

# Reve 2.0: An Image Model That Builds a Layout First, Then Draws

Kabui, Charles

2026-06-15

---

[Read at ToKnow.ai](#)

---



**Reve 2.0 Builds the Layout Before Drawing Pixels**

An editable map of every element, rendered at native 4K

**#2**  
Arena text-to-image rank  
Behind only GPT Image 2

**4K**  
Native render resolution  
Print-ready, no upscaling

**10x**  
Fewer GPUs to train  
Beat far larger rivals

Images as code: humans and agents edit the same layout

June 15, 2026

ToKnow.ai

Reve, a Palo Alto research lab, released [Reve 2.0](#), an image model that refuses to turn your prompt straight into pixels. Instead, it **first builds a layout**: a structured map where every object, text block, and region has its own position, size, and description. Only then does it render the layout as a native 4K image. Plain text is ambiguous: tweak one word and the whole picture can change. A layout is explicit and editable, like HTML is to a web page. The

model is a single Large Layout Model, built by training [open-source Qwen language models](#) on billions of images for spatial reasoning. It debuted at [#2 on the Arena text-to-image leaderboard](#), a 125-point jump over the previous version, while training on 10x fewer GPUs.

With a normal generator, fixing one detail in an otherwise good image means re-rolling the whole prompt and hoping the good parts survive. Reve 2.0 treats every element as a separate, self-contained piece, so you can move a subject, rewrite the sign on a wall, or swap a background while the rest of the scene stays put. A designer can fix an ad's headline without regenerating the whole image, and AI agents can read and edit the layout directly, the way they edit code.

This reframes image generation as something closer to writing editable code than rolling dice. The bet is that the bottleneck is no longer raw image quality but control. On the same leaderboard, Reve 2.0 edged out Google's [Nano Banana 2](#), which also renders native 4K but still turns a prompt straight into pixels.

Sources:

- [The Layout Bet \(Reve technical write-up\)](#)
- [Introducing Reve 2.0 \(Reve announcement\)](#)
- [Arena Text-to-Image Leaderboard](#)
- [Qwen3 Language Models \(Alibaba\)](#)

---

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