

How to Install and Use Comfy Desktop

Running ComfyUI locally on your own hardware — from download to your first generated image
Windows walkthrough · Comfy Desktop 1.0.38

This guide walks through the complete process of installing Comfy Desktop on Windows, creating a ComfyUI instance, loading a text-to-image workflow, downloading the required models, and generating a test image. Every step is shown with a screenshot; the item you need to click is circled in red.

Before you begin


- A Windows PC with a dedicated GPU. This walkthrough uses an NVIDIA card, which Comfy detects automatically.
- Roughly 75 GB of free disk space: about 500 MB for the application, ~5 GB for the ComfyUI instance, and ~69 GB for the Flux.2 Dev model files used in this example.
- A stable internet connection. The model download is the longest part of the process by far.
- Time. Budget an hour or more depending on your connection speed.

Note: You do not need a Comfy account, a subscription, or credits to follow this guide. Everything here runs locally and free on your own hardware.

Part 1 · Download and install Comfy Desktop

Step 1 — Download the installer

Go to the Comfy website and click **DOWNLOAD DESKTOP**. This is the free, run-on-your-own-hardware version — it includes the full ComfyUI engine.




Click **DOWNLOAD DESKTOP** on the Comfy Desktop landing page.

Tip: Ignore the "Install from GitHub" option unless you specifically want to build from source. The desktop installer handles all dependencies for you.

Step 2 — Run the installer

Find the downloaded file in your Downloads folder — it will be named something like Comfy-Desktop-Setup-... .exe and is roughly 155 MB. Right-click it and choose **Open** (or just double-click it).




Right-click the downloaded installer and choose **Open**.

Note: You do not need to run this as administrator if you are installing for yourself only, which is the recommended option in the next step.

Step 3 — Choose who to install for


The setup window asks whether to install for everyone or just you. Leave **Only for me** selected and click **Next**.



Keep "Only for me" selected, then click Next.

Step 4 — Confirm the install location

Setup suggests installing to your user AppData folder. The application itself needs about 502 MB. Unless you have a specific reason to change this, leave the default and click **Install**.




Accept the default destination folder and click Install.

Important: This location is only for the application. The much larger model files are stored separately, in the instance folder you choose in Step 8.

Step 5 — Wait while files extract


The installer works through three stages: the Microsoft Visual C++ Redistributable, the Comfy Desktop application files, and final setup. If you already have an up-to-date Visual C++ Redistributable, that step is skipped automatically. No action is needed here.



The installer extracting application files. This takes a minute or two.

Step 6 — Finish setup

When setup completes, leave **Run Comfy Desktop** and **Add a desktop shortcut** checked, then click **Finish**. Comfy Desktop will launch on its own.




Click Finish to close setup and launch Comfy Desktop.

Part 2 · Create your first ComfyUI instance

Comfy Desktop is a launcher that can manage several independent ComfyUI environments, called **instances**. Each instance has its own Python environment and model folders, so you can keep projects separate. You need to create one before you can generate anything.

Step 7 — Create a new instance


On the launcher home screen, click **New Instance** to set up a fresh ComfyUI environment. (The Comfy Cloud tile next to it runs on Comfy's hardware and requires credits — skip it for a local setup.)



Click New Instance to build a local ComfyUI environment.

Step 8 — Name the instance and confirm settings

Give the instance a name, or leave the default ComfyUI. Check that the **Detected GPU** field shows your graphics card — if it does, Comfy will install the correct GPU-accelerated packages. Confirm the install location has enough free space (about 4.83 GB is needed at this stage), then click **Continue**.

A dark-themed installation window for ComfyUI. It features three main sections: 'Name' with a text box containing 'ComfyUI'; 'Detected GPU' with a text box showing 'Detected GPU: NVIDIA'; and 'Install Location' with a text box showing a path 'C:\Users\...AppData\Local\Comfy-Desktop\ComfyUI-Inst...' and a 'Browse...' button. Below the path, it displays '643.15 GB free · ~4.83 GB required'. At the bottom left is a '> Advanced' link, and at the bottom right is a yellow 'Continue' button circled in red.

Name

ComfyUI

Detected GPU

Detected GPU: NVIDIA

Install Location

C:\Users\...AppData\Local\Comfy-Desktop\ComfyUI-Inst... Browse...

643.15 GB free · ~4.83 GB required

> Advanced


Continue

Confirm the name, detected GPU, and install location, then click Continue.

Tip: If you plan to download many large models, point the install location at a drive with plenty of room. Models live inside this folder, and they add up fast — the single workflow in this guide pulls down nearly 70 GB.

Step 9 — Skip the starter workflow

Comfy offers to install a starter workflow along with its models. Click **Skip & Install** instead. It is cleaner to finish setting up first and then pick your workflow deliberately, which is what Part 3 does.



Click Skip & Install rather than choosing a starter workflow here.


Why skip? Some of the offered workflows require cloud credits, and the recommended one downloads about 53 GB of video models you may not want. Skipping keeps the initial install small and fast.

Part 3 · Load a workflow and download its models

ComfyUI now opens with an empty canvas. Rather than wiring up nodes by hand, start from a template — a pre-built workflow that already has the right nodes connected in the right order.

Step 10 — Open the template browser and filter to Image

Open the **Templates** panel. In the left sidebar under **GENERATION TYPE**, click **Image** to narrow the list to image-generation workflows.




In the Templates panel, select Image under Generation Type.

Note: If the panel reads "No templates found" before you filter, that is normal — choosing a category or typing in the search box populates the list.

Step 11 — Choose the Flux.2 Dev text-to-image template


Type **Flux.2 Dev** into the search box, then click **Flux.2 Dev Text to Image**. This is a straightforward prompt-in, image-out workflow — a good first test.



Search for Flux.2 Dev and select the "Flux.2 Dev Text to Image" template.

Step 12 — Expect errors on first load


The workflow opens and immediately reports **4 errors found**. This is expected and not a problem: the workflow references model files that are not on your machine yet. Click **View details** to see what is missing.



Four errors appear because the required models are not downloaded yet. Click View details.

Step 13 — Download all the missing models

The **Workflow Overview** panel lists each missing model under **Missing Models** — in this case a LoRA, a VAE, a text encoder, and the main diffusion model. You can fetch them one at a time, but the simplest route is to click **Download all**. For this workflow that is about 69 GB.




Click "Download all (68.97 GB)" to fetch every missing model at once.

Heads up: This is the long part. Expect a substantial wait on most connections, and make sure the destination drive has the space free before you start.

Step 14 — Track the download progress


A loading bar appears under the app icon in the menu bar. It often moves quickly at first and then slows down — this is normal. Simply wait for the bar to disappear.

To see exactly how many files are left, click the **downloads** button in the top-right corner of the window.



The downloads button in the top-right corner. The badge shows how many files remain.

The Downloads panel lists each file with its progress. Completed items show a checkmark and a "Show in Explorer" link; the one still in flight shows a percentage. Notice that the error count in the Workflow Overview drops as each model lands.



The Downloads panel: three models complete, one still transferring at 60%.

Tip: You can leave this running and come back later. Once the error count reaches zero, the workflow is ready to run.

Part 4 · Generate your first image

Step 15 — Enter a test prompt


Use the scroll wheel to zoom in on the middle node — **Text to Image (Flux.2 Dev)** — and replace the existing text in the prompt box with the test prompt below. It is deliberately detailed, which makes it easy to judge whether the model is working properly.

Interior photograph of a small concrete chapel on a rocky coastline. Board-formed concrete walls with visible tie holes, a single clerestory slot washing the east wall, oak pews, worn brick floor laid in stretcher bond. Late-afternoon winter light raking low through the slot; overcast sky as soft fill. Shot on a 24mm tilt-shift lens, verticals corrected, tripod height 1.5m, Kodak Portra 400. Muted palette of grey-green and warm oak. Empty, quiet, slightly cold. No people, no furniture beyond the pews.

Leave the other parameters at their defaults — width and height at 1024, and the model fields already filled in by the template.

Step 16 — Press Run


Click **Run** in the top right. The prompt is queued and generation begins.



With the prompt in place, click **Run** in the top right.

Step 17 — Watch the progress


Click the **1 active** button to follow progress. The active node is outlined in green and reports its current stage and percentage. The first run of a session is slower than later ones, because the model has to be loaded into memory.



Click "1 active" to see live progress. The running node is highlighted in green.

Step 18 — Review the result

When generation finishes, the image appears in the **Save Image** node at 1024 × 1024. Results vary from run to run — each generation uses a different random seed — but you should get something in this neighbourhood:



A finished 1024 × 1024 generation from the test prompt.

Generated images are also written to disk automatically, using the `filename_prefix` shown on the Save Image node.

Troubleshooting and next steps

The workflow still shows errors

- Open the Workflow Overview panel and check the Missing Models list. Use the refresh icon next to it to re-scan after downloads finish.
- If a download failed, click Download again on just that file rather than restarting the whole set.
- If a model was downloaded manually, it must sit in the correct subfolder of the instance's models directory — the Model Links note in the workflow shows the expected layout.

Generation is very slow or runs out of memory

- Confirm the Detected GPU field found your card during instance setup. If it fell back to CPU, generation will be extremely slow.

- Lower the width and height on the Text to Image node (try 768×768) to reduce memory use.
- Close other GPU-heavy applications before running.

Where to go from here

- Try your own prompts in the same node — no other changes needed.
- Change the noise seed to get a different variation of the same prompt.
- Browse the other template categories (Video, Audio, 3D Model) to see what your hardware can handle.
- Create additional instances from the launcher home screen to keep separate projects and model sets isolated from each other.