

ComfyUI Won – But A1111 Users Should Switch to Forge Neo Instead

February 2, 2026 · by Mark Bartlett

[Download this guide as PDF](#)

Quick Answer: ComfyUI is the best choice for most people in 2026. It's the fastest, uses the least VRAM, has the best Flux support, and gets new model support first. The learning curve is real – expect 2-3 weeks before you're comfortable – but the Desktop app has made installation painless. If you want A1111's familiar interface but need Flux support and better performance, use Forge Neo (a maintained fork). Fooocus is the easiest to start with but only supports SDXL and is no longer actively developed.

 **More on this topic:** [Stable Diffusion Locally](#) · [Flux Locally](#) · [What Can You Run on 8GB VRAM](#) · [Planning Tool](#)

If you want to generate images locally, the first real decision isn't which model to use – it's which interface. There are three main options, and they're very different from each other.

This guide compares them honestly: what each does well, where each falls short, and which one you should install based on what you actually want to do.

The Three Options (Plus One Fork)

ComfyUI – A node-based workflow editor. Think visual programming for image generation. Steep learning curve, but the most powerful and fastest option. 102k GitHub stars, very actively developed.

Automatic1111 (A1111) – The original Stable Diffusion web UI. Traditional form-based interface with dropdowns and sliders. 160k GitHub stars, but development has effectively stalled. No Flux support.

Forge – A fork of A1111 by llyasviel (the creator of ControlNet). Same familiar interface, 30-75% faster, less VRAM, and adds Flux support. The original Forge is dormant, but **Forge Neo** is actively maintained by the community.

Foocus – A minimal, Midjourney-like interface by the same developer as Forge. Zero configuration required. SDXL only, maintenance mode, no longer adding features.

Head-to-Head Comparison

	ComfyUI	A1111	Forge / Forge Neo	Foocus
Interface	Node-based (visual graph)	Traditional (forms/sliders)	Traditional (same as A1111)	Minimal (prompt + generate)
Speed (SDXL 1024px)	~8 sec	~11 sec	~5-6 sec (8GB GPU)	~18-20 sec
VRAM (SDXL)	~9.2 GB	~10.7 GB	~8-9 GB	~4 GB minimum
Flux support	Excellent (first-class)	None	Yes (NF4, GGUF)	None
Video generation	Yes (Wan 2.2, SVD)	No	Yes (Forge Neo)	No
Active development	Very active (weekly releases)	Stalled	Forge Neo: active	Maintenance only
Learning curve	2-4 weeks	2-3 hours	2-3 hours	Minutes
Workflow sharing	Yes (JSON, embedded in images)	No	No	No
API access	Yes (HTTP + WebSocket)	Yes (<code>--api</code> flag)	Yes	No
GPU support	NVIDIA, AMD, Intel, Apple Silicon	NVIDIA primary, AMD partial	NVIDIA only	NVIDIA only

ComfyUI: The Power User's Choice

ComfyUI is a visual workflow editor where you connect nodes – prompt, model loader, sampler, VAE decoder – into a pipeline. It looks intimidating at first. It's worth learning anyway.

Why People Switch to ComfyUI

It's the fastest. ComfyUI generates SDXL images 25% faster than A1111 and handles complex workflows (ControlNet + upscaling) 60% faster. Batch generation of 50 images: 28 minutes vs A1111's 38 minutes.

It uses less VRAM. Smart memory management loads models only when needed and unloads immediately after. SDXL peaks at ~9.2 GB vs A1111's ~10.7 GB. With `--lowvram` mode, you can run SDXL on a 6 GB card. Flux GGUF Q4 runs on **8 GB VRAM**.

New models land here first. When Flux launched, ComfyUI had support on day one. Same for Wan 2.2 video, Hunyuan, SD3.5 – if it's new, ComfyUI gets it first. A1111 sometimes catches up in weeks, sometimes months, sometimes never.

Workflows are shareable and reproducible. Every workflow saves as JSON. Every generated image embeds the full workflow as metadata – drag a ComfyUI-generated PNG back into the editor to get the exact setup that created it.

The Learning Curve Is Real

The node interface is genuinely confusing at first. Community surveys report 2-3 weeks before most people feel comfortable, and 20-30 hours of use before matching their A1111 skill level.

The good news: the **ComfyUI Desktop app** (released late 2024, improved throughout 2025) bundles Python, auto-updates, and comes with ComfyUI Manager pre-installed. Installation is now 3 clicks.

ComfyUI Manager is essential – it auto-detects missing custom nodes in shared workflows and installs them with one click. Install this first.

Best starter workflows

- Drag any ComfyUI-generated image into the editor to load its workflow
- Click "Load Default" on first launch for a working txt2img setup
- Browse community workflows at OpenArt, Civitai, or comfyworkflows.com

Who Should Use ComfyUI

- Anyone serious about image generation in 2026
- Users who want Flux, video generation, or the latest models
- Production/API workflows and automation
- Users on 6-8 GB GPUs who need aggressive VRAM optimization

Install

Desktop app (recommended): Download from comfy.org – Windows and macOS. Handles everything.

Manual:

```
git clone https://github.com/comfyanonymous/ComfyUI.git
cd ComfyUI
pip install -r requirements.txt
python main.py
```

Automatic1111: The Legacy Standard

A1111 is what most people started with. It has a traditional web interface — type a prompt, adjust some sliders, click generate. It's easy to understand and has the most tutorials of any tool.

What It Still Does Well

Easiest to learn. If you've used any web form, you can use A1111. No nodes, no graphs, no visual programming. The extension ecosystem adds features through a familiar install-and-enable model.

Massive knowledge base. 160k GitHub stars means every question has been asked and answered. YouTube tutorials, Reddit guides, extension documentation — there's more help available for A1111 than any other tool.

Good SD 1.5 and SDXL support. If all you need is Stable Diffusion 1.5 or SDXL with ControlNet, LoRA, and inpainting, A1111 still works fine.

Why People Are Leaving

No Flux support. This is the dealbreaker. Flux is the best open image model available, and A1111 doesn't support it. No native support exists and none is planned.

Development has stalled. The last meaningful update was mid-2024 (v1.10.0). Bug reports from January 2026 go unanswered. Many extensions are unmaintained.

It's the slowest and most VRAM-hungry option. 33% slower than ComfyUI on identical tasks. Peaks at 10.7 GB for SDXL where ComfyUI needs 9.2 GB. Struggles on 8 GB cards without special flags.

Who Should Still Use A1111

- You have an existing A1111 setup with specific extensions you depend on
- You only generate SD 1.5 or SDXL images and don't need Flux
- You're not ready to learn nodes and don't want Forge's forks

Honest advice: if you're starting fresh, don't install A1111. Use ComfyUI or Forge Neo instead.

Forge / Forge Neo: A1111 but Better

Forge is what A1111 should have become. Same interface, dramatically better performance.

What Forge Gives You Over A1111

GPU	Speed Improvement	VRAM Reduction
RTX 3070 (8 GB)	30-45% faster	700 MB - 1.3 GB less
RTX 3060 (6 GB)	60-75% faster	800 MB - 1.5 GB less
RTX 4090 (24 GB)	3-6% faster	1-1.4 GB less

The VRAM improvements are most dramatic on smaller GPUs. Without any special flags, Forge runs SDXL on 4 GB VRAM and SD 1.5 on 2 GB VRAM — the old A1111 `--medvram` and `--lowvram` flags are eliminated entirely.

And critically: **Forge supports Flux**. NF4 quantized Flux runs 1.3-4x faster than FP8 on 6-8 GB GPUs. See our [Flux guide](#) for model setup details.

The Fork Situation

The original Forge by llyasviel hasn't been updated since early 2025. But the community picked it up:

- **Forge Neo** (by Haoming02) — the recommended fork. Actively maintained, supports Wan 2.2 video, Flux-Kontext, and newer models. NVIDIA only. Has one-click Windows installers.
- **reForge** (by Panchovix) — another active fork with additional samplers and Python 3.12 support.

If you go the Forge route, **install Forge Neo**, not the original.

Who Should Use Forge

- A1111 users who want Flux support without learning nodes
- Users who prefer traditional interfaces over node-based editors
- Anyone with a 6-8 GB GPU who needs better VRAM management

Install

```
git clone https://github.com/Haoming02/sd-webui-forge-classic.git
cd sd-webui-forge-classic
git checkout neo
# Windows: run webui-user.bat
# Linux: run webui.sh
```

Or use the one-click Windows installer from the Forge Neo releases page.

Foocus: The “Just Works” Option

Foocus strips everything down to the essentials: type a prompt, click generate, get an image. No samplers to choose, no CFG to tune, no steps to set. It handles all of that internally.

When Foocus Makes Sense

- You want to test whether AI image generation interests you at all
- You have 10 minutes and want to see results immediately
- You need a dead-simple SDXL generator and nothing more

When It Doesn't

- You want to use Flux (not supported)
- You want ControlNet, custom workflows, or API access
- You want a tool that's still being developed

Foocus is in maintenance mode. SDXL only. No new features coming. Community forks (FoocusPlus, SimpleSDXL) add Flux support, but at that point you're better off with ComfyUI or Forge Neo.

Install

Windows: Download the release ZIP from GitHub, extract, run `run.bat`. Models download automatically on first launch. 4 GB VRAM minimum.

Feature Support Matrix

Feature	ComfyUI	A1111	Forge Neo	Foocus
SD 1.5	Yes	Yes	Yes	No
SDXL	Yes	Yes	Yes	Yes
Flux	Yes	No	Yes	No
SD3/SD3.5	Yes	Partial	Yes	No
Video (Wan 2.2, SVD)	Yes	No	Yes	No
ControlNet	Yes (all models)	Yes (extension)	Yes	Internal only
LoRA	Yes	Yes	Yes	Yes (SDXL)
Inpainting	Yes	Yes	Yes	Yes
Img2Img	Yes	Yes	Yes	Limited
Upscaling	Yes	Yes (extension)	Yes	Yes
Batch generation	Yes	Yes	Yes	Limited
API	Yes (HTTP/WS)	Yes (<code>--api</code>)	Yes	No
Workflow sharing	Yes (JSON)	No	No	No

Performance & VRAM Comparison

Generation Speed (SDXL 1024x1024, 20 steps)

Tool	Time	vs ComfyUI
Forge	~5-6 sec (8GB GPU)	Faster on low VRAM
ComfyUI	~8.2 sec	Baseline

Tool	Time	vs ComfyUI
A1111	~10.9 sec	33% slower
Foocus	~18-20 sec	~2.4x slower

VRAM Usage (SDXL)

Tool	Peak VRAM	Minimum Workable
Forge	~8-9 GB	4 GB
ComfyUI	~9.2 GB	4 GB (with <code>-lowvram</code>)
A1111	~10.7 GB	8 GB (with <code>-medvram-sd-xl</code>)
Foocus	Varies (swap)	4 GB

Flux VRAM Requirements (ComfyUI / Forge)

Flux Format	VRAM Needed
FP16	24+ GB
FP8	12+ GB
NF4 / GGUF Q4	6-8 GB

For Flux model download and setup details, see our [Flux guide](#).

Choose This If...

Choose ComfyUI if:

- You want the fastest, most capable tool available
- You plan to use Flux, video generation, or future models
- You're willing to invest 2-3 weeks in the learning curve
- You want reproducible, shareable workflows

Choose Forge Neo if:

- You want A1111's familiar interface with better performance
- You need Flux support but don't want to learn nodes

- You have a 6-8 GB GPU and need the VRAM savings
- You're migrating from A1111 and want minimal friction

Choose Fooocus if:

- You want to generate your first AI image in under 10 minutes
- You only need SDXL and don't care about Flux
- You want zero configuration and zero learning curve
- You understand it's a dead end for future features

Choose A1111 if:

- You already have a working setup with extensions you depend on
- You only use SD 1.5 or SDXL and don't need Flux
- Honestly, there aren't many good reasons to choose A1111 for a new install in 2026

The Practical Path

Most people in the community follow this progression: start with Fooocus or Forge to see quick results, then migrate to ComfyUI when they want more control. If you know you'll stick with image generation, skip straight to ComfyUI – the Desktop app makes installation easy, and the time invested in learning nodes pays off quickly.

Stability Matrix (a free launcher) lets you install ComfyUI, Forge, A1111, and Fooocus from one GUI and share model files between them. It's the easiest way to try everything without committing.

Bottom Line

The local image generation landscape has consolidated around ComfyUI. It's faster, uses less VRAM, supports every model, and gets updates weekly. The learning curve is its only real downside, and the Desktop app has made that significantly less painful.

If you're choosing today:

1. **Start with ComfyUI Desktop** – install, load the default workflow, generate your first image
2. **Install ComfyUI Manager** – it comes pre-installed in Desktop; essential for custom nodes
3. **Follow our [Stable Diffusion getting started guide](#)** for model setup
4. **When you're ready for Flux**, follow our [Flux guide](#)

If nodes genuinely aren't for you, Forge Neo gives you 90% of the capability with A1111's familiar interface. That's a perfectly good choice too.

Get notified when we publish new guides.

[Subscribe – free, no spam](#)

Source: <https://insiderllm.com/guides/comfyui-vs-automatic1111-vs-fooocus/>

Free guides for running AI locally