

PaddleOCR-VL-1.6: A 0.9B Document Parser That Beats Models Hundreds of Times Its Size

Kabui, Charles

2026-06-15

[Read at ToKnow.ai](#)

PaddleOCR-VL-1.6:
A 0.9B Parser That Beats Giant Models

Turning documents into clean data in 109 languages

96.33%
top score on
OmniDocBench v1.6

0.9B
params, beats
Gemini 3 Pro

109
languages parsed,
Arabic to Thai

June 15, 2026

ToKnow.ai

Baidu's PaddlePaddle team released PaddleOCR-VL-1.6, a compact model that turns document images into structured Markdown and JSON, pulling out text, tables, formulas, charts, and seals across 109 languages. It pairs a layout detector with a 0.9B vision-language model that reads each region, built on a dynamic-resolution image encoder and the small ERNIE-4.5-0.3B language model. On [OmniDocBench v1.6](#), a standard document-parsing test, it scores

96.33%, the top result on the leaderboard. That puts a 0.9B model ahead of Gemini 3 Pro (92.91%), GPT-5.2 (86.59%), the 235B-parameter Qwen3-VL (89.78%), and the 1.2B MinerU2.5-Pro (95.75%).

Because the model is so small, you do not need a giant cloud model or a paid per-page API to convert documents into clean data. Quantized builds run locally in tools like Ollama and LM Studio, and the Apache 2.0 license lets you fine-tune and ship it commercially. For anyone digitizing contracts, invoices, medical records, or old books, that means accurate extraction of nested tables and dense math at close to zero cost, in scripts from Arabic to Thai.

The gains did not come from more data. The team instead hunted for the old version's weak spots: inputs where a tiny visual change flips the output, rare layouts the data barely covered, and training examples with wrong labels, checked against independent parsers like MinerU2.5-Pro. Fixing just those spots, plus a short reinforcement-learning pass, beat brute-force scaling. As with [MinerU2.5-Pro](#), document AI now turns on data quality, not model size.

Sources:

- [PaddleOCR-VL-1.6 Technical Report \(arXiv\)](#)
- [PaddleOCR-VL-1.6 Model \(HuggingFace\)](#)
- [OmniDocBench Benchmark](#)
- [PaddleOCR-VL Base Model Paper \(arXiv\)](#)
- [PaddleOCR-VL-1.6 Online Demo \(HuggingFace\)](#)

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. Read more: [/terms-of-service](#)