BraintankAI ASIA

EXPLAINER · VISUAL GENERATION LEVELS · 5 MIN

BEIJING · Tsinghua University

A five-level ladder for judging how intelligent image generators really are

A 2026 survey sorts visual generators from prompt-to-picture models up to agents and world models, then probes two closed systems to find where they break.

12 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Shane Rounce

Image generators can produce pictures that are hard to tell from photographs, and people often say they are getting more intelligent. More intelligent in what sense? The same photorealistic systems can still lose track of a character's face from one picture to the next. They can garble a line of text, get the spatial relations between objects wrong, or show a scene that no physical process could produce.

Keming Wu and a large group of co-authors from Tsinghua University, LMMs-Lab and other institutions posted a survey to arXiv in 2026 that tries to give that question some structure. It is a preprint. The authors argue that calling a model more intelligent should mean spelling out which capability has improved. Familiar automatic scores such as FID and CLIP score rate how realistic an image looks and how well it matches its prompt, and the authors say these mostly test the simplest abilities. Relying on them, they argue, leads researchers to overestimate progress.

The idea

The authors propose a five-level ladder, inspired by OpenAI's staged roadmap for general AI. Each level adds one competence. Systems are classed by what they can do rather than by how they are built. The first three levels concern how much a single forward pass can absorb. A forward pass is one run of the model from input to output. Level 1 makes a picture from a text prompt alone. Level 2 also follows one explicit control, such as a sketch or a reference face. Level 3 handles rich context in one go, such as several reference images or the history of earlier edits.

The top two levels are about orchestrating many calls to the model rather than one. A Level 4 system has an external controller and meets four criteria drawn from the language-model literature. It chooses its own actions, works in a loop with feedback from its environment, persists towards a long-range goal and decides for itself when to stop. Applying Anthropic's distinction between workflows and agents, the authors class systems such as GEMS and Gen-Searcher as workflows, while JarvisArt counts as an agent. Level 5, world modelling, adds a requirement for grounding in physical and causal knowledge. A system at this level would predict what would really happen, not just what looks plausible. The authors mark this level as lying in the future.

How this paper puts it to work

The paper introduces no new model. It breaks every system into four parts. An encoder or tokeniser converts images into a form the network can work on. A backbone does the main computation. One module takes in conditions such as reference images, and another fuses them with the rest. Through that lens the authors read ten industrial technical reports from 2025 and 2026 side by side, among them HunyuanImage 3.0, Qwen-Image and Wan-Image. Six of the ten run text-to-image generation, image-to-image work and instruction-based editing through the same backbone. The authors take this as evidence that generation and editing have merged into one architecture.

Closed commercial systems disclose little about how they are built. For these, the authors offer four conjectures that they describe as speculative but falsifiable. Their main suggestion is that the gap between open and closed systems lies mostly in the machinery around the image renderer, not in the renderer itself. The first conjecture places a top-tier vision-language model upstream of the renderer; such a model is an AI that reads images and text together. Taken together, the conjectures imply that closed systems most plausibly operate at Level 4, while the strongest open ones the authors discuss are capped at Level 3 by construction. On training, the authors argue that gains increasingly come from data quality and curation rather than from parameter scaling alone. HunyuanImage 3.0, by its makers' account, keeps under 45 per cent of a pool of more than ten billion images.

The hands-on part is a set of stress tests in eight dimensions, run on two closed generators, Nano Banana and GPT-Image-2, using identical prompts and settings. Each dimension is mapped to the level or levels it mainly challenges. In one test, the model is asked to reassemble a jigsaw cut from a photo of hot-air balloons. In another, it is given a photo of a citrus slice floating on water and asked to explain the physics, then to show what would happen if the slice sank. Other tests ask it to

solve a Chinese entrance-exam-style problem on electromagnetic induction directly on the image, edit a picture over many turns, estimate depth in a cluttered real-world photo and draw a plan of Tang Dynasty Chang'an.

What it shows, and what it does not

The tests suggest where the systems break. In the jigsaw test, the model invented plausible content instead of rigidly putting the pieces back. In the buoyancy test, it added bubbles and deformation that looked causal but reflected no real physics. The authors describe the model's approach in that case as understanding first and rendering second. Solving the exam on the image was productive but fragile. In multi-turn editing, each change looked fine on its own. However, the authors infer that the systems in effect respond only to the previous image and the new instruction, so small errors piled up unnoticed. Depth maps caught object outlines but gave many different objects the same shade. The authors describe this as rewriting from prior expectations rather than measurement. The Chang'an plan had an accurate layout and grid of city wards but unclear or duplicated Chinese labels.

The survey also questions how the field keeps score. Using a vision-language model as a judge, it reports, typically yields Spearman correlations of 0.57 to 0.75 with human raters. Spearman correlation measures how similarly two rankings order the same items; a score of 1 means identical. DPG-Bench is a test built on 1,000 long prompts. On it, scores reported in recent industrial technical reports cluster between 84 and 89, and the Relation and Other categories are the ones that still separate frontier models.

The evidence has clear limits. This is a survey and roadmap, not a controlled experiment. The stress tests are individual qualitative cases on two closed systems, with no reported scoring rubric, repeated trials or independent raters. Where the frontier sits is therefore the authors' interpretation of selected examples. The six-in-ten tally rests on reports the authors selected, and the industrial numbers the survey cites come from labs' own technical reports. The authors call their own reading of closed systems speculative. They mark Level 5 as future and place video world models at the edge of their image-focused scope. They also acknowledge that some work may be missing from a fast-growing literature. What the paper offers is a vocabulary. It gives readers a way to ask not whether a generator is impressive, but which rung it stands on.

PAPER	*Visual Generation in the New Era: An Evolution from Atomic Mapping to Agentic World Modeling*
AUTHORS	Keming Wu, Zuhao Yang, Kaichen Zhang, Shizun Wang, Haowei Zhu, Sicong Leng, Zhongyu Yang, Qijie Wang, Sudong Wang, Ziting Wang, Zili Wang, Hui Zhang et al.
VENUE	arXiv preprint — not yet peer reviewed · 30 April 2026
LINK	arxiv.org/abs/2604.28185
INDEX	stacks.braintank.ai
CODE	github.com/EvolvingLMMs-Lab/Evolving-Visual-Generation