



PRESENTS

ComfyUI From Zero

Understand the machine before you copy a single workflow.

🎁 Part 1 • Free

10 modules

about 2 hours



ComfyUI From Zero

PART 1 • FREE

M00

Intro



m00-01

04

• MODULE 1 • 12 MIN

What ComfyUI actually is

Before we touch a single node: what the machine is, what it isn't, and why it looks like spaghetti.

-  The model vs the interface, in one sentence
-  What a workflow is and why it runs left to right
-  The map of the whole course


 VRAM • none

 Free



- BY THE END OF PART 1

You'll read this at a glance.



- THE ONE-SENTENCE VERSION

The **model** is the photographer. **ComfyUI** is the studio.



THE MODEL

The photographer

A big file of numbers that learned what pictures look like. It does the actual drawing.

Z-Image

Flux

Qwen


...hundreds more

Different photographers, different styles, different day rates (paid in VRAM).

THE INTERFACE

The studio


Where you brief the photographer and collect the photo. ComfyUI is one studio among several.




Swap the photographer. **Keep the studio.**

- WHY IT LOOKS LIKE THIS

Most apps hide the machine. ComfyUI hands you the **wiring**.



😊 Easy on day one 🗄 Stuck on day ten



👻 Scary on day one 🗄 Free on day ten


This course is for the people who want to see the **screws**.



• IT'S NOT AN IMAGE APP

One studio. **Very different photographers.**

YOUR RESULT HERE



Images

Words in, picture out.

YOUR RESULT HERE



Edits

Change a photo with a sentence.


YOUR RESULT HERE



Video

A picture that learned to move.

YOUR RESULT HERE



Audio

Voices, music and sound effects.


YOUR RESULT HERE



3D

One image in, a model out.

YOUR RESULT HERE




Prompts & LLMs

Let a language model write the brief.



• WORKFLOW

A workflow is a set of connected boxes that runs **left to right**.




DATA FLOWS THIS WAY, LIKE READING



What each box does? That's [Modules 5 and 6](#). Today, just notice the shape.

Four ideas carry the whole course.



IDEA 1

Noise

Every image starts as TV static.


"a cat in a tiny hat"



IDEA 2

Conditioning

Your prompt steers what the static becomes.



IDEA 3

Sampling

The static gets cleaned up one step at a time.



IDEA 4

Structure vs style

Where things go vs how they look. Two separate levers.

• HOW THIS COURSE WORKS

Part 1 is **free**. Parts 2 and 3 are one paid bundle.

PART 1 · FREE

YOU ARE HERE

Understand the machine

- What ComfyUI is
- Your hardware tier
- Install and first launch
- Your first image
- The words
- Build a workflow from scratch
- Reading model files
- Saving and updating safely
- The eight errors

PART 2 · PAID



Open the hood

- Using LoRAs properly
- Edit, inpaint and outpaint
- Video models
- Low-VRAM tiers tested on real cards
- Prompting formulas
- Cloud GPUs, step by step

PART 3 · PAID



Build the engine

- Train your own LoRA locally
- Write a custom node
- Automate with the API
- Merge your own models
- Licences for commercial work



● MODULE 1 COMPLETE

You can now...

- ✓ Say what a **model** is and what the **studio** is
- ✓ Explain why ComfyUI shows you **every cable**
- ✓ Recognise a **workflow** and which way it flows
- ✓ Name the **four ideas** that carry the course

NEXT UP

02

Can my computer run this?

Find your tier before you download 20 GB of anything.



02

• MODULE 2 • 10 MIN

Can my computer run this?

Find your tier before you download 20 GB of anything.

-  VRAM vs RAM, finally explained
-  Which tier you're in and what it can do
-  Your path: Windows, Mac, AMD or cloud

 VRAM • none

 Free



• VRAM VS RAM

VRAM is the **desk**. RAM is the filing cabinet in the **hallway**.



ON THE GRAPHICS CARD

VRAM

8 GB desk

⚠ desk full


2 GB

2 GB

2 GB

2 GB

Fast, small, expensive. Where the photographer actually works.



ON THE MOTHERBOARD

System RAM

2 GB

2 GB

Big, cheap, and a walk away.


Too big for the desk? It spills into the hallway. **Slower, not broken.**



● 30 SECONDS

Find your number before anything else.

- 1 Press **Ctrl** **Shift** **Esc** to open Task Manager
 - 2 Click **Performance**, then **GPU**
 - 3 Read **Dedicated GPU memory**. That's your VRAM.
- 🍏 On a Mac: **Apple menu** → **About This Mac**.
Memory is shared, so the GPU gets most of it, not all.




👻 **Shared GPU memory** is RAM wearing a costume. Ignore it.



● FIVE TIERS

Everyone fits in one of these.

6–8 GB



TYPICAL

RTX 3050, 2060, 4060
and most laptops


IN PART 1

Everything, using FP8 or
GGUF files

LATER

Images yes. Short video at
low resolution, slowly.

12–16 GB



TYPICAL

RTX 3060 12 GB, 4070,
5060 Ti, 5070 Ti


IN PART 1

Everything, comfortably

LATER

Most video workflows at
720p

24 GB+



TYPICAL

RTX 3090, 4090, 5090


IN PART 1

Everything, full-size files if
you like

LATER

Everything in this course

Mac • AMD



TYPICAL

Apple Silicon, Radeon
cards


IN PART 1

Images via the Desktop
app. Some FP8 files won't
load.

LATER

Video is limited. Cloud
helps.

No GPU



TYPICAL

Any laptop

IN PART 1

Rent a cloud GPU by the
hour


LATER

Same, just a bigger rented
card



● PICK YOUR PATH

Four questions. One answer.




● RULE OF THUMB

Model file size + 2 GB $\approx$ the VRAM you need.

Full size • FP16


example model



✓ 16 GB card

Half size • FP8


same model



✓ 8 GB card

GGUF • Q8


near-perfect copy



✓ 12 GB card

GGUF • Q4

rougher, but tiny



✓ 6 GB card

6 GB

8 GB

12 GB

16 GB

24 GB

Doesn't fit? Step down one size until it does.



• NO GPU?

Borrow one by the hour.



Same ComfyUI

The studio runs on a rented computer. You use it in your browser, exactly like on this screen.



Pay while it's on

It bills for every minute it's running, even while you're only moving boxes around. **Switch it off.**



Full walkthrough later

Part 2 covers templates, storage and costs step by step. For Part 1, any ComfyUI cloud service works.

You can now...

- ✓ Tell **VRAM** from **RAM** and why overflow is slow, not broken
- ✓ Find your **real VRAM** number in 30 seconds
- ✓ Name your **tier** and your **install path**
- ✓ Estimate if a model fits: **file size + 2 GB**

NEXT UP

03

Install and first launch

Get ComfyUI running, and learn which window you must never close.



03

• MODULE 3 • 12 MIN

Install and first launch

Get ComfyUI running, and learn which window you must never close.

-  The two official ways to install
-  The three install failures and their fixes
-  Engine room vs dashboard

 VRAM • any

 Free



• TWO OFFICIAL WAYS


Pick your studio's front door.



EASIEST START

ComfyUI Desktop

- ✓ An installer, like any other app
- ✓ Windows and Apple Silicon Macs
- ✓ Updates from inside the app
- ✓ The calmest first day



MOST FLEXIBLE

Portable

- ✓ A folder you unzip and run
- ✓ Windows with an Nvidia card
- ✓ Move it, copy it, keep two side by side
- ✓ Great once you're comfortable

📺 On screen in this course: ComfyUI Desktop

Every node works exactly the same in both.






Install it with me.


Pause the video whenever you need to. I'll wait.

- 1 Download from **comfy.org**. Only from there.
- 2 Pick a drive with **30 GB+ free**, ideally an SSD
- 3 Choose **NVIDIA**, or Apple Silicon on a Mac
- 4 Let it download everything. Go make a coffee.



- WHY IS IT SO BIG?

It's downloading the photographer's **equipment truck**.



Python · the language it's written in

PyTorch · the maths engine

CUDA · talks to your Nvidia card



No models inside yet. The photographer arrives in **Module 4**.



• WHEN IT GOES WRONG

The three install failures, and their fixes.

FAILURE 1

The nervous bouncer

⚠ What you see

Windows protected your PC

🔑 The fix

Only for the installer you got from **comfy.org**: click **More info**, then **Run anyway**.

FAILURE 2

The overprotective antivirus

⚠ What you see

A ComfyUI file was quarantined

🔑 The fix

Restore the file, then add your **ComfyUI folder** to the antivirus exclusions.

FAILURE 3

The path from hell

⚠ What you see


C:\Users\You\~~OneDrive\New folder~~
(2)\ComfyUI

🔑 The fix


Use a short path with no spaces, like **C:\AI\ComfyUI**. Never OneDrive, never Program Files.

- TWO WINDOWS

ComfyUI is an **engine** and a **dashboard**.



The engine room · does the work




The dashboard · where you work

 **Close the engine room and the dashboard goes dark.**

• 60-SECOND TOUR

Four places. That's the whole interface.



- 1 Canvas.** Where the boxes live. Drag to move around, scroll to zoom.
- 2 Run.** Sends your workflow to the engine room.
- 3 Sidebar.** Nodes, models, workflows and the queue.
- 4 Bottom panel.** What the engine is saying.
- ×** **Templates popup.** Close it. For now.

• NOT ON WINDOWS + NVIDIA?

Your path, in sixty seconds each.



Mac

- ✓ ComfyUI Desktop, Apple Silicon (M1 or newer)
- ✓ It uses the Mac's GPU automatically
- ✓ Some FP8 model files won't load. Grab the FP16 or GGUF version.

📖 Written guide in the description



AMD

- ✓ Windows: a ZLUDA build of ComfyUI
- ✓ Linux: ComfyUI with ROCm
- ✓ Expect extra setup. Follow the guide step by step.

📖 Written guide in the description



Cloud

- ✓ Pick a ComfyUI template at your GPU provider
- ✓ Open the web address it gives you
- ✓ From Module 4 on, everything is identical

📖 Written guide in the description



● MODULE 3 COMPLETE

You can now...

- ✓ Install ComfyUI the **official** way
- ✓ Fix the **three** most common install failures
- ✓ Tell the **engine room** from the **dashboard**
- ✓ Find the **canvas, Run, sidebar** and **bottom panel**

NEXT UP

04

Your first image

An image before any theory. Then we break it on purpose.



04

• MODULE 4 • 10 MIN

Your first image

An image before any theory. Then we break it on purpose.

-  Run a real workflow in five minutes
-  Read your first red node like a pro
-  The five folders you actually touch

 VRAM • 6 GB

 Free






Load workflow 01.


It's in the description, free, with every download link inside.

- 1 Download **workflow-01** from the description
- 2 Drag the **.json** file onto the canvas
- 3 Read the yellow **note**. It lists every file you need.



● PRESS RUN

It fails. On purpose.



⚠ Prompt outputs failed validation

CheckpointLoaderSimple:
Value not in list: ckpt_name:
'z_image_turbo_aio.safetensors' not in []



A warning light, not a crash.

The car is telling you a bulb is missing.

Nothing is broken. **A file is missing.** Workflows never include the models.

• THE FIX

Download. Drop it in the folder. Press **R**.

1




Download

Use the link in the note node.

`z_image_turbo_aio.safetensors`

2



Drop it in the folder

ComfyUI/

models/

checkpoints/

`.safetensors`

3



Press R

R

Refreshes the lists, so the file shows up in the dropdown.

Load Checkpoint

MODEL

ckpt_name `z_image_turbo_aio.safetensors`


▶ Run again

The dropdown finally has something in it.



● RUN

Green nodes. Your first image.



output/ComfyUI_00001_.png



Now change the prompt to something silly and run it again.

• STILL RED?

Didn't show up? Check these, in this order.



CHECK 1

Press R

Refresh the lists. Most of the time, that's it.



STILL
MISSING?




CHECK 2

Restart ComfyUI

Close the engine room and the dashboard, then start again.



STILL
MISSING?



CHECK 3

Right folder?

`checkpoints/` for all-in-one files.
`diffusion_models/` for bare models.



STILL
MISSING?



CHECK 4

Finished?

It must end in `.safetensors`, not `.crdownload` or `.part`.



Nine times out of ten it's **check 1** or **check 3**.

• WHERE THINGS LIVE

The five folders you'll actually touch.

ComfyUI/

models/

every model file you download

checkpoints/

all-in-one models

diffusion_models/

bare models

text_encoders/ · vae/ · loras/

the helpers (Module 5)

input/

pictures you load in

output/

everything you generate

custom_nodes/

add-ons from GitHub (Module 8)

user/default/workflows/

your saved workflows



● MODULE 4 COMPLETE

You can now...

- ✓ Load a workflow by **dragging a .json** onto the canvas
- ✓ Read a **red node** and fix a missing model
- ✓ Make an image and **find it** in the output folder
- ✓ Know the **five folders** that matter

NEXT UP

05

The words

41 words in 5 groups. You don't have to memorise a single one.



05

• MODULE 5 • 15 MIN

The words

41 words in 5 groups. You don't have to memorise a single one.

- 📖 Files, nodes and settings in plain English
- 📷 The photographer version of every term
- 📄 A cheat sheet to keep

📦 VRAM • none

📺 Free



• HOW THIS MODULE WORKS

41 words. Five groups. Zero memorising.

A

The files you download

10 WORDS

model · diffusion model · text encoder · VAE · latent · LoRA ...

B

The nodes in a basic workflow

12 WORDS

node · link · workflow · Load Checkpoint · the three loaders · CLIP Text Encode ...

C

Inside the KSampler

6 WORDS

seed · steps · CFG · sampler · scheduler · denoise

D

Canvas words

8 WORDS

Group · Bypass · Ctrl+B · Mute · Ctrl+M · Subgraph · Template · Custom node ...

E

Hardware words

5 WORDS

VRAM · RAM · Offloading · Attention · sage · flash · pytorch · CUDA

📄 What it is



📷 The photographer version



📌 Remember

on every slide



There's a two-page **cheat sheet** in the description. Grab it now.





● GROUP A • 10 WORDS • ABOUT 5 MINUTES

The files you download

model

diffusion model

text encoder

VAE

latent

LoRA

ControlNet

upscaler

FP16 / FP8 / GGUF

safetensors

Model

also called: checkpoint, ckpt, base model

WHAT IT IS

A big file of numbers that learned what images look like. It's the part that actually draws.



THE PHOTOGRAPHER VERSION

The photographer you hire. A checkpoint means she brings her own translator and darkroom.



REMEMBER

No model, no image. Different models are different photographers with different day rates, paid in VRAM.

↓ loaded here

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Checkpoint

MODEL

CLIP

VAE

ckpt_name `z_image_turbo_aio.safetensors`

models/checkpoints

Diffusion model

also called: UNet, DiT, the bare model

WHAT IT IS

The same kind of model, shipped without its helpers. Most new models arrive this way, so you load three files instead of one.



THE PHOTOGRAPHER VERSION

She shows up with just her camera. You hire the translator and book the darkroom yourself.



REMEMBER

Goes in diffusion_models, not checkpoints. The wrong folder is the number one cause of 'value not in list'.

↓ Load Checkpoint gets swapped for three loaders

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Diffusion Model

MODEL

unet_name z_image_turbo_fp8.safetensors

weight_dtype default

models/diffusion_models

Text encoder

shown as CLIP in ComfyUI • the file can be CLIP, T5, Qwen or Gemma

WHAT IT IS

Turns your words into numbers the model understands.
ComfyUI still calls the slot CLIP, whatever file is inside.



THE PHOTOGRAPHER VERSION

The translator. You speak human, the photographer speaks maths.



REMEMBER

Built into checkpoints. For diffusion models, load one separately and set Load CLIP's type to match the model family.

↓ loaded here, used by both prompts

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load CLIP

CLIP

clip_name qwen_3_4b.safetensors

type lumina2

models/text_encoders

VAE

Variational AutoEncoder

WHAT IT IS

Converts between the small compressed image the model works on and real pixels. Decode goes to pixels, encode comes back.



THE PHOTOGRAPHER VERSION

The darkroom. She shoots on a negative; the darkroom develops it into a photo.



REMEMBER

Built into checkpoints, separate for diffusion models. A VAE from the wrong family gives odd colours or a shape error.

↓ loaded here, used here

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load VAE

VAE

vae_name

ae.safetensors

VAE Decode

samples

IMAGE

vae

models/vae

Latent

latent image · latent space · 'samples'

WHAT IT IS

The compressed working version of an image, about eight times smaller on each side. All the generating happens in here.



THE PHOTOGRAPHER VERSION

The negative. You can't really look at it, but everything happens on it.



REMEMBER

Empty Latent Image is a blank negative of the size you pick. Keep sizes in multiples of 8.

↓ the latent travels through these three

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

pixels

VAE Encode →

← VAE Decode

128 × 128

1024 × 1024

8× smaller each side · 64× less to compute



LoRA

Low-Rank Adaptation

WHAT IT IS

A small add-on file that nudges a model toward a style, a face or a skill. It rides on top of the model and must match its family.



THE PHOTOGRAPHER VERSION

The photographer took a weekend class. Same camera, new trick.



REMEMBER

Put its trigger word in the prompt. Start around strength 0.6 to 1.0. Using LoRAs is Part 2, training them is Part 3.

↓ **plugs in between these two**

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

● LoraLoaderModelOnly

● model MODEL ●
lora_name **watercolor_style.safetensors**
strength_model **0.80**

📁 **models/loras**



ControlNet

Canny · Depth · OpenPose · Union

WHAT IT IS

A helper network that steers the picture with another picture: a pose skeleton, an edge drawing or a depth map.



THE PHOTOGRAPHER VERSION

You hand her a reference photo: 'stand exactly like this.'



REMEMBER

The prompt says WHAT. ControlNet says WHERE.

↓ sits between the prompts and the sampler

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Apply ControlNet

- positive positive
 - negative negative
 - control_net
 - image
- strength 0.80
- start_percent 0.000
- end_percent 1.000

models/controlnet

Upscaler model

ESRGAN-style models • 4x upscalers

WHAT IT IS

A small model that makes an image bigger and invents believable detail, instead of just stretching the pixels.



THE PHOTOGRAPHER VERSION

The print lab that blows your photo up to poster size without it going mushy.



REMEMBER

Not the same as Upscale Latent, which only resizes the negative before the shoot.

↓ added after decoding, before saving

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Upscale Model

UPSCALE_MODEL

model_name

4x_upscaler.pth

Upscale Image (using Model)

upscale_model

IMAGE

image

models/upscale_models

FP16 • FP8 • GGUF

precision and quantisation • Q8 • Q6 • Q4

WHAT IT IS

The same model saved with more or fewer decimal places. Fewer decimals means a smaller file and less VRAM, slightly rougher results.



THE PHOTOGRAPHER VERSION

The same photo saved at 4K, 1080p and 720p. She isn't dumber in the small file, just blurrier.



REMEMBER

File size plus about 2 GB is roughly the VRAM you need. Step down until it fits. GGUF files need GGUF loader nodes.

↓ **decides which loader node you use**

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Same model • relative file size

FP16 / BF16 100%

FP8 ~50%

GGUF Q8 ~53%

GGUF Q6 ~41%

GGUF Q4 ~30%

Q8 is hard to tell apart from FP16. Q4 is visibly rougher.

safetensors

.safetensors versus .ckpt · .pt · .pth

WHAT IT IS

The modern file format for model weights. It holds numbers only, so it can't hide code that runs when you load it.



THE PHOTOGRAPHER VERSION

A sealed envelope of photos, instead of a parcel that might contain a surprise.



REMEMBER

Prefer .safetensors whenever you have the choice. It's safer, not faster.

↓ applies to anything a loader opens

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI



.ckpt · .pt · .pth

can contain Python that runs the moment you load it



.safetensors

numbers only. Nothing runs.

Same weights inside. Different level of trust.

● GROUP B • 12 WORDS • ABOUT 5 MINUTES

The nodes in a basic workflow

node

link

workflow

Load Checkpoint

the three loaders

CLIP Text Encode

conditioning

Empty Latent Image

KSampler

VAE Decode / Encode

Save / Preview Image

Load Image

Node

a box

WHAT IT IS

A box that does one job. Inputs come in on the left, outputs leave on the right, and its settings sit in the middle.



THE PHOTOGRAPHER VERSION

One crew member with one task.



REMEMBER

Double-click an empty spot on the canvas to search for any node by name.

↓ every box here is a node

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

↓ title: what it does

Upscale Image By

image	IMAGE
upscale_method	lanczos
scale_by	1.50

↑ inputs
come in left

outputs ↑
leave right

↑ widgets = its settings

Link

also called: wire · cable · noodle

WHAT IT IS

A wire that carries data from an output to an input. Its colour is the data type, and ComfyUI won't let you plug it into the wrong socket.



THE PHOTOGRAPHER VERSION

The cables on set. The yellow one only fits a yellow socket.



REMEMBER

One output can feed many inputs, like a power strip. One input takes exactly one cable, like a wall socket.

↓ the lines between the boxes

Load Checkpoint

Prompt +

Prompt -


Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI



Workflow

graph · .json file

WHAT IT IS

All the nodes and links together. It saves as a small JSON text file, and ComfyUI hides a copy inside every PNG it saves.



THE PHOTOGRAPHER VERSION

The shot list for the whole day.



REMEMBER

Workflows never include the models. Drag a ComfyUI PNG onto the canvas to get its workflow back.

↓ the whole thing

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

workflow-01.json

```
{
  "nodes": [
    { "id": 4, "type": "CheckpointLoaderSimple",
      "widgets_values": ["z_image.safetensors"] },
    { "id": 3, "type": "KSampler", ... }
  ],
  "links": [
    [1, 4, 0, 3, 0, "MODEL"], ...
  ]
}
```

📄 saved as .json

🖼️ also hidden in every PNG

Load Checkpoint

CheckpointLoaderSimple

WHAT IT IS

Loads a checkpoint file and splits it into three outputs, MODEL, CLIP and VAE, each wired to the node that needs it.



THE PHOTOGRAPHER VERSION

Hire the photographer. She arrives with her translator and her darkroom.



REMEMBER

Purple goes to the KSampler, yellow to both prompts, red to VAE Decode.

↓ **always first**

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Checkpoint

MODEL 
CLIP 
VAE 

ckpt_name `z_image_turbo_aio.safetensors`

The three loaders

Load Diffusion Model • Load CLIP • Load VAE

WHAT IT IS

Three nodes that do Load Checkpoint's job for models shipped in pieces. Same three outputs, from three files.



THE PHOTOGRAPHER VERSION

Hire the photographer, the translator and the darkroom separately.



REMEMBER

Load CLIP's type must match the model family. GGUF files use Unet Loader (GGUF) and CLIP Loader (GGUF) instead.

↓ these three replace this one

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Diffusion Model

MODEL

unet_name `z_image_turbo_fp8.safetens...`

weight_dtype `default`

Load CLIP

CLIP

clip_name `qwen_3_4b.safetensors`

type `lumina2`

Load VAE

VAE

vae_name `ae.safetensors`

CLIP Text Encode

the prompt boxes

WHAT IT IS

Turns a prompt into conditioning. You usually have two: one for what you want and one for what you don't.



THE PHOTOGRAPHER VERSION

The brief in one hand, the 'please, no' list in the other.



REMEMBER

The two nodes are identical. The KSampler socket you plug into decides which one is positive.

↓ one positive, one negative

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

CLIP Text Encode (Positive)

clip

CONDITIONING

text

a cat in a tiny hat, golden hour

CLIP Text Encode (Negative)

clip

CONDITIONING

text

blurry, text, watermark

Conditioning

the orange cable

WHAT IT IS

The encoded prompt travelling along the orange cable. The KSampler steers toward the positive and away from the negative.



THE PHOTOGRAPHER VERSION

The translated brief the photographer actually reads.



REMEMBER

Turbo models ignore the negative, so their workflows plug in Conditioning Zero Out as a polite empty brief.

↓ made here, read here

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

● Conditioning Zero Out

● conditioning

CONDITIONING ●

Empty Latent Image

newer models have their own version, e.g. Empty SD3 Latent Image

WHAT IT IS

Creates a blank latent with the width, height and batch size you pick. Every text-to-image run starts here.



THE PHOTOGRAPHER VERSION

Choosing the blank photo paper. Bigger paper takes longer to shoot.



REMEMBER

To start from an existing picture, swap it for Load Image followed by VAE Encode.

↓ the starting point

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Empty Latent Image

LATENT

width	1024
height	1024
batch_size	1

KSampler

the sampler node

WHAT IT IS

The node that actually generates. It takes the model, both prompts and a latent, removes noise step by step, and hands back a finished latent.



THE PHOTOGRAPHER VERSION

The shoot itself.



REMEMBER

Every setting people argue about lives in here. Section C walks through them one at a time.

↓ where the work happens

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

- model LATENT
- positive
- negative
- latent_image
- seed 42
- control after generate fixed
- steps 8
- cfg 1.0
- sampler_name euler
- scheduler simple
- denoise 1.00

VAE Decode • Encode

latent to pixels • pixels to latent

WHAT IT IS

VAE Decode turns the finished latent into pixels near the end of every workflow. VAE Encode does the reverse.



THE PHOTOGRAPHER VERSION

The darkroom, working in both directions.



REMEMBER

Both need the red VAE cable. A 'missing VAE' error means that wire isn't connected.

↓ **decode lives here**

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

VAE Decode

samples

IMAGE

vae

VAE Encode

pixels

LATENT

vae

Save • Preview Image

the end of every branch

WHAT IT IS

Save Image writes the result to the output folder. Preview Image shows it on the canvas and keeps it in temp, which gets cleared.



THE PHOTOGRAPHER VERSION

Deliver the photo to the client, or just show her the Polaroid.



REMEMBER

A workflow with neither gives 'prompt has no output'.
Saved PNGs carry the workflow inside.

↓ **always last**

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Save Image

images

filename_prefix

ComfyUI

Preview Image

images

output/ • temp/

Load Image

your own picture, in

WHAT IT IS

Brings a picture from the input folder onto the canvas. It outputs the image and a mask; right-click it for the mask editor.



THE PHOTOGRAPHER VERSION

The photo you bring to the shoot as a starting point.



REMEMBER

Load Image into VAE Encode replaces Empty Latent Image when you want image-to-image.

↓ replaces this for image-to-image

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

Load Image

IMAGE

MASK

image

cat.png

VAE Encode

pixels

vae

LATENT

input/



● GROUP C • 6 WORDS • ABOUT 3 MINUTES

Inside the KSampler

seed

steps

CFG

sampler

scheduler

denoise



Seed

seed + control after generate

WHAT IT IS

The number that decides the random starting noise. Same seed and same settings give the same image, every time.



THE PHOTOGRAPHER VERSION

Which roll of the dice the shoot starts from.



REMEMBER

Control after generate can be fixed, increment or randomize. Fix the seed whenever you compare settings.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

- model LATENT
- positive
- negative
- latent_image
- seed 42
- control after generate fixed
- steps 8
- cfg 1.0
- sampler_name euler
- scheduler simple
- denoise 1.00

Steps

sampling steps

WHAT IT IS

How many passes the sampler makes to clean up the noise. More steps look cleaner and take longer, until they stop helping.



THE PHOTOGRAPHER VERSION

How many times you wipe the foggy window.



REMEMBER

Classic models like 20 to 30. Turbo models like 4 to 8. Use the number on the model page.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

model	LATENT
positive	
negative	
latent_image	
seed	42
control after generate	fixed
steps	8
cfg	1.0
sampler_name	euler
scheduler	simple
denoise	1.00

CFG

classifier-free guidance

WHAT IT IS

How hard the sampler is pushed toward the positive prompt and away from the negative one.



THE PHOTOGRAPHER VERSION

The tension on the rope: the positive pulls, the negative pushes.



REMEMBER

Turbo models want CFG 1, where the negative lets go of the rope. Too high gives burnt colours and extra fingers.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

- model LATENT
- positive
- negative
- latent_image
- seed 42
- control after generate fixed
- steps 8
- cfg 1.0
- sampler_name euler
- scheduler simple
- denoise 1.00

Sampler

sampler_name · euler, dpmp2m and friends

WHAT IT IS

HOW the noise is removed at each step. Euler, DPM++ 2M and the rest are different methods for the same job.



THE PHOTOGRAPHER VERSION

Her walking style down the stairs.



REMEMBER

Use what the model page recommends. Euler is the safe default.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

- model LATENT
- positive
- negative
- latent_image
- seed 42
- control after generate fixed
- steps 8
- cfg 1.0
- sampler_name euler
- scheduler simple
- denoise 1.00

Scheduler

simple · karras · beta · and friends

WHAT IT IS

WHEN the noise is removed: how big each step is from start to finish. Different schedulers spread the same work differently.



THE PHOTOGRAPHER VERSION

The height of each stair.



REMEMBER

Sampler and scheduler usually come as a pair on the model page. Copy both.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

model	LATENT
positive	
negative	
latent_image	
seed	42
control after generate	fixed
steps	8
cfg	1.0
sampler_name	euler
scheduler	simple
denoise	1.00

Denoise

denoise strength

WHAT IT IS

How much of the starting latent is allowed to change, from 0 to 1. It only matters when you start from an existing image.



THE PHOTOGRAPHER VERSION

'Redraw this sketch, but keep most of it.'



REMEMBER

1.0 ignores your input image completely, the classic mistake. Try 0.3 to 0.7 for image-to-image.

↓ a KSampler setting

Load Checkpoint

Prompt +

Prompt -

Empty Latent

KSampler

VAE Decode

Save Image

IN COMFYUI

KSampler

model	LATENT
positive	
negative	
latent_image	
seed	42
control after generate	fixed
steps	8
cfg	1.0
sampler_name	euler
scheduler	simple
denoise	1.00


Canvas words

the words you'll hear while moving boxes around



Group

A coloured rectangle that moves nodes together. Purely for tidiness.



Bypass • Ctrl+B

Purple. The wire runs straight through and the node is skipped. Perfect for before/after tests.



Mute • Ctrl+M

Grey. Scissors on the wire: everything after it errors.

The node called in sick and told nobody.



Subgraph

Several nodes folded into one box. Double-click to look inside.



Template

A saved copy-paste of nodes you reuse across workflows.



Custom node

A folder of Python from GitHub, living in custom_nodes. Delete the folder to uninstall.



Manager

The app store for custom nodes. 'Install Missing Custom Nodes' fixes red nodes.



Run • Queue

Run adds the workflow to a queue. Only nodes whose inputs changed run again.



double-click = search • **R** refresh files • **Ctrl** + **B** bypass • **Ctrl** + **M** mute • **Alt** + drag clone

Hardware words

the words behind 'why is this so slow?'



VRAM

Memory on the graphics card. Everything fast happens here.

The photographer's desk.



RAM

System memory. Overflow lands here and everything slows down.

The filing cabinet in the hallway.



Offloading

ComfyUI shuffling model pieces between VRAM and RAM for you. Leave it alone in Part 1.



Attention • sage • flash • pytorch

Different maths engines for the same step. Faster ones can break some models: black images, buzzing audio. When in doubt, use the default.



CUDA

Nvidia's language for talking to the GPU. 'CUDA out of memory' means the desk is full.




Rule of thumb: model file size + about 2 GB $\approx$ the VRAM you need.

Grab the cheat sheet.

ComfyUI cheat sheet

The files you download and the nodes that load them



Part 1 - Free - page 1/2
plain meaning - the photographer version in yellow - what to remember in grey

A - Files

Model / checkpoint — the big file of numbers that actually draws. A checkpoint bundles model + text encoder + VAE. *The photographer you hire; she brings her helpers.* models/checkpoints

Diffusion model — the same model shipped alone, without helpers. Most new models come this way. *She shows up with just her camera.* models/diffusion_models

Text encoder (CLIP) — turns your words into numbers. "CLIP" is the slot name; the file may be T5, Owen, Gemma... *The translator between you and the photographer.* models/text_encoders

VAE — converts between the small working image (latent) and real pixels. Encode = in, Decode = out. *The darkroom.* models/vae - wrong family = odd colours or shape error

Latent — the compressed working image, ~8x smaller per side. Generation happens here. *The negative.* sizes in multiples of 8 (some models 16/32/64)

LoRA — a small add-on that nudges a model toward a style, face or skill. Must match the model family. *She took a weekend class. Same camera, new trick.* models/loras - trigger word - strength 0.5-1.0

ControlNet — steer with an image (pose, edges, depth) instead of only words. *You hand her a reference photo: 'stand like this!'* prompt - WHAT, ControlNet - WHERE

Upscaler model — enlarges an image and invents plausible detail. *The print lab.* models/upscale_models

A - Formats

FP16 / FP8 / GGUF Q8-Q4 — same model, fewer decimals. Smaller file, less VRAM, slightly rougher. Q8 = perfect, Q4 visibly rough. *Same photo saved at 4K / 1080p / 720p.* file size + 2 GB - VRAM needed - GGUF needs the GGUF loader nodes

safetensors — modern model format that cannot hide code inside. Prefer it over .ckpt. It is safer, not faster

Read the filename — family - version - precision (fp8) - variant (scaled) - packager tag (KJ) - extension. Only precision and extension change what you need.
z_image_turbo_fp8_scaled_KJ.safetensors

B - Nodes in a basic workflow

Node — a box that does one job. Inputs left, outputs right, settings inside. *One crew member, one task.* double-click the canvas to search

Link — a wire from an output to an input. Colour = data type; wrong colour won't plug, output = power strip (many) - input = wall socket (one)

Workflow — all nodes + links, saved as JSON. Also hidden inside every PNG ComfyUI saves. *The shot list.* workflows never contain models - drag a PNG in to restore it

Load Checkpoint — one node, three outputs: MODEL, CLIP, VAE. *Hire the photographer with her helpers.*

Load Diffusion Model / Load CLIP / Load VAE — the three-node version for models shipped in pieces. Load CLIP's 'type' must match the model family. *Hire everyone separately.* GGUF files: Unet Loader (GGUF) + CLIP Loader (GGUF)

CLIP Text Encode +2 — prompt → conditioning. One for what you want, one for what you don't. *The brief, and the 'please no' list.* the two nodes are identical; the KSampler socket decides

Conditioning — the encoded prompt (orange cable). turbo models ignore the negative → 'Conditioning Zero Out' as an empty plug

Empty Latent Image — a blank latent of the size you pick. *Choosing the blank photo paper.* swap for Load Image → VAE Encode to start from a picture

KSampler — the node that generates: model + positive + negative + latent → finished latent. *The shoot.* all the settings you'll argue about live here

VAE Decode / Encode — latent → pixels at the end; pixels → latent when starting from an image. *The darkroom, both directions.*

Save Image / Preview Image — Save writes to output/. Preview shows it (temp/, cleared). 'prompt has no output' - you forgot one

Load Image — a picture from input/. Outputs IMAGE and MASK; mask editor on right click.

The seven-node pipeline

- 1 Load Checkpoint — hire
- 2 CLIP Text Encode (+ / -) — brief the translator
- 3 Empty Latent Image — choose the paper
- 4 KSampler — the shoot
- 5 VAE Decode — the darkroom
- 6 Save Image — deliver

The folders you actually touch

models/ — checkpoints - diffusion_models - text_encoders - vae - loras - controlnet - upscale_models wrong subfolder = 'value not in list' - press R after adding

input/ output/ temp/ — what you load - what you save - previews (cleared)

custom_nodes/ — one folder per node pack (Python from GitHub). Delete folder = uninstall. red node → Manager → Install Missing Custom Nodes → restart

user/default/workflows/ — where your saved JSON workflows live.

First-week errors

Value not in list — file not in the folder, or not refreshed (R).

Red node — missing custom node → Manager.


CUDA out of memory — too big: lower size/batch or use FP8/GGUF.

Shape mismatch — wrong VAE / encoder / LoRA family.

Error at node X — X is where it failed, not why. Walk one node upstream.

ComfyUI cheat sheet

The settings, the canvas, and the hardware



Part 1 - Free - page 2/2
plain meaning - the photographer version in yellow - what to remember in grey

C - Inside the KSampler

Seed — the random starting noise as a number. Same seed + same settings = same image. *Which roll of the dice you started with.* control after generate: fixed / increment / randomize - fix it when comparing

Steps — how many denoising passes. More = cleaner, slower, then no better. *How many wipes of the foggy window.* classic 20-30 - turbo 4-8 - use the model page's number

CFG — how hard the model is pushed toward the prompt and away from the negative. *Tension on the rope: positive pulls, negative pushes.* turbo models: CFG 1 (negative lets go) - too high = burnt colours, six fingers

Sampler — HOW noise is removed each step (Euler, DPM++ 2M...). *Her walking style down the stairs.* Euler is the safe default

Scheduler — WHEN noise is removed: how much per step (simple, karas, beta...). *The height of each stair.*

Denoise — how much the input may change, 0-1. Only matters when starting from an image. *Redraw this sketch but keep 70% of it.* 1.0 - ignores the input (the usual mistake) - 0.3-0.7 for img2img

Batch size vs batch count — size - images in one go (more VRAM). count - run N times (more time).

How to test anything

The compare ritual — fix the seed → change ONE thing → run → compare with an Image Compare node. Bypass the node you're testing (Ctrl+B) for a clean before/after. *Lab coat optional.*

D - Canvas words

Group — a coloured rectangle that moves nodes together. Cosmetic.

Bypass (purple, Ctrl+B) — wire runs through, node skipped. Downstream still works. use for A/B tests

Mute (grey, Ctrl+M) — scissors on the wire; downstream errors. *The node called in sick and told nobody.* rarely useful

Subgraph — nodes folded into one box. Double-click to open.

Template — a saved copy-paste of nodes.

Custom node / Manager — Python from GitHub in custom_nodes/. Manager is the app store for them. Install Missing Custom Nodes → restart

Queue / Run — Run adds the workflow to a queue. Only nodes whose inputs changed re-run.

Shortcuts — double-click: search - R: refresh files - Ctrl+B / Ctrl+M: bypass / mute - Alt-drag: clone - Ctrl+Shift+V: paste with links - Space+drag: pan

Keeping it working

Update rules — close ComfyUI first - never before a deadline - if it works, don't touch it - rollback exists. *It updates like a toddler: daily, and everything breaks.*

PNG trick — every saved PNG carries its workflow. Drag it onto the canvas. Discord and Twitter strip it.

E - Hardware words

VRAM — memory on the graphics card. Everything fast happens here. *The photographer's desk.* file size + 2 GB = what you need

RAM — system memory. Overflow goes here and slows down. *The filing cabinet in the hallway.* 32 GB comfortable, 64 GB rescues small cards

Offloading / dynamic VRAM — ComfyUI shuffling pieces between desk and hallway. Leave it alone in Part 1.

Attention (sage / flash / pytorch) — different maths engines for the same job. Faster ones can break some models. black image or buzzing audio → try the default

CUDA — Nvidia's language for the GPU. 'CUDA out of memory' - desk is full

Your tier

6-8 GB — everything in Part 1 with FP8 or GGUF.

12-16 GB — everything. FP8 comfortably; most video at 720p later

24 GB+ — everything. FP16 if you like.

Mac / AMD — images via the Desktop app; fp8 files may not load.

No GPU — cloud (RunPod / RunningHub), pay per hour.

Model classes

Text → image - Edit - Inpaint - Video - Audio - LLM — six kinds of model. Most 'it doesn't work' is asking an edit model to inpaint, or a text-to-image model to edit. *Pick the right photographer for the job.* Part 2 goes deep on this

You don't need to remember any of this. You need to know where to look. - ComfyUI From Zero, Part 1 (free)

You don't need to remember any of this. You need to know where to look.



006

• MODULE 6 • 20 MIN

Meet the photographer

Seven nodes, one metaphor, and the why behind every setting.

-  Build text-to-image from an empty canvas
-  How noise turns into a picture
-  Seed, steps, CFG, sampler, scheduler


 VRAM • 6 GB

 Free



• THE WHOLE PIPELINE

Six jobs, one photographer.






Now build it yourself.

Empty canvas. No downloads. Just the seven boxes.

- 1 Double-click the canvas, search **Load Checkpoint**
- 2 Drag from **CLIP**, drop, pick **CLIP Text Encode**. Twice.
- 3 Add **Empty Latent Image** and **KSampler**
- 4 Finish with **VAE Decode** and **Save Image**, then Run



- THE PROMPT BOXES

Two identical nodes. The socket decides.



Swap the two cables and the prompts **swap jobs**.


• THE WIRING RULES

Power strips out. Wall sockets in.




One output, many cables

Like a power strip.




One input, one cable

Like a wall socket. A second plug won't fit.



Colours are types

Yellow only fits yellow. ComfyUI is protecting you.



• HOW IT ACTUALLY WORKS

Every image starts as **TV static**.

- 🌊 **Start:** pure random noise, like a TV with no signal.
- 📺 **Early steps:** big shapes and colours come out of the fog.
- 🌟 **Late steps:** small details: whiskers, the hat, the light.




It's a foggy window. **Every step wipes a bit more fog away.**

step 30 / 30




● STEPS

How many times you wipe the foggy window.




1 step



2 steps



4 steps



8 steps



12 steps



20 steps



30 steps



Turbo models: 4 to 8

Trained to look finished fast. More steps won't help much.



Classic models: 20 to 30

Need more wiping. Past that, you're just waiting longer.

● SEED

The seed is your roll of the dice.



a photo



the exact same photo



a different shot



Comparing settings? **Fix the seed first.**

- CFG

CFG is the tension on the rope.

Positive pulls toward what you asked for. Negative pulls away from what you banned.



POSITIVE



NEGATIVE

CFG 15+

Everyone pulls too hard.

Burnt colours, crunchy edges, extra fingers.



• SAMPLER VS SCHEDULER

Sampler is **how**. Scheduler is **when**.



SAMPLER • HOW

Her walking style



- ▶ euler • steady steps
- ▶ dpmpp_2m • looks back, then steps
- ▶ ...and friends



SCHEDULER • WHEN

The height of each stair



simple • even stairs

karras • big drops first



Copy the pair from the model page. **Euler + simple** is the safe default.

• WHAT DID IT ACTUALLY LEARN?

Visual grammar. **Not a scrapbook.**

WHAT PEOPLE THINK

A folder of stolen photos



WHAT IT REALLY LEARNED

How pictures work

- ✓ cats have two ears on top
- ✓ light comes from one direction
- ✓ shadows fall away from the light
- ✓ far things look smaller
- ✓ fur catches the light at the edges

Like learning a language: you learn grammar, not every sentence ever written.

• YOUR LAB HABIT

Change one thing. **Compare.**

-  **Fix the seed.** Otherwise you're comparing luck.
-  **Change exactly one thing.** Steps, CFG, a word in the prompt.
-  **Compare side by side.** An Image Compare node makes it a slider.



IMAGE COMPARE



The four ideas, now with faces.

Noise



Empty Latent Image

seed

The blank paper and the dice roll decide what static you start from.

Conditioning



CLIP Text Encode

Your prompt, translated, steering what the static becomes.

Sampling



KSampler

steps · sampler · scheduler

How the fog gets wiped away, and how many times.

Structure vs style



prompt today

LoRA in Part 2

Where things go and how they look. Today the prompt does both.

● MODULE 6 COMPLETE

You can now...

- ✓ Build **text-to-image** from an empty canvas
- ✓ Explain **positive vs negative** and the cable rules
- ✓ Say what **steps, seed, CFG, sampler** and **scheduler** do
- ✓ Test anything with the **compare ritual**

NEXT UP

07

Reading model files

Look at any model page and know exactly what to download and where it goes.



07

• MODULE 7 • 14 MIN

Reading model files

Look at any model page and know exactly what to download and where it goes.

-  FP16, FP8 and GGUF without the panic
-  Read a filename like a label
-  The three sites everything comes from

 VRAM • none

 Free



• FP16 • FP8 • GGUF

Same photographer. Different file sizes.



FP16

100%


Full size. The original.



FP8

~50%

Half the size. Barely different.



GGUF Q4

~30%

Tiny. Visibly rougher.

She isn't dumber in the small file. **Just a bit blurrier.**



- ONE FILE OR THREE?

A checkpoint brings her helpers. A diffusion model comes **alone**.

ALL-IN-ONE


Checkpoint



MODEL



TEXT ENCODER




VAE

models/checkpoints/

Load Checkpoint

IN PIECES

Diffusion model + helpers



the_model.safetensors




diffusion_models/



text_encoder.safetensors



text_encoders/



vae.safetensors



vae/

Load Diffusion Model · Load CLIP · Load VAE



Wrong drawer = "value not in list". Every time.

- FILENAMES

Read a model filename like a food label.

z_image_turbo_fp8_scaled_KJ .safetensors

FAMILY + VERSION

which photographer

PRECISION

how big, how
rough

VARIANT

tuned for quality

PACKAGER

who shared it.
Changes nothing.

FORMAT

safe file, normal loader

SAME IDEA, GGUF EDITION

z_image_turbo-Q4_K_M.gguf

Q4 = 4-bit, rougher

K_M = medium flavour

.gguf = needs the GGUF loader



• LOADERS

Match the **loader** to the file.

📄 the_model.safetensors

● Load Diffusion Model



Normal file, normal loader.

📄 the_model-Q4_K_M.gguf

● Unet Loader (GGUF)



GGUF file, GGUF loader.

📄 the_model-Q4_K_M.gguf

● Load Diffusion Model



Wrong loader. It can't even see the file.

📄 GGUF model

📄 safetensors text encoder

● GGUF loader + Load CLIP




Mix and match. Each file picks its own loader.

Pick the sealed envelope.




.ckpt • .pt • .pth



The old formats can hide **code that runs** the moment the file is loaded. A parcel that might contain a surprise.



.safetensors



Holds **numbers only**. Nothing can run when you load it. A sealed envelope of photos.

It's **safer**, not faster. Always pick it when you have the choice.

• WHERE TO GET THINGS

Three addresses cover almost everything.



Civitai

civitai.com

The style agency. Someone taught a look and packed it into a **LoRA**.

✦ **LoRAs** and fine-tunes

🔍 Filter by **base model**

🔗 `models/loras`



Hugging Face


huggingface.co

The warehouse. **Base models** and their pieces, from the people who built them.

📦 Checkpoints, encoders, **VAEs**

📄 **Files** tab, pick precision

🔗 `models/checkpoints`



GitHub

github.com

The tool shop. **Custom nodes**, and ComfyUI itself, written by other people.

🔗 **Node packs**, not models

🔧 Install via the **Manager**

🔗 `custom_nodes`




Take models from the creator's **own page**, and only install node packs you would trust with your whole computer.



Read a model page in 30 seconds.


- 1 **Model card:** what it is, recommended settings, licence
- 2 **Files:** pick the precision that fits your tier
- 3 Grab the **helpers** too, if it ships in pieces
- 4 **Licence:** check it before you sell anything




- NEW MODEL, NEW NODES?

How do you know what a new model needs?

- ② **The problem:** some models need a special node, and the model page doesn't always say so.
- 🗑️ **The answer:** open its official ComfyUI template. Workflow menu → Browse Templates.
- 📄 **Copy it.** Its nodes and settings are the source of truth, not a random screenshot.



Six kinds of model. Hire the right one for the job.




Text → image

Blank paper, describe it, get a photo.



Edit

Give it a photo and an instruction. It redraws the whole picture.



Inpaint

Mask one area and change only that. The rest stays untouched.



Video

Same idea, many frames at once. Hungry for VRAM.



Audio

Voices, music and sound effects.



LLM

Writes your prompts and describes your images.

Most "it doesn't work" is the right photographer for the **wrong job**.

[🔒 Part 2 goes deep](#)

You can now...

- ✓ Choose **FP16, FP8 or GGUF** for your card
- ✓ Put **checkpoints** and **pieces** in the right folders
- ✓ Read a **filename** and pick the right **loader**
- ✓ Find a new model's nodes in its **official template**

NEXT UP

08

Saving, sharing, not breaking things

Keep your work, share it, and update without wrecking a setup that works.



08

• MODULE 8 • 12 MIN

Saving, sharing, not breaking things


Keep your work, share it, and update without wrecking a setup that works.

-  Save, Save As and Export
-  The PNG that remembers everything
-  Custom nodes, updates, bypass vs mute



• KEEPING YOUR WORK

Save. Save As. Export.



SAVED WORKFLOWS LIVE IN


📁 ComfyUI/user/default/workflows/



Save

overwrite

Keeps the same name. The old version is gone.



Save As

new copy

A second workflow with a new name. Try things safely.



Export

send anywhere

A .json file you choose where to put. For sharing and backups.



- THE BEST TRICK IN COMFYUI


Every image you save **remembers its workflow.**



🖼️ ComfyUI_00001_.png

```
Notepad · ComfyUI_00001_.png

%PNG
IHDR 1024 x 1024
tEXt prompt {"3": {"class_type": "KSampler",
"inputs": {"seed": 42, "steps": 8, ...
tEXt workflow {"nodes": [{"id": 4,
"type": "CheckpointLoaderSimple", ...
```



→ Drag the PNG onto the canvas and the whole workflow comes back, same seed and all.

⚠️ Discord and most social sites strip it

© your prompt is in there too

● CUSTOM NODES

A custom node is just a folder of code.



someone / ComfyUI-Handy-Nodes

__init__.py
nodes.py
requirements.txt
README.md

★ 1.2k

updated 3 days ago




ComfyUI/

custom_nodes/

ComfyUI-Handy-Nodes/

← the whole thing

another-node-pack/



Delete the folder = uninstalled.




It's code that runs on your machine. **Only install packs you trust.**

- RED BOXES YOU CAN'T OPEN

Missing node? Let the Manager fetch it.

• Handy Resize

"When loading the graph, the following node types were not found."



1 · Manager

2 · Install Missing Custom Nodes

3 · Restart

4 · Red box gone



• UPDATING

ComfyUI updates like a **toddler**.

1



Close it first

Quit ComfyUI completely before you update. Both windows.

2



Never before a deadline

Updating the night before a client job is how nights get long.

3



Works? Don't touch it

Update when you need a new model or a fix, not because a button is glowing.



Broke anyway? **Roll back** to the version that worked. It's in the Manager.

• TURNING NODES OFF


Bypass vs mute. Learn it once, side by side.

BYPASS

Ctrl

B

the wire runs straight through · downstream still works




MUTE

Ctrl

M

scissors on the wire · everything after it errors



Use **bypass** for before/after tests. Use **mute** almost never.

Three ways to tidy up.

Group

Prompts

Prompt +

Prompt -

A coloured rectangle. Drag it and the nodes inside come along. **Purely cosmetic.**

Subgraph

 Upscale (5 nodes)

double-click to open

Several nodes folded into one box. **Double-click** to open it. Common in downloaded workflows.

Template

Loader

KSampler

Save

A saved copy-paste of nodes. Right-click the canvas to drop your favourite setup in again.

● MODULE 8 COMPLETE

You can now...

- ✓ Choose **Save**, **Save As** or **Export** on purpose
- ✓ Get a workflow back from any **ComfyUI PNG**
- ✓ Install, fix and remove **custom nodes** safely
- ✓ Update without breaking things, and use **bypass** for tests

NEXT UP

09

The eight errors

You will see every one of these. Better to meet them here than alone at midnight.



● MODULE 9 • 12 MIN

You will see every one of these. Better to meet them here than alone at midnight.

- ⚠ Eight errors, eight fixes
- 👣 Walk upstream to the real cause
- 👤 Ask for help and actually get it

VRAM • any

 **Free**



Every error tells you three things.

❗ **Prompt outputs failed validation**

VAEDecode:

Required input is missing: vae

1 What kind of problem

Validation means it failed before running. Good news: nothing heavy happened.

2 Where it noticed

The node that complained. Not always the node that's wrong.

3 The clue

The actual reason, often with the exact name of what's missing.

Red node: missing custom node

a downloaded workflow opens with red boxes

WHAT YOU SEE

❗ Missing node types

When loading the graph, the following node types were not found:

HandyResize · HandySharpen

● HandyResize

● HandySharpen



WHAT IT MEANS

The workflow uses a node pack you haven't installed yet. **Nothing is broken.**

THE FIX

- 1 Open the **Manager**
- 2 **Install Missing Custom Nodes**
- 3 **Restart** ComfyUI

● ERROR 2 OF 8

Value not in list

a loader can't find its file

WHAT YOU SEE

❗ Prompt outputs failed validation

```
CheckpointLoaderSimple:  
Value not in list: ckpt_name:  
'z_image_turbo_aio.safetensors' not in []
```

● Load Checkpoint

	MODEL
ckpt_name	null



WHAT IT MEANS

The model isn't in the folder this loader looks in, or ComfyUI hasn't noticed it yet.

THE FIX

- 1 Download the file the error names
- 2 Put it in the **right folder**
- 3 Press **R**, or restart



• ERROR 3 OF 8

CUDA out of memory

the job is bigger than your graphics card


WHAT YOU SEE

ComfyUI · server

```
torch.OutOfMemoryError: CUDA out of memory.  
Tried to allocate 2.00 GiB. GPU 0 has a total  
capacity of 8.00 GiB of which 312.00 MiB is free.
```



VRAM



WHAT IT MEANS

The job doesn't fit on the desk. Your GPU ran out of VRAM.

THE FIX

- 1 Lower the **resolution** or **batch size**
- 2 Use a smaller file: **FP8** or **GGUF**
- 3 Close other apps using the GPU and run again



• ERROR 4 OF 8

Shape mismatch

the scariest-looking error is usually the simplest

WHAT YOU SEE

ComfyUI · server

```
RuntimeError: mat1 and mat2 shapes  
cannot be multiplied (154x2048 and 768x320)
```

📷 model: family A



🧩 helper: family B



WHAT IT MEANS

Parts from **different model families** are wired together: a VAE, text encoder, LoRA or ControlNet made for another model.

THE FIX

- 1 Check every helper matches the **model family**
- 2 Check **Load CLIP's type** setting
- 3 When unsure, use the files from the **official template**



● ERROR 5 OF 8

Prompt has no outputs

you pressed Run and nothing happened

WHAT YOU SEE

❗ **Prompt outputs failed validation**

Prompt has no outputs

● KSampler

● VAE Decode

Save Image?
nowhere to go



WHAT IT MEANS

Nothing in the workflow saves or shows a result, so ComfyUI refuses to run it.

THE FIX

- 1 Add a **Save Image** or **Preview Image** node
- 2 Connect it to the **IMAGE** output
- 3 Run again

The file won't show up

it's in the folder, but not in the dropdown

WHAT YOU SEE

• Load Diffusion Model

	MODEL
unet_name	(none)
weight_dtype	default

IN THE FOLDER

models/diffusion_models/

the_model-Q4_K_M.gguf

← right there



WHAT IT MEANS

It's a **GGUF** file, and the normal loader only lists .safetensors files.

THE FIX

- 1 Install the **GGUF node pack**
- 2 Use **Unet Loader (GGUF)** or **CLIP Loader (GGUF)**
- 3 Leave the file where it is

Black image or buzzing audio

it runs, finishes, and gives you nothing

WHAT YOU SEE



ComfyUI_00007_.png



speech.wav · bzzzzz



WHAT IT MEANS

Usually a **speed-up setting**, like sage attention, that this particular model doesn't agree with.

THE FIX

- 1 Start ComfyUI **without** sage attention, using the default launcher
- 2 Run the exact same workflow again
- 3 Turn speed-ups back on only for models that like them

• ERROR 8 OF 8

Required input is missing


a socket with no cable

WHAT YOU SEE

⚠ Prompt outputs failed validation

VAEDecode:

Required input is missing: vae



WHAT IT MEANS


A socket on that node has no cable plugged in. The message even tells you which one.

THE FIX

- 1 Find the node the error names
- 2 Find the socket it names, here **vae**
- 3 Connect the cable and run again

• THE GOLDEN RULE

The error shows **where** it failed. Not **why**.




Q VAE Decode was just the messenger. The **wrong VAE** two nodes back was the problem.

🔧 Fix: load the VAE that matches the model family

• STILL STUCK?

Ask for help in a way that **gets answered.**

-  **Your machine:** GPU, VRAM, RAM, operating system
-  **Your ComfyUI:** Desktop or portable, when you last updated, a **screenshot** of the workflow
-  **The error:** the full text, and what you already tried



● MODULE 9 COMPLETE

You can now...

- ✓ Read any error: **kind, where, clue**
- ✓ Fix the **eight** errors beginners actually see
- ✓ **Walk upstream** to the real cause
- ✓ Ask for help with **everything** a helper needs

NEXT UP

10

What you can do now

A real recap, three challenges, and where the road goes next.



10

• MODULE 10 • 6 MIN

What you can do now

A real recap, three challenges, and where the road goes next.

- ✓ Everything you can do now
- 🎯 Three homework challenges
- 🔓 What Parts 2 and 3 unlock

📦 VRAM • none

📁 Free



• THE REAL RECAP

Ten things you can do **right now**.




Explain **model, studio** and **workflow**




Know your **VRAM tier** and install path



Install and launch **without panic**




Make an image and **find** it



Name every word on the **cheat sheet**



Build **text-to-image** from an empty canvas



Tune **steps, seed, CFG, sampler, scheduler**



Pick the right **model file** and **folder**



Save, share and **update safely**



Fix the **eight errors** and walk upstream



• HOMEWORK

Three challenges. No peeking.

1



warm-up

Rebuild it blind

Empty canvas. Build text-to-image from memory and make one image.

2




medium

Swap the file

Switch to the GGUF version of the model. Same seed. Compare both with an Image Compare node.

3



spicy

Fix workflow 03

It's broken in three different places on purpose. Find all three.



Post your fixes in the comments.



• WHAT'S NEXT

Machine understood. Now **open the hood.**

PART 2

🔒 paid bundle

Open the hood

- LoRAs, properly
- Edit, inpaint, outpaint
- Low-VRAM tiers, tested
- Cloud GPUs, step by step
- Video models
- Prompting formulas
- Real project workflows

PART 3

🔒 same bundle

Build the engine

- Train your own LoRA
- Write a custom node
- Merge your own models
- Build datasets that work
- Automate with the API
- Licences for client work

🔗 Parts 2 + 3 • link in the description

Part 1 stays free. Understanding should be.





Thanks for watching.

See you in the studio.

