



INO AI LAB

[AI EDUCATION • COURSE MATERIAL]

Advanced Visual AI: Flux, SDXL, ComfyUI

Open-source image pipelines

LEVEL	DURATION	LESSONS
Advanced	5 hours	9



[00]

Table of Contents

LESSON 01 **Why Open Models Still Matter in 2026**

LESSON 02 **ComfyUI: Node-Based Workflows**

LESSON 03 **ControlNet, IP-Adapter, and LoRAs**

LESSON 04 **Production-Grade Image Pipelines**



[LESSON 01]

Why Open Models Still Matter in 2026

Closed tools like Midjourney and DALL-E are convenient, but open models — Flux Pro, SDXL Turbo, Stable Cascade — give you control, privacy, unlimited generations, and the ability to fine-tune on your own brand. For agencies, product designers, and anyone generating at volume, the unit economics flip in favor of open within a few thousand images per month.

Flux from Black Forest Labs is the current open SOTA — photorealism rivaling MJ, dramatically better prompt adherence, and a true commercial license on the Pro variant. Run it locally on a 4090, on RunPod for \$0.30/hr, or via Replicate for pay-per-image. Pick based on volume: occasional use → Replicate, daily use → RunPod, every day → local GPU.

// KEY TAKEAWAYS

- › Open models = control, privacy, unlimited generations.
- › Flux Pro is the 2026 open SOTA.
- › Choose infra by volume, not by hype.

[LESSON 02]

ComfyUI: Node-Based Workflows

ComfyUI is the professional node editor for image generation. Where Midjourney offers a chat box, ComfyUI offers a graph: load model → encode prompt → sample → decode → upscale → save. The learning curve is steep for a week, then everything becomes trivially composable. Most production studios in 2026 ship final assets through ComfyUI, not chat interfaces.

Start with the official example workflows and modify them. Install ComfyUI Manager first — it handles missing node dependencies that otherwise take hours to track down. Save workflows as templates per use case (product photo, character portrait, hero banner). Share inside your team via the workflow JSON file so quality stays consistent across designers.

// KEY TAKEAWAYS

- › ComfyUI is the pro tool for image workflows.
- › Install Comfy Manager day one.
- › Save workflows as templates per use case.

[LESSON 03]

ControlNet, IP-Adapter, and LoRAs

ControlNet conditions generation on structure — pose, depth, edges, segmentation. Pass a stick figure, get a perfectly posed person. IP-Adapter conditions on style or identity from a single reference image, often outperforming prompt-only style transfer. LoRAs are tiny fine-tunes (5–200 MB) that inject a specific style, person, or concept; you can stack multiple LoRAs at different weights for compound effects.

The skill is combining them with restraint. Stacking three LoRAs at full weight produces incoherent mush. Start with one ControlNet + one LoRA at 0.6 weight, generate, then layer additions one at a time. Train your own LoRAs on 20–50 well-tagged images of your subject — it takes 30 minutes on a 4090 and pays for itself the first time you need a consistent brand mascot.

// KEY TAKEAWAYS

- › ControlNet = structure; IP-Adapter = style; LoRA = concept.
- › Combine with restraint — start at 0.6 weights.
- › Train custom LoRAs on 20–50 tagged images.

[LESSON 04]

Production-Grade Image Pipelines

A production pipeline isn't just generation — it's batching, watermarking, brand-safety filtering, metadata tagging, and storage. Wire ComfyUI behind a FastAPI server, queue jobs with Redis, store outputs in S3/R2 with embedded EXIF (prompt, model, version, operator), and run every output through a CLIP-based brand-safety classifier before release. This catches the 1-in-500 generation that would have embarrassed the brand.

Set up automated regression tests. Pick 20 'canary' prompts that represent your brand style, run them weekly, and diff against last week's outputs. When a model update silently changes results — which happens — you catch it in hours, not after a campaign launches. This single discipline separates serious image teams from the ones constantly fighting drift.

// KEY TAKEAWAYS

- › Pipeline = generate + tag + filter + store.
- › Canary prompts catch silent model drift.
- › EXIF metadata = your audit trail.



[NEXT]

Keep Going

You've completed this course material. The real learning starts when you apply what you've read. Pick one idea from this PDF and run a small experiment this week. Document what worked and what didn't. Share your findings with the community.

Explore more free courses, daily AI tips, and curated tools at:

innovationailab.com

Have feedback or want to suggest a topic? We read every message.

hello@innovationailab.com

— Innovation AI Lab Team —

// part of LumiLife Tech