



FREE GUIDE • FOR CREATORS

The Beginner's Guide to Local AI Image Generation

How creators generate images privately on their own machine — without cloud limits, confusing tools, or expensive trial-and-error.

● TITANXT

What's Inside

01	Why Local AI Is Taking Off	The shift away from cloud services
02	Cloud vs Local — What Actually Changes	An honest side-by-side comparison
03	What You Need to Get Started	GPUs, RAM and the good news about gaming PCs
04	Models Explained Without the Jargon	Checkpoints, LoRAs and VAEs in plain English
05	Prompting for Local AI Tools	Structure, negative prompts, seeds
06	Five Practical Workflows	Products, social, thumbnails, concepts, ads
07	Common Beginner Mistakes	Six pitfalls and their fixes
08	A Simpler Way to Work Locally	Where Logen fits in
09	Next Steps	Start generating today



CHAPTER 01

Why Local AI Image Generation Is Taking Off

Something has shifted. While millions still type prompts into cloud services like Midjourney, DALL-E and Leonardo, a growing number of creators are going local — running AI image models directly on their own computers.

Why? Because local generation solves problems that cloud services create:

- **Complete creative control.**
No content filters removing valid artistic choices. No style restrictions. No prompt censorship. Your machine, your rules.
- **Zero ongoing costs.**
Cloud services charge per image or per month. Local generation costs nothing after setup. Generate one image or ten thousand — same price.
- **Total privacy.**
Your prompts, your images, your creative process — none of it touches a third-party server. For businesses generating product concepts, this matters enormously.
- **Speed without queues.**
No waiting behind thousands of other users. Your GPU processes your request immediately.
- **Ownership and repeatability.**
Save exact settings, seeds and workflows. Reproduce results perfectly, without wondering if the service changed its model overnight.

THE BIGGER PICTURE

Two years ago, local AI image generation required command-line expertise and hours of configuration. Today, accessible tools put professional-grade generation within reach of any creator with a decent computer.



CHAPTER 02

Cloud vs Local — What Actually Changes

Cloud generation

You type a prompt into a web interface, their hardware processes it, you receive an image.

- Monthly subscriptions (£8–£60/month)
- Per-image credits that run out
- Moderation blocks legitimate work
- Images stored on third-party servers
- Queues at peak hours; unannounced model swaps

Local generation

Models run on your own GPU. Prompts processed locally, images saved on your machine.


- No recurring costs after hardware
- Complete privacy — nothing leaves your machine
- Full parameter control, reproducible seeds
- Works offline
- Needs a capable GPU + initial setup

The real-world comparison

Factor	Cloud (Midjourney / DALL-E)	Local AI
Monthly cost	£8–£60	£0
Per-image cost	Credit-based	Free
Privacy	Images on their servers	Everything stays local
Content freedom	Filtered / moderated	Unrestricted
Speed	Queue-dependent	GPU-dependent
Consistency	Updates change outputs	You control the model
Offline use	No	Yes
Learning curve	Low	Medium (and improving)

IN PRACTICE

Most serious creators end up using both — cloud for quick exploration, local for production work where control, privacy and cost matter.



CHAPTER 03

What You Need to Get Started

AI image generation runs on your graphics card, not your main processor. The key specification is **VRAM** — the GPU's dedicated memory.

Minimum — 6 GB: slower generation, smaller resolutions.

Recommended — 8–12 GB: comfortable for most workflows; standard images in seconds.

Ideal — 16 GB+: high-resolution output, larger batches, fast iteration.

Which GPUs work?

GPU	VRAM	Performance level
NVIDIA RTX 3060	12 GB	Good — solid entry point
NVIDIA RTX 3070 / 3080	8–10 GB	Good — fast generation
NVIDIA RTX 4060 Ti	16 GB	Excellent — handles everything
NVIDIA RTX 4070 / 4080	12–16 GB	Excellent — fast and capable
NVIDIA RTX 4090	24 GB	Overkill for most — professional speed
AMD RX 7900 XTX	24 GB	Good — improving support


NVIDIA is recommended — the CUDA ecosystem has the most mature software support. AMD works but needs extra configuration.

RAM, storage and OS

- **RAM:** 16 GB minimum, 32 GB recommended
- **Storage:** 50–100 GB free; models are 2–7 GB each — SSD strongly recommended
- **OS:** Windows 10/11 (best supported) · Linux (excellent, often faster) · Apple Silicon (growing support)

THE GOOD NEWS

If your computer can run modern games at medium-to-high settings, it can almost certainly run local AI image generation.



CHAPTER 04

Models Explained Without the Jargon

A model is a file (typically 2–7 GB) containing everything the AI has learned about generating images. Think of it like a camera lens: same body, different lens — different results.

The models you'll encounter

- **Stable Diffusion XL (SDXL)**
— the workhorse. High quality at 1024×1024, massive community support. Where most beginners start.
- **Flux**
— newer architecture. Excellent prompt-following and text rendering; wants 12 GB+ VRAM.
- **SD 1.5**
— older but useful. Runs on 4 GB VRAM with a huge library of community styles.

Checkpoints, LoRAs and VAEs

Checkpoint — a complete model file; what people mean by "download a model." CivitAI and Hugging Face are the main sources. Popular picks: **Juggernaut XL** (photorealism), **DreamShaper XL** (versatile), **RealVisXL** (portraits).

LoRAs — small add-ons (10–200 MB) that teach a base model specific things: an art style, a character, a product, a lighting technique. Stack several for combined effects.

VAEs — handle colour accuracy and sharpness. Modern models build them in; you only care when colours look washed out.

WHAT BEGINNERS ACTUALLY NEED

Start with **one good SDXL checkpoint** and explore LoRAs for the styles you want. Don't download 50 models on day one — learn one well before expanding.



CHAPTER 05

Prompting for Local AI Tools

Local tools pass your text more directly to the model than cloud services do. That means more control — and slightly more precision required.

Basic prompt structure

1. **Subject** — what's in the image (a woman, a landscape, a product)
2. **Style** — how it looks (photorealistic, oil painting, anime, minimalist)
3. **Details** — lighting, camera angle, time of day
4. **Quality boosters** — detailed, sharp focus, professional

EXAMPLE

"Product photograph of a ceramic coffee mug, white background, soft studio lighting, sharp focus, commercial photography, 4K, detailed texture"

Negative prompts

Uniquely powerful in local tools — they tell the model what to avoid, preventing common artifacts:

EXAMPLE NEGATIVE

"blurry, low quality, distorted, watermark, text, oversaturated, extra fingers, deformed"

Seeds and repeatability

Every image has a seed — same prompt + same seed + same settings = an identical image, every time. That's how you iterate on a composition, build consistent product series, and reuse successful outputs.

Practical tips

- **Be specific.** "A dog" is random; "a golden retriever puppy on a wooden floor, warm afternoon sunlight, shallow depth of field" is consistent quality.
- **Front-load** important elements — models weight the start of a prompt more heavily.
- **Separate concepts with commas** — each phrase gets weighted independently.
- **Iterate.** Your first prompt won't be perfect. Adjust, regenerate, refine.

Five Practical Local AI Workflows

1 Product mockups

1. Photograph your product with even lighting
2. Remove and replace the background with AI
3. Generate studio, lifestyle or seasonal contexts
4. Export at your platform's required resolution

Why local matters: 50 variations without credit limits — and unreleased designs stay private.

2 Social media images

1. Define your brand style as a prompt template + LoRA combo
2. Generate images matching your content calendar
3. Keep consistent aspect ratios per platform
4. Batch-generate a week's content in one session

Why local matters: consistent style via saved seeds; no subscription for social content.

3 YouTube & blog thumbnails

1. Create a base template style (colours, composition, energy)
2. Generate the background or hero image
3. Leave space for text overlays
4. Maintain series consistency with saved settings

Why local matters: test multiple concepts in minutes instead of waiting on a designer.

4 Concept art & moodboards

1. Describe the concept in prompt form
2. Generate variations with different seeds
3. Combine favourites into a moodboard
4. Refine with collaborator feedback

Why local matters: rapid iteration and complete privacy for unreleased ideas.

5 Ad creative testing

1. Define ad formats and safe zones
2. Generate multiple visual approaches per message
3. Export at platform specs and split-test
4. Double down on winners with more variants

Why local matters: 20 variants in an hour at zero per-image cost — test aggressively.

Common Beginner Mistakes

1 Installing too many tools

ComfyUI, Automatic1111, Fooocus and InvokeAI all at once — different formats, interfaces, learning curves.

Fix: Pick one tool. Learn it well. Expand later if needed.

2 Downloading every model

Dozens of CivitAI models before understanding any of them = confusion and wasted disk space.

Fix: One recommended SDXL checkpoint for a week. Learn its strengths before trying others.

3 Ignoring image organisation

A week in, you'll have hundreds of images and no idea which prompt made which.

Fix: Use a tool that saves generation metadata. Project folders and meaningful names from day one.

4 Overcomplicating prompts

200-word kitchen-sink prompts perform worse than focused ones.

Fix: Start short and specific. Add detail only when the output needs it.

5 Wrong resolution for the model


SDXL wants 1024×1024; SD 1.5 wants 512×512. Wrong resolutions produce artifacts.

Fix: Use your model's recommended resolution as baseline; upscale afterwards.

6 Giving up too early

Your first local images may trail cloud quality — cloud services quietly enhance prompts for you.

Fix: A week of daily practice. By week two you'll beat your old cloud results.



CHAPTER 08

A Simpler Way to Work Locally

The most powerful local tool — ComfyUI — uses a node interface built for technical users. Other tools simplify but sacrifice power. You end up choosing between "easy but limited" and "powerful but overwhelming."

What if you didn't have to choose?

Logen is a desktop application built for creators who want local AI power without the complexity:

- **One-click installation**
 - no Python, no command line, no dependency management
- **ComfyUI power under the hood**
 - same models, same quality, streamlined interface
- **Automatic model management**
 - download, organise and switch models without file wrangling
- **Guided workflows**
 - pre-built pipelines for product photos, portraits, concept art
- **Private by design**
 - everything runs on your machine; nothing phones home

Who Logen is for

- Designers who want AI as a creative tool, not a technical project
- Small business owners generating product and marketing imagery
- Content creators who need consistent, branded visuals
- Anyone frustrated by ComfyUI's complexity but unwilling to accept cloud limits

Next Steps

You now understand the landscape — how local generation works, what you need, and the pitfalls to avoid. The next step is action.

- **Try Logen**
— download free at titanxt.io/download and generate your first local AI image in minutes. No command line, no configuration.
- **Join the newsletter**
— weekly tutorials, workflow breakdowns, model recommendations and early access to new features.
- **Explore the blog**
— regular tutorials on specific workflows, model comparisons and creative techniques at titanxt.io/blog.

The creators adopting local AI today are building the skills and workflows that will define creative production for years. Cloud services will keep a place for quick exploration — but for serious work, local generation delivers control, privacy and economics no subscription can match.

Welcome to local AI. Your machine. Your images. Your rules.

Generate your first local AI image today

Logen is free to download for Windows and Linux. Install, pick a model, and create — the setup wizard handles everything else.

titanxt.io/download