

Ideogram 4.0: A Downloadable Image Model That Gets Text and Layout Right

Kabui, Charles

2026-06-29

[Read at ToKnow.ai](#)

Ideogram 4.0:
An Open Image Model
That Gets Text Right

Downloadable model built for typography and layout

- #1**
Top open-weight model on the Design Arena board
- 9.3B**
Parameters, beats 80B models at text rendering
- 2K**
Native 2048px output, no upscaling step

Trained on structured JSON prompts for exact layout and color control

June 29, 2026

ToKnow.ai

The graphic features a dark blue background with a grid pattern and a faint line graph in the top right corner. The text is primarily white and light blue, with key statistics highlighted in orange, yellow, and light blue.

Ideogram, the Toronto image company, released Ideogram 4.0 on June 3, 2026, its first model with downloadable weights. It is a 9.3-billion-parameter generator trained from scratch, not a fine-tune. Instead of a text-only prompt reader, it uses Qwen3-VL-8B, a model that reads both words and images. And it was trained entirely on structured JSON prompts, so you can set the composition, pass exact hex color codes, and draw bounding boxes that place each

text line or object exactly. It outputs up to 2048 pixels (2K) directly and ranks as the top open-weight model on the [Design Arena](#) leaderboard, behind only closed models from OpenAI and Google.

The everyday frustration with AI image tools is that they misspell words on a logo or poster and ignore requests for a specific layout. Ideogram 4.0 is built for that job. In a blind typography test judged by ten professional designers, it was picked best of four models 47.9% of the time, ahead of Google's Nano Banana 2 at 30% and FLUX.2 at 15.5%. It does this at 9.3 billion parameters, beating the text rendering of models up to 80 billion. You can run it on your own graphics card or call the hosted API from \$0.03 an image.

One catch: the free weights carry a non-commercial license, so commercial use needs a paid license or the API. The larger bet is that a model built for design beats the general multimodal systems that were supposed to make specialized image models obsolete. [Reve 2.0](#) made a similar layout-first wager earlier this month.

Read More: [Nano Banana 2, the Google image model Ideogram is benchmarked against](#)

Sources:

- [Ideogram 4.0 technical details \(Ideogram blog\)](#)
- [Ideogram 4.0 press release \(Ideogram\)](#)
- [Ideogram 4.0 model page \(Ideogram\)](#)
- [ideogram-ai/ideogram-4-nf4 weights \(Hugging Face\)](#)
- [Ideogram 4.0 overview \(ImagineArt\)](#)

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