

Best Anime and Stylized Checkpoints for Local Image Generation (2026)

March 1, 2026 · by Mark Bartlett

[Download this guide as PDF](#)

Quick Answer: For SDXL anime on 8GB+ VRAM: NoobAI-XL (v-prediction variant) is the best overall anime checkpoint in 2026. It has the deepest character knowledge, best prompt adherence, and shares LoRAs with the Illustrious XL ecosystem. Animate XL 4.0 Opt is a close second with cleaner out-of-box output. Pony Diffusion V6 XL is still worth using for its massive LoRA library and furry/cartoon content. For SD 1.5 on 4-6GB VRAM: Anything V5 and Counterfeit V3 still work well. For Flux: the anime ecosystem is thin. Use SDXL instead unless you need Flux's hand accuracy or text rendering. For stylized non-anime (oil painting, watercolor, comics): ClassipeintXL (oil), Painter's Checkpoint (oil), Nova Comic XL (western comics), and style LoRAs on any SDXL base.

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

Photorealism checkpoints are fine-tuned on photographs. Anime checkpoints are fine-tuned on illustrations, typically scraped from Danbooru and similar image boards. The prompting is different, the quality tags are different, and choosing the wrong checkpoint for your goal wastes more time than any other mistake.

The anime checkpoint ecosystem is also more fragmented than the photorealism side. There are two major model families (Illustrious and Pony) with incompatible LoRA ecosystems, plus legacy SD 1.5 models that still have the largest variety of character LoRAs. Choosing a checkpoint means choosing an ecosystem, not just a model file.

I've tested the top anime and stylized checkpoints across SDXL, SD 1.5, and Flux, and ranked them by what actually matters: output quality, prompt adherence, LoRA compatibility, and VRAM requirements.

Quick comparison table

Checkpoint	Base	File Size	Min VRAM	Best For	Sampler	Status
NoobAI-XL (v-pred)	SDXL	6.5 GB	8 GB	Anime, widest character knowledge	Euler / DDIM	Final release
Illustrious XL v2.0	SDXL	6.5 GB	8 GB	Anime base model, LoRA training	Euler a	Active
Animagine XL 4.0 Opt	SDXL	6.5 GB	8 GB	Clean modern anime, polished output	Euler a	Active
Pony Diffusion V6 XL	SDXL	6.9 GB	8 GB	Anime/furry/cartoon, huge LoRA library	Euler a	Mature
Anything V5	SD 1.5	2 GB	4 GB	Classic anime on low VRAM	DPM++ SDE Karras	Legacy
Counterfeit V3	SD 1.5	2.1 GB	4 GB	Soft painterly anime, bold colors	DPM++ 2M Karras	Legacy
AbyssOrangeMix3	SD 1.5	2.1 GB	4 GB	2.5D anime-realism hybrid	DPM++ 2M Karras	Legacy
AnimePro FLUX	Flux	12+ GB	12 GB	Anime with accurate hands	Euler	Early
Painter's Checkpoint	SDXL	6.5 GB	8 GB	Oil painting style	DPM++ 2M Karras	Stable
Nova Comic XL	Illustrious	6.5 GB	8 GB	Western comic style	Euler a	Active

The Illustrious XL ecosystem

This is the most important section if you're choosing an anime checkpoint in 2026. Illustrious XL is an ecosystem, not a single model. The base checkpoint, its community fine-tunes, and their shared LoRA library all work together.

How we got here

OnomaAI Research released Illustrious XL v0.1 in mid-2024 as an intentionally unfinished base model trained on Danbooru data. The strategy worked. The community built hundreds of fine-tunes on top of it, and v0.1 became the foundation for the strongest anime models available today.

The lineage matters because LoRA compatibility follows it. LoRAs trained on Illustrious v0.1 work on v1.0 and most derivatives. LoRAs trained on Pony V6 do not work on Illustrious models, and vice versa. Choosing a checkpoint means choosing a LoRA ecosystem.

Illustrious XL v2.0 (the official base)

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

The latest official release from OnomaAI. Trained on Danbooru data through August 2024 with better natural language support than v0.1. Supports native resolution up to 1536x1536 (v1.0+ feature).

Illustrious XL v2.0 has the highest TrueSkill rating among official Illustrious releases. It understands both Danbooru tags and natural language prompts (about 50/50). Tag fidelity is the best of any anime checkpoint, and it recognizes more characters than Pony V6 because Danbooru's character tagging is more thorough.

In practice, most people don't run the base Illustrious model directly. They use one of the community fine-tunes below, which add aesthetic polish on top of Illustrious's strong foundation.

Recommended settings:

Setting	Value
Sampler	Euler a
Steps	28-35
CFG	5.5-7
Resolution	1024x1024 or 1248x1824
Clip Skip	2

NoobAI-XL: the community favorite

[CivitAI](#) | 6.46 GB (pruned fp16)

NoobAI-XL is built on Illustrrious v0.1 and trained on the complete Danbooru and e621 datasets, roughly 13 million images. That's 5x the training data of base Illustrrious, and it shows in character recognition and style range.

Two variants exist, and the choice matters:

Epsilon prediction (eps): Standard noise prediction. Compatible with Karras schedulers, most existing workflows, and the widest range of tools. The safe default.

V-prediction (v-pred): Newer noise schedule. Produces wider color gamut, stronger lighting and shadow contrast, and better prompt adherence. The output has a distinctive look that many users prefer. The catch: it does not support Karras schedulers. Use Euler or DDIM only.

For most people, v-pred is the better pick. The color and lighting improvements are immediately visible. Just remember to switch your sampler.

Recommended settings (v-pred):

Setting	Value
Sampler	Euler or DDIM (NOT Karras)
Steps	25-40
CFG	5-7
Resolution	1024x1024+

Sample prompt:

```
masterpiece, best quality, 1girl, silver hair, red eyes, school uniform,
sitting on rooftop, sunset, wind blowing hair, city skyline background,
detailed face, detailed hands
```

Other Illustrrious fine-tunes worth knowing

WAI-Illustrrious-SDXL (v16.0) ([CivitAI](#)) is a popular ready-to-use fine-tune with integrated VAE. Handles NSFW content well with rating tags. Same settings as base Illustrrious.

Hassaku XL is bright and style-agnostic. Good for users who want vivid colors without committing to a specific anime substyle.

Cat Tower (ChenkinNoob) is a NoobAI fine-tune that adds extra polish. If you like NoobAI's output but want slightly more refined detail, try this.

Animagine XL 4.0

[CivitAI](#) | [HuggingFace](#) | 6.5 GB

Animagine XL is developed by Cagliostro Research Lab. It produces the most polished output of any SDXL anime checkpoint without tuning. Version 4.0 was retrained from scratch on 8.4 million anime images with a knowledge cutoff of January 2025.

Three variants exist: the base v4.0, the **Opt** (optimized) version with improved anatomy and noise reduction, and the community-refined **Opt Clear** with higher contrast and richer colors. Install Opt Clear if you want the best version.

Animagine is a standalone model, not an ecosystem base. The LoRA library trained specifically for Animagine is smaller than Illustrious's. What you get instead is clean, modern anime art with less prompting effort. Drop in a character description and it looks good without much tuning.

The tradeoff is character knowledge. Animagine knows fewer specific characters than NoobAI-XL, and its Danbooru tag comprehension, while good, isn't as precise as Illustrious-based models. If you need to generate a specific obscure character, NoobAI will recognize it more often.

Recommended settings:

Setting	Value
Sampler	Euler a
Steps	25-28
CFG	5-7 (6 sweet spot)
Resolution	1024x1024 or SDXL buckets

Pony Diffusion V6 XL

[CivitAI](#) | [HuggingFace](#) | 6.9 GB + 335 MB VAE

Pony V6 was the dominant SDXL anime model through most of 2024. It's trained on 2.6 million images split roughly 50/50 between SFW and NSFW, covering anime, cartoon, furry, and anthro content. The result handles a wider range of art styles than any pure-anime checkpoint.

The score tag system

Pony V6 uses a unique quality scoring system instead of the standard “masterpiece, best quality” tags. You prepend score tags to your prompt:

```
score_9, score_8_up, score_7_up, source_anime, rating_safe,
1girl, blue hair, school uniform, classroom
```

There’s a known issue with this system. The model learned `score_9, score_8_up, score_7_up` as a single concatenated phrase rather than understanding each tag independently. Using just `score_9` alone has a weaker effect than using the full chain. It’s a training artifact that was recognized too late to fix.

You also need content-type tags (`source_anime` , `source_furry` , `source_cartoon`) and rating tags (`rating_safe` , `rating_questionable` , `rating_explicit`) to steer the output.

Why people still use it

The LoRA library. Pony V6 was dominant long enough that CivitAI has thousands of Pony-specific LoRAs for characters, styles, and concepts. If you need a LoRA for a specific character and it only exists for Pony, that’s your answer.

Pony V6 is also the strongest option for furry and anthro content. The Illustrious ecosystem is primarily anime-focused; Pony’s training data explicitly includes furry and cartoon content.

Where it falls short

Prompt adherence is weaker than Illustrious-based models. You need more generations to get what you asked for. Hands are worse than NoobAI-XL, and anime-specific character knowledge is shallower. The score tag system is also annoying compared to standard quality tags.

Recommended settings:

Setting	Value
Sampler	Euler a
Steps	25
CFG	7
Resolution	1024x1024

A note on Pony V7

Pony V7 exists but is a completely different architecture (AuraFlow, 7B parameters). V6 LoRAs don't work with it. The tooling ecosystem is immature. Unless you're specifically interested in AuraFlow, stick with V6 for now.

SD 1.5 anime checkpoints

These still have a place. If you have 4-6GB VRAM (GTX 1060, 1650, RTX 3050), SD 1.5 anime checkpoints are your only local option. They generate at 512px natively and need hires.fix for larger output, but the models are small, generation is fast, and the LoRA ecosystem is enormous.

Every major SD 1.5 anime checkpoint creator has moved to SDXL. These models are no longer updated. But they work, and they run on hardware that SDXL can't touch.

Anything V5

[CivitAI](#) | ~2 GB (pruned fp16)

The most popular SD 1.5 anime checkpoint. Clean linework, vibrant colors, good at a wide range of anime styles from cel-shaded to semi-realistic. Responds well to Danbooru tags and has strong character recognition for popular series.

Setting	Value
Sampler	DPM++ SDE Karras
CFG	6-8
Steps	20-30
Resolution	512x768 base

Counterfeit V3

[CivitAI](#) | ~2.13 GB (fp16 fix)

Counterfeit has a distinctive look: bold colors, soft shading, slightly desaturated palette. The output resembles visual novel CGs or light novel cover art more than typical anime illustration. If you want that dreamy, painterly quality at SD 1.5 resolution, this is the pick.

Setting	Value
Sampler	DPM++ 2M Karras
CFG	5-10
Steps	25-30

AbyssOrangeMix3

[CivitAI](#) | ~2.13 GB

AOM3 occupies a space between anime and realism. The output has a “2.5D” quality, with anime-style faces and more realistic lighting and texturing on clothing and backgrounds. Three variants exist (A1, A1B, A3) with different balances between anime and realism. A1 is the most anime-leaning.

Setting	Value
Sampler	DPM++ 2M Karras
CFG	6-7
Steps	20-30

MeinaMix V12

[CivitAI](#) | ~3.5 GB (baked VAE)

MeinaMix leans toward a clean, slightly stylized anime look with good backgrounds. The baked VAE means one less thing to configure. Good for landscape-heavy anime scenes.

Setting	Value
Sampler	DPM++ 2M Karras
CFG	4-9
Steps	25-30

Flux for anime: honest assessment

Flux is the best image model for photorealism. For anime, the picture is more mixed.

Where Flux wins

Hands. Flux produces accurate hands 80-90% of the time without any special prompting or LoRAs. On SDXL anime models, you're lucky to get 50% without hand-fixing workflows. If bad hands are your biggest frustration, Flux solves it.

Prompt adherence. Complex multi-part prompts with specific character positions, clothing details, and scene composition work more reliably on Flux than SDXL.

Text rendering. If you need text in your image (signs, logos, speech bubbles), Flux handles it. SDXL doesn't.

Where Flux falls short

The anime ecosystem barely exists. SDXL has thousands of anime checkpoints and tens of thousands of character LoRAs. Flux has a handful of fine-tunes and maybe a few hundred anime-related LoRAs. The 12B parameter count makes fine-tuning expensive, so the community hasn't produced the depth of options that SDXL has.

Photorealistic bias. Base Flux defaults hard toward photorealism. Even with anime LoRAs applied, the output often looks semi-realistic rather than pure illustrated. You have to fight the model's instincts. SDXL anime checkpoints were trained specifically to produce illustrated output and don't have this problem.

VRAM. Flux needs 12GB+ for comfortable anime generation. SDXL anime models run well on 8GB. On 6GB with optimizations, SDXL is workable; Flux is painful.

Speed. Flux is slower per image. When you're iterating on a prompt and generating dozens of variations, this adds up.

The practical options

AnimePro FLUX ([CivitAI](#)) is a Flux.1 Schnell fine-tune for anime. Works at 4-8 steps, ships in FP8 quantization. The most usable Flux anime checkpoint available.

Dark Katana LoRA ([CivitAI](#)) is the highest-rated anime LoRA for Flux with 20,000+ likes. Produces a dark, atmospheric cyberpunk-anime style. Trigger: "DarKata style", weight 0.8.

AnimEasy Flux ([CivitAI](#)) is available in 4-bit quantization for lower VRAM setups.

The verdict

For pure anime and illustrated styles, SDXL with Illustrrious or NoobAI is still the better choice in early 2026. The ecosystem is deeper, the models are purpose-built, and the VRAM requirements are lower. Use Flux for anime only if you specifically need accurate hands, text rendering, or complex compositions that SDXL struggles with. Most serious users maintain both workflows.

Stylized non-anime checkpoints

Not all illustration is anime. Here are the best options for other styles.

Oil painting

Painter's Checkpoint ([CivitAI](#)) is a dedicated SDXL oil painting checkpoint. Produces convincing brushwork, canvas texture, and color mixing. Good for landscapes and portraits in a classical painting style.

ClassipeintXL ([CivitAI](#)) is an SDXL LoRA rather than a full checkpoint. Apply it at weight 0.7-1.2 on any SDXL base for an oil painting effect. More flexible than a dedicated checkpoint since you keep your base model's other capabilities.

For Flux: **flux_oilscape** ([HuggingFace](#)) with trigger "in the style of Oilstyle002".

Watercolor

Watercolor Style SDXL ([CivitAI](#)) is a LoRA with trigger "ral-wtrclr" that adds watercolor texture and bleeding effects to SDXL output.

For Flux: **WATERCOLOR (FLUX)** ([CivitAI](#)) and **Expressive Watercolor Ink Wash** ([CivitAI](#)) both work well.

Ink and lineart

Lineart cross-platform LoRA ([CivitAI](#)) works across SD3.5, Flux, SDXL, and Pony. Weight 0.5-1.0 for clean black-and-white lineart.

LineAniRedmond ([CivitAI](#)) is specifically for manga-style lineart on SDXL, with modes for detailed, minimalist, and colorful output.

Western comics

Nova Comic XL v2 (CivitAI) is built on Illustrious and produces western comic-style art. The best option for superhero-style illustrations and graphic novel panels.

Comic Diffusion (CivitAI) is an SD 1.5 checkpoint for classic western comic art. Runs on 4GB VRAM.

Arthemy Comics Flux (CivitAI) brings western illustration style to Flux, still in beta.

VRAM requirements

Your GPU determines which checkpoints you can run. This table covers anime and stylized models specifically.

Your VRAM	Best anime checkpoint	What to expect
4 GB	Anything V5 or Counterfeit V3 (SD 1.5)	512x768 base. Fast generation, huge LoRA library. Hires.fix for larger output.
6 GB	Anything V5 (SD 1.5) or SDXL with -lowvram	SD 1.5 comfortable. SDXL possible but tight at 1024x1024.
8 GB	NoobAI-XL or Animate XL 4.0 (SDXL)	Full SDXL anime at 1024x1024. LoRA stacking works. The sweet spot.
10 GB	NoobAI-XL + LoRAs (SDXL)	SDXL with room for LoRA stacking and 1536px output.
12 GB	SDXL anime or Flux (FP8)	Flux anime possible with FP8 quantization. Quality compromise vs SDXL.
16 GB+	Flux (GGUF Q8) with anime LoRAs	Near-full-quality Flux with anime LoRAs. Best hands and composition.

Prompting for anime checkpoints

Anime checkpoints use Danbooru-style tags, not the natural language prompts that work best on Flux and photorealism models. Getting this right matters more than which checkpoint you pick.

Tag ordering

The standard Danbooru prompt structure:

```
[quality tags], [character count], [character name], [series],
[physical features], [clothing], [expression/pose], [setting/background]
```

Example:

```
masterpiece, best quality, 1girl, hatsune miku, vocaloid,
long turquoise twintails, detached sleeves, thigh boots,
singing, open mouth, concert stage, spotlight, crowd
```

Quality tags by model

Different checkpoints expect different quality tags. Using the wrong ones has no effect or actively hurts output.

Model	Quality tags	Notes
Illustrious / NoobAI	masterpiece, best quality	Standard Danbooru quality tags
Pony V6	score_9, score_8_up, score_7_up	Must use the full chain for best effect
Animagine XL 4.0	masterpiece, best quality	Also responds to very aesthetic
SD 1.5 models	masterpiece, best quality, highres	The classic set

Negative prompts

Negative prompts matter less on SDXL than SD 1.5, but they still help. A solid baseline for SDXL anime:

```
worst quality, low quality, normal quality, lowres, bad anatomy,
bad hands, extra fingers, fewer fingers, text, watermark, signature
```

For SD 1.5 anime, negative prompt embeddings make a bigger difference. Install these as textual inversions:

- **EasyNegative** covers most common artifacts
- **BadHandV4** specifically targets hand deformities
- **bad-image-v2-39000** is a general quality filter

These embeddings only work on SD 1.5. They have no effect on SDXL or Flux models.

Prompting differences from photorealism

Camera references (“Canon EOS R5, 85mm f/1.4”) that work well for photorealism do nothing useful on anime checkpoints. Instead, use art style tags: `detailed background`, `dynamic angle`, `dramatic lighting`, `from below`, `depth of field`.

Anime checkpoints respond to artist style tags (`[artist name] \ (style\)`) though I’d avoid relying on these since many have been removed from newer model training data for copyright reasons.

Character-specific tags are where anime checkpoints shine. If a character exists on Danbooru with enough tagged images, Illustrious and NoobAI will recognize them by name. Pony V6 is weaker here due to its smaller anime training subset.

LoRA recommendations for anime

Character LoRAs

The main reason people pick a specific checkpoint ecosystem. A character LoRA trained on Pony V6 only works on Pony V6 and its derivatives. Same for Illustrious. Check the LoRA’s page on CivitAI for which base model it was trained on.

CivitAI has the largest collection. Filter by base model (SDXL, Pony, Illustrious, SD 1.5) and sort by rating or downloads.

Style LoRAs worth trying

Detail Tweaker XL ([CivitAI](#)) works on any SDXL checkpoint. Positive weight adds detail (linework, texture). Negative weight smooths. Range -3 to +3, start at 1.0.

Flat Color LoRA for cel-shaded anime looks. Several versions exist for different base models.

Anime Screencap style LoRAs emulate the look of actual anime screenshots rather than illustration art. Useful for a different aesthetic than the “pixiv illustration” default.

Stacking rules

Same as photorealism: keep it to 2-3 LoRAs maximum. When stacking, reduce individual weights to 0.4-0.7 each. Test each LoRA alone first, then combine. Total combined weight under

1.5 avoids the face distortion and color bleeding that happens when competing LoRAs fight each other.

Where to download and how to install

Download sources

CivitAI (civitai.com) has the largest selection. Filter by “Anime” in the model type, and check the base model tag (SDXL, Pony, Illustrious, SD 1.5). Preview images show output quality before you download.

HuggingFace (huggingface.co) hosts official releases from model creators. Better for base models like Illustrious XL and Animate.

Always download `.safetensors` files, not `.ckpt`.

Installation: ComfyUI

```
ComfyUI/models/checkpoints/    <- checkpoint .safetensors files
ComfyUI/models/loras/          <- LoRA files
ComfyUI/models/embeddings/     <- negative prompt embeddings (SD 1.5 only)
```

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

Installation: A1111 / Forge

```
stable-diffusion-webui/models/Stable-diffusion/ <- checkpoints
stable-diffusion-webui/models/Lora/             <- LoRA files
stable-diffusion-webui/embeddings/             <- embeddings
```

For more on choosing between UIs: [ComfyUI vs A1111 vs Fooocus](#).

Which checkpoint should you install first

Have 4-6GB VRAM? Install Anything V5. It's the best SD 1.5 anime checkpoint and runs on anything.

Have 8GB VRAM and want anime? Install NoobAI-XL (v-prediction). It has the deepest character knowledge, the best prompt adherence, and access to the growing Illustrious LoRA ecosystem. If NoobAI's output feels too raw, try Animate XL 4.0 Opt Clear for more polished results out of the box.

Have 8GB VRAM and want stylized non-anime? Install your SDXL checkpoint of choice and add style LoRAs. ClasspaintXL for oil painting, Watercolor Style SDXL for watercolor, Nova Comic XL for western comics.

Have 12GB+ VRAM? You can run Flux with anime LoRAs, but the anime ecosystem is thin. Install NoobAI-XL for serious anime work and keep Flux for tasks that need accurate hands or text rendering. Maintaining both workflows is what most experienced users do.

Coming from Pony V6? Don't switch immediately if you have Pony-specific LoRAs you depend on. But for new projects, the Illustrious ecosystem produces better anime output with less prompting effort. Consider migrating gradually.

Next steps:

- [How to install and run Stable Diffusion](#) if you're starting from scratch
 - [ComfyUI vs A1111 vs Fooocus](#) to pick your interface
 - [Best photorealism checkpoints](#) for the companion to this guide
 - [AI upscaling locally](#) to upscale your generations
-

seo: title: "Best Anime and Stylized Checkpoints for Local Image Generation (2026) | InsiderLLM" meta_description: "Illustrious XL, NoobAI, Animate, Pony V6, and SD 1.5 anime checkpoints ranked. VRAM needs, Danbooru prompts, LoRA picks, settings for ComfyUI and A1111." slug: "best-anime-stylized-checkpoints-local-image-generation" primary_keyword: "best anime checkpoint stable diffusion" secondary_keywords: ["illustrious xl vs noobai", "pony diffusion v6 settings", "anime stable diffusion locally", "best sdxl anime model 2026", "danbooru prompting guide"] 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: "Best Photorealism Checkpoints" anchor_text: "photorealism checkpoint guide" - topic: "AI Upscaling Locally" anchor_text: "AI upscaling guide" image_alt_texts: - "Side-by-side comparison of NoobAI-XL vs Animate XL vs Pony V6 anime output" - "VRAM requirements chart for anime checkpoints across SD 1.5, SDXL, and Flux" - "Illustrious XL ecosystem diagram showing model lineage and LoRA compatibility"

Get notified when we publish new guides.

[Subscribe – free, no spam](#)

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

Free guides for running AI locally