BEIJING · Xiaomi

Video model restyles one object or its background while largely sparing the rest

Xiaomi's LociArt pairs a mask that encodes both where and which style with 100,000 synthetic training pairs, though a rival fine-tuned on the same data comes close.

17 November 2026 · arXiv preprint, not yet peer reviewed

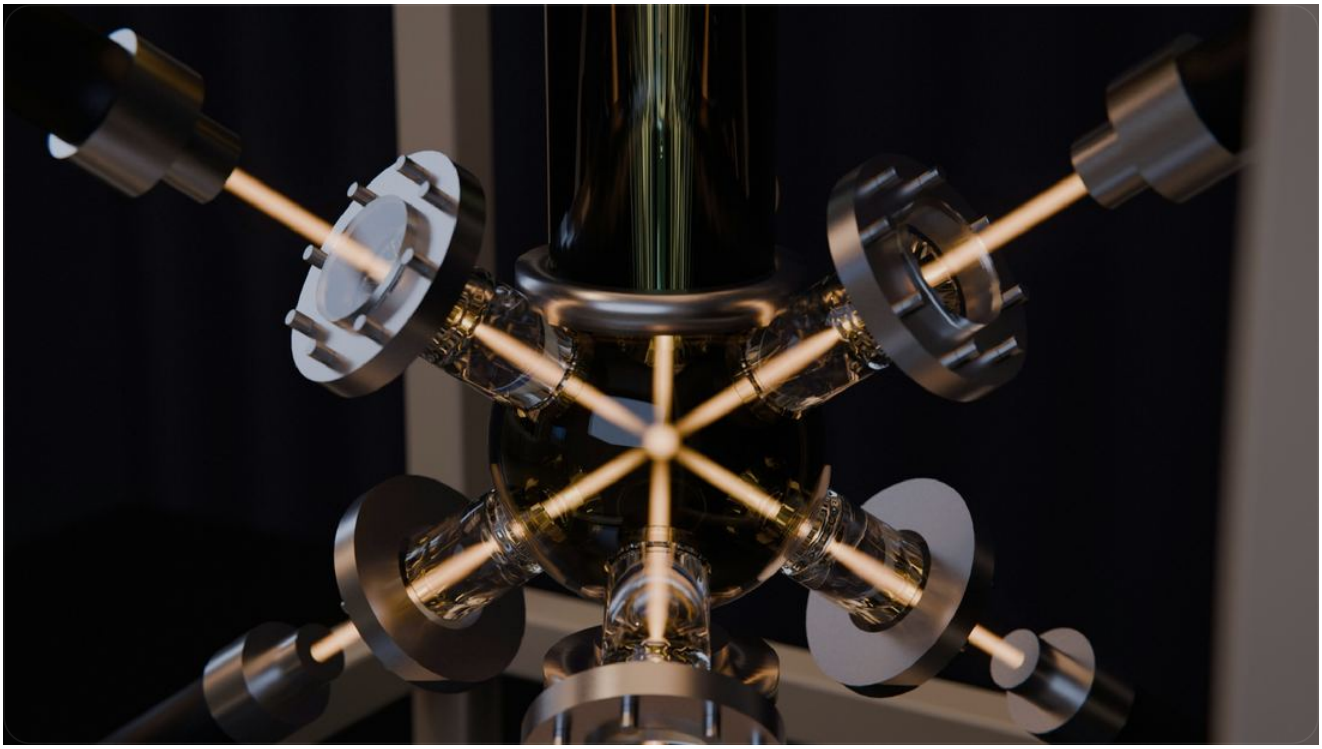


Photo: Brecht Corbeel

Most video stylisation repaints the whole frame. LociArt, from researchers at Xiaomi and Xi'an University of Electronic Science and Technology, targets a narrower job: restyling only a chosen part of a clip, either a moving object or everything around it, in a requested art style while preserving the rest of the picture. The 2026 paper is a preprint that has not yet been peer reviewed, and it bundles three things: a training dataset, a model called GRNStyle and a benchmark for judging such edits.

The authors identify three obstacles. A binary mask, a black-and-white map marking which pixels to edit, says where to edit but not which style to apply, so the model must link a separately written style description to the right region. Ways of feeding a mask into a video model also trade context against precision: compressing the mask the way video is compressed blurs its edges, while simply

shrinking the raw mask keeps its shape but lacks the learned context across frames. And existing paired datasets cover whole-frame restyling or broader editing, not a clip before and after a local restyle.

GRNStyle builds on Generative Refinement Networks, a generator that represents video as binary bits and repeatedly refines an initially random set of bits into the output. Its main addition is a 'polar style-aware mask'. Each style is given a fixed angle on a circle, and inside the edit region the mask stores the cosine and sine of that angle in two channels, with a third channel simply marking the region. Every style code therefore has the same length and differs only in direction, and compared with a plain mask the scheme adds no extra channels or tokens. A second component passes the mask through the same frozen encoder as the video, subtracting the encoding of a blank mask so that 'no edit' produces exactly zero signal, and also supplies a downsampled raw mask that gates where that signal is injected.

The training data, called LociArt-100K, comes from an automated pipeline. Qwen3-VL-32B names prominent objects in videos drawn from the SA-V and Ditto-1M collections, SAM3 tracks them, and SVOR erases each object to give a clean background clip. Both versions are then restyled in full by Ditto and an internal fine-tuned version of the VACE video model, with GPT-5.5 filtering the outputs. VideoMaMa estimates alpha mattes, soft per-pixel transparency maps, so the stylised and original pieces can be blended into foreground-only and background-only targets. Human annotators checked the results, leaving 100K pairs across 18 style categories.

The authors trained a 2B-parameter model on eight NVIDIA H200 chips and tested it on their benchmark of 168 cases. On their main metric, Style Prototype Similarity (SPS), which compares the edited region with reference examples of the style, GRNStyle scored 0.407 for foreground edits and 0.705 for background edits. The runner-up was VACE-1.3B fine-tuned on LociArt-100K, at 0.386 and 0.676; the commercial Seedance 2.5 scored 0.347 and 0.358. Before fine-tuning, VACE-1.3B managed only 0.156 and 0.191, which suggests much of the improvement comes from the dataset rather than the architecture. GRNStyle does not lead everywhere. Unmodified VACE-1.3B kept the untouched area most intact, at 32.35 and 29.17 dB PSNR against GRNStyle's 25.92 and 26.15, but it also barely restyled anything. Seedance 2.5 led on matching foreground edits to the style's name and, narrowly, on frame-to-frame consistency, at 0.984 against 0.983.

One test favours the polar design clearly. When a model trained on ten styles was taught five new ones without revisiting old training data, its SPS on the old styles fell by 0.022 with polar codes and by 0.102 with ordinary binary masks.

The limits are considerable. 'Free-style' means a closed set of learned categories, and new styles need further fine-tuning. Training targets inherit the quality and biases of other generators, one of them not public. The benchmark and SPS are the authors' own, the benchmark is small, and AI judges are used both to filter data and, via Gemini 3.6 Flash, to score results. A user study involved 20 participants and does not report the baselines' mean scores in the text. The benchmark tables give no error bars, and the striking ability to restyle several separate regions in different styles at once is shown only through examples.

The problem of steering a video generator with extra spatial inputs also appeared in braintankAI's coverage of a KAIST AI preprint that guides generation with surface-shape maps so a subject stays consistent as the camera circles it (earlier story). There the extra input fixes what a subject looks like; here it fixes where a style may go.

PAPER	*LociArt: A Unified Paradigm for Free-Style Local Video Stylization*
AUTHORS	Yuxuan Chen, Feng Xie, Jiagao Hu, Fuhao Li, Zepeng Wang, Fei Wang et al.
VENUE	arXiv preprint — not yet peer reviewed · 27 October 2026
LINK	arxiv.org/abs/2610.36456
INDEX	stacks.braintank.ai