

# Best Photorealism Checkpoints for Local Image Generation (2026)

March 1, 2026 · by Mark Bartlett

[Download this guide as PDF](#)

**Quick Answer:** For SDXL on 8GB+ VRAM: Juggernaut XL Ragnarok is the most versatile photorealism checkpoint (18M+ downloads, actively updated). RealVisXL V5.0 beats it for tight portrait work with a natural film look. For SD 1.5 on 4-6GB VRAM: Realistic Vision V6.0 still holds up. For Flux on 12GB+: base Flux.1 Dev is already photorealistic enough that you don't need a fine-tuned checkpoint – add XLabs RealismLora for the last 5% of skin detail. Download from CivitAI, drop into your checkpoints folder, and start generating.

 **Related:** [ComfyUI vs A1111 vs Fooocus](#) · [Stable Diffusion Locally](#) · [Flux Locally](#) · [AI Upscaling Locally](#)

There are hundreds of checkpoints on CivitAI claiming to be “the most photorealistic.” Most are mediocre merges of the same handful of models. A few are genuinely good. And which one to pick depends on your GPU, your subject matter, and whether you care more about speed or fine detail.

I've tested the top-downloaded photorealism checkpoints across SDXL, SD 1.5, and Flux, and ranked them by what they're actually good at, with the settings and VRAM numbers that most checkpoint lists leave out.

---

## Photorealistic vs Stylized: The 30-Second Version

---

A “photorealistic” checkpoint is a Stable Diffusion model fine-tuned on real photographs rather than illustrations, paintings, or anime. The training data shapes the output. Feed a model millions of stock photos and portraits, and it learns to produce images with natural skin texture, realistic lighting, proper shadow falloff, and accurate color grading.

Stylized checkpoints (anime, digital art, oil painting) are trained on different data. Using the wrong checkpoint for your goal is the most common beginner mistake. No amount of prompt engineering turns an anime checkpoint into a portrait photographer.

The architecture matters too. SD 1.5 checkpoints generate at 512px natively. SDXL generates at 1024px. Flux generates at 1024px with a 12B parameter transformer that understands prompts at a level SDXL can't match. Bigger model, better results, more VRAM.

## Quick Comparison Table

Checkpoint	Base	File Size	Min VRAM	Best For	CFG	Updated
<b>Juggernaut XL Ragnarok</b>	SDXL	6.6 GB	6 GB	Portraits, cinematic, all-around	3-7	Feb 2026
<b>RealVisXL V5.0</b>	SDXL	6.5 GB	8 GB	Portraits, analog film look	1-3	2024
<b>Realities Edge XL 3.0</b>	SDXL	6-7 GB	8 GB	Art/photo hybrid, flexibility	Varies	Iterating
<b>CyberRealistic XL V9</b>	SDXL	6.5 GB	8 GB	Clean product shots, portraits	3-5	Active
<b>Realistic Vision V6.0</b>	SD 1.5	2-4 GB	4 GB	Portraits on low VRAM	7	Feb 2025
<b>CyberRealistic Classic V5</b>	SD 1.5	2 GB	4 GB	Clean portraits, LoRA base	5-7	Final release
<b>Flux.1 Dev</b>	Flux	24 GB (12 GB Q8)	8 GB (Q8)	Everything photorealistic	3.5-4.5	Active
<b>Flux.1 Schnell</b>	Flux	24 GB (12 GB Q8)	6 GB (Q8)	Fast iteration, commercial	1-3	Active

## SDXL Checkpoints

These need 8GB VRAM minimum (6GB with Lightning variants). They generate at 1024px natively and are the current sweet spot between quality and hardware requirements.

### Juggernaut XL Ragnarok (v13) – Best All-Around

[CivitAI](#) | 18M+ downloads | By KandooAI / RunDiffusion

Juggernaut XL is the most downloaded SDXL checkpoint for a reason. Version 13 (Ragnarok, released February 2026) improved anatomy: hands and full-body poses are noticeably better than v11, and it handles both concise and long descriptive prompts well.

Where Juggernaut excels is flexibility. It does portraits, full-body shots, landscapes, cinematic compositions, and even stylized digital painting without needing to switch checkpoints. It's the "one checkpoint for everything" pick. Most other SDXL checkpoints specialize; Juggernaut doesn't have to.

The tradeoff is that it doesn't beat the specialists at their own game. RealVisXL produces more natural-looking portraits. CyberRealistic XL does cleaner commercial shots. Juggernaut is the best generalist, not the best at any single thing.

### Recommended settings:

Setting	Standard	Lightning
Sampler	DPM++ 2M Karras	DPM++ SDE Karras
CFG	3-7	1.5-2
Steps	30-40	5-7
Resolution	832x1216	832x1216
VAE	sdxl-vae-fp16-fix	Baked in

### Sample prompt:

Portrait of a woman in her 40s, natural window light, soft shadows on face, slightly messy hair, wearing a linen shirt, shallow depth of field, Canon EOS R5, 85mm f/1.4, grain

## RealVisXL V5.0 – Best for Portraits

[CivitAI](#) | [HuggingFace](#) | By SG161222

RealVisXL is the portrait specialist. It produces a natural analog film quality that Juggernaut can't quite match: images look like they came from a real camera sensor rather than a render engine. Skin has the right amount of subsurface scattering, eyes catch light naturally, and the color palette leans warm without being oversaturated.

V5.0 fixed the biggest complaint from V4: small faces in mid-distance shots used to lose detail. That's gone. Hand anatomy also improved.

The catch is CFG sensitivity. If you're coming from SD 1.5, you're used to CFG 7-8. RealVisXL wants CFG 1-3 on the Lightning variant and 5-7 on standard. Run it at CFG 8 and your output will look plastic and oversaturated. This trips up a lot of new users.

RealVisXL also pulls hard toward photorealism regardless of your prompt. Want stylized or painterly output from the same checkpoint? Won't happen. It's a one-trick model, but the trick is very good.

### Recommended settings:

Setting	Standard	Lightning (BakedVAE)
Sampler	DPM++ SDE Karras	DPM++ SDE Karras
CFG	5-7	1-3
Steps	30-50	6-11
Resolution	1024x1024 or 832x1216	832x1216
VAE	sdxl-vae-fp16-fix	Baked in

### Sample prompt:

Elderly man sitting on a park bench, overcast afternoon light, reading a newspaper, wool coat, pigeons nearby, natural skin wrinkles and age spots, Fujifilm X-T5, 56mm f/1.2, shallow depth of field

## Realities Edge XL 3.0 – Best for Versatile Styles

[CivitAI](#) | [HuggingFace](#)

Realities Edge is built differently from the others on this list. It's a heavy merge model: the creator incorporated community LoRAs directly into the checkpoint through block merging (MBW) across 17+ iterations. The result can swing from photorealistic portraits to oil paintings to sci-fi concept art depending on how you prompt it.

For photorealism specifically, it pairs well with skin and face detail LoRAs. The base output leans slightly more "artistic" than Juggernaut or RealVisXL. You need to prompt more deliberately to lock in a photographic look rather than drifting into painted territory.

The Lightning + Turbo variant generates in 3-4 steps with acceptable quality, making it the fastest photorealism-capable SDXL checkpoint for quick iteration.

### Recommended settings:

Setting	Value
Sampler	DPM++ 3M SDE Karras
Clip Skip	2
Steps	20-30 (standard), 3-4 (Lightning)
Resolution	768x1344 or 1024x1296

## CyberRealistic XL V9 – Best for Clean Commercial Shots

[CivitAI](#) | By Cyberdelia

CyberRealistic XL produces the cleanest output of any SDXL photorealism checkpoint. Where Juggernaut and RealVisXL add character and grain, CyberRealistic gives you polished, evenly-lit images that look like they came from a commercial photography studio.

That's its strength for product shots, headshots, and anything that needs a "stock photo" level of polish. It's also the most LoRA-friendly SDXL checkpoint I've tested. The creator designed it specifically to play well with external LoRAs and textual inversions.

The weakness is personality. CyberRealistic images can feel sterile if you're going for editorial or documentary-style photography. The lighting is almost too consistent. If you want images with mood and atmosphere, Juggernaut or RealVisXL will serve you better.

### Recommended settings:

Setting	Value
Sampler	DPM++ 2M SDE Karras
CFG	3-5
Steps	30+
Resolution	832x1216 or 896x1152
Clip Skip	1
VAE	Baked in

## SD 1.5 Checkpoints – Still Relevant on Low VRAM

---

SDXL needs 8GB minimum. If you have a 4-6GB GPU (GTX 1060, 1650, RTX 3050), SD 1.5 is still your path to photorealistic local generation. The results won't match SDXL at the same resolution, but with hires.fix and the right settings, they're surprisingly good.

The creators of every major SD 1.5 checkpoint have moved on to SDXL or Flux. But the models still work, the LoRA ecosystem is massive (thousands of specialized LoRAs that don't exist for SDXL), and generation is fast even on older cards.

### Realistic Vision V6.0 – The SD 1.5 Standard

[CivitAI](#) | [HuggingFace](#) | By SG161222

Same creator as RealVisXL, and you can tell. Realistic Vision is the SD 1.5 equivalent: natural skin tones, realistic eyes, good facial detail at 512px. It's the most downloaded SD 1.5 photorealism checkpoint (1.7M+ downloads, 8,500+ reviews rated "Overwhelmingly Positive") and the community consensus pick.

Why people still use it in 2026: it runs on 4GB VRAM, the LoRA ecosystem is unmatched, and generation takes 1-2 seconds on a decent card. If your GPU can't run SDXL, this is the checkpoint.

The limitation is resolution. SD 1.5 generates at 512x768 natively. For anything printable, you need hires.fix to upscale to 1024x1536 or beyond, which doubles generation time. Full-body shots at 512px lack facial detail: use ADetailer or inpaint faces after generation.

#### Recommended settings:

Setting	Value
Sampler	DPM++ 2M Karras
CFG	7
Steps	20-30
Resolution	512x768 base, hires.fix to 1024x1536
VAE	vae-ft-mse-840000-ema-pruned

#### Sample prompt:

(RAW photo:1.2), close-up portrait of a young man, natural sunlight, freckles, green eyes, messy brown hair, slightly overexposed background, Nikon D850, 105mm macro, f/2.8

## CyberRealistic Classic V5.0 – Clean and Stable

[CivitAI](#) | By Cyberdelia

CyberRealistic Classic V5.0 is the final SD 1.5 release from Cyberdelia (development moved to the SDXL version). It produces clean, controlled photorealistic output with minimal prompting: good for users who want consistent results without fiddling with settings.

Where it differs from Realistic Vision: CyberRealistic Classic has slightly cooler tones and more even lighting. Realistic Vision leans warmer and more organic. Pick whichever matches your preference. Both are mature, stable, and well-tested with the SD 1.5 LoRA ecosystem.

---

## Flux – When Your VRAM Allows It

If your GPU can handle it, Flux makes SDXL checkpoints feel like workarounds. Base [Flux.1 Dev](#) (12 billion parameters, transformer architecture) is more photorealistic out of the box than any SDXL checkpoint with any amount of fine-tuning. Skin texture, lighting, hand anatomy, text rendering, prompt comprehension: Flux handles all of it better because the model is 5x larger and was trained on a bigger, more curated dataset.

The reason SDXL needs fine-tuned checkpoints for photorealism is that base SDXL (2.6B parameters) has a noticeable “AI look,” with slightly plastic skin, inconsistent lighting, and bad hands. Fine-tuned checkpoints like Juggernaut compensate for those weaknesses. Flux doesn’t have them to begin with.

### Do You Need a Flux Photorealism Fine-Tune?

For most people, no. Base Flux.1 Dev already passes the “does this look like a photograph” test. The community has focused on subject and style LoRAs rather than “make it more realistic” checkpoints, because the base model doesn’t need that push.

A few options exist if you want to squeeze out extra realism:

**LoRAs (applied on top of Flux.1 Dev):**

- [XLabs RealismLora](#) – the standard Flux realism LoRA. Enhances skin pore detail and subtle lighting. Weight 0.7-0.9.
- [Flux Super Realism LoRA](#) – pushes face and skin realism further than XLabs.
- [UltraRealistic LoRA Project V2](#) – good anatomy stability and flexible quality control.

**Full checkpoint fine-tunes:**

- [UltraReal Fine-Tune V4](#) – one of the few actual Flux checkpoint fine-tunes (not just a LoRA). Exceptional shadow and lighting detail. Works well as a base for stacking additional LoRAs.

**Flux VRAM Reality Check**

Format	File Size	VRAM (Comfortable)
Flux.1 Dev bf16	24 GB	20-24 GB
Flux.1 Dev GGUF Q8	12 GB	12-16 GB
Flux.1 Dev GGUF Q4	8 GB	8-10 GB
Flux.1 Schnell bf16	24 GB	16 GB
Flux.1 Schnell GGUF Q8	12 GB	8-12 GB

If you have 12GB+ VRAM, Flux in GGUF Q8 is the best photorealism option available for local generation. Below 12GB, stick with SDXL checkpoints. Below 8GB, SD 1.5.

**Flux.1 Dev recommended settings:**

Setting	Value
Sampler	Euler
CFG	3.5-4.5
Steps	20-28
Resolution	1024x1024

**Sample prompt (Flux responds best to natural language, not keyword lists):**

A woman in her 30s sitting at a cafe table in Paris, afternoon golden hour light coming through the window, cup of espresso on the table, she's looking slightly off camera with a relaxed expression,

wearing a dark blue blazer, shallow depth of field with the street blurred behind her

## VRAM Requirements by Checkpoint

This is the table other checkpoint guides skip. Your GPU determines which checkpoints you can actually run.

Your VRAM	Best Checkpoint	What to Expect
4 GB	Realistic Vision V6 (SD 1.5)	512x768 base, hires.fix for larger. Slow but works.
6 GB	Juggernaut XL Lightning	SDXL at 832x1216 in 5-7 steps. Tight fit.
8 GB	Juggernaut XL or RealVisXL	Full SDXL at 832x1216, 30-40 steps. Comfortable.
10 GB	Flux.1 Dev GGUF Q4	Flux with some quality loss from quantization.
12 GB	Flux.1 Dev GGUF Q8 or SDXL + LoRAs	Flux at near-full quality. SDXL with LoRA stacking.
16 GB	Flux.1 Dev GGUF Q8 (comfortable)	Flux with room for LoRAs and larger resolutions.
24 GB	Flux.1 Dev bf16	Flux at full precision. Best possible local output.

## Where to Download and How to Install

### Download Sources

**CivitAI** ([civitai.com](https://civitai.com)) – Largest selection of checkpoints and LoRAs. Most models have preview images so you can see output quality before downloading. Free account required.

**HuggingFace** ([huggingface.co](https://huggingface.co)) – Official releases from model creators. Better for Flux models and official variants. No account required for most downloads.

Download the `.safetensors` file (not `.ckpt`). Safetensors is the safer format).

### Installation: ComfyUI

Drop files into these folders inside your ComfyUI directory:

```
ComfyUI/models/checkpoints/ ← checkpoint .safetensors files
ComfyUI/models/vae/         ← VAE files
ComfyUI/models/loras/       ← LoRA files
```

Restart ComfyUI. The checkpoint appears in the “Load Checkpoint” node dropdown.

If you also use A1111 or Forge and don’t want duplicate 6GB files, edit `ComfyUI/extra_model_paths.yaml` (rename from `.yaml.example`) and point it to your A1111 models folder. Both UIs share the same files.

## Installation: A1111 / Forge

```
stable-diffusion-webui/models/Stable-diffusion/ ← checkpoints
stable-diffusion-webui/models/VAE/             ← VAE files
stable-diffusion-webui/models/Lora/            ← LoRA files
```

Restart the web UI. Select the checkpoint from the dropdown in the top-left. If you’re using Forge, the folder structure is identical. For more on choosing between UIs: [ComfyUI vs A1111 vs Fooocus](#).

## The Right VAE Matters

Using the wrong VAE is one of the most common reasons SDXL images look washed out, pink-tinted, or have strange color artifacts.

**For SDXL checkpoints:** Use [sdxl-vae-fp16-fix](#) by madebyollin. The default SDXL VAE produces NaN errors and corrupted pixels when running in fp16 (which most consumer GPUs default to). This fixed version is numerically stable, runs faster, and uses less VRAM. Output is visually identical to the original at fp32.

Exception: checkpoints labeled “BakedVAE” (like RealVisXL Lightning and CyberRealistic XL V9) already have the VAE built in. Don’t load an external VAE on top of these.

**For SD 1.5 checkpoints:** Use `vae-ft-mse-840000-ema-pruned`. Most SD 1.5 photorealism checkpoints expect this VAE. If your checkpoint filename includes “noVAE,” you need to load this separately.

**Important:** SD 1.5 VAEs and SDXL VAEs are not interchangeable. Using the wrong one produces garbled output.

## LoRA Stacking for Extra Realism

---

LoRAs are small add-on models (typically 50-200MB) that modify a checkpoint's output. For photorealism, skin detail and lighting LoRAs make the biggest difference.

### Recommended LoRAs for SDXL

**Detail Tweaker XL** – The most universal SDXL LoRA. Works with any checkpoint. Positive weight adds detail (pores, fabric texture, hair strands). Negative weight smooths. Start at weight 1.5, adjust from there. Range: -3 to +3.

**PhotorealTouch SDXL V2** – Adds high-frequency skin details: pores, vellus hair, natural imperfections. Makes skin look photographed rather than rendered. Weight 0.6-0.8.

**Skin Realism SDXL** – More aggressive than PhotorealTouch. Adds acne, moles, pores, and natural imperfections. Good for editorial and documentary-style realism where perfection looks fake.

### Stacking rules

- **Keep it to 2-3 LoRAs max.** Beyond three, competing weight signals start degrading coherence. Faces distort, textures conflict.
- **Total combined weight under 1.5.** If you're using Detail Tweaker at 1.0 and PhotorealTouch at 0.7, that's already 1.7. Back one of them down.
- **Test each LoRA alone first.** Add one at weight 0.5, generate, see what it does with your checkpoint. Then add the next. Don't stack three untested LoRAs and wonder why the output looks wrong.
- **Detail/texture LoRAs first, style LoRAs second.** In some implementations, load order affects output.

---

## Common Mistakes That Kill Photorealism

---

**CFG too high.** The number one mistake. SD 1.5 users are trained on CFG 7-8. SDXL photorealism checkpoints want CFG 3-7. RealVisXL wants CFG 1-3 on Lightning. Flux wants CFG 3.5-4.5. High CFG on these models produces oversaturated, plastic, obviously-AI output. If your images look "too perfect" and unnaturally vibrant, drop CFG first.

**Wrong or missing VAE.** If SDXL output looks washed out or has weird color bands, you're probably running the broken default VAE in fp16. Load [sdxl-vae-fp16-fix](#). Takes 30 seconds. Fixes everything.

**Keyword-stuffing prompts.** SDXL and Flux were trained on natural language captions, not the "masterpiece, best quality, ultra realistic, 8k, raw photo, detailed skin" keyword lists that SD 1.5 LoRAs taught everyone to write. On newer models, keyword stuffing can actively hurt output. Describe the scene as if you're talking to a photographer, not tagging a Danbooru image.

**Wrong sampler for the architecture.** Euler a was popular for SD 1.5. On SDXL, DPM++ 2M Karras or DPM++ 2M SDE Karras at 20-30 steps gives better results. On Flux, plain Euler at 20-28 steps is what Black Forest Labs recommends.

**Ignoring face detail in full-body shots.** SDXL generates at 1024px, but a face in a full-body shot might only occupy 80x80 pixels. That's not enough for realistic detail. Use ADetailer (A1111/ Forge extension) or face inpainting nodes in ComfyUI to fix faces after generation. This single step makes more difference than switching checkpoints.

**Wrong Clip Skip.** Most SDXL checkpoints want Clip Skip 1. Realities Edge XL wants Clip Skip 2. Using the wrong value shifts how the model interprets your prompt. Check the CivitAI page for your checkpoint; the creator usually specifies this.

## Sample Prompts: Same Scene, Different Architectures

Prompt style matters as much as the checkpoint. Each architecture responds to different prompting conventions:

### SD 1.5 (Realistic Vision) – tag-based with quality boosters:

```
(RAW photo:1.2), woman in her 30s, natural window light,
(minimal makeup:1.1), small smile, wearing a white cotton t-shirt,
blurred living room background, Canon EOS R5, 85mm, shallow depth
of field, film grain, 8k uhd
```

Negative: (worst quality, low quality:1.4), blurry, bad anatomy, extra fingers, text, watermark, cartoon, anime, illustration, 3d render

DPM++ 2M Karras, 30 steps, CFG 7, 512x768. Hires.fix: 4x-UltraSharp, 2x, denoise 0.35.

### SDXL (Juggernaut XL) – natural language with light tags:

Close-up portrait of a middle-aged man with weathered skin, salt-and-pepper stubble, deep laugh lines around his eyes, wearing a worn flannel shirt, standing on a dock at sunrise, warm golden light, Canon EOS R5, 85mm f/1.4, shallow depth of field, morning mist in background

Negative: cartoon, illustration, anime, blurry, low quality, watermark

DPM++ 2M SDE Karras, 35 steps, CFG 4, 832x1216.

### Flux.1 Dev – full conversational English, no negative prompt:

A candid photograph of a woman in her early 30s standing by a window in her apartment. Morning light is coming through sheer curtains, creating soft shadows on her face. She's wearing a simple white t-shirt and has a small, natural smile. The background is a slightly out-of-focus living room. Shot on Canon EOS R5 with an 85mm f/1.4 lens.

Euler, 25 steps, guidance 3.5, 1024x1024.

### Prompting tips that apply across all architectures:

Camera references help. "Shot on Canon EOS R5, 85mm f/1.4" produces more realistic output than "photorealistic." The models learned what real camera output looks like from training data tagged with EXIF info.

Add imperfections. "Skin pores, freckles, fine facial hair, flyaway hair, subtle wrinkles" break the "AI perfect" look that screams generated.

Lighting terms are specific. "Golden hour," "Rembrandt lighting," "soft diffused window light," "rim lighting" each produce distinct, realistic patterns. Vague terms like "good lighting" give vague results.

---

## Which Checkpoint Should You Install First

---

**Have 4-6GB VRAM?** Install Realistic Vision V6.0. It's the only good photorealism option at this tier.

**Have 8GB VRAM?** Install Juggernaut XL Ragnarok. It covers the most ground. Add RealVisXL later if you do a lot of portrait work and want that analog film quality.

**Have 12GB+ VRAM?** Install Flux.1 Dev in GGUF Q8 format. It beats everything else in this article for raw photorealism, and adding XLabs RealismLora gets you 95% of what dedicated Flux fine-tunes offer. Keep Juggernaut XL around for faster batch generation when Flux is overkill.

**Have 24GB VRAM?** Run Flux.1 Dev at full bf16 precision. No SDXL checkpoint competes at this level. Use SDXL when you want speed and don't need Flux-tier quality.

**Next steps:**

- [How to install and run Stable Diffusion](#) if you're starting from scratch
  - [ComfyUI vs A1111 vs Fooocus](#) to pick your interface
  - [Flux locally](#) for the full Flux setup walkthrough
  - [AI upscaling locally](#) to upscale your generations to print resolution
-

**seo: title: "Best Photorealism Checkpoints for Local Image Generation (2026) | InsiderLLM" meta\_description: "Juggernaut XL, RealVisXL, Realistic Vision, and Flux ranked for photorealism. VRAM requirements, settings, prompts, and install instructions for ComfyUI and A1111." slug: "best-photorealism-checkpoints-local-image-generation" primary\_keyword: "best photorealism checkpoint stable diffusion" secondary\_keywords: ["juggernaut xl vs realvisxl", "photorealistic stable diffusion settings", "best sdxl checkpoint 2026", "stable diffusion photorealism", "realistic ai images locally"] internal\_links: - topic: "ComfyUI vs A1111 vs Fooocus" anchor\_text: "ComfyUI vs A1111 comparison" - topic: "Stable Diffusion Locally" anchor\_text: "Stable Diffusion setup guide" - topic: "Flux Locally" anchor\_text: "Flux locally guide" - topic: "AI Upscaling Locally" anchor\_text: "AI upscaling guide" image\_alt\_texts: - "Side-by-side comparison of Juggernaut XL vs RealVisXL portrait output" - "VRAM requirements chart for photorealism checkpoints across SD 1.5, SDXL, and Flux" - "Screenshot of ComfyUI workflow with photorealism checkpoint loaded"**

---

Get notified when we publish new guides.

[Subscribe – free, no spam](#)

---

Source: <https://insiderllm.com/guides/best-photorealism-checkpoints-local-image-generation/>

Free guides for running AI locally