Heakl et al., 2025), *arabic-img2md*, and SARD are dominated by synthetic or automatically extracted ground truth, *Misraj-DocOCR* is a fully expert-verified benchmark that mixes real and synthetic pages and covers structured layouts, making it a stricter test of layout-aware transcription.

However, a critical challenge remains: the generalization of these models to languages with scripts that are fundamentally different from Latin-based languages. Arabic serves as a prime example of this challenge. Its inherent characteristics include cursive script, context-sensitive character shapes, optional but meaningful diacritics, right-to-left orientation, and a wide variety of fonts and styles. These characteristics often cause state-of-the-art document OCR systems to degrade sharply in performance when applied to Arabic texts.

To the best of our knowledge, the application of modern MLLM frameworks to the specific, challenging problem of Arabic document OCR remains a largely unexplored area. Although several multilingual and multimodal models include Arabic in their training, they are not optimized for the script-specific and structural challenges posed by Arabic documents. This work aims to bridge this critical gap by introducing *Baseer*, a system specif-

ically engineered for the complexities of Arabic documents. Our results demonstrate that this specialized approach yields substantial improvements over general-purpose baselines on the benchmarks we evaluate.

3 Data

A primary motivation for our work is the severe scarcity of high-quality Arabic data for document OCR, a challenge that firmly positions Arabic as a low-resource language within this specific research domain. To overcome this critical bottleneck and address the shortcomings of existing resources, we dedicated substantial effort to sourcing, curating, and annotating an original, high-quality dataset tailored for complex document understanding.

To support effective document OCR, it is essential to represent textual content in a format that preserves both structure and semantics. In our dataset, the text corresponding to each image is formatted in Markdown, providing a clean and standardized representation of content. Tables are represented in HTML to accurately capture diverse table structures and complex layouts. Furthermore, specialized tags were introduced to mark specific elements within the text including watermarks, page numbers, and the presence of images enabling precise supervision for layout-aware OCR and document parsing tasks.

The dataset itself was constructed as a hybrid collection, combining a large corpus of synthetically generated documents with a carefully curated set of real-world publications, supplemented by targeted data covering tables, and non-text imagery. Each of these sources is described in detail below.

3.1 Synthetic Data

The first component of our dataset was generated synthetically using an in-house pipeline, designed to capture the diverse formatting and layout variations commonly found in word-processing documents. The foundation for this synthetic data is a corpus of markdown-formatted documents, which were downloaded and filtered from the Common Crawl archive. To ensure the quality and relevance of the source material, the raw data were subjected to the following preprocessing filters:

1. **Perplexity Filtering:** An in-house language model based on KenLM (Heafield, 2011) was employed to calculate perplexity scores, re-

taining only the most linguistically cohesive text samples.

2. **Table Sparsity Filtering:** To ensure structural integrity, documents containing markdown tables with more than 25% empty cells were identified and discarded.

The filtered markdown documents were then converted into image-text pairs. From the pool of generated images, 200,000 samples were randomly selected and underwent random transformations to consist an augmentation data. To prevent redundancy, the original, pre-augmentation versions of these images were discarded. In total, this synthetic pipeline produced **400,000** high-quality image-text pairs, comprising 200,000 clean rendered samples and 200,000 augmented variants designed to simulate diverse real-world document conditions.

3.2 Open-Source Books and Magazines

The second component of our dataset was sourced from real-world documents, including a diverse collection of books, magazines, educational documents, and academic papers. In contrast to synthetic data, these samples reflect authentic publishing environments, capturing genuine layout complexities and typographic conventions.

To ensure maximum diversity, the selected pages span a broad spectrum of layout complexities, identified using vision-based algorithms. Specifically, bounding boxes were first detected at the paragraph level, and their alignment and overlap were analyzed to capture challenging structures such as tables, figures, index pages, and skewed layouts. Page color distributions were also examined to include samples with embedded images, colorful backgrounds, and multi-colored text.

Ground-truth text for the real-world documents was obtained using a state-of-the-art vision-language model, *Gemini-2.5-pro*. To ensure the highest level of fidelity in our training data, we implemented a rigorous human-in-the-loop validation stage following the initial VLM extraction. For each sample, one of 20 trained annotators performed a side-by-side review of the original document image and the model-generated output. In this phase, the annotator’s role was strictly binary: to either *accept* or *reject* the sample based on the presence of any errors in the extracted text or structural formatting.

Given our commitment to high-quality supervision, we maintained a strict threshold for accuracy;

consequently, approximately 35% of the processed data was rejected due to inconsistencies or extraction errors. While these rejected samples were excluded from the current training cycle to maintain data purity, they represent a valuable resource for future work. In subsequent research, we plan to implement a corrective pipeline to refine these samples rather than discarding them, further expanding the dataset and iteratively improving the model’s robustness. From this source, we curated **257,000** document images paired with their corresponding ground-truth text.

3.3 Table-Centric Data

To enhance tabular recognition and reconstruction, we developed a dedicated dataset of **111,000** samples using two complementary pipelines: First, *Template-Based Generation from Real Financial Documents*, Human-curated Jinja templates derived from real-world financial documents were populated with 20 unique LLM-generated instances each. A diversity-aware prompting strategy was employed to minimize redundancy. Second, *Translation and Augmentation of HTML Tables*, structured HTML tables were translated into Arabic using 12 stylistic prompts. These variations included full content translation, preservation of technical terms, and conversion between Western and Arabic numerals to reflect real-world document diversity.

3.4 Non-Text Images

To mitigate hallucination tendencies and reduce the model’s bias toward assuming that every input image contains extractable text, we incorporated a set of **~1,500** images that contain no textual content. This subset includes blank pages, natural scenes, and other non-document imagery. Exposing the model to such samples during training encourages it to correctly produce an empty output when presented with text-free inputs, rather than fabricating non-existent content.

Collectively, the four data sources described above yield a total of approximately **769,000** image-text pairs used for training our model. A detailed breakdown of the dataset distribution across sources is shown in Figure 1.

4 Misraj-DocOCR: An Arabic Document OCR Benchmark

The evaluation of OCR models for Arabic text requires robust and accurate benchmarks. Our ini-




Figure 1: Distribution of training samples across the four data sources (synthetic, real-world books and magazines, table-centric, and non-text), totalling approximately 769K image–text pairs. The synthetic portion is split into clean rendered pages and their augmented variants.

tial investigation involved assessing existing benchmarks, such as the KITAB-bench (Heakl et al., 2025) pdf-to-markdown dataset. During this analysis, we identified significant shortcomings that compromise its reliability for model evaluation. A primary issue discovered was the presence of numerous errors in the ground truth data. We observed multiple instances of hallucinatory text, where the ground truth contained phrases not present in the source documents likely originating from a data creation or annotation tool rather than authentic content[‡]. Furthermore, our review revealed that many examples lacked corresponding page numbers, and small-font text was frequently omitted from the ground truth. These inaccuracies suggest that the dataset may not have undergone a thorough verification process after the initial data extraction. To address these deficiencies and provide a more reliable resource for the research community, we undertook a comprehensive correction of the KITAB-bench PDF-to-markdown dataset. This

corrected version, with all identified errors rectified, has been made publicly available for academic use[§].

Beyond the inaccuracies, our examination of existing resources also indicated a lack of diversity in the style and type of documents. To foster more generalized and robust model development, a benchmark should encompass a wide variety of real-world scenarios. Therefore, we introduce *Misraj-DocOCR*, a new, comprehensive benchmark specifically designed for evaluating Arabic Document OCR models.

The primary contributions of this benchmark are:

- **Diverse and Comprehensive Content:** The benchmark consists of 400 high-quality images, curated to include a wide variation of document types, layouts, and fonts, and comprising both synthetic (150) and real-world (250) pages. The selection deliberately includes challenging conditions such as small-font and low-contrast text, dense tables, multi-column layouts, and diacritized passages, so

[‡]For example, one entry included the English sentence: "You're right - let me write it exactly as it appears in the image, maintaining the right-to-left direction:"

[§]https://huggingface.co/datasets/Misraj/KITAB_pdf_to_markdown_reviewed

that the benchmark stresses layout-aware transcription rather than clean single-column text alone.

- **Expert-Verified Ground Truth:** To ensure the highest level of accuracy, every image in the dataset has been meticulously reviewed by human experts. This verification process guarantees that both the transcribed text and the document structure are correct, eliminating the types of errors found in previous benchmarks.
- **Open Access:** Misraj-DocOCR is open-source and publicly available to all researchers. By providing this resource, we aim to facilitate further advancements and foster reproducible research in the field of Arabic OCR.

We evaluate many models on this benchmark and the corrected version of KITAB-Bench, all results on the section 7.

5 Methodology

The overall process for developing Baseer, involved a comprehensive data collection stage followed by a targeted fine-tuning stage.

The development of our model, Baseer, followed a two-stage methodology designed to tailor a powerful, pre-trained foundation model to our specific needs. The first stage involved the comprehensive collection and curation of a high-quality dataset, the details of which are described in Section 3. The subsequent stage, which is the focus of this section, consisted of fine-tuning the selected base model to align with our data and enhance its capabilities for Arabic document processing. For the base architecture of Baseer, we selected the *Qwen2.5-VL-3B-Instruct* model (Bai et al., 2025). This decision was predicated on its robust and state-of-the-art performance on multimodal tasks, particularly its demonstrated proficiency with the Arabic language compared to other open-source alternatives.

Despite its advanced capabilities, our preliminary analysis revealed that the base Qwen2.5-VL-3B-Instruct model exhibited certain limitations relevant to our use case. These included occasional reversion to left-to-right text generation, suboptimal handling of diacritized Arabic text, and other performance artifacts. Our fine-tuning strategy involved updating all model parameters, except for

the vision encoder, which remained frozen. This approach allows the model to adapt its language and reasoning capabilities to our specialized dataset while preserving the powerful, generalized visual features learned during its original pre-training. The specific hyperparameters, hardware used, and other details of the training procedure are provided in Table 1.

Parameter	Value
Training Epochs	3
Learning Rate Schedule	Cosine decay
Learning Rate	1e-4
Batch Size	640
Weight Decay	0.01
Warm-up Steps	100
Optimizer	AdamW
Max Sequence Length	8192
GPU	8xH100

Table 1: Training Hyperparameters for *Baseer* Model

6 Experiments and Results

This section details the series of experiments conducted to systematically determine the optimal architecture and training configuration for Baseer. Our experimental process was designed to isolate variables and build upon the findings of each preceding stage.

6.1 Base Model Selection

The initial experiment was focused on selecting the most suitable base model for our task. To this end, we conducted a qualitative evaluation of several prominent open-source vision-language models. A curated set of representative examples, designed to test key capabilities in Arabic document understanding, was used as the input.

The outputs from each model were then subjected to a rigorous manual review by our evaluation team. The models were assessed based on criteria such as text recognition accuracy, preservation of right-to-left directionality, and overall coherence. This qualitative analysis concluded that Qwen2.5-VL-3B-Instruct demonstrated superior performance on Arabic-language tasks compared to the other candidates, making it the clear choice for our foundation.

6.2 Fine-Tuning Strategy Evaluation

After selecting the base model, our next objective was to identify the most effective fine-tuning strat-

egy. We designed a controlled experiment to compare three distinct approaches:

1. **Full Fine-Tuning (Baseer-Full):** All model parameters, including the vision encoder, were made trainable.
2. **Decoder-Only Fine-Tuning (Baseer-Decoder):** Only the parameters of the language decoder were updated, while the vision encoder remained frozen.
3. **Parameter-Efficient Fine-Tuning (Baseer-LoRA):** Low-Rank Adaptation (LoRA) was employed to update a small subset of parameters. LoRA adapters were applied to the linear projections of the language decoder (with the vision encoder frozen), using rank $r = 256$, scaling factor $\alpha = 128$.

To ensure a fair comparison, each of these strategies was tested on a 50,000-sample subset of our training data for two epochs, holding all other hyperparameters constant. We evaluated the models using ChrF, which measures OCR accuracy at the character level and captures text transcription quality. Table 2 summarizes the performance of the different fine-tuning strategies on the Baseer model.

Model	Trainable part	ChrF $\uparrow$
Baseer-Full	Full model	84.79
Baseer-Decoder	Language-decoder	89.79
Baseer-LoRA	LoRA weight	85.52

Table 2: Performance comparison of different fine-tuning strategies on Baseer model

As shown in Table 2, the results from our test set indicate that the decoder-only fine-tuning approach (Baseer-Decoder) significantly outperformed the other methods. This suggests that preserving the generalized features of the pre-trained vision encoder while adapting the language model to our specific data yields the best performance.

6.3 Impact of Sequence Length

Building on the previous finding, we adopted the decoder-only fine-tuning strategy and proceeded to investigate the effect of input sequence length on model performance. All training configurations

Context Length	ChrF $\uparrow$
2048	82.69
4096	89.79
8192	88.52

Table 3: Performance comparison of different context lengths on Baseer model.

were fixed while we experimented with three sequence length variants: 2048, 4096, and 8192 tokens.

While the preliminary baseline metrics presented in Table 3 initially suggest an optimal performance at a sequence length of 4096, extensive real-world inference testing revealed a critical limitation. We observed that models constrained to 4096 tokens frequently failed to extract the complete context of highly dense visual data. This bottleneck was especially pronounced for images containing complex tables, which inherently necessitate a significantly larger token budget to accurately generate the exhaustive HTML table tags required for structural fidelity. To ensure comprehensive data extraction, we ultimately trained our final model using the 8192 context length, which substantially mitigates truncation issues on dense inputs and yields more robust results during inference.

7 Evaluation

We evaluate our model on our proposed Misraj-DocOCR benchmark and a corrected version of KITAB-Bench PDF-to-Markdown, alongside several open-source and commercial models. Text extraction performance is assessed using Word Error Rate (WER) and Character Error Rate (CER), which measure word and character-level transcription errors, BLEU for n-gram overlap, and ChrF, a character-level F-score suited for morphologically rich languages like Arabic. Structural and layout fidelity is measured with Tree Edit Distance Similarity (TEDS), capturing hierarchical document structures, and MARS (Heakl et al., 2025), which evaluates layout-aware alignment between predicted and reference renderings. Table 4 presents the results, showing that Baseer achieves state-of-the-art performance across both text and structural metrics, despite being smaller than competing models.

7.1 Evaluation Protocol

To ensure fair comparison across models, models designed for document understanding were evalu-

ated using their respective recommended system prompts, while general-purpose Multimodal Large Language Models (MLLMs) were provided with a single fixed instruction prompt requesting Markdown output with HTML tables. Because prompting can materially affect MLLM behaviour, we report the exact prompt templates used for every model in the Appendix, and all models were run with greedy decoding (temperature 0) to make outputs deterministic and reproducible. All outputs were then standardized using the following post-processing steps:

1. Remove HTML tags outside table structures
2. Convert Markdown tables to HTML format for consistency
3. Normalize horizontal line representations (—, ***, etc. → —)
4. Standardize header formatting
5. Unify formatting tags within HTML tables (, →)

This standardization is critical because different models may produce semantically equivalent but syntactically different outputs, which would unfairly penalize models based on formatting choices rather than content accuracy.

Table 4 presents a comparative evaluation of different OCR and vision-language models using multiple metrics. The results indicate that **Baseer** achieves the best performance across all metrics, attaining the lowest WER and CER, as well as the highest TEDS and MARS scores. The **Gemini-2.5-pro** model follows closely, obtaining the highest BLEU and CHRF scores and ranking second overall. Notably, Baseer consistently outperforms large commercial systems such as GPT-based models and Azure AI Document Intelligence,[¶] underlining its robustness in both text and structure recognition. This is especially significant given that the evaluation dataset, Misraj-DocOCR, was deliberately designed to be highly diverse and challenging, with wide variation in layout and typography. The results also highlight a sharp performance gap between the top-performing systems and smaller or

less specialized models (e.g., Gemma3, AIN, Aya-vision), underscoring the difficulty of this benchmark. Overall, Baseer and Gemini-2.5-pro emerge as the strongest systems in this comparison.

Table 5 reports the results on the KITAB-Bench PDF-to-Markdown dataset, which was carefully reviewed and corrected by domain experts to ensure high-quality ground truth annotations. For this corrected subset we restrict the comparison to open-source systems, so as to provide a fully reproducible, cost-free reference point on a small set of samples; commercial systems are instead evaluated on the larger and more diverse Misraj-DocOCR benchmark (Table 4). **Baseer** achieves the strongest performance across all metrics, attaining the best WER, CER, BLEU, CHRF, TEDS, and MARS scores, with **Dots.ocr** consistently ranking second across text-centric metrics. It is also worth noting that the KITAB-Bench subset is relatively small, consisting of only 30 samples, which makes every misprediction more impactful on the reported scores. In contrast, on the larger and more challenging **Misraj-DocOCR** benchmark with 400 diverse examples, Baseer’s advantage over both open-source and commercial systems becomes more pronounced, highlighting its robustness across varied document types and layouts.

8 Conclusion

In this paper, we introduced Baseer, a vision-language model tailored for Arabic Document OCR, and presented Misraj-DocOCR, a high-quality benchmark designed for rigorous evaluation. By training on a diverse dataset of approximately 770K document-image pairs, we demonstrated that decoder-only fine-tuning is a powerful strategy that enables Baseer to achieve superior performance compared to a wide range of existing systems. Our detailed experimental analysis highlighted the importance of sequence length, fine-tuning scope, and dataset diversity in achieving robust performance. Notably, Baseer consistently achieved the best or near-best scores across Word Error Rate, Character Error Rate, and structure-aware metrics such as TEDS and MARS, surpassing both open-source and proprietary alternatives. These positive results underscore the value of domain-specific adaptation of general-purpose MLLMs, and provide new insights into how tailored data and efficient training strategies can push the boundaries of OCR for complex scripts. We

[¶]<https://azure.microsoft.com/en-us/products/ai-services/ai-document-intelligence>

Model	WER ↓	CER ↓	BLEU ↑	CHRF ↑	TEDS ↑	MARS ↑
Baseer (ours)	0.17	0.1	<u>77.28</u>	<u>89.29</u>	61.27	75.28
Gemini-2.5-pro	<u>0.37</u>	<u>0.31</u>	77.92	89.55	<u>52</u>	<u>70.775</u>
Gemini-3-pro	0.69	0.6	60.39	78.6	59	68.8
Azure AI	0.44	0.27	62.04	82.49	42	62.245
Dots.ocr	0.50	0.40	58.16	78.41	40	59.205
Nanonets	0.71	0.55	42.22	67.89	37	52.445
Qari	0.76	0.64	38.59	64.50	21	42.750
Qwen2.5-VL-32B	0.76	0.59	37.62	62.64	41	51.820
GPT-5	0.86	0.62	40.67	61.6	48	54.8
Qwen2.5-VL-3B-Instruct	0.87	0.71	25.39	53.42	27	40.210
Qwen2.5-VL-7B	0.92	0.77	31.57	54.70	27	40.850
Gemma3-12B	0.96	0.80	19.75	44.53	33	38.765
Gemma3-4B	1.01	0.85	9.57	31.39	28	29.695
GPT-4o-mini	1.36	1.1	22.63	47.04	26	36.52
AIN	1.23	1.11	1.25	2.24	21	11.620
Aya-vision	1.41	1.07	2.91	9.81	26	17.905

Table 4: Comparison of models across multiple evaluation metrics on Misraj-DocOCR. Best values are highlighted in **bold** and the second-best values are underlined.

Model	WER ↓	CER ↓	BLEU ↑	CHRF ↑	TEDS ↑	MARS ↑
Baseer (ours)	0.22	0.17	77.24	89.39	63.79	76.59
Dots.ocr	<u>0.39</u>	<u>0.28</u>	<u>59.28</u>	<u>83.16</u>	<u>55</u>	<u>63.08</u>
Nanonets	0.51	0.40	51.37	77.45	33	55.225
Qari	0.65	0.48	44.61	71.45	43	57.225
Qwen2.5-VL-3B-Instruct	0.70	0.57	40.44	66.78	31	48.89
Qwen2.5-VL-7B	0.76	0.63	36.76	62.45	24	43.225
Gemma3-12B	0.85	0.69	27.56	52.09	43	53.545
Gemma3-4B	0.95	0.82	12.94	31.72	27	29.36
Aya-vision	1.27	0.96	5.58	16.19	26	21.095
AIN	1.18	1.08	2.61	3.99	24	13.995

Table 5: Comparison of models across multiple evaluation metrics on KITAB-BenchPDF-to-Markdown dataset. Best values are highlighted in **bold** and the second-best values are underlined.

believe that this work establishes a strong baseline for future research and will accelerate the development of practical, high-accuracy OCR solutions for Arabic and other morphologically rich languages.

Limitations

While Baseer sets a strong baseline, several limitations should be noted. First, our evaluation relies on two benchmarks that we ourselves constructed or corrected: Misraj-DocOCR (400 expert-verified pages) and the corrected KITAB-Bench PDF-to-Markdown subset (60 pages). Both are modest in size, and the KITAB-Bench subset in particular is small enough that individual mispredictions can noticeably shift aggregate scores. We report single-run point estimates without statistical significance

testing; the differences between the top systems on the small subset should therefore be read as indicative rather than definitive. Evaluation on independent, third-party Arabic OCR benchmarks would further strengthen these conclusions and is left to future work. To support scrutiny, we release both benchmarks, the evaluation scripts, and the exact prompts; we nonetheless caution that MLLM results are sensitive to prompt wording and post-processing, and that different choices may change the relative ordering of the general-purpose baselines.

References

- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibong Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Saurabh Dash, Yiyang Nan, John Dang, Arash Ahmadian, Shivalika Singh, Madeline Smith, Bharat Venkitesh, Vlad Shmyhlo, Viraat Aryabumi, Walter Beller-Morales, et al. 2025. Aya vision: Advancing the frontier of multilingual multimodality. *arXiv preprint arXiv:2505.08751*.
- Kenneth Heafield. 2011. [KenLM: Faster and Smaller Language Model Queries](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197, Edinburgh, Scotland. Association for Computational Linguistics.
- Ahmed Heakl, Abdullah Sohail, Mukul Ranjan, Rania Hossam, Ghazi Shazan Ahmad, Mohamed El-Geish, Omar Maher, Zhiqiang Shen, Fahad Khan, and Salman Khan. 2025. Kitab-bench: A comprehensive multi-domain benchmark for arabic ocr and document understanding. *arXiv preprint arXiv:2502.14949*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Hugo Laurençon, Andrés Marafioti, Victor Sanh, and Léo Tronchon. 2024a. Building and better understanding vision-language models: insights and future directions. *arXiv preprint arXiv:2408.12637*.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024b. What matters when building vision-language models? *Advances in Neural Information Processing Systems*, 37:87874–87907.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Zhang Li, Yuliang Liu, Qiang Liu, Zhiyin Ma, Ziyang Zhang, Shuo Zhang, Zidun Guo, Jiarui Zhang, Xinyu Wang, and Xiang Bai. 2025. Monkeyocr: Document parsing with a structure-recognition-relation triplet paradigm. *arXiv preprint arXiv:2506.05218*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916.
- Souvik Mandal, Ashish Talewar, Paras Ahuja, and Prathamesh Juvatkar. 2025. Nanonets-ocr-s: A model for transforming documents into structured markdown with intelligent content recognition and semantic tagging.
- Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakkka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, et al. 2025. Smolvlm: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*.
- Ahmed Nassar, Andres Marafioti, Matteo Omenetti, Maksym Lysak, Nikolaos Livathinos, Christoph Auer, Lucas Morin, Rafael Teixeira de Lima, Yusik Kim, A Said Gurbuz, et al. 2025. Smoldocling: An ultra-compact vision-language model for end-to-end multi-modal document conversion. *arXiv preprint arXiv:2503.11576*.
- Mohamed Rashad. 2024. [Arabic-nougat: Fine-tuning vision transformers for arabic ocr and markdown extraction](#). *arXiv preprint arXiv:2411.17835*.
- Prince Sultan University RIOTU Lab. 2025. [Sard: A large-scale synthetic arabic ocr dataset for book-style text recognition](#). *arXiv preprint arXiv:2505.24600*.
- Andreas Steiner, André Susano Pinto, Michael Tschanen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, et al. 2024. Paligemma 2: A family of versatile vlms for transfer. *arXiv preprint arXiv:2412.03555*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Ahmed Wasfy, Omer Nacar, Abdelakreem Elkhateb, Mahmoud Reda, Omar Elshehy, Adel Ammar, and Wadii Boulila. 2025. Qari-ocr: High-fidelity arabic text recognition through multimodal large language model adaptation. *arXiv preprint arXiv:2506.02295*.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. 2025. Internv13: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.