

dots.ocr-1.5 landed on February 16, 2026 and immediately entered OCR conversations because it is not framed as only a document OCR model. It is positioned as a broader visual-language parser that also covers web pages, scene text, and SVG-oriented tasks.

That makes comparison harder and more useful at the same time. If your goal is pure document OCR, GLM-OCR and PaddleOCR-VL-1.5 still set a strong baseline. If your goal is one model across OCR and non-document parsing tasks, dots.ocr-1.5 may be worth a serious challenger slot.

1 Evidence tier first: what is strongest today?

Before comparing scores, separate evidence quality:

Model	Evidence maturity (Feb 16, 2026)	Practical implication
PaddleOCR-VL-1.5	Strongest: paper + model card + official Paddle docs	Best option when you want the most complete public technical write-up
GLM-OCR	Strong: official repo + model card; technical report not yet published	Good production candidate, but some claims still rely on repo/model-card reporting
dots.ocr-1.5	Early: official repo + model card release notes for 1.5; no dedicated 1.5 paper yet	High-upside challenger, but treat current results as early-cycle evidence

2 What changed in dots.ocr-1.5

Based on the official release notes/model card, dots.ocr-1.5 adds:

- A 3B model profile (1.2B visual encoder + 1.7B language model).
- Expanded task framing beyond document OCR: scene text OCR, object grounding, UI/web parsing, and image-to-SVG style parsing.
- Reported gains over prior dots.ocr on both classic OCR and broader parsing benchmarks.
- Longer output behavior emphasis (release notes mention max output length up to 16k tokens).

3 Comparison at a glance

Dimension	dots.ocr-1.5	GLM-OCR	PaddleOCR-VL-1.5
Model size	3B	0.9B	0.9B
Core positioning	OCR + broader visual parsing	OCR-first specialist	OCR-first specialist
Output emphasis	Text/Markdown plus richer parsing modes (including SVG-oriented workflows)	OCR text, markdown, bbox/layout outputs	OCR + document parsing outputs across complex elements
Deployment paths highlighted by authors	HF/Transformers + vLLM guidance in release docs	vLLM, SGLang, Ollama, and SDK flow	Paddle ecosystem and HF model distribution
License	MIT	Model MIT, repo Apache-2.0	Apache-2.0

4 Benchmark signals without overclaiming

4.1 Reported numbers from official sources

Metric family	dots.ocr-1.5	GLM-OCR	PaddleOCR-VL-1.5
OmniDocBench aggregate headline	No single 94.x headline in 1.5 card text; card reports TextEdit 0.031 and ReadOrderEdit 0.029	94.62 (reported)	94.5 (reported)
Elo on OmniDocBench v1.5 (author setup)	1025.8	963.1	976.4
Elo on olmOCR-Bench (author setup)	1089.0	1008.8	1010.0
Elo on XDocParse (author setup)	1157.1	933.9	932.5

4.2 How to read those numbers correctly

- GLM-OCR and PaddleOCR-VL-1.5 headline figures are usually quoted as direct benchmark scores in their own sources.
- dots.ocr-1.5 highlights an Elo-based comparison stack (Gemini 3 Flash as judge in the model card setup), which is directional but not identical to independent leaderboard methodology.
- Conclusion: dots.ocr-1.5 has strong early signals, but you should still verify with your own stratified set before promoting it to primary.

5 Where each model is usually the better first test

If your workload is mostly...	First model to test	Why	Second model
High-volume document OCR with compact deployment constraints	GLM-OCR	Strong reported accuracy in 0.9B size and practical serving VL-1.5 paths	PaddleOCR-
Messy captures, scans, and robustness-heavy document pipelines	PaddleOCR-VL-1.5	Robustness-oriented evaluation framing in the paper/release	GLM-OCR
One stack for OCR plus web/screen/scene/SVG parsing	dots.ocr-1.5	Broadest task envelope in this trio	GLM-OCR
Risk-controlled rollout for regulated production	PaddleOCR-VL-1.5 or GLM-OCR	Stronger OCR-specific evidence maturity today	dots.ocr-1.5 as challenger lane

6 Recommended rollout pattern (2 weeks)

1. Keep GLM-OCR and PaddleOCR-VL-1.5 as the two baseline tracks.
2. Add dots.ocr-1.5 as a challenger track on the same 50-page stratified protocol from the main OCR post.
3. Score three groups separately: text fidelity, reading order, and task coverage outside plain OCR.
4. Promote dots.ocr-1.5 only if its gains are stable on your hardest slice, not just average pages.

Use this companion together with the main OCR market map:

https://instavar.com/blog/OCR_SOTA_Feb_2026_Open_Document_AI_Leaderboard

7 Bottom line

dots.ocr-1.5 is not just another OCR score update. It is a broader product direction: one model handling OCR plus adjacent visual parsing tasks.

For pure document OCR, GLM-OCR and PaddleOCR-VL-1.5 remain high-confidence anchors. For teams consolidating OCR with UI/screen/scene parsing, dots.ocr-1.5 is the most important new challenger to evaluate now.

Sources

- [dots.ocr repository](#)
- [dots.ocr-1.5 model card](#)
- [dots.ocr paper \(base model\)](#)
- [GLM-OCR repository](#)
- [GLM-OCR model card](#)
- [PaddleOCR-VL-1.5 paper](#)
- [PaddleOCR-VL-1.5 model card](#)