Moebius: A 0.2B Image-Inpainting Model That Rivals 10B Giants

Kabui, Charles

2026-06-23

Researchers from Huazhong University of Science and Technology and VIVO AI Lab released Moebius, an image inpainting model that erases objects, restores faces, and repairs damaged photos. Inpainting means filling or replacing a chosen part of an image so the edit looks natural. Moebius uses 0.22 billion parameters, under 2% of the 11.9 billion in FLUX.1-Fill-Dev, the model it competes with, yet matches or beats it and SD3.5 Large-Inpainting across

six benchmarks of natural scenes and faces. Each step takes about 26 milliseconds on one GPU, and full edits finish over 15 times faster. It summarizes context into small fixed-size matrices instead of comparing every pixel pair, then trains the compact model to copy a larger teacher, the team's earlier PixelHacker.

Because the model is small and its weights are open under a permissive license, this editing can run on a laptop or phone, not a data center. Someone can remove a photobomber or fix an old family photo locally, for free, without sending private pictures to the cloud. That size gap separates a paid-server feature from one that ships inside an app.

Moebius also pushes back on the idea that better results always need bigger models. By treating inpainting as one narrow job and training a specialist through distillation, it reaches a giant's quality far more cheaply. For a defined task, careful design is outrunning raw scale.

Read More: Nano Banana 2 does the same inpainting in a big cloud model, the reverse of Moebius's tiny, local one.

Sources:

- Moebius: 0.2B Lightweight Image Inpainting Framework with 10B-Level Performance (arXiv, June 17, 2026)
- Moebius Project Page (HUST Vision Lab)
- Moebius Code and Training Recipe (GitHub, ECCV 2026, Apache-2.0)
- Moebius Open Model Weights (Hugging Face)
- FLUX.1-Fill-dev, the 11.9B Baseline (Black Forest Labs, Hugging Face)
