

TL;DR Omni-Effects unifies promptable and spatially controllable visual effects inside one CogVideoX-based pipeline. LoRA-MoE experts keep per-effect quality high, Spatial-Aware Prompts with Independent-Information Flow isolate each mask, and the Omni-VFX dataset plus released checkpoints make composite VFX runs practical for in-house teams.

What is Omni-Effects?

Omni-Effects is a research framework for controllable VFX generation that was unveiled on 11–12 August 2025 via arXiv, GitHub, and a LinkedIn announcement. Instead of training one LoRA per effect, the team introduces a unified diffusion pipeline that produces multiple effects at once while holding spatial constraints.

The system fine-tunes CogVideoX (image-to-video) backbones with two core ideas: a LoRA-based Mixture of Experts (LoRA-MoE) that routes prompts to effect-specific adapters, and a Spatial-Aware Prompt (SAP) format that injects mask layouts into the text stream. An Independent-Information Flow (IIF) block keeps those control signals from bleeding across effects. The release arrives with Omni-VFX, a curated VFX dataset distilled from Open-VFX assets, Remade-AI clips, and First–Last Frame-to-Video synthesis, plus CogVideoX checkpoints finetuned on that corpus.

Links:

- Paper: <https://arxiv.org/abs/2508.07981>
 - Project page: <https://amap-ml.github.io/Omni-Effects.github.io/>
 - GitHub: <https://github.com/AMAP-ML/Omni-Effects>
 - Hugging Face weights: <https://huggingface.co/GD-ML/Omni-Effects>
 - Omni-VFX dataset: <https://huggingface.co/datasets/GD-ML/Omni-VFX>
-

Key ideas

- LoRA-MoE routing: groups expert LoRAs per effect category so the unified model can blend creative prompts without cross-task interference (paper + README).

- Spatial-Aware Prompting: appends binary masks to the text tokens, letting users paint regions for “Melt it”, “Levitate it”, “Explode it”, “Anime style”, or “Winter scene” in the same clip (project page + README).
 - Independent-Information Flow: isolates control signals for each mask, stopping leakage when multiple effects fire simultaneously (paper abstract).
 - Omni-VFX corpus: assembles edited assets, Remade-AI distillations, and FLF2V-generated clips to supply diverse training data and benchmarking splits (README).
 - Released finetunes: CogVideoX-1.5 (prompt-guided) and CogVideoX-5B (single + multi-VFX) checkpoints, plus LoRA weights for spatial control (README updates).
-

Model lineup & availability

- CogVideoX1.5-5B-I2V-0mniVFX: prompt-guided VFX finetune backed by Omni-VFX.
- Omni-Effects LoRA bundles for single-VFX and multi-VFX runs (pairs with CogVideoX-5B).
- Omni-VFX dataset with prompt lists and region masks for mask-guided training/evaluation.
- Example prompts (VFX-prompts.txt) and shell scripts for repeatable inference in scripts/.

Artifacts are hosted on Hugging Face under the GD-ML organisation; the repo provides download and organisation instructions.

Quickstart (from repo docs)

Clone and install:

```
git clone https://github.com/AMAP-ML/Omni-Effects.git
cd Omni-Effects
conda create -n OmniEffects python=3.10.14
pip install -r requirements.txt
```

Grab checkpoints:

```
# place weights under ./checkpoints
pip install "huggingface_hub[cli]"
```

```
huggingface-cli download GD-ML/Omni-Effects --local-dir checkpoints/omni-effects
huggingface-cli download GD-ML/Omni-VFX --local-dir datasets/omni-vfx
```

Run prompt-guided VFX generation:

```
sh scripts/prompt_guided_VFX.sh # edit the prompt + source image in the script
```

Execute mask-guided spatial control:

```
# Single effect (e.g., Melt it)
sh scripts/inference_omnieffects_singleVFX.sh
```

```
# Composite effects (multiple masks/effects)
sh scripts/inference_omnieffects_multiVFX.sh
```

Control & tuning tips

- Spatial masks: craft binary masks per region; SAP merges them with text so each effect stays inside its polygon.
 - Effect menu: current release supports five presets-Melt, Levitate, Explode, Anime style, and Winter scene. Custom slots require retraining the relevant expert LoRA.
 - Prompt structure: follow “scene description → effect description → cinematic notes” for stronger prompt grounding before the SAP data is appended.
 - Performance tuning: adjust batch size, mask resolution, and diffusion steps within the provided scripts; CogVideoX finetunes remain 16:9 by default.
 - Dataset reuse: Omni-VFX ships image pairs, masks, and prompts; reuse them for fine-tuning bespoke effects or evaluation baselines.
-

Practical production notes

- Hardware: multi-VFX inference piggybacks on CogVideoX-5B; expect similar VRAM needs (~48GB for comfortable batching) when running without model parallel tweaks.
- Mask authoring: draw masks in the same resolution as the input image; mismatched scaling can cause halo artefacts along boundaries.
- Consistency vs. stylisation: LoRA-MoE keeps base fidelity, but extreme prompts may still blend effects-watch for SAP mask overlap.
- Deployment: scripts run via `torchrun`; integrate into pipelines by swapping prompts/masks and capturing the generated MP4 assets from the output folder.

- Roadmap signals: LinkedIn launch mentions continued expansion of effect categories and production-ready presets-monitor the repo for new expert LoRAs and higher-resolution checkpoints.
-

References

- arXiv: <https://arxiv.org/abs/2508.07981>
- Project: <https://amap-ml.github.io/Omni-Effects.github.io/>
- GitHub: <https://github.com/AMAP-ML/Omni-Effects>
- Hugging Face (models): <https://huggingface.co/GD-ML/Omni-Effects>
- Hugging Face (dataset): <https://huggingface.co/datasets/GD-ML/Omni-VFX>

Notes: Details reflect public material as of 12 August 2025. Check the repository for newly added effect experts, higher-res checkpoints, and updated SAP formats.