BraintankAI ASIA

EXPLAINER · STREAMING VIDEO DISTILLATION · 5 MIN

HEFEI · University of Science and Technology of China

How reward scores tell a fast video model where to focus its training

Stream-R1 weights each training video, frame and region by a reward model's verdict, edging past its direct baseline on benchmarks at the same speed.

16 November 2026 · arXiv preprint, not yet peer reviewed

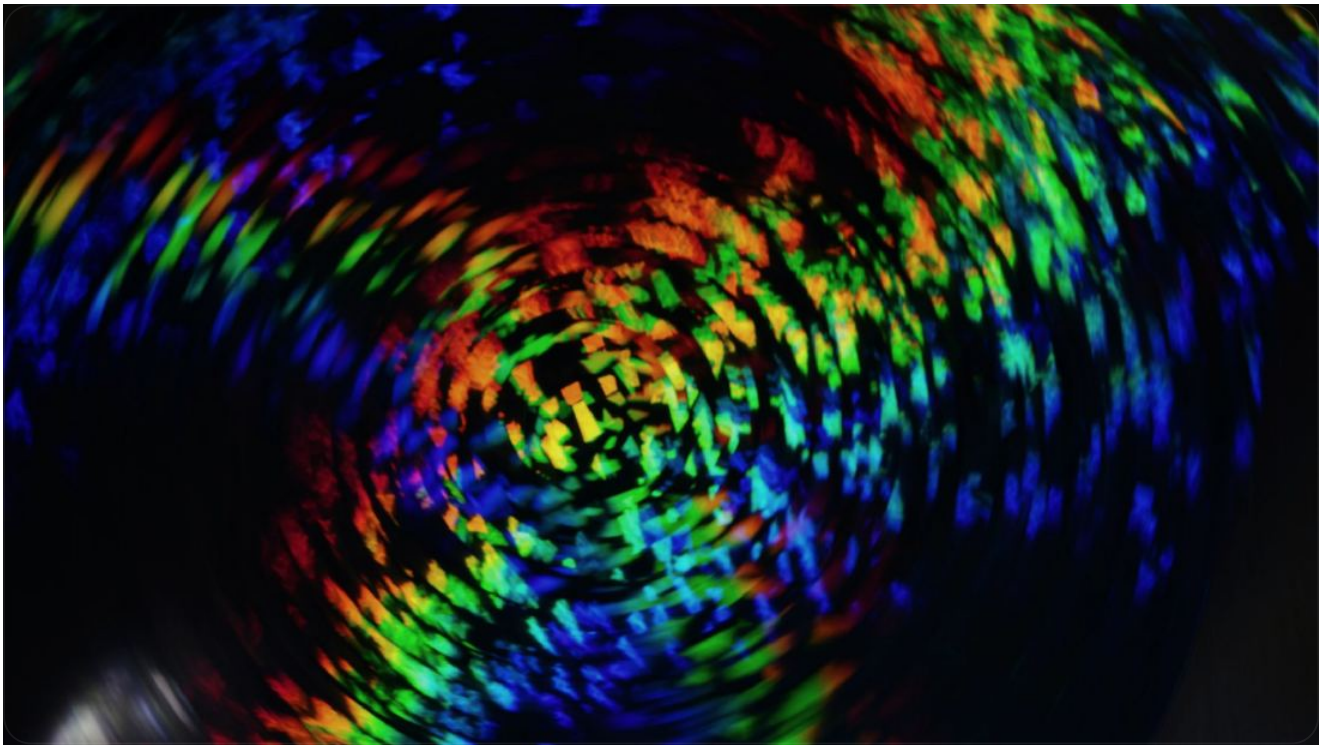


Photo: Flyd

The video models that produce the most convincing clips are slow. They start from random noise and clean it up over many passes. In the authors' comparison, one such model, Wan2.1 with 1.3 billion parameters, managed 0.78 frames per second. That is acceptable for a clip you can wait for, but no use for video that has to keep flowing, such as an interactive scene or a stream with no fixed end.

The usual workaround is to train a fast model to imitate a slow one and have it produce video in short chunks, one after another. A study by Bin Wu and colleagues at the University of Science and Technology of China and FrameX.AI, posted to arXiv in 2026 as a preprint that has not yet been peer reviewed, asks a narrower question about that training. Should every part of every practice video count equally when the fast model learns? Their method, Stream-R1, answers no.

The idea

In autoregressive streaming generation, a diffusion model produces video chunk by chunk. Each chunk is conditioned on the ones before, so frames can stream out live and the video can in principle run indefinitely. To make each chunk quick, a few-step 'student' is trained to match a slow, many-step 'teacher', commonly through distribution matching distillation, or DMD. The student generates a training video, called a rollout. Two helper networks then estimate how a noisy version of that rollout should be cleaned up. One is derived from the teacher, and the other is trained continuously on the student's own outputs. The gap between their estimates tells the student which way to change.

The authors place several streaming systems in this family, among them CausVid, Self-Forcing, LongLive and Rolling-Forcing. They argue that all of them treat every rollout, frame and pixel as equally good evidence.

That, they argue, caps quality in two ways. The first they call Inter-Reliability: whole rollouts differ in how trustworthy their signal is. The teacher-derived network acts as a local corrector, and the direction of its correction depends on where a sample already sits. On a rollout far from the teacher's best work, its guidance is therefore less dependable. The second they call Intra-Perplexity, a term borrowed loosely from language modelling. It refers to differences within one video in how much room each region or frame has to improve. Uniform weighting spends effort on parts that are already as good as they will get. The authors argue for both premises but do not directly measure either.

How this paper puts it to work

Stream-R1 is built on an earlier method, Reward Forcing, which already scales the DMD loss using a reward model. A reward model is a pretrained network that scores a video, here on visual quality, motion quality and match to the text prompt. Stream-R1 averages the three scores and subtracts a penalty when the axes have been improving unevenly over recent training steps. It then multiplies the rollout's loss by an exponential of the result, so rollouts the reward model rates highly count for more. The authors recast this weight as a measure of how far to trust each rollout.

The part that goes beyond Reward Forcing works inside each video. Each score is traced back to the input pixels to produce a gradient saliency map. The map records how much the score would change if a given pixel changed slightly. The three maps are blended so that the weakest-scoring axis counts most, then split in two: a weight for each frame, and a weight for each location within a frame. Each has a floor, 0.20 for frames and 0.15 for locations, so no part is ignored entirely. The product of the two multiplies the DMD error location by location. The weights are used only during training, and the student's architecture and speed are unchanged.

The student is the 1.3-billion-parameter Wan2.1 and the teacher is the 14-billion-parameter version. The student produces five-second clips at 832 by 480 pixels, three compressed 'latent' frames per chunk, using four denoising steps. Training ran for 1,000 steps on eight A100 graphics processors and took about 56 hours.

What it shows, and what it does not

The main test is VBench, a standard video-generation benchmark. Its Total score combines a Quality sub-score, covering visuals and motion, with a Semantic sub-score for prompt adherence. Stream-R1 scores 84.40, the highest among the models compared. Reward Forcing, its direct baseline, scores 84.13, and both run at 23.1 frames per second. Stream-R1 gains 0.30 on Quality and 0.12 on Semantic. The multi-step 1.3-billion-parameter Wan2.1 scores 84.26 overall and slightly higher on Quality (85.30 against 85.14), while running about 30 times slower. That comparison is with the smaller Wan2.1, not the 14-billion-parameter teacher used in training.

The secondary tests are mixed. On videos up to three minutes long, the authors report Stream-R1 ahead of Reward Forcing on six metrics at every length, with the gap widening at two and three minutes. This appears only as a chart, against that one baseline. Qwen3-VL, a large vision-language model that reads both images and text, scored one-minute videos from 1 to 5. It rated Stream-R1 higher than Reward Forcing for visuals (4.92 against 4.82) and text match (4.11 against 4.04), but lower for dynamics (4.04 against 4.18). A human study used 50 one-minute videos and five annotators, comparing against Reward Forcing only. Stream-R1 had an overall win rate of 57%, with ties counted as half. On temporal consistency it was almost even, 25 wins to 24 losses.

The margins are small, and they are reported without error bars. The ablation adds components one at a time to a DMD baseline without any reward. It never isolates the whole-rollout reward weighting that Stream-R1 shares with Reward Forcing, so it cannot show how much each of Stream-R1's own additions contributes beyond that shared weighting. Most of the gain arrives with the per-frame weighting, which lifts Total from 83.67 to 84.40. Yet raising that floor from 0.20 to 0.40 drops Total to 83.42, just below the no-reward baseline. Spatial weighting alone made long-video drift, the decay in quality over a long clip, worse. Drift fell below the baseline only once per-frame weighting was added. Some figures quoted in the paper's text do not match its own table. And because the reward model scores much the same qualities the benchmarks measure, part of the gain may reflect training toward the yardstick. In a single controlled test with artificial blur, the frame weights did rise steadily as more of the picture was blurred, from 0.587 to 2.117. A fair reading is a small improvement on most measures at no cost in speed. The case for why it works has yet to be made.

PAPER	*Stream-R1: Reliability-Perplexity Aware Reward Distillation for Streaming Video Generation*
AUTHORS	Bin Wu, Mengqi Huang, Shaojin Wu, Weinan Jia, Yuxin Wang, Zhendong Mao, Yongdong Zhang et al.
VENUE	arXiv preprint — not yet peer reviewed · 5 May 2026
LINK	arxiv.org/abs/2605.03849
INDEX	stacks.braintank.ai
CODE	github.com/FrameX-AI/Stream-R1