

PLATE 01 CHAPTER A4

AI STUDIO · CHAPTER 4

Images at home

You can make your persona’s art on your own computer, with no monthly bill and

nothing sent to a service. What decides it is your graphics card, your disk and the licence of the model you pick.

ENTRY A4 CATEGORY AI STUDIO CHECKED 7 NOVEMBER 2026 READING 7 MIN
SECTIONS 6

EVIDENCE ◆ Verified ◆ Vivian's field notes ◆ Unverified

IN 30 SECONDS

01 Local generation runs on your graphics card. Its memory decides which models load and how fast they run.

02 ComfyUI is a free program that runs open image models on your own computer, even offline.

03 Every model comes with a licence. Some allow commercial use, some don't, and it can change between versions of the same model.

04 One model file can take 20 to 40 GB, so plan a drive for them.

05 Work in small rounds: describe, generate several, compare, change one thing, write down what worked.

YOU WILL LEARN

- ✓ What a GPU and its memory change for image generation.
- ✓ What ComfyUI is and why its files can carry your recipe.
- ✓ Which open model families make anime art, and what their licences allow.
- ✓ How much disk space local generation needs.

- ✓ A method for working with AI that gets you to the picture you want.

PART 01

What your computer needs

Local generation means the model runs on your own computer. Nothing is sent to a service, and there's no bill per image. The part that matters most is the graphics card, the GPU, and its own memory, called VRAM. More VRAM lets bigger models load in one piece, and that makes them faster. ◆

You need less than you'd think. ComfyUI says it can run even the biggest open models on 4 GB of VRAM and 8 GB of normal memory, by moving parts of the model in and out while it works. ◆ Its Windows app recommends a dedicated NVIDIA or AMD card, but it doesn't require one. ◆ Without a card, it runs on the processor, more slowly. ◆

What the makers say

8 GB

Stability AI says SDXL 1.0 works on consumer cards with 8 GB of VRAM.

◆

86%

Share of a 24 GB card that the lighter version of Qwen-Image used in ComfyUI's own test.

◆

Most model makers don’t publish a minimum card. ◆ In my experience, a small card works but each picture takes longer, and slow rounds make you test less. ◆

PART 02

What ComfyUI is

ComfyUI is free software that runs image models on your computer. You build the job as a graph of boxes joined by wires: load a model, write the prompt, generate, save. ◆ Its core runs fully offline and downloads nothing unless you ask. ◆ Some optional boxes call paid online models, and those can be switched off. ◆

ComfyUI can save a whole workflow and recover it, seeds included, from a picture it made. ◆ That’s handy as a recipe, and it’s also a leak: a file that carries its recipe shows your prompt and settings to anyone who opens it in the same tool. ◆ Post a clean export, never the original file. See Stay anonymous.

PART 03

The open model families

A model family is a base model and the versions built on it. For anime art, most people start from a family trained on anime pictures. ◆

◆	Made for	Licence of the model	The pictures you make
SDXL 1.0 (Stability AI)	General, and the base of many anime models	CreativeML Open RAIL++-M	The licensor claims no rights in them
Illustrious XL (OnomaAI)	Anime, trained on Danbooru2023	Changes by version	Posting them isn't restricted (first release)
Pony Diffusion V6 XL (PurpleSmartAI)	Anime and stylised art	Modified Fair AI Public License	Commercial use: ask the makers
FLUX.1 [dev] (Black Forest Labs)	General	Non-commercial	Allowed for commercial use
FLUX.1 [schnell]	General, fast	Apache 2.0	Allowed for commercial use
FLUX.2 [dev]	General and editing	Non-commercial, filters required	Allowed for commercial use, as its licence describes

	Made for	Licence of the model	The pictures you make
Qwen-Image (Alibaba)	General, anime included	Apache 2.0	Not stated separately on its card

SDXL. It's released under the CreativeML Open RAIL++-M licence. ◆ That licence claims no rights in the pictures you make, and makes you accountable for them. ◆

Illustrious XL. It's fine-tuned on the Danbooru2023 dataset of tagged anime pictures. ◆ Its first release, v0.1, uses a fair public AI licence that discourages using it to make money. ◆ The makers don't restrict posting the pictures, but you can't sell a closed, modified copy of the model. ◆ Later versions carry different licence tags: "sdxl-license" for v1.0, "creativeml-openrail-m" for v2.0. ◆

Pony Diffusion V6 XL. Its licence is a modified Fair AI Public License. ◆ It can't run on any website or app that makes money, and commercial use means writing to its makers first. ◆

FLUX. The FLUX.1 [dev] model is for non-commercial use only. ◆ The pictures it makes can still be used for any purpose, commercial included, unless the licence forbids it elsewhere. ◆ Using the model itself to earn money doesn't count as non-commercial. ◆ FLUX.1 [schnell] is under Apache 2.0 and can be used commercially. ◆ FLUX.2 [dev] has 32 billion parameters, and its files come with NSFW filters. ◆ Its licence says filters or manual review must be used. ◆

Qwen-Image. It's licensed under Apache 2.0, and its makers say it handles many styles, anime included. ◆ Its model card says nothing separate about

the pictures you make. ◆

Two licences to read

1

The model.

What you may do with the model file: share it, change it, sell access to it, or run it on a paid site. ◆

2

The pictures.

What you may do with what you make: post it, sell it as a custom, use it on a paid page. ◆

Running a model at home doesn't take you outside its licence. Read it for the exact version you download, because the licence can change from one version to the next. ◆ None of these licences names findom. ◆

PART 04

Disk space

Models are big files, and you'll end up keeping several. ◆

Sizes the makers publish

4.85 GB

ComfyUI's Windows app alone, per install, before any model.

23 GB

FLUX.1 [dev], original file, about 12 billion parameters.

20.4 GB

Qwen-Image, fp8 version.

40.9 GB

Qwen-Image, bf16 version.

fp8 and bf16 are two ways of storing the same model. The fp8 file is smaller and lighter on the card, and ComfyUI says its quality is slightly lower than the full version. The size of an SDXL-based model file isn't on any official page I read. Add your pictures and the small add-on files you'll collect, and a separate drive with a few hundred GB free is a comfortable start.

PART 05

What my own tool makes

I built my own image tool with this method, and one picture takes me about a minute. Here are six pictures of Vivian it made, as they came out.

Six pictures from my own tool

6 PLATES


02


03


04


05


06


07

Made by Vivian with an image tool she built herself, about a minute per picture. ◆

The tool came from the same steps as below: a plain description, many small tries, a board to compare them and a recipe sheet for every keeper. ◆

PART 06

How I work with AI

This is my method. It works with a model at home or a service online, for pictures and for words. ◆

TUTORIAL

My method, step by step

1 Say what you want in plain words

Describe the picture the way you'd tell a friend: who, where, the mood, the colours, what she's doing. Then ask an AI assistant to turn it into a prompt. ◆

Do this. Write one paragraph before you touch any setting.

> *Common mistake.* Copying a long list of keywords from someone else's picture. You won't know which word does what.

2 Start small

Use a small size and a fast setting while you look for the idea. Go big only once you like it. ◆

Do this. Keep the first rounds quick, so you can afford many of them.

3 Generate several

One picture tells you nothing. Make four to eight with the same prompt, so you see what the prompt does and what's luck. ◆

4 Compare on a board

Put the results side by side in one folder or on one page, and mark the ones that work. Look at them the next day too. ◆

Do this. Pick for the face first. Hands and background are easier to fix later.

5 Keep a recipe sheet

For every keeper, write the date, the model, the prompt, the settings and one line on why it worked. ◆

Common mistake. Trusting your memory. Three weeks later you'll never find that face again.

6 Fix one thing at a time

Change one word or one setting per round. If you change three things and it gets better, you don't know which one did it. ◆

7 Keep the originals

Save the full-size file and its recipe in a private folder. Post copies only. ◆

Common mistake. Editing your only copy. You lose the version you could have gone back to.

TRY IT

From words to a prompt

Tap every message that shows a red flag. There are two.

Example conversation. Invented for this exercise.

YOU

I want my persona on a dark red velvet sofa, looking down at the viewer, calm and a little amused. Drawn, anime style, not explicit. Turn that into a prompt.

THEM

Anime illustration, a woman with long black hair, seated on a dark red velvet sofa, looking down at the viewer, slight amused smile, dim warm light, deep red and black palette. One question before you run it: portrait or square?

Red flag 1. It kept her words and asked one question instead of guessing.

YOU

Portrait. I made six. The face is right in four, the hands are wrong in all six.

THEM

Keep everything else as it is and add one thing: “hands resting on the armrest”. A simple pose is easier to draw right. Run six again with the same settings.

Red flag 2. One change per round, so the next six show what that change did.

TRY IT

A line for the recipe sheet

Tap every message that shows a red flag. There is one.

Example conversation. Invented for this exercise.

YOU

I kept the third one. Write me a recipe line for it.

THEM

Date, model, prompt, size, settings, then: “Worked: armrest pose fixed the hands. Next time: try a softer light from the left.” Paste it next to the file.

Red flag 1. The why goes in the sheet, so the next session starts from what worked.

VIVIAN’S FIELD NOTES

I use the AI chat for words and planning, and the image tool for pictures. The chat helps me describe, compare and decide. The pictures come from rounds, never from one perfect prompt. I never paste anything real into either one, as Pick your AI explains.

◆

TRY IT

Selling customs made at home

You want to sell custom pictures of your persona made on your own computer.

A Any model is fine, since it runs on my own computer

Not this one. The licence still applies at home. Some forbid paid sites or need you to ask first.

B Read the licence of the exact version, for the model and for the pictures

The right answer. Right. Note which version you used on your recipe sheet, too.

C Open a paid site where subs make pictures with Pony Diffusion V6 XL

Not this one. Its licence forbids running it on any website or app that makes money.

REMEMBER IT

- VRAM decides what loads and how fast. You can start with a small card.
- ComfyUI runs offline, but its files can carry your prompt. Post clean exports.
- Read the licence of the exact version, for the model and the pictures.
- Plan a drive: one model can take 20 to 40 GB.
- Describe, generate several, compare, change one thing, write it down.

SOURCES AND EVIDENCE

Check it

Every claim with its evidence level and its sources.

EVIDENCE LEVELS

◆ Verified	Checked against the original source.
◆ Reported	Stated by a serious secondary source read in full.
◆ Community	Widely said in the scene; no document proves it.
◆ Vivian's field notes	First-hand experience, not a study.
◆ Unverified	Plausible, not proven yet.

01 Nothing is sent to a service, and there's no bill per image.

◆

02 The part that matters most is the graphics card, the GPU, and its own memory, called VRAM.

◆

03 More VRAM lets bigger models load in one piece, and that makes them faster.

◆

04 ComfyUI says it can run even the biggest open models on 4 GB of VRAM and 8 GB of normal memory, by moving parts of the model in and out while it works.

◆

ComfyUI README (https://raw.githubusercontent.com/comfyanonymous/ComfyUI/master/README.md)

Accessed 7 November 2026

05 Its Windows app recommends a dedicated NVIDIA or AMD card, but it doesn't require one.

◆

Install ComfyUI Desktop on Windows (https://docs.comfy.org/installation/desktop/windows)

Accessed 7 November 2026

06 Without a card, it runs on the processor, more slowly.

◆

ComfyUI system requirements (https://docs.comfy.org/installation/system_requirements)

Accessed 7 November 2026

07 Stability AI says SDXL 1.0 works on consumer cards with 8 GB of VRAM.

◆

Announcing SDXL 1.0 (https://stability.ai/news-updates/stable-diffusion-sdxl-1-announcement)

Accessed 7 November 2026

08 Share of a 24 GB card that the lighter version of Qwen-Image used in ComfyUI's own test.

◆

ComfyUI Qwen-Image tutorial (https://docs.comfy.org/tutorials/image/qwen/qwen-image)

Accessed 7 November 2026

09 Most model makers don't publish a minimum card.

◆

10 In my experience, a small card works but each picture takes longer, and slow rounds make you test less.

◆

11 You build the job as a graph of boxes joined by wires: load a model, write the prompt, generate, save.

◆

ComfyUI README (https://raw.githubusercontent.com/comfyanonymous/ComfyUI/master/README.md)

Accessed 7 November 2026

12 Its core runs fully offline and downloads nothing unless you ask.

◆

ComfyUI README (https://raw.githubusercontent.com/comfyanonymous/ComfyUI/master/README.md)

Accessed 7 November 2026

13 Some optional boxes call paid online models, and those can be switched off.

◆

ComfyUI README (https://raw.githubusercontent.com/comfyanonymous/ComfyUI/master/README.md)

Accessed 7 November 2026

14 ComfyUI can save a whole workflow and recover it, seeds included, from a picture it made.

◆

ComfyUI README (https://raw.githubusercontent.com/comfyanonymous/ComfyUI/master/README.md)

Accessed 7 November 2026

15 That's handy as a recipe, and it's also a leak: a file that carries its recipe shows your prompt and settings to anyone who opens it in the same tool.

◆

16 For anime art, most people start from a family trained on anime pictures.

◆

17 It's released under the CreativeML Open RAIL++-M licence.

◆

SD-XL 1.0-base model card (https://huggingface.co/stabilityai/stable-diffusion-xl-base-1.0/raw/main/README.md)

Accessed 7 November 2026

18 That licence claims no rights in the pictures you make, and makes you accountable for them.

◆

SDXL 1.0 base licence (CreativeML Open RAIL++-M) (https://huggingface.co/stabilityai/stable-diffusion-xl-base-1.0/raw/main/LICENSE.md)

Accessed 7 November 2026

19 It's fine-tuned on the Danbooru2023 dataset of tagged anime pictures.

◆

Illustrious XL early release v0 model card (https://huggingface.co/OnomaAIResearch/Illustrious-xl-early-release-v0)

Accessed 7 November 2026

20 Its first release, v0.1, uses a fair public AI licence that discourages using it to make money.

- **Illustrious XL early release v0 model card** (https://huggingface.co/OnomaAIResearch/Illustrious-xl-early-release-v0)

Accessed 7 November 2026

21 The makers don't restrict posting the pictures, but you can't sell a closed, modified copy of the model.

- **Illustrious XL early release v0 model card** (https://huggingface.co/OnomaAIResearch/Illustrious-xl-early-release-v0)

Accessed 7 November 2026

22 Later versions carry different licence tags: "sdxl-license" for v1.0, "creativeml-openrail-m" for v2.0.

- **Illustrious XL v1.0 model page** (https://huggingface.co/OnomaAIResearch/Illustrious-XL-v1.0)

Accessed 7 November 2026

- **Illustrious XL v2.0 model page** (https://huggingface.co/OnomaAIResearch/Illustrious-XL-v2.0)

Accessed 7 November 2026

23 Its licence is a modified Fair AI Public License.

- **Pony Diffusion V6 XL** (https://purplesmart.ai/pony/content6x)

Accessed 7 November 2026

24 It can't run on any website or app that makes money, and commercial use means writing to its makers first.

- **Pony Diffusion V6 XL** (https://purplesmart.ai/pony/content6x)

Accessed 7 November 2026

25 The FLUX.1 [dev] model is for non-commercial use only.

◆

FLUX.1 [dev] Non-Commercial License (https://raw.githubusercontent.com/black-forest-labs/flux/main/model_licenses/LICENSE-FLUX1-dev)

Accessed 7 November 2026

26 The pictures it makes can still be used for any purpose, commercial included, unless the licence forbids it elsewhere.

◆

FLUX.1 [dev] Non-Commercial License (https://raw.githubusercontent.com/black-forest-labs/flux/main/model_licenses/LICENSE-FLUX1-dev)

Accessed 7 November 2026

27 Using the model itself to earn money doesn't count as non-commercial.

◆

FLUX.1 [dev] Non-Commercial License (https://raw.githubusercontent.com/black-forest-labs/flux/main/model_licenses/LICENSE-FLUX1-dev)

Accessed 7 November 2026

28 FLUX.1 [schnell] is under Apache 2.0 and can be used commercially.

◆

FLUX.1 [schnell] model card (https://huggingface.co/black-forest-labs/FLUX.1-schnell)

Accessed 7 November 2026

29 FLUX.2 [dev] has 32 billion parameters, and its files come with NSFW filters.

◆

FLUX.2 [dev] model card (https://huggingface.co/black-forest-labs/FLUX.2-dev)

Accessed 7 November 2026

30 Its licence says filters or manual review must be used.

◆

FLUX.2 [dev] model card (https://huggingface.co/black-forest-labs/FLUX.2-dev)

Accessed 7 November 2026

31 It's licensed under Apache 2.0, and its makers say it handles many styles, anime included.

◆

Qwen-Image model card (https://huggingface.co/Qwen/Qwen-Image/raw/main/README.md)

Accessed 7 November 2026

32 Its model card says nothing separate about the pictures you make.

◆

33 What you may do with the model file: share it, change it, sell access to it, or run it on a paid site.

◆

34 What you may do with what you make: post it, sell it as a custom, use it on a paid page.

◆

35 Read it for the exact version you download, because the licence can change from one version to the next.

◆

36 None of these licences names findom.

◆

37 Models are big files, and you'll end up keeping several.

◆

38 ComfyUI's Windows app alone, per install, before any model.

◆

Install ComfyUI Desktop on Windows (https://docs.comfy.org/installation/desktop/windows)

Accessed 7 November 2026

39 FLUX.1 [dev], original file, about 12 billion parameters.

◆

ComfyUI Flux.1 text-to-image tutorial (https://docs.comfy.org/tutorials/flux/flux-1-text-to-image)

Accessed 7 November 2026

40 Qwen-Image, fp8 version.

◆

ComfyUI Qwen-Image tutorial (https://docs.comfy.org/tutorials/image/qwen/qwen-image)

Accessed 7 November 2026

41 Qwen-Image, bf16 version.

◆

ComfyUI Qwen-Image tutorial (https://docs.comfy.org/tutorials/image/qwen/qwen-image)

Accessed 7 November 2026

42 The fp8 file is smaller and lighter on the card, and ComfyUI says its quality is slightly lower than the full version.

◆

ComfyUI Flux.1 text-to-image tutorial (https://docs.comfy.org/tutorials/flux/flux-1-text-to-image)

Accessed 7 November 2026

43 The size of an SDXL-based model file isn't on any official page I read.

◆

44 Add your pictures and the small add-on files you'll collect, and a separate drive with a few hundred GB free is a comfortable start.

◆

45 I built my own image tool with this method, and one picture takes me about a minute.

◆

46 Here are six pictures of Vivian it made, as they came out.

◆

47 Made by Vivian with an image tool she built herself, about a minute per picture.

◆

48 The tool came from the same steps as below: a plain description, many small tries, a board to compare them and a recipe sheet for every keeper.

◆

49 It works with a model at home or a service online, for pictures and for words.

◆

50 Then ask an AI assistant to turn it into a prompt.

◆

51 Go big only once you like it.

◆

52 Make four to eight with the same prompt, so you see what the prompt does and what’s luck.

◆

53 Look at them the next day too.

◆

54 For every keeper, write the date, the model, the prompt, the settings and one line on why it worked.

◆

55 If you change three things and it gets better, you don’t know which one did it.

◆

56 Post copies only.

◆

57 I use the AI chat for words and planning, and the image tool for pictures. The chat helps me describe, compare and decide. The pictures come from rounds, never from one perfect prompt. I never paste anything real into either one, as Pick your AI explains.

◆